UNIVERSITY OF SOUTHAMPTON

Telling Your Story: Autobiographical Metadata and the Semantic Web

by

Mischa Moussavian Tuffield

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science
Intelligence, Agents, Multimedia Group

July 2010

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE
INTELLIGENCE, AGENTS, MULTIMEDIA GROUP

Doctor of Philosophy

by Mischa Moussavian Tuffield

Given the current explosion of user-generated content driven by the ever-decreasing price of sensing and storage hardware the dream of capturing and archiving the entirety of a human life is slowly being realised. The Semantic Web, a discipline of Computer Science, aims to support the sharing and interoperation of knowledge using the Web's infrastructure. This thesis aims to roadmap a framework utilising the principles and technologies underpinning the Semantic Web, enabling the vision of global knowledge sharing, in an open and policy aware manner, with the end aim of supporting a network for the exploitation of personal information. This sharing is facilitated through the adoption of a *lingua franca*, shared conceptualisations for domain knowledge, and some core design principles. The main focus of Semantic Web research has been the development of a web-scale knowledge-base whereby information is stored and exposed in a machine-readable format with the ultimate aim of aggregating information from disparate sources, allowing for statements to be contextualised with respect to others culminating in a web-scale knowledge resource accessible through standard protocols.

The current popularity of social computing – *Web 2.0* – where users post personal information to online communities is eluding to the fact that information, linked and shared within a social-context presents added value to the end-user. Given the sensitive nature of personal information, one may not wish to expose all of the information about them self to the World Wide Web, but may wish to benefit by linking to knowledge residing on this shared resource. This ability to store personal information privately, in ones own personal web-space and not on a third party server, whilst at the same time connecting to the publicly available information is presented as key challenge facing the Computer Science community today. Specific information pertaining to one aspect of a user's activities, such as their picture taking habits or their geographic log, may not present a detailed account of a user's actions, but as more information is pushed into the public domain and aggregation technologies mature individuals and their day-to-day activities will be easier to track.

As more and more of our personal lives are pushed into the public domain, the notion of an *online-persona* is becoming more and more applicable to the average person.

This thesis presents an infrastructure for the capturing and archival of autobiographical metadata, whereby information from multiple sensors is aggregated and stored in a personal *Lifelog*. The surrender of digital identity has become commonplace, for purposes ranging from commerce, marketing, social networking, government, receipt of services, travel or security, Lifelogging has the potential to reaffirm the individual's control of his or her own digital identity. The Lifelog is a constructed identity that outweighs the others simply by weight of evidence, complexity and comprehensiveness.

This thesis presents an infrastructure for the capture and exploitation of personal metadata to drive research into context aware systems. The aim is to expose ongoing research in the areas of capture of personal experiences, context aware systems, multimedia annotation systems, narrative generation, all set in the context of enabling and supporting the Semantic Web Vision. The thesis details the work underway towards the goal of creating a multi-domain contextual log, and is followed by a discussion of how such a log can be used to drive the development of detailed Lifelog and an investigation into the amount of personal information being pushed into the public domain.

Contents

List of Figures

List of Tables

Acknowledgements

I would like the thank my supervisors Professor Nigel R. Shadbolt and Dr David E. Millard, for their irreplaceable guidance. My sister Joy gets a special mention for putting so much time and effort into proof reading my thesis – now you know what your brother gets up to! Since starting my PhD I have had access to a wealth of knowledge and have been given the opportunity to immerse myself within a vibrant and exciting research community, and this would not have been possible without the support of my supervisors and the research group as a whole. A big thanks to all of my friends in the Intelligence, Agents, Multimedia (IAM) group for making me feel at home in Southampton, without you none of this would have been worthwhile. I would also like to thank the AKTors for all of the stimulating conversations I was fortunate to be a part of.

Here I will stress the importance of my family who have supported me from day one. A big thank you to my mother Minoo, without you I would have been nothing, my sister Joy who has brought me nothing but, and my late father Paul for helping me put everything into context. Thank you all for allowing me the opportunity to grow up around such wonderful people.

Following on from this I would also like to give special thanks Dr Antonis Loizou, Dr Mike O Jewell, Dr Faith Lawrence, Dr Dave P Dupplaw, Dr Harith Alani, Dr Martin Szomszor, and Dr Kieron O'Hara for their collaborative efforts during my studentship. I should also thanks Dr Mark Weal, Dr Yves Raimond, Parastou Marashi for helping proof my thesis and for being so lovely. I would also like to thank Professor Dame Wendy Hall for her support and guidance throughout the years, and would like to thank Susan Davies for everything she has done for me over the years.

I should thank all of my friends for making my life enjoyable and eventful: Sebastien Francois for being there whenever need be, Dr Paul Groth for friendship and a critical eye, Dr Steve Munroe, Dr Simone Scaringi and Dr Antonis Loizou for making my house feel like a home. Finally, I would like to thank Dr Ashley Smith for his cynicism and company throughout.

Finally, I should also thank the people at Garlik Ltd, my current employer, who have been very supportive of my studies and my interests in the Semantic Web, Privacy, and Personal Information. And finally I should give a big thank you to Steve Harris, for showing me an enthusiasm for Web technologies, engineering, software design, and for helping me take my interests forward.

This work was supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01.

To My Family, You Rule…

Chapter 1

Introduction

The huge growth of the Internet and the World Web Web has meant that in the last few years more information has been generated and disseminated than had previously been thought possible. The ability to link heterogeneous information in the form of web accessible documents hosted on computers all over the world has forever changed the way knowledge will be shared and propagated through society. The Semantic Web (SW) (Berners-Lee et al., 2001), the proposed future complement of the document web, aims to do for knowledge what the World Wide Web (WWW) did for documents, that is, link and share it on a global scale. The Semantic Web vision is to facilitate the sharing of knowledge in a machine-readable manner, on a global scale, so that humans can deploy software agents on their behalf to undertake tasks in a knowledge rich domain.

Technological advances have continuously shaped the way we have gone about our everyday lives. Language, the ability to communicate, allowed humans to share knowledge they had acquired with one another, allowing for an individual to gather wisdom from more than just their own experiences. Following on from this came the advent of the written word, allowing for innovators to document their findings, freeing up time previously spent on spreading knowledge to that of discovery. The printing press further revolutionised the dissemination of knowledge to the masses, allowing for books to be printed in volumes not previously viable, catering for knowledge to be shared by a larger audience. This automation of the book production process, in turn freed up time previously allocated to scribing.

More recently, since the birth of the World Wide Web in 1989[1] the speed in which knowledge can travel from one side of the world to the other has become almost instantaneous. But this speed is not the only revolutionary aspect about the WWW. The fact that anyone can publish whatever they want, and to an audience spanning the entirety of the globe has started affecting the way day to day events are recorded, forever shaping the way that history will be documented. This paradigm shift in the manner in which

[1] Tim Berners Lees original proposal http://www.w3.org/History/1989/proposal.html

events are documented can be witnessed in the way that news agencies have embraced the user-generated content in recent years. The recent trend of Social Networking (SN) and the advent of numerous Social Networking Sites (SNS), such as Facebook[2], Flickr[3], MySpace[4], Delicious[5], and so on, has brought about a new era in Internet participation whereby the barrier to entry has been significantly, dropped resulting in increased participation. This phenomena coined *Web 2.0* has highlighted peoples' willingness to share information about themselves with the rest of the world and is powered by the apparently willingness for the publication and sharing of user generated content. For example, now the photo sharing and microblogging communities are now two of the most active communities around, both of which promote the sharing and publication of personal information, in the shape of photographs and personal memoires. We are living in an age where an ever increasing proportion of young and old are maintaining profiles and posting personal information to social networking sites, lifelogging is introduced as a method of capturing and archiving this information personal gain.

The Semantic Web vision of enabling the sharing and linking of machine readable knowledge through the adoption of shared conceptualisation of domains via Web's architecture is a grand one. The Web's infrastructure, most importantly, the notion of Uniform Resource Identifiers (URI)s[6] and the Hypertext Transfer Protocol (HTTP)[7] presents the whole world with a set of unique identifiers whereby information about a particular resource is acquired by resolving the URI via a web-browser. In its simplest form the Semantic Web envisions that knowledge will be represented via the Resource Description Framework (RDF)[8] where statements about things are made in the form of *triples*. These triples take the form of a *Subject*, *Predicate*, and *Object*, e.g. *Nigel Shadbolt* (subject) *has an email address* (predicate) that is *nrs@ecs.soton.ac.uk* (object) or *The British Library* (subject) *is based near* (predicate) *Kings Cross train station* (object). This manner of representing knowledge is presented as being very similar to simple sentences used in the English language, that of *"noun, verb, noun"*. This manner of representing knowledge is one of the key building blocks of the Semantic Web.

1.1 Overview

Since the inception of the first Social Networking Site in 1997, and the subsequent explosion of similar social websites (boyd and Ellison, 2007) we have seen a rise in the amount of user generated content and an increase in the amount of personal metadata

[2]Facebook SNS http://www.facebook.com/
[3]Flickr Photosharing Site http://www.flickr.com/
[4]MySpace SNS http://www.myspace.com/
[5]Delicious SNS http://delicious.com/
[6]Cool URIs don't change : http://www.w3.org/Provider/Style/URI
[7]The HTTP specification at the W3C: http://www.w3.org/Protocols/
[8]RDF page on the W3C http://www.w3.org/RDF/

found on the World Wide Web. This apparent willingness to publish personal information to the Internet coupled with the fact that most people in the western world spend a good proportion of their time interacting with networked devices means that we are leaving behind vasts trails of *digital traces*[9] (O'Hara et al., 2009). These network devices include, but are not limited to: *laptop computers, mobile-telephones, digital cameras, global positioning systems* (GPS)s to name a few. One notable example of how ubiquitous networked computers are becoming more and more intertwined with our everyday lives is the advent of the world's first networked pacemaker, which was fitted in in the United States in August 2009[10]. This advance in technology will forever change what is said to make up the Internet of Things, as we move closer to world where even our bodily functions will leave *digital traces*.

Given that we live in an age where our personal information is stored and shared in more and more places, sometimes without our knowledge, we are witnessing a trend where our identity is pushed into the digital and public realm. This somewhat chilling view of the world we live in is one of the key motivating factors of this work. In order to help understand what personal information is being published about people, the concept of *Lifelogging* is presented as a means to educate users to what information about them is likely to be found if someone were to ever bother looking. Furthermore the proposed method of capturing autobiographical facts is said to empower end-users with an unabridged collection of personal facts which could potentially be used as an alibi if ever there was a case of mistaken identity which would leave someone falsely accused (Coughlin, 2007).

As more forms of personal information are pushed into the online communities, there is a need to roadmap a user- centric approach for the sharing of personal information in a social space. Positioning knowledge in a social environment allows for data to be placed within a context. This contextualisation of facts with respect to other facts is presented as one of key enablers of the phenomenon that is social networking, whereby people can upload personal information to the Internet, e.g. personal photos, whilst allowing for their pictures to be contextualised with respect to pictures their friends take. This thesis aims to detail a framework for capturing personal information whilst identifying mechanisms that would allow for users to maintain control of their personal information whilst still being able to interact with the Social Web.

This thesis is said to be a first step toward the dream of capturing a digital lifelog or a "Memory for Life". This work is intended to describe a framework for capturing personal information, and by virtue of the nature of the information detailed in this line of research it is important to stress how the data captured by the framework should only

[9]Digital Traces Definition on Wikipedia: http://en.wikipedia.org/wiki/Digital_traces

[10]First Wi-Fi pacemaker in US gives patient freedom http://www.reuters.com/article/newsOne/idUSTRE5790AK20090810

ever be accessed by the curator themself and should never be made accessible to anyone but the curator.

1.2 Thesis Statement and Contributions

This thesis presents the Semantic Framework for the capturing, archival, and exploitation of personal information. As more and more personal information is pushed out to the Internet, coupled with the fact that industry and the state are becoming more and more privy to this fruitful source of knowledge, it is important for users to know what information about them left behind in the form of a *digital trace* and what is available to interested parties. The focus of this work is to help empower users to capture their own Lifelog, which by virtue of curator and the fact that data was collected at its source, should outweigh any mutterings said about the user in the public domain. Further to the description of the Semantic Framework, this thesis sets forward two use-cases that demonstrate how the knowledge captured in the lifelog can be exploited.

These two complementary pieces of work are described in practical terms, focusing on the areas of data integration, infrastructure, and personal photo-annotation, with a focus on enabling the Semantic Web vision of the seamless sharing of knowledge. The framework aims to present a sound method for integrating a personal Lifelog with the Semantic Web, so that facts stored can be contextualised with respect to the rest of the Web. The first use case presented shows how the information captured in a lifelog can be combined with content-based image analysis to enrich one's personal photo-collection. The second use case takes this concept a step further by showing how distinct lifelogs can share information between themselves so as to benefit from being a part of a social network. The novelty in this approach is exhibited by how personal information can be shared in a social network without the need to use a central service like current Web 2.0 sites, that require users to upload their personal information to a third-party. The approach illustrates how individual peers, representing users' lifelogs, can share data in a distributed peer-2-peer network. This is said to allude to what form future distributed social networks may take, whereby users maintain ownership of their personal data whilst still being able to reap the benefits of being in a social network.

1.2.1 Thesis Statement

"People leave digital traces that can be captured at source as a Semantic Lifelog. Information captured in the proposed Semantic Lifelog can be used to enrich a personal multimedia collection. Personal information does not need to be given up to a third party service in order to benefit from interactions between peers in a social network."

1.2.2 Contributions

This thesis makes the following contributions:

- The definition of a number of motivations for capturing an autobiographical Lifelog (see chapter 3)
- The identification of the issues surrounding the capture of a Lifelog (see section 3.1.2)
- The definition of the concepts of privacy, data portability, and online identity in relation to digital personal information (see section 3.1.3)
- A demonstration of how stories and Narrative structures can be represented on the Semantic Web (see chapter 4)
- The specification of an open-framework for the capture, archival, and sharing of personal information (see chapter 5)
- The identification of best practises and standards which can be adopted to facilitate the storage of a Lifelog (see chapter 5)
- The definition of what is an easily captured *digital trace* (see section 5.3)
- A demonstration of the utility of data captured by the proposed Lifelog to aid the task of multimedia knowledge management (see chapter 6)
- An example application that shows how personal information can shared between trusted peers in a decentralised manner (see section 7)
- A roadmap for an open, distributed, decentralised, and socially contextualised Lifelog (see chapter 8)

1.3 Thesis Structure

Below is a breakdown of the contents of the various chapters.

- **Chapter 2 – Background & Related Work** divulges a critique on the existing literature related to this thesis. The foundations presented in the following chapter will include the relationship of the work to the Semantic Web (SW) Vision, a review of existing methods of representing autobiographical knowledge, and insight into the state of the art in the field of multimedia asset management. This chapter aims to provide insight into the relevant work to further contextualise the contributions made.

- **Chapter 3 – Motivating a Digital Autobiographical Log** discusses the motivations behind the work completed, setting the scene for the rest of the thesis. This chapter aims to contextualise the work by presenting the motivational factors, sources of inspiration followed by use cases aimed to help justify the research carried out. The chapter motivates the desire to capture an autobiographical log, i.e. a *Lifelog*, the reasons for wanting this log to interact in the Social Space that is the World Wide Web, along with reasons why the particular technology set was deemed most appropriate.

- **Chapter 4 – Narratives, Stories, & Ontologies** presents a machine readable model for representing events in a narrative structure. This chapter introduces the notion of representing knowledge in the form of a sequence of events, and how this form of modelling knowledge helps drive the subsequent research. This chapter motivates the use of narrative structures as a form of knowledge transfer, and presents OntoMedia – an ontology for representing heterogeneous multimedia.

- **Chapter 5 – Semantic Logger: An Autobiographical Contextual Log** defines a framework for the capture and archival of personal information. This chapter starts out by enumerating the sources of information harnessed by the framework, and presents insight into how the information captured could be shared in an open/standard compliant manner. Information pertaining to the technologies used to the develop the framework are presented followed by practical examples of how to exploit the knowledge gathered by the framework.

- **Chapter 6 – Photocopain: Context Driven Multimedia Asset Management** demonstrates the first use-case showing how the lifelogged data can be exploited to help automate the task of personal photo-annotation. The choice of metadata gathered is grounded in the experimental work of enriching metadata about one's personal photo-collection. This chapter describes the technologies exploited by the approach taken, and presents the manner that the infrastructure facilitates the interrogation and integration of the captured metadata. This chapter aims to the ground the framework with a practical use case, whilst providing guidelines to how the framework could be extended in the future.

- **Chapter 7 – Social Semantic Logger : Sharing Annotations with OpenKnowledge** this second use case sets forward the potential for personal knowledge bases to interact in a decentralised distributed environment is presented as the second use case for the Semantic Logger. Given that popularity of Social Networking Sites, and the perceived value added of interacting within a social network, the Semantic Logger has been ported to the OpenKnowledge peer-2-peer knowledge sharing framework. This chapter illustrates how users can benefit from sharing information within their social network without the need to surrender any of their personal information to a third party. This highlights the notion that one should not have to relinquish ownership of their personal information to a third party service in order to benefit from being a part of a social network.

- **Chapter 8 – Conclusions & Further Work** identifies the future direction of the work capitalising on the work completed, identifying the scope of our future research, while stressing the end product of this proposed research endeavour.

1.4 Declaration

The author declares that this thesis and the work presented within are his own and has been generated by the author as the result of his original research.

The following list enumerates parts of this thesis that were undertaken as collaborative efforts, identifying the collaborators and with links to the relevant chapters :

- Chapter 3, was undertaken in collaboration with Dr Kieron O'Hara
- Chapter 4, was undertaken in collaboration with Dr Mike O Jewell and Dr Faith K Lawrence
- Chapters 5, 6, and 7 were undertaken with Dr Antonis Loizou and Dr David P Dupplaw

Additionally, the author confirms that:

- This work was done wholly while in candidature for a research degree at the University of Southampton.
- Where the author has consulted the published work of others, this is always clearly attributed.
- Where the author has quoted from the published work of others, the source is always given. With the exception of such quotations, the thesis is entirely the author's own work.
- The author has acknowledged all main sources of help (see list above)

1.5 Publications

O'Hara, K., Tuffield, M. and Shadbolt, N. (2009) **Lifelogging: Privacy and Empowerment with Memories for Life**. *Identity in the Information Society*, 1 (2).

Loizou, A, Tuffield, M. M., Dupplaw, D. P., Lewis P. H., Shadbolt N. R., Dasmahapatra, S. (2009) **Context and Content in Multimedia Knowledge Management: An OpenKnowledge Approach (*Submitted*)**. *IEEE Transaction on Multimedia Special Issue on Integration of Context and Content for Multimedia Management*, Jan 2009.

O'Hara, K., Tuffield, M. M., Shadbolt, N. R. (2008). **Lifelogging: Issues of Identity and Privacy with Memories for Life**. *In The First International Workshop on Identity in the Information Society*, May 2008, Arona Italy.

Dupplaw, D., Croitoru, M., Loizou, A., Dasmahapatra, S., Lewis, P., Tuffield, M. and Xiao, L. (2007) **Multimedia Markup Tools for OpenKnowledge**. *In 1st Workshop on Multimedia Annotation and Retrieval enabled by Shared Ontologies*, 5th December 2007, Genoa, Italy.

Potter, S., Kalfoglou, Y., Alani, H., Bachler, M., Buckingham Shum, S., Carvalho, R., Chakravarthy, A., Chalmers, S., Chapman, S., Hu, B., Preece, A., Shadbolt, N., Tate, A. and Tuffield, M. (2007) **The Application of Advanced Knowledge Technologies for Emergency Response**. *In 4th International Information Systems for Crisis Response and Management (ISCRAM 07)*, May 2007, Delft, The Netherlands.

Tuffield, M., Harris, S., Dupplaw, D. P., Chakravarthy, A., Brewster, C., Gibbins, N., O'Hara, K., Ciravegna, F., Sleeman, D., Wilks, Y. and Shadbolt, N. R. (2006) **Image annotation with Photocopain**. *In First International Workshop on Semantic Web Annotations for Multimedia (SWAMM 2006) at WWW2006*, May 2006, Edinburgh, United Kingdom.

Tuffield, M. M., Dupplaw, D. P., O'Hara, K., Shadbolt, N. R., Chakravarthy, A., Brewster, C., Ciravegna, F. and Wilks, Y. (2006) **Photocopain - Annotating Memories For Life (*Poster*)**. *In Memories for Life: The Future of our Pasts*, 12/12/06, British Library.

Tuffield, M. M., Loizou, A., Dupplaw, D., Dasmahapatra, S., Lewis, P. H., Millard, D. E. and Shadbolt, N. R. (2006) **The Semantic Logger: Supporting Service Building from Personal Context**. *In Capture, Archival and Retrieval of Personal Experiences (CARPE) Workshop at ACM Multimedia*, Oct 06, Santa Barbara.

Lawrence, K. F., Tuffield, M. M., Jewell, M. O., Prugel-Bennett, A., Millard, D. E., Nixon, M. S., schraefel, m. c. and Shadbolt, N. R. (2005) **OntoMedia - Creating an Ontology for Marking Up the Contents of Heterogeneous Media**. *In Ontology Patterns for the Semantic Web ISWC-05 Workshop*, Nov 05, Galway, Ireland.

Jewell, M. O., Lawrence, K. F., Tuffield, M. M., Prugel-Bennett, A., Millard, D. E., Nixon, M. S., schraefel, m. c. and Shadbolt, N. R. (2005) **OntoMedia: An Ontology for the Representation of Heterogeneous Media**. *In Multimedia Information Retrieval Workshop (MMIR 2005) SIGIR*, 08/05, Brazil.

Tuffield, M. M., Millard, D. E. and Shadbolt, N. R. (2005) **Towards the Narrative Annotation of Personal Information and Gaming Environments**. *In Hypertext 05, Workshop on Narrative, Musical, Cinematic and Gaming Hyperstructure*, 09/05, Salzburg, Austria.

Tuffield, M. M., Shadbolt, N. R. and Millard, D. E. (2005) **Narrative as a Form of Knowledge Transfer: Narrative Theory and Semantics**. *In 1st AKT Doctoral Symposium*, June 2005, Milton Keynes, UK.

Tuffield, M. M., Millard, D. E. and Shadbolt, N. R. (2006) **Ontological Approaches to Modelling Narrative**. *In 2nd AKT DTA Symposium*, January 2006, Aberdeen University, UK.

Tuffield, M. M., Millard, D. E. and Shadbolt, N. R. (2006) **An Investigation into Automatically Captured Autobiographical Metadata, and the Support for Autobiographical Narrative Generation (*Mini-thesis: PhD upgrade report*)**.

Tuffield, M. M., Millard, D. E. and Shadbolt, N. R. (2005) **Narrative as a Form of Knowledge Transfer, Narrative Theory and Semantic: Present Challenges - Future Possibilities (*9 month Report*)**.

Chapter 2

Background & Related Work

This thesis presents a framework for the capture and archival of personal information, in the form of Semantic metadata detailing one's interactions with their digital hardware. The information gathered and supported by the framework is then used to contextualise and enumerate the various forms of personal information which people are posting to the public domain. In an age where more and more of a person's identity is being pushed into the *digital* realm, one must be aware of their *digital persona*. A person's *digital persona* is presented as the summation of a user's *digital traces*. In the context of this thesis the terms *digital footprint* and *digital persona* will be used to refer to the accumulation of a person's various digital, traces which are generated as they interact with their computer or their peers via the medium of the Internet. The capture of an autobiographical Lifelog is presented as a form of empowering users when it comes to knowing what information about them can be gathered to approximate their *digital persona*. In an age where employers and one's peers potentially have access to a wealth of knowledge about someone, the notion that one's *digital persona* is an asset and worth being aware of is presented as a key motivating factor of this line of research.

2.1 The Semantic Web and Personal Metadata

The work carried out has built upon a number of Semantic Web enabling technologies (see section 5.2). The information identified and captured as a result of our metadata acquisition system is represented by models adhering to as many W3C[1] recommendations as possible. This decision was taken to ensure that the metadata published would be exploited by as many web services as possible, building on the Semantic Web Vision (Berners-Lee et al., 2001) of interoperable web accessible resources. Based on the attempts to help fulfill the Semantic Web Vision, the Resource Description Framework

[1]The World Wide Web Consortium (W3C): http://www.w3c.org/

(RDF)[2] has been adopted as this work's knowledge representation language. Key to the development of the framework was the use and adoption of a SPARQL[3] (World Wide Web Consortium, 2005) compliant triple store[4].

The main requirements in selecting an appropriate triple-store implementation were efficiency and consistency. 3store is a system bench-marked against other RDF storage and query engines such as Jena (McBride, 2001), Sesame (Broekstra et al., 2002) and Parka (Stoffel et al., 1997) and shown to outperform in terms of both efficiency and scalability (Streatfield, 2005; Lee, 2004). Since the creation of the Semantic Logger system 3store's successor 4store (Harris et al., 2009) has been open-sourced and available to the public. 4store has made significant progress in the scalability front, allowing for the storage of up to 15 Giga triples as apposed to 3store's upper limit of 150 million triples. Insight will be presented at the end of the chapter to how much personal data we could potentially store given the bleeding edge of scalable RDF storage technology. 4store was released after the work was undertaken, but would have been the triple-store of choice if it was available at the time.

The following figure 2.1, known as the "Semantic Web Layer Cake" is used to illustrate the technology stack associated with the Semantic Web vision.

This work integrates a variety of tools (see section 5.2) from the Advanced Knowledge Technologies project[5] (AKT), a six-year interdisciplinary collaboration working in the general area of technology and infrastructure development for the Semantic Web. AKT's aim is to extend and develop integrated technologies and services, using the SW as a unified basic framework, for acquiring, retrieving and publishing content over the whole of the knowledge life cycle.

This research is presented as a means to populate the Semantic Web with personal metadata, by exposing information in a structured and standard form, i.e. by using RDF accessible through SPARQL endpoints (World Wide Web Consortium, 2005), via the SPARQL-Protocol[6].The system uses a Universal Resource Identifier (URI[7]) to point to a user's Friend of a Friend (FOAF)[8] file, and subsequently adopts the given user's `foaf:Person`[9] URI was a WebID[10]. In the case that the user does not have a FOAF file the system presented helps the user generate a basic one. Each users FOAF file serves as a unique identifier for their RDF data. The user's FOAF URI is employed to log

[2] W3C's RDF resource `http://www.w3.org/RDF/`
[3] SPARQL RDF query language `http://www.w3.org/TR/rdf-sparql-query/`
[4] Triple store description on Wikipedia `http://en.wikipedia.org/wiki/Triplestore`
[5] The AKT project: `http://www.aktors.org/`
[6] SPARQL Protocol `http://www.w3.org/TR/rdf-sparql-protocol/`
[7] Names and Addressing `http://www.w3.org/Addressing/`
[8] Friend of a Friend: `http://www.foaf-project.org/`
[9] The notion presented here is an abbreviation for the class `http://xmlns.com/foaf/0.1/Person`, where `foaf:` is the FOAF namespace of `http://xmlns.com/foaf/0.1/`
[10] Description of a WebID `http://esw.w3.org/topic/WebID`

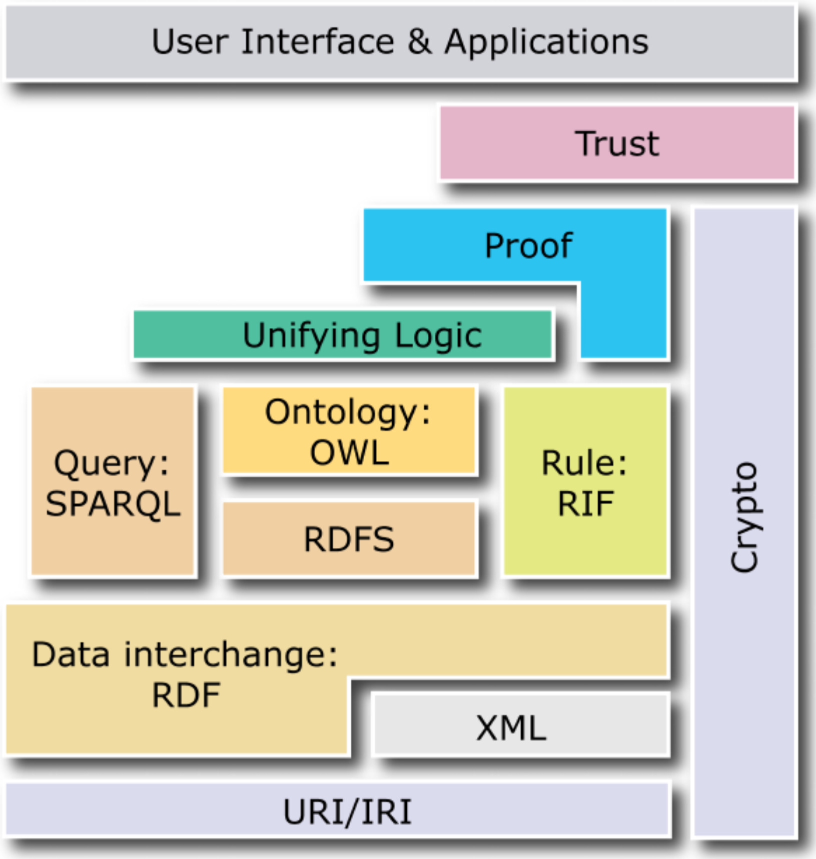

FIGURE 2.1: The Semantic Web Layer Cake: http://www.w3.org/2007/03/layerCake.png

the provenance of any RDF generated. For a more detailed description of how FOAF is used in this thesis see sections 5.4 and 7.3.2.

In an attempt to realise some of promises road-mapped by the Semantic Web community: the seamless integration of heterogeneous data, and that of services exploiting existing machine-accessible knowledge (Shadbolt et al., 2006), a decision to create an easy to use infrastructure allowing users to store, update, visualise, and query their own personal knowledge base(s) through the web, seemed a pragmatic course of action and is described in full in chapter 5.

User interface issues arise when designing how potential users may interrogate their contextual log and their annotated photo collections. SPARQL is not the only way envisaged for users to access the knowledge gathered, the incorporation of a faceted browser such as mSpace (m. c. schraefel et al., 2003) (see 5.5), taking heed of the lessons learned from Haystack (Huynh et al., 2002), and Lloyd Rutledge's work on Making RDF presentable (Rutledge et al., 2005), have been applied to help the user navigate through the generated knowledge space. Section 5.5 puts forward work that allows for mSpaces to be generated for any given SPARQL endpoint with arbitrary RDF. The Photocopain (see chapter 6) system also allows for the incorporation of the

AKTive Media (Chakravarthy et al., 2006) user interface to ease the burden of manual annotation. Our initial investigation has suggested that avoiding the use of *deep tree-lists*, as employed by MindSwap's PhotoStuff system (Halaschek-Wiener et al., 2005), and the adoption of a simple web interface similar to that of Flickr, is seen as a step in the right direction.

The work undertaken (see section 5.5) to incorporate the mSpace faceted browser is envisaged to be more than just a user interface for the knowledge bases generated by the Semantic Logger and the Photocopain system. It is presented as a means to perform future evaluations of the personal metadata generated by the two aforementioned systems. Any such future experiment would follow from the experiments undertaken at Stanford[11] by Mor Naaman (Naaman et al., 2004b) which employed the use of the Flamenco faceted browser (Yee et al., 2003) and presented as a contribution to the utility of metadata to support human memory.

2.2 Lifelogging and Autobiographical Knowledge Management Environments

MyLifeBits (Gemmel et al., 2002) and SemanticLIFE (Admed et al., 2004) can be regarded as the modern seminal systems in this area building on the ideas put forward by Vannevar Bush (Bush, 1945) in his Memex device. While these have proven to be a valuable source of inspiration for this project, numerous others have been undertaken under the field of context based system (Cayzer and Castagna, 2005; Heath et al., 2005; Iofciu et al., 2005; Molle and Decker, 2005; Richter et al., 2005; Xiao and Cruz, 2005). The domain of interest of such systems is limited to the publishing, browsing, and sharing of information in tailored knowledge representations, where as this approach aims to employ as many existing models as possible. For further insight into Lifelogging systems and issues surrounding them see sections 3.1 and 3.2. The infrastructure presented (see chapter 5) attempts to cater for any RDF data-structure uploaded to it, promoting flexibility.

The SemanticLIFE project uses a variety of techniques to extract metadata from personal communications, and considers issues relating to provenance and trust; these are presented in Weippl *et al* (Weippl et al., 2004). The SemanticLIFE project focuses on the capture and annotation of a user's 'work related life', where as our focus of interest is not confined to that particular subsection of a person's life.

MyLifeBits aims to organise and retell personal experiences using narratives as a structuring mechanism; (Gemmell et al., 2005) presents valuable insight on how stories can be generated based on information such as location and time. However, neither of

[11] Stanford University `http://www.stanford.edu/`

these aforementioned systems strive to surreptitiously amalgamate as much contextual metadata as identified in this work. The metadata logging system presented below is presented as a low-cost metadata collection system, where as MyLifeBits is presented primarily as a method to to aid the human authoring and annotation of personal experiences.

SemanticLIFE is preoccupied with allowing users to set-up an information repository to provide enhanced querying capabilities, while MyLifeBits introduces the notion of automatically producing annotations by exploiting co-occurring events. This section sets out to identify the principal differences between the Semantic Logger (see chapter 5), an autobiographical contextual log framework, and such previously developed systems.

Such systems have engineered over-ranging knowledge representations to support the functionality they provide. The Semantic Logger makes no attempt to homogenise data that is heterogeneous by nature; this is left for applications that will use the system as a platform, as per their requirements. The rationale is that different mappings will be appropriate for different applications. A caveat worth mentioning is that data has to be in RDF.

The Semantic Logger aims to aggregate as much available personal information into a central knowledge base allowing for context-based systems (Falkovych and Nack, 2006; Davis et al., 2005; Adomavicius et al., 2005; Tuffield et al., 2006a,b) to exploit it as needed. The Semantic Logger is currently being used as a platform for recommender system research, the work is elaborated upon in Antonis Loizou's work (Loizou, 2009), as well as a number of other projects see section 5.6 in chapter 5.

Another development worth mentioning is the NEPOMUK project[12], a European Framework Six Programme[13] funded collaboration of industrial and academic partners and industrial end-users. The project brings together various previous semantic desktop implementations, and focuses on knowledge integration in shared peer-to-peer environments, supporting automated community recognition. Detailed information on this has not been made available, however it seems that the focus is put once again in providing a solid platform for such sharing, rather than the ease of adding services to the system. For a full overview of the state of the art in developments in the field of semantic desktop research the interested reader is pointed to (Sauermann et al., 2005). The output of this work has since been adopted by the open-source KDE [14] graphical desktop manager.

Another important effort in the lifelogging space was *LifeLog* (IPTO 2003), sponsored by the American Defense and Advanced Research Projects Agency (DARPA [15]), which became somewhat notorious, this is touched upon the following chapter of the thesis.

[12]NEPOMUK Project: http://nepomuk.semanticdesktop.org/
[13]EU FP6 call: http://cordis.europa.eu/fp6/
[14]KDE Free Desktop http://www.kde.org/
[15]DARPA http://www.darpa.mil/

Others include Total Recall (Cheng et al., 2004) and SemanticLIFE (Admed et al., 2004). The most famous and comprehensive Lifelogging experiment is the ongoing attempt by Gordon Bell of Microsoft to create a digital archive, MyLifeBits[16].

2.3 Multimedia Asset Management: Image Annotation

Photographs can be viewed as externalised additions to human memory; many human memory management tasks rely on collections of photographs, even if only shoe-boxes of photographic prints. Furthermore, as digital technology has dramatically increased the numbers of photographs taken (it has been estimated that up to 375 petabytes of information is created in photographic form annually), the problems associated with archiving and managing these photographs have become more pressing (describing their content, storing and retrieving them, and developing narratives to link them together). Image management is a highly labour-intensive task, and the Photocopain system (chapter 6) has been designed with the intent of alleviating such burdens. Photocopain is presented as an example multimedia asset management system that utilises the contextual information captured by the Semantic Logger.

The generation of annotations from multimedia content can be seen as an application of image classification of the type used in content-based indexing and retrieval (Smeulders et al., 2000). Systems such as MAVIS2 (Joyce et al., 2000) and SIMPLIcity (Wang et al., 2001) use collections of image classifiers (embodied as agents, in the case of MAVIS2) to add descriptions, in the form of semantic annotations, to images. Photocopain uses a combination of *content* and *context* based services to annotated personal photo-collections. The study of content based image retrieval techniques is still in its infancy, and are far from semantic, for a full discussion on the limitations of content based approach please see (Veltkamp and Tanase, 2000). The approach presented in this thesis combines annotations extraction from the content of an image, with the contextual information in the Semantic Logger, in order to annotate personal photo collections.

A series of experiments have been performed looking to identify the usefulness of various forms of metadata given the task of searching and browsing through personal photo collections. In (Naaman et al., 2004b) an evaluation of a number of different types of metadata is presented with regard to the personal photo search and browsing. Naaman's research aims to provide insight into which metadata categories are useful for mentally recalling and finding photographs, and is presented as a multimedia asset management system. Personal photo libraries usually have little discernible structure, and enriching the individual photos with metadata is presented as a means of organisation. The work captured the following metadata categories to enrich the photo collections: year, time of the day, location, elevation, season, light status, weather status, and temperature.

[16]MyLifeBits http://research.microsoft.com/en-us/projects/mylifebits/

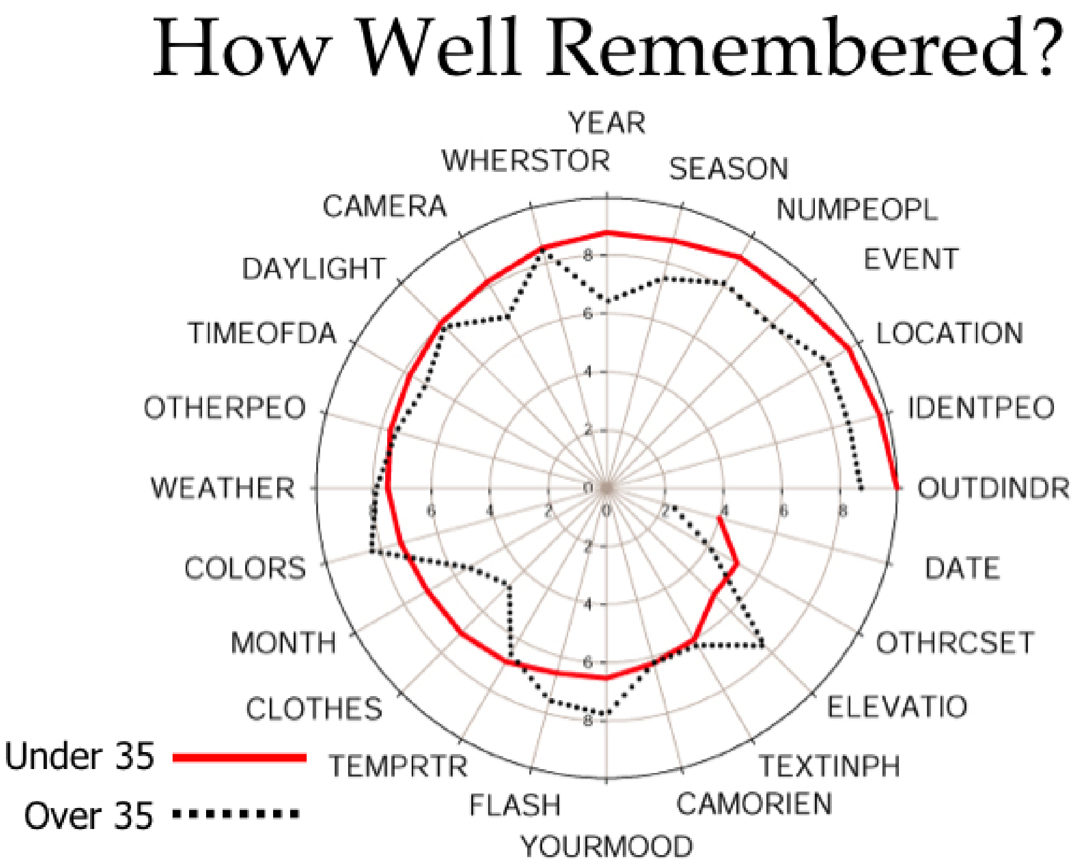

FIGURE 2.2: Mor Naaman's ZoneTag 'How Well Remembered' results

Section 5.3 outlines the sources of information captured, and presents them as complementary to the work undertaken by Mor Naaman. Mor's findings can be illustrated in the following figures 2.2 and 2.3. Where 2.2 shows how well the two age groups studied did at remembering certain facts about specific photos, and 2.3 illustrates which forms of metadata where of most use when trying to find a specific images from a collection of pictures.

Work undertaken by Professor Marc Davis (Davis et al., 2004) proposes three main categories of contextual information: spatial, temporal, and social. The Semantic Logger also attempts capture and archive all three of these forms of information, the infrastructure has been designed to allow for annotations and knowledge to be shared between peers. Work has also been published presenting the various aspects of digital photography (House and Davis, 2005) that of self-presentation, self-expression, memory archival, and that of social documentary. These findings are mimicked in our motivation of supporting human memory management.

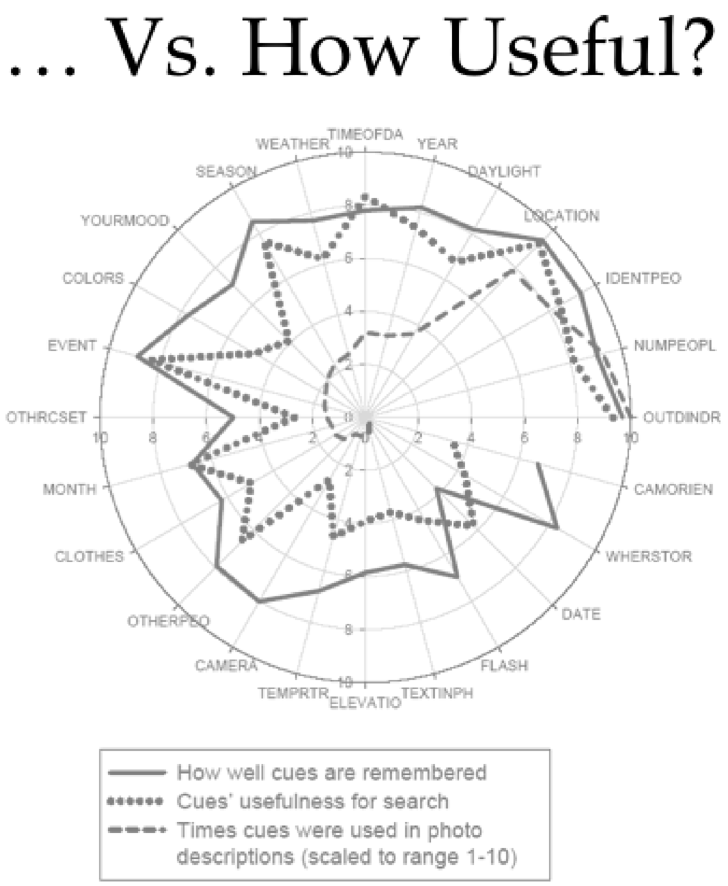

FIGURE 2.3: Mor Naaman's ZoneTag 'How Useful for Recall' results

A specific problem faced by the multimedia knowledge management community, is that of personal multimedia management. The approach described in (Naaman et al., 2004b) aims to tackle the problem of indexing and managing user generated content by acquiring and exploiting contextual metadata pertaining to a specific person. The work details techniques for the capture of contextual data, followed by examples of how such information can be used in our application scenario. There is a focus on taking advantage of Semantic Web technologies, namely the Resource Description Framework (RDF)[17] and SPARQL query language, as well as adopting shared conceptualisations such as those illustrated in the Linking Open Data (LOD) project[18].

Given the current trend in social technologies, whereby users hand over personal information to online services, reaping the benefits of social computing raises also issues of privacy, personal identity, and that of *online personae*. Personal information is being posted in various forms such as, personal-photos on Flickr[19], social relationships

[17]W3C RDF Specification http://www.w3.org/TR/rdf-concepts/
[18]Linking Open Data http://richard.cyganiak.de/2007/10/lod/
[19]Flickr http://www.flickr.com/

on Facebook[20], and time-stamping microblogging such as Twitter[21]. The nature of the private information posted to these social networking sites does not seem to deter people from using them, as the network effect provides the necessary added value. The information posted to each of these sites in isolation may not seem that sensitive in nature, but by aggregating such data information pertaining to a user's interests and day-to-day activities may become more transparent. This thesis proposes that the aggregation of personal information hosted on various social networking sites will empower users to better understand their online presence, educating them with respect to what information about them is posted in the public domain.

2.4 Multimedia Ontologies

The following ontologies are presented as relevant to the engineering of the OntoMedia ontology (see chapter 4).

2.4.1 The Functional Requirements for Bibliographic Records Model

The existence of textual evidence in cultural heritage had lead to work on mapping the CIDOC CRM to the Functional Requirements for Bibliographic Records (FRBR) model. FRBR was created by the International Federation of Library Associations and Institutions to "produce a conceptual model that would serve as the basis for relating specific attributes and relationships (reflected in the record as discrete data elements) to the various tasks that users perform when consulting bibliographic records" (Saur, 1998).

FictionFinder[22] is a prototype system, which applied the FRBR model to the Online Computer Library Center (OCLC) WorldCat database. This system is notable for allowing search on fictional characters and imaginary places as well as the more common author, setting, genre, summary, title, and subject. The Columbus Metropolitan Library Fiction Finder[23] has a similar although less developed system, in so far that only some of their records have metadata about characters associated with them. This information on what might be seen as content rather than bibliography is taken, in the case of the OCLC system, from the section of the FRBR model referred to as the Group 3 Entities. These are `Concept`, `Object`, `Event` and `Place` and relate to the FRBR object Work through the `has-as-subject` relationship. A work may also have as a 'subject' entities from Group 1 (other media objects) and Group 2 (people and corporations).

[20]Facebook http://www.facebook.com/
[21]Twitter http://twitter.com/
[22]Project url: http://www.oclc.org/research/projects/frbr/fictionfinder.htm
[23]Library Fiction Finder: http://www.columbuslibrary.org/cmlradv/browse2.cfm

The Group 3 entities defined in the FRBR can be considered equivalent to the top level of the OntoMedia (see section 4.3) ontology using the following mapping:

FRBR Group 3 Entities	**OntoMedia Core Classes**
Concept	Abstract-Item
Object	Physical-Item
Event	Event
Place	Space

Expending this mapping, if we consider the FRBR Groups 1 and 2 as if they appeared as subjects in a narrative the mapping could be made as below:

Group 1 Entities As Subject	**OntoMedia Core Classes**
Work	Context
Expression	Abstract-Item
Manifestation	Collection/Physical Item
Item	Physical-Item
Group 2 Entities As Subject	**OntoMedia Being Classes**
Person	Being/Character
Corporate Body	Organization

Where the FRBR approach differs from OntoMedia is the meaning with which these ‘subjects’ are imbued. The FRBR has no model of time or narrative flow. For example, the attribute for an event is simply the term used for that event, i.e. ‘the Second World War’. It is an identifying label rather than an object with its own meaning. OntoMedia expands on this metadata as it does with the CIDOC CRM to allow exploration of the events and entities which the media object contains within in conceptual framework. From this, the subject as defined by its FRBR can be directly drawn or inferred. Future research may even allow such bibliographic categorisations such as genre and summary to be suggested if not generated by querying of the OntoMedia data.The CIDOC CRM, FRBR, and OntoMedia models work as complementary vocabularies. They overlap enough to be mappable between each other, and the differences in their scopes and strengths implies that they best be applied for different purposes and subsequently linked through mapping.

2.4.2 The ABC Ontology

In designing the OntoMedia (see section 4.3) ontology a variety of existing techniques for media annotation were taken into consideration. Of particular note is the ABC Ontology by Lagoze and Hunter (Doerr et al., 2003; Lagoze and Hunter, 2001). This was developed as part of the Harmony international digital library project with an aim to provide a level of interoperability between existing metadata ontologies, primarily for the cataloguing community. The OntoMedia Core is based on similar principles, particularly the separation of spatial and temporal classes. OntoMedia further adds the capability for trait and motivation representation, which ease the annotation of attributes and intent to provide specializing classes for application to fiction. Furthermore, Hunter proposes a technique to represent MPEG-7[24] using a DAML+OIL representation, whereas we make use of a VLit[25] adaptation combined with a geometry ontology to reference sections of source media.

2.4.3 The Action Ontology

A further related ontology is that of Action (Feinberg and Shaw, 2004), a taxonomy focusing on the representation of events. As such, it highlights the physical effects of events, the activities involved, and more abstract characteristics such as style. OntoMedia carries several of these ideas into the Event class, which may have preconditions and postconditions to describe the causes and effects. Furthermore, this has the capability of being combined with the motivation attributes that are assignable to participants to infer more information regarding the intent of those involved.

2.4.4 The Event Ontology

The Event Ontology[26] is presented as an ontology for the representation of chronological events. The Event Ontology's primary intention was to be used was a method of representing the succession of events in a piece of music (Raimond et al., 2007), but has subsequently been adopted by the community as a generic, lightweight method of representing chronological phenomenon in RDF. Figure 2.5, illustrates the way that events are represented in the ontology.

The above example (figure 2.4), taken from the Event Ontologies description page[27] is an RDF document describing a performance, including a performer, and a Santur in London on 15/10/2007.

[24] MPEG-7 is an ISO/IEC standard developed by MPEG (Moving Picture Experts Group) `http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm`

[25] VLit Transclusion Support `http://www.eprints.org/documentation/tech/php/vlit.php`

[26] Event Ontology Description `http://motools.sourceforge.net/event/event.html`

[27] The Event Ontology :`http://motools.sourceforge.net/event/event.html`

```
@prefix event: <http://purl.org/NET/c4dm/event.owl#>.
@prefix mit: <http://purl.org/ontology/mo/mit#>.
@prefix foaf: <http://xmlns.com/foaf/0.1/>.
@prefix tl: <http://purl.org/NET/c4dm/timeline.owl#>.
@prefix xsd: <http://www.w3.org/2001/XMLSchema#>.

:performance
a event:Event;
event:factor mit:Santur;
event:agent [
a foaf:Person;
foaf:name "P. H.";
];
event:place <http://sws.geonames.org/2643744/>;
event:time [
a tl:Interval;
tl:at "2007-10-15T12:00:00"^^xsd:dateTime;
tl:duration "PT1H"^^xsd:duration;
];
.
```

FIGURE 2.4: Example event using the Event Ontology

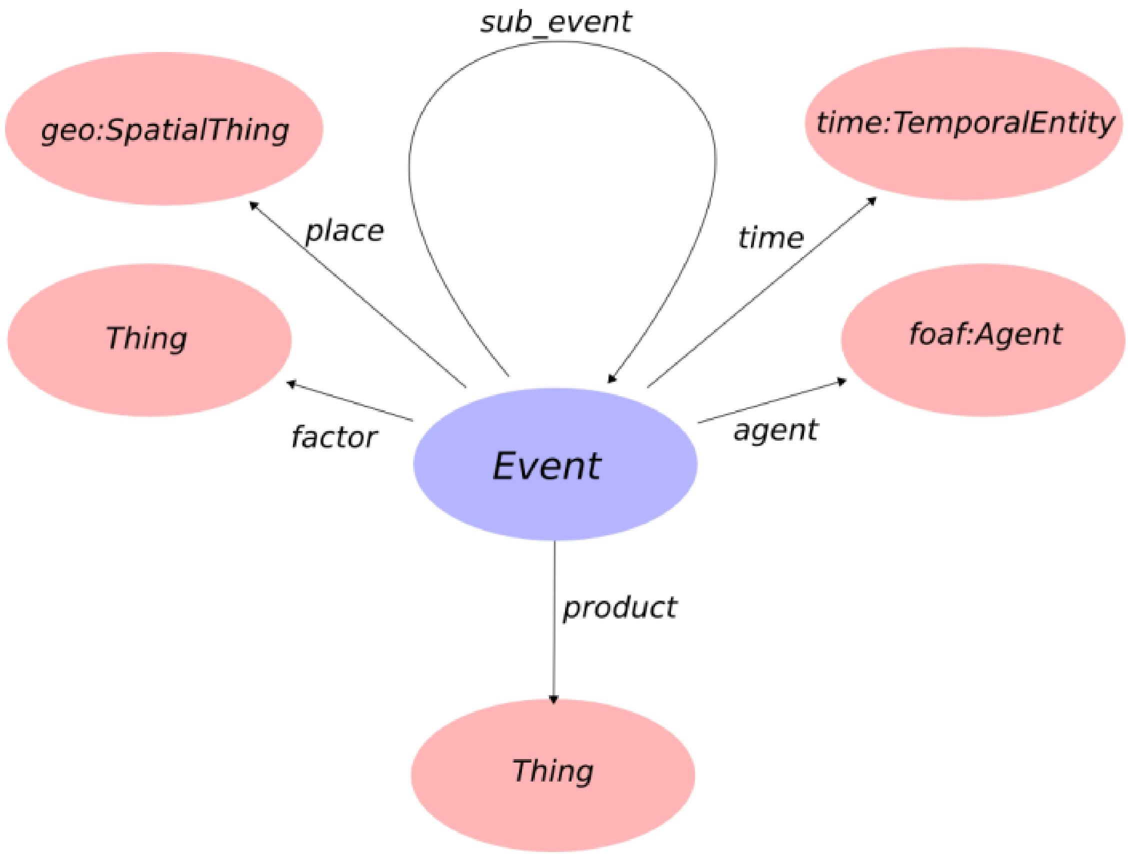

FIGURE 2.5: Yves Raimond's Event Ontology

2.5 Narratives and the Semantic Web

Throughout human history narratives, in one form or another, have been used as a mechanism of transferring and propagating knowledge through society, and its subsequent generations. Insight into the evidence of narrative use, in its many forms, and the evolution of narratives across different mediums can be found in Joseph Campbell's book 'A Hero with a Thousand Faces' (Campbell, 1949).

From a very young age, people are exposed to narratives. They tell stories, whether real or made-up, to teach children about the intricacies of life. The traditions of oral storytelling that have evolved into our contemporary modes of narrative have been recognised as core to the transfer of knowledge within society (Campbell, 1949). Work has been undertaken to illustrate the transfer of knowledge within societies through the use of narrative mediums and the similarities between these differing modes (Campbell, 1949; Murray, 1998; Bal, 1997). Research into the ways that we make sense of our world have resulted in the term *Narrative Intelligence* (NI) being coined, describing how humans organise events into more-or-less familiar narratives (Blair and Mayer, 1997).

NI has been identified as one of the main synergies around which Artificial Intelligence (AI) research into narrative has been brought together (Mateas and Sengers, 1999). At the time when this work was done Semantic Web technologies were still in their infancy, the remainder of this chapter presents how advances in knowledge representation and the enrichment of information through explicit semantics can be utilised to develop meaningful narratives. For further information regarding the role which narratives have to play see chapter 4.

2.5.1 Narrative Theory for the Semantic Web

Narratives and the study of narratology, have been a central theme of the social sciences for a long time and have recently become increasingly popular in the field of knowledge technologies (Alani et al., 2003b; Geurts et al., 2003; Little et al., 2002; Silva and Henderson, 2005; Gemmell et al., 2003). Narratology has focused on representing and defining one of the core modes of human communication. As a result of these efforts, there exists a growing set of different narrative paradigms, or ways of conceptualising narrative spaces, such as formalism, structuralism, post-structuralism, and post-modernity (Schärfe, 2004). Chapter 4 presents work undertaken attempting to model narrative structures in RDF. The aim of the work presented in chapter 4 is not to try and develop a new paradigm or even to create a taxonomy of narrative approaches, this is best left to narrative theorists, but to learn from and harness existing methods to aid the representation of narrative structures on the Semantic Web.

There are a growing number of areas where strategies and viewpoints put forward from narratology have been adopted and utilised within computer science research, such as dynamic multimedia presentations (Geurts et al., 2003; Little et al., 2002), summarisation (Alani et al., 2003b), and storytelling and interactive drama (Murray, 1998). One example of an adopted theory is Bal's (Bal, 1997) view of narrative (Little et al., 2002) (for detailed discussion see section 4.2.2). In abstract terms this states that narrative can be viewed to consist of 3 layers, the lowest being the *Fabula*, which represents the raw chronological events; the *Story*, where given a fabula one could derive a number of different stories, and at the third and highest level the *Narrative*. The narrative is said to be the final form of the rendered material. The reason why a distinction is made between the *Narrative* and the *Fabula* level is due to the fact that the collection of facts from the *Fabula* arranged into a given *Story* by the author, or authoring system, only becomes a *Narrative* once consumed by a reader. This final understanding of the collection of events presented to the consumer is said to be the *Narrative*, this is described in further detail in section 4.2.2 of this thesis.

2.5.2 The Analogy: SW meets Narratology

Semantic Web enabling technologies (see section 2.1) allow for the annotation of multimedia items, resulting in a corpus of available "knowledge nuggets".

This collection of meaningful knowledge is presented as the *fabula* (section 4.2.3); story grammars, that are most commonly implemented in the form of templates (Alani et al., 2003b), present the structural design used to portray the desired *story*, and the resulting output, whether the form be textual or a composition of a variety of multimedia is presented as the *narrative*.

Given the above analogy, the major shortfall of existing narrative generation, for dynamic presentations, are the use of story-grammars. These static structures have to be defined by developers before the deployment of a system. This method of pre-empting the *story* of the narratives will limit a system's ability to discover any new, and possibly previously unknown, relationships to render a narrative. The next section will propose one method of enhancing the current template-based methods for two differing, but not disjoint domains.

2.5.3 Memories for Life

Memories for Life (see section 3.4) is being discussed as a grand challenge for UK computing. It aims to address the challenges of storing autobiographical knowledge in the form of multimodal electronic media, and to identify any issues that may arise from such a situation. We believe that SW technologies could be adopted to help realise the potential of such a vision. Given a system that could store a comprehensive collection of a lifetime's worth of acquired electronic media a unified method of marking up this inherently heterogeneous data-set is needed. A vocabulary of terms and their relationships are presented as a means to annotate these "memory nuggets", to encapsulate the semantics of these autobiographical memories. The OntoMedia ontology is a possible candidate. This ontology has been designed to allow the mark-up of literature, film, and other forms of narrative at the *fabula* level.

The Advanced Knowledge Technologies (AKT) project has produced tools to automate the annotation of textual media, (e.g. Armadillo (Ciravegna et al., 2004)), as well as a framework for semi-automating annotating photographs (Tuffield et al., 2006a). These tools and techniques are being applied to as many of the readily available sources of metadata as possible to create as rich a memory bank as possible while keeping the cognitive overhead to a bare minimum.

Chapter 3

Motivating a Digital Autobiographical Log

This chapter of the thesis will present the motivations behind the design and realisation of a framework for the capture and archival of autobiographical data. As technology advances and the Internet becomes more ubiquitous in nature, a greater proportion of our daily lives is being moved into cyberspace. Discussions and concerns about how more and more of our identity is being pushed online is presented and put forward as a key motivating factor for this thesis. The notion of *lifelogging* is then outlined as a method of empowering people to understand what information they are leaving behind in the form of *digital traces*. Within the scope of this thesis the term lifelogging is used to denote the proactive collection of personal information for personal gain and consumption. Discussions around the notion of privacy and identity in our networked world are outlined and presented as concepts which people should be educated about in order to promote the active safeguarding of peoples' now networked identities. This is then followed by a section on how we can take the principles of lifelogging and use them to build up a collection of autobiographical facts about one's life, in an extensible manner, which can be queried retrospectively to answer specific questions. Motivations to why people may want to collect their own autobiographical log are presented and are contextualised with respect to concepts of identity, and privacy in an ever more online world.

The Internet has enabled immediate access to a wealth of information changing the problem from one of a lack of access to pertinent data/information about a particular topic to that of trying to find relevant information from an ever expanding pot of heterogeneous data. This problem is often referred to as *infosmog* (Shadbolt and O'Hara, 2003), and discussion follows to how the Semantic Web vision aims to tackle the problem of organising and retrieving structured data on a web-scale. Insight is then provided into the other key drivers for adopting the technologies presented in this thesis. Following a discussion of how the Semantic Web aims to tackle the problem of *infosmog* are further justifications for using the aforementioned technology (see section 2.1) in terms of why one should adopt an open, non-proprietary data format given the intention of longevity and archival.

Further to the concepts of empowerment and personal identity, the notion of being able to make use of autobiographical facts to enrich one's personal multimedia collection is presented as another key motivating factor to the study at hand. As digital storage costs decrease and digital capturing devices such as cameras, video recording devices, and location tracking hardware, to name but a few, become cheaper and can be seen to be slowly merging into a single portable device [1], we have witnesses a phenomenal rise in the generated of digital multimedia. The motivations behind capturing a personal lifelog to be used to add context to a personal multimedia collection is presented, followed by a discussion around how personal multimedia items are moving out of shoe-boxes and are being situated in the social space that is the Internet.

This chapter then outlines the activities of a few key research initiatives which have informed and inspired the work undertaken. Here discussion will focus on the proposed impact of the listed research initiatives, and how the author thinks they will influence concepts such as identity, and privacy in an online world rich of personal information.

Finally, this chapter is rounded off with use cases aimed to illustrate the potential benefits of capturing an autobiographical log.

3.1 Lifelogging and the Disappearance of the Body

This section will examine the technologies behind Lifelogging, and discuss the concerns relating to the capture and storage of privately held autobiographical logs. Many of the concerns surrounding the capture of the entirety of a human's life have been discussed within the context of the "Memories for Life" UK Grand Challenge for Computer Science (M4L) which is presented towards the end of this chapter (see 3.4). It is also clear that there are likely to be privacy concerns about Lifelogging practise and technology (Allen, 2008). The structure of the section is as follows. The following section will set out some of the principles of Lifelogging, and highlight some of the more prominent efforts in this space (further detail can be found in the background section of this thesis 2). Then I will examine pragmatic aspects, such as the sorts of information that Lifelogging is likely to draw upon, the uses to which it might be put and so on. Next, I shall examine some of the issues surrounding privacy, identity, and empowerment for the Lifelogger.

[1] Apple's iPhone `http://www.apple.com/iphone/`, or Nokia's N95 `http://www.nokia.co.uk/find-products/all-phones/nokia-n95` are two perfect examples of such devices

3.1.1 Lifelogging: An Introduction

As more of our transactions take place remotely, the need for our physical presence has declined, and increasingly many interactions are carried out by digital avatars of ourselves, in a general trend termed by sociologists 'the disappearance of the body' (O'Hara et al., 2006, p. 1-24). Hence our daily lives leave behind more evidence that we can collect and curate, and the indiscriminating collection of such evidence is what we call Lifelogging. Lifelogging can be passive – one stores the by-products of the life one would have lived anyway, or active, one surrounds oneself with sensors and information capture tools to create as rich a picture of one's life as possible.

Lifelogging has been the subject of investigation as long as it has been clear that digital technologies were going to be playing a large part in our lives. Steve Mann has spent a long time investigating wearable computing, in particular wireless cameras to record his daily existence, and has enjoyed media status as 'the world's first cyborg' (Mann and Niedzviecki, 2001). Jennifer Ringley also achieved a level of celebrity and indeed notoriety when she set up JenniCam, a webcam that recorded events in her living space, and made available on the Internet images of events in her daily life ranging from the mundane to the pornographic (Jimroglou, 1999). Initially she filtered out very private moments, while also sometimes 'performing' for the camera; but for most of the experiment, which ran from 1996-2003 the images were unfiltered. At its height, JenniCam attracted four million viewers daily and is said to be the beginnings of reality television.

One important effort in the lifelogging space was *LifeLog* (IPTO 2003), sponsored by the American Defense and Advanced Research Projects Agency (DARPA [2]), which became somewhat notorious, see section 3.1.2. Others include Total Recall (Cheng et al., 2004) and SemanticLIFE (Admed et al., 2004). The most famous and comprehensive Lifelogging experiment is the ongoing attempt by Gordon Bell of Microsoft to create a digital archive, MyLifeBits[3].

MyLifeBits, also referred to as Gordon Bell's surrogate brain, has provided some of the tools needed to compile a lifelong digital archive. It has been found that digital memories allow one to vividly relive an event with sounds and images, enhancing personal reflection in as much as the Internet has aided scientific investigations. Every word one has ever read, whether in an e-mail, an electronic document, or a Web based resource, can be found again with just a few keystrokes. Computers can analyse digital memories to help with time management, pointing out when you are not spending enough time on your highest priorities. Your locations can be logged at regular intervals, producing animated maps that trace your peregrinations. Perhaps most important, digital memories can enable people to tell their life stories to their descendants in a compelling, detailed fashion that until now has been reserved solely for the rich and famous (Bell

[2]DARPA http://www.darpa.mil/
[3]MyLifeBits http://research.microsoft.com/en-us/projects/mylifebits/

and Gemmell, 2007, p. 40-42). This ability to preserve a digital account of a person's life in the form of a legacy has started to become more commonplace whereby commercial offers such as Webwill[4] and FuneralOne[5] are attempting to market the digital identities of the deceased as valuable assets in our ever more online world. This notion of preserving a person's online persona has also started moving into the mainstream whereby social networking providers are starting to implement policies for converting profile pages to testimonial pages given a user's passing. This is illustrated by Facebook's testimonial[6] service that caters for a profile page being turned into a moderated testimonial of a loved one. Given that social networking is a new phenomenon there are no best practises for what should happen to data on a social network given the passing of a user. Facebook's approach differs drastically from the manner in which Last.fm[7], the social radio station, handles data belonging to the deceased. Last.fm presents its users with music recommendations which are made based on the listening habits of other users with similar music taste. Given that Last.fm is a community effort, and in theory gets better as more and more people partake in the community effort of sharing their musical preferences Last.fm do not offer an option to delete/remove data pertaining to the deceased upon their passing. Whereas others sites such as Facebook have policies in place to remove information regarding a deceased user.

MyLifeBits began in 2001, but its roots were in Bell's attempt from 1998 onwards to go completely paperless, where he scanned all documents (including logos on freebie conference mugs and t-shirts), videoed lectures, voice tapes etc, a high-cost strategy that required full-time personal assistants. MyLifeBits was then set up to enable Bell to make sense of the enormous repository of information, to support querying, retrieval and search. MyLifeBits also provides new capture tools, automatically recording telephone calls he makes and television programmes he watches. It copies every web page he visits, transcribes every instant message he sends or receives, records the files he opens, the songs he plays, the searches he performs, the windows which are in the foreground of his computer at any time and the movements of the mouse, and uploads his position from the Global Positioning System (GPS) tracker that he wears constantly (Bell and Gemmell, 2007, p. 45). Out of all of the innovations brought forward by Gordon Bell's endeavours on the MyLifeBits project is the development of SenseCam (Hodges et al., 2006). SenseCam is a wearable camera with a wide-angled lens that periodically take pictures of what the user in seeing.

Unsurprisingly he has a large archive, of about 150 gigabytes, which he has used for tracing people he wants or needs to contact, finding citations for his academic papers, providing his medical history to his doctor and finding material for an obituary of a friend. MyLifeBits has also been important for experimentally discovering what extra

[4]Webwill http://www.webwill.se/en/
[5]FuneralOne http://www.funeralone.com/
[6]Facebook Testimonial Pages: http://www.facebook.com/help/contact.php?show_form=deceased
[7]Last FM http://www.last.fm/

tools are important for commercialising Lifelogging, including face recognition to help with annotating photographs, speech-to-text software for transcribing or searching telephone calls and cleverer search and retrieval to classify documents in advance (Bell and Gemmell, 2007, p. 45).

3.1.2 Lifelogging: Discussion

As Lifelogging tools are currently a matter for scientific research rather than commercial application, there is a possibly brief window for reflection about privacy issues. Certainly the activity is rare enough that legislation is hardly required now. There have been some unrealistic suggestions, such as that of Dodge and Kitchin to program imperfections into the system, deliberate error that would prevent the Lifelog from being veridical, and therefore invading privacy (Dodge and Kitchin, 2007), and some possibly unfounded optimism, as with Cheng and colleagues, who are sanguine about the use of Lifelog data by the judicial system (Cheng et al., 2004). This notion of being able to anonymise data gathered off of the Internet has been shown to be ineffective due to the social nature of interactions on the web. A researcher at Microsoft Research (Xie et al., 2009) suggests techniques for de-anonymising identifiers found on the internet and reports high levels of accuracy. Furthermore, research has been undertaken to evaluate whether current techniques of data anonymisation employed by social networking providers in order to facilitate the trade/sale of their user-data are at all effective. Research undertaken by Narayanan & Shmatikov (Narayanan and Shmatikov, 2009) suggests that the anonymised userdata which companies end up sharing with advertising companies, data-mining experts, and application developers, can be de-anonymised with high levels or accuracy. They showed that they could de-anonymise Twitter users with less than 12% error, and subsequently showed how they could perform a similar job on the Netflix prize data (Narayanan and Shmatikov, 2006), and finally generalised their work into a framework for de-anonymising users in sparse data sets (Narayanan and Shmatikov, 2008).

There are certain ethical parameters that suggest themselves, as listed by Anita Allen (Allen, 2008) that no one should be required to keep a Lifelog. No one should be suspected for not keeping a Lifelog. Personal Lifelogs should be deemed the property of the person or persons who create them. No one should record or photograph others for a Lifelog without consent of the person or their legal guardian. A counter-technology to block Lifelog surveillance should be designed and marketed along with Lifelogging tools. The owner/subject of a Lifelog should be able to delete or add content at will. No one should copy a Lifelog or transfer a Lifelog to a third party without the consent of its owner. Tools to gather data will have obvious applications for surveillance, and this is a worry, as evinced by the fate of the DARPA LifeLog project. It is essential that one is not forced (either by law or social pressure) to keep tabs on oneself. But discussions about privacy and Lifelogging have made unwarranted assumptions which skew the privacy arguments. Many presume that the data is necessarily personal, though it may be public domain. And why think that lifelogged information should be kept private? Especially given that the publication of the data might just be the point of capturing it in the first place.

There are also practical considerations that need to be addressed, the most obvious of which is that we cannot easily control the appearance of others in our photographs. And furthermore, it is important not to develop solutions that apply to outdated technologies. For instance, large amounts of an individual's information are migrating from their personal standalone devices, and are being stored instead in what is becoming known as 'the cloud', a set of computing resources operated by a third-party provider (such as Google[8]) and located in data centres, while the user is unconcerned with the actual technology being used, and instead hires the information storage or processing service. Given that the current trend of employing 'cloud-based' computing solutions, research has been undertaken at Cambridge University whereby the terms of conditions of 40 popular social networking sites were examined (Bonneau and Preibusch, 2009), and defined behaviour was compared to the actual behaviour of the sites. One of the findings presented in the "Privacy Jungle" (Bonneau and Preibusch, 2009) was that many of current social networking sites use what are known as content delivery networks, for example Facebook use Akamai[9] to ensure that content is cached all over the world so that content can be accessed at high-speed regardless of the user's geo-location. The cloud-based solution to the problem of content delivery is the main reason why content can remain on Facebook for up to 30 days if and when it is deleted by a user for it has to be flushed from a whole number of "content- delivery caches" around the world. With respect to removing/deleting information Facebook's privacy policy[10] states that "Removed information may persist in back-up copies for a reasonable period of time but will not be generally available to members of Facebook.", the ambiguity in the statement lies with respect to what is said to be a "reasonable period of time" and is almost certainly brought about as a result of the cloud-based solution adopted by Facebook.

In this more nuanced environment, Lifelogging tools can provide the user with somewhat more control, gathering together his or her appearances in the public domain. Knowing what can be seen is an important first step in the preservation of privacy, and it is important that restrictions or distortions are not brought into Lifelogging tools that prevent this empowerment of the user. One's online identity has public and private aspects, and focus on the private aspects will be misleadingly partial. Furthermore, the social aspects of the Web are where its interest lies for many people. Restrictions on their ability to construct their own identity would be a severe curtailment of their

[8]Google Inc. `http://www.google.com/`
[9]Akamai `http://www.akamai.com/`
[10]Facebook Privacy Policy `http://www.facebook.com/policy.php`

personal freedom. The advent and rise in popularity of the Social Networking Sites, for an illustration please refer to danah boyd's work entitled "Social networking sites: Definition, history, and scholarship" (boyd and Ellison, 2007), has brought about a higher level of user engagement with the Internet. The rise in the phenomena that is social networking has seen an increase in the levels of Web participation from that of most people being simply consumers of information/data from the web, to that of high levels of participation culminating in high levels of user-generated content on the web. This shift has been facilitated by advances in Web-based user interfaces (UI) which allow for non-technical users to be able to post content to the Web. This notion of higher levels of user participation is illustrated by the fact that since the social networking site MySpace[11] has quietly rolled-out email addresses for all of its users, it has now managed to position its self as the second largest email service provider[12].

A final point: Lifelogging sounds like a somewhat recondite pursuit, but personal knowledge management is an issue for anyone with a significant Web presence, or who uses digital technologies. Recent commercial propositions, such as Garlik Ltd[13], aim to capitalise on the notion that one's online persona is a valuable asset and something users should be aware of. Lifelogging is an extreme case, but the tools and interfaces that support it will also support querying of and retrieval from smaller repositories of personal data collected using more discriminating methods, a fact that is captured by the term 'Personal Information Management'. To that extent, Lifelogging tools are tools for everyone to exert more control over their personal data, their public presence online and their digital identity.

3.1.3 Privacy, Empowerment, and Identity

Privacy is of course a serious issue for the individual who wishes to amass data. The course of DARPA's LifeLog project is instructive here. LifeLog was conceived as an experiment in life-long information capture, a fairly mainstream Lifelogging effort: "an ontology-based (sub)system that captures, stores, and makes accessible the flow of one person's experience in and interactions with the world in order to support a broad spectrum of associates/assistants and other system capabilities". The objective of this 'LifeLog' concept was to be able to trace the 'threads' of an individual's life in terms of events, states, and relationships by aggregating raw data into a timeline that is an "episodic memory" (IPTO 2003).

Patterns of events in the timeline support the identification of routines, relationships, and habits. Preferences, plans, goals, and other markers of intentionality are at the highest level (IPTO 2003).

[11] MySpace `http://www.myspace.com/`

[12] The new MySpace Mail Quietly Emerges as a big-time email competitor `http://www.techcrunch.com/2009/07/30/the-new-myspace-mail-quietly-emerges-as-a-big-time-email-competitor/`

[13] Garlik: Online Identity Experts `http://www.garlik.com/`

However, the solicitation of bids from a defence research agency was seen as worrying and many focused on the aim to build a database of all the transactions, including credit card details and phone calls, of an individual. DARPA had recently been under fire for its controversial Total Information Awareness system, which licensed too much surveillance for American voters even in the name of security, (O'Hara and Shadbolt, 2008, p. 39), and the FutureMap system which appeared to encourage people to bet on the likelihood of terror attacks. LifeLog was pulled in 2004 under pressure from civil libertarians[14].

LifeLog's difficulties are indicative of the strength of privacy fears when the funding body is not trusted by a potential user community. Anita Allen (Allen, 2008) has reviewed a number of privacy issues that are likely to result from widespread Lifelogging, but at least some of the worries are overblown. First, there may be many occasions when the preservation of 'memories' is not what is required. Misfortunes or misjudgements would be preserved, possibly at the expense of average behaviour, creating in effect a false picture of a time period. But this is a danger of any kind of record photographs of one's youth tend to cluster around parties and set-piece meetings with other family members, neither of which activity actually looms very large at all on a day-to-day level. And the release of some sensitive information about the past e.g. criminal records is unlikely to be an issue with Lifelogging. Practical obscurity is on the way out ((O'Hara and Shadbolt, 2008), 81-101), and Lifelogging per se is not the most guilty party. A concrete example of how the notion of practical obscurity is slowly being swept aside, is the fact that in a majority of the states in the United States of America, implement what is informally known as Megan's Law[15], or the Sexual Offender (Jacob Wetterling) Act of 1994. An example of implementation of this Law in the State of Georgia means that personal information regarding pardoned sexual-offenders, i.e. offenders which have served their sentence, get posted to the Internet in a machine readable manner. Upon release of this information emerged a Google Maps Mash-up of the released sex-offenders[16]. It should also be noted that as the mobile web starts to mature, coupled with the development of more and more powerful devices, we will see a rise in the ubiquity of such information. One can already purchase an iPhone application which will use the user's geo-location information to given them real time information regarding the location of known sex offenders within their given local[17].

The extreme example presented above is the only justification for stating that the notion of practical obscurity is a dying concept. One area which causes concern is the notion that the advent of sex offenders data on the web will set a president for other forms of crimes to be posted to the public domain. One can easily imagine a future where

[14] Wired News April 2004 http://www.wired.com/politics/security/news/2004/02/62158

[15] Megan's Law http://en.wikipedia.org/wiki/Megan%27s_Law

[16] Georgia Sex Offenders Mashup http://www.georgia-sex-offenders.com/maps/offenders.php

[17] iPhone App Tracks Sex Offenders http://www.telegraph.co.uk/technology/apple/5918923/iPhone-app-tracks-sex-offenders.html

all crimes committed are posted to the web regardless of the severity of the felonies. For example, high-school student Bob gets arrested for shop-lifting and gets a minor punishment that could be community service or something of a similar vain. Alice, a classmate of Bob's finds this so funny that she posts it to whatever hip social network she is currently a member of, pushing it into the public domain. Now after Bob has served his sentence in a world pre-internet this information would have been practically obscure, it would have been logged in a filing cabinet/database system in some local magistrate court, and unless you had the impetus to seek out this information you would probably never have found out about it. Alice would have been able to communicate the "funny story" to her social network, but those conversations would not have been in the public domain. And given the current climate of mass-participation in social networking sites, one can almost expect mutterings between friends to get into the public domain.

Secondly, Allen highlights complex issues to do with mental health and trauma, and worries that Lifelogging might encourage pathological rumination about the past. It may be hard to persuade a patient that "Data captured by Lifelogging is not fixed, 'hard' evidence of an important whole story, rather than ... something partial, ambiguous, unimportant and interpretable". As Allen points out, pathological rumination by those obsessed with the past can happen in the absence of any reliable memory at all, but one can go further and argue that the sheer wealth of detail of Lifelogging might actually support therapy that stresses ambiguity. It is harder to maintain that one embarrassing photo tells the whole story of a gathering when another hundred are easily available. Similarly, Allen's point that forgiving may be of limited value "when there is a diminished capacity to forget" cuts both ways. Forgiveness can be easier when there is a specific commitment to remember; reconciliation has been possible in Germany and South Africa partly because of a commemoration of the horrors of the holocaust and apartheid. However, Mayer-Schönberger (Mayer-Schönberger, 2009), argues the opposite, he outlines scenarios where recalling past events can open wounds that were otherwise literally both forgiven and forgotten.

Allen's third worry is pernicious surveillance, and this of course is an important privacy issue. There are three routes by which Lifelogging might become surveillance. First, Lifelog data may feature the actions of others, who might appear in photographs, telephone calls, email exchanges, and so on. Second, the tools for gathering data about oneself might also be the tools for gathering data about others – which to be frank can be seen in work undertaken by Garlik Ltd. Third, governments have a lot of power to insist that information that exists is made available to them; as Allen points out, current laws give the US government, brought about by a post 9/11 neo-conservative agenda, "access to virtually all means of communications and data storage" and there is no reason to believe that it would stop at Lifelogs, or that it could be designed out of the technology (Allen, 2008). These routes deserve serious consideration (though this is beyond the scope of this thesis).

The W3C's Social Web Incubator Group (SWXG)[18] has been put together by the W3C to look at issues surrounding the Social Web. The Social Web is used to describe sites which support social interaction between end users. The remit of the incubator group is roadmap best practises and to detail considerations that will help develop user centric policies for the future of Social Networking. The notion of the privacy on the social web has been championed as one of the key areas of research that ought to be developed to ensure that personal information posted social websites are handled and treated with care in order to safeguard their users from having too much of their data exposed to the public domain.

3.1.3.1 The empowerment of the individual

The privacy argument is clearly real. There is, however, also a flip side of such an argument. While it is clear that some people will get pleasure from storing such data, and that others will value improvements in data management, there are also genuine goods that come from constructing and maintaining an identity using Lifelogging techniques. First, one has access to representations of the past which may come in useful. Of course, in any legal situation data would need to be authenticated, but one could prove one's whereabouts if need be. An art professor at Rutgers, Hasan Elahi[19], who was arrested and subjected to some heavy treatment by the Federal Bureau of Investigation (FBI) in 2002 despite being absolutely innocent of any criminal or terrorist activity, has taken to Lifelogging and posting the information on the Web as a pre-emptive alibi (Coughlin, 2007).

Second, the technology is of course helpful for the practise of what has been called 'sousveillance', community-based recording of events to democratise the process of surveillance. Rather than traditionally owned and controlled surveillance techniques being used to monitor a community, sousveillance supports the monitoring of the authorities, for instance searching for and reporting misdeeds by police forces, or electoral fraud, in a distributed way by a community. There are pros and cons to sousveillance (Mann and Niedzviecki, 2001) (O'Hara and Shadbolt, 2008, p. 181-183), but for an individual it can be empowering to recall interactions whose nature is disputed. A contemporary example of how fraud can be tackled by way of community is the way that the Guardian newspaper[20] recruited the masses to help them annotate members of the UK Parliament's expense claims[21]. Following the recent scandal surrounding UK Members of Parliament (MPs), the government took the initiative to release PDFs (Portable Document Format) of all of the MPs expenses, in an attempt to come across as transparent as possible. This gesture was seen as a token one, as PDFs are not machine readable, and it was thought

[18] W3C's Social Web Incubator Group http://www.w3.org/2005/Incubator/socialweb/
[19] Hasan Elahi's Tracking Transience http://trackingtransience.net/
[20] Guardian UK Newspaper http://www.guardian.co.uk/
[21] Guardian's MP-expenses annotation tool http://mps-expenses.guardian.co.uk/

unfeasible and to time consuming for anyone to ever bother going through and verifying the PDFs put into the public domain by the government. Due to extremely high levels of public interest, and the innovative tools created by technologists at the Guardian, a vast amount of the expense claims have now been annotated by members of the public, and can be found in a machine readable format on the site. The instance of crowd-sourcing can be seen as direct example of how sousveillance is starting to propagate throughout society.

Third, the construction and maintenance of an identity by an individual could act as a counterpoint to initiatives by formal and informal authorities to impose identities. There are many sources of unwanted identities, whether or not it is the creation of a formal system of ID cards (which may trespass on sensitivities, for example, by insisting on the use of a given name rather than the name in common use, or by failing to respect gender images of transgendered people), or an informal family insistence that one conform to social norms with respect to dress or sexual behaviour. The Lifelog, for the Lifelogger, might constitute the "real" person.

In the context of this thesis the Semantic Logger (see chapter 5) framework has created a focus for the digital identity by the adoption of a Uniform Resource Identifier (URI) or a WebID[22] to refer to the user. Setting up a log requires the user to create a URI which will be associated with all the personal information logged. A FOAF document, i.e. an RDF documented developed using the FOAF ontology to model users and social networks, is imported into the personal knowledge base (KB) of the user, in which the primary subject of the document is the user's URI. The FOAF document can then point to other pieces of information about the user or his or her friends, so creating a linked structure of information available on the Web. Details of the manner in which the Semantic Logger stores and associates personal information is expanded upon in chapter 5 of this thesis.

The result is an amalgamation of data about the user from distributed online sources in a single KB, providing the user with a global view of the personal information published on the Web. Information about the user can be gathered and associated, though not necessarily stored in one place, as separation is useful to help identify the provenance of a particular statement. Although of course the published information is in the public domain, users can see the information collected, and can make informed decisions about whether to attempt to withdraw or amend items (Tuffield et al., 2006b). Commercial systems in this space are beginning to appear, such as Garlik's foaf.qdos.com[23], and the experimental site that is foaf.me[24].

[22]ESW wiki WebID: http://esw.w3.org/topic/WebID
[23]Garlik's foaf.qdos.com, the beginnings of a decentralised social network: http://foaf.qdos.com/
[24]foaf.me http://foaf.me/

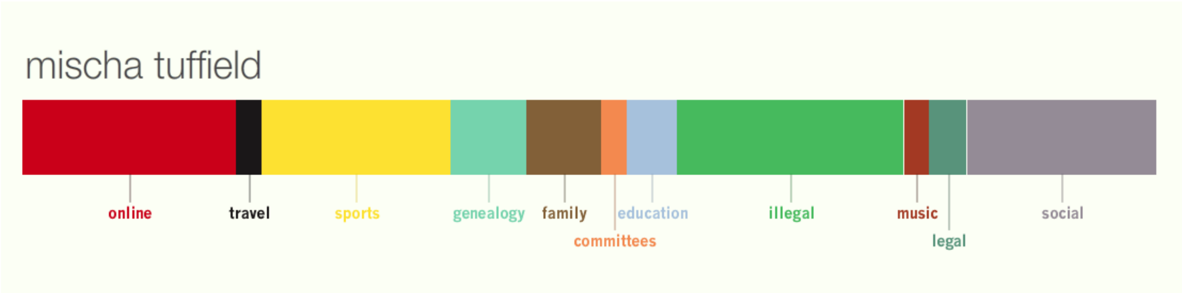

Figure 3.1: Mischa Tuffield's online persona, as per MIT's Personas on 24/08/09

3.1.3.2 Alibis and unstructured data

Given that we are only starting to witness the birth of methods allowing for queries to be made of the Internet which in turn end up presenting the query issuer with a breakdown of a person's web-accessible digital life, one can observe the relative immaturity of such services. Figure 3.1 illustrates the concerns surrounding the parsing of non-structured data off of the web using current natural language processing techniques. Given that all of the major internet search engines employ natural language processing techniques to extract information from web accessible documents, and the severity of false accusations which are being punted around the internet (see Hasan Elahi example in section 3.1.3.1), therein lies another motivating factor to why someone may want to capture their own lifelog.

Massachusetts Institute of Technology (MIT), has recently put out a service which presents the user with a pictorial description of their "online persona". The service is called Personas[25], and 3.1 is the result of the query "Mischa Tuffield". The system, aggregates results from web searches and attempts to classify the documents into classes, which in turn influences the output of the final pictorial representation. The immediacy of the of the service illustrates how much personal information can be gathered about an individual in a matter of seconds, the implications of such services are bound to be of interest to a vast majority of commercial institutions from web-search providers, e-commerce, the world of advertising, to name a few.

Figure 3.1 suggests that the three most prominent characteristics of Mischa Tuffield's online persona are : *online, sports, and illegal*. The concept of online makes perfect sense, whereas the other two aspects are slightly confusing to start with but understandable upon investigation. The concept of "sports", is due to an article published in the motoring section of the of the Daily Telegraph[26], and the notion of "illegal" is due to the fact that personas thinks that Mischa Tuffield is an identity thief. It seems that it is down to the misinterpretation of a blog post, entitled : "Garlik releases FOAF services"[27]

[25]Personas Project at MIT http://personas.media.mit.edu/

[26]Total Recall http://www.telegraph.co.uk/motoring/2754679/Total-recall.html

[27]Garlik releases FOAF services http://tomheath.com/blog/2008/05/garlik-launches-foaf-services/

written by Tom Heath[28]. The blog post includes the following sentence "...was trying to steal my identity (presumably because he had a fragment of RDF about me in his FOAF file)..." that seems to be content which triggered the usage of the term "illegal". This is a perfect example of how natural language processing can fail, and how much more sophisticated metrics must be used if we are to identify accusations, opinions, or any more complex statements from free text. This notion of opinion mining is central to the work being undertaken by the LivingKnowledge EU project[29]. Tom's blog post was actually referring to how Google's Social Graph API[30] failed to understand his FOAF file, merging Mischa Tuffield and Tom Heath into one person. This example highlights the fact that unstructured data, published to the web can easily be misinterpreted and often extremely misleading. The notion that Mischa Tuffield, who now works in an identity fraud protection company, would be labelled an "identity thief" could be very damaging, although in this case it is easy to spot where the misunderstanding occurred. This is presented as motivating factor for the generation of structure data on the web.

3.1.3.3 Logging in public and private

There are clearly complex balances to be struck. But there is a clear distinction between information about one that is out there, and personal information. Personal items, such as Web browsing habits, geolocation data or emails, can be stored in a personal knowledge base, while public domain information, including personal information that the user has deliberately exposed on social networking sites, can be stored separately. Although Lifelogging involves gathering data in a relatively non-discriminating manner, that does not mean that one cannot discriminate in the treatment of the data. In particular, the distinction between publicly available data and data one does not wish to expose becomes more pressing as 'techy' lifestyles such as Lifelogging and social networking are pursued simultaneously. In 1990, Alan Dix (Dix, 1990) wrote about privacy highlighting the fact that in order to keep facts private one should focus on not recording things that should not be remembered or stored. This thesis aims to provide techniques and insight into how one can empowering users to not give away their privacy whilst still being able to interact and make use of the various technological advances made in our highly interconnected social online world.

The Semantic Logger (see chapter 5) uses a two-way system to produce a very basic implementation of an intuitive trust model. It consists of a central KB where public data is held and published, whilst one password protected KB is created for each user – their private KB. The private knowledge base is presented as one which the user can grant their friends access to so that they can make use of each other's lifelogged information. An example use for such information could be as follows : a Semantic

[28] Dr Tom Heath http://tomheath.com/id/me
[29] LivingKnowledge EU FP7 Project http://livingknowledge-project.eu/
[30] Google Social Graph API http://code.google.com/apis/socialgraph/

Logger user could be identified to be in the same place as one of the friends, by virtue of geolocation information, and if their friend has got a calendar entry for that specific event at that time and place, the user whose calendar does not hold any information for that event could end up adopting their annotations. This model could be extended so that each user is given two private knowledge bases, the first to hold the user's private data, while access to the second can be granted to friends. Hence – as with life offline – one's online identity as determined by the Lifelog can be decomposed into entirely personal aspects, aspects reserved for one's intimates, and a public face (Tuffield et al., 2006b). It should be noted that there would be nothing stopping a user from creating multiple accounts for his/her different personas. This naïve approach of implementing two different personal knowledge bases adopted by the Semantic Logger is presented as a first pass at what a decentralised social network may look like, further discussion around future of social networking is presented in the future work chapter of this thesis (see chapter 7).

Given the privacy issues presented thus far in this thesis the Lifelogging framework presented, namely the Semantic Logger (see chapter 5), must cater for the below characteristics.

- **Total User Control** : Any lifelogging system must allow for the user, and only the user to dictate what information is shared with other parties.
- **The User as the Primary Topic** : Lifelogging tools, should be focused on capturing personal information about the user. Ideally information pertaining to other people should be linked to, as apposed to stored with a given personal knowledge base.
- **Best Practises** : Given that Lifelogging as a concept is currently novel and limited to the technologically adept, best practises regarding the sharing and storing of lifelogged personal information should be shared in order for the discipline to mature.

3.2 Lifelogging: The Realisation

The MyLifeBits model of Lifelogging is centralised and labour-heavy. If Lifelogging is to take off as a pastime or as a way of life, the MyLifeBits model needs revisiting. This section discusses some more general points about what Lifelogging is likely to entail.

The growth of information acquisition, storage and retrieval capacity has led to the development of the practise of Lifelogging, the indiscriminating collection of information concerning one's life and behaviour. There are potential problems in this practise, but equally it could be empowering for the individual, and provide a new locus for the construction of an online identity. This section looks at the technological possibilities and constraints for Lifelogging tools, and refer back to the important privacy, identity, and empowerment-related issues. In the previous section it is put forward that some of the privacy concerns are overblown, and the major issues will be concerned with surveillance. Furthermore it is important to stress that much of the research and commentary on Lifelogging has made the unrealistic assumption that the information gathered is for private use, whereas, in a more socially-networked online world, much of it will have public functions and will be voluntarily released into the public domain. Furthermore this section will present motivations behind the capture of lifelogged data in a structure manner whereby knowledge about the day-to-day activities of a user are captured as a collection of events. The motivation to encode lifelogged information as events is presented in chapter 5 of this thesis and directly influenced the development of the OntoMedia ontology, presented in chapter 4.

3.2.1 Lifelogging and Technology

The growth of computing capacity as predicted (and driven) by Moore's Law (see figure 3.2[31]) has meant that there are few barriers to the storage of information. In particular, a person can now store significant quantities of information about him- or herself. It has been estimated that, in the normal course of events during a life of normal duration, transactions involving a person will create something of the order of 100 gigabytes of information as a by-product. Gordon Bell and Jim Gemmell have estimated that a terabyte of storage would hold all the books, emails, recorded conversations, music tracks and photographs that one is likely to accumulate over sixty years, which a typical desktop PC will probably be able to hold by 2010, and personal digital assistants (PDAs) will manage by 2015 (Bell and Gemmell, 2007). Alan Dix has noted that even 70 years of high-quality video recording would require something less than 30 terabytes of storage (equivalent to under 375 of Apple's largest iPods, which store 80 gigabytes and cost in the order of £150 each cf. (Dix, 2002; O'Hara et al., 2006)). Since the publication of Dix's calculations, the proposed size of continuous recording of 70 years of a human's life will now fit on 256 of Apple's current instantiation of their largest iPod, which can store 120 gigabytes of data costing £175 pounds each.

These technological possibilities have led to the development of so-called Lifelogging technologies and tools, to support the practise of exploiting digital storage systems to record information about a person, or group of people, automatically and persistently. Typical types of information to be logged include emails, documents, digital photographs and video, diaries/calendars, geolocation data using the Global Positioning System (GPS), music downloads, listening habits, blog entries and Web browser bookmarks and navigation history. The result for the user is a large store of information much of which

[31]Picture taken from Keio University's online lecture: "What's in a computer?" http://web.sfc.keio.ac.jp/~rdv/keio/sfc/teaching/architecture/architecture-2008/lec01.html

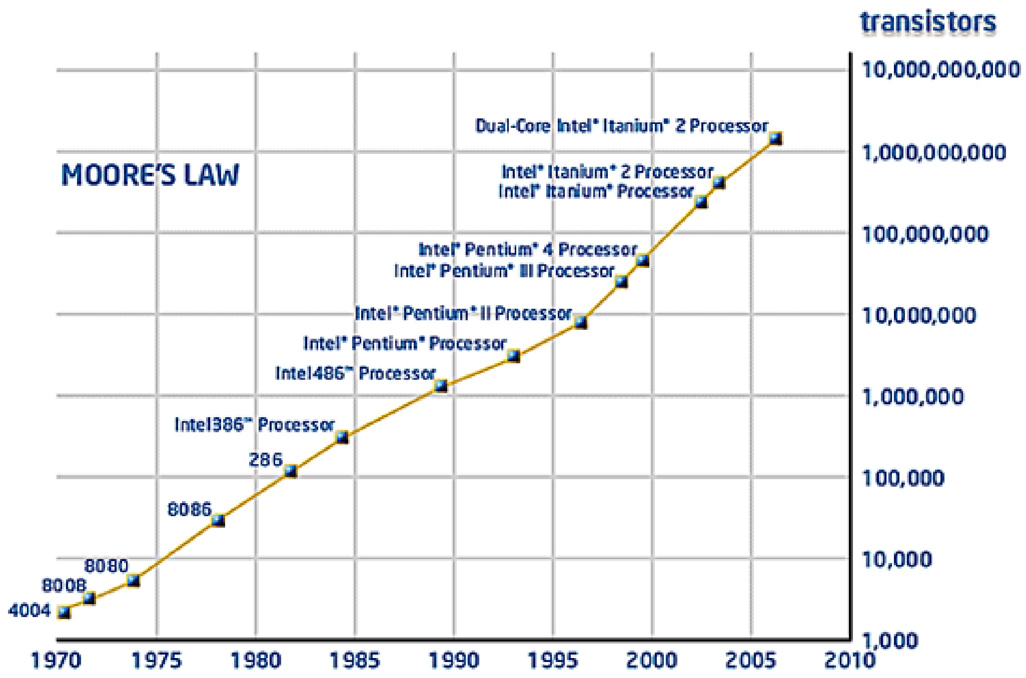

FIGURE 3.2: Moore's Law

will be trivial or ephemeral. It is important to note that it may not always to clear at the point of capture whether information will be simply ephemeral, as future technological advances may make use of information once deemed useless. One example of such information that may have seemed to be deemed trivial but now can be used to contextualise location data is that of IP-addresses[32] of networks which a user has joined, this is referred to as a Network Gazetteer in section 5.3 of this thesis.

When the Semantic Logger was implemented, in 2006, there was no publicly accessible service which would translate IP addresses into geolocation information. The Semantic Logger proposed the capture of such information, which could then be associated to geo data collected by the user in order to associate an IP address with a location. This task is no longer needed due to the advent of the W3C's Geolocation API[33], which in turn employs the Skyhook service[34] to assign geolocation data to a given IP address.

The purpose of storage of such information can vary, and may not be clear even to the Lifelogger at the point of storage. However, in an information-intensive age where the surrender of digital identity is a commonplace, for purposes ranging from commerce, marketing, social networking, government, receipt of services, travel or security, Lifelogging has the potential to reaffirm the individual's control of his or her own identity. The Lifelog is a constructed identity that outweighs the others simply by weight of evidence, complexity, curator, and comprehensiveness. It is the thoroughness of the captured information which is presented as a method of empowering the user with knowledge about

[32]Internet Protocol Address or IP-address http://www.ietf.org/rfc/rfc791.txt
[33]W3C Geolocation API http://www.w3.org/TR/geolocation-API/
[34]Skyhook Network Gazetteer http://www.skyhookwireless.com/

their digital persona. It is likely to include other identities, amalgamate and supplement them, and contain links to information both held in the public domain and to other lifelogs.

3.2.2 Lifelogging and Memory

Human memory is strangely unreliable. One remembers pointless stuff while forgetting someone's name, who won the cup last year, how to make *coq au vin*, someone's telephone number, that one has a dental appointment at three, where one has put one's car keys. Much forgetting is particularly linked to memory function, so deleting irrelevant, unused or out of date information is good housekeeping for the head (Schacter, 2001), but unfortunately in the nature of the case it is impossible to forget all and only those things that one does not need.

Technology has been used to aid memory since the written word. In Plato's Phaedrus [35], Socrates worries about the effects that the new literacy will have on Athenians' memories. But in the 21st century computing power has provided us with the means for augmenting our memories to a degree hardly dreamt of even ten years ago. We can store information in enormous quantities, so we no longer have to be selective. Removing the requirement for selectivity makes storage much less heavy in its use of other resources; e.g. one does not have to view it all to decide which to keep and which to delete. Gadgets such as cheap, small sensors make it increasingly simple to extract large quantities of data automatically from the environment. Increasing use of digital technologies means that many records of meetings or communications are in digital form already. And better search, retrieval and mining techniques mean that we are better able find the important signals in noisy data (O'Hara et al., 2006, p. 81-108).

We are getting to the stage when all information, interesting or otherwise, generated in a lifetime by a single person can be assembled in a giant autobiographical silo, and queried relatively efficiently, creating a need for Personal Information Management (Jones, 2008). This can be seen as a step-change in the relation between the technology of information storage and human memory, as identified by the 'Memories for Life' (M4L) network. M4L demands interdisciplinary research not only on technological infrastructure but also on social, legal, political, philosophical, psychological and medical aspects of the technological development.

Technological advances in search technology brought about by the open-source implementation of Google's PageRank algorithm (Page et al., 1999), based on Kleinberg's Hubs and Authorities paper (Kleinberg, 1999), called Nutch (R. Khare and Rifkin, 2004), and an implementation of Google's distributed filesystem called Hadoop (Borthakur, 2007) which have both been championed by the Apache Project[36] allows for off the shelf

[35]Wikipedia Entry for Phaedrus : http://en.wikipedia.org/wiki/Phaedrus_%28dialogue%29
[36]Nutch http://lucene.apache.org/nutch/

indexing of a large numbers of documents, allowing lifeloggers the ability to search and rank their documents. This Java based implementation of PageRank allows for Google like free text querying of personal documents. It should be noted that the maturity of the Nutch/Hadoop codebase is evident from the fact that Yahoo Inc[37], run majority of their search services on top of the Hadoop infrastructure.

But why would we want M4L? Why should anyone want to keep emails whose point, if any, vanished the moment the computer was shut down? Surely there is a limit to how much information we would like archived. There is surprisingly little research on memory prostheses, although (Kalnikaité and Whittaker, 2007) have made some interesting discoveries of when exactly people are tempted to reach for the PDA rather than try to remember something unaided. Some people with severe memory impairments may rely on prostheses, and important work is appearing to show that patients with Alzheimer's can use automatically-created photographs of apparently trivial events in their recent past to boost their medium-term memories more effectively than other methods such as keeping meticulous diaries (Berry et al., 2007). Further insight into the challenges posed by the M4L network are highlighted at the end of this chapter (see section 3.4).

A compelling reason often overlooked is the issue of association. Memory links are often associative – one is drawn to one topic by links to another, and recall is enhanced if background conditions can be recreated. So, for instance, suppose you were looking for a particular document on your computer without noticeable success. If you remembered that when you read it last you were listening to Metallica's S&M on your MP3 player, and if you had stored information from the player and a properly integrated interface to all your personal data, you could restrict the search on the desktop to all documents that were open at the times at which Metallica's S&M was playing on the player. There is nothing fundamentally interesting about the times at which particular pieces are played on the player, but there may be wider application. Information clarifying the context of an event answering *what, where, when*, and *who* are key to supporting tasks of recall. This notion of associative linking is one of the key motivators in the development of the Framework at hand and the adoption of Semantic Web technologies. The schemaless nature of storing data in RDF, the defacto knowledge representation language for the Semantic Web, allows for new data to be found, and links to be added on an ad-hoc basis. Further insight into the Semantic Web technologies are presented in chapter 2. Another motivational factor for the aggregation and storage of one's digital traces in a central all encompassing repository is the ability to mine and interrogate the data in order to find associative links. The associations found between the various aspects of one's digital life may present insight into what information can be mined about people from the outside. Given that more and more information is being posted to Social Networking Sites, it may be hard to imagine the implications of posting information to each site in isolation,

[37] Yahoo Inc http://www.yahoo.com/

whereas storing all of the data in one place may suggest what picture can be painted about an individual from the information they are leaving on on the Internet.

3.2.3 Lifelogging and Social Networking

In dana boyd's paper entitled "Social Network Sites: Definition, History, and Scholarship" (boyd and Ellison, 2007) a social networking site (SNS) is defined as any web-based system that caters for three interactions which follow. This definition of a social networking site will be used within the context of this thesis.

- The ability to construct a public or semi-public profile within a bounded system
- The ability to articulate a list of other users with whom they share a connection
- The ability to view and traverse their list of connections and those made by others within the system

The original rationale for Lifelogging was as a personal tool to manage one's own information. However, since early experiments in this field, the practise of social networking has developed, where users generate and share information with others. This information can be quite specific in type (delicious.com allows people to share their Web bookmarks) or form (Flickr allows sharing of photographs), or can be quite general (Facebook and MySpace allow people to connect and interact, revealing as much information about themselves as they care to post). Meanwhile, other practises such as blogging admit conversation, information and discussion into the public space.

In this context, information gathered by Lifelogging practises could be shared, or enhanced by integration or cross-referenced with information from others. As the recreation of context enriches information, there is no reason in principle why the information sources upon which such context-recreation draws should be restricted to ones controlled by oneself. Someone else's calendar might be as informative as one's own when it comes to retracing events in one's life.

This suggests the imperative to integrate the information describing people and their social relationships which is exposed by social networking sites. Portability of data across applications is already an important concern (Szomszor et al., 2008) as people wish to carry data and personal profiles or identities across sites. Much of the information is non-sensitive and its creators are keen to share it, for instance, about film ratings and music downloads.

Hence it is probable, given the current profile of those who spend a significant proportion of their lives online, that the activities of Lifelogging and social networking will intersect, and some social networkers will use Lifelogging techniques to generate large quantities of information for their own use, and who will not be shy about sharing it with friends or like-minded people.

3.2.4 Lifelogging and Services

The rationale for collecting large quantities of generally unfiltered information is that it cannot be specified in advance which information will be useful and which not, or what tasks the information might be used for. Hence the ideal architecture for a Lifelogging system should be open not only from the perspective of information sources, but also from that of service provision. The Semantic Logger uses a knowledge base (storing RDF) as a persistent repository for the system, and to mediate interactions between information sources and service outputs (Tuffield et al., 2006b). On such a model, a software service could be devised for almost any purpose to which the stored information could be put. The flexibility that RDF based solutions offer in terms of their ability to index and house arbitrary data, whilst exposing access to the knowledge stored within in a standard manner is a key motivator for the implementation presented both in the background chapter 2 and in the chapter dedicated to the Semantic Logger (see chapter 5). Likely examples include the following.

- **Queries** Simple querying over the integrated data set, with questions such as those set out the 'Use Cases' section at the end of the Semantic Logger chapter (see section 5.7), which is the approach explored by MyLifeBits.
- **Recommendations** Where past behaviour is used to suggest items of interest. One's music downloads could suggest other music one might enjoy; academic papers saved to disc could suggest other items in the literature; the ratings one has given films could be used to suggest future films to watch or DVDs to buy (Loizou, 2009). Relevant information might come from a wider social group as well as the Lifelogger.
- **Search & Retrieval** The provision of metadata about information in order to facilitate search and retrieval (either by oneself or others), a strategy particularly suited to multimedia. One example, Photocopain, helps with the traditionally labour-heavy task of annotation of digital photographs (Tuffield et al., 2006a).
- **Social Networking Population** Creating and populating some kind of avatar, knowledgeable about oneself, which could act as an interface between oneself and other online actors, making decisions (for instance related to privacy and the revelation of information) on one's behalf. Such an avatar would be an interesting facet of one's identity that could, for instance, be used to construct narratives about one's life (which could easily take place after one was dead), to populate social networking sites, or partially to automate one's interactions with governments

or companies (Wilks, 2006). The beginnings of such technologies are starting to pop up, one such example is RDF Pushback[38]. RDF Pushback aims to turn the current 'read-only' Semantic Web into a 'read/write' Semantic Web but building open-source APIs to write data back to the various Web 2.0 services currently so prevalent on the Web.

- **Medical History** The provision of medical history. Possibilities include: (a) bodily sensor data to monitor changes in current health of the individual; (b) community-wide effort to prevent or monitor an epidemic; (c) the use of technology accurately to determine, e.g. someone's actual diet, a notoriously difficult thing to measure; (d) monitoring the use of household gadgets (e.g. kettles or fridges) to signal that the user (perhaps an elderly person) remains in good health ((O'Hara and Shadbolt, 2008), 15-16). Google has a service to allow people to upload medical records[39], to follow Microsoft's HealthVault[40].

- **Memory Prosthesis** An aid to forgetting. Ironically, for a technology often seen as the antithesis of forgetting, Lifelogging can be used to measure which pieces of information are recalled directly or used indirectly in associative recall. As well as being of interest from the information management point of view, this information could also be used to decide which information could be junked with least harm. The notion of forgetting, or in technical terms, deleting information is slowing starting to fade (Dib, 2008). This could be due to the proliferation of ranking algorithms based on the Kleinberg's original Hubs and Authorities algorithm (Kleinberg, 1999) and the fact that storage costs, as driven by Moore's Law, are decreasing by the day. The idea around Hubs and Authorities is based on the fact that all the items indexed by such an algorithm will have a probability, albeit very small, to be returned given a query. Every time a new item is added to the index, the probabilities are recalculated in order to recalculated the Hubs and Authorities in the document pool. The manner in which these probabilities are recalculated vary based on the given implementation but they all seem to share two common properties. The probability of a given document tends to be based on both the number of other documents which reference it, and the amount of times the document has been accessed. The fact that documents tend to never be removed or deleted, but just demoted to a state where they are very unlikely to be returned, naïvely, can be said to be very similar to way human memory works. Memories tend not to be deleted from our memory but in turn seem to be pushed to the back of our consciousness whereby they are recalled given exposure to an event/fact that then triggers the association in our brains, bringing the memory back into the foreground of our thoughts.

[38] RDF Pushback Project: http://code.google.com/p/pushback/
[39] Google Health http://www.google.com/Health
[40] Mircosoft's HealthVault http://www.healthvault.com/

Another piece of work worth mentioning is the PhD thesis entitled "Who Controls the Past Controls the Future - Life Annotation in Principle and Practice" (Smith, 2008). The thesis describes a portable GPS-based logging system, carried by the researcher for a two year period, that automatically captured contextual data in order to generate a lifelog. The thesis states that one type of contextual data, geo-location information, is not enough to paint an accurate picture of one's life, but illustrates how merging the data with other forms of contextual information could be worthwhile.

3.3 Information Overload: Personal Information, *Infosmog* & the Semantic Web

A key problem faced by the information management community is that of information overload, or *infosmog* (Shadbolt and O'Hara, 2003). The ever growing and evolutionary nature of the World Wide Web (Web) is not making the task of information management any easier. The web is populated with an unmanageable amount of heterogeneous data, in forms as diverse as image, sound, video, and text. This coupled with the shift in the consumer electronics market from the analogue to the digital recording medium and the mass-market, low cost nature of the hardware industry has produced a rich *corpus* of personal multimedia artifacts.

The current trend of publishing personal information to the web, in the form of calendars, photo-collections, Global Positioning System (GPS) tracklogs, and the adoption of standards such as iCalendar (iCal)[41] files, the MBOX family of file formats[42] used for email correspondence, and the abundance of web accessible services such as Flickr or Plazes[43], is presented as liberation of personal information. This liberation is seen as a social shift towards the self publishing and archival of personal information to the social environment that is the Web.

Since the rise of social networking, estimated to be in 1997 (boyd and Ellison, 2007), we have witness a dramatic rise in numbers of people which have been writing to the Internet as apposed to being simply consumers of information. Given the maturity of interactive web-site building technologies, such as Javascript[44], JSON[45], and AJAX[46] and the ease to with they allow people to upload information to the web, the generation of content is no-longer restricted to big publishing houses or governmental institutions. This low barrier to entry, whether it be the easy of setting up of a blog, or creating your

[41] RFC 2445: Internet Calendaring and Scheduling Core Object Specification http://tools.ietf.org/html/rfc2445
[42] Mbox Specification: http://www.qmail.org/man/man5/mbox.html/
[43] Plazes network caching site: http://www.plazes.com/
[44] Javascript http://en.wikipedia.org/wiki/JavaScript
[45] JavaScript Object Notation http://www.json.org/
[46] Asynchronous JavaScript and XML http://en.wikipedia.org/wiki/Ajax_%28programming%29

Rank	Site URL
1	http://www.google.com/
2	http://www.yahoo.com/
3	http://www.facebook.com/
4	http://www.youtube.com/
5	http://www.live.com/
6	http://wikipedia.org/
7	http://blogger.com/
8	http://www.msn.com/
9	http://www.baidu.com/
10	http://www.yahoo.co.jp/
11	http://myspace.com/
12	http://www.google.co.in/
13	http://www.google.de/
14	http://twitter.com/
15	http://qq.com/
16	http://www.rapidshare.com/
17	http://www.microsoft.com/
18	http://www.google.fr/
19	http://www.wordpress.com/
20	http://www.bing.com/

TABLE 3.1: Alexa Top 20 most accessed websites as per August 2009

own personal online photo-collection is illustrated by the increase in amount of email addresses currently found on the web (see section 3.1.2). This coupled with the fact that out of the top 20 most visited sites on the Internet, as per Alexa[47] in August 2009, see table 3.1, all of them are either social networking sites, search engines, or one of either Wikipedia or Microsoft. All of which, less the search engines and Microsoft site, are ones which promote the creation and sharing of user-generated content. This illustrates the fact that more and more people are using the internet to share and build communities around their personal information.

3.3.1 The Web and the abundance of Personal Information

As eluded to above we are currently faced with a dramatic change in the way people publish personal multimedia artifacts. The popularity of social software sites like Flickr, Plazes, MySpace[48], Facebook, Delicious[49], and Last.fm[50] can be perceived as a new found desire to publish personal information on publicly accessible websites. The benefits of this social nature, i.e. the interaction amongst peers, of the aforementioned sites have been presented as one of the key drivers behind their success (Marlow et al., 2006). Other factors such as: the ease of publishing, and vanity have also been cited in the

[47] Alexa Web Traffic Monitor http://alexa.com/
[48] MySpace http://www.myspace.com/
[49] Social Bookmarking site, Delicious.com: http://delicious.com/
[50] Last.fm: http://www.last.fm/

literature (House and Davis, 2005). The number of people adopting shared practises to document certain phenomena, Flickr for images, iCal for calendars, is increasing by the day. This apparent willingness to post personal information on the web, is said to be a key motivator for this work.

The ever increasing production of affordable hardware such as digital cameras, global positioning systems (GPS), hard-disks and so on has led to a phenomenon of mass generation and archiving of multimedia, the likes of which have never been seen before (Gillmor, 2008). In addition to the immense volume of generated multimedia data, the phenomenon is also new in the sense that the publishers and curators of this data are members of the public. This change inevitably introduces a much greater variety of representations used to describe such objects, since users of the Web cannot be expected to have the classification skills of trained librarians. The tasks of searching for, navigating through, and maintaining awareness regarding what personal multimedia are available but be re-purposed to suit this social phenomenon.

3.3.2 The Semantic Web: An Introduction

The requirement to query heterogeneous information implies that important underlying technologies for Lifelogging will be those associated with the Semantic Web (Shadbolt et al., 2006). One recent effort to develop low-effort Lifelogging tools, the Semantic Logger, was aimed explicitly at using as many World Wide Web Consortium Semantic Web recommendations as possible (Tuffield et al., 2006b). The Semantic Logger is able to exploit Semantic Web formalisms as a *lingua franca* for representing information from large-scale, distributed and heterogeneous sources, which is the ultimate purpose of the Semantic Web. Such formalisms include the knowledge representation language RDF, querying language SPARQL, the framework of Universal Resource Identifiers (URIs) and basic structuring of information using the Friend of a Friend (FOAF) ontology (Tuffield et al., 2006a).

The work presented in this thesis attempts to build upon a number of key Semantic Web enabling technologies (see section 5.2), which have gained maturity over the last few years. In an attempt to adhere to the initial motivation of capturing personal information in a machine-readable, non-proprietary format the information chosen to be captured by the Semantic Logger will be stored and archived in a manner which attempts to implement as many W3C[51] recommendations as possible. This decision was taken to ensure that the metadata published would be exploited by as many web services as possible, building on the Semantic Web Vision (Berners-Lee et al., 2001) of interoperable, standard compliant web accessible resources.

[51]The World Wide Web Consortium: http://www.w3c.org/

The work builds on the ideas brought forward in the original *Scientific American* Semantic Web article, with a particular focus on the notion of assembling, and integrating personal information into web accessible resources (Shadbolt et al., 2006). At his keynote speech during the International Semantic Web Conference 2003 (Berners-Lee, 2003) Tim Berners-Lee identified the *'Killer App for the Semantic Web'*, not as a single application but the successful integration of information, or to use his blunt words, *'Its the integration, stupid!'*.

This thesis presents work that integrates a number of sources of information (identified in Section 5.3), to build up personal metadata chronology. It is important to allow users to select how much information they wish logged, or wish exposed to the rest of the Semantic Web. This thesis is presented as a means to roadmapping how to populate the Semantic Web with personal metadata. It comes down to some basic principles regarding the exposing of information in a structured and standard manner, i.e. by using the Resource Description Framework (RDF) (Manola and Miller, 2004) accessible through SPARQL endpoints (World Wide Web Consortium, 2005). The framework described also makes use of the universal naming scheme of web accessible Universal Resource Identifiers (URI) to point to a user's Friend of a Friend (FOAF)[52] representation. In the case that the user does not have a FOAF file the system can generate a basic one upon registration, allowing them to edit it as they see fit. Each user's FOAF Person URI serves as a unique identifier for data about them. The user's FOAF URI is employed to log the provenance of all the information asserted in both the personal and the public KBs. A URI associated to a person has become known as a WebID[53], and this term will be used in the context of this thesis to mean – a URI adopted by an individual to represent themselves on the Internet. Further discussion into the nature of the FOAF ontology, WebIDs, and the manner in which the Semantic Logger frameworks makes use of unique identifiers will be presented during the course of this thesis.

In an attempt to realise some of promises road-mapped by the Semantic Web community: the seamless integration of heterogeneous data, and that of services exploiting existing machine-accessible knowledge (Shadbolt et al., 2006), a decision to create an easy to use infrastructure allowing users to store, update, visualise, and query their own personal knowledge base(s) through the web, seemed a pragmatic course of action and is described in full in chapter 5.

3.3.3 The Semantic Web and Personal Information

One of the goals of this research is to identify readily available sources of information (see section 5.3), and combine them into a structured autobiographical log. One key motivation to undertaking this task of knowledge elicitation was to define a framework for the

[52]Friend of a Friend: `http://www.foaf-project.org/`
[53]WebID description on the ESW Wiki `http://esw.w3.org/topic/WebID`

capture of personal information requiring minimal effort by the end user. The Semantic Squirrels Special Interest Group[54] (SSSIG), based at the University of Southampton, set out to identify methods of logging available raw data, (referred to as *'nuts'*), that describe aspects of a person's interactions with their computer and their physical environment (GPS track-logs and photos). The group was also concerned with identifying potential uses for the information gathered and this work is presented as just that, an exploitation of personal data.

A number of *squirrels* have been developed to gather the *nuts* available and it is these that have been propagated into our metadata logging infrastructure[55]. An ethos of the group is to preserve this raw data in order to retain any unforeseen potentials for exploiting the information gathered, this approach has also been adopted here. The group is keen to not process the information at all in order to transcend issues pertaining to platform, knowledge representation, and application restrictions, and this is where this thesis moves away from the work undertaken in the SSSIG.

Most Lifelogging projects, such as MyLifeBits, tend to engineer over-arching knowledge representation formats to integrate information, but it is arguable that a simpler route is to retain the heterogeneity of information sources, so that applications using the information can use the most appropriate mappings between information sources, depending on which sources the application is currently exploiting (Tuffield et al., 2006a).

It is clear from the examples given in the discussion of MyLifeBits that the range of information that could be gathered is almost limitless. But of course most Lifeloggers will be unable to afford a personal assistant to scan their coffee mug logos. Instead, the information that is likely to be gathered will be relevant to the Lifelogger's main interests, and/or virtually costless to gather. Information sources which are cheap to capture and likely to be popular include:

Based on initial discussions within the SSSIG, the following collection of personal information sources were select for capture by the framework presented. The nature of the data captured, and the resulting knowledge stored in the user's personal knowledge-base is presented in chapter 5.

- **Email** Sent and received.
- **Calendar entries** The user's calendar, giving plain text information regarding the user.
- **Geodata** Taken from GPS tracklogs and a Network Gazetteer.
- **Music** Listening habits.

[54] The Semantic Squirrels SIG: http://www.semantic-squirrel.org/
[55] The Semantic Logger Downloads Page: http://akt.ecs.soton.ac.uk:8080/downloads.php

- **Web browsing information** Including bookmarks, downloads, and navigation history.
- **File system information** Including document access information.
- **Photo-based information** The user's photos are examined to provide insight into events in their lives.
- **Community-generated information** Information about a user's peers extracted from Social Networking Sites.

Information of this sort is relatively straightforward to gather, and likely to be of interest in itself (to the individual who generated it), or to help with associative searching when looking for items in their personal multimedia collection. Key to decisions taken is the fact that capturing the list presented above requires little to no effort on behalf of the users. Furthermore it is important to stress that all of the sources of information presented above make up a user's *digital footprint* and get produced as a by-product of daily computer usage. The information presented above was selected due to the potential of being able to help answer the key questions about the user's activities touching on – *what, where, when, and with who* has the user been interacting. The question of who was relevant to a particular event might be answered with the help of emails, community tags, calendar entries, online accounts of events and friends' GPS records. Figure 3.3[56] shows schematically how information from different sources can be used associatively. For example, information available from an address book links the concepts of people and places; a calendar links people and time; events link people, time, places and transfers of money, and so on. The diagram below presented by Tim Berners-Lee (Berners-Lee, 2007) is presented as one of the motivators of writing this thesis.

3.3.4 The Linked Data Initiative

In 2007 a grassroots movement within the Semantic Web community, named the Linked Open Data (LOD) movement[57], took it upon itself to start eating their own dog food and started to publish resolvable RDF data – mainly converted from existing online data-sets with resolvable URIs. Core to this activity of minting re-useable, resolvable URIs for things was the development of DBpedia [58], an RDF representation of Wikipedia[59]. DBpedia has formed the locus for linked data on the web (Bizer et al., 2009), by providing both application developers and researchers with a set of URIs for the wide range of concepts defined on Wikipedia. Furthermore, due to the collaborative nature of Wikipedia,

[56]Tim Berners-Lee's Semantic Web Metro Diagram: http://www.w3.org/DesignIssues/diagrams/SemWebAppMetro.png
[57]Linked Open Data Community: http://linkeddata.org/
[58]DBpedia machine readable version of Wikipedia http://dbpedia.org/
[59]Wikipedia the free Encyclopedia http://www.wikipedia.org/

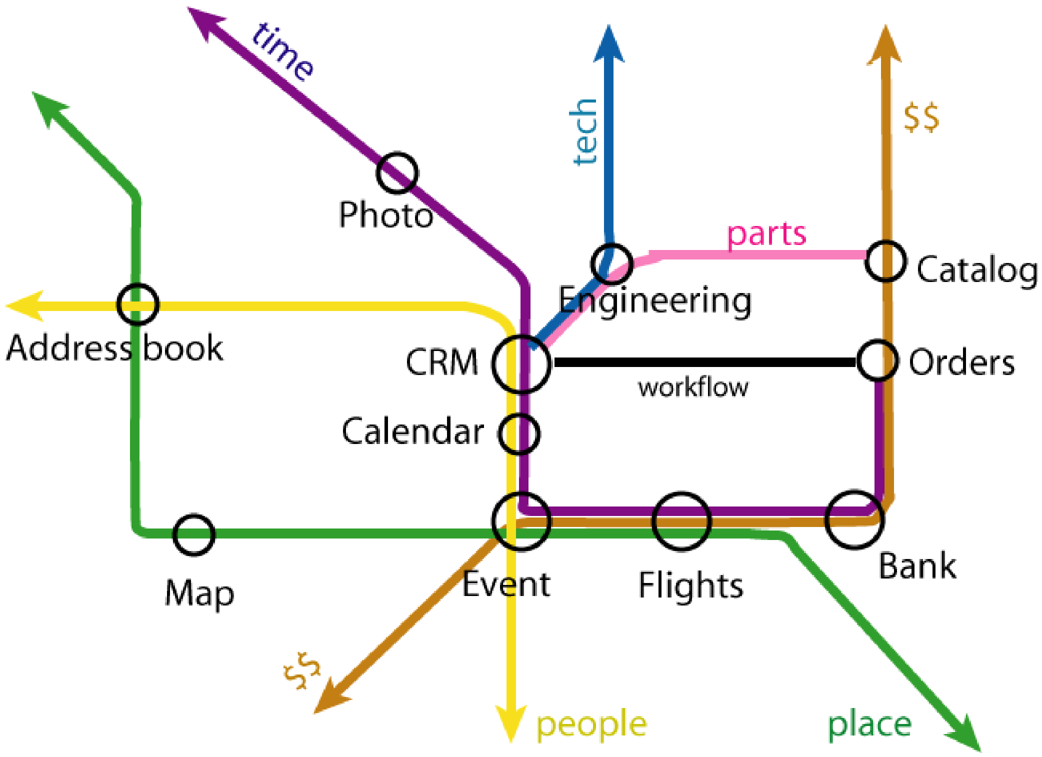

FIGURE 3.3: Integrated Data and Services - 'The Semantic Web Metro'

and the large number of contributors editing Wikipedia, the definitions presented are for the most part the consensus view, which in turn helps overcome the age old problem of experts not agreeing on shared definitions for concepts. To see a pictorial description of the Linked Open Data available on the web refer to figure 3.4, this pictorial description has been labeled the LOD Cloud. For a chronological view of how the LOD cloud has evolved since its first inception please refer to appendix A.

The Linked Open Data movement attempts to set out a number of best practises[60][61][62] for the publishing of data to the Semantic Web, and it is important to stress that none of the principles set out by this grassroots movement contradict any of the basic principles set out by the original semantic web vision. The simple matter of road-mapping best practises by providing tutorials and accessible examples was all that was needed by this grassroots movement. The best practises set out by the Linked Data community can be summarised as follows :

- The Use of Resolvable URIs
- The Use of RDF

[60] Tim Berners-Lee's talk: Cool URIs don't change http://dig.csail.mit.edu/2007/Talks/0108-swuri-tbl

[61] W3C Document describing URI styles and best practices http://www.w3.org/Provider/Style/URI

[62] Best Practises for Publishing RDF Vocabularies http://www.w3.org/2001/sw/BestPractices/VM/http-examples/

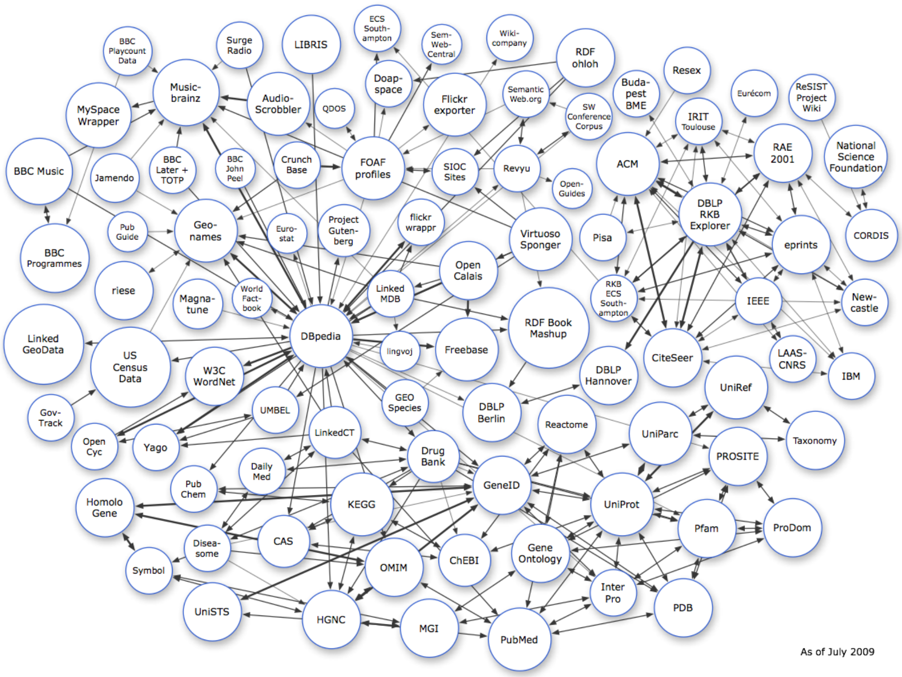

FIGURE 3.4: LOD Dataset 2008-07-14 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2008-07-14.png

- Reuse of Existing Ontologies

By virtue of minting new URIs which are within domains owned by the data publishers, they can ensure that whenever one stumbles upon an RDF fragment referring to a given URI, one can always resolve the URI to get information about it. This notion of minting URIs which when resolved return RDF about the concept has said to put the "Web back into the Semantic Web", in turn helping to fore fill the original vision of a machine-readable linked data on a global scale.

Given that for the most part, this thesis aims to tackle issues surrounding the capture and exploitation of sensitive personal information it should be noted that since the development of the first LOD cloud personal information in the form of FOAF profiles have made up a significant proportion of RDF data found on the Web. This high proportion of person related information is also evident upon inspection of the data harvested for the Billion Triples Challenge[63] at the International Semantic Web Conference in (ISWC) 2009[64]. The best-practises around the notion of hosting and serving RDF data on the Web, are taken and expanded upon with in this thesis whereby suggestions are presented to how these principles can work alongside emerging access control mechanisms to give us a privacy aware Semantic Web (see section 8.2.6 in the future work chapter).

[63]Billion Triples Challenge 2009 http://challenge.semanticweb.org/
[64]ISWC 2009 http://iswc2009.semanticweb.org/

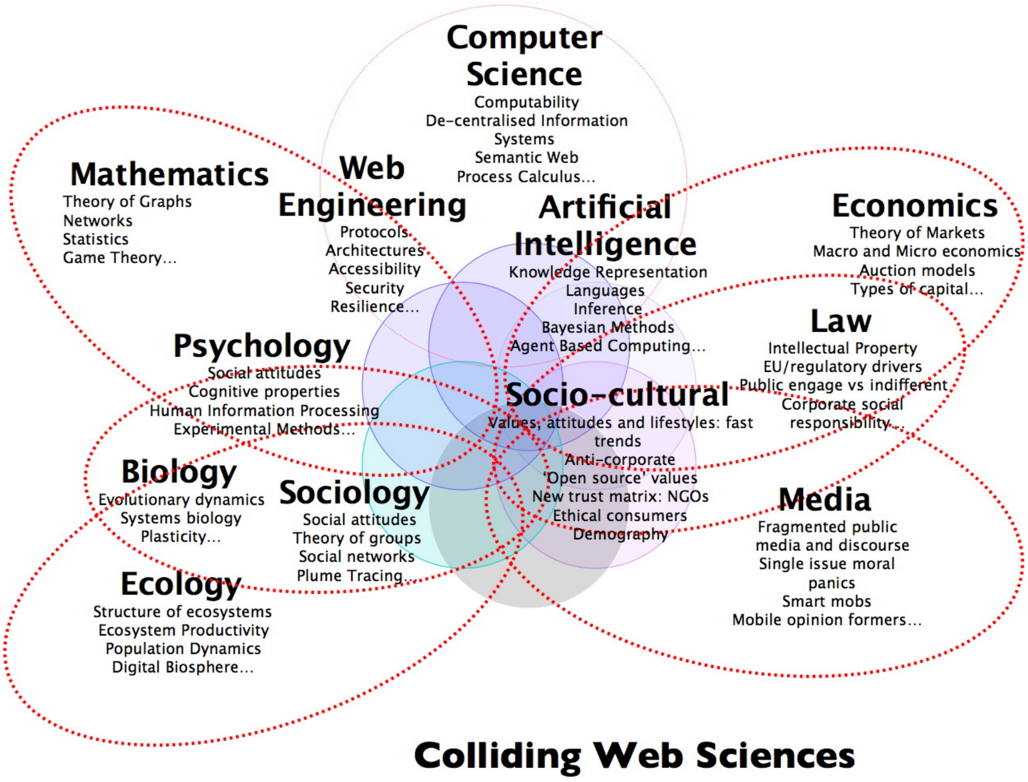

FIGURE 3.5: Web Science – Collision of Disciplines

3.4 Memories for Life and Web Science

Memories for Life (M4L) has been championed as a grand challenge for UK computing. M4L focuses on the exploitation of technology to support human memory. M4L draws together a number of recent technological and scientific advances from both life sciences and computer science in the area of human and artificial memory. M4L focuses on the use of technology alongside human memory in context, to provide support for memory management. The focus of this work in relation to M4L is that of addressing the challenges of capturing, storing, and exploiting autobiographical metadata to support memory management. Given a system that could store a comprehensive collection of a lifetimes worth of acquired electronic media we detail a unified method of marking up this inherently heterogeneous data-set. This is studied to advance insight into the availability and utility of metadata as a way to sort through one's personal multimedia collection. Each of the sources of information modeled in our approach are presented as instances of events making up the user's *digital persona*.

The work undertaken towards the development of this thesis is presented as a engineering task that attempts to push forward the boundaries of personal information management. The adoption of the somewhat indiscriminate ethos of lifelogging is re purposed within the more structured and orderly world of personal information management. This thesis is presented as a synergy of computer science, the discipline of engineering, and the humanities and is said to shed insight into the way that both personal information and privacy will be perceived in the years to come. The interdisciplinary nature of this work coupled with the fact that any study/attempt to develop a widely used Lifelogging system would have to touch on the fields of ethics, privacy, anthropology, as well as web-based engineering and computer science is why it is said to fall is said to fall under the banner of Web Science[65]. The Web is the largest human information construct in history, and it is growing by the day. Given the increasing levels of participation on the Web, and proliferation of the mobile web one does not have to look far to witness how profoundly it is influencing society. Given the fact that one cannot travel far in the Western world without being in close proximity to a networked device, the Web Science Research Initiative has been championed by the University of Southampton[66] and Massachusetts Institute of Technology (MIT)[67] in order to :

- Understand what the Web Is
- Engineer its Future
- Ensure its Social Benefit

As a newly formulated discipline Web Science sets out a research agenda whereby the study of the networked environments, such as the Internet, are said to energise synergies between the various traditional fields of study by predominantly acting as a enablers for collaborative research. This notion of uniting the arts and the sciences is illustrated in figure 3.5.

[65]Web Science Research Initiative `http://webscience.org/`
[66]The University of Southampton `http://www.soton.ac.uk/`
[67]MIT `http://www.mit.edu`

Chapter 4

Narratives, Stories, & Ontologies

This chapter presents narrative, specifically the study of narrative structures, as an important mode of knowledge representation. It is the form most familiar to people and the form most adopted by our personal and social archives. Ontological representations of knowledge may be of high importance to machines, but if those machines are to successfully access human records and communicate to human users, then it would be advantageous for them to have an understanding of narrative. Given the personal nature of lifelogged information, the notion of logging events within a narrative structure is presented as the core contribution of this chapter. For completeness and consistency, and adherence to the motivations put forward in this thesis it is presented that the most appropriate form for this understanding is itself ontological. In this chapter there will be some discussion around the study of ontological narrative structures followed by the outcome of the work – the OntoMedia[1] (Jewell et al., 2005a; Lawrence et al., 2005) ontology. The event based structure presented in this chapter will be applied to the knowledge captured in a given user's lifelog in the subsequent chapters.

4.1 Towards Ontological Narrative

Given the Semantic Web vision of encoding knowledge in a machine readable manner, this chapter aims to motivate the adoption of narrative structures as a key method of modelling information. Given the sheer amount of data being published to the Semantic Web (see section 3.3.4) and the description below of the use of narrative structures to transmit knowledge between people, the OntoMedia ontology is presented. Given that the machine-readable data of interest to the concept of lifelogging is one of fact related to human-beings, and that humans are intended as the end consumer of the

[1]The original version of the OntoMedia ontology has been removed from a now defunct server. OntoMedia has been given a permanent home at the following address http://purl.org/ontomedia/

data, OntoMedia presents a method of representing machine-readable as a succession of events on a given timeline.

Narratives have long been considered a primary way in which human beings communicate with one another. The traditions of oral storytelling that have evolved into our contemporary modes of narrative have been recognised as the basis of transferring knowledge within societies (Campbell, 1949). Narratives have also been identified as a central part of how humans learn to make sense of the world around them (Nelson, 1989), and to interact in social situations. The word narrative itself stems from the Latin root *gna*, which also is the root of the word knowledge.

The transfer of knowledge through narrative has been illustrated by studies of different narrative mediums and the identification of similarities between the modes of transfer (Campbell, 1949; Murray, 1998; Bal, 1997).

Other research into the ways that we make sense of our world has resulted in the term *Narrative Intelligence* (NI) being coined by Michael Travers and Marc Davis at the MIT Media Lab. This is the notion that humans organise and make sense of events by placing them into more-or-less familiar narratives (Blair and Mayer, 1997). NI has been identified as one of the main synergies around which Artificial Intelligence (AI) research into narrative has been brought together (Mateas and Sengers, 1999).

At the same time the Semantic Web (SW) vision (Berners-Lee et al., 2001), and technologies (Shadbolt et al., 2004), have challenged the manner in which authors publish information; from the classic method of developing a document that is intended to convey a message to a human reader, to the publishing of "nuggets" of raw knowledge in the form of annotated multimedia items that are linked together in a structured and meaningful manner for machine communication.

The Semantic Web defines the necessary relational models for describing resources with context independent standards, such as the Resource Description Framework (Manola and Miller, 2004) (RDF), but it is the use of ontologies that forms the cornerstone of SW interoperability.

Ontologies are paving the foundations for the realisation of the Semantic Web (SW) vision (Berners-Lee, 1999) by capturing knowledge in a machine understandable language, such as the Web Ontology Language (OWL) (McGuinness and v. Harmelen, 2004). These conceptualisations of different domains are being harnessed to annotate documents for a variety of tasks. The OntoMedia ontology aims to provide a meaningful set of classes and relationships to facilitate the annotation of the semantic content of heterogeneous multimedia items[2] (see section 4.3). OntoMedia presents a method of encoding the story told within a piece of multimedia.

[2]The term multimedia items is used to refer to text documents, video and audio streams, pictures, etc.

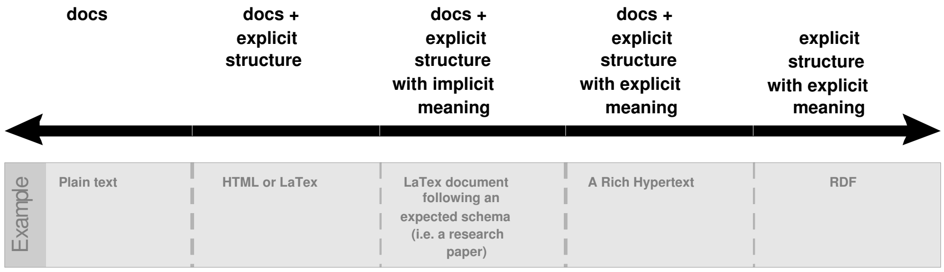

FIGURE 4.1: Narrative and the Scale of Formality

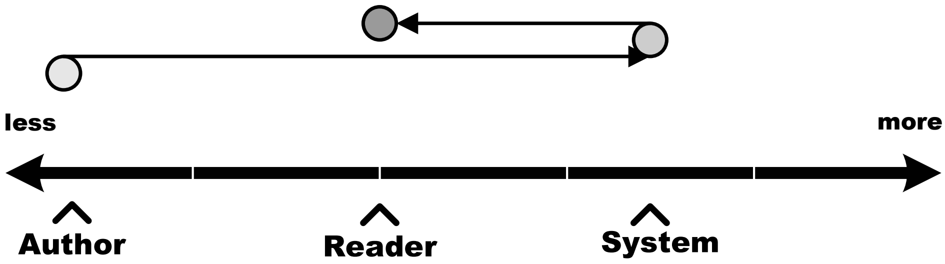

FIGURE 4.2: The ArtEquAKT System

Given the widely adopted Stuber *et al* definition of an ontology: "An ontology is a formal, explicit specification of a shared conceptualisation" (Stuber et al., 1998), this chapter will present the phenomenon that OntoMedia aims to model highlighting how this differs from existing models.

This machine to machine paradigm is a challenge to the way in which human beings have learnt throughout their history to express knowledge, as it requires them to formalise their intentions in ways that may seem quite alien. For example, two of the most influential methods of interacting with SW enabled data-sets are faceted browsers, such as mSpace (m. c. schraefel et al., 2003), and knowledge data navigation techniques presented by systems like Haystack (Huynh et al., 2002) and work undertaken by Lynda Hardman's group at CWI[3] (Rutledge et al., 2005). All of these methods require an understanding of the underlying data-structure, and a grasp of "categorisation by concept" that cannot be expected from all end-users.

This tension between semantics as expressed by people and as expressed by machine has been described as the Semantic Gap (Millard et al., 2005). One of the ways in which the Semantic Gap can be bridged is by creating systems that are able to deal with knowledge at both a human and machine level. Such systems may reason about knowledge in the form of ontologies and tightly defined semantic networks, but express that knowledge to human users in the form of traditional narratives.

Semantic Web research has been primarily focused on the development of mechanisms for machine to machine communication. The high level of formalism required for knowledge

[3] Centrum voor Wiskunde en Informatica http://www.cwi.nl/

to be shared between machines poses problems where human intervention is needed. The two key stages of human intervention are at the authoring and user level. There is a growing trend in systems that towards a "translation to narrative" approach. This is when a system's knowledge base is encoded formally (e.g. in an ontology) and then converted to a multimedia presentation. This reduces the assembled narrative to an interface, not a genuine way of structuring information hence accommodating end-users that do not wish be to exposed to the formal notations underlying such a system.

ArtEquAKT (Alani et al., 2003b) extracts information from unstructured narratives, the World Wide Web, populates an ontology and subsequently combines the extracted knowledge into an adaptive hypermedia document. Figure 4.1 presents a qualitative scale of the formality of various document formats, and is presented along side figure 4.2 that shows how the ArtEquAKT system processes knowledge. The illustration highlights the form of knowledge from the author's, system's, and reader's point of view. That is, in the ArtEquAKT system the author generating the biographies can interact with the system using plain text, whilst the system has a semantic understanding of the facts inside its knowledge base, and the reader or consumer of the output gets displayed with semi-structured HTML document.

The ArtEquAKT system generates biographies of artists from information gathered from the Web. Not strictly a hypermedia application in the traditional sense, it does use both hypermedia and semantic web technologies as part of a larger composite architecture. Figure 4.2 shows quite a distance in placement between the three noted points on the scale. The 'authoring' process involves extracting information from documents on the web (the text rather than the networks of links). The distance between Author representations (basic web pages and text) and System representation (explicitly structured knowledge with linked textual fragments) represents the process of knowledge extraction that is taking place. Implicit structure and meaning in the web document's text is being automatically extracted and recorded into explicit knowledge as an ontologically structured knowledge base. When this is presented back to the reader, the knowledge is combined and published as a hypermedia document with a range of adaptive features. This process provides document and explicit structure to the reader and some explicit meaning becomes implicit in the content of the document produced (via stretching and dimming text).

To do this, machine systems are required to have an understanding of how narratives are structured, and how "nuggets" of knowledge might be sensibly combined. For this reason Narratology, the study of narrative, has become increasingly popular in the field of knowledge technologies (Alani et al., 2003b; Geurts et al., 2003; Little et al., 2002; Silva and Henderson, 2005; Gemmell et al., 2003; Rutledge et al., 2005; Bilasco et al., 2005), and knowledge management (Connell et al., 2005).

Here it will be presented that the best way for ontological machine systems to parse or produce narratives is for them to have an ontological understanding of narrative itself.

This chapter is an attempt to move toward this ontological understanding. Narratology has produced a growing set of different narrative theories, ways of conceptualising narrative spaces, such as formalism, structuralism, post-structuralism, and post-modernity (Schärfe, 2004). The aim of this work is not to develop new theories, this is best left to narrative theorists, but instead to learn from and harness existing methods to aid narrative generation and management in ontological environments.

The following section takes a look at the ways in which narratives have been supported in computational systems in the past, based on existing narrative theory. Following on from this is a discussion around a view on the different levels of narrative that could be represented in a machine-readable format. Finally, OntoMedia is presented as an effort to model one of these layers this coupled with the design decisions taken are presented along with example uses cases of the ontology in practise.

4.2 Computational Models of Narrative

Up until now discussion has been focused around the already overloaded term *Narrative*. This section aims to present an overview of computer science disciplines that have harnessed insights from narratology to motivate future research in the field.

The underlying importance of narratives to human memory and communication has long inspired Artificial Intelligence (AI) research. Roger Schank's group at Yale pioneered work on story generation and parsing in the 1970s and 80s. Generation systems like TAIL-SPIN (Meechan, 1976) and natural language processing systems like PAM (Wilensky, 1981), produced insight into what form of explicit knowledge is needed to comprehend natural language (Schank et al., 1980). Techniques developed by Schank's group highlighted the fact that the meaning of a sentence cannot be appreciated without having knowledge of the context it is in.

During the "AI Winter" (Russell and Norvig, 2003, p. 24) funding for knowledge-intensive systems was few and far between. The problems these systems attempted to tackle were presented in fit-to-purpose closed world domains, arguably making them not feasibly scalable for real world problems. This assertion is what brought about the lack of funding opportunities within the AI community.

After the funding drought of the AI winter, new momentum reinvigorated narrative studies in AI research. One example is, Interactive Fiction (IF) (Murray, 1998). This research field attempts to create interactive stories, allowing for user decisions to take place within a coherent narrative. Various methods have been pursued to achieve this ambitious task of generating interactive narrative. Techniques span from multi-agent based approaches, where modelling of human-like characters deployed within a given environment in hope of observing emergent narratives (Mateas, 1997; Project, 1989-2002; Cavazza et al., 2002; Riedl and Young, 2003), to knowledge based approaches that attempt to define narrative models that facilitate interactivity (Szilas, 2001).

The field of human computer interaction (HCI) has produced a number of arguments for the use of narrative as a mode of interfacing. Systems have employed techniques from oral storytelling to help organise multimedia interfaces and have also used analytic categories of Aristotelian dramatic theory to influence interface design (Laurel et al., 1991). The field of deliberative agent design have also sought insight from narrative theory. Research has suggested that agents will become more intelligible if their memories are structured into narrative constructs (Sengers, 1999; Meech, 1999). The understanding and identification of narrative structures have also been explored to aid the dissemination of knowledge within organisations (Schank, 1997; Connell et al., 2005). Knowledge management has identified narratives as one of the ways that tacit knowledge is transfered within organisations (Ball and Ragsdell, 2003).

Advances in techniques for the assembling of multimedia presentations based on ontological structures has been documented in recent works (Little et al., 2002; Kim et al., 2002; Geurts et al., 2003; Rutledge et al., 2003). The assembly of multimedia into structured narratives from a collection of knowledge elements will only be enriched by the wide adoption of the Semantic Web vision, and the subsequent availability of more annotated "knowledge nuggets". Ontologies provide shared conceptualisations of given domains, allowing for these resources to be reasoned upon from heterogeneous sources. Advances in information extraction techniques (Ciravegna, 2003; Ciravegna et al., 2004), the World Wide Web, and the uptake of ontologies are presenting AI researchers with access to a vast amount of knowledge. AI research has now shifted from tailor made domains, to "open-world" scenarios, where issues such as provenance, inconsistency, and validity of data have to also be considered.

4.2.1 A Simple Taxonomy

This section presents a method of classifying approaches taken towards narrative generation. This simple taxonomy is by no means the only way of conceptualising the efforts of the community, it is employed to illustrate future research directions.

Character Based Systems that generate narratives from character based approaches do so by modelling the intricacies of complex characters, simulating interactions between them, in order to bring about emergent narratives. Agent-based computing is the most common paradigm used when designing these rational entities. This approach often fails to generate interesting narratives. This is due to the fact that modelling human characteristics in an agent's reasoning system is not a tractable task. Reports suggest that systems endeavoring emergent narratives often result in unexciting and undirected stories (Matheas and Stern, 2003), this is usually put down to the fact that the systems do not contain an explicit model of a narrative structure to direct and maintain consistency of a story arc (Mateas, 2000).

Plot Based or narrative-structure based approaches are methodologies that attempt narrative generation/understanding by adapting and proceduralising narrative theories. These systems are built around explicit narrative structures. Rule-based methods are the most common way of representing such knowledge (Szilas, 2002). A number of different narrative structures have been implemented computationally to aid narrative generation. These include rhetorical structure theory used to generate video documentaries (Bocconi et al., 2005) and to aid technical writing (Silva and Henderson, 2005). Propp's functions (Propp, 1968) a procedural formalisation of fairytales has also been used to steer interactive fiction along a consistent and dramatic storyline (Szilas, 2001). This approach is a knowledge intensive one, and has been referred to by Szilas (Szilas, 2002) as the "temporal unfolding of a non-temporal structure".

User Modeling Based This method of narrative generation, like the plot based approach, is a knowledge based one. The difference being in the knowledge that is modeled. In the plot based approach the explicit conceptualisation is of the narrative structure, whereas in this case the specialised knowledge is to do with the end-users preferences, or "user-profile". These systems usually incorporate an explicit narrative model and utilise it along side any available knowledge from the "user-model" to set the context and drive the outcome towards a targeted narrative (Bailey, 1999; Bilasco et al., 2005; Bernstein, 2001).

4.2.2 Ontological Narrative

In order to represent narrative ontologically it is first necessary to have an understanding of what aspect of narrative is being modeled. Possibilities include (amongst other things) the events depicted in the narrative, the structure of the story itself, the intended meaning of the narrator, or the perceived understanding of the reader.

Bal's layered view of narrative (Bal, 1997) is a useful way in which to understand what is being modeled. This states that narrative can be viewed as consisting of three layers, these are depicted in figure 4.3.

The lowest level is the *Fabula*, this represents the raw chronological events that are being depicted. The second level is the *Story*, this is the subset of the Fabula restructured into

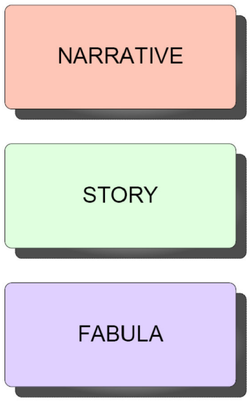

FIGURE 4.3: Bal's Three Layers of Narrative

a new sequence for a particular effect, for example, to create plot lines, to build tension or expectation, or to inform the user about the background of a topic. For any given fabula one could derive a number of different stories. At the third and highest level is the *Narrative* itself. The narrative is the story given form, with all the added semantics of the form itself. Any given story could be turned into many different narratives, for example a monologue, novel, film or multimedia presentation.

These three levels have analogies to the three methodologies identified previously in section 4.2.1 within computational narrative systems. Systems that deal with objects and events can be said to be concerned with Fabula, systems that deal with the structure or arrangement of content can be said to be concerned with Story, and systems that attempt to model the effect and impact of the final article on the user could be said to be concerned with narrative. It should be stressed that plot based systems, Story structure driven, must also have an understanding of the fabula, in order to make sense of its knowledge base.

It should be stressed that non-linear stories do also fit into Bal's three layers of narratives, as Bal's model does not require a story to be a chronological succession of events. A non-linear story, like it's linear counterpart, is said to be a collection of facts, taken from a given fabula and arranged in a manner used to tell a given story.

4.2.3 Ontological Models of the Fabula

Semantic Web technologies allow for the annotation of multimedia items, resulting in a corpus of available knowledge nuggets. If this annotation describes the entities and events within the multimedia items, then they would become a fair representation of the fabula.

As an example, one may annotate a short video sequence with details of who appeared in the video and when, and what their interactions were. Our description of the order of events would be independent of how they appeared in the video and would relate to their actual chronology (whether real or fictional). We could also consider the markup of news websites, identifying the entities and events described in each article. These annotated media items could be re purposed and assembled to present new stories, similar to the Bocconi's work on the generation of video documentaries (Bocconi et al., 2005). Other work in this space was undertaken where annotations of the events within feature length movies where made in order to cater for the automatic generation of movie soundtracks, this work is presented in (Jewell, 2007).

Ultimately the plan is to describe occurrences in such a way that they could form a non-derivative fabula, which may lead to new stories, and thus the generation of new multimedia items.

Modeling at this low level is advantageous as it provides a base point from which new stories can be generated and existing ones analysed. It is also useful as it is the level at which raw information is expressed without being polluted by authorial intention. All of the systems described in this paper rely on methods of describing the knowledge bases, fabula. Ontologies and Semantic Web technologies provide the backbone for annotating heterogeneous media, *the fabula*, presenting AI researchers with the challenge of problem solving in a distributed "open-world" scenario.

4.2.4 Ontological Models of the Story

Ontological models of the middle layer, the story, are concerned with the structure and thus the purpose of the arrangement of fabula items. Readers have expectations about the way in which stories are arranged, often based on genre (Shneiderman, 1997), this structural knowledge is what needs to be modeled (Falkovych and Bocconi, 2005). In order for such models to work it is important to make sure that necessary semantic threads can be found in the fabula. The advent of finding and threading available semantic relationships within a knowledge base will be dependent on the quality of annotations. That is, if a knowledge base only contains metadata regarding the time and date of its media items, the only relationships that could be found will be chronological in nature.

A common way of enforcing this higher level structure, and thus managing the expectations of readers, is to use story grammars – most commonly implemented in the form of templates (Alani et al., 2003b). Genuine grammars contain rules about how parts of a story may be arranged for a given genre, templates circumvent these rules by defining a rigid structure that must be populated by items from the fabula. A grammar based on the Toulmin model (Toulmin et al., 1984) has been implemented to generate rhetorical argumentation for video documentaries (Bocconi et al., 2005).

A major shortfall of existing narrative generation systems is the use of story templates as grammars. These static structures have to be defined by developers before the deployment of a system. This limit's a system's ability to adapt to the content of the fabula and in particular to discover any new relationships to render into a *narrative*. Motivations into the identification of genres and their salient features, have been highlighted as key to the design of story grammars in (Falkovych and Bocconi, 2005), allowing for systems to be less domain specific.

4.2.5 Ontological Models of the Narrative

Once a story has been defined based on a fabula it must still be presented through some medium. This final layer is the narrative perceived by the human reader. Even at this level their will be semantic effects resultant from presentation choices. For example, in cinematography different cuts are known to imply different things to a viewer, such as a slow fade being used to indicate the passage of time.

Ontologies of this upper layer will be dependent on the form, and may even be derivative of the story genre (for example, one should expect different presentation methods in a documentary than they do in a drama). Scott McCloud's (McCloud, 1994) six point categorisation of the different panel-to-panel transitions used by comic book writers to tell stories is currently under investigation as a method of juxtapositioning images to aid narrative generation (Tuffield et al., 2005a). OntoMedia is presented as a method of annotating arbitrary stories, and as a result can be used to describe any secession of events, whether real-world or fictional. In the future work section, discussion around how a methodology based around OntoMedia could be used to not only annotate stories but to generate them from a corpus of events (see section 8.2.2).

4.2.6 Modeling the Fabula

Information-seeking on the Web is currently mainly done through search engines; one normally "googles" a given subject, and subsequently traverses a list of related documents, looking for the best match. There is no method of querying the Web with a topic that would generate a narrative, something that approaches the rich and engaging overviews that a human may deliver.

It is possible that this short fall is due to the fact that information posted on the web does not contain the necessary semantics, in an explicit machine-readable manner. This shortfall hinders a computer's ability to reason upon, disambiguate, and infer relationships from this vast pool of information, and present it in a structured and targeted way.

In reaction to the personal information overload problem Memories for Life (M4L) has set out the task of road mapping the issues surrounding the vision of capturing a lifelog. It aims to address the challenges of storing and presenting autobiographical knowledge in the form of multi modal electronic media, and to identify any issues that may arise from such a situation. This chapter presents that ontological methods could be adopted to help solve this problem. Given a system that could store a lifetime's worth of acquired electronic media, a unified method of marking up this inherently heterogeneous data-set is needed. The flexibility of RDF, along with the open natured of the format are key factors in presenting it as a viable solution for the capture of a lifelog.

This chapter attempts to tackle these problems with an ontological model of Fabula, an ontology that is capable of modelling the objects and events described within any narrative found in a collection of multimedia items, so that they can be searched and reasoned upon.

4.3 OntoMedia

The term narrative is used to describe the story that an item of rendered media is presenting. The Semantic Web vision is challenging the manner in which we are publishing content, from a manner suitable for solely human consumption, to the publishing of items of raw knowledge in the form of annotated multimedia items, linked together by a common model, in a machine processable manner. The availability of such semantically enriched artifacts would allow for narratives to be generated in a manner targeted to the user preferences (Bilasco et al., 2005). A discussion of the relevance of the application of narrative theory to Semantic Web enabling technologies is presented in (Tuffield et al., 2005b). OntoMedia is presented as an ontology which allows for the representation of arbitrary stories in RDF. The work is presented with focus on its flexibility to annotate all sorts of stories, and will be used in future chapters to present how OntoMedia fits in with Lifelogged data.

Of the ontologies used in the Linked Data Web (see appendix A.1) there are many which are to do with bibliographic data which are used to add metadata about published work, categorising it and attributing the work accordingly (Saur, 1998). OntoMedia attempts to tackle a different but complimentary problem of annotating the events within a story, where the story could be anything from a plot of a hollywood movie to the mundane intricacies of day to day life.

4.3.1 Application Domains

This section presents a few application domains which have made use of OntoMedia to represent their specific narrative needs. The section will point to the related literature on this topic to give the reader a feel for the scope of the OntoMedia ontology.

4.3.1.1 FicNet

FicNet (Lawrence, 2007) was a human-computer interaction project undertaken at the University of Southampton to investigate the best way to present RDF data found on the Semantic Web to end users. The application domain brought forward by the FicNet work was that of facilitating online amateur writing (Lawrence and m. c. schraefel, 2005), which took both the form of collaborative and individual writing. Further to acquiring insight into the best mode of presenting RDF data in a web based graphical user interface, the work also tested the robustness of the OntoMedia ontology and its suitability for the representation of arbitrary stories.

From the study above a number of requirements were drawn up that fed directly into the early development of the OntoMedia ontology. The first of these was the need to describe the media objects that were created by this community in terms of both bibliographic detail and content. While the majority of these media objects were textual works of fiction others included images, video and occasionally music. The bibliographic data could be described by any one of the many vocabularies that already exist such as the Dublin Core or the Functional Requirements for Bibliographic Records (see 2.4.1) but none of these were designed with the intent of describing the internal content of the media. While it is possible to use these models to include information concerning the contents of the media item it produces a less than ideal situation since this was not the primary purpose for which they were designed. The second of these requirements was for the ontology to acknowledge that some of the metadata records could be considered sensitive information in that they would give away important plot information. This spoiler related information is now included in the Fiction extension of OntoMedia.

4.3.1.2 State-Based Sequencer

Concentrating on video-related annotation, SBS (Jewell et al., 2005b; Jewell, 2007) (State-Based Sequencer) is a project for the automatic composition of film soundtracks. The composition process is parameterised using a marked video and a 'composer representation' which denotes how aspects of a film should be represented in the music. For example, it can be specified that a certain colour should signify the introduction of a different instrument into the resultant music. OntoMedia is being utilised for the annotation process, and this is then mapped using SerQL queries into the final parameter file for input into a set of composing algorithms. The culmination of this research can be found in (Jewell, 2007).

4.3.1.3 Cultural Heritage

OntoMedia is also being directed as a result of both current and prior cultural heritage projects. This area is a significant driver of ontologies and annotation techniques, and the Sculpteur (Addis et al., 2003) and eChase projects have both been influential towards the design of the OntoMedia classes. The Sculpteur project, itself extending the Artiste (Allen et al., 2000) project, provided metadata for museum collections (specifically 3D items, such as a sculpture) and hence enabled access through a semantic layer. More recently, the eChase project, one which considers OntoMedia and its aim to provide access to cultural heritage material, provides an apt opportunity. Working with the eChase team has resulted in a number of additions to the ontology through extension classes, such as flexible means for denoting regions of media and additional attributes to provide details of ownership and creator. OntoMedia further augments the cultural heritage representation with the availability of a timeline which allows for the placing of events and items within a temporal context - and hence allowing for the generation of historical narratives. Similarly, the Story Fountain system (Mulholland et al., 2004) uses annotated multimedia to produce story paths from the historical archives of Bletchley Park.

A similar of using annotated multimedia is the Story Fountain system (Mulholland et al., 2004) that produces story paths from the historical archives of Bletchley Park.

4.3.2 The OntoMedia Ontology: An Overview

There are currently many overlapping ontologies on the Semantic Web. A Swoogle[4] search for the term "character" would get approximately 95 matches from their repository. This is because people tend to represent the same phenomenon from different view points and not always for differing domains. To help justify the creation and deployment of an ontology, the abstract model needs to present a novel view point of a given phenomenon or a representation of an altogether new domain, for otherwise an existing ontology should be employed or re-factored. This chapter proposes a vocabulary, in the form of an ontology (Noy and McGuiness, 2001) for the annotation of multimedia documents. This has been labeled the OntoMedia[5] ontology (Jewell et al., 2005a). The resulting annotation, or markup, applied to a piece of media will provide, a semantically rich description of the content, in the form of machine-readable metadata. One of the motivations behind making such knowledge explicit, as opposed to simple keyword labeling, common to the current Web, is to capture and describe knowledge that is implicit within the content of the given media unit (Klamma et al., 2005), in a manner less ambiguous to natural language keyword labeling.

[4]MindSwap's Semantic Web search engine, `http://swoogle.umbc.edu/`

[5]Contextus Project URL, taking OntoMedia forward `http://www.contextus.net/`

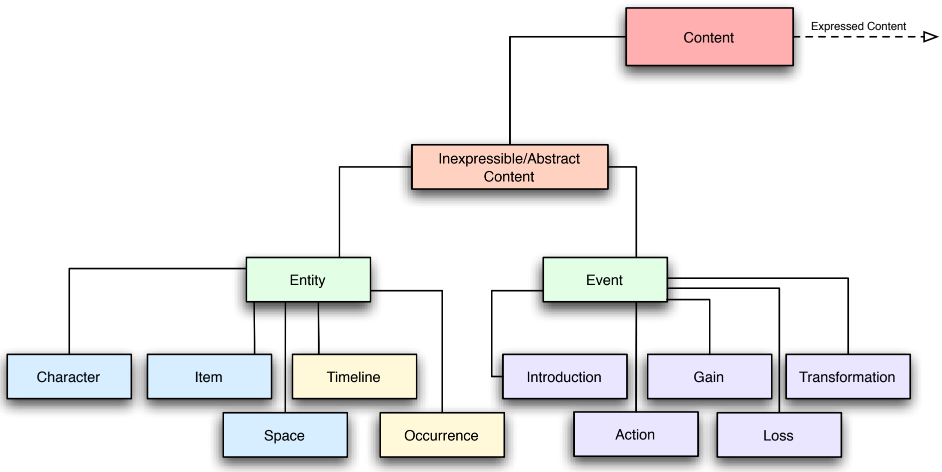

FIGURE 4.4: The OntoMedia Ontology

The OntoMedia ontology aims to provide a meaningful set of classes and relationships to facilitate the annotation of the semantic *content* of heterogeneous multimedia items.

The metadata produced by annotating multimedia with respect to the OntoMedia ontological vocabulary will allow for search and navigation by concept (Jewell et al., 2005a), an example of such a system is Sculpteur (Addis et al., 2005). Sculpteur is a system that allows users to search and navigate semantically enriched museum multimedia; the system also demonstrates how metadata, from a shared vocabulary, in this case the CIDOC Conceptual Reference Model, can be used to search across and navigate through heterogeneous media.

The scope of the OntoMedia ontology is the representation of heterogeneous media through the description of the semantic content of that media item. The representation may be limited to the description of some or all of the elements contained within the source and may include information regarding the narrative relationships that these elements have, both to the given media and to each other.

OntoMedia is presented as a General/Common ontology (Mizoguchi et al., 1995; van Heist et al., 1997) for multimedia annotation, intended for re-use across domains. OntoMedia's structure is based around the core concepts of *entities* and *events*. These two concepts describe the elements present inside a media item, and where appropriate, the interactions between participating elements. This high-level abstraction allows for OntoMedia's flexible ability to describe multimedia. Figure 4.4 shows an overview of OntoMedia's classes, with the division between entities and events made explicit. For full details of the ontology please refer to (Jewell et al., 2005a), and (Lawrence et al., 2005).

4.3.3 The OntoMedia Ontology: The Details

At the center of the OntoMedia design are entities and events. These two classes represent the elements present in a media and, where appropriate, the situations in which they are present. This high level of abstraction is capable of encompassing a wide variety of media, whether factual or fictional, and is not restricted to modelling interactions between people. Figure 4.4 shows an overview of the OntoMedia classes, with the division between entities and events apparent.

4.3.3.1 Event Modelling

An event within a given piece of media consists of one or many interactions between participating entities, and may be instantaneous (happening at a specific moment) or continuous (happening throughout a set period). As shown in figure 4.4 there are three core events, namely `Gain` and `Loss` (e.g. a character learning some information or losing money), and `Transformation` (e.g a character becoming older). The classes of events are based on those contained in the 'typical' story, as discussed by Bal (Bal, 1997) and Chatman (Chatman, 1978). While these may seem simplistic initially, they may be extended - so `Loss` can be extended to `Destruction` or `Betrayal` (the loss of a bond) and `Transformation` can contain `Travel` (locational transformation). The formal definition of an event in OntoMedia is "an interaction between one or more entities during which zero or more traits of those entities are modified and/or a new entity is created".

Furthermore, each event may have preconditions and postconditions which are required for the event to take place or be judged as complete. For example, the event in which a character loses money requires that the character has money initially. This information, though not so useful for inference, is ideal for the markup process as it is then possible to ensure that events only occur when it is possible for them to happen.

To allow for the modelling of event chains, such as a situation that is likely to arise as a result of another, OntoMedia provides ordering properties. An event may cause, or be caused by, another event (not necessarily preceding). Note, however, that these orderings do not impose any timing information on the event objects, as this extra layer is provided by occurrence representations. Occurrences place events into a temporal context. This allows for the same event to occur multiple times in the same media, possibly in several timelines. An occurrence is a straightforward class that provides the location specifier of the start of the event within the medium, and a location specifier of the ending. Location specifiers provide a means to reference portions of media in an extensible and multimodal manner, and hence allows for events to occur within media which are either spatial (such as photographs or comics) or temporal (such as audio or video). The occurrence also contains references to the Event which is occurring, and the

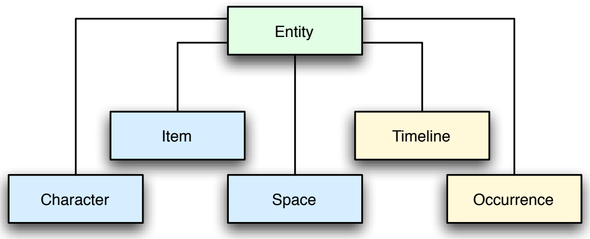

Figure 4.5: The OntoMedia Entity Hierarchy

timeline in which it occurs, and a link to the physical location (see section 4.3.3.3) in which the event was held if need be.

4.3.3.2 Entity Modelling

The entities in a media, which make up the other half of the core OntoMedia classes, represent the items or concepts which participate within the contained events. As such they include physical entities, such as characters or props, and abstract entities, such as language or culture (see figure 4.5). In figure 4.5 the yellow boxes are used to illustrate concepts to do with time and how events can occur on more than one timelines in the case of non-linear stories. The ability to encode event occurrence on various timelines in the OntoMedia ontology caters for the ability to encode different points of view, or different timelines using events common to various actors, whether real or fictional.

The base-level entity provides a few key properties which are inherited by both the abstract and physical subclasses. These include container information, allowing for one entity to be contained by another, location information, which refers to a custom location ontology, and a collection of 'traits'. These are fundamental to the OntoMedia representation, as they embody the characteristics and properties of entities within the media.

OntoMedia defines traits which cover the most common attributes which were found to occur in fictional media. These include personal information, such as age and faith, physical information, such as build and distinguishing marks, and state-based attributes, such as being and form. As mentioned previously, traits may only be altered or added as the by-product of an event, so it is feasible to denote a character's physical appearance altering as the result of a transformation event.

The final, and possibly most powerful, trait within the OntoMedia description is that of 'motivation'. This defines the state that the entity is aiming to achieve in order to gain fulfilment. The trait contains zero or more event instances as well as zero or more entity instances which represent the goal state which the subject aims to realise. By analysing

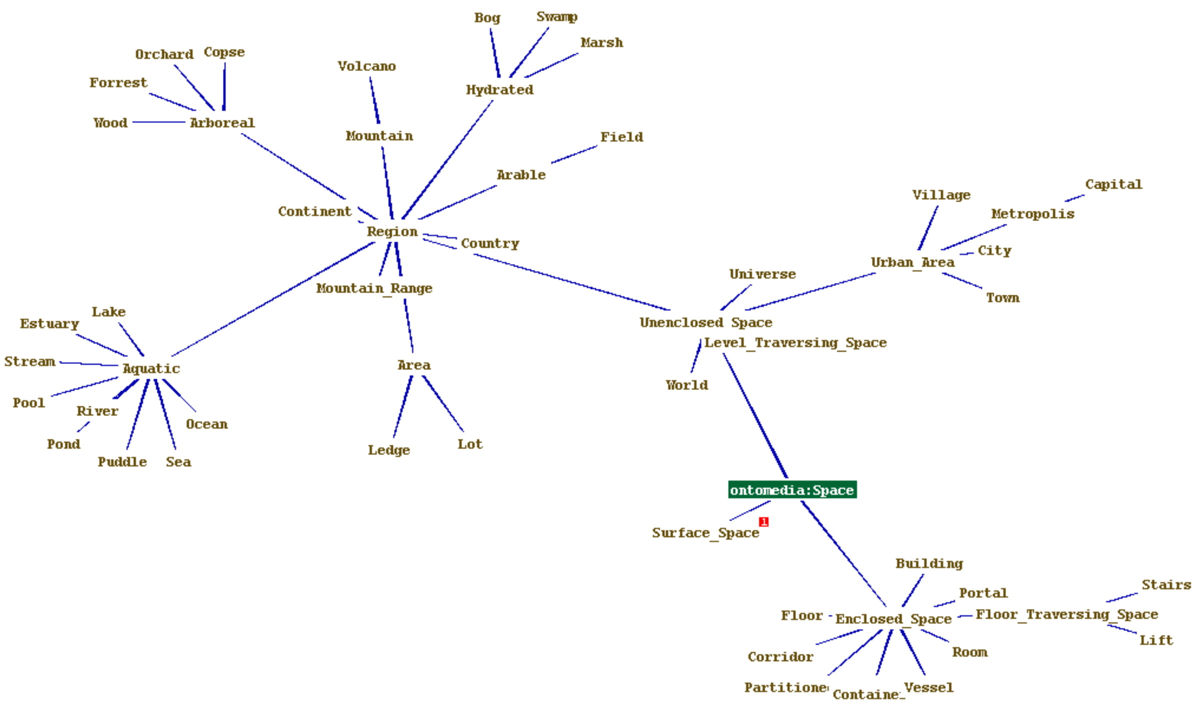

Figure 4.6: The OntoMedia extended location ontology

the state of the medium with regards the motivation of a character, it is therefore possible to determine whether the entity has fulfilled its goals by the culmination of the piece.

4.3.3.3 Extensibility

With the aim of modelling the contents of fiction, regardless of media format, it was necessary to allow for extensibility of the framework. For example, OntoMedia incorporates the location ontology created for the Signage project (Millard et al., 2004) to provide a basic spatial model. This ontology provides the requisite level of detail which arose from example cases that were analysed during the course of the design process. OntoMedia's extensibility is further illustrated by an explanation to how it has been extended by the profession model, which was initially a part of the main OntoMedia ontology.

Focusing on the needs raised by those examples the team extended this ontology where it was insufficient for purpose. The top division between `Enclosed_Space` and `Unenclosed_Space` proved suitable with only the addition of a `Surface_Space` class necessary to cover the those cases where the intended area was two dimensional rather than three. Further classes were also added under `Unenclosed_Space` to match the granularity that existed within the `Enclosed_Space` tree. The extended version of this section of the ontology can be seen in figure 4.6.

The examples used to test the integration of the location ontology were taken from both multimedia sources and literature. However, because the requirements and restrictions were deliberately set at the `ontomedia:Space`[6] level a more or less detailed spatial model could be substituted if required. This is illustrated more clearly in the case of the humanoid body parts ontology. The ontology itself subclasses both `ontomedia:Physical_Item` and `ontomedia:Surface_Space` and was envisioned to be used primarily in conjunction with the `ontomedia:Character_Description` and `ontomedia:Distinguishing_Mark` traits. A work of fiction with a medical setting might require a much more accurate model of the human body while other works might require more detail in certain areas of anatomy.

The location and body parts ontologies were always constructed separately to the OntoMedia model, whereas the profession model was originally created as part of the main ontology. Having tested this arrangement it quickly became apparent that it was not ideal. While having a generic profession ontology was useful as a time saving measure the overlap between the professions required by any two works of fiction was so small and the range of possible professions so large that even the basic break-down into profession types was too cumbersome for the main ontology.

This concept of creating reusable models that could be included when necessary was one that was explored further with our use of the `ontomedia:Context` class.

4.3.3.4 Contexts

The `ontomedia:Context` class was created to separate the many different versions of the same entity that may exist. This is a particular issue when considering the contents of fiction, especially when those works have been reinterpreted across media, within the same work or after a period of time. As these different interpretations may be physically distinctive, for example when a character is portrayed by different actors or given different personality traits, it becomes necessary to recognise that there are occasions when their differences are as important as their similarities. Further discussion around the use of differing contexts within the fictional domain please refer to (Lawrence, 2007). It is suggested that context information could be used to encode different viewpoints when considering the annotation of real-world events, allowing events to be described by different people with potentially different observations to the goings on in an event.

[6] The namespace abbreviation `ontomedia:` is used to denote http://purl.org/ontomedia/core/

Character	Weighted Score
Douglas Quail	105
Interplan	15
Kirsten Quail	10
Martian Politician	7
Alien Invaders	6
McClain	2
Keeler	2
Lowe	2
Shrink	2
SS1	1

TABLE 4.1: Table of characters in the story 'We Can Remember It For You Wholesale', with scores representing their involvement in the storyline.

4.3.4 An example story annotation

Initial tests of the OntoMedia ontology focused on the short story "We Can Remember It For You Wholesale"[7]. As well as containing interesting characterizations and plot direction, the story was also chosen for its changing timelines (some sections are written as the dreams of the main character) and cross-media possibilities (due to its movie adaptation).

Using an annotated version of the story imported into a triplestore (Harris and Gibbins, 2003), several SPARQL queries were carried out. These included the retrieval of events featuring specific characters, the automatic construction of a cast list, the identification of all characters of a specified gender, and the location of 'key' scenes based on spoiler-specific metadata. All of these tests were successfully performed using the OntoMedia representation. The annotated story can be found as a Linked Data resource[8], with a complementary SPARQL endpoint[9].

Building on these already encouraging initial results, a more complex problem was designed. In order to identify important events within a media, it was felt that it would first be necessary to discover presence of key characters. To achieve this, three queries were constructed; the first to enumerate the occasions in which a character was the *subject* of an event, the second to enumerate those cases where a character was the *object*, and the third to identify characters present in a scene but not actively participating. Once attained, these totals where weighted to reflect the significance of the characters' presence in the events, with subject cases multiplied by 4, object cases by 2, and 'involved' cases remaining unchanged.

[7]'We Can Remember It for You Wholesale' is a novelette by Philip K. Dick first published in *The Magazine of Fantasy & Science Fiction* in April 1966

[8]Linked Data Version of 'We can Remember It for You Wholesale': `http://arcadia.ecs.soton.ac.uk/wholesale`

[9]SPARQL Endpoint for 'We can Remember It for You Wholesale `http://arcadia.ecs.soton.ac.uk/wholesalesparql/`

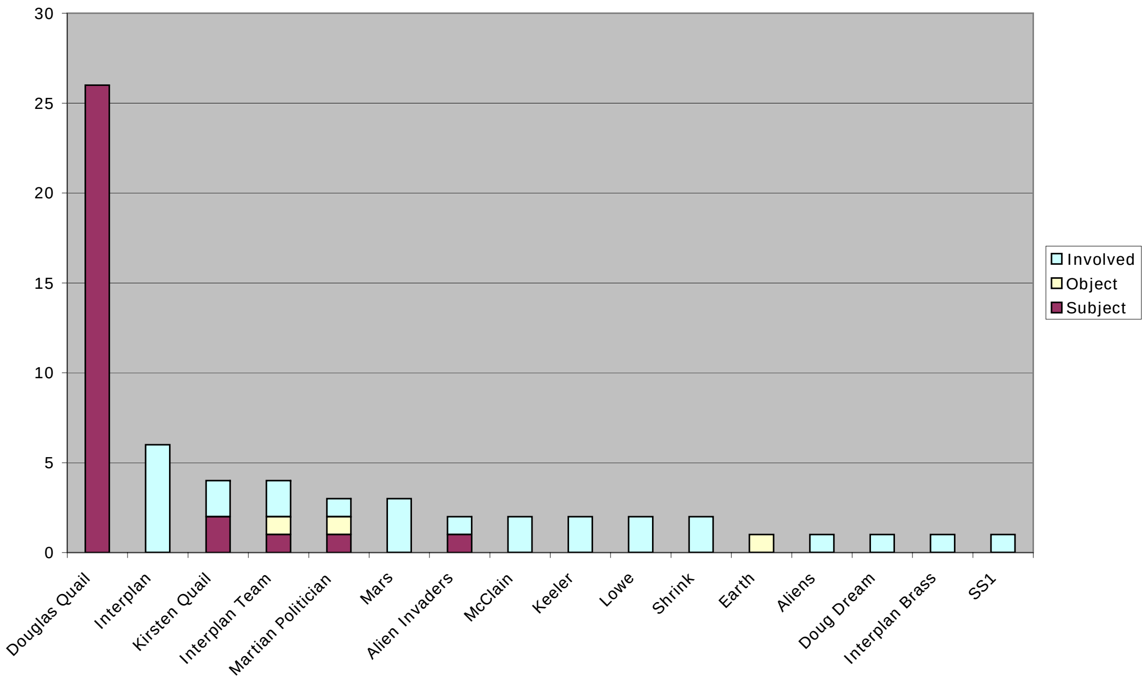

FIGURE 4.7: Chart of characters in the story 'We Can Remember It For You Wholesale', with scores representing their involvement in the storyline.

The results of this calculation applied to 'We Can Remember It For You Wholesale' can be seen in Table 4.1 and in figure 4.7. It is immediately evident from these figures that Douglas Quail is the key character in the story, which is logical as the plot unfolds from his point of view, but it is also straightforward to identify the other major characters (Interplan, Kirsten Quail, the Martian Politician, and the Alien Invaders) as well as those deemed secondary (McClain, Keeler, et al).

4.4 OntoMedia: Summary

An extract of the OntoMedia's formalisation, "OntoMedia Core" is listed in Appendix A of this thesis (B). OntoMedia is split into three parts, the core ontology of which the main section is below, extensions e.g. Fiction, and finally stand alone additions. The ontology can be viewed in its entirety on the project's google code checkout[10]. And a linked data version of the ontology can be found internally within the Southampton's private

[10] OntoMedia Google Code repository `http://code.google.com/p/contextus/`

network. For a description of all of the various extensions to the core of OntoMedia, refer to the project website. From this point onwards this thesis will be concerned by OntoMedia core, and will attempt to apply the principles to the task of lifelogging.

4.5 OntoMedia: Conclusions

Thus far it has been presented that narratives are an important form of knowledge representation, in that they are the existing expression of choice for human authors. They have been identified as key to machine accessible knowledge in both understanding existing human works and expressing new knowledge to human users.

This chapter has explored the way in which narrative systems of the past have tackled the problem of modelling narrative and shown that there are three key approaches: modelling content, modelling story and modelling the user.

The work was grounded in Bal's layered view of narrative, and suggested that any complete ontological model of narrative must address each of the layers: *The Fabula* that describes the objects and events and their chronological interactions, *the Story* that describes their arrangement for a purpose, and *the Narrative* that describes how this is realised in a particular media or form.

The work explored the ontological modelling of the Fabula, and produced the OntoMedia ontology that untangles a Story into a Fabula representation using asynchronous time lines, independent characters and objects, and the transformations that happen to them.

Ontologies are an important method of knowledge representation, but they have serious shortcomings in terms of their ability to capture succinctly the meaning of human to human communication, or *narratives*. It is unrealistic to hope to model everything that is implied in a narrative, granularity is key, how would one model the mind set of Captain Ahab, or the relationship between Macbeth and Lady Macbeth? Ontologies built around existing narrative theory offer a powerful way to tackle this problem at a more pragmatic level, without encumbering end users with additional overheads of conceptualising explicit semantics.

The OntoMedia ontology provides a powerful annotation vocabulary which is both multimodal and extensible. As well as building upon the work done in the ABC and CIDOC projects, it has the capability to handle more elaborate media formats and content. The use of location specifiers provides possible ties into photography, audio, and film, as well as textual information.

Core to the OntoMedia philosophy is the ability to extend a generic framework in such a way to accommodate requirements for media markup. This chapter has highlighted the use of metadata specific to fiction, but it is equally possible to represent the metadata for factual media. Furthermore, the constructs provided for motivation and trait description may be employed both for annotation and analysis.

With the emergence of the Semantic Web, it is increasingly important to annotate data into machine-readable formats. The OntoMedia ontology provides a shared conceptualisation of this domain which will act as an enabling technology for this evolution of the web. It should be noted that although OntoMedia's focus is on the annotation of fictional stories, including its ability to represent non-linear stories, autobiographical stories as said to be subset of fictional ones. These autobiographical stories are said to be linear ones which focus on a given individual relating to activities in their day to day lives.

Chapter 5

Semantic Logger: An Autobiographical Contextual Log

The Semantic Logger (SL)[1] (Tuffield et al., 2006a) is presented as a web-accessible means to capture, store, browse, and interrogate a user's contextual log. This is seen as the first step in the liberalisation of personal information on the web, by empowering users to store and hold their own data.

As mentioned earlier in the thesis the contextual log acquired by the Semantic Logger is grounded in a multimedia asset management system, namely Photocopain (Tuffield et al., 2006b), an image annotation service. Photocopain combines contextual information stored in the Semantic Logger, with content based information extracted from the images to generate metadata for one's personal photo-collection. Details of the aggregation of the *context* and *content* based information are presented in chapter 6, where a personal image annotation service is built upon the Semantic Logger framework.

The capture of a rich and complete lifelog is presented as a task which requires the aggregation of a rich corpus of contextual information about a person. The ability to successfully capture metadata about an event in a contextual log is presented as a way of evaluating the quality and scope of the data gathered. This chapter describes a framework for the capture and archival of a lifelog, making use of web-based standards in an ever online world. Subsequently in chapter 6, the Photocopain application is presented as an example application which exploits data from a user's lifelog to help with the task of image annotation.

The Semantic Logger attempts to address the interoperation of personal metadata by making it available in a standard machine readable form, ensuring the easy of future data portability. It attempts to do this by enabling its users the option to capture and archive an in-depth and thorough collection of facts pertaining to their life, which in

[1] The Semantic Logger: http://akt.ecs.soton.ac.uk:8080/

turn should present a picture of how many digital traces the user is potentially leaving behind.

An implementation of the Semantic Logger operating in a social environment is described in chapter 7 of this thesis. This is but a framework and a prototype whose future potential is expanded upon the in the final chapter of this thesis (chapter 8). The Semantic Logger is presented as a proof of concept for the design of a Semantic Web based solution for storing personal data, and is said to be the predecessor to initiatives such as the foaf.qdos.com system[2]. Thus, for an overview of the foaf.qdos.com system please refer to the following slides http://foaf.qdos.com/slides/london090909/index.html.

5.1 Overview : The Semantic Logger

The Semantic Logger, currently live on the web, is a system for the importing, housing, and harnessing of personal information. The Semantic Logger's utility is grounded in two context-based applications, namely a photo-annotation tool (see chapter 6), and a recommender system. A description of the recommender system designed is not the focus of this thesis, more information can be found in (Loizou, 2009). Furthermore the Semantic Logger framework has been adopted as one of the technology demonstrators[3] in the OpenKnowledge EU Project[4], which is presented in chapter 7 of this thesis.

Upon registration of a Semantic Log the user is presented with client-side tools that allow for the capture and uploading of personal information, and server-side functionality has been implemented to support the archival of the lifelogged data as well as APIs to allow users to interact with their knowledge bases (KBs), namely the SPARQL protocol. The list of information sources is far from an exhaustive one, and is not intended to limit the functionality of the system. The Semantic Logger has been designed in a manner to allow information, in various forms of RDF, to be posted to the user's knowledge base (KB). The sources of information identified and implemented are rationalised by the nature of the services currently provided by the system, and are merely presented as inspiration for future development.

Given the flexible and extensible nature of the framework this thesis argues that, by virtue of knowledge integration alone, added value emerges (as road-mapped by Tim Berners-Lee Semantic Tube Map figure 3.3). The principal support for this argument stems from the power of enabling the application of SPARQL queries on the available information, to answer questions that would be unfeasible under representations of singular domains. Design decisions whereby individual instances are encoded in isolation, i.e. data sources are not coupled with each other, along with the adoption of both

[2]FOAF services http://foaf.qdos.com/
[3]OpenKnowledge Demonstrators http://www.cisa.informatics.ed.ac.uk/OK/drupal/demos
[4]Open Knowledge EU Project http://www.openk.org/

home-brewed and well known ontologies for logging events, are presented as the key contributions brought forward by the Semantic Logger Framework.

The notion of logging information from the various sources identified below in isolation, is due to the fact that people may not wish to log all of the types of information presented in this chapter, and if they wish to log only one of them, it was important to realise that the data made sense in isolation. For example, if a user only wishes to log their geolocation information they should be able to make use of it without requiring the logging of any other sources of information. This loose coupling very key to the design of the RDF generated by the Semantic Logger Framework. This along with the inferential capabilities of the knowledge representation are key to our approach of information integration.

The Semantic Logger does not require the user to produce hand crafted annotations. The existence of various domains in the knowledge base supports the automatic creation of such metadata. For example, iCal entries referring to the same time period as GPS location data can be used to provide suggestions for the name of the place with the specified coordinates. This inferential capability is enabled via the use of RDF as a knowledge representation language. RDF and the adoption of commonly used shared vocabularies, or ontologies will allow for queries similar to the ones below to be presented to one's lifelog.

It should be noted that the RDF generated by the Semantic Logger makes specific claims about a given person being involved at a given event, an event may be generated from an iCal entry, stating that the Semantic Logger's own was at a given event or may be generated more explicitly by a user carrying a GPS unit, an example RDF fragment describing an event is presented in 5.5.

How many users of the system attended the same events as me between time X and Y?

This can be achieved by first selecting all events attended by the user between X and Y, using the iCal data, and then selecting all users with similar entries. If geo data is also available, it can be used to extend and target the query.

How many hyperlinks did I receive in email correspondence that I have yet to visit?

A single query can be used to tackle this, by querying the email and browser history representations.

What document was I reading on the way to event X?

What was the name of the band I discovered while on holiday in Y?

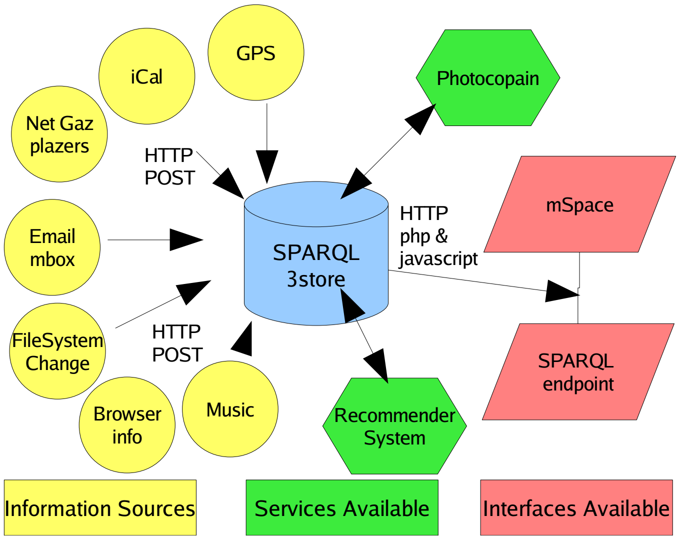

FIGURE 5.1: Overview of the Semantic Logger

5.2 Infrastructure: Semantic Logger

The Semantic Logger is based on a service-based architecture, as shown in figure 5.1, and has been designed so that new services may join on in an ad-hoc manner. The interactions between Web Services have been implemented using HTTP requests, namely HTTP POST, while the interactions with the central RDF triplestore make use of the SPARQL query language (World Wide Web Consortium, 2005), and is subsequently accessed via the SPARQL-PROTOCOL [5]. Given the adoption of Semantic Web best practises, whereby similar ontologies and knowledge representation styles are used across the various domains to encode information, a decision was taken to adopt the Event ontology (Raimond et al., 2007) to represent individual events captured by the Semantic Logger framework.

At the heart of the system is the AKT Project's SPARQL-compliant RDF triplestore 3store (Harris, 2005). The key role of the triplestore is to act as a persistent store for the system, and to mediate the interactions between the other system components. The cornerstone in designing this architecture has been to develop an open and accessible system, so that third parties can exploit the knowledge stored. I have chosen to expose two distinct methods of interacting with the system, namely in a public and private manner. A number of distinct knowledge bases are maintained: A system-wide shared one - the *public kb*[6], (see figure 5.2) and one for each user - *private*[7] (see figures 5.3 and 5.4), which is created automatically for the user upon registration. The screenshots presented are of the backend user interface will allows users to download source code

[5]SPARQL Protocol http://www.w3.org/TR/rdf-sparql-protocol/
[6]The SPARQL endpoint can be found here : http://akt.ecs.soton.ac.uk:8080/
[7]The SPARQL endpoint can be accessed post log in here http://akt.ecs.soton.ac.uk:8080/userHome.php

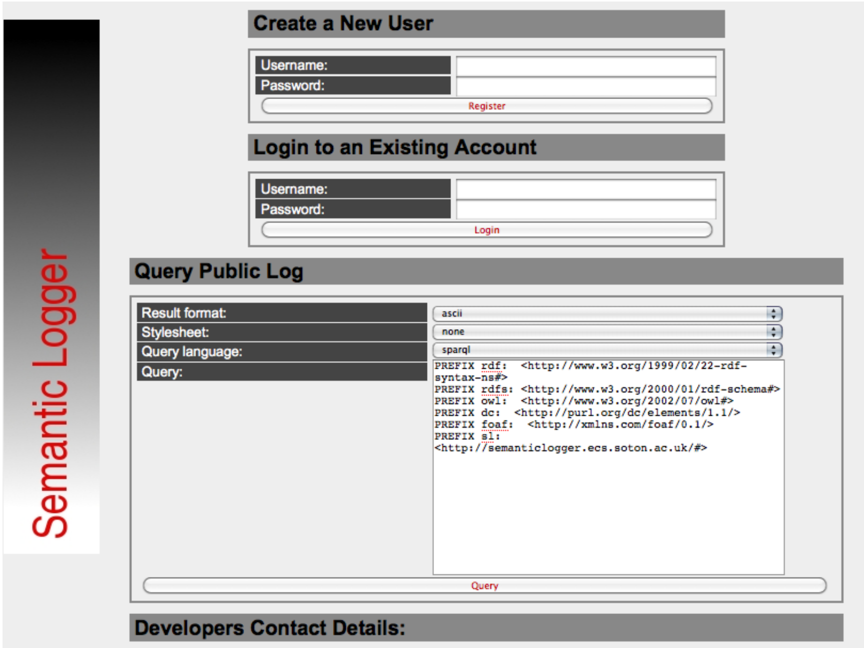

FIGURE 5.2: Semantic Logger's homepage: public SPARQL endpoint and query interface, new user registration, and login

for their personal machines, and allows for SPARQL queries to be issued to both the user's private and the public knowledge bases. The user interface shown is used to support developers who wish to programmatically upload and access their Semantic Logs. When information is imported into the system, users are able to specify whether or not it should be publicly accessible. If this is the case, the information is added to both the shared and private knowledge bases. Both are exposed through web-based user interfaces to allow SPARQL queries on the data and the import, and removal of new knowledge. Furthermore, user interfaces have been designed to support browsing of the knowledge space, as it cannot be expected of system end users to be fluent in SPARQL. It should be noted that ideally all the information would be added to one knowledge base organised by virtue of graph URIs which in turn determine access control. Due to the infancy of access control languages for the Semantic Web at the time when the Semantic Logger was created, the above method of having one central public knowledge base and per user private knowledge bases has the advantage of allowing for SPARQL queries to be issued over everyones public data.

Some thought has been given to the implementation of a similar service which makes use of one public knowledge base and one private knowledge base, where private data is restricted by virtue of model URIs. It is noted that this form of having one private and one public knowledge base would require the service provider to limit full SPARQL access to the private domain so that users can not access other people's private data. There are two ways which one could potentially envisage how a service provide would support a single private knowledge base for its users – a) By providing SPARQL access which in turn limits all queries to the user's graph, or b) by providing API calls to the

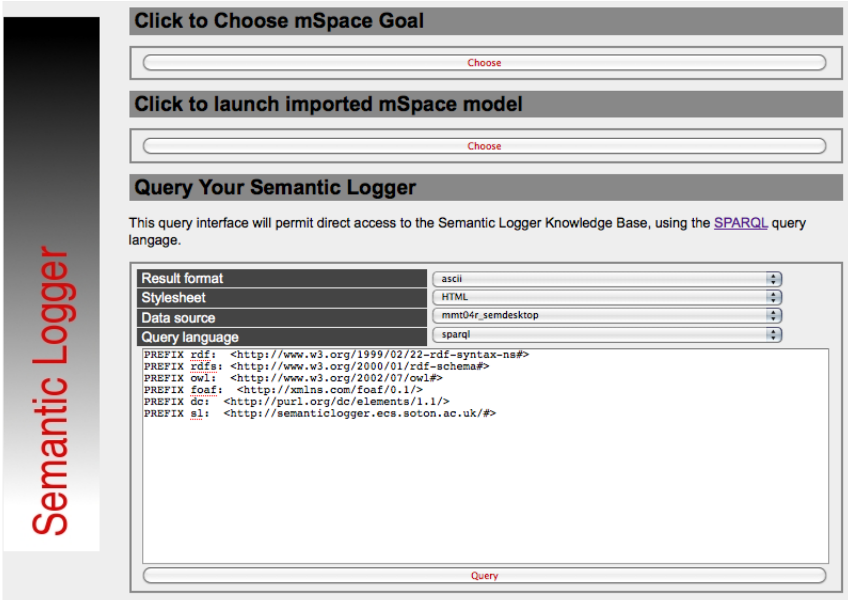

FIGURE 5.3: Semantic Logger's user homepage: private SPARQL endpoint, and mSpace launcher

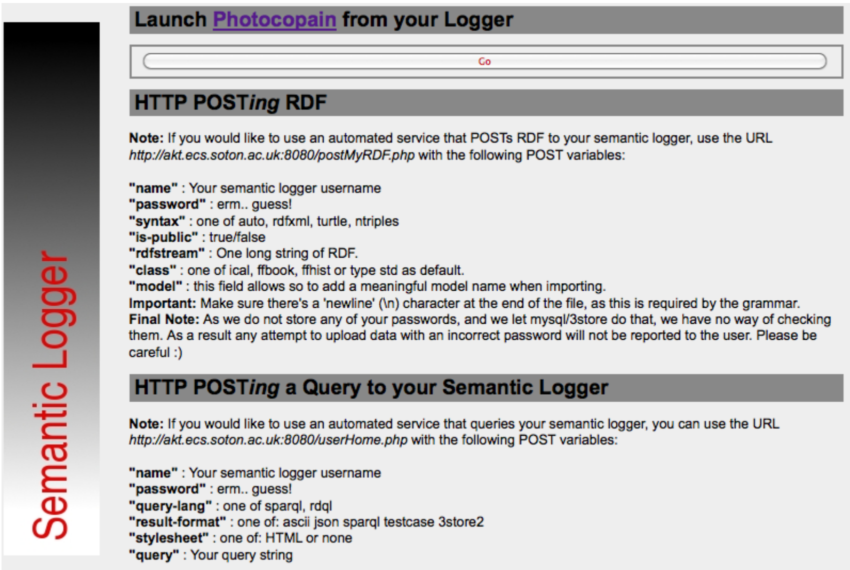

FIGURE 5.4: Semantic Logger's user homepage: private SPARQL endpoint and query interface, and SPARQL protocol

private KB which are engineered to not let users ever have access to other users' private data. This second approach could be used to mine relationships between friends. This is enabled due to the fact that the application developer can write queries over everyone's private data, allowing for a simple way to incorporate social annotation and the sharing of private data.

It is crucial for the Semantic Logger to impose the minimum burden on a user joining the system. Focus has been placed on allowing the import of knowledge described in heterogeneous, widely used vocabularies, to avoid the need for prior semantic agreement. The lack of an overall representation however, introduces the need for alternative means of knowledge integration. Where it is possible, this is to be achieved via automated means, such as the S-MATCH algorithm, developed by the University of Trento (Giunchiglia et al., 2004). Alternatively, where disagreement is too complex to be resolved in an automated fashion, mappings will be hard-coded into applications that use the Semantic Logger as a knowledge source, in *ad-hoc* fashion as per their requirements.

The richness of the metadata acquired, enables the system to be used as a platform for Community of Practice identification (Garcia et al., 2009). For example, named entity recognition can be applied to email correspondence to identify closely related groups while co-authorship and co-reference of scholarly articles can be analysed as shown in (Alani et al., 2003a). Co-location at various events can be inferred from geo-data and calendar entries, while the latter, in combination with the analysis of locally stored multimedia files (e.g. music and video files) can aid in identifying common interests. The utility of geolocation context for the task the photo-annotation is presented in (Naaman et al., 2004a).

5.3 Knowledge Acquisition

The information sources presented in the knowledge acquisition phase are a result of the work undertaken in the Photocopain project, the discussions brought forward from the Semantic Squirrels SIG, and through analysis of the results presented in (Naaman et al., 2004b). Our motivations and interests are similar to that of the Smilie Project at MIT[8], where they host a number of RDFizers[9], tools for converting information into RDF. This work was the foundation for the Piggy Bank project (Huynh et al., 2005) . The methodology put forward by the Semantic Logger differs from MIT's work insofar as it is not focused on supporting web-browser based harvesting of RDF. This work presents an easy to use, scalable, SPARQL compliant, accessible framework for the housing and query of RDF data. This infrastructure is presented along side a number of services to capture and upload contextual information belonging to a given user. The sources

[8]Simile Project: `http://simile.mit.edu/`
[9]RDFizers: `http://simile.mit.edu/RDFizers/`

of information presented in this section are said to cover the sources identified by the aforementioned work but have been constrained by one key motivating factor, of that requiring little to no effort from the part of the user. Manual annotations are presented as cumbersome, and something that the author of this thesis has little time for.

The Semantic Logger collects, and propagates the following types of information into RDF representations, details of the various propagators can be found at Semantic Logger's download page http://akt.ecs.soton.ac.uk:8080/downloads.php. The metadata sources listed below are the foundations for the autobiographical contextual log captured by the Semantic Logger:

- **Calendar entries**

  The Semantic Logger has adopted the W3C recommendation for representing calendar entries in RDF [10]. A client-side application is available for download from the Semantic Logger site to automate the export of iCal (Dawson and Stenerson, 1998) files (commonly used and platform independent) into this representation. In addition to querying capabilities as before, calendar entries can serve as context indicators for geographical locations (described below), enabling to an extent the resolution of co-location.

- **Geo-Data**

  In an attempt to build up a log of a user's geographical data, a two pronged approach is taken. For research purposes we have been carrying around GPS units to log our data, this information is extracted and parsed into an RDF representation, taken from http://www.hackdiary.com/. The RDF model builds on-top of the dublin core namespace[11], and W3C's recommendation for geographical data[12].

  GPS information is being used to track a user's change of location, but is not always a suitable method of tracking, for tall buildings, and movement between buildings within close proximity is hard to track, so a decision was taken to start employing a network gazetteer. Initially the network gazetteer Plazes was employed by the Semantic Logger. Plazes supplies the end user with client side applications that pick up a laptop's current network connection and provides information about the location if information has been entered for that WiFi[13] hotspot. Plazes provides a comprehensive API, and RSS 1.0 feeds, that export parsable RDF, of a users activity.

  The combination of the GPS information, a user's network gazetteer, and his/her iCal file, allow us to infer a user's geographical context. Plazers has become slightly

[10] iCal RDF representation: http://www.w3.org/2002/12/cal/ical

[11] Dublin Core namespace: http://dublincore.org/documents/dces

[12] W3C Geo-Data Namespace: http://www.w3.org/2003/01/geo/

[13] Wireless Internet Network http://en.wikipedia.org/wiki/Wi-Fi

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix event: <http://purl.org/NET/c4dm/event.owl#> .
@prefix geo: <http://www.w3.org/2003/01/geo/wgs84_pos#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix tl: <http://purl.org/NET/c4dm/timeline.owl#> .
<> a foaf:Document .
<> foaf:primaryTopic <http://mmt.me.uk/foaf.rdf#mischa> .
<http://mmt.me.uk/foaf.rdf#mischa> a foaf:Person .
_:mischaBnode0 a event:Event .
_:mischaBnode0 event:agent <http://mmt.me.uk/foaf.rdf#mischa> .
_:mischaBnode0 event:place _:mischaBnode1 .
_:mischaBnode1 a geo:Point .
_:mischaBnode1 geo:lat "50.899837"^^<http://www.w3.org/2001/XMLSchema#decimal> .
_:mischaBnode1 geo:long "-1.395606"^^<http://www.w3.org/2001/XMLSchema#decimal> .
_:mischaBnode0 event:time _:mischaBnode2 .
_:mischaBnode2 tl:at "2009-07-18T11:51:29 01:00"^^<http://www.w3.org/2001/XMLSchema#dateTime> .
_:mischaBnode0 dc:description "Finishing thesis" .
```

Figure 5.5: RDF fragment describing an event

redundant given the definition of the W3C Geolocation API[14], which makes use of the same Skyhook[15] service to translate IP address to Geolocations in an open and standard manner. An example service which makes use of the W3C Geolocation API to output event RDF which logs a `foaf:Person URI` to a geolocation from where the users access the service can be found at the following URL http://mmt.me.uk/geo, see figure 5.6[16].

The RDF fragment[17] presented in figure 5.5 is a `foaf:Document`, whose `foaf:primaryTopic` is http://mmt.me.uk/foaf.rdf#mischa. The above RDF also defines an event, which states that the primary topic was involved with at the event, which has a geographical location and a time stamp.

- **Music playcount information**

  Audioscrobbler[18], is a music identification and logging service. Audioscrobbler runs off of a large collection of music profiles and uses Musicbrainz identifiers for artists, albums, and tracks. The music listen habits or "scrobbles" are captured by users via the installation of plugins for the user's media player, that propagates the information to the audioscrobbler web service. Audioscrobbler makes use of MusicBrainz URIs to uniquely identify Artists, Albums, and Tracks. Musicbrainz[19], a freely accessible dataset for describing the domain of music, which as of August 2009 includes 483,000 unique artists, 725,000 unique albums, and 8,300,000 unique

[14] W3C Geolocation API http://www.w3.org/TR/geolocation-API/
[15] Skyhook Service http://www.skyhookwireless.com/
[16] It should be noted that that a is the short hand notation for the relationship http://www.w3.org/1999/02/22-rdf-syntax-ns#type
[17] Which can be recreated by the following Web Service call http://mmt.me.uk/services/FOAFEvent?lat=50.899837\&long=-1.395606\&webid=http://mmt.me.uk/foaf.rdf%23mischa\&datetime=2009-07-18T11:51:29+01:00\&doing=Finishingthesis
[18] Audioscrobbler: http://www.audioscrobbler.net/
[19] Musicbrainz: http://musicbrainz.org/mm/mm-2.1

FIGURE 5.6: IP Address to Geolocation Service, http://mmt.me.uk/geo

tracks. The Semantic Logger allows for the parsing of a user's Audioscrobbles into an RDF representation. Audioscrobbler is used by Last.fm, the social radio station to recommended music based on communities of practises identified by peoples' music listening habits. Libre.fm[20] is an open-source initiative, which is still in its infancy, that allows for users to run there own server to log their music listening habits, as apposed to having to use Last.fm's servers. Libre.fm would have been adopted for this thesis, but did not exist when the Semantic Logger was implemented. Libre.fm uses FOAF to represent user profile data and makes use of the Event[21] and Music[22] ontology to log songs played. Libre.fm takes inspiration from the dbtune web service which creates an RDF view of a user's Last.fm profile.

- **Web browsing habits**

  By virtue of its cross-platform and open source nature, Mozilla's Firefox has been selected as our web-browser of choice. Firefox exposes the download information in RDF form[23] and thus can be easily imported to the system. Scripts have been developed to parse the bookmarks and history data into RDF. The RDF model uses namespaces taken from the Mozilla developers centre [24].

- **Email**

  A simple RDF schema has been constructed to describe email correspondence [25] as at the time of development one of satisfactory quality was not available. A client-side application has been developed to parse and convert the widely used MBOX

[20] Libre.fm open-source music habit logging system http://libre.fm/
[21] Event Ontology http://purl.org/NET/c4dm/event.owl#
[22] Music Ontology http://purl.org/ontology/mo/
[23] Netscape namespace: http://home.netscape.com/NC-rdf#
[24] Mozilla namespace: http://developer.mozilla.org/en/docs/XUL_Tutorial:_RDF_Datasources
[25] Semantic Logger email namespace: http://semanticlogger.ecs.soton.ac.uk/email/#

representation into an RDF representation. The intended use of this information, in addition to the ability to query one's records is to support the identification of communities of practise, under a predefined temporal context. For example, the representation allows for users to interrogate their lifelog to answer questions regarding the frequency of mail correspondence between different people at various points in time. Given that full-text searching within RDF triple-store technology is still in its infancy, and not currently that usable, Email correspondence allows for the identification of people which the user has conversed with, allowing for events to be contextualise around events such as high levels of email contact between student and supervisor.

- **File System Information**

  Beagle[26] search indexes every file found on a user's computer. This is achieved by combining specialised analysis tools for extracting content from different file types. This creates a personal information space describing a computer at the file-system level. The information is parsed into a simple ontology and can be loaded into a user's Semantic Log. This enables services to detect the presence and usage of files, giving an indication to a user's interests. At the time of developing the Semantic Logger, beagle was the pinnacle of desktop search on linux desktop environments and has since been superseded by the work undertaken by the NEPOMUK project[27], which is now shipped as part of the newest version of the KDE desktop environment[28], KDE 4.3. This Semantic Desktop solution is powered by a SPARQL compliant triplestore.

The information presented above was inspired by the discussions in the SSSIG, and motivated by Tim Berners-Lee illustration presented in section 3.3 of this thesis.

5.4 Semantic Logger and FOAF

A final feature of the Semantic Logger worth mentioning is the way the logger makes use of the FOAF model. A user's FOAF file is used to allow a user to publish data about themselves, using a URI, allowing for the user's data to be referred to from any dataset, or from within any context. The notion of using a `foaf:Person` URI as a unique identifier for a person, has been highlighted in note on the ESW Wiki[29], outlining the concept of a 'person URI' or a WebID. A WebID is presented as a unique web accessible resource used to describe a person. The Semantic Logger presents users with an ability to generate a simple FOAF file if they do not already have one at sign up see figure 5.7.

[26]The Beagle Project: http://beagle-project.org
[27]NEPOMUK, the Social Semantic Desktop http://nepomuk.semanticdesktop.org/
[28]KDE open-source desktop environment http://www.kde.org/
[29]ESW Wiki WebID: http://esw.w3.org/topic/WebID

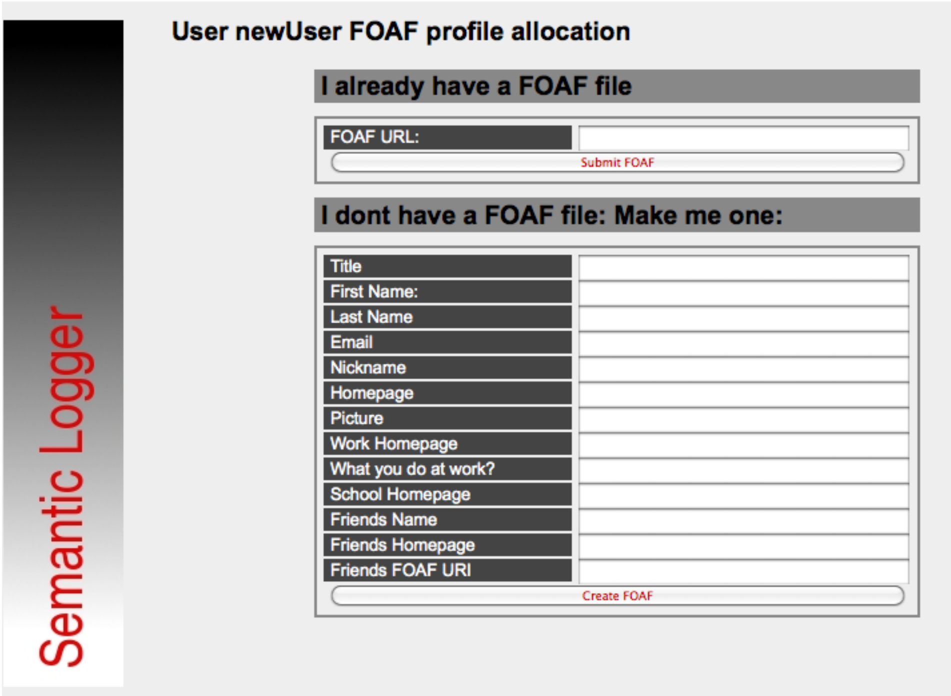

FIGURE 5.7: Basic FOAF generating web interface built into the Semantic Logger

The concept of a web-based UI to aid the generation of a FOAF file has been taken one step further with foaf.qdos.com's FOAF builder UI[30].

Here it is important to stress the distinction between a `foaf:Document` and a `foaf:Person`, which should have two separate URIs, as a `foaf:Document` is an information resource which is an RDF document on the web whereas a `foaf:Person` URI is a non-information resource used to represent a person on the web. A common problem when minting new URIs for RDF data is that of selecting suitable URIs. The distinction between information and non-information resources are key when writing RDF. A URI for an information resource is one which can be sent across the wire, i.e. an RDF document, or an HTML document, whereas non-information resources represent things in the world which can not be sent across the wire, i.e. a URI for a person, or a URI for place. The author's `foaf:Document` URI is `http://mmt.me.uk/foaf.rdf` is an information resource whose primary topic is a `foaf:Person` URI `http://mmt.me.uk/foaf.rdf#mischa`, the WebID of the author. This notion of distinguishing the document (information resource) from the concepts presented within the document is best illustrated by Ren Magritte's painting entitled "*Ceci n'est pas une pipe*" see figure 5.8.

The adoption of a WebID will empower users with the ability to uniquely identify themselves on the web, allowing a given user the ability to associate public data on the web to themselves. Another advantage of the adoption of personal FOAF files is the ability for a user to define his/her friends, allowing for further connections to be made when using

[30]FOAF Builder `http://foafbuilder.qdos.com/`

FIGURE 5.8: *Ceci n'est pas une pipe* Ren Magritte - The Treachery of Images, 1928-1929

```
<http://mmt.me.uk/foaf.rdf> a foaf:Document .
<http://mmt.me.uk/foaf.rdf> foaf:primaryTopic <http://mmt.me.uk/foaf.rdf#mischa> .
<http://mmt.me.uk/foaf.rdf#mischa> a foaf:Person .
<http://mmt.me.uk/foaf.rdf#mischa> foaf:homepage <http://mmt.me.uk/> .
```

FIGURE 5.9: Example FOAF triples

the system to identify communities of practise. This feature supports the incorporation of social annotations as described in Marc Davis's work (Davis et al., 2006, 2004).

Both 3store (Harris and Gibbins, 2003) and 4store (Harris et al., 2009) are quad-stores[31], and not simply triple stores, all of the RDF triples held within the KBs are contextualised with respect to where the RDF data came from. For example, the below triples (see figure 5.9) present a URI for a `foaf:Document`, which has a primaryTopic which is a `foaf:Person`, who has a `foaf:homepage` of http://mmt.me.uk/.

The following quads (see figure 5.10) would be stored within a quad-store which in turn identifies the URI of the document where each of the above assertions originate from. The below notation is based on the non-standard n-quads notation as specified by Richard Cyganiak et al at http://sw.deri.org/2008/07/n-quads/. Whereby the fourth URI is the URI of the document where the triples originated from, this URI is known as a 'model' URI or the provenance URI. This allows the quad-stores, or knowledge bases, the ability to evaluate the legitimacy of the triples based on where the information came from. In the below example we can see that all of the Subject URIs of

[31]Named Graph W3C http://www.w3.org/2004/03/trix/

```
<http://mmt.me.uk/foaf.rdf> a foaf:Document <http://mmt.me.uk/foaf.rdf> .
...
<http://mmt.me.uk/foaf.rdf#mischa> a foaf:Person <http://mmt.me.uk/foaf.rdf> .
<http://mmt.me.uk/foaf.rdf#mischa> foaf:homepage <http://mmt.me.uk/> <http://mmt.me.uk/foaf.rdf> .
```

FIGURE 5.10: Example FOAF quads

Count	Entity Type
http://xmlns.com/foaf/0.1/Person	19, 007, 420
http://xmlns.com/foaf/0.1/Document	4, 918, 936
http://purl.org/goodrelations/v1#ProductOrServiceModel	1, 408, 940
http://purl.org/goodrelations/v1#UnitPriceSpecification	784, 786
http://blogs.yandex.ru/schema/foaf/Posts	585, 188
http://purl.org/goodrelations/v1#BusinessEntity	540, 861
http://purl.org/goodrelations/v1#Offering	439, 130
http://www.w3.org/2002/07/owl#Ontology	436, 809
http://purl.org/goodrelations/v1#TypeAndQuantityNode	421, 255
http://purl.org/goodrelations/v1#ProductOrServicesSomeInstancesPlaceholder	363, 863

TABLE 5.1: Top Ten most prevalent RDF types as per http://pingthesemanticweb.com/ on 14/07/2009

the triples are either the same as the model URI or a fragment of the model URI, and as a result we can safely trust all of the statements made.

Table 5.1 presents statistics taken from the "Ping the Semantic Web" service which maintains a list of URLs for RDF resources found of the web. This tables is used to illustrate how much FOAF data exists on the web, and a big part of the Semantic Web is made up of data pertaining to people. Up-to-date statistics can be found at the following URL http://pingthesemanticweb.com/stats/types.php.

5.5 Making use of mSpace

When data is represented in an RDF graph, by virtue of the representation, there exists multiple dimensions in which the data may be indexed and viewed. The mSpace interface (m. c. schraefel et al., 2003) has the ability to organise such data, in multipane browsers. In addition, the edges of the graph are allowed to be reordered, using dimensional sorting independent of the hierarchical nature of the representation, allowing for a number of such trees to be visualised and browsed.

mSpace requires the definition of a *default column* and a *target column*, along with the path through the ontological relationships (edges in the graph), to create a multi-columned re-arrangeable browser. While in the current implementation of mSpace[32] (0.6.2.3) these have to be made explicit by a knowledge engineer, an algorithm has been

[32] mSpace API definition: http://www.mspace.fm/

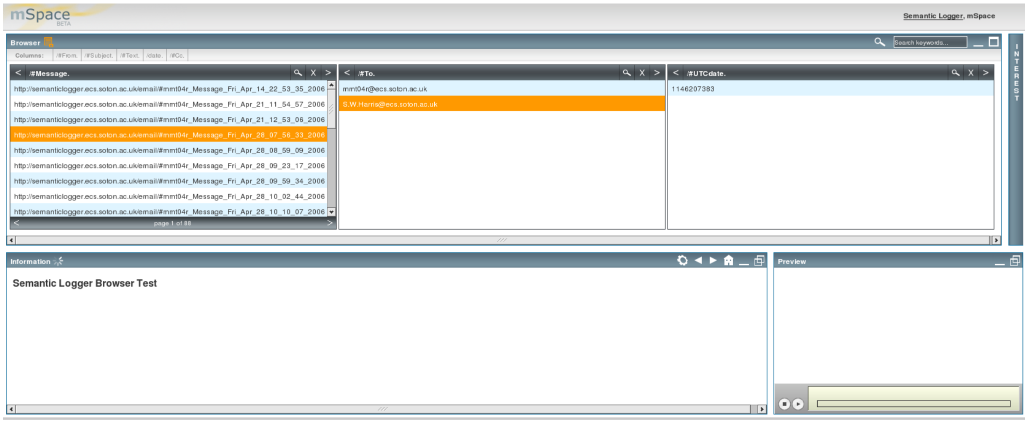

FIGURE 5.11: An example mSpace interface

implemented to automate this procedure allowing users to choose their target column, i.e. the class they wish to navigate to. Furthermore, it is important to stress that this browsing ability is greater than that achieved through representations of singular domains, since all the information logged by the system will be interconnected in automatically inferred or hard-coded ontologies. An example mSpace real estate model is shown in figure 5.11. From the Semantic Logger's UI, one can add arbitrary data to their KB, and can then subsequently choose to navigate the data, by launching an mSpace. Upon launch of mSpace from the Semantic Logger's UI, one gets presented with a list of all of the RDF classes within the KB, and is asked to select one which they desire to navigate towards from the faceted browser.

5.6 Services: Semantic Logger

The following section presents some services that integrate the information stored in the Semantic Logger. These services are proposed as a method to enrich the knowledge found in both the public knowledge base and the individual private knowledge bases.

It is assumed that the human selection process is better modeled through a dynamic function that operates on some weighted subset of an artifact's physical and contextual attributes. Defining this subset statically at the outset is expected to have a negative effect on the quality of the knowledge base. The Semantic Logger architecture employs a variety of components, each capable of performing a subroutine of the knowledge integration. These are combined at run-time to find correlations in the knowledge housed in our infrastructure.

- **Clustering algorithms**

  Clustering algorithms will be used to partition the dataset into groups of similar items and users. For users, this is achieved through exploiting subsets of the information available in their profile, while in the case of items the clustering is carried out by considering a subset of the descriptive features available for them. A wide variety of such algorithms is needed to facilitate the architecture and instances are chosen based on their past performance under similar contexts, as logged by the system. The Semantic Logger infrastructure allows for a number of novel approaches to such clustering, such named entity recognition in email correspondence, co-authorship, co-location inferred from GPS data, event attendance from calendar entries, file system similarity and so forth. This clustering algorithms will be used to highlight the salient/recurring events logged in a user's Semantic Log.

- **Ontology aligners/Aggregators**

  Heterogeneity exists between the representations of different types of resources. In order to assess similarity the system will need to acquire the relevant partial translations from those representations to a temporary shared one, which will be discarded after the process is facilitated. Since it would be unfeasible to define a representation to which any user-defined ontology can be translated to, a variety of such components will be implemented to enable different modes of generalisation or specialisation. Work has been undertaken to identify same place as, and same day, week, month as relationships in the Knowledge Bases. As a part of the OpenKnowledge Project (Siebes et al., 2007) the Semantic Logger was adapted as a technology demonstrator to show how knowledge, specifically personal metadata, could be shared between a number of trusted peers (Dupplaw et al., 2007). For a description of the work undertaken please see chapter 7.

- **Recommender Systems**

  These are the components responsible for evaluating the context of a recommendation need and for selecting the components that will be used to produce that recommendation. Recommender systems will also receive predicted ratings computed by aggregators and augment them according to the recommendation context. Different recommender systems may use other component selection, and ranking strategies to improve performance in specific contexts. The bias in choosing a particular instance is again determined by its past performance. These are elaborated upon in (Loizou, 2009). Loizou showed how Wikipedia and its link structured could be adopted as a universal vocabulary of things and how it could be used to power a multi-domain context driven recommender system.

5.7 Use Cases

It is likely that the competitive advantage of Lifelogging is that the Lifelogger can cross-reference information to provide greater power for his or her information retrieval. So for instance a digital photograph will be timestamped. The date and time of the photograph can be cross-checked against the Lifelogger's personal calendar to discover what he or she was scheduled to be doing at that particular time. If the Lifelogger has kept a GPS tracker, then the exact position of the camera can also be located. This information can be used to annotate the photograph with metadata to help in searching for and retrieving the picture (Tuffield et al., 2006a).

The following queries are presented as use cases to help motivate the capture of a Lifelog. Each use case is followed by a description of how the information needed can be extracted from a user's Lifelog.

- **What pictures have friends of mine taken from the cities which I am going to be visiting this year?**

  By combining calendar information found in one's Semantic Logger with calendar information from their friends lifelog's one can generate a list of geotagged images taken by their friends, from places that they are scheduled to go on a given year.

- **Which events did my peers attend last year which I missed?**

  This can be answered by combining information from a user's Lifelog with their friends' by comparing calendar information.

- **What are the URLs of PDFs that I have yet to visit, which friends of mine who attended one of the events I went to last year have seen?**

  In order to answer this question one would have to combine information from users' browsing histories, with information from their calendars, and information based on their file system changes. This would allow for a list of URLs of PDFs to be presented to query issuer.

- **Who sent me an email with a link to the RDF specification while I was writing my paper tuffield_www05.tex and listening to Metallica?**

  This would involve querying a user's personal Lifelog, combining information from their email communications and the contents of their emails, with information from their file system, and their music listening habits. This process would end up generating a list of email addresses that would uniquely identify a set of people.

- **What pictures did I take from Edinburgh is 2005**

  This would involve looking for photos taken by the user, combining that with information from their geolocation data, and with information from their calendar (for times when the user could not make use of their GPS device).

These are all plausible queries which have two vital things in common. First, they require integrated search of two or more data stores. And second, although they are (in the right context) serious questions, finding the answer through associative query and search demands keeping information that is of no intrinsic interest separately. If storage space was limited, this information upon which these queries rely on would probably be the first to be discarded.

Chapter 6

Photocopain: Context Driven Multimedia Asset Management

This chapter introduces the problem of personal photo annotation as a task that can be aided by the adoption of a integrated and accessible semantic lifelog – namely the Semantic Logger. Photocopain is presented as a use-case describing the application of the knowledge captured by the Semantic Logger Framework to a real world problem.

Photo annotation is a resource-intensive task, yet is increasingly essential as image archives and personal photo-collections grow in size. There is an inherent conflict in the process of describing and archiving personal experiences, because casual users are generally unwilling to spend large amounts of effort on creating the annotations which are required to organise their collections. This chapter describes the Photocopain system (see figure 6.1), a semi-automatic image annotation system which combines information about the context in which a photograph was captured, from the Semantic Logger, along with information from other readily available sources in order to generate outline annotations for that photograph that the user may further extend or amend. Given the availability of low cost hardware the task of managing personal photo collections is becoming a non trivial and time consuming process. Photocopain is as a live web-service currently running off of the Semantic Logger making use of the context based information, social metadata, and content-based image analysis. This task of annotating images by exploiting personal contextual information is motivated by the work undertaken by Naaman et al (Naaman et al., 2004a).

6.1 Photocopain: System Architecture

In Tuffield *et al* (Tuffield et al., 2006b) the Photocopain photo annotation system was presented as a stand alone system that utilises context and content based methods to

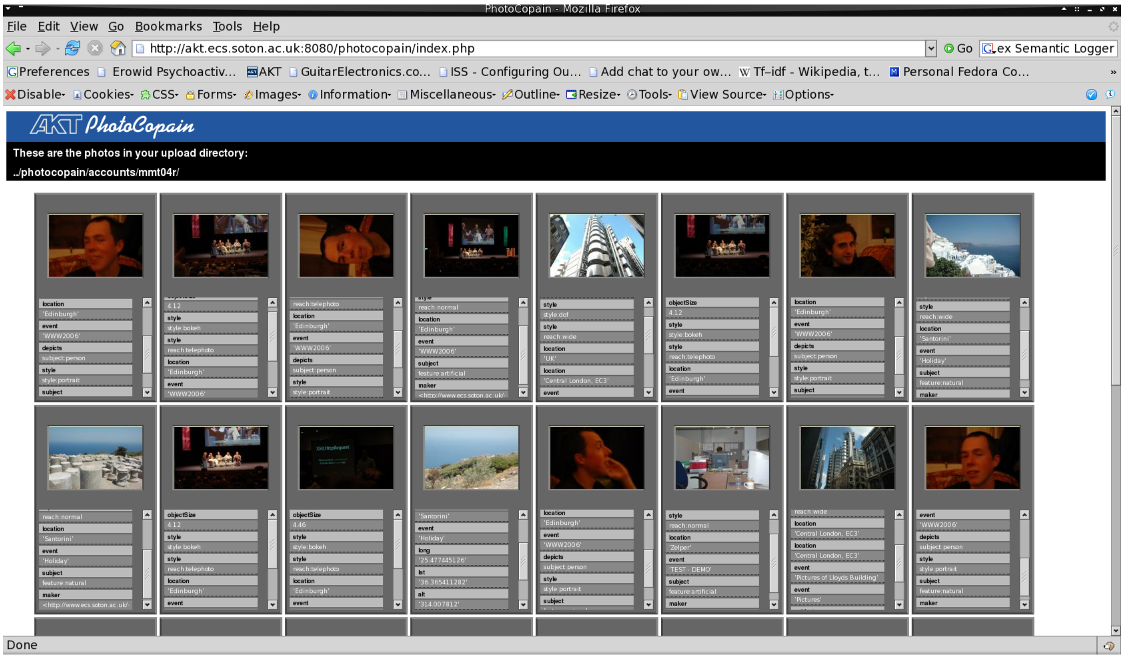

FIGURE 6.1: Screenshot of the Photocopain System

generate metadata to enrich one's personal photo-collection. The semi-automatic nature of the service is stressed, identifying the need to allow a user to author any proposed annotations, highlighting the *'Gold Standard'* of any manual annotations. The integration of a number of sources of highly heterogeneous data, along with the combination of low-level content based feature vectors, allows us to suggest annotations to the user.

The remainder of this discussion is focused around the Photocopain system, this will include a list of information sources utilised, and will be followed by an insight into the advances made to the system by incorporating it with the Semantic Logger. As with the Semantic Logger, Photocopain performs best when presented with many sources of information. The utility of Photocopain running off a user's Semantic Log is proportionate to how much knowledge is stored in the personal Knowledge Base that relies on how pro-active a user is to logging his/her digital traces.

6.2 Photocopain: Annotation Sources

Work has been undertaken to examine the potential sources of information that can be used to produce annotations, with a particular emphasis placed on ambient contextual information which can be applied to photographs with minimal effort on the part of the user. Up until this stage, the annotation task has been considered in abstract terms, aiding the making of the conscious decision to examine the inputs to the annotation task, rather than the structure of the task itself, for two main reasons:

First, with this research, the intention is to examine mechanisms and architectures for acquiring and organising the available information that are cheap, in that they do not require a great deal of human intervention. Our initial research milestone was to discover which information sources *are* cheap, as opposed to appearing cheap; easily acquired information may have hidden costs associated with its representation, or its integration with other information.

Secondly, it is clear that a user carrying out the image annotation task has a number of different priorities. He/She will be interested in different aspects of the photography depending on the specific task they are performing: annotating a photo library for the press; organising a personal collection of photographs; or selecting closed-circuit television (CCTV)[1] stills during a security operation. Rather than developing a special-purpose system that supports a specific task, a flexible system has been developed; it may be that the readily-available information that have been identified here is more useful for some tasks than others. The intention was not to introduce biases by making assumptions about the specific goals of the end user.

The following sources of information are harnessed alongside the metadata stored in the Semantic Logger:

- **Camera metadata**

  Exchangeable Image File (EXIF) (EXIF, 2002) metadata records camera parameters at the time that the photograph was taken. These parameters include: aperture setting; focal length of the lens; exposure time; time of photo; flash information; camera orientation (portrait/landscape); and focal distance. We can also derive other information from these metadata, such as the average scene brightness of an image. The EXIF is extracted from the images, presented to the Photocopain system, and then uploaded to the 3store, in a RDF representation.

- **Global positioning data**

  GPS data can be recorded in EXIF if the camera is equipped with the required hardware, or alternatively a GPS tracklog matched with a photo's timestamp can be used to determine location accurately. This is primarily of use when the camera is used outdoors. As described in *'Geo-Data'*, section 5.3 the GPS, the Network Gazetteer, the Getty Gazetteer, and the iCal information can be used to piece together a geographical log of a given user. All of aforementioned sources of information can be found in a user's Semantic Log.

  Given the rise in popularity of location aware smartphones like Apple's iPhone, Nokia N95 (see figure 6.2), more and more people are now carrying location aware devices. Based on current trends, location based digital traces will become more

[1] CCTV http://en.wikipedia.org/wiki/Closed-circuit_television

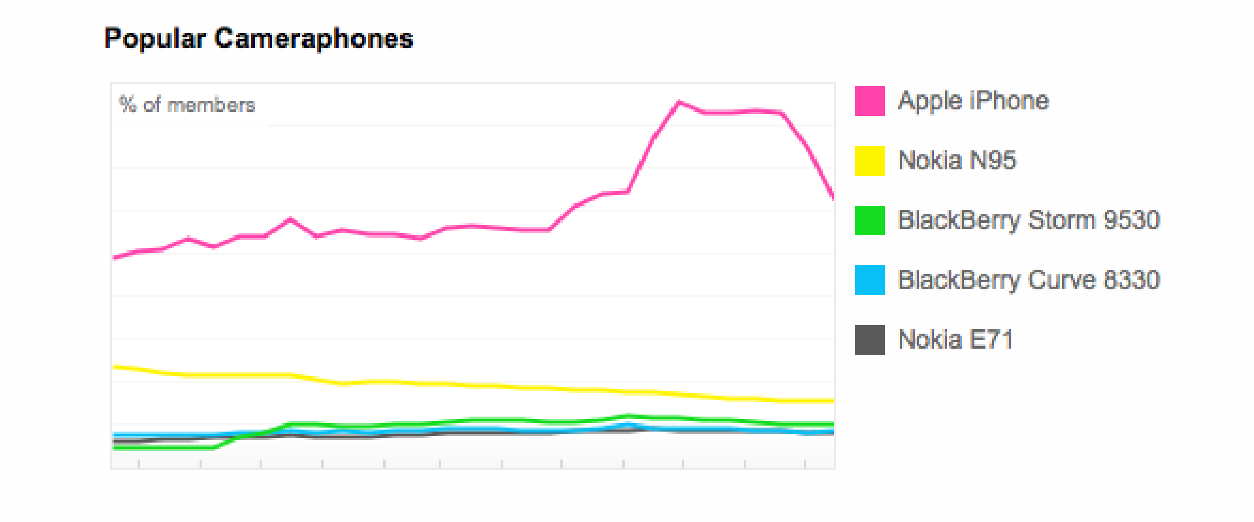

FIGURE 6.2: Most used camera phones as per Flickr 16/09/2009 http://www.flickr.com/cameras/

and more prevalent. This notion of ubiquitous geo-logging devices is the advent of Nike running shoes with an integrated GPS device[2].

The X-axis of figure 6.2 shows what percentage of Flickr users that have uploaded pictures using a given camera phone, with the scale starting at 0% at the origin through 100%. This illustrates that at around July 2009 approximately 95% of Flickr users had uploaded at least one photograph taken from the geo-location aware Apple iPhone. The Y-axis is broken up into 12 months, this show how each camera phone's usage statistics has changed over the 12 months prior the figure being generated. The graph is normalised to deal with the fact that number of Flickr users changes from day to day. The graph moving up or down indicates a change in the camera's popularity relative to all other cameras used by Flickr members. This should also be noted that the graphs are only representative of pictures taken with cameras that can be automatically detect.

- **Image analysis, Classification, and Flickr**

  A selection of image analysis techniques, such as the CIELab[3] *Hue, Intensity, Texture (HIT) Map*, and the edge direction coherence vector have been used to propose annotations for image content (see section 6.3.2 for further details). A number of classifiers have been trained using Flickr's image pool as our source of training data, these are elaborated upon in (Tuffield et al., 2006b). Flickr users may associate images with a number of free text tags (e.g. TimBL, WWW2006, Edinburgh); I use the photographs associated with certain tags as training sets for our image analysis algorithms. For example one hundred and fifty images of

[2] Nike Plus : http://nikerunning.nike.com/nikeplus/
[3] CIELab or Lab is often used as an abbreviation for the CIE 1976 (L*, a*, b*) color space colour-map

the tag ‘flower’ were taken from Flickr via its API[4], and then images that are not flowers were also downloaded from Flickr, in order to train a classifier.

This process was automated by first identifying what words have been clustered together inside Flickr (getRelated function via Flickr api), and then listing the words have been clustered in conjunction with the list of flower related tags. This list was then combined with the terms related to flower extracted from Wikipedia’s Categories[5], and was used as a filter when randomly collecting a set of one hundred and fifty images I assume to be the class of ‘not flower’. The RDF representation of Wikipedia categories can be accessed through the Semantic Logger using the username and password *wiki*.

A handful of Flickr’s most popular tags will be the initial content-based annotations (vocabulary). These include: landscape, cityscape, portrait, groupphoto, architecture, seascape, and flower. The decision to use this dataset has ensured that any proposed annotations are grounded within Flickr’s shared conceptualisation of these terms. For example, if Photocopain proposes the annotation ‘landscape’, what it actually means is ‘this image is similar to images tagged landscape by the Flickr community’ as opposed to the developers’ understanding of the word.

The web service based architecture developed for Photocopain, allows easy integration of new image analysis algorithms, and/or new clustering algorithms as needed.

- **Community Annotations from the Semantic Logger**

  Given that Photocopain has been re-purposed to work with the Semantic Logger, a piece of social software, the scope now allows for annotations to be shared within communities. Friend lists or a user’s social graph can be exploited in order for annotations to be shared within communities. Given that a friend of your’s took a picture at the same time and place will allow for annotations to be proposed by the system. This notion of using other peoples’ annotations was inspired by the undertakings of the ZoneTag project [6] at Yahoo!/Berkeley.

Photocopain shows how the information found inside the Semantic Logger can be used to enrich one’s personal media library. And as mentioned before Photocopain is presented as a *multimedia asset management system*. It uses *content*, *context*, and *community* based knowledge in order to generate as much metadata as possible. Photocopain presents the user with a number of annotations for each image submitted, while these are in turn corrected by the user, and then uploaded back to the users Semantic Log. This process of importing this photo-specific information back to the Semantic Logger, adds another dimension to be exploited by the aforementioned recommender system.

[4]The Flickr API: http://www.flickr.com/services/api/
[5]Wikipedia Categories: http://en.wikipedia.org/wiki/Wikipedia:Browse
[6]ZoneTag API: http://zonetag.research.yahoo.com/zonetag/

These information sources are quite basic, and will not allow the creation of sophisticated annotations such as 'Yet another shot of Briony falling into the pool', they can be combined to derive other relevant information, such as is was summer's day, and that the calendar suggests the family was on holiday. Work undertaken my Marc Davis at Berkeley provides insight into how context can be combined with content to aid the identification of faces inside photographs (Davis et al., 2006). For example, the focal length, focal distance and sensor size could be used to calculate the size of an object in the focal plane, which might enable us to differentiate between a leaf and a tree, even though both are in the same location and have a similar colour. The aim of our system is not to create annotations *per se*, but to generate hints which guide the user as they create the annotations. As a further example, consider how colour information can be used combined with the technical information available from the camera and other devices to help identify 'things' in the photograph. If one knows that the picture was taken with a wide angle lens, the GPS log shows that the photo was taken in a city, and the picture is largely grey, one may be able to infer that the picture is of a building or a piece of architecture. Given such an inference, and possibly some user-provided confirmation, one might then deploy an image analysis algorithm to spot sharp edges.

Understanding of photographic genres is certainly relevant here. A grasp of the standard settings employed by photographers will help with the identification of different types of pictures. For instance, portraits are typically taken at focal lengths of around 100mm, with wide aperture (f/4 or above), fill flash (reduced power), and the subject of the photo, a head/shoulders, will be of a standard size. The background at top left and top right will be out of focus (low sharpness), and the photo will have a portrait aspect ratio. These are fairly standard settings which are used by the majority of photographers; they are also the settings that an automatic camera set to 'portrait' will usually use. In the latter case, the photograph may record the automatic setting of the camera, in which case the annotation system may not need to search for the individual settings of the camera. This allows one to annotate all of their photos of type portrait, to add extra metadata for future retrieval.

6.3 Photocopain: Architecture

The Photocopain system has a service-based architecture, as shown in Figure 6.3, and has been designed so that new annotation or classification services may be added on in an ad-hoc basis. The connections between system components have been implemented as simple Web Services (making use of HTTP's PUT and POST verbs) wherever possible, and interactions with the central RDF triplestore make use of the SPARQL RDF query language (World Wide Web Consortium, 2005). Below is a description of the components of the Photocopain system, illustrating their interactions with a simple workflow.

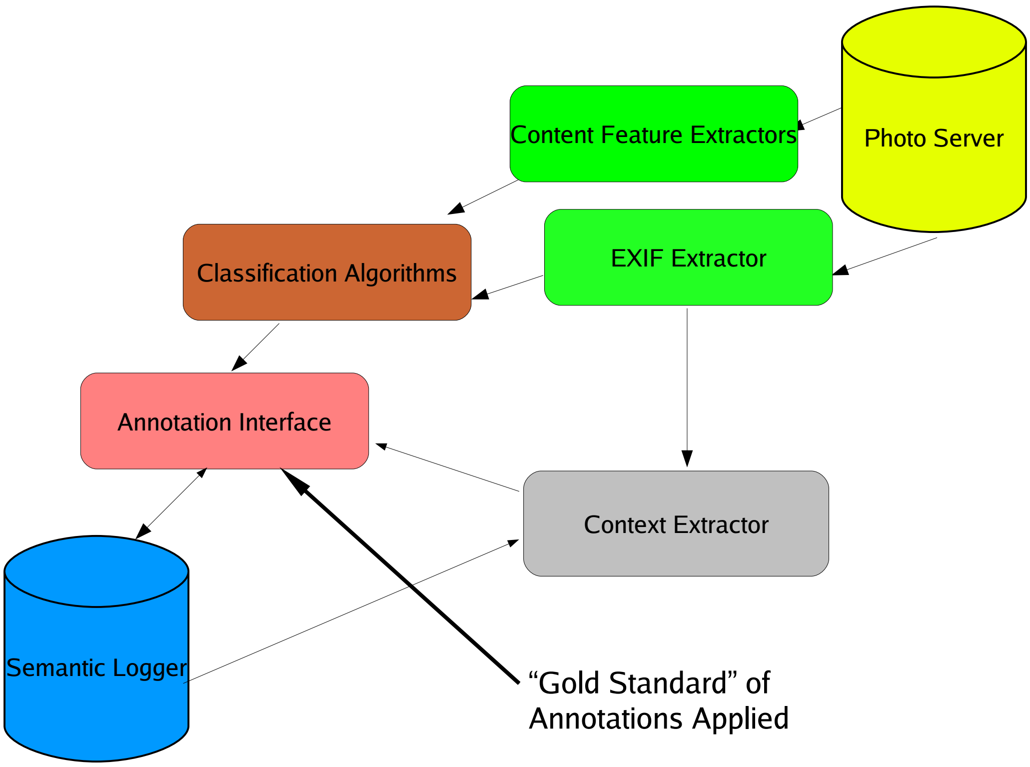

FIGURE 6.3: The Photocopain Architecture

6.3.1 Annotation Server

The central component of the system is a SPARQL-compliant RDF triplestore (Harris and Gibbins, 2003), the Semantic Logger, which is used to store the image annotations, both candidate and user-vetted. The key role of the triplestore is to act as a persistent store for the system, and to mediate the interactions between the other system components. Due to the schemaless nature of RDF, triplestores are seen as an extensible method of data storage, allowing for new sources of information to be trivially added if desired.

All of the contextual information harnessed by Photocopain is stored in the Semantic Logger, and subsequently after the annotation process, all of the extra metadata produced by Photocopain is imported back into the user's Semantic Log, enriching the user's contextual log.

6.3.2 Image Feature Extraction

Most of the current work on semantic annotation of images focuses on attempting to solve the computer vision problem. Our guiding principle with Photocopain has been simplicity; study into the use of simple, well-understood techniques with additional metadata has been core to the development. The system is extensible so that new

FIGURE 6.4: CityLandScape Detector

feature extraction algorithms can be included to extend its breadth of recognition, or its classification accuracy. For prototyping purposes four feature extraction techniques have been implemented. These were selected from the literature to see how they perform with the extra contextual information that our approach is able to provide.

A simple CIELab colour-map is used alongside other feature extraction methods for increasing the quality of a match. The colour map lays a 10x10 grid over the image and takes the average colour in the perceptually-uniform CIELab colour-space for each grid element. This spatial-colour map can be used to qualify classifications (such as providing more evidence towards a 'landscape' annotation due to the top of the image being the colour of sky, for example).

A face detection module provides an estimate of the number of faces that appear within an image (see figure 6.5). This uses the Hue, Intensity, Texture (HIT) Map (Mateos and Chicote, 2001) technique to detect face-like areas. The method performs a combination of thresholded hue, and intensity channels to obtain a 'skin' image. Clearly this is quite

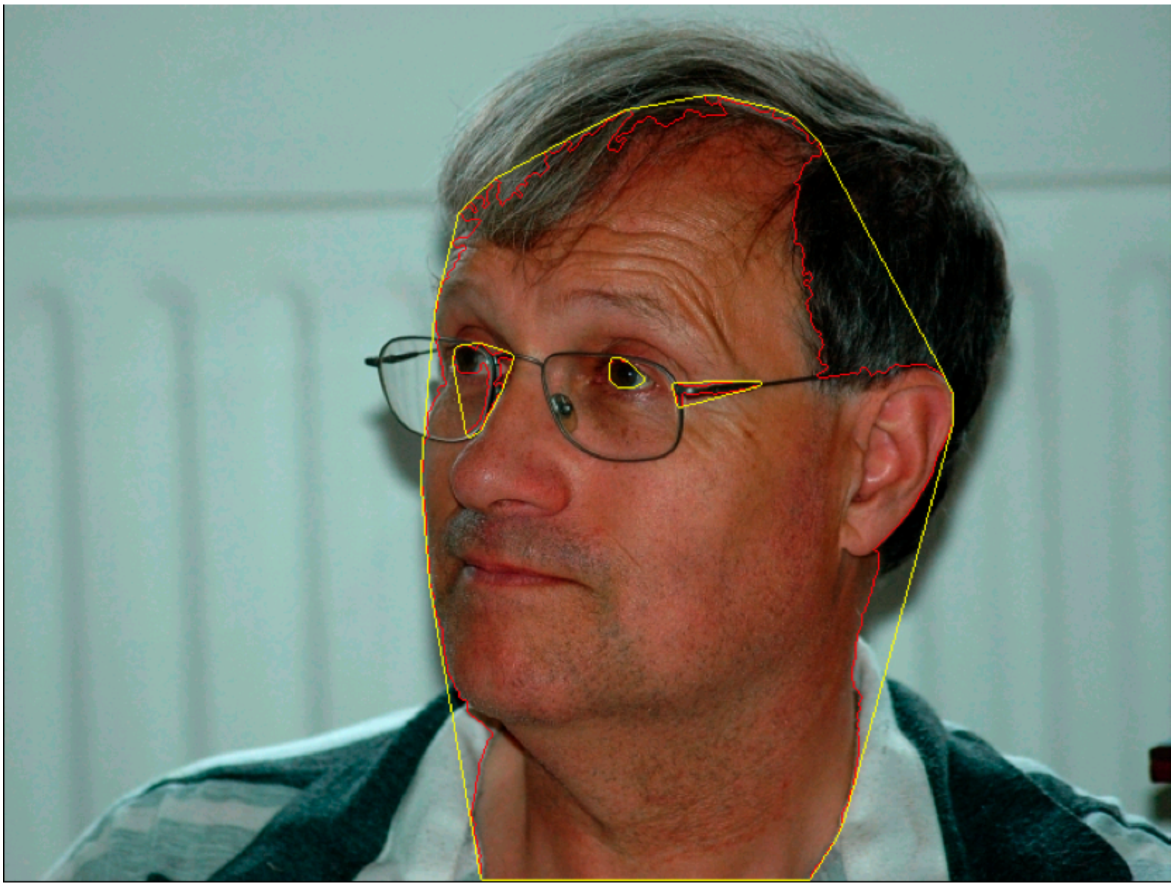

FIGURE 6.5: Number of Faces Detector

a simplification of the description of a face, but it works to identify skin regions well. After connected-component labelling and segmentation, regions are filtered based on shape and size allowing the module to return only those images which are likely to be faces. However, this process often leads to false positives, so combining this with the EXIF data from the camera lens, an approximate physical size for the possible face region can also be calculated. As faces can only ever be certain sizes, this provides more evidence for identifying whether a region within the image is a face. Combined with other EXIF-based metadata it is possible to annotate images as being portraits, or group photos.

The artificial *vs.* natural feature extractor uses classification of the edge direction coherence vector (Vailaya et al., 1998) to classify image content, based on the assumption that artificial structures tend to be created with straight-edges, whereas natural structures do not. A Canny edge-detector (Canny, 1986) is used to filter the image and short edges are binned into a histogram. An edge-tracing algorithm is used to find those edges that are coherent, and these are also binned into another histogram. The two histograms make up the edge direction coherence vector, which can be classified against, or is able to provide a measure of 'artificialness'. In our early prototype Photocopain exclusively classified an image as artificial or natural, however, it became clear that an image can

FIGURE 6.6: Focus Detector: Final Representation

be both artificial and natural simultaneously (for example a photo of a cityscape behind a park) (see figure 6.4). Creating two classifiers, rather than one, allows images to be annotated with both labels if necessary.

A focus-map performs Fourier transforms over the image in a grid. These transforms are filtered to obtain a measure of high and low frequency components within each grid element (see figures 6.6 and 6.7). The measures are then linearly combined at each location to provide an overall map for the image, of where focus lies. In images that have a shallow depth of field, and therefore contain unfocused regions, the map is able to provide both evidence of this, and also information on where the feature extractions should be targeted, if possible. Combined with other techniques, this extraction may be able to provide more evidence towards portrait annotations (photos of which tend to have shallow depth of fields at a medium focal length), or landscape annotations (photos of which tend to have a large depth of field, focused near the hyperfocal distance).

6.3.3 AKTiveMedia Annotation Tool

The annotations produced by the image classifiers and the context-based annotators as treated as candidate annotations; while human-authored annotations are expensive, they are seen as a 'gold standard', and have considerable value by virtue of their provenance. The AKTive Media Tool (see figure 6.8) is used as an interface to easy the burden of annotating the images at hand. AKTive media (Chakravarthy et al., 2006) image annotation interface was integrated in the Photocopain demo to highlight the importance of 'Gold Standard' annotations.

FIGURE 6.7: Focus Detector: Focus Map

6.3.4 Metadata Output

Once a set of annotations have been generated for an image, they may be exported as RDF, and put back into the Semantic Logger. These metadata may also be sent back to to a web gallery (Flickr) ; the RDF is used to add annotations to the image, allowing other systems to search and process the images. The mSpace interface has been chosen as our primary method of interaction with the images and their metadata.

One of our long term goals for this work, is the creation of narrative structures which present photographs in context, in order to better archive personal experiences. The annotations form a foundation for these structures, and enable chronological and narrative based queries over the data such as: "Where did we go in the last week of our holiday in the Peloponnese?" or "Who pulled me out of the crevasse I fell into in the mountains?" or even "What happened next?". This form of meta-level query is a valuable way of unlocking memories. We use the narrative structure to facilitate the most

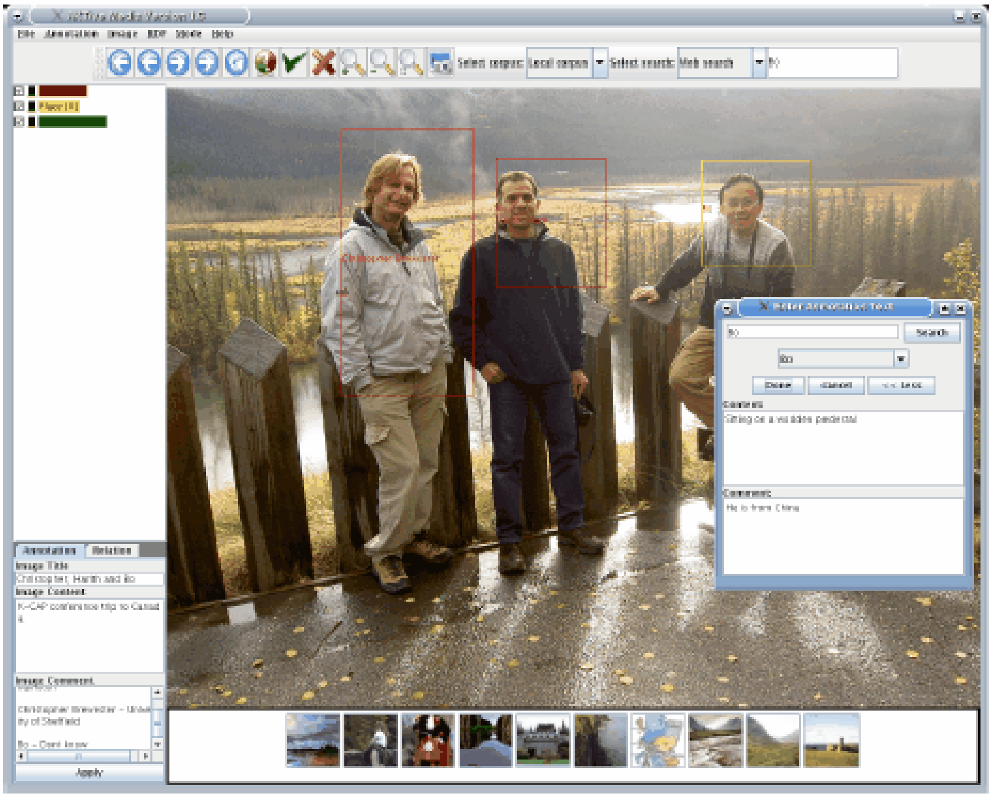

FIGURE 6.8: The AKTive Media Interface

effective presentation of the data to the user whether that be as a slide show, themed collection (such as an online album), a structured tree (Spinelli, 1999), or multimedia presentation (Alani et al., 2003b), this is presented in the future work of this thesis (see section 8.2).

6.3.5 Workflow

Broadly speaking, there are two parallel workflows shown in Photocopain: one concentrates on annotations which can be gleaned from the content of the image, while the other creates annotations based on its context, see figure 6.3.

The content-based annotation workflow, contains an EXIF extraction service which retrieves the camera metadata from the image, and a set of feature extraction services which analyse the content of the image. These metadata and features are passed to the image classifiers, which perform simple classification on the images (e.g. natural *vs.* artificial, indoors *vs.* outdoors, and so on). Finally, these classifications are passed as candidate annotations to the user annotation interface.

The context-based workflow is executed along-side the content based services. The time-stamping information found in the EXIF metadata is used to identify contemporary events in the user's Semantic Logger account. The public Semantic Logger Knowledge

Base is then queried to find any possible annotations used by other users at the same place and time. The spatial, temporal, social, and content-based metadata is then passed to the annotation interface in order to be verified by the user. Upon submission of the 'Gold Standard' annotations, the knowledge is then uploaded to the user's Semantic Logger account. For an example of the annotation generated by Photocopain please refer to appendix D.

6.4 Evaluation

The evaluation of Photocopain (see table 6.1) is based upon on the performance of the individual automatic annotation components. I have evaluated the classifiers for identifying indoor/outdoor images, for spotting artificial/natural environments, and for face detection. In all three cases, I used a training set of 150 images, with 75 instances representing each classification, and a test set of 30 images, it should be noted that the 30 test images were not used in the training set, but were then incorporated into the final classifiers used within the system. All combinations of the various metadata forms extracted by Photocopain (CIELab/HIT/Canny Edge Detector/EXIF) were tested in order to identify which combinations yielded the highest levels of precision for three desired automatic classifiers.

For the classification of images as indoor or outdoor, I found that a combination of information extracted from the EXIF data and the CIELab colour map without any form of dimension reduction performed the best, yielding only 6 errors in the 30 tests. The classification of natural and artificial environments yielded its best results with a combination of the edge direction coherence vector and the CIELab colour map clustered, using the nearest neighbour algorithm, giving 4 errors in the 30 tests. Finally, the face detection algorithm was developed by combining information generated by the Hue, Intensity, Texture (HIT) algorithm with information extracted from the picture's EXIF data. This combination of metadata was used to successfully generate a high level of precision, in terms of correctly identifying images with faces in them, by successfully classifying all but 3 of the images returned.

It should be stressed that due to the nature of these automatic annotation components and the task at hand, out of the two statistical measures provided precision is deemed the most important. Given the amount of photographs taken, and that any form of machine generated annotations must be 'semi-automatic' in nature, i.e. must be complemented by a process of human intervention to make sure that the metadata generated is of a 'Gold Standard', it is important that the tools operate with a high level of exactness, to maintain minimal user intervention. Furthermore, it is noted that in any future versions of Photocopain, the system should learn from the changes made by the end-user, i.e.

Classifier	Features Used	Precision	Recall
Cityscape/Landscape	CIELab / Exif	80	70
Natural/Artificial	Edge Direction Coherence Vector	87	77
Face Detector	HIT Map / EXIF	90	64

TABLE 6.1: Precision and Recall for Photocopain's classifiers

from the 'Gold Standard', by incorporating the classified images into the algorithm's collection of annotated images.

6.5 Conclusion

This section, presented the Photocopain system, which integrates a number of existing and well-understood technologies that address the task of photo annotation, both those which rely on the content of the images, and those which examine their context.

Organising non-textual artifacts such as photographs in such a way as to facilitate retrieval is a difficult problem; the usual retrieval strategy of exploiting keywords is of course unavailable without annotation. However, annotation requires significant effort on the part of the user, and so I have focused on ways of automating or semi-automating the image annotation task. Photocopain takes an unusual route of using information that is readily and, in many cases, cheaply available. This information is extremely heterogeneous, and much of it of low value (certainly compared to manually-generated annotations). Although the benefits may be small, the low costs of acquiring it may mean that some of these sources are still worth exploiting. It is argued that it is possible to process and integrate this sort of information, and that basic processing might yield further information. Such information, and inferences over it, could be used to help suggest annotations or annotative strategies for users, particularly in sharply circumscribed domains, and in relatively routine annotation tasks.

Photocopain has been described as a real-world application that makes use of a user's lifelog data to help automate the task of personal photo-annotation. This integration of contextual information is described as an example use-case for the capture of a semantic lifelog. Furthermore, it should be noted that one could process their photos using Photocopain to generate RDF annotations from their Semantic Log and subsequently choose to upload the annotations to a social networking site such as Flickr. RDF Pushback, as described in section 3.2.4, is a perfect technological solution for such a task, allowing for RDF to be written back into Flickr via its API. There are privacy issues surrounding such an interaction which should be noted, even though the proposed approach of first processing one's images via Photocopain and then uploading the outputted annotation to Flickr is less obtrusive than giving one's Semantic Log over to Flickr allowing for it to annotate the photos, one should always be aware of how much personal information could be harvested from their public Flickr profile page.

Chapter 7

Social Semantic Logger: Sharing Annotations with OpenKnowledge

This chapter presents a use case demonstrating how the Semantic Logger, an autobiographical metadata acquisition system (Tuffield et al., 2006b), has been taken forward by the OpenKnowledge (OK) peer-to-peer (p2p) platform. Furthermore this chapter will aim to highlight the opportunities that emerge when utilising context to support multimedia knowledge management in a social environment. This chapters aims to contextualise the advances in peer-to-peer knowledge sharing, by presenting its ability to incorporate and enhance existing context and content based techniques. The application scenario of generating and sharing narrative structures will be used in an environment where one can guarantee the interoperability of heterogeneous peers – the OpenKnowledge system.

7.1 OpenKnowledge : An Introduction

OK[1] is a p2p framework for sharing knowledge and is built around the execution of Interaction Models (IMs), expressed in the Lightweight Coordination Calculus (LCC) (Robertson, 2005). The goal is to allow knowledge to be shared freely and reliably, regardless of the source or the consumer. Reliably sharing semantic data requires either a shared conceptualisation (where there is consensus), or the mapping of semantic terms between locally defined terms. However, the process of mapping needs a context in which it can be performed, and tasks formalised by interaction models provide just that. Section 7.2 describes the basic principles of the OpenKnowledge platform, detailing its key features and capabilities.

[1]The OpenKnowledge Project http://openk.org/

Insight into how multimedia data can be used in OK, by first deploying some multimedia annotation tools onto the network and then providing interaction models in which to use them, is detailed in (Dupplaw et al., 2007).

This chapter shows how the collected knowledge can be shared to facilitate the annotation of, and semantic mapping between multimedia objects, utilising both contextual clues and content analysis. Furthermore, it is argued that novel services can be built to exploit such a knowledge corpus, and present specific examples of such services that interact within a social environment.

7.1.1 Lifelogging and Personal Metadata in a Social Context

As stated previously, the development of the contextual lifelogging framework that is the Semantic Logger has been influenced by the activities road-mapped by UK's *Memories for life* Network (O'Hara et al., 2006).

Advances in our understanding of human memory have made distinctions between the concepts of 'semantic' and 'episodic' memory (Tulving, 2002). Semantic memory refers to one's general knowledge of the world, and episodic memory is said to refer to one's ability to recollect events pertaining to their own past. Episodic memory is said to be the memory of autobiographical events including the times, places, associated emotions, and any other piece of contextual knowledge. Examples of semantic memory include facts such as *a window will break if too much pressure is exerted on it*, whereas episodic memories refer to facts such as, *I went to National Gallery 6 months ago to chat about RDF*.

The concept of keeping track of a user's activities in order to aid the contextualisation of events was thought of as far back as Vannevar Bush's concept of the *Memex* in 1945 (Bush, 1945). Recently, work has been undertaken on the *MyLifeBits* (Gemmel et al., 2002) project[2] where Gordon Bell, a researcher at Microsoft Research, has undertaken an experiment to capture as much information relating to his own life as possible. Such information includes books, memos, videotapes, phone conversations, and so on.

Two different experiments have both reported results on the utility of contextual metadata to aid the tasks of search and recall. Mor Naaman (Naaman et al., 2004a) performed a large-scale user-based evaluation on images annotated with a rich set of contextual annotations to identify which forms of metadata were the most useful for search and recall. The second study focuses on how captured contextual metadata can be used to aid memory recall for patients with varying impairments (Sellen et al., 2007; Hodges et al., 2006). Both of these sets of research have implied that information relating to answering the questions *what, where, when*, and *who* with respect to an event are key to support tasks

[2] MyLifeBits `http://research.microsoft.com/barc/MediaPresence/MyLifeBits.aspx`

of recall. This research provides insight into the workings of the human autobiographical memory system. It is clear that the remit of current life-logging technologies is to capture and utilise relationships between gathered metadata in an attempt to enhance autobiographical memory.

7.1.2 Supporting Personal Metadata via Social Context

The ability to mine social networks to aid the task of multimedia annotation has been reported in a number of fields. Metadata referring to *what, where*, and *when* can be extracted through the use of life-logging technologies, but questions pertaining to *who* else was at an event with you is somewhat harder. Reports (Davis et al., 2005; Apostoloff and Zisserman, 2007) suggest that the hard task of face-recognition can be made easier by considering the social context in which the picture was taken. These recent advances are seen as a step towards the capture of further autobiographical metadata.

7.1.3 Peer-to-Peer Knowledge Sharing

Peer-to-Peer systems appear analogous to social software insofar as a subscription to a common ontological view-point is not a requirement. The OpenKnowledge p2p system is novel in that the focus is shifted to sharing the definitions of interactions carried out between peers, rather than the actual data or services.

OK adopts a subscription-based paradigm, where peers may join interactions by subscribing to advertised roles. Only when all necessary roles within a specific interaction have been filled will the bootstrapping begin, to execute the task defined by the model.

This bootstrapping involves the negotiation between peers which is then used to infer the optimal set of peers to carry out the interaction. Peers playing roles will be called to execute specific functions that their role requires; these are encoded as constraints on message operations in the model definitions (Siebes et al., 2007).

Constraints, encountered in the interaction definition, are that data items are assumed to be annotated by the components in the network that produced them. In order for such data to be used in the context of other interactions, or by different peers, mappings must be acquired both between local representations and the context of the interaction.

In the majority of cases this can be handled by ontology mapping based on textual representations of concepts.

Providing similar functionality for multimedia data items poses extra challenges due to implicit semantics. This is the space occupied by the work presented herein to highlight the benefits of an 'open' approach to multimedia data manipulation.

$$
\begin{aligned}
Framework &:= \{Clausei, ...\} \\
Clause &:= Agent::Def \\
Agent &:= a(Type, Id) \\
Def &:= Agent \,|\, Message \,|\, Def\ then\ Def \,|\, Def\ or\ Def \,|\, Def\ par\ Def \,| \\
&\quad\ null \leftarrow C \\
Message &:= M \Rightarrow Agent \,|\, M \Rightarrow Agent \leftarrow C \,|\, M \Leftarrow Agent \,| \\
&\quad\ M \Leftarrow Agent \leftarrow C \\
C &:= Term \,|\, C \wedge C \,|\, C \vee C \\
Type &:= Term \\
M &:= Term
\end{aligned}
$$

Figure 7.1: Syntax of LCC interaction framework

7.2 OpenKnowledge

The open sharing of knowledge in a decentralised, heterogeneous environment is achieved through shifting the emphasis to sharing definitions of the interactions between peers.

Interactions are specified formally in Interaction Models expressed in the LCC language. The functionality in such models is represented by constructs called *constraints*, which are implemented locally by peers as plug-in components. The following sections provide some detail of how these elements are used in order for interactions to take place in the OK environment.

7.2.1 Lightweight Coordination Calculus (LCC)

LCC (Robertson, 2005) is a variant of π-calculus, with an asynchronous semantics and extended to express p2p interactions within multi-agent systems, thus eliminating the need for central control mechanisms. The following definition presents the abstract syntax of LCC (see figure 7.1).

There are five key syntactic categories in the definition: *Framework, Clause, Agent, Definition* (abbreviated to *Def*), and *Message*. A *Framework* consists of a set of clauses. Each *Clause* corresponds to a role in the interaction, carried out by an *Agent*. In turn, each *Agent* is assigned a unique identifier, *Id*, and a *Type* which declares the name of the role the agent will play.

The interactions that the agent must perform are given by a definition, *Def*. Definitions may be composed as sequences (*then*), choices (*or*), or parallelisation (*par*). The actual interactions between agents are carried out through message passing, while *null* is used to denoted events which do not require the exchange of messages. Messages involve the

$\mathbf{a}(\text{data\_client},C)$ ::
$query(Q) \Rightarrow \mathbf{a}(\text{data\_server},S) \leftarrow \mathbf{getQuery}(Q)$ *then*
$data(Q,D) \Leftarrow \mathbf{a}(\text{data\_server},S)$
or $unauthorised(Q) \Leftarrow \mathbf{a}(\text{data\_server},S)$

$\mathbf{visual}(getQuery(Q),input(Q))$

$\mathbf{a}(\text{data\_server},S)$ ::
$query(Q) \Leftarrow \mathbf{a}(\text{data\_client},C)$ *then*
$null \leftarrow \mathbf{resolveQuery}(Q,D)$ *then*
$data(Q,D) \Rightarrow \mathbf{a}(\text{data\_client},C) \leftarrow \mathbf{authorised}(C,D)$
or $unauthorised(Q) \Rightarrow \mathbf{a}(\text{data\_client},C)$

FIGURE 7.2: Example Interaction Model in LCC

sending ($\Rightarrow$) or receiving ($\Leftarrow$) of structured terms M from another agent, and whether these exchanges are carried out may be predicated by the truth value of constraint C. $Term$ refers to a structured term in the Prolog[3] syntax, while Id is either a variable or a unique agent identifier. Finally there is the concept of

7.2.2 Interaction Models (IMs)

This section puts forward the notion of Interaction Models within the OK framework. Given the proposed task of peers requesting annotations from other peers, the below description focuses around the notion of a client asking a server for data. In our proposed lifelogging use case, one could imagine a peer (client) attempting to annotate a photo asking all of the peers available (servers) for metadata pertaining to an event at a given time and place.

Since OK uses interaction models as the network's main currency every message propagated through the network always has some context (an IM) within which the messages are constrained (Siebes et al., 2007). As it stands, Interaction Models (IMs) are discovered using simple keyword search.

LCC ensures the coherence of interactions between agents by imposing constraints relating to the messages they send and receive in their chosen roles. The clauses are arranged so that although the constraints on each role are independent of others, the ensemble of clauses operates to give the desired overall behaviour.

The following Interaction Model (see figure 7.2) defines the interaction carried out for querying a database server.

[3]Prolog programming language syntax and semantics http://en.wikipedia.org/wiki/Prolog_syntax_and_semantics

First, a *query* message is sent to the seller with the variable Q representing the query to be resolved. The message is only sent if the client agent, C, is able to compile a query. This is carried out via the *getQuery* constraint; the model hints that this should be a visual clause that involves the user. Such generic clauses are used to suggest default visualisations for arbitrary constraints, and in this case, *input* proposes that the user is asked to provide input, by showing a text box for example. The query is then stored in the variable Q of our example that was passed to the constraint in the IM. Once the *query* message is received by the server, it is resolved to data that will be stored in variable D, by solving the *resolveQuery* constraint. Subsequently, the server will assess whether client C is authorised to view data D by evaluating the *authorised* constraint. If so, the data is sent in a *data* message. Otherwise the client receives an *unauthorised* message.

A constraint in LCC can thus be viewed as a predicate, whose arguments can be either instantiated or not. The task is to instantiate all the arguments and return a binary value indicating whether the constraint has been successfully satisfied.

7.2.3 Open Knowledge Components (OKCs)

As seen in the above section, constraints in an LCC protocol are pre– or post– conditions on message passing. As such, the way in which they are solved is not specified by the IM but through plug-in components called Open Knowledge Components (OKCs). OKCs are local to the peer and expose methods that can be mapped to constraints in the interactions the peer participates in. The architecture (Siebes et al., 2007) allows peers to participate in shared interactions using local components to carry out the constraint solving. As such, the functionality described in an IM can be achieved among heterogeneous peers, independently of operating system or other requirements. To better illustrate this point, consider the example of finding images relevant to a query: one peer could require a human to browse through a collection of images and identify the relevant ones – a costly yet accurate process. Another could make use of an automated algorithm to yield faster, but less accurate results.

7.3 Personal Multimedia Knowledge Management

The Semantic Logger Framework has been extended to use online data to generate further contextual information, as shown in figure 7.3. The Semantic Logger now makes use of the ever increasing amount of HTTP[4] resolvable RDF data.

[4]HTTP - Hypertext Transfer Protocol definition from the W3C http://www.w3.org/Protocols/

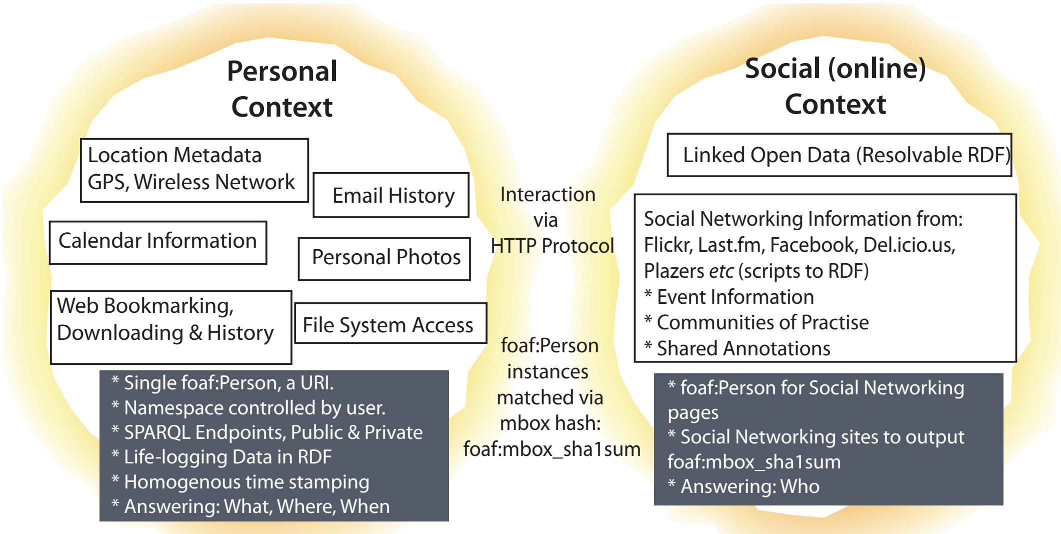

FIGURE 7.3: Overview of the Social Semantic Logger Architecture

A distinction is made between *personal metadata* – information captured and stored directly by the user, and information harvested from the Internet, referred to as *social metadata*.

When setting up a Semantic Log, a user is asked to identity a Uniform Resource Identifier (URI) or a WebID[5], which will be associated to all the personal information logged. The Friend-Of-A-Friend ontology (FOAF) has been developed to model users, their friends, and their online accounts. This is adopted by the Semantic Logger system by importing a FOAF document into the personal knowledge-base (KB) of each user. An important side note is that the primary subject of the FOAF document must be the user's URI[6] in order to avoid ambiguities.

The FOAF document is used to describe and point to other available RDF which may hold information about the user, or the user's friends. This is achieved through the `foaf`[7]`:knows`, and `rdfs`[8]`:seeAlso` relationships respectively. These are subsequently used to link to other social networking information available on the web. The document is either uploaded from a web-accessible Uniform Resource Location (URL). A discussion[9] on the differences between URIs and URLs is beyond the scope of this thesis.

The adoption of a URI to refer to the user, or a WebID, is key to the approach. This is presented as a method of empowering the user with respect to their *digital identity*. The amalgamation of data regarding the user from disparate online resources into a single KB provides the user with a global view of the amount of personal information published on the World Wide Web. While information posted to the Web is in the

[5]For example: http://id.ecs.soton.ac.uk/person/6914 or http://mmt.me.uk/foaf.rdf#mischa
[6]The author's WebID at time of writing http://users.ecs.soton.ac.uk/mmt04r/foaf.rdf#me
[7]FOAF namespace: http://xmlns.com/foaf/spec/
[8]RDFS namespace http://www.w3.org/TR/rdf-schema/
[9]Naming and Addressing: URIs, URLs, ... from the W3C http://www.w3.org/Addressing/

public domain, the Semantic Logger system enables the users to hold all of their personal information in a single place, so that they can then make informed decisions on whether to withdraw or amend such information. Each of the online accounts held by the user are in turn described as a `foaf:Person`, that is an instance of the type Person from the FOAF ontology. These have their own URI's, such as http://www.last.fm/user/MischaTuffield or http://www.flickr.com/people/MischaTuffield, each of them can be seen as a different persona for a given user. In the FOAF document the user's main URI is also cast to the type `foaf:Person`, and links to a user's other online accounts must be made explicit. A site that attempts to bundle up and expose, in RDF, a list of people's social networking data is QDOS[10]. The site presents the user with a URI (e.g. http://qdos.com/user/5acc361496df109a7c2967760d5d9792), which itself is an instance of `foaf:Person`, that has a list of online accounts. Section 7.3.2 details the manner in which the data captured from social-networking sites is related to the user's URI.

7.3.1 The Semantic Logger - Personal Metadata

In figure 7.3, the items grouped under the label, *Personal Context*, are sources of contextual information logged directly by the user. The choices are informed by the work presented in Section 7.1.1. The use of a common time format is essential to generating an autobiographical contextual log. The iCal[11] ontology has been employed to represent the occurrence of events in our KB.

The following sources of personal information sources have been parsed into an RDF format, associated to a user's URI, and stored in a personal KB.

- **Web Browsing habits** This information includes web-browsing, bookmarking, and downloading. These are parsed into a local namespace with concepts described in the Netscape ontology[12].
- **Location Information** GPS tracklogs are converted into geographic points, expressed in the *Basic Geo (WGS84 lat/long) Vocabulary*[13]. Location information is also tracked by logging wireless access points a user's laptop connects to. This is done using the online service Plazes, which is then converted into a local RDF namespace.
- **Email History** The user's MBOX is parsed into a local namespace, detailing e-mail activity. These can in turn be matched up to public `foaf:mbox_sha1sums` to indicate relationships between people.

[10] QDOS is a social networking site, that outputs RDF describing people by listing their various social networking profiles http://qdos.com/
[11] The W3C's iCal namespace: http://www.w3.org/2002/12/cal/
[12] Mozilla's RDF schema definition: http://www.mozilla.org/rdf/doc/
[13] The W3C's Geo-location namespace: http://www.w3.org/2003/01/geo/

- **Music Listening Habits** The Audioscrobbler ontology[14] and the Musicbrainz URIs[15] are used to detail music listened by the user.

- **Personal Photos** Taken from a user's Photocopain account are used to ground events inside the contextual log. The act of taking a photo is seen to imply the existence of an event. Please refer to chapter 6 for further information regarding Photocopain.

7.3.2 The Semantic Logger - Social Metadata

There are a number of efforts to integrate information describing people and their social relationships, exposed by online networking sites. One of the major efforts in this space is Google's OpenSocial API[16], which attempts to unify existing social networking sites by detailing a common set of Representational State Transfer interfaces based around exchanging data in the Extensible Markup Language[17] (XML) and the JavaScript Object Notation [18] (JSON).

Another approach is that taken by the Open, each of them can be seen as a different persona for a given user. For example, one's work persona may be hosted on the social networking site LinkedIn[19], and their personal life may be hosted on Facebook. Linked Data Community, to employ existing ontologies, namely FOAF, SIOC[20], and SKOS[21], to generate linked HTTP-resolvable RDF. These are just two initiatives working to unite disparate information online and the Semantic Logger adopts the latter.

The act of publishing factual information to the Web is seen as indicative of the occurrence of an event, which can then be annotated using the assembled context.

Information harvested from the Internet, and uploaded to the Semantic Logger, is what is referred to as *social metadata*. The sources of information selected to be captured from the Web have been chosen to complement the sources of *personal metadata* described earlier by utilising the social aspect of the sites. The information gathered by the system includes, but is not limited to:

- **Web Browsing Information** The site delicious.com, provides information regarding Web bookmarking. This information is parsed into a local timestamped namespace. Information regarding a user's friends bookmarking activity is also stored in the Semantic Logger.

[14]DBTune's Last-fm RDF data http://dbtune.org/last-fm/
[15]Musicbrainz, URIs for music http://musicbrainz.org/
[16]OpenSocial Project: http://code.google.com/apis/opensocial/
[17]XML specification from the W3C http://www.w3.org/XML/
[18]The JavaScript Object Notation: http:/www.json.org/
[19]LinkedIn the professional social networking site http:///www.linkedin.com/
[20]Semantically-Interlinked Online Communities http://sioc-project.org/
[21]Simple Knowledge Organisation System http://www.w3.org/2004/02/skos/

- **Location Information** Location information is taken from the photo-sharing site Flickr, in the form of geo-tagged images, and from the social networking site Facebook. Through the use of linking events, based on iCal and Facebook events, one is able suggest that two Facebook friends may have been at the same location. This piece of social information would allow for images taken at the same time to share tags, and hence increase the available contextual information surrounding that piece of multimedia.

  Flickr exposes user-based information in HTTP-resolvable RDF, using the SIOC ontology[22]. These fragments contain friend information, as well as information regarding user groups, and are imported to the Semantic Logger. It should be noted that this page does not expose an `foaf:mbox_sha1sum` as this information is not publicly available. This means that people surfing a public Flickr page (or a publicly available RDF representation of a Flickr page) will not be able to connect the information with the owner's public FOAF URI, unless they already know it. RDF versions of Flickr users can be found on http://foaf.qdos.com/, for example http://foaf.qdos.com/flickr/people/MischaTuffield is the RDF version of the author's Flickr account.

- **Music Playlist Information** Last.fm is used to capture information regarding both the user's and their friends' music listening preferences. This information is currently published by the site dbtune.org, which creates an RDF representation of a user's Last.fm page. The following URL is an RDF version of the author's Last.fm account http://dbtune.org/last-fm/MischaTuffield.rdf.

As stated earlier, a user is set to adopt a URI, which is used to dereference any information stored in the Semantic Log. Each online account held by the user has its own URI, and recent work(Passant, 2008) suggests the use of the `owl`[23]`:sameAs` property to unify a user's various FOAF instances. This approach has not been taken in the Semantic Log since the transitive nature of the `owl:sameAs` relationship means that any contextual information will be lost. Thus one would not be able to issue a query to discriminate between one's Flickr friends, and the friends related to them through their primary FOAF URI. That said, at times one may want a list of URIs for all their online friends, regardless of the domain in which they are based. At this point, the work suggests the use of SPARQL and the `foaf:mbox_sha1sum` property to achieve this. The `foaf:mbox_sha1sum` is a hash[24] of the user's email address (e.g. *mailto:mmt04r@ecs.soton.ac.uk*) that can only be decrypted if the email address is already known. It should be noted that the FOAF ontology caters for users to have as many email address as they wish, allowing for different identifiers to be used across different social networking sites. A user could choose to use a different email address for each

[22]SIOC namespace http://rdfs.org/sioc/spec
[23]http://www.w3.org/2004/OWL/
[24]W3C Document describe SHA1http://www.w3.org/TR/1998/REC-DSig-label/SHA1-1_0

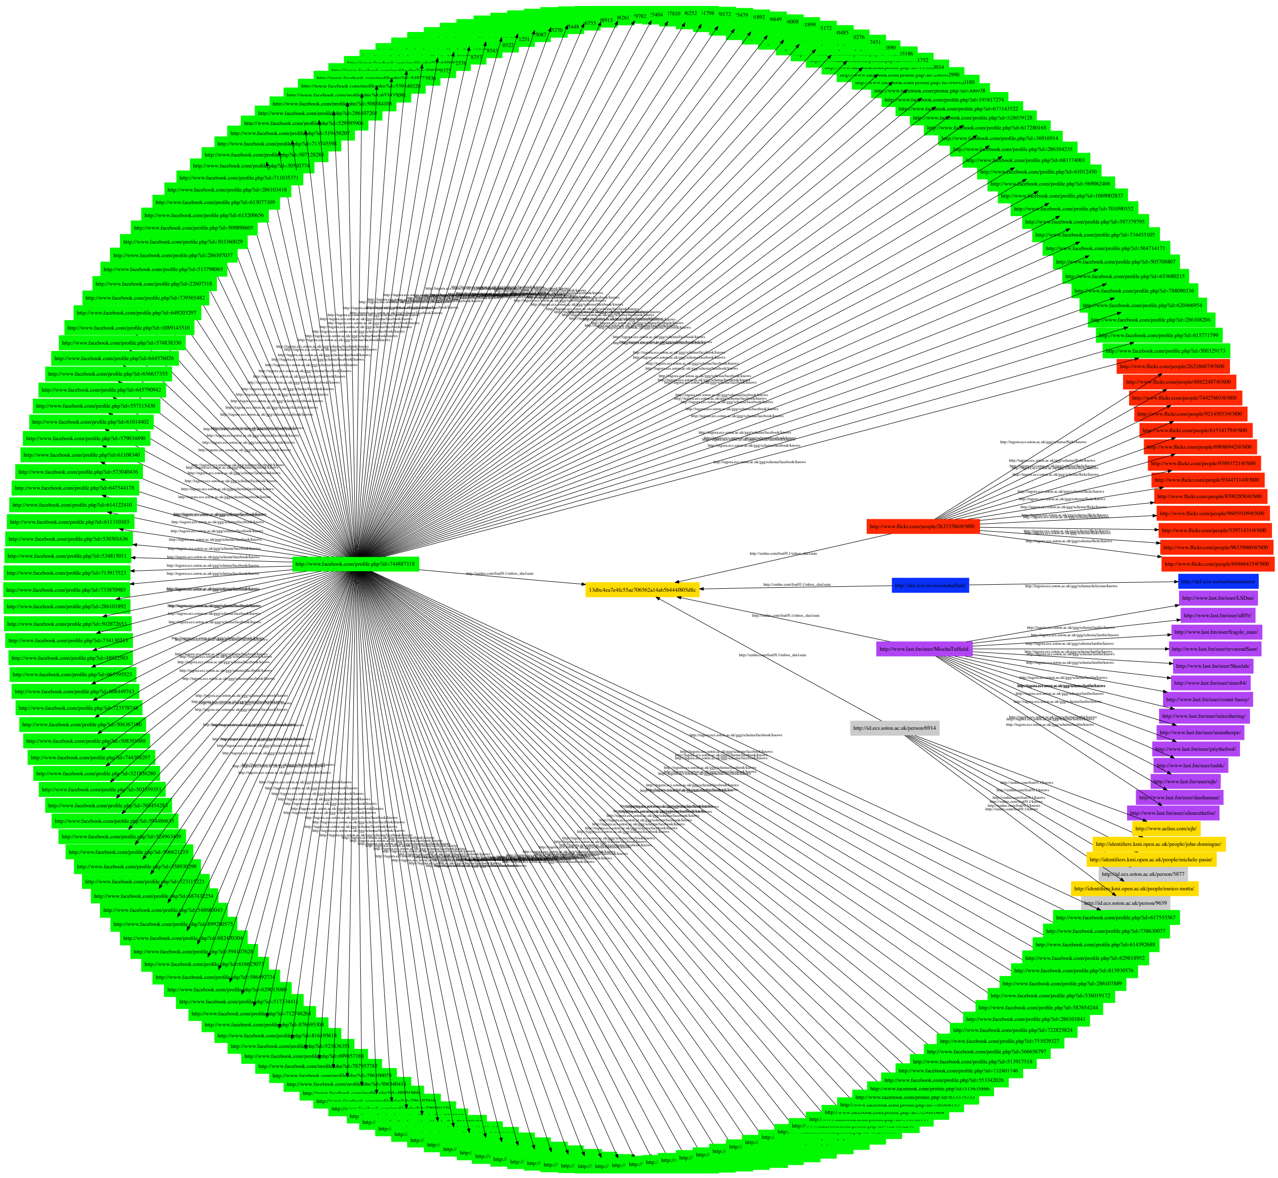

FIGURE 7.4: Graph of the author's friends, when there is an owl:sameAs relationship

social networking site, and in turn only they themselves would know that the various online profiles are different personae for the same person. The FOAF ontology makes use of this relationship, as do many social networking sites including Last.fm, Flickr, and QDOS, thus allowing for multiple URIs to be queried as one without committing to the powerful semantics of the `owl:sameAs` relationship. Figures 7.4 and 7.5 illustrate how one's various online accounts can be kept distinct through the use of shared identifiers (7.5), or how they can be merged into one overriding online persona (7.4).

At this point in time is it worth presenting the reader with a note of caution. Researchers from China have found (Wang et al., 2005b,a) they can find collisions in the sha1sum hashing algorithm with a complexity of less than 2^{69} hash operations, as apposed the theoretical bound of 2^{80}. This in turns changes the game with respect to how secure

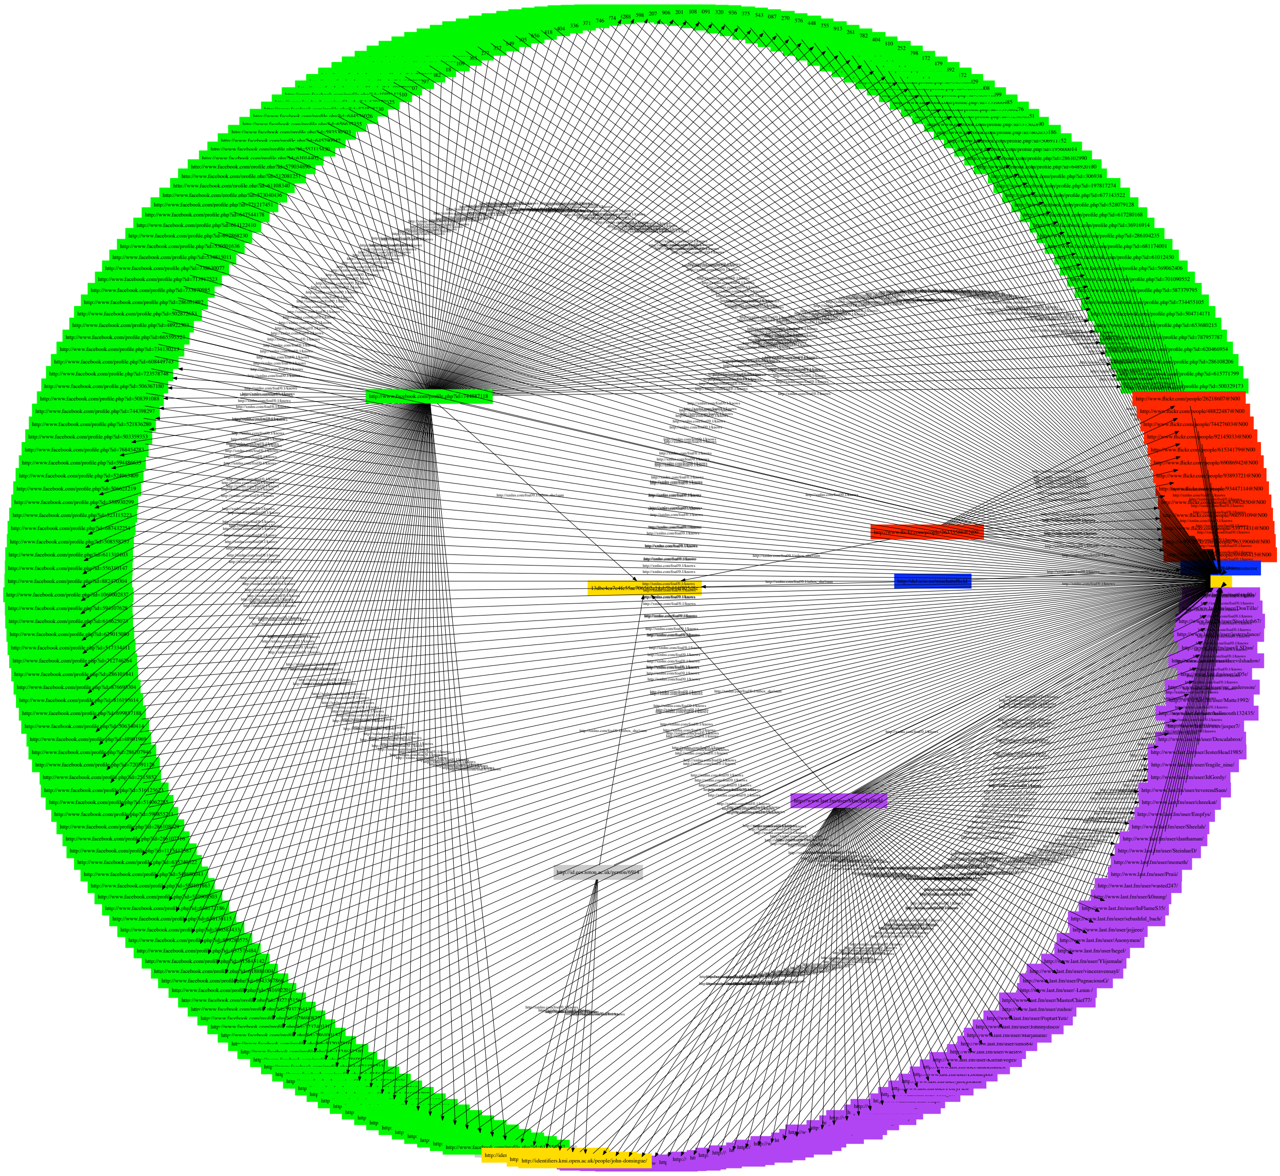

FIGURE 7.5: Graph of the author's friends, with distinct online personae

the hashing algorithm actually is. This fact is put forward as a caution, and if someone wanted to keep their email address completely private, they should be wary of using it's sha1sum as a unique identifier.

As stated before, the Semantic Logger consists of a central public KB, where all users publish their public data, along with a KB for each user which is password protected. This is presented as a simplistic implementation of a FOAF model of trust, and should be extended upon in any future work.

7.3.3 Image Annotation

Images are related to the user using the FOAF ontology (instances of `foaf:Image`). Each of the user's personal photographs are analysed by a number of content-based techniques in order to output high-level annotations. The details of the content-based annotation generated by our system, and how these are combined with Exchangeable image file format (EXIF) information can be found in (Tuffield et al., 2006a). The approach describes simple content-based techniques to acquire high-level annotations such as indoor/outdoor, landscape, portrait photo, group-photo, night/day, thus adding further contextual information.

7.3.4 Utilising OpenKnowledge to Generate and Share Narrative Structures

This application scenario aims to describe how both multimedia and knowledge can be shared in OpenKnowledge. Below presents a system where users interact with their Semantic Log to generate a narrative of their past experiences in the form of a multimedia presentation. A narrative is presented as a collection of related facts represented as events telling a story about a person's life. Practically this autobiographical narrative is presented as an ordered collection of the multimedia objects that appear in a user's lifelog.

The context assembled in the log, together with the rich semantic relationships between events are exploited to provide this ordering. The 'openness' of the OK platform provides the ideal environment for deploying a plethora of different narrative generation techniques to cater for different needs, thus providing a personalised service.

Figure 7.6 provides an example of an annotated image. A collection of these images taken in one day are automatically annotated as above, using the context assembled and content-based analysis. The representation provides a powerful distinction between the *instances* and their *semantics*. The distinction allows the extraction of an abstract model for the generated narrative. This is the SPARQL query issued to the KB to return a list of annotated images as above. The abstract models facilitates the sharing of *narrative structures* over the p2p network, which can then be re-instantiated by heterogeneous peers with other multimedia objects.

The annotated image is related to an `ical:Vevent` in the user's calendar, and has been associated to a geographic location, both by virtue of time. Due to analysis of the EXIF information, the system can tell that the object depicted in the image is not that far away from the camera, and as a result the image has been associated to the geo-location. If the image was of an item far away from the camera, the image's topic would not have

```xml
<?xml version="1.0" encoding="utf-8"?>
<!DOCTYPE rdf:RDF [
<!ENTITY xsd "http://www.w3.org/2001/XMLSchema#">
]>
<rdf:RDF
 xmlns:dc="http://purl.org/dc/elements/1.1/"
 xmlns:exif="http://www.kanzaki.com/ns/exif#"
 xmlns:foaf="http://xmlns.com/foaf/0.1/"
 xmlns:ical="http://www.w3.org/2002/12/cal/ical#"
 xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
 xmlns:geo="http://www.w3.org/2003/01/geo/wgs84_pos#"
 xmlns:image="http://jibbering.com/vocabs/image/#"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema#">
<rdf:Description rdf:about="http://akt.ecs.soton.ac.uk:8080/photocopain
/accounts/mmt04r/DSC_2063.JPG">
  <dc:title>Lloyds Building</dc:title>
  <dc:description>Picture taking trip central London</dc:description>
  <dc:date>2006-10-08T14:52:02</dc:date>
  <exif:location>Central London, EC3</exif:location>
  <exif:make>NIKON CORPORATION</exif:make>
  <exif:model>NIKON D70s</exif:model>
  <exif:fNumber rdf:datatype="&xsd;decimal">6.3</exif:fNumber>
  .
  .
  .
  <foaf:maker rdf:resource="http://id.ecs.soton.ac.uk/person/6914"/>
  <rdf:type rdf:resource="http://xmlns.com/foaf/0.1/Image"/>
  <foaf:depicts>
   <ical:Vevent>
    <ical:uid>419A1550-AF49-436C-BDBF-BD6381DA1A97</ical:uid>
   </ical:Vevent>
  </foaf:depicts>
  <foaf:topic rdf:parseType="Resource">
   <geo:lat>35.647488888889</geo:lat>
   <geo:long>139.73964166667</geo:long>
   <geo:alt>54.451782</geo:alt>
  </foaf:topic>
</rdf:Description>
</rdf:RDF>
```

Figure 7.6: Example Annotated Image

been the coordinates, but the fact that the image was taken from that location would be encoded instead.

By sharing this through the p2p network, other peers who log the same activities can then re-purpose the narrative generator to their personal multimedia collection and KB. This is an interactive and recursive process, allowing for further fields of metadata to be utilised to narrow the set of images returned. The task is essentially utilising SPARQL to partition the user's RDF graph and generate personalised photo-albums. The number of annotations inferred is proportional to the amount of information exposed by the user along with the amount of HTTP-resolvable RDF about them.

The following Interaction Model (figure 7.7) is used to create the narrative.

Provided that the *narrative_requestor* peer has a Semantic Log available, it can send a *narr_struct_req* message to a peer in the *narrative_generator* role. Upon receipt of the message, the *narrative_generator* uses a private method to create a narrative structure, S, which is returned to the *narrative_requestor* in a *narr_structure* message. Note that the generation process is encapsulated in a single constraint, *genStructure* and it is left up to individual peers to decide how to satisfy it. In the simplest of cases this can be a static SPARQL query to obtain information about pre-specified events, while more complex scenarios could include the automatic assessment of the importance of arbitrary events to provide more appropriate narrative structures. Once the *narrative_requestor* peer receives the structure, provided they can instantiate the narrative from their Semantic Log, L, the narrative can be presented using a local visualisation method. This allows for the deployment of simple, generic visualisations, alongside elaborate task-specific ones. A transformation of the RDF representation

a(narrative_requestor,R) ::
  narr_struct_req(S) ⇒ **a**(narrative_generator,G)←**getSemanticLog**(L) *then*
  narr_structure(S) ⇐ **a**(narrative_generator,G) *then*
  null←**instantiate_narrative**(S,L,N) *then*
  null←**show**(N)
 or **a**(narrative_requestor,R)

visual(*show*(N),$NarrativeVisualiser(N)$)

a(narrative_generator,G) ::
  narr_struct_req(S) ⇐ **a**(narrative_requestor,R) *then*
  narr_structure(S) ⇒ **a**(narrative_requestor,R)←**genStructure**(R,S) *then*
  a(narrative_generator,G)

FIGURE 7.7: Interaction Model for generating narratives in the OpenKnowledge framework

of the narrative to hyper-text using XLS Transformations (XSLT)[25] technology, with links representing semantic relationships is presented as a simple means to visualise a narrative.

By generating more elaborate narrative structures, one can then dynamically deploy systems tailored to support specific memory processes. A narrative summarising a person's daily activities can support better time management and aid in the scheduling of future events. This can be carried out relatively easily without actually allowing any third parties to access the information. The narratives can thus be presented as a generic method of organising and sharing knowledge that can then be exploited for various purposes. One could even speculate that the framework could be used to support a scenario where patients with memory impairments are monitored through the system by health professionals. Narrative structures can then be hand tailored to cater to each patient's individual needs, while using the same data collection system.

7.4 Conclusions & Future Work

Firstly, OpenKnowledge's Semantic Logger Demonstrator can be downloaded from the project demo page http://www.cisa.informatics.ed.ac.uk/OK/drupal/demos. The work presented in this chapter identifies specific techniques to obtain context-based annotations for multimedia items to complement existing content-based techniques. However, by populating the p2p network with a corpus of annotated multimedia, opportunities emerge for sharing such annotations between peers. For example, if a peer seeks to annotate a specific image, but lacks the required contextual information to do so, the image could be compared to annotated ones using content-based feature extraction and

[25] XSLT Definition from the W3C: http://www.w3.org/TR/xslt

the annotations from similar images reused. This notion of one peer being able to ask others whether or not they have information that can be used to annotated one of their multimedia items is presented as work towards a future whereby user's can store information in their own personal knowledge base whilst still being able to benefit from a social environment.

The distributed nature of the OpenKnowledge platform allows for users to have their own personal knowledge base, e.g. in the form of a Semantic Log, which stores all of their personal contextual information, to subsequently allow their friends access to their personal store, so that annotations can be shared within a social context – without the need to surrender all of their personal information to a third party such as Facebook or Flickr. All of the current, popular, social networking services require a given user to upload their personal information to a central website in order to facilitate social annotations, whereas OpenKnowledge allows data to be shared between peers in a distributed environment, hence empowering users by allowing them to own their personal information and to store it on a server which they own. The distributed, p2p nature of the OpenKnowledge platform along with a Semantic Log, and a user's FOAF network allows for sharing of data between friends. This sharing will allow for annotations to be propagated, for example, if two friends are known to be at the same event, via an iCal entry, and one adds further annotation such as 'WWW08' to their images or one has geolocation information provided by their camera, the second user can have the same annotations suggested for the images they took during the event's time frame. This is presented as the key contribution of the work described in this use case. This notion of sharing annotations between trusted peers is said to be the foundation of an ongoing work to build a decentralised social network that is described in the future work section of this thesis (see section 8.2.6).

Chapter 8

Conclusions and Future Work

8.1 Summary

Given that more and more of our lives are moving into the digital realm, from the notion of maintaining relationships on social networking sites or keeping online personal diaries in the form of blogs, to managing professional collaborative efforts via online calendaring systems and online documents editing suites (such as Google Calendar[1] and Docs[2]). This thesis presents a framework for capturing one's *digital traces*, by introducing the Semantic Logger, which is followed by use cases which build upon the principles enumerated during the course of the thesis.

The thesis starts off by presenting the concept of *Lifelogging*, which is seen as the proactive capture and archival of digital personal information, and then identifies some key motivations to why one may wish adopt it. In an age when more and more of our public persona can be accessed on the Web, lifelogging techniques are presented as a method of empowering users with respect to what information about them could be captured if someone was attempting to do so. Chapter 3 defines the notion of lifelogging, digital personae, digital trails, and outlines how one's online identity is an asset that is becoming more and more valuable in our online world. Lifelogging technologies are presented (section 3.2.1) along with insight into how the culture is being driven by advances in both sensory and storage hardware. Lifelogging is contextualised with respect to its potential future role in supporting human memory (section 3.2.2), which is followed by insight into the amount of personal information that is being posted to the Web. The motivations then present a number of examples whereby personal information has been procured from the Web and has been used to tell misleading stories about people. Lifelogging and the capture of personal information from the source is presented as a manner of empowering people with detailed knowledge of their perceived actions. This chapter

[1]Google Calendar http://www.google.com/calendar
[2]Google Documents http://docs.google.com/

further lays forward issues pertaining to privacy (section 3.1.3), and highlights the fact that no one should be forced to keep a lifelog and if they decide to the lifelogger must be in a position to determine any access control privileges themself. Finally, in section 3.3 the concept of the *infosmog* is elaborated within the context of the personal information on the Web, whereby the Semantic Web vision is introduced as a major factor in the way future generation will manage their data in ever more networked world (section 3.3.3).

In chapter 4 the notion of capturing metadata in the form of stand alone events that can be serialised into a narrative structure is presented as a desirable means of representing knowledge. The chapter starts with insight into the field of narrative (section 4.1) and its relationship to the field of computer science (section 4.2). Narrative structures are presented as one of the more natural methods of conveying information between humans, and is subsequently presented as a way of organising machine readable information (section 4.2.2). OntoMedia 4.3, an ontology for the annotation of heterogeneous multimedia is then introduced, described, and presented along with a number of application which have made use of the ontology (section 4.3.1). OntoMedia's flexibility and ability to represent stories is then discussed, and it complemented with an example (section 4.3.4). Given that lifelogging technologies aim to capture and archive autobiographical information, the knowledge engineering task of generating an ontology for the description of stories educated future design decisions made when developing the Semantic Logger Framework. The core principles set out by the OntoMedia ontology, specially the knowledge encapsulated in OntoMedia Core, guided the decision to adopt an event based representation when capturing information for one's lifelog. This along with the notion of capturing all of the data in the lifelog as individual self contained events is presented as the main contribution of the OntoMedia knowledge engineering exercise.

The Semantic Logger framework is introduced in chapter 5. The Semantic Logger framework includes the enumeration of a number of Semantic Web Technologies (section 5.2) that cater for portable and extensible data storage. The knowledge acquisition section of the Semantic Logger chapter (section 5.3) present a number of source of personal metadata which are presented as cheap to capture. Bases on the use of RDF as the Semantic Logger's knowledge representation language the fact that the knowledge base is schemaless, the framework presented can be easily modified and extended in the future. The chapter also described the implemented Semantic Logger, presenting screen-shots of the work completed where applicable. It is also important to stress that all of the technologies adopted by the Semantic Logger are open, and the data format employed is not proprietary. The W3C employs a strict patent-free policy when it goes about defining standards, and this is presented as key to the longevity of any lifelog. This openness means that access to one's personal data stored in the Semantic Logger will never be subject to copyright or patent related issues in future years.

Finally, chapters 6 and 7 present two applications, in the form of technology demonstrators which build upon the Semantic Logger. Chapter 6 presents the Photocopain system, which was presented as one of the key technology demonstrations of the AKT project. Photocopain shows how personal information, in the form of an autobiographical contextual log, can be employed to help annotate images. In a time where digital cameras with high levels of storage capacity are common place, Photocopain presents a method of exploiting one's lifelog in combination with content-based imaging techniques to generate automatic annotations. Photocopain combines off-the-shelf content-based image annotation techniques with easy to capture contextual information to enrich one's personal multimedia collection. Given the ever increasing ubiquity of the Web, Photocoapin shows how a Lifelog can help one organise and in turn search through one's personal multimedia collection based on simple annotations. In chapter 7, a demonstrator for the OpenKnowledge project is presented as whereby the Semantic Logger was ported to a peer-2-peer knowledge sharing environment to illustrate how trusted peers could share metadata between themselves. In the OpenKnowledge demonstrator, each user is allocated their own Semantic Logger peer, which acts as an agent for the given user. Based on a user's social graph, as dictated by the FOAF profile, peers may grant access to other peers to allows for contextual information to be shared. In this context a user's social graph is used to determine which peers, if any, are to be considered trusted. This second demonstrator illustrates how services which are enriched through social interactions can operate in a decentralised manner. This illustration of separate knowledge bases interacting with each other and sharing information, given the social graph permits so, is said to be a clear indication of the future of social networking, whereby user's can store all of their own information in their own knowledge base, whilst still being able to reap the benefits of acting in a social environment.

8.2 Future Work

This future work section is broken down into subsections relating to various pieces of work undertaken in this thesis and how the work is being pursued by either the author himself or by the research community at large. Before continuing with the future work breakdown, it is important to stress that the author aims to finish packaging up the work into an open-source repository so that others can download and install their own Semantic Logger if they wish to start lifelogging their own personal information, a screenshot of this standard alone system can be seen in figure 8.1.

8.2.1 Privacy on the Social Web

Following on from the concerns and issues presented in chapter 3, specifically around the notions of data-ownership, data-portability, and privacy on the Web, the author is

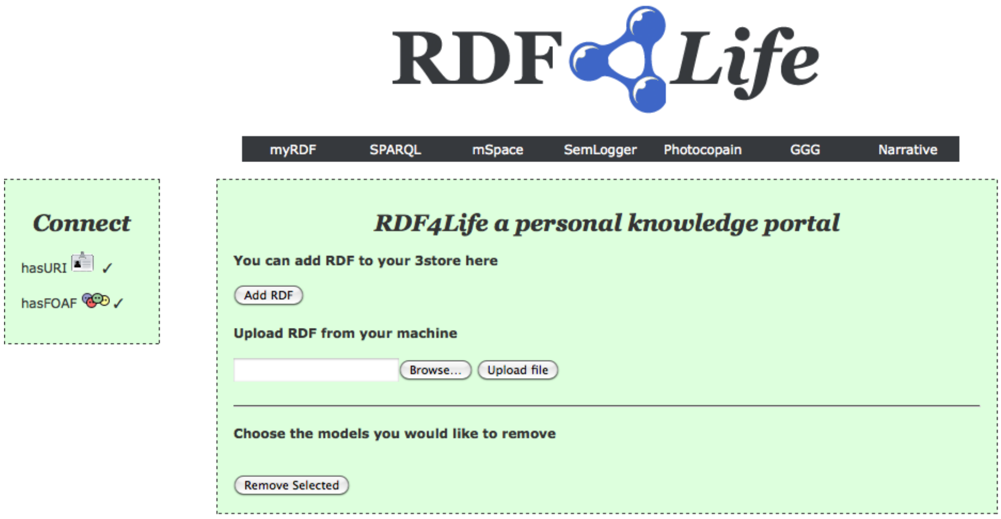

FIGURE 8.1: RDF4LIFE, a stand alone single user Semantic Logger, for open-source release, currently running on http://arcadia.ecs.soton.ac.uk/

now involved in looking at how the W3C can generate best-practices when it comes to privacy on the Social Web. The author is currently working on the W3C's Social Web Incubator Group (SWXG)[3] focusing his attention on the Privacy Task Force, and the Contextual Information Task Force. As per SWXG's website, the incubator group's mission "*is to understand the systems and technologies that permit the description and identification of people, groups, organizations, and user-generated content in extensible and privacy-respecting ways.*".

8.2.2 OntoMedia

The various pieces of work developed using the OntoMedia ontology (section 4.3.1) have inspired a number of people within the BBC to start viewing their data as events that take place with a story arc. As a result the ontology is being adopted in an attempt to experiment with representations for both factual (currently, not in the public domain) and a fictional work created by the BBC. Work is underway to annotate BBC's collection of Dr Who episodes, and OntoMedia has been selected as the ontology of choice. The following two articles highlight the interest and applicability of OntoMedia form representation narratives – http://www.r4isstatic.com/?p=54 and http://www.r4isstatic.com/?tag=semantic-weblinked-data-roots. These pages link to graphical representations to how OntoMedia can happily represent stories with obscure timelines[4].

[3]SWXG Incubator Group at the W3C: http://www.w3.org/2005/Incubator/socialweb/

[4]An example of multiple timelines in Dr. Who http://www.r4isstatic.com/linkeddata/dw/resource/blink_2.jpg

Following on from the community's interest in OntoMedia, work is being undertaken to build a Linked Data version of the ontology, as it is currently only available via svn[5], or from within the University of Southampton's network. Furthermore, a digital narrative forum (Contextus)[6] has been setup by the creators of OntoMedia, which seems to have caught the interest of the community at large. Plans are afoot to create a web resource to stimulate interest in the notion of defining, using, and sharing narrative structures in a machine-readable manner on the Semantic Web.

8.2.3 The Semantic Logger: Future sources of contextual information

Given that the Semantic Logger framework is extensible in nature the following extensions are presented as future work.

- **Weather Information**

  Weather information is proposed as a means of putting an event into context. The weather service that captures weather information given a time and location is still under development. Work is underway to harvest data from `http://www.weather.com/`.

- **News Headlines**

  The capturing of News Headlines is also presented as another method of enriching personal context. Work is underway to harvest New Headlines from the BBC website[7]. This is presented as another means of placing events into context.

- **Online Activity Streams**

  The capture of a user's online activity stream (the term used by the NoTube project to refer to a user's digital traces) is the subject of the on-going EU project "NoTube"[8]. NoTube has a number of academic institutions partnered up with the BBC in an attempt to semantically capture and annotate one's TV watching habits. TV watching habits are also presented as another source of contextual information which could be consumed by one's lifelog. This intention is illustrated by the BBC's recent initiative whereby they expose broadcast information regarding their programmes in a machine-readable format. An example of such machine-readable information is the programme information for the "The Daily Politics Show" which can be captured as RDF from `http://www.bbc.co.uk/programmes/b006mjxb.rdf`. One could imagine that anyone who wished to pursue lifelogging, and who listened to the radio would find this source of context very useful.

[5] OntoMedia ontology via Google code `http://contextus.googlecode.com/svn/trunk/ontomedia`
[6] Contextus : Digital Narrative forum `http://contextus.net/`
[7] BBC News website: `http://www.bbc.co.uk/news/`
[8] NoTube EU project :`http://www.notube.tv/`

- **Named Entity Recognition**

  In order to aid the discovery of ontological matches in the sources of information identified in section 5.3, a *named-entity recognition* approached needs to be implemented. This will allow for correlations to be found from within the knowledge bases, for example it would be nice to be able to recognise the fact that when my calendar file states that I have a meeting with Nigel Shadbolt, this is referring to the same Nigel Shadbolt found in my email correspondence, hence making an explicit connection in the KB.

8.2.4 Photocopain

Future work to enhance and the content based imaging techniques adopted by Photocopain are presented as future work. Since the development of the Photocopain service tools for the automatic identification of faces has matured tremendously. Google's photo-sharing network Picasa[9] and Apple's iPhoto[10] software now both cater for the tagging of people faces. Both systems can be presented with a handful of pictures of a given face, and then subsequently learn to auto-tag new pictures of the same face with the learnt annotation. They are both said to do this using the Eigenface feature extraction methodology[11]. This is presented as a major breakthrough in content-based image analysis, and is presented as a key method of facilitating the sharing of annotation within social graphs. For example, if Alice's Photocopain was able to identify Bob, who also had a Semantic Logger account, in a photo-graph Alice's Photocopain service could ask the Bob's Semantic Logger for any annotation available to add extra metadata to Alice's multimedia collection. Currently, Photocopain implements a face-detection algorithm which goes as far as identifying a human-face, but it does not attempt to identify the individual depicted.

8.2.5 Narrative Generation in the OpenKnowledge environment

The concept of sharing narrative structures between peers in the OpenKnowledge environment is under investigation. This attempts to share information in the same way as presented in chapter 7. As mentioned before, the semantic gap between content-based image features and concept-based image descriptions is a major barrier to the understanding and processing of multimedia data. Below is a real life scenario to illustrate the automatic construction of semantic graphs by using the Semantic Logger framework, which can be taken a step further to generate narratives.

[9] Google's Picasa Photo-Sharing http://picasa.google.com/
[10] Apple iPhoto software http://www.apple.com/ilife/iphoto/
[11] Eigenface Algorithm description on Wikipedia http://en.wikipedia.org/wiki/Eigenface

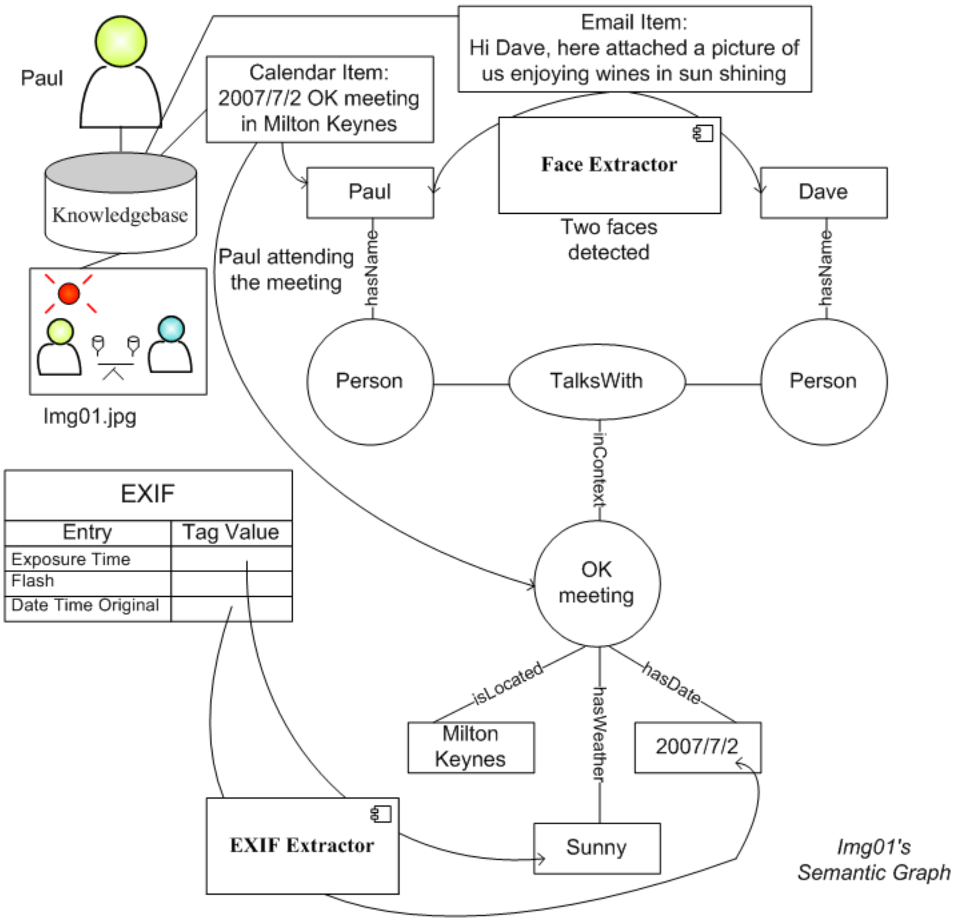

FIGURE 8.2: The scenario using multimedia annotations for Semantic Graphs

Suppose a picture will be taken in which Paul talks with Dave over the OpenKnowledge meeting and the picture is stored in Paul's knowledge base. The picture's EXIF information reveals that the picture is taken on 2007/7/2, a sunny day (low level image features). Then the calendar agent looks for the diary items put for that day by Paul and it is found that an OK meeting in Milton Keynes has been arranged on that day which Paul plans to attend (high level semantic descriptions). Having the knowledge so far it is believed that this picture is about Paul attending a meeting. Furthermore, the Email agent detects an email Paul sends to Dave with the picture attached (high level semantic descriptions) and the Face Detector finds two human faces in the picture (low level image features). As a result, a fact is established that the picture has both Paul and Dave on it, further supported by the semantic descriptions given in the email. Eventually, a semantic graph is incrementally built-up with knowledge from low level image features and high level semantic descriptions used in combination, shown in Figure 8.2. In conclusion, the complementary natures of low level image features and high level semantic descriptions (collectable by various propagators in the Semantic Logger framework) are exploited in the semantic graph and it is may be demonstrated that, if they are combined in use for multimedia annotations they provide better expressiveness.

8.2.6 Making use of Public and Private Personal RDF data

Finally, given that the second use case of sharing knowledge within the OpenKnowledge platform, the notion of catering for public and private FOAF data is presented as a future challenge, and the first step in enabling what could become a decentralised social network. This future work involves porting the behaviour illustrated in the OpenKnowledge demonstrator (see chapter 7) from a bespoke peer-2-peer environment to a standard Web based one. The work undertaken by Garlik Ltd on foaf.qdos.com, is presented as a first step in towarding a decentralised social network. foaf.qdos.com offers users with the ability to store both personal public and private data in a SPARQL endpoint. foaf.qdos.com take the approach that public information is exposed as linked data, and not a SPARQL endpoint, so that anyone can grab the public FOAF data and do with it as they wish. foaf.qdos.com presents the users with the ability to store private data behind an OAuth[12] endpoint, and subsequently allow access to the OAuth endpoint to their trusted peers, which in foaf.qdos.com's terms is the list of the user's friends taken from their public FOAF data. Currently, the foaf.qdos.com system requires that the user users the private OAuth endpoint hosted at the following URL `http://private.qdos.com/` but plans are afoot to release the code so that people can elect to host their own OAuth server on their machine if they are that way inclined. Further work is required before the foaf.qdos.com system can perform the message passing illustrated in the chapter 7 of this thesis required for distinct personal knowledge bases to share contextual information. Current investigations are underway to look at the applicability of technologies such as Extensible Messaging and Presence Protocol (XMPP) [13], to create a means for trusted peers to exchange private metadata. An example of such an approach is the 'Jabber chat query services' (JQBus)[14] as presented by Dan Brickley, one of the original creators of the FOAF ontology. For insight into the services available on the foaf.qdos.com site, please refer to the following slides `http://foaf.qdos.com/slides/london090909/index.html`.

8.3 Conclusions

As the internet is becoming more and more intertwined with our daily lives, and the price of digital sensory hardware falls, people are leaving behind more and more digital traces. The thesis puts forward a framework for the capture and archival of this incidental personal information in the form of a Semantic Lifelog. The framework is then applied to two real-world use-cases illustrating how information gathered can be exploited to support the labour intensive task of personal multimedia annotation. The work is rounded-off with an example use-case alluding to how information can be shared

[12] OAuth Specification `http://oauth.net/`
[13] XMPP Specification `http://xmpp.org/`
[14] JQBus `http://svn.foaf-project.org/foaftown/jqbus/intro.html`

between distributed, distinct lifelogs allowing for users to take advantage of their social network without being tied into a third-party service.

Appendix A

The Linked Data Clouds

This appendix illustrates a pictorial progression of how the Linked Data community has since its conception in late 2007. These pictures have been taken from the Richard Cyganiak's website[1] and are available under the Creative Commons License Attribute / Share Alike 3.0[2].

[1]Richard Cyganiak's homepage: http://richard.cyganiak.de/

[2]Creative Commons license http://creativecommons.org/licenses/by-sa/3.0/

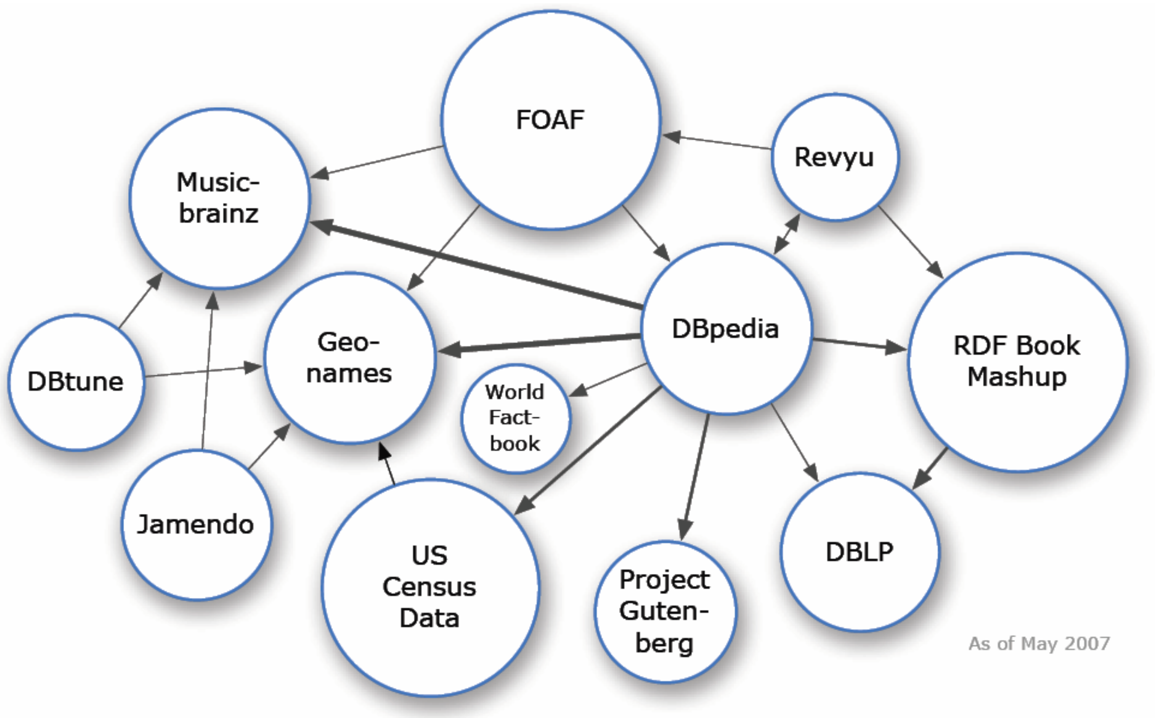

FIGURE A.1: LOD Dataset 2007-05-01 http://homepages.cwi.nl/~troncy/Talks/2009-03-06-mozcamp/lod-datasets_2007-05.png

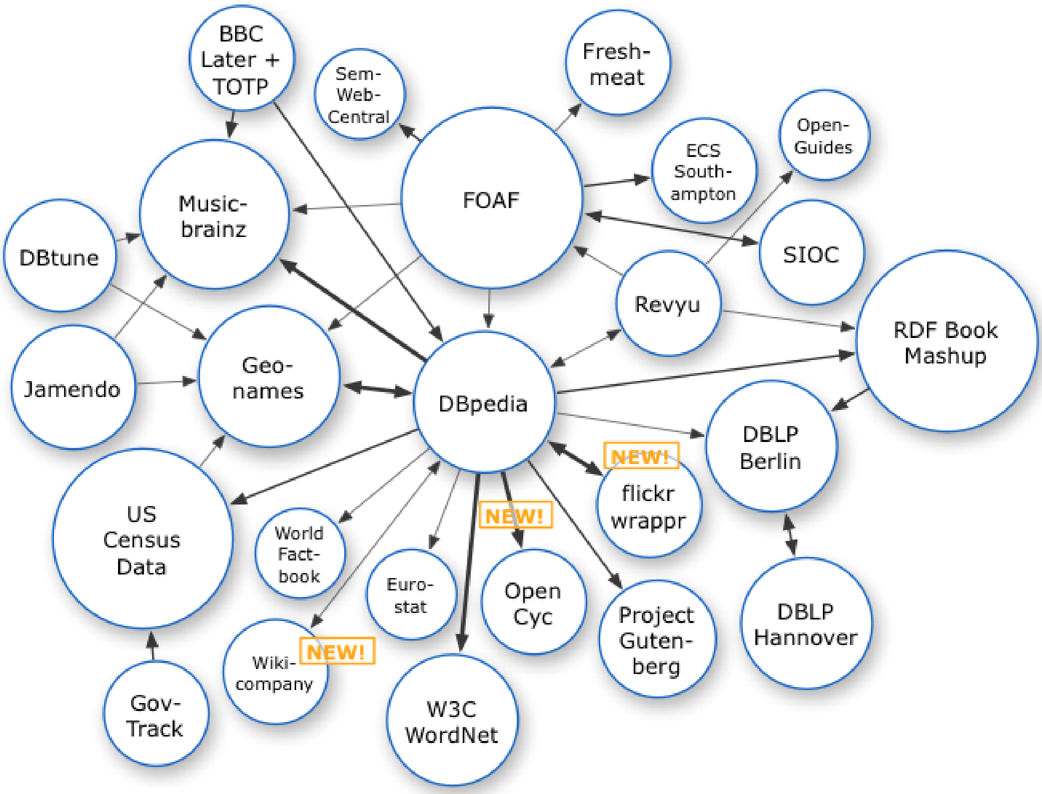

Figure A.2: LOD Dataset 2007-10-08 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2007-10-08.png

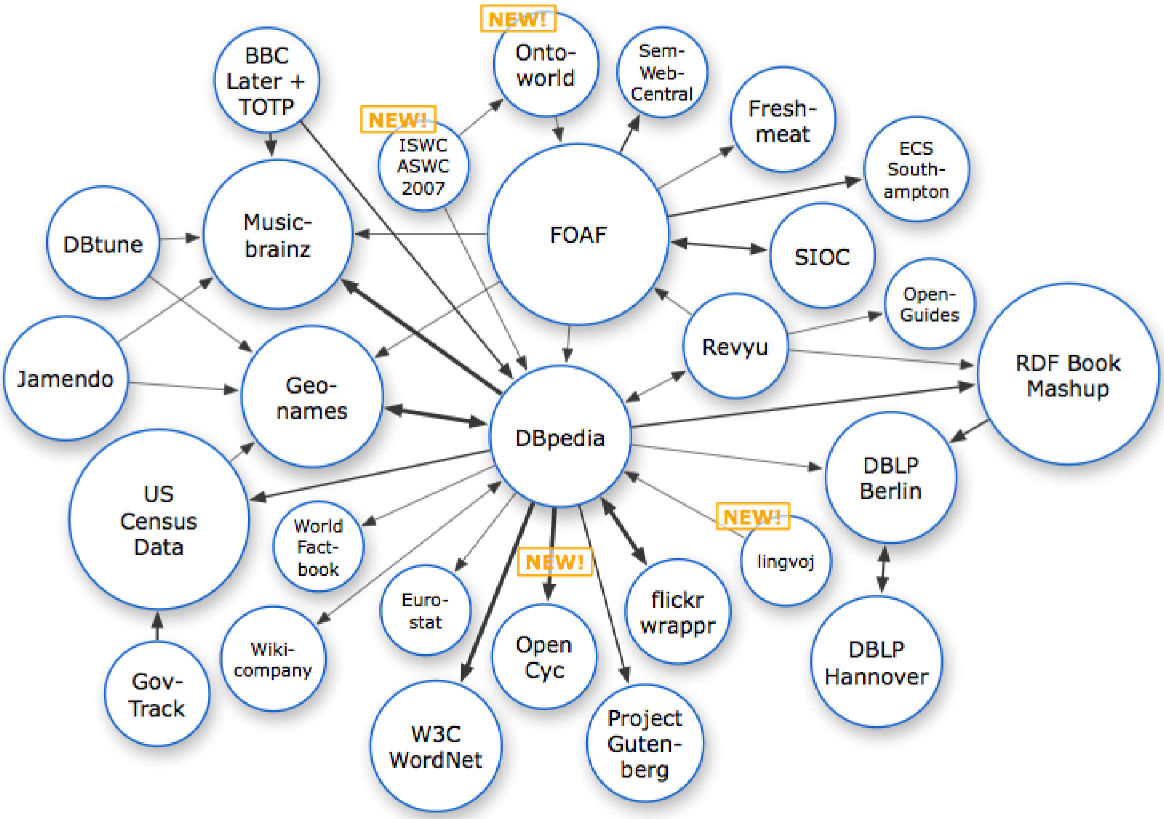

Figure A.3: LOD Dataset 2007-11-07 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2007-11-07.png

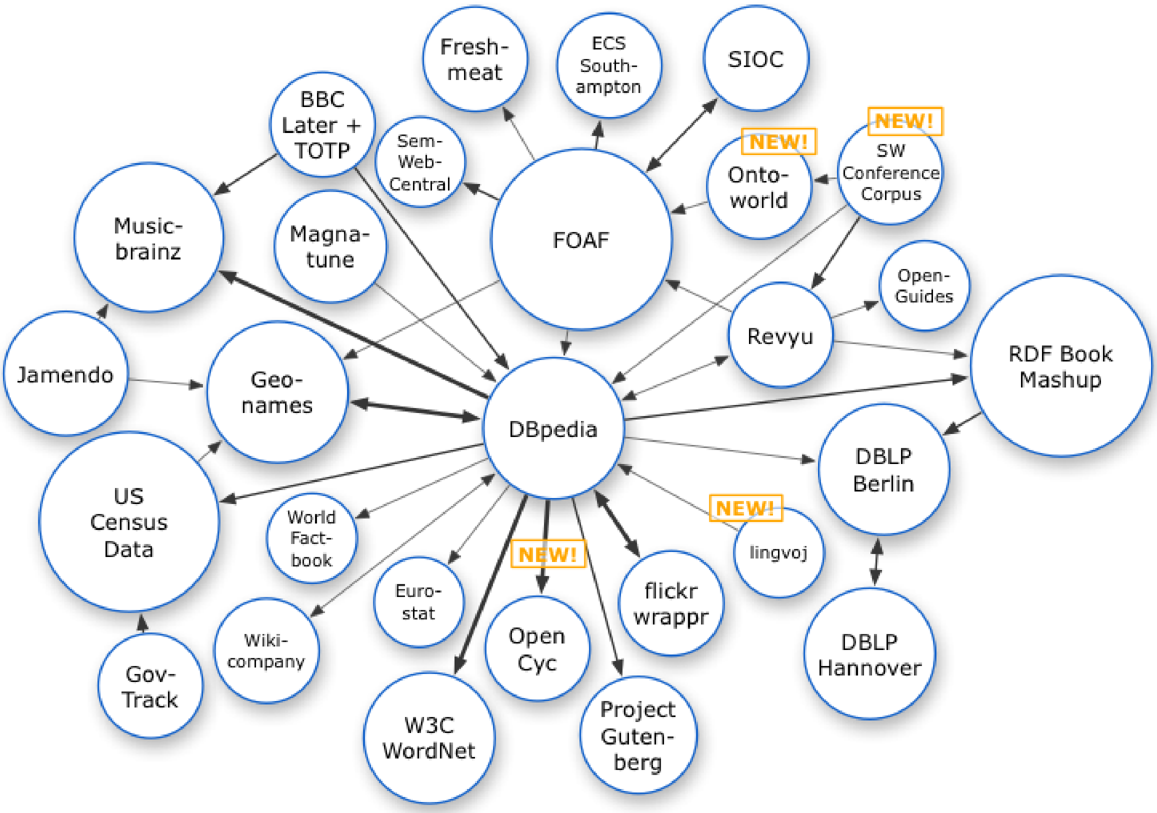

FIGURE A.4: LOD Dataset 2007-11-10 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2007-11-10.png

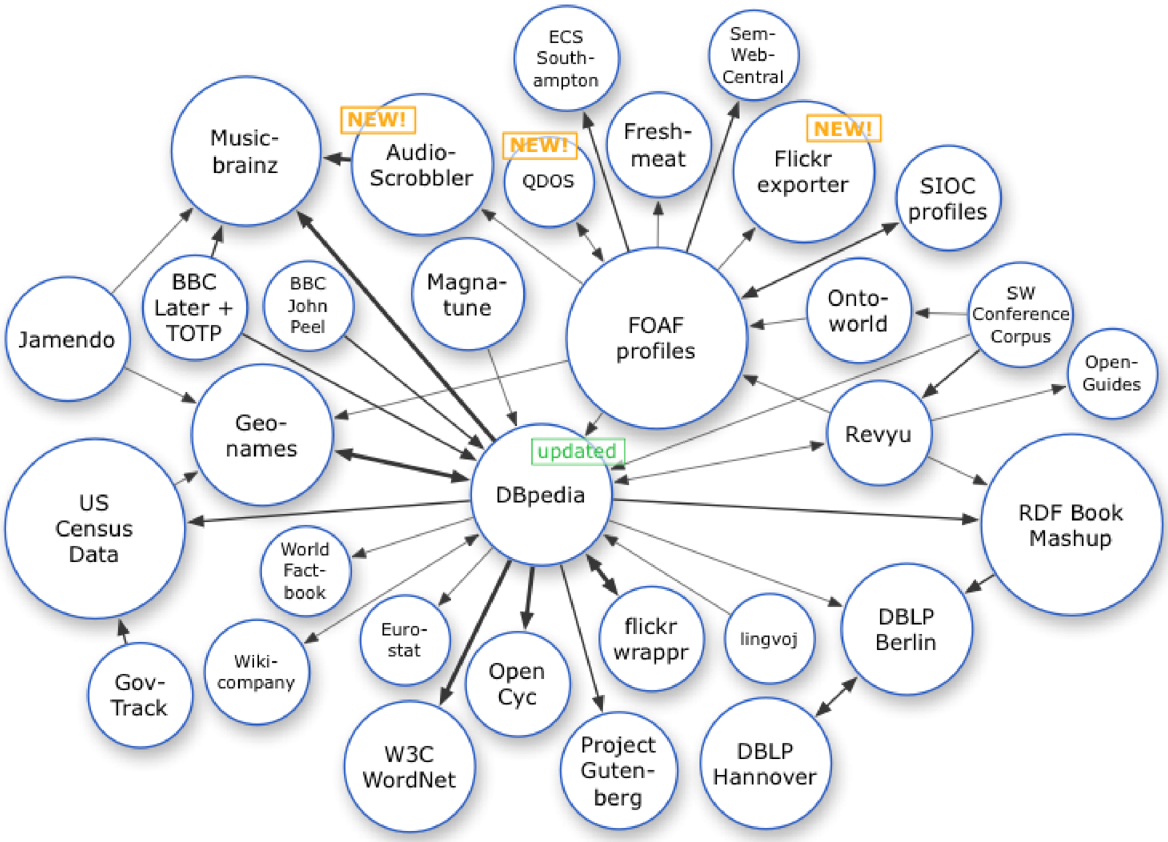

FIGURE A.5: LOD Dataset 2008-02-28 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2008-02-28.png

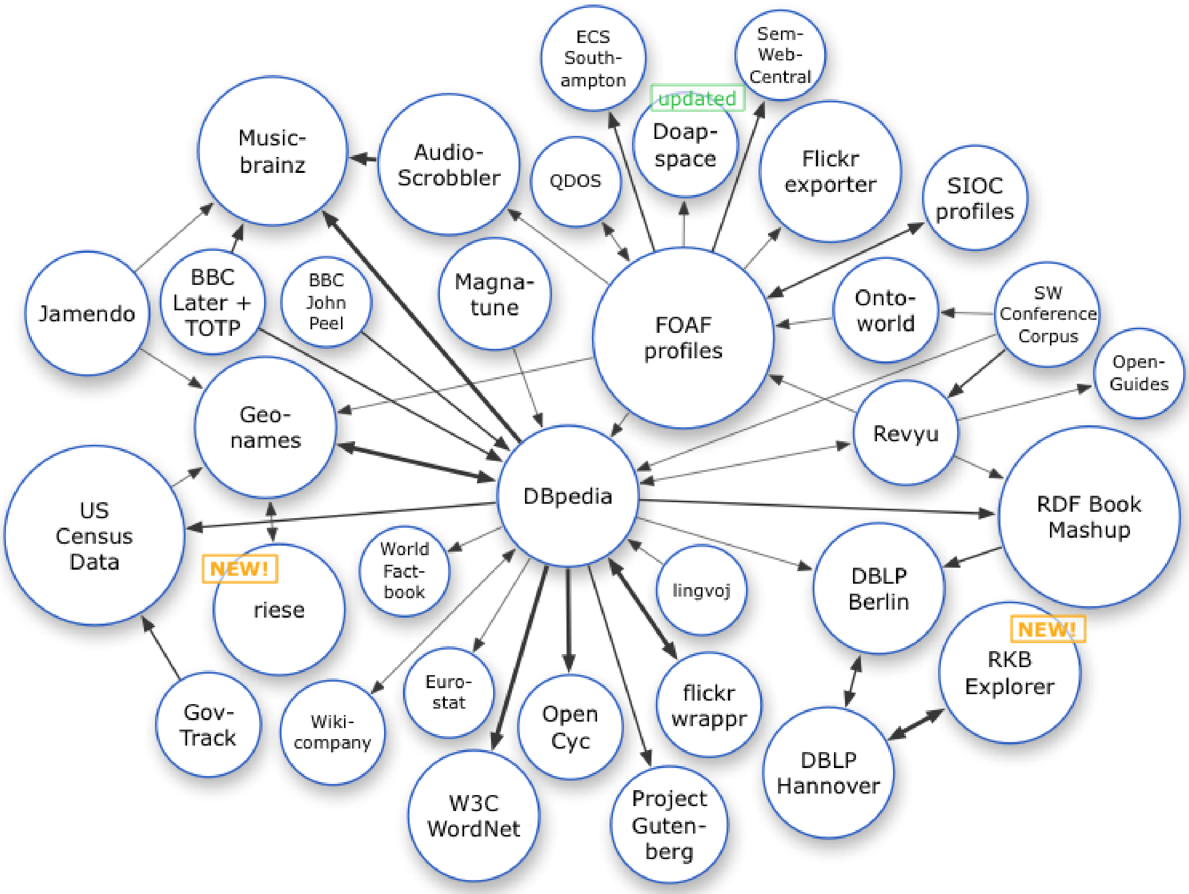

FIGURE A.6: LOD Dataset 2008-03-31 http://richard.cyganiak.de/2007/10/lod/lod-datasets_2008-03-31.png

Appendix B

OntoMedia Ontology

```
<?xml version="1.0"?>

<!DOCTYPE owl [
...
  <!ENTITY base "http://http://purl.org/ontomedia/core/expression">
...
]>

<rdf:RDF xmlns:rdf="&rdf;"
...
 xmlns:dct="&dct;">

  <owl:Ontology rdf:about="&base;">
    <rdfs:label>OntoMedia Core</rdfs:label>
    <dc:title xml:lang="en">OntoMedia Core</dc:title>
    <dc:description xml:lang="en">OntoMedia (Ontology for Media) has been designed to describe
     the interactions occurring in multimedia.</dc:description>
    <dc:creator>Michael O. Jewell (mailto:moj@ecs.soton.ac.uk)</dc:creator>
    <dc:creator>K Faith Lawrence (mailto:kf03r@ecs.soton.ac.uk)</dc:creator>
    <dc:creator>Mischa M Tuffield (mailto:mmt04r@ecs.soton.ac.uk)</dc:creator>
    <dct:created>2005-05-03</dct:created>
    <owl:versionInfo>0.3</owl:versionInfo>
    <owl:imports rdf:resource="http://signage.ecs.soton.ac.uk/ontologies/location" />
  </owl:Ontology>

<!-- Core -->

  <owl:Class rdf:ID="Expression">
    <rdfs:label>Expression</rdfs:label>
    <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string">This class
     represents a piece of information conveyed through a media</rdfs:comment>
  </owl:Class>

  <owl:ObjectProperty rdf:ID="inspired-by">
    <rdfs:label>inspired by</rdfs:label>
```

```
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates that the
   expression was inspired by another</rdfs:comment>
  <rdfs:domain rdf:resource="#Expression"/>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="inspired"/>
  </owl:inverseOf>
  <rdfs:range rdf:resource="#Expression"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="has-shadow">
  <rdfs:label>has shadow</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates that the
   expression is a variation on another, typically darker in nature</rdfs:comment>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="is-shadow_of"/>
  </owl:inverseOf>
  <rdfs:range rdf:resource="#Expression"/>
  <rdfs:domain rdf:resource="#Expression"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="has-spin-off">
  <rdfs:label>has spin off</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates that the
   expression has developed from another</rdfs:comment>
  <rdfs:domain rdf:resource="#Expression"/>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="is-spin-off-of"/>
  </owl:inverseOf>
  <rdfs:range rdf:resource="#Expression"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="is-potentially">
  <rdfs:label>is potentially</rdfs:label>
  <rdfs:range rdf:resource="#Expression"/>
  <rdfs:domain rdf:resource="#Expression"/>
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates that the
   expression is potentially another. For example, it may be a possible future
   version</rdfs:comment>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="is">
  <rdfs:label>is</rdfs:label>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="is-not">
      <rdfs:comment rdf:datatype="&xsd;string">This property indicates that
       the expression is entirely different to another</rdfs:comment>
    </owl:ObjectProperty>
  </owl:inverseOf>
  <rdfs:domain rdf:resource="#Expression"/>
  <rdfs:range rdf:resource="#Expression"/>
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates that the
```

```
    expression is exactly the same as another</rdfs:comment>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:ID="in-context">
    <rdfs:label>in context</rdfs:label>
    <rdfs:comment rdf:datatype="&xsd;string">This property specifies the context
     in which this expression lies.</rdfs:comment>
    <owl:inverseOf>
      <owl:ObjectProperty rdf:ID="includes-expression"/>
    </owl:inverseOf>
    <rdfs:range rdf:resource="#Context"/>
    <rdfs:domain rdf:resource="#Expression"/>
 </owl:ObjectProperty>

  <owl:Class rdf:ID="Entity">
    <rdfs:label>Entity</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Expression" />
  </owl:Class>

  <!-- Entity Subclasses -->
  <!-- Items -->

  <owl:Class rdf:ID="Item">
    <rdfs:comment rdf:datatype="&xsd;string">This class represents an entity which
     may participate in an event within the media. An Item may be abstract or
     physical</rdfs:comment>
    <rdfs:label>Item</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Entity" />
  </owl:Class>

  <owl:Class rdf:ID="Physical-Item">
    <rdfs:comment rdf:datatype="&xsd;string">This class represents a physical entity
     which may participate in an event within the media</rdfs:comment>
    <rdfs:label>Physical Item</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Item" />
  </owl:Class>

  <owl:Class rdf:ID="Abstract-Item">
    <rdfs:comment rdf:datatype="&xsd;string">This class represents an abstract entity
     which may participate in an event within the media</rdfs:comment>
    <rdfs:label>Abstract Item</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Item" />
  </owl:Class>

  <!-- Abstract-Item Subclases -->

  <owl:Class rdf:ID="Context">
    <rdfs:comment rdf:datatype="&xsd;string">This class represents the context in
     which an event or entity exists</rdfs:comment>
    <rdfs:label>Context</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Abstract-Item" />
```

```
</owl:Class>

<owl:Class rdf:ID="Collection">
  <rdfs:comment rdf:datatype="&xsd;string">This class represents a collection of
   entities</rdfs:comment>
  <rdfs:label>Collection</rdfs:label>
  <rdfs:subClassOf rdf:resource="#Abstract-Item" />
</owl:Class>

<!-- Temporal -->

<owl:Class rdf:ID="Timeline">
  <rdfs:comment rdf:datatype="&xsd;string">This class contains a sequence of occurring
   events</rdfs:comment>
  <rdfs:label>Timeline</rdfs:label>
  <rdfs:subClassOf rdf:resource="#Entity" />
</owl:Class>

<owl:Class rdf:ID="Occurrence">
  <rdfs:comment rdf:datatype="&xsd;string">This class represents a single occurrence
   of an event, placing it at a position in a timeline</rdfs:comment>
  <rdfs:label>Occurrence</rdfs:label>
  <rdfs:subClassOf rdf:resource="#Entity" />
</owl:Class>

<owl:ObjectProperty rdf:ID="final-event">
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:range rdf:resource="http://www.w3.org/2002/07/owl#Class"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="initial-event">
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:range rdf:resource="http://www.w3.org/2002/07/owl#Class"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="precedes">
  <rdfs:label>precedes</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property defines the occurrence which
   immediately follows this occurrence</rdfs:comment>
  <rdfs:range rdf:resource="#Occurence"/>
  <rdfs:domain rdf:resource="#Occurence"/>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="follows"/>
  </owl:inverseOf>
</owl:ObjectProperty>

<!-- Events -->

<owl:Class rdf:ID="Event">
  <rdfs:label>Event</rdfs:label>
  <rdfs:subClassOf rdf:resource="#Expression" />
```

```
</owl:Class>

<owl:ObjectProperty rdf:ID="has-subject-entity">
  <rdfs:label>has subject entity</rdfs:label>
  <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string">This property
   specifies the entity which carries out the aim of the event</rdfs:comment>
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:range rdf:resource="#Entity"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="has-object-entity">
  <rdfs:label>has object entity</rdfs:label>
  <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string">This property
   specifies the entity which is the target of the event</rdfs:comment>
  <rdfs:range rdf:resource="#Entity"/>
  <rdfs:domain rdf:resource="#Event"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="has-occurrence">
  <rdfs:label>has occurrence</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property defines any occurrences
   of this event</rdfs:comment>
  <rdfs:range rdf:resource="#Occurrence"/>
  <rdfs:domain rdf:resource="#Event" />
  <owl:inverseOf>
    <owl:FunctionalProperty rdf:ID="occurrence-of"/>
  </owl:inverseOf>
</owl:ObjectProperty>

<owl:DatatypeProperty rdf:ID="summary">
  <rdfs:label>summary</rdfs:label>
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:range rdf:resource="&xsd;string"/>
  <rdfs:comment rdf:datatype="&xsd;string">This property is a plain-text description
   of what occurs in the event</rdfs:comment>
</owl:DatatypeProperty>

<owl:ObjectProperty rdf:ID="precondition">
  <rdfs:label>precondition</rdfs:label>
  <rdfs:range>
    <owl:Class>
      <owl:unionOf rdf:parseType="Collection">
        <owl:Class rdf:about="#Entity"/>
        <owl:Class rdf:about="#Event"/>
      </owl:unionOf>
    </owl:Class>
  </rdfs:range>
  <rdfs:domain rdf:resource="#Event"/>
  <rdfs:comment rdf:datatype="&xsd;string">This property is a state that must exist
   before the event can occur</rdfs:comment>
</owl:ObjectProperty>
```

```
<owl:ObjectProperty rdf:ID="postcondition">
  <rdfs:label>postcondition</rdfs:label>
  <rdfs:range>
    <owl:Class>
      <owl:unionOf rdf:parseType="Collection">
        <owl:Class rdf:about="#Entity"/>
        <owl:Class rdf:about="#Event"/>
      </owl:unionOf>
    </owl:Class>
  </rdfs:range>
  <rdfs:comment rdf:datatype="&xsd;string">This property contains the state which
   should occur as a consequence of this event</rdfs:comment>
  <rdfs:domain rdf:resource="#Event"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="involves">
  <rdfs:label>involves</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property specifies the entities involved
   in this event. Note that this includes the subject and object.</rdfs:comment>
  <rdfs:range rdf:resource="#ontomedia_Entity"/>
  <rdfs:domain rdf:resource="#ontomedia_Event"/>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="involved-in" />
  </owl:inverseOf>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="causes">
  <rdfs:label>causes</rdfs:label>
  <rdfs:comment rdf:datatype="&xsd;string">This property indicates the instigating
   factor of an event, whether it be an item, event, or collection.</rdfs:comment>
  <rdfs:range>
    <owl:Class>
      <owl:unionOf rdf:parseType="Collection">
        <owl:Class rdf:about="#Event"/>
        <owl:Class rdf:about="#Entity"/>
      </owl:unionOf>
    </owl:Class>
  </rdfs:range>
  <rdfs:domain>
    <owl:Class>
      <owl:unionOf rdf:parseType="Collection">
        <owl:Class rdf:about="#Event"/>
        <owl:Class rdf:about="#Entity"/>
      </owl:unionOf>
    </owl:Class>
  </rdfs:domain>
  <owl:inverseOf>
    <owl:ObjectProperty rdf:ID="caused_by"/>
  </owl:inverseOf>
</owl:ObjectProperty>
```

```
<!-- Events Subclasses -->

  <owl:Class rdf:ID="Gain">
    <rdfs:label>Gain</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Event" />
    <rdfs:comment rdf:datatype="&xsd;string">This event class results in an overall
     increase of the entities related to the primary subject or subjects of the
     event</rdfs:comment>
  </owl:Class>

  <owl:Class rdf:ID="Introduction">
    <rdfs:label>Introduction</rdfs:label>
    <rdfs:comment rdf:datatype="&xsd;string">This event class denotes the introduction
     of an entity to the media</rdfs:comment>
    <rdfs:subClassOf rdf:resource="#Event" />
  </owl:Class>

  <owl:Class rdf:ID="Loss">
    <rdfs:label>Loss</rdfs:label>
    <rdfs:comment rdf:datatype="&xsd;string">This event class results in an overall
     reduction of the entities related to the primary subject or subjects of the
     event</rdfs:comment>
    <rdfs:subClassOf rdf:resource="#Event" />
  </owl:Class>

  <owl:Class rdf:ID="Transformation">
    <rdfs:comment rdf:datatype="&xsd;string">This event class results in no gain or loss
     of attributes or entities, merely alteration</rdfs:comment>
    <rdfs:label>Transformation</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Event" />
  </owl:Class>

  <owl:Class rdf:ID="Action">
    <rdfs:comment rdf:datatype="&xsd;string">This event class describes an action sequence
     (ie no plot)</rdfs:comment>
    <rdfs:label>Action</rdfs:label>
    <rdfs:subClassOf rdf:resource="#Event" />
  </owl:Class>

  <owl:ObjectProperty rdf:ID="from">
    <rdfs:label>from</rdfs:label>
    <rdfs:comment rdf:datatype="&xsd;string">This property specifies the entity which
     is being transformed</rdfs:comment>
    <rdfs:range rdf:resource="#Entity"/>
    <rdfs:domain rdf:resource="#Transformation"/>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:ID="to">
    <rdfs:label>to</rdfs:label>
    <rdfs:comment rdf:datatype="&xsd;string">This property specifies the resultant
```

```
    entity</rdfs:comment>
    <rdfs:range rdf:resource="#Entity"/>
    <rdfs:domain rdf:resource="#Transformation"/>
  </owl:ObjectProperty>

  <!-- Unsorted -->

  <owl:ObjectProperty rdf:ID="has_parody">
    <rdfs:domain rdf:resource="#Expression"/>
    <rdfs:range rdf:resource="#Expression"/>
    <owl:inverseOf>
      <owl:ObjectProperty rdf:ID="is_parody_of"/>
    </owl:inverseOf>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:ID="occurs">
    <rdfs:range rdf:resource="#locspec_Location_Specifier"/>
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#FunctionalProperty"/>
    <rdfs:domain rdf:resource="#Instant_Occurence"/>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:about="#allows_existance_of">
    <rdfs:domain rdf:resource="#Context"/>
    <owl:inverseOf rdf:resource="#exists_in"/>
    <rdfs:range rdf:resource="#Expression"/>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:ID="start_point">
    <rdfs:domain rdf:resource="#Period_Occurence"/>
    <rdfs:range rdf:resource="#locspec_Location_Specifier"/>
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#FunctionalProperty"/>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:ID="duration">
    <rdfs:domain rdf:resource="#Period_Occurence"/>
    <rdfs:range>
      <owl:Class>
        <owl:unionOf rdf:parseType="Collection">
          <owl:Class rdf:about="#locspec_Location_Specifier"/>
          <owl:Class rdf:about="#Dimension"/>
        </owl:unionOf>
      </owl:Class>
    </rdfs:range>
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#FunctionalProperty"/>
  </owl:ObjectProperty>

  <owl:ObjectProperty rdf:about="#contains">
    <owl:inverseOf rdf:resource="#contained_by"/>
    <rdfs:range rdf:resource="#Expression"/>
    <rdfs:domain>
      <owl:Class>
```

```
      <owl:unionOf rdf:parseType="Collection">
        <owl:Class rdf:about="#Expression"/>
        <owl:Class rdf:about="#Expression"/>
      </owl:unionOf>
    </owl:Class>
  </rdfs:domain>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:about="#follows">
  <rdfs:range rdf:resource="#Occurence"/>
  <rdfs:domain rdf:resource="#Occurence"/>
  <owl:inverseOf rdf:resource="#precedes"/>
  <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string">Follows should
   specify both timeline and event IDs where there is more than one timeline or over
   two events</rdfs:comment>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="timeline_ref">
  <rdfs:range rdf:resource="#Timeline"/>
  <rdfs:domain rdf:resource="#Occurence"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:about="#is_parody_of">
  <rdfs:range rdf:resource="#Expression"/>
  <owl:inverseOf rdf:resource="#has_parody"/>
  <rdfs:domain rdf:resource="#Expression"/>
</owl:ObjectProperty>

<owl:ObjectProperty rdf:ID="end_point">
  <rdfs:range rdf:resource="#locspec_Location_Specifier"/>
  <rdfs:domain rdf:resource="#Period_Occurence"/>
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#ObjectProperty"/>
</owl:ObjectProperty>

<owl:FunctionalProperty rdf:ID="TPQ">
  <rdfs:range rdf:resource="#locspec_Location_Specifier"/>
  <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string"
  >Terminus Post Quem</rdfs:comment>
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#ObjectProperty"/>
  <rdfs:domain rdf:resource="#Occurence"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:ID="type">
  <rdfs:range rdf:resource="http://www.w3.org/2001/XMLSchema#string"/>
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#DatatypeProperty"/>
  <rdfs:domain rdf:resource="#Item"/>
</owl:FunctionalProperty>

<owl:FunctionalProperty rdf:about="#occurence_of">
  <rdfs:domain rdf:resource="#Occurence"/>
  <rdfs:range rdf:resource="#Expression"/>
```

```
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#ObjectProperty"/>
    <owl:inverseOf rdf:resource="#has_occurence"/>
  </owl:FunctionalProperty>

  <owl:FunctionalProperty rdf:ID="TAQ">
    <rdfs:domain rdf:resource="#Occurence"/>
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#ObjectProperty"/>
    <rdfs:comment rdf:datatype="http://www.w3.org/2001/XMLSchema#string"
    >Terminus Ante Quem</rdfs:comment>
    <rdfs:range rdf:resource="#locspec_Location_Specifier"/>
  </owl:functionalproperty>

  <owl:functionalproperty rdf:id="initial_event">
    <rdfs:comment rdf:datatype="http://www.w3.org/2001/xmlschema#string"
    >the first event which begins this sequence</rdfs:comment>
    <rdfs:range rdf:resource="#event"/>
    <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#objectproperty"/>
    <rdfs:domain rdf:resource="#event"/>
  </owl:functionalproperty>

</rdf:rdf>
```

Appendix C

Semantic Logger : Data Capture tools

Below presents the user with two pieces of code which post data to the Semantic Logger. All of the software used to capture autobiographical data can be found on the Semantic Logger download page http://akt.ecs.soton.ac.uk:8080/downloads.php

How to automatically import an RDF file from your local machine to your Semantic Log:

Java program to post RDF files through HTTP POST requests http://akt.ecs.soton.ac.uk:8080/downloads/postRDFjava.tar.bz2

Usage:

```
java postRDFjava username password is-public syntax fileName class
```

username : Your semantic logger username

password : Your semantic logger password

is-public : true/false

syntax : one of auto, rdfxml, turtle, ntriples

rdfstream : A parsable RDF file

class : one of std, ical, ffbook, ffhist.

Perl version of POST RDF http://akt.ecs.soton.ac.uk:8080/downloads/postRDF.pl

An equivalent but faster perl program to post RDF files through HTTP POST.

Usage:

```
perl postRDF username password is-public syntax fileName
```

username : Your semantic logger username

password : Your semantic logger password

is-public : true/false

syntax : one of auto, rdfxml, turtle, ntriples

rdfstream : A parsable RDF file

class : one of std, ical, ffbook, ffhist.

Appendix D

Photocopain Example

The following example RDF was generated by Photocopain. First group of triples relate to figure D.1.

```
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix exif: <http://triplestore.aktors.org/additional-exif-rdf-ns#> .
@prefix department: <http://aktors.org/photocopain/departments#> .
@prefix city: <http://aktors.org/photocopain/cities#> .
@prefix area: <http://aktors.org/photocopain/areas#> .
@prefix style: <http://aktors.org/photocopain/photo-styles#> .
@prefix feature: <http://aktors.org/photocopain/photo-features#> .
@prefix subject: <http://aktors.org/photocopain/photo-subjects#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix reach: <http://aktors.org/photocopain/lens-reach#> .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_3276-20070719145001.JPG>
        exif:objectSize 3.27 ;
```

FIGURE D.1: The Lloyds building in London http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_2063.JPG

```
        exif:style style:bokeh ;
        exif:style reach:telephoto ;
        exif:location \'Building 32 Seminar Room\\, Level 3\' ;
        exif:event \'Demo-Session - ISTIC\' ;
        exif:subject feature:artificial ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_0613.JPG>
        exif:objectSize 0.49 ;
        exif:style style:bokeh ;
        exif:style reach:telephoto ;
        exif:location 'Edinburgh' ;
        exif:event 'WWW2006' ;
        exif:depicts subject:person ;
        exif:style style:portrait ;
        exif:subject feature:natural ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_3274-20070719133228.JPG>
        exif:objectSize 0.92 ;
        exif:style reach:normal ;
        exif:location \'Building 32 Seminar Room\\, Level 3\' ;
        exif:event \'Demo-Session - ISTIC\' ;
        exif:depicts subject:person ;
        exif:subject feature:artificial ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_0622.JPG>
        exif:objectSize 7.21 ;
        exif:style reach:normal ;
        exif:location \'Edinburgh\' ;
        exif:event \'WWW2006\' ;
        exif:subject feature:artificial ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_2063.JPG>
        exif:objectSize 16.03 ;
        exif:style style:dof ;
        exif:style reach:wide ;
        exif:location \'Central London\\, EC3\' ;
        exif:location \'Central London\\, EC3\' ;
        exif:event \'Pictures of Lloyds Building\' ;
        exif:event \'Pictures\' ;
        exif:long \'-0.080895424\' ;
        exif:lat \'51.512296200\' ;
        exif:alt \'54.451782\' ;
        exif:subject feature:natural ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_3278-20070719145001.JPG>
        exif:objectSize 2.19 ;
```

```
        exif:style style:bokeh ;
        exif:style reach:telephoto ;
        exif:location \'Building 32 Seminar Room\\, Level 3\' ;
        exif:event \'Demo-Session - ISTIC\' ;
        exif:subject feature:artificial ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
<http://akt.ecs.soton.ac.uk:8080/photocopain/accounts/mmt04r/DSC_2425-2006121295933-20070207132442.jpg>
        exif:objectSize 2.27 ;
        exif:style reach:normal ;
        exif:location \'British Library\' ;
        exif:event \'Memories for Life: The Future of our Pasts\' ;
        exif:depicts subject:person ;
        exif:subject feature:artificial ;
        foaf:maker <http://www.ecs.soton.ac.uk/~mmt04r/foaf/mischa.rdf> ;
        a exif:Photo .
```

Bibliography

M Addis, M Boniface, S Goodall, P Grimwood, S Kim, P Lewis, K Martinez, and A Stevenson. SCULPTEUR: Towards a new paradigm for multimedia museum information handling. In *Proceedings of Semantic Web ISWC 2003*, pages 582–596, 2003.

M. Addis, P. H. Lewis, K. Martinez, J. Stevenson, and F. Giorgini. New ways to search, navigate and use multimedia museum collections over the web. In *In Proceedings of Museums and the Web 2005*, 2005.

M. Admed, H. H. Hoang, M. S. Karim, S. Khusro, M. Lanzenberger, K. Latif, E. Michlmayr, K. Mustofa, H. T. Hguyen, A. Rauber, A. Schatten, M. N. Tho, and A. M. Tjoa. 'SemanticLIFE' - A framework for managing information of a human lifetime. In *In Proceedings of the 6th International Conference on Information Integration and Web-based Applications and Services*, Jakarta, Indonesia, September 2004.

G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *j-TOIS*, 23(1):103–145, January 2005. ISSN 1046-8188.

H. Alani, S. Dasmahapatra, K. O'Hara, and N. R. Shadbolt. Identifying communities of practice through ontology network analysis. *IEEE Intelligent Systems*, 18(2):18–25, 2003a.

Harith Alani, Sanghee Kim, David E. Millard, Mark J. Weal, Wendy Hall, Paul H. Lewis, and Nigel R. Shadbolt. Automatic ontology-based knowledge extraction from web documents. *IEEE Intelligent Systems*, 18(1):14–21, January 2003b.

Anita L. Allen. Dredging up the past: Lifelogging, memory and surveillance. *University of Chicago Law Review*, 75:47–74, 2008.

P. J. Allen, R. Vaccaro, and G. Presutti. ARTISTE: An integrated art analysis and navigation environment. *Cultivate Interactive*, 1, 2000.

N. E. Apostoloff and A. Zisserman. Who are you? real-time person identification. In *Proceedings of the British Machine Vision Conference*, 2007.

P. Bailey. Searching for storiness: Story generation from a reader's perspective. In *Narrative Intelligence Symposium AAAI 1999 Fall Symposium Series*. AAAI, 1999.

M. Bal. *Introduction to the theory of Narrative.* University of Toronto Press, Toronto:Canada, 2nd edition, 1997.

N. R. Ball and G. Ragsdell. Telling knowledge management stories: creating and implementing a new technology-society knowledge base. *OR Insight*, 16(4):5–12, 2003.

G. Bell and J. Gemmell. A Digital Life. *Scientific American*, pages 40–47, February 2007.

T. Berners-Lee. *Weaving the Web*. Orion Business Books, 1999.

T. Berners-Lee. Keynote speech: Semantic web status and direction. In *International Semantic Web Conference 2003(IWSC-03)*. Springer, Berlin, Heidelberg, 2003.

T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5), May 2001.

Tim Berners-Lee. Looking back, looking forward: The process of designing things in a very large spaces. Inaugural lecture, Southampton University, 2007.

Mark Bernstein. Card shark and thespis: exotic tools for hypertext narrative. In *HYPERTEXT '01: Proceedings of the twelfth ACM conference on Hypertext and Hypermedia*, pages 41–50, New York, NY, USA, 2001. ACM Press. ISBN 1-59113-420-7.

E. Berry, N. Kapur, L. Williams, S. Hodges, P. Watson, G. Smyth, J. Srinivasan, R. Smith, B. Wilson, and K Wood. The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis. In *Encephalitis: Assessment and Rehabilitation Across the Lifespan (Special Issue of Neuropsychological Rehabilitation)*, volume 17 (4/5), pages 582–681, 2007. ISBN 978-1-84169-836-6.

I. M. Bilasco, J. Gensel, and M. Villanova-Oliver. Stamp: A model for generating adaptable multimedia presentations. *Int. Journal Multimedia and Applications*, 25 (3):361–375, 2005.

Chris Bizer, Sren Auer, Georgi Kobilarov, Jens Lehmann, Christian Becker, and Sebastian Hellmann. Dbpedia - querying wikipedia like a database and an interlinking-hub in the web of data, April 2009. Querying Wikipedia Like a Database (4/4/2009) FU Berlin, Universitt Leipzig.

D. Blair and T. Mayer. Tools for interactive virtual cinema. In R. Trappl and P. Petta, editors, *In Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents*. Springer, 1997.

S. Bocconi, F. Nack, and L. Hardman. Using rhetorical annotations for generating video documentaries. In *Proceedings of the IEEE International Conference on Multimedia and Expo.* (ICME), July 2005.

Joseph Bonneau and Sren Preibusch. The privacy jungle: On the market for data protection in social networks. In *The Eighth Workshop on the Economics of Information Security (WEIS 2009)*, 2009.

Dhruba Borthakur. *The Hadoop Distributed File System: Architecture and Design*. The Apache Software Foundation, 2007.

d. m. boyd and N. B. Ellison. Social networking sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication*, 13(1):11, 2007.

J. Broekstra, A. Kampman, and F. van Harmelen. Sesame: A generic architecture for storing and querying RDF and RDF schema, 2002.

Vannevar Bush. As We May Think. *The Atlantic Monthly*, 176(1):101–108, July 1945.

J. Campbell. *The Hero with a Thousand Faces*. Fontana Press, London:UK, 1949.

J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, November 1986. ISSN 0162-8828.

Marc Cavazza, Fred Charles, and Steven J. Mead. Emergent situations in interactive storytelling. In *SAC '02: Proceedings of the 2002 ACM symposium on Applied computing*, pages 1080–1085, New York, NY, USA, 2002. ACM Press. ISBN 1-58113-445-2.

S. Cayzer and P. Castagna. How to build a snippet manager. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

Ajay Chakravarthy, Vitaveska Lanfranchi, and Fabio Ciravegna. Requirements for multimedia document enrichment, 2006.

S. Chatman. *Story and Discourse*. Cornell University Press, New York, 1978.

William C. Cheng, Leana Golubchik, and David G. Kay. Total recall: are privacy changes inevitable? In *CARPE'04: Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pages 86–92, New York, NY, USA, 2004. ACM. ISBN 1-58113-932-2.

F. Ciravegna. Designing adaptive information extraction for the semantic web in amilcare. *Annotation for the Semantic Web*, 2003.

F. Ciravegna, S. Chapman, A. Dingli, and Y. Wilks. Learning to harvest information for the semantic web. In *Proceedings of the 1st European Semantic Web Symposium*, Heraklion, Greece, May 2004.

N. Connell, J. Klein, and E. Meyer. Narrative approaches to the transfer of organisational knowledge. *Knowledge Management Research and Practice*, 2:184–193, 2005.

Kevin Coughlin. Tracking himself, so the FBI won't have to. Digital Life with the Star Ledger, 28th October 2007, 2007.

Marc Davis, Simon King, Nathan Good, and Risto Sarvas. From context to content: leveraging context to infer media metadata. In *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia*, pages 188–195, New York, NY, USA, 2004. ACM Press.

Marc Davis, Michael Smith, John Canny, Nathan Good, Simon King, and Rajkumar Janakiraman. Towards context-aware face recognition. In *MULTIMEDIA '05: Proceedings of the 13th annual ACM international conference on Multimedia*, pages 483–486, New York, NY, USA, 2005. ACM Press.

Marc Davis, Michael Smith, Fred Stentiford, Adetokunbo Bambidele, John Canny, Nathan Good, Simon King, and Rajkumar Janakiraman. Using context and similarity for face and location identification. In *Proceedings of the IS&T/SPIE 18th Annual Symposium on Electronic Imaging Science and Technology Internet Imaging VII*. IS&T/SPIE Press, 2006.

F. Dawson and D. Stenerson. Internet calendaring and scheduling core object specification (iCalendar), 1998.

Lina Dib. Memory as concept in the design of digital recording devices. *Alterites*, 5(1): 38–53, 2008.

A. Dix. The ultimate interface and the sums of life? *Interfaces*, (50):16, 2002.

Alan Dix. Information processing, context and privacy alan dix. In *HumanComputer Interaction INTERACT'90*, pages 15–20. North-Holland, 1990.

Martin Dodge and Rob Kitchin. Outlines of a world coming into existence: pervasive computing and ethics of forgetting. In *Environment and PLanning B: Planning and Design*, volume 34, pages 431–445, 2007.

M. Doerr, J. Hunter, and C. Lagoze. Towards a core ontology for information integration. *Journal of Digital Information*, 4(1), April 2003.

David P Dupplaw, Madalina Croitoru, Antonis Loizou, Srinandan Dasmahapatra, Paul Lewis, Mischa M. Tuffield, and Liang Xiao. Multimedia markup tools for OpenKnowledge. In *1st Workshop on Multimedia Annotation and Retrieval enabled by Shared Ontologies*, December 2007.

EXIF. Exchangeable image file format for digital still cameras: EXIF version 2.2. Technical report, Japan Electronics and Information Technology Industries Association, 2002.

Kateryna Falkovych and Stefano Bocconi. Creating a Semantic-based Discourse Model for Hypermedia Presentations: (Un)discovered Problems. In *Workshop on Narrative, Musical, Cinematic and Gaming Hyperstructure*, Salzburg, Austria, September 2005.

Kateryna Falkovych and Frank Nack. Context aware guidance for multimedia authoring: harmonizing domain and discourse knowledge. *Multimedia Systems*, 11(3):226–235, 2006.

Melanie Feinberg and Ryan Shaw. Action: An annotation framework for events in video. In *Proceedings of Semantic Web ISWC 2004*, 2004.

Andres Garcia, Martin Szomszor, Harith Alani, and Oscar Corcho. Preliminary results in tag disambiguation using dbpedia. In *Knowledge Capture (K-Cap'09) - First International Workshop on Collective Knowledge Capturing and Representation - CKCaR'09*, September 2009.

J. Gemmel, G. Bell, R. Lueder, S. Drucker, and C. Wong. MyLifeBits: fulfilling the memex vision. In *MULTIMEDIA '02: Proceedings of the 10th ACM international conference in Multimedia*, pages 235–238, 2002.

J. Gemmell, A. Aris, and R. Lueder. Telling stories with MyLifeBits. *ICME 2005*, 8: 6–9, July 2005.

J. Gemmell, R. Lueder, and G. Bell. The MyLifeBits lifetime store. In *Proceedings of ACM SIGMM 2003 Workshop on Experiential Telepresence (ETP 2003)*, Berkeley, CA, November 2003.

J. Geurts, S. Bocconi, J. van Ossenbruggen, and L. Hardman. Towards ontological-driven discourse: From semantic graphs to multimedia presentations. CWI Technical Report INS-R0305, 2003.

Dan Gillmor. Bloggers and mash. *New Scientist*, 197(2647):44+, March 2008.

F. Giunchiglia, P. Shvaiko, and M. Yatskevich. S-match: An algorithm and an implementation of semantic matching. In *European Semantic Web Symposium*, pages 61–75, 2004.

Christian Halaschek-Wiener, Andrew Schain, Michael Grove, Bijan Parsia, and Jim Hendler. Management of digital images on the semantic web. In *Proceedings of the International Semantic Web Conference ISWC05*, Galloway, Ireland, 2005.

S. Harris and N. Gibbins. 3store: Efficient bulk RDF storage. In *PSSS'03*, pages 1–15, Sanibel, Fl, 2003.

S. Harris, N. Lamb, and N. Shadbolt. 4store: The design and implementation of a clustered rdf store. In *ISWC 2009 SSWS Workshop 09*, Washington DC, 2009.

Stephen Harris. SPARQL query processing with conventional relational database systems. In *WISE Workshops*, pages 235–244, 2005.

T. Heath, E. Motta, and M. Dzbor. Context as a foundation for a Semantic Desktop. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

Steve Hodges, Lyndsay Williams, Emma Berry, Shahram Izadi, James Srinivasan, Alex Butler, Gavin Smyth, Narinder Kapur, and Kenneth R. Wood. SenseCam: A retrospective memory aid. In Paul Dourish and Adrian Friday, editors, *Ubicomp*, volume 4206 of *Lecture Notes in Computer Science*, pages 177–193. Springer, 2006.

Nancy A. Van House and Marc Davis. The social life of cameraphone images. In *Proceedings of the Pervasive Image Capture and Sharing: New Social Practices and Implications for Technology Workshop (PICS 2005) at the Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*. UbiComp 2005, 2005.

D. Huynh, D. Karger, and D. Quan. Haystack: A platform for creating, organising, and visualising information using RDF. In *Semantic Web Workshop, The Eleventh World Wide Web Conference.*, 2002.

David Huynh, Stefano Mazzocchi, and David R. Karger. Piggy bank: Experience the semantic web inside your web browser. In Yolanda Gil, Enrico Motta, V. Richard Benjamins, and Mark A. Musen, editors, *International Semantic Web Conference*, volume 3729 of *Lecture Notes in Computer Science*, pages 413–430. Springer, 2005.

T. Iofciu, C. Kohlshutter, W. Nejdl, and R. Paiu. Keywords and RDF fragments: Integrating metadata and desktop search in beagle++. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

IPTO 2003. Lifelog proposer information pamphlet. Defense Advanced Research Projects Agency Information Processing Technology Office document PIP_03-30, 2003.

M. O. Jewell, F. Lawrence, M. M. Tuffield, M. S. Nixon, A Prügel-Bennett, N. R. Shadbolt, D. E. Millard, and m. c. schraefel. OntoMedia: An ontology for the representation of heterogeneous media. In *Proceedings of MultiMedia Information Retrieval (MMIR) Workshop at SIGIR*. SIGIR, August 2005a.

Michael O Jewell. Motivated music: Automatic soundtrack generation for film. Mike O. Jewell's PhD thesis at the University of Southampton, 2007.

Michael O Jewell, Mark S. Nixon, and Adam Prügel-Bennett. State-based sequencing: Directing the evolution of music. In *International Computer Music Conference*, 2005b.

Krissi M. Jimroglou. A camera with a view: JenniCam, visual representation and cyborg subjectivity. *Information, Communication and Society*, 2:439–453, 1999.

William Jones. *Keeping Found Things Found: The Study of Practise of Personal Information Management*. Morgan Kaufmann, Burlington MA, 2008.

D. Joyce, P. H. Lewis, R. Tansley, M. Dobie, and W. Hall. Semiotics and agents for integrating and navigating through multimedia representations of concepts. *Proceedings of Storage and Retrieval for Media Databases*, 3972(3972):120–131, 2000.

Vaiva Kalnikaité and Steve Whittaker. Software or wetware?: discovering when and why people use digital prosthetic memory. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 71–80, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9.

S. Kim, H. Alani, D. E. Millard, M. J. Weal, P. H. Lewis, W. Hall, and N. R. Shadbolt. Generating tailored biographies with automatically annotated fragments from the web. In *In Proceedings of Workshop on Semantic Authoring , Annotation, and Knowledge Markup*, pages 1–6. 15th European Conference on AI (ECAI), 2002.

R. Klamma, M. Spaniol, and M. Jarke. MECCA: Hypermedia capturing of collaborative scientific discourses about movies informing science. *Int. Journal of an Emerging Discipline: Special Series on Issues in Informing Clients using Multimedia Communications*, 8:3–38, 2005.

Jon M. Kleinberg. Hubs, authorities, and communities. *ACM Comput. Surv.*, page 5, 1999. ISSN 0360-0300.

C. Lagoze and J. Hunter. The ABC ontology and model. *Journal of Digital Information*, 2(2), November 2001.

B. Laurel, J. Bates, A. Don, and R. Strickland. Interface and narrative arts: Contributions from narrative, drama, and film. In *Proc. of CHI-91*, pages 381–383, New Orleans, LA, 1991.

K. F. Lawrence. The web of community trust - amateur fiction online: A case study in community focused design for the semantic web. Faith Lawrence's PhD thesis at the University of Southampton, 2007.

K. F. Lawrence and m. c. schraefel. Amateur fiction online - the web of community trust. In *Proceedings of the 1st AKT Doctoral Colloquium*. AKT, June 2005.

K. F. Lawrence, M. M. Tuffield, M. O. Jewell, M. S. Nixon, A. Prügel-Bennett, N. R. Shadbolt, D. E. Millard, and m. c. schraefel. OntoMedia: Creating an ontology for marking up the contents of heterogeneous media. In *Ontology Patterns for the Semantic Web Workshop at ISWC05*. ISWC, November 2005.

Ryan Lee. Scalability report on triple store applications. Technical report, MIT, 2004.

S. Little, J. Geurts, and J. Hunter. Dynamic generation of intelligent multimedia presentations through semantic inferencing. In *Proceedings of 6th European Conf on Research and Advanced Techniques for Digital Libraries*, pages 159–189. Springer - ECDL, 2002.

Antonis Loizou. How to recommend music to film buffs: Enabling the provision of recommendations from multiple domains. Antonis Loizou's PhD thesis at the University of Southampton, January 2009.

m. c. schraefel, M. Karam, and S. Zhao. mSpace: interaction design for user-determined, adaptable domain exploration in hypermedia. In P. De Bra, editor, *Proceedings of AH 2003: Workshop on Adaptive Hypermedia and Adaptive Web Based Systems.*, pages 217–235, Nottingham, UK, 2003. Springer.

F. Manola and E. Miller. RDF primer: W3C recommendation. http://www.w3.org/TR/rdf-primer/, 2004.

Steve Mann and Hal Niedzviecki. *Cyborg: Digital Destiny and Human Possibility in the Age of the Wearable Computer*. Random House, New York, 2001.

Cameron Marlow, Mor Naaman, Danah Boyd, and Marc Davis. HT06, tagging paper, taxonomy, Flickr, academic article, to read. In *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia*, pages 31–40, New York, NY, USA, 2006. ACM Press. ISBN 1-59593-417-0.

M. Mateas. An Oz-centric review of interactive drama and believable agents. Technical report, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1997.

M. Mateas. A neo-aristotelian theory of interactive drama. In *AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment*. AAAI, 2000.

M. Mateas and P. Sengers. An introduction to narrative intelligence. In *Narrative Intelligence Symposium AAAI 1999 Fall Symposium Series*. AAAI, 1999.

G. G. Mateos and C. V. Chicote. A unified approach to face detection, segmentation and location using HIT maps. In *In Syposium Nacional de Re-conocimiento de Formas y Anlisis de Imgenes*, Benicasim, Castelln, 2001.

M. Matheas and A Stern. Faade, an experiment in building fully realised interactive drama. In *Games Design Track*. Games Developers Conference, 2003.

Viktor. Mayer-Schönberger. *Delete : the virtue of forgetting in the digital age / Viktor Mayer-Schnberger*. Princeton University Press, Princeton, N.J. ; Woodstock :, 2009. ISBN 9780691138619 0691138613.

B. McBride. Jena: Implementing the RDF model and syntax specification. In *Proceedings of the 2nd International Workshop on the Semantic Web, SemWeb*, 2001.

S. McCloud. *Understanding Comics, The Invisible Art.* HarperPerennial, NY:USA, 1994.

D. L. McGuinness and F. v. Harmelen. OWL Web Ontology Language Overview: W3C Recommendation. http://www.w3.org/TR/owl-features/, 2004.

J. F. Meech. Narrative theories as contextual constraints for agent interaction. In *Proceedings of the AAAI '99 Fall Symposium on Narrative Intelligence*. AAAI, 1999.

J. Meechan. *The metanovel: writing stories by computer*. PhD Thesis, Yale University, 1976.

D. E. Millard, N. E. Gibbons, and M. J. Weal D. T. Michselides. Mind the semantic gap. In *In Proceedings of ACM Hypertext 2005*, pages 54–63. SIGWEB, Sep 2005.

I. Millard, D. De Roure, and N. R. Shadbolt. The use of ontologies in contextually aware environments. In *In Proceedings of First International Workshop on Advanced Context*, pages 42–47, 2004.

R. Mizoguchi, J. Vanwelkenhuysen, and M. Ikeda. Task ontology for reuse of problem solving knowledge. *Towards very Large Knowledge Bases: Knowledge Building and Knowledge Sharing*, pages 46–59, 1995.

K. Molle and S. Decker. Harvesting desktop data for semantic blogging. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

P. Mulholland, T. Collins, and Z. Zdrahal. Story fountain: intelligent support for story research and exploration. In *Proceedings of 9th international conference on Intelligence User Interfaces*, pages 62–69. International Conference on Intelligent User Interfaces, 2004.

J. Murray. *Hamlet on the Holodeck*. MIT University Press, Cambridge:Mass, 1998.

Mor Naaman, Susumu Harada, Qianying Wang, Hector G. Molina, and Andreas Paepcke. Context data in geo-referenced digital photo collections. In *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia*, pages 196–203, New York, NY, USA, 2004a. ACM Press.

Mor Naaman, Susumu Harada, QianYing Wang, and Andreas Paepcke. Adventures in space and time: Browsing personal collections of geo-referenced digital photographs. Technical report, Stanford University, April 2004b. Available at http://dbpubs.stanford.edu/pub/2004-26.

Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006.

Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy*, pages 111–125. IEEE Computer Society, 2008.

Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. Mar 2009.

K. Nelson. *Narratives from the Crib*. Harvard University Press, Cambridge:Mass, 1989.

N. Noy and D. McGuiness. Ontology development 101: A guide to creating your first ontology. Technical report, Stanford University, 2001.

Kieron O'Hara, Richard Morris, Nigel R. Shadbolt, Graham J. Hitch, Wendy Hall, and Neil Beagrie. Memories for life: A review of the science and technology. *Journal of the Royal Society Interface*, 3(8):351–365, June 2006.

Kieron O'Hara and Nigel R. Shadbolt. *The Spy in the Coffee Machine: The End of Privacy as we Know It*. Oneworld, Oxford, UK, 2008.

Kieron O'Hara, Mischa Tuffield, and Nigel Shadbolt. Lifelogging: Privacy and empowerment with memories for life. *Identity in the Information Society*, 1(2), 2009.

Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web, 1999.

Alexandre Passant. :me owl:sameas flickr:33669349@n00 . In *Proceedings of Linked Data on the Web Workshop (LDOW2008) at WWW 2008*. ACM Press, 2008.

Oz Project. Carnegie mellon university. http://www.cs.cmu.edu/afs/cs.cmu.edu/project-/oz/web/oz.html, 1989-2002.

V. Propp. *Morphology of a Folktale*. University of Texas Press, Texas:USA, 2nd edition, 1968. Translated by L. Scott and L. Wagner.

K. Sitaker R. Khare, D. Cutting and A. Rifkin. Nutch: A flexible and scalable open-source web search engine. CommerceNet Labs, CN-TR-04-04, November 2004.

Yves Raimond, Samer Abdallah, Mark Sandler, and Frederick Giasson. The music ontology. In *ISMIR 2007: 8th International Conference on Music Information Retrieval*, Vienna, Austria, September 2007.

J. Richter, M. Volkel, and H. Haller. DeepaMenta - a semantic desktop. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

M. O. Riedl and R. M. Young. Character-focused narrative generation for execution in virtual worlds. In *Proceedings of the Second International Conference on Virtual Storytelling*, Toulouse France, 2003.

D. Robertson. A lightweight coordination calculus for agent systems. *Lecture Notes in Computer Science*, 3476:183–197, 2005.

Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, New Jersey, 2nd edition, 2003. ISBN 0-13-790395-2.

L. Rutledge, M. Alberink, R. Brussee, S. Pokraev, W. van Dieten, and M. Veenstra. Finding the story - broader applicability of semantics and discourse for hypermedia generation. In *Proceedings of 14th ACM conference on Hypertext and Hypermedia*, pages 67–76, Nottingham, UK, 2003.

Lloyd Rutledge, Jacco van Ossenbruggen, and Lynda Hardman. Making RDF presentable: integrated global and local semantic web browsing. In *WWW '05: Proceedings of the 14th international conference on World Wide Web*, pages 199–206, New York, NY, USA, 2005. ACM Press.

L. Sauermann, A. Bernandi, and A. Dengel. Overview and outlook on the semantic desktop. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

K. G. Saur. Functional requirements of bibliographic records: final report. Technical report, IFLA Study Group on the Functional Requirements of Bibliographic Records, München, 1998.

Daniel L. Schacter. *The Seven Sins of Memory: How the Mind Forgets and Remembers.* Houghton Mifflin, New York, 2001.

R. Schank. *Virtual Learning: A Revolutionary Approach to Building a Highly Skilled Workforce.* McGraw-Hill, 1997.

Roger C. Schank, Janet L. Kolodner, and Gerald DeJong. Conceptual information retrieval. In *SIGIR*, pages 94–116, 1980.

H. Schärfe. Narrative ontologies. In C. G. Cao and Y. F. Sui, editors, *In Proceedings of Knowledge Economy Meets Science and Technology*, pages 19–26. Tsinghua University Press (KEST2004), 2004.

Abigail J. Sellen, Andrew Fogg, Mike Aitken, Steve Hodges, Carsten Rother, and Ken Wood. Do life-logging technologies support memory for the past?: an experimental study using sensecam. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 81–90, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9.

P. Sengers. Designing comprehensive agents. In *Proceedings of 16th International Joint Conference on Artificial Intelligence*, volume 2, pages 1227–1232, 1999.

N. Shadbolt and K. O'Hara. AKTuality: An overview of the aims, ambitions and assumptions of the advanced knowledge technologies interdisciplinary research collaboration. *AKT Selected Papers 03*, pages 1–11, 2003.

N. R. Shadbolt, F. Ciravegna, J. Domingue, W. Hall, E. Motta, K. O'Hara, D. Robertson, D. Sleeman, A. Tate, and Y. Wilks. Advanced knowledge technologies at the midterm: Tools and methods for the semantic web. *AKT Selected Papers 04*, pages 1–60, 2004.

Nigel R. Shadbolt, Wendy Hall, and Tim Berners-Lee. The Semantic Web: Revisted. *IEEE-Intelligent System*, 21(3):96–101, May 2006.

B. Shneiderman. Designing information-abundant web sites: Issues and recommendations, 1997.

Victor Shoup, editor. *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, 2005. Springer. ISBN 3-540-28114-2.

Ronny Siebes, Dave Dupplaw, Spyros Kotoulas, Adrian Perreau de Pinninck, Frank van Harmelen, and David Robertson. The OpenKnowledge system: An interaction-centered approach to knowledge sharing. In *OTM Conferences (1)*, pages 381–390, 2007.

N. De Silva and P. Henderson. Narrative support for technical documents: Formalising rhetorical structure theory. In *In Proceedings of Enterprise Information Systems (In Press)*. ICEIS, 2005.

A. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(12):1349–1380, 2000.

Ashley Smith. Who controls the past controls the future - life annotation in principle and practice. Ashley Smith's PhD thesis at the University of Southampton, July 2008.

Gabriella Spinelli. A multimedia system to collect, manage and search in narrative productions. In Hans-Jörg Bullinger and Jürgen Ziegler, editors, *HCI (1)*, pages 451–455. Lawrence Erlbaum, 1999. ISBN 0-8058-3391-9.

Kilian Stoffel, Merwyn G. Taylor, and James A. Hendler. Efficient management of very large ontologies. In *AAAI/IAAI*, pages 442–447, 1997.

Michael Streatfield. Report on Summer Internship Work For the AKT Project: Benchmarking RDF Triplestores. Technical report, University of Southampton, 2005.

R. Stuber, V. R. Benjamins, and D. Fensel. Knowledge engineering: Principles and methods. *Data and Knowledge Engineering*, 25:161–197, 1998.

N. Szilas. A new approach to interactive drama: From intelligent characters to an intelligent virtual author. In *AAAI Spring Symposium on AI and Interactive Entertainment Tech. Rep. SS-01-02*, pages 72–76, 2001.

N. Szilas. Structural models for interactive drama. In *The 2nd International Conference on Computational Semiotics for Games and New Media*, Augsburg, Germany, 2002.

Martin Szomszor, Harith Alani, Ivan Cantador, Kieron O'Hara, and Nigel R. Shadbolt. Semantic modelling of user interests based on cross-folksonomy analysis. In *International Semantic Web Conference 2008*, 2008.

S. Toulmin, R. Rieke, and A. Janik. *Introduction to Reasoning*. MacMillan Publishing Company, New York, 2nd edition, 1984.

M. M. Tuffield, D. E. Millard, and N R Shadbolt. Towards the narrative annotation of personal information and gaming environments. In *Proceedings of 1st Workshop on Narrative, Musical, Cinematic and Gaming Hyperstructure ACM Hypertext 05*, Salzburg, Austria, September 2005a. ACM Hypertext 2005.

M. M. Tuffield, N. R. Shadbolt, and D. E. Millard. Narrative as a form of knowledge transfer: Narrative theory and semantics. In *Proceedings of the 1st AKT Doctoral Colloquium*. AKT, June 2005b.

Mischa M. Tuffield, Stephen Harris, David P. Dupplaw, Ajay Chakravarthy, Christopher Brewster, Nicholas Gibbins, Kieron O'Hara, Fabio Ciravegna, Derek Sleeman, Nigel R. Shadbolt, and Yorick Wilks. Image annotation with Photocopain. In *Proceedings of the First International Workshop on Semantic Web Annotations for Multimedia (SWAMM)*, Edinburgh, May 2006a. held as part of 15th World Wide Web Conference (22-26 May, 2006).

Mischa M. Tuffield, Antonis Loizou, David P. Dupplaw, Sri Dasmahapatra, Paul H. Lewis, David E. Millard, and Nigel R. Shadbolt. Semantic Logger: Supporting service building from personal context. In *Proceedings of Capture, Archival and Retrieval of Personal Experiences (CARPE) Workshop*. ACM MultiMedia, October 2006b.

Endel Tulving. Episodic memory: from mind to brain. *Annual review of psychology*, 53: 1–25, 2002.

A. Vailaya, A. Jain, and H. Zhang. On image classication: City images vs. landscapes. In *Pattern Recognition.*, page 19211935, 1998.

G van Heist, T Schreiber, and B Wielinga. Using explicit ontologies in KBS. *International Journal of Human-Computer Studies*, 46:183–292, 1997.

R. Veltkamp and M. Tanase. Content-based image retrieval systems: A survey, 2000.

James Ze Wang, Jia Li, and Gio Wiederhold. SIMPLIcity: Semantics-sensitive integrated matching for picture LIbraries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(9):947–963, 2001.

Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding collisions in the full SHA-1. In Shoup (2005), pages 17–36. ISBN 3-540-28114-2.

Xiaoyun Wang, Hongbo Yu, and Yiqun Lisa Yin. Efficient collision search attacks on SHA-0. In Shoup (2005), pages 1–16. ISBN 3-540-28114-2.

Edgar Weippl, Alexander Schatten, Shuaib Karim, and A. Min Tjoa. SemanticLIFE collaboration: Security requirements and solutions - security aspects of semantic knowledge management. In *PAKM*, pages 365–377, 2004.

R. Wilensky. Pam. In R. Schank and C. Riesbeck, editors, *Inside Computer Understanding: Five Programs Plus Miniatures*, New Jersey, 1981.

Yorick Wilks. Artifical companions as a new kind of interface to the future internet. Oxford Internet Institute Research Report 13, 2006.

World Wide Web Consortium. SPARQL query language for RDF, working draft. Technical report, World Wide Web Consortium, 2005.

H. Xiao and I. F. Cruz. A multi-ontology approach for personal information management. In *Proc. of Semantic Desktop Workshop at the ISWC*, 2005.

Yinglian Xie, Fang Yu, and Martin Abadi. De-anonymizing the internet using unreliable ids. *SIGCOMM Comput. Commun. Rev.*, 39(4):75–86, 2009. ISSN 0146-4833.

Ka Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. Faceted metadata for image search and browsing. In *Proceedings of the conference on Human factors in computing systems*, pages 401–408. ACM Press, 2003. ISBN 1-58113-630-7.