UNIVERSITY OF SOUTHAMPTON
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science
Intelligence, Agents, Multimedia Group

A progress report submitted for continuation towards a PhD

Supervisor: Professor Nigel R. Shadbolt
2nd Supervisor: Dr David E. Millard
Examiner: Dr Kieron O'Hara

Narrative as a Form of Knowledge Transfer, Narrative Theory and Semantics: Present Challenges - Future Possibilities

by Mischa M. Tuffield

August 8, 2005

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

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

A progress report submitted for continuation towards a PhD

by Mischa M. Tuffield

An increase in digital multimedia data, has presented computer scientists with the key problem of large-scale information management. The Semantic Web vision proposes the publishing of information as semanticly accessible resources. Current methods of querying and browsing such rich knowledge bases require an understanding of the domain at hand that can not be asked of end-users. This progress report aims to present how Narratives, that are core to our method of conceptualising our environment, can be used to present targeted knowledge to an end-user. Ontology definitions are promoted as the fundamental building blocks for annotating these multimedia resources and are presented as key to this research. The main contribution of this paper is the identification of the shortcomings of existing narrative generation systems, followed by the identification of techniques that could be adopted to overcome these deficiencies.

Contents

Acknowledgements

I would like the thank my supervisors Nigel Shadbolt and David Millard, for the irreplacable guidance they have given me over this last year. I would also like to thank all the AKTors, and my friends in the IAM group for making me feel at home in Southampton.

This work is 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.

Chapter 1

Introduction

1.1 Information Overload

A key problem faced by the information management community is that of information overload, or *infosmog* (Shadbolt and O'Hara, 2003). The ever growing, evolving nature of the World Wide Web (Web) is not making information retrieval any easier. The web is populated with an unmanageable amount of heterogeneous data, in forms as diverse as image, sound, video, and text. 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 brought about this rich *corpus* of digital media.

The Web facilitates searching 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 and infer relationships from this vast pool of information, and present it in a structured and targeted way.

1.2 The Semantic Web

The Semantic Web (SW) vision (Berners-Lee, 1999) and technologies (Shadbolt et al., 2004) are challenging 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[1] that are linked together in a structured and meaningful manner for machine communication. This new paradigm of publishing would allow for context-based multimedia to be exploited in aid of generating dynamic narratives from the available knowledge elements in a manner best suited to the user's profile (Bilasco et al., 2005).

The Semantic Web defines the necessary relational models for describing web resources with context independent standards, such as the Resource Description Framework (Manola and Miller, 2004) (RDF). RDF is currently used to annotate vast amounts of heterogeneous data, and is defined to be applicable to the description of multimedia artefacts.

Advances in the techniques and tools for the semantic enrichment of information have allowed for relationships in data to be made explicit in knowledge rich domains. One of the cornerstones of developing SW applications is the development of ontologies. Advances in ontological driven multimedia presentation assembly from semantically enabled knowledge elements have been reported in recent works (Geurts et al., 2003)(Little et al., 2002)(Bilasco et al., 2005). This assembly of structured multimedia presentations 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". Summarisation (Alani et al., 2003), document generation systems (Silva and Henderson, 2005)(Domingue et al., 2001), and systems that aid biography tailoring (Kim et al., 2002) are just a few systems that have benefited from the proposed method of semantic enrichment.

The metadata produced by annotating multimedia with respect to a shared ontological vocabulary will allow for searching and navigation by concept (Jewell et al., 2005), 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 heterogeneous media. The CIDOC CRM was created as a "semantic approach to integrated access" (Doerr et al., 2003) for cultural heritage data, and will be discussed in section 3.2.

1.3 Narratives

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 the basis of transfering knowledge within societies (Campbell, 1949). Work has been undertaken to illustrate the transfer of knowledge within societies through the

[1] The term multimedia items is used to refer to text documents, videos, audio, images, etc, ...

use of narrative mediums and the similarities between these differing modes (Campbell, 1949)(Murray, 1998)(Bal, 1997). Research into one of the proposed ways that we make sense of our world has resulted in the term *Narrative Intelligence* (NI) being coined. This term was based on research suggesting 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 time when the aforementioned NI research was undertaken SW technologies were still in their infancy, our research aim to presents how advances in knowledge representation and the enrichment of information through explicit semantics can be utilised to develop meaningful narratives, that we are all so familiar with.

1.4 Report Outline

This progress report discusses how this new paradigm of publishing information might be utilised to allow for narratives to be generated dynamically from available annotated knowledge items. The structure of this report is as follows; following this introduction, in chapter 2 the problem at hand is elaborated to provide an insight into the motivating factors behind this research. Next, chapter 3 will present the initial findings and the work completed. Chapter 4 will propose areas of future work and present the conclusions reached.

Chapter 2

The Problem

2.1 Narrative Theories and Knowlege Engineering

Narratives and the study of narratology, have been an area of interest to the social sciences for a long time and have recently become increasingly popular in the field of knowledge technologies (Alani et al., 2003)(Geurts et al., 2003)(Little et al., 2002)(Silva and Henderson, 2005)(Gemmell et al., 2003) and knowledge management (Connell et al., 2005). Narratology has focused on representing and reproducing one of the core modes of human communication. As a result a growing set of different narrative theories, ways of conceptualising narrative spaces, such as formalism, structuralism, post-structuralism, and post-modernity (Schärfe, 2004) have been developed. The aim of this research does not involve developing a new theory or even proposing a taxonomy of narrative approaches, this is best left to narrative theorists, but learning from and harnessing existing methods to aid narrative generation from knowledge rich environments.

There are a growing number of areas where strategies and viewpoints put forward by narratology have been adopted and utilised within computer science research, such as dynamic multimedia presentations (Kim et al., 2002)(Geurts et al., 2003)(Little et al., 2002)Rutledge et al. (2003), summarisation (Alani et al., 2003), and storytelling and interactive drama (Murray, 1998).

One example of an adopted theory is Bal's (Bal, 1997) view of narrative (Kim et al., 2002). 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.

2.2 The Analogy: Semantic Web meets Narratology

SW enabling technologies allow for the annotation of multimedia items, resulting in a corpus of available knowledge nuggets (see section 1.2).

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

Given the above analogy, the major shortfall of existing narrative generation systems, for dynamic presentations, is the use of story templates as grammars. The shortcomings of existing techniques are presented in tandem with this project's initial findings in section 3.1.1. 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 into a *narrative*.

2.3 Interfacing with the Semantic Web

The possibility of generating naratives is being explored as an alternative to the current methods of human computer interaction with semantic knowledge bases. Two of the most influential methods of interacting with SW enabled data-sets are facetted browsers, like mSpace (m. c. schraefel et al., 2003), and knowledge data navigation techniques presented by systems like Haystack (Huynh et al., 2002). Both of these methods require an understanding of the underlying data-structure, and a grasp of "categorisation by concept" that can not be expected from all end-users. Some proposed advantages for wanting to have narratives generated from semantic data-sets are presented, within the context of an application domain, in section 3.1.1.

2.4 The Research Problem

The advances made in knowledge representation through the development of SW enabling technologies are allowing for more knowledge about a given media item, in the form of metadata, to be made explicit. Given the motivation to have information rendered in the form of a targeted narrative and the knowledge rich nature of semanticly annotated knowledge bases, research is sought to confirm whether this metadata could be utilised to convey meaningful knowledge to an end-user. This envisioned method of

communicating information stored in a knowledge base to human user has been refered to as "knowledge transfer" (Tuffield et al., 2005).

Firstly, my research aims to provide insight into what knowledge is required to be made explicit to facilitate knowledge transfer. The nature of the annotation will obviously be determined by the problem at hand, one such application domain is detailed in section 3.1. Insight into the form and content of the desired metadata will be identified upon completion of the proposed future work (4.1).

Secondly, an indication into how semanticly annotated, ontology driven, multimedia can be reasoned upon, to generate narratives. Methods that take advantage of the relationships that exist within a knowledge base will be key to the development of such a service. This endeavour coupled with the identification of instances of narrative theory, will provide insights into the reasoning processes needed for automatic construction of narratives.

Following on from this, my research aims to present how useful narratives boad as a mode of knowledge transfer from semanticly enabled knowledge bases.

Chapter 3

Work Completed

3.1 Memories for Life

Memories for Life[1] (M4L) is being discussed as a grand challenge for UK computing. It aims to address the challenges of storing and presenting 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 inheritantly 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[2] ontology (Jewell et al., 2005) is proposed as a possible candidate (see section 3.2). This ontology has been designed to allow the mark-up of literature, film, and other forms of narrative at the *fabula* level, detailing the chronological events inside a piece of media.

The Advanced Knowledge Technologies (AKT)[3] project has produced tools to automate the annotation of textual media, (e.g. Armadillo (Ciravegna et al., 2004)), and is currently looking at methods of semiautomating photograph annotation (Harris et al., 2005) (see section 4.1.1). These tools and techniques are being applied to extract knowledge from as many of the readily available sources of accessible metadata to create as rich a memory bank as possible while keeping the cognitive overhead to a bare minimum.

My research aims to produce a system that will dynamically generate a narrative given a query from a collection of annotated memories.

[1] Memories for Life website: www.memoriesforlife.org, describes the ambitions and directions of the network

[2] Project website: http://ontomedia.ecs.soton.ac.uk

[3] Project website: www.aktors.org

3.1.1 Narratives in the form of Dynamic Presentations

Given the start of work on the Le PhotoCOPAIN system (Harris et al., 2005) (see section 4.1.1) methods of generating autiobiographical narratives from a personal collection of photos are currently being investigated.

A desired narrative could be a collection of photos taken on holiday of landscapes in foreign countries, or a collection of conference photos taken in an academic year. Assuming a story is the presentation of a succession of events related to a given topic, this work aims to generate such narratives based on the resources available to it, namely the photos and any relevant information extracted from available sources like emails, and iCal files (Harris et al., 2005).

In M4L a two-layered approach is being adopted in order to allow maximum automation and full flexibility. The first layer will be an automatically generated time line, which will depend entirely on the date stamping of the images and the media available. Human intervention at this level is limited to altering the granularity of the time segments should the user feel this necessary. Automatically integrated at this level are the location annotations, derived from the absolute positioning via GPS, or coarse positions through GSM phone information (Harris et al., 2005), and any clues to the user's location presented from their iCal file. The combination of location and chronology provide useful contextual cues for generating stories to help refresh the users memory (Gemmell et al., 2005).

The second layer will aim to organise photos into a "narrative" structure based on any other metadata available. Another core research issue is to determine the appropriate balance between user intervention and automatically generated narrative structures, this will be tackled when appropriate. In essence, the annotated images will function as the instances in the underlying, the *fabula*, which will allow for sequential organisation of the instance events into a cognitively coherent narrative. This notion of dynamically generating narratives from concepts other than location and time, is identified as a research problem that has yet to be tackled and is the focus of this research. The aim is to enhance the static story-template based method of generating narratives (see section 3.1.1), by employing existing narrative techniques, and learning from the work presented in generating narratives from leverages such as time and location (Gemmell et al., 2005).

A number of major projects are working towards similar visions, MyLifeBits (Gemmell et al., 2003), and SemanticLIFE (Admed et al., 2004) are the two most noteworthy. The aim of this research differs from that of the above projects due to the novel manner in which the proposes M4L system will strive to accumulate as much low-cost metadata as possible. The methods that the SemanticLIFE project uses to extract metadata from the personal communications, i.e. emails and the surreptious analysis of web browsing, and the related issues surrounding trust are presented in Weippl *et al* (Weippl et al., 2004),

this research presents insight into methods of accumulating personal memory orientated metadata. It is important to stress that their system does not attempt to generate a narrative from the harvested knowledge. The story telling mechanisms in MyLifeBits (Gemmell et al., 2005) presents valuable insight on how stories can be generated based on leverages such as location and time, but they fail to generate narratives based on any further concepts, like the examples above (3.1.1).

The cultural heritage domain has also produced systems that generate narratives from annotated media. The ArtEquAKT system (Alani et al., 2003) and Geurts *et al*'s system (Geurts et al., 2003) both generate biographical narratives about artists, and Story Fountain (Mulholland et al., 2004) does so from the historical records of Bletchley Park. It is important to mention that none of the systems listed above do any form of dynamic narrative construction. Instead, they all populate story-templates.

Research in being undertaken into ways of extending the typical static template-based approach of narrative generation through adopting techniques used by comic book writers to tell stories through a succession of images (McCloud, 1994), and how the use of simple natural language could be used to strengthen the *story* presented by the rendered *narrative*. An example of one of McCloud's proposed panel-to-panel transitions could be, the presentation of a photograph taken at heathrow aiport on departure day followed by a picture taken on the beach, to tell the story of travelling to get on holiday. This transition takes the reader from subject-to-subject while staying within a scene or an idea, and is proposed as one of six different panel-to-panel transitions employed by comic book writers.

3.2 OntoMedia

Ontologies are paving the foundations to the realisation of the Semantic Web (SW) vision by capturing knowledge in a machine understandable language, such as the Web Ontology Language (OWL) (McGuinness and van 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.

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), (Jewell et al., 2005) presents the phenomenon that OntoMedia aims to model highlighting how it will be used to aid multimedia information retrieval.

There are currently many overlapping ontologies on the Semantic Web. A Swoogle[4] search for the term "character", you 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 refactored.

The remainder of this section will present an overview of the unique view that OntoMedia takes to the annotation of multimedia items with respect to existing models.

In (Jewell et al., 2005) we propose a vocabulary, in the form of an ontology (Noy and McGuiness, 2001) for the annotation of multimedia documents. 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 labelling.

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 3.1 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., 2005), and (Lawrence et al., 2005).

The annotation of multimedia into an ontological form has been investigated from a variety of angles by prior researchers, most significantly by Lagoze and Hunter's ABC Ontology (Lagoze and Hunter, 2001). This was designed primarily for the cataloguing community, with a focus on factual information representation, such as object provenance and rights management. ABC and OntoMedia share the separation of entity and temporal classes, whilst the OntoMedia ontology augments this core with classes that allow for narrative sequences to be annotated.

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

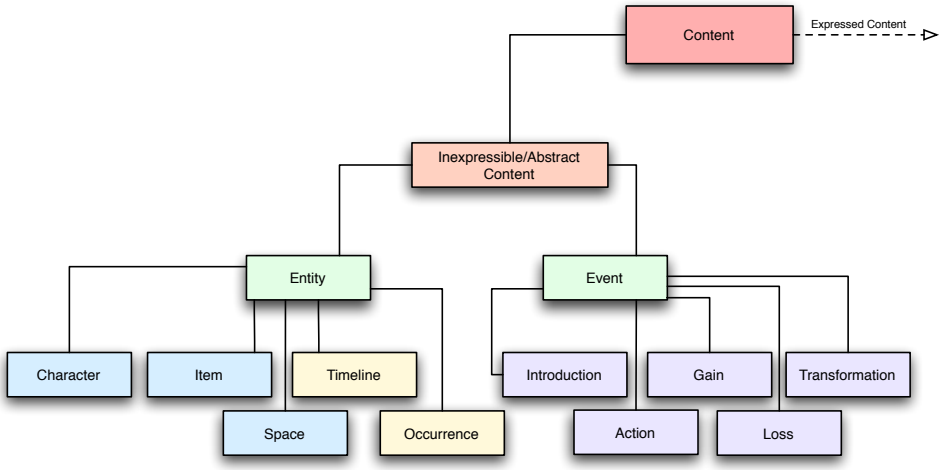

FIGURE 3.1: The structure of the OntoMedia ontology

The summarised scope of the CIDOC Conceptual Reference Model (CRM) is "the curated knowledge of museums" (Crofts et al., 2005). In their most recent documentation this is expanded to describe the intended scope of the CRM "as all information required for the exchange and integration of heterogeneous scientific documentation of museum collections". The CRM works with an entity & event model that in many ways resembles the OntoMedia ontology. Despite their initial similarities the CIDOC CRM is still primarily concerned with the physical object equivalent of bibliographic data. It can describe in detail the condition, provenance, and attributes of an object, such as a red-figured vase. As part of this description it will include information on the decorations, without describing the depictions as a narrative sequence. A more detailed discussion of the salient differences between OntoMedia and other relevant existing ontologies is presented in (Lawrence et al., 2005).

Chapter 4

Future Work and Conclusions

4.1 Future Work

My research aims to produce a generic system that given a domain ontology, and a proposed narrative paradigm will be able to assemble any relevant semanticly accessible resources into a targered narrative. The remainder of this section will identify three application domains, that we believe are sufficiently similiar in nature to fall under the same line of research. The desired outcome of my research will be the identification of the required metadata and the reasoning proposes needed to facilitate the ability to generate narratives from large collections of heterogeneous data.

4.1.1 Le PhotoCOPAIN

With regards to the M4L research agenda, the development of the Le PhotoCOPAIN has been proposed as a first step in reaching this goal. The setting up of a triple store (Harris and Gibbins, 2003) for data annotated using OntoMedia, in order to get acquainted with another underlying SW technology, has been identified as the first step towards implementing the proposed system. A training set of annotated photos has recently been developed, and is about to be subject of some data-mining. A selection of different machine learning algorithms, pca, neural networds, baynesian networks, linear parameter models, etc, will be invoked on the data set, in order to find the best possible clustering algorithm. Once the best algorithm is found, this data set set will be used to suggest annotations to new photographs present to the classifier.

Following the completion of the data-set, a first attempt at using Scott McClouds' (McCloud, 1994) theory on how to assemble images to produce narratives, will be implemented. This will be integrated into a system that can generate narratives, in the form of a dynamic presentation, based on time, location, and other concepts embedded in the annotations. This is elaborated in the postion paper (Harris et al., 2005).

4.1.2 eChase

*e*Chase or the electronic Cultural Heritage made Accessible for Sustainable Exploitation[1] is the follow up project to Sculpteur. After some inital meetings with the *e*Chase team about the use of OntoMedia within their project, news that they are creating a web-service to map metadata from the CiDOC CRM to OntoMedia has come as good news. Enquiries into the availability of annotated multimedia for use as test data on the proposed narrative generation system will be followed by the proposal of a narrative generation system for *e*Chase.

4.1.3 Narratives in Computer Games

Over the course of the last few decades we have seen an abundance of academic work in computer game worlds that have attempted to combine narratives with interactivity (Murray, 1998)(Matheas, 1997)(Matheas and Stern, 2003). Research in this task has produced valuable insight into the need for identifying the middle ground between interactive freedom, and coherent dramatic narratives. Systems have taken a number of different approaches with the aim of finding this elusive middle ground. Some have taken agent-based approaches, where complex autonomous believable agents (Matheas, 1997) are deployed in simulated environments, such as the Oz project (Project, 1989-2002), limitations of which include the lack of focus on the user's actions and the systems inability to scale. Others have taken a more global approach, inspired by AI planning techniques.

These approaches tend to focus more on the users' actions, their shortfalls come from the need for the stories to be carefully planned, to the degree that the game must ensure the chances of interrupting these plans are kept to a bare minimum. Façade (Matheas and Stern, 2003) like the Oz project aims to find the middle-ground between interactivity and a coherent story, by incorporating more explicit knowledge in terms of the overall narrative structure and keeping track of the human-players juxtaposition in the narrative arc. One of the original aims of Façade was to empower the player to have complete control over the story. However, the final result still only allowed for players to be "interactive observers" rather than protagonists (Matheas and Stern, 2003).

4.1.3.1 Massively Multiplayer Online Role Playing Games

This section aims to present how similar methods of narrative generation, as described for M4L, could be extended to MMORPGs to "stimulate" dramatic narratives. As of July 2005, the game World of Warcraft[2] has sold more than 3.5 million copies world wide,

[1] Project website: www.echase.org
[2] Official website: www.worldofwarcraft.com

1.5 million of these being sold since the Chinese launch on 7th July 05. The publishing house, Blizzard[3] states that there are over half a million subscribers on online at any given time.

This uptake of MMORPGs, were humans take up the *personna* of an imaginary character depicted by an avatar in a massive online virtual world, has been tremendous. Game play in MMORPGs are not dissimilar in form to that of computer based role playing games commonly associated with narrative research. A player takes on a character and subsequently "role plays", in order to gain experience that in turn makes their characters more powerful, hence achieve status in the online community. Players are not following any narratives in so far as wanting to achieve their local goal of completing a given dungeon or finding a certain artifact, as opposed to traditional adventure games, were a player undertakes tasks in order to progress through a predefined story.

Given that MMORPGs are closed worlds, every entity, whether human player, item, or location, can be accounted for, and their relationship could be made explicit given the rules of games. These worlds are presented as knowledge rich domains, containing information about characters, their activities, and their belongings; all of which can be described with an ontology like OntoMedia. An ontological description of all of the character classes and races, items could also be developed. Given the server based nature of these games, a snap-shot of the state of the world could be taken at any given time, presenting the environment as an *accessible* one.

The method envisaged to stimulate dramatic narratives in MMORPG, is by generating a narrative based on a snap-shot of the online world every time a user logs in to the game. The proposed form of the narrative would work on the following three levels. Firstly, a high level narrative is proposed, describing what the most powerful characters are doing, what relics they have accumulated, what lands they have conquered, etc. This is intended to promote a desire to play more, for all avid gamers want to be recognised as elite. The second level, consists of a narrative describing the major events that have taken place within a players vicinity, promoting local actions that may result in immediate reward. And finally another narrative based on the characters vicinity, describing characters that oppose the player, that need to be eliminated, and possible accomplices. This "newsletter" is proposed as a method of instigating actions that would result in dramatic events occurring in a MMORPG, effectively emergent narratives.

Unfortunately stories about real life seldom fit into the grand narratives of myth or fairytale and simple heuristics working on these semantics, such as chronology or human relationships, would not produce satisfying narratives. Work is currently being undertaken to adapt Bremond's (Bremond, 1974) extension of Propp's functions (Propp, 1968) into a manner that could be used to generate suitable narratives in MMORPGs. Other descriptions of story structures are also being considered (Colby, 1973)(Tolkien, 1964).

[3]Official website: www.blizzard.com

4.2 Conclusions

This progress report aims to motivate narratives as a natural way of conveying knowledge to people. My future research will present whether semanticly annotated multimedia and the relationships explicit in the metadata could be harnessed to generate narrratives. This paper has identified the shortcomings of existing systems, and will be followed by the development in the proposed domains (4.1) in order to confirm whether the employment of ontology driven metadata and reasoners could demonstrate the expressive power of the Semantic Web infrastructure.

Bibliography

M. Addis, P. 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.

H. Alani, S. Kim, D. E. Millard, M. J. Weal, P. H. Lewis, W. Hall, and N. R. Shadbolt. Automatic ontology-based knowledge extraction from web documents. *IEE Intelligent Systems*, 18(1):14–21, January 2003.

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

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

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.

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.

C. Bremond. *Logique du récit.* Seuil, Paris, 1974.

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

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.

B. Colby. A partial grammar of eskimo folktales. *American Anthropologist*, 75:645–662, 1973.

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

Nick Crofts, Martin Doerr, Tony Gill, Stephen Stead, and Matthew Stiff (eds). Definition of the cidoc crm conceptual reference model. Reference document, International Council of Museums, March 2005.

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

J. Domingue, E. Motta, S. Buckingham Shum, and S. Vargas-Vera. Ontology-driven document enrichment within communities of practice. In *In Proceedings of 1st Int. Conference on Knowledge Capture*, pages 30–37. International Conference on Knowledge Capture, 2001.

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.

S. Harris and N. Gibbins. 3store: Efficient bulk rdf storage. In *In Proceedings of 1st International Workshop on Practical and Scalable Semantic Systems*, pages 1–15. (PSSS'03), 2003.

S. Harris, K. O'Hara, M. Tuffield, C. Brewster, N. Gibbins, W. Hall, N. Shadbolt, I. Simpson, D. Sleeman, and Y. Wilks. Le photocopain: A low cognitive overhead photo annotation system. In *(In Press)*. SIGMM CARPE Workshop, 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.*, 2002.

M O Jewell, Faith 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 2005.

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.

C. Lagoze and J. Hunter. The ABC ontology and model. *Journal of Digital Information*, 2(2), November 2001.

Faith 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 *(In Press)Ontology Patterns for the Semantic Web Workshop at ISWC05*. ISWC, November 2005.

S. Little, J. Geurts, and J. Hunter. Dynamic generation of intelligent multimedia presentations through semantic inferencing. In *In Proceedings of 6th European Conf on Research and Advanced Techniques for Digital Libraries*, pages 159–189. Springer - ECDL, 2002.

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.

M. Mateas and P. Sengers. Narrative intelligence. In *In Proceedings of AAAI Fall Symposium on Narrative Intelligence*. AAAI, 1999.

M. Matheas. An oz-centric review of interactive drama and believable agents. Technical report, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1997.

M. Matheas and A Stern. Facade, an experiment in building fully realised interactive drama. In *Games Design Track*. Games Developers Conference, 2003.

S. McCloud. *Understanding Comics, The Invisible Art.* HarperPerennial, NY:USA, 1994.

D. L. McGuinness and F. van Harmelen. Owl web ontology language overview: W3c recommendation. http://www.w3.org/TR/owl-features/, 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.

P. Mulholland, T. Collins, and Z Zdrahal. Story fountain: intelligent support for story research and exploration. In *In Proceedings of 9th international conference on Intelligence User Interfaces*, pages 62–69. International Conference on Intelligent User Interfaces, 2004.

J. Murray, editor. *Hamlet on the Holodeck*. MIT University Press, Cambrigde:Mass, 1998.

N. Noy and D. McGuiness. Ontology development 101: A guide to creating your first ontology. Technical report, Stanford University, 2001.

Oz Project. Carnegie mellon university. http://www.cs.cmu.edu/afs/cs.cmu.edu/project-/oz/web/oz.html, 1989-2002.

V. Propp, editor. *Morphology of a Folktale*. University of Texas Press, Texas:USA, 2nd edition, 1968. Translated by L. Scott and L. Wagner.

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.

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.

N. 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.

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. 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.

R. Stuber, V. R. Benjamins, and D. Fensel. Knowledge engineering: Principles and methods. *Data and Knowledge Engineering*, 25:161–197, 1998.

J. R. R. Tolkien. *Tree and Leaf*. George Allen & Unwin Ltd, London, UK, 1964.

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 2005.

G van Heist, T Schreiber, and B Wielinga. Using explicit ontologies in kbs. *International Journal of Human-Computer Studies*, 46:183–292, 1997.

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.