

Ontological Approaches to Modelling Narrative

Mischa M Tuffield[†], David E Millard[†], Nigel R Shadbolt[†]

[†]Intelligence, Agents, Multimedia Group
University of Southampton, Southampton, UK
{mmt04r, dem, nrs}@ecs.soton.ac.uk

Abstract. We outline a simple taxonomy of approaches to modelling narrative, explain how these might be realised ontologically, and describe our continuing work to apply these techniques to the problem of Memories for Life.

1 Introduction

Narratives have long been considered a primary way in which human beings communicate with one another [1]. 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 [2]. Narrative has also been identified as a central part of how we learn to approach the world [3, 4].

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 Davies 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 [5]. NI has been identified as one of the main synergies around which Artificial Intelligence (AI) research into narrative has been brought together [6].

The Semantic Web (SW) vision, and resulting technologies, 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 (RDF), but it is the use of ontologies that forms the cornerstone of SW interoperability.

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. Influential techniques of interacting with SW enabled data, such as mSpace [7], Haystack [8] and work undertaken at CWI [9] are attempting to alleviate the burden of requiring an understanding of the underlying data-model.

This tension between semantics as expressed by people and as expressed by machine has been described as the Semantic Gap [10]. 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.

The high level of formalism required for knowledge to be shared between machines poses problems whenever human intervention is needed. The two key stages of human intervention are authoring and viewing. There is a growing trend in systems towards a “translation to narrative” approach [10]. This is when a system’s knowledge base is encoded formally (e.g. in an ontology) and then converted to a multimedia presentation for display. This treats narratives as an interface, not a genuine way of structuring information, and accommodates end-users that do not wish to be exposed to the formal notations underlying such a system.

For example, the ArtEquAKT system [11] assembles biographies of artists from gathering information from web accessible resources. This system extracts information from unstructured narratives, the World Wide Web, populates an ontology and subsequently combines the extracted knowledge into an adaptive hypermedia document, or what we call a narrative.

But even to do this interface level work 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 [11, 12, 9, 13], and knowledge management.

We believe 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 paper is an attempt to move toward this ontological understanding.

2 A Simple Taxonomy of Narrative Generation

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 modeling the intricacies of complex autonomous beings, stimulating 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 modeling human characteristics in an agent’s reasoning system is not a tractable task. Reports suggest that systems endeavoring to generate emergent narratives often result in unexciting and undirected stories [14], 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.

Plot Based: Are methodologies that attempt narrative generation/understanding by adapting and proceduralising narrative theories. These systems are built around explicit narrative structures on plot and sequential structure. Rule-based methods are the most common way of representing such knowledge [15]. A number of different narrative structures have been implemented computationally to aid narrative generation. These include rhetorical structure theory used to generate video documentaries [16] and to aid technical writing. Propp’s functions [17] a procedural formalisation of fairytales has also been used to steer interactive fiction along a consistent and dramatic storyline. This approach is a knowledge intensive one, and has been referred to by Szilas [15] 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 perception of the unfolding narrative, or “user-profile”. These systems usually incorporate an explicit narrative model and utilise it along side any available knowledge from a “user-model” to set the context and drive the outcome towards a targeted narrative [18, 13].

3 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 [4] is a useful way in which to understand what is being modeled. This states that narrative can be viewed as consisting of three layers. 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 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 we have already identified 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, must also have an understanding of the fabula, in order to make sense of its knowledge base.

3.1 Ontological Models of the Fabula

Semantic Web technologies allow for the annotation of multimedia items, resulting in a corpus of available knowledge nuggets. If these annotations describe the entities and events within the multimedia items, then they can be seen as a type of fabula.

As an example, we 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 repurposed and assembled to present new stories, similar to the Bocconi's work on the generation of video documentaries [16].

Ultimately we might 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. 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.

3.2 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 [19], it is this structural knowledge that needs to be modeled [20]. In order for such models to work it is necessary that story threads can be found in the fabula. The likelihood of finding and threading available semantic relationships within a knowledge base will be dependent on the quality of annotations. For example, 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 [11]. 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 [21] has been implemented to generate rhetorical argumentation for video documentaries [16].

However, there are problems with using templates in place of a more flexible grammar. Static template structures have to be defined by developers before the deployment of a system. This limits 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 [20], allowing for systems to be less domain specific.

3.3 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 there 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, we expect different presentation methods in a documentary than we do in a drama). Scott McCloud's [22] 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 [1].

4 Future Work: Modeling the Higher Layers

OntoMedia [23] has been proposed as the first step in modelling narrative as presented in this paper. OntoMedia has been created to model narrative by focusing on Fabula level events. We are currently applying the OntoMedia ontology to the Memories for Life ¹ research challenge, in the development of the Le PhotoCOPAIN system. The objective here is to use photograph analysis techniques to extract information from sets of photographs that will be annotated using OntoMedia. The assumption is that each photo is owned, time stamped and positioned at the time it is taken using GPS (a reasonable assumption considering the current trend towards convergent pervasive personal devices). Techniques are also being investigated to extract other forms of metadata from accessible resources such as iCal files, weather reports and news feeds from the web. The task involves the identification of as much readily available metadata that can be extracted from the web in order to generate narratives as rich as possible. This research will provide insight into what kind of metadata, and to what granularity, are necessary to generate meaningful narratives.

A training set of annotated photos has recently been developed, and a set of data-mining tools have been developed to the analysis. A selection of different machine learning algorithms, principle component analysis (PCA), neural networks, bayesian networks, linear parameter models, etc, will be invoked on the

¹ www.memoriesforlife.org

data set, in order to find the best possible clustering algorithm. Once the best algorithm is found, this data set will be used to suggest annotations for any new photographs present to the classifier.

We are also interested in developing higher level ontologies that will help to structure these photographs into stories. When designing an ontology that is scoped over a long-period time, it is important to consider issues of ontology evolution and version tracking. Issues of how story ontologies may change over a lifetime are also central to this research effort. Consider personal photo albums, these tend to focus on identifiable life events such as weddings or the arrival of a new baby.

There have been a number of attempts to deal with Bal's story level. Most of these have used templates, however templates hide a multitude of complex decisions regarding the arrangement of content, given the context of a particular genre. Since templates are rigid structures they also adapt badly to changes in the content being arranged. Work is currently being undertaken [20] into the identification of genres and their salient features, in order to help design grammars for template generation.

An alternative strategy is to use grammar's; rules that determine how content can be arranged. Grammar's are flexible and extensible but extremely difficult to encode accurately, and there is no assurance that the resulting structure will be sensible.

There are a number of structures that grammar's could be built around, including classic work such as Propp's functions [17] or Campbells observations on the Hero with a Thousand Faces [2], however such mythic structures will not always translate to everyday situations. It is more likely that in scenarios such as M4L we will need everyday grammars for describing common situations such as day trips, holidays, weddings, etc.

Once the pictures have been arranged it will be necessary to arrange them for a user to view. At this point we shall consider ontological models of the narrative layer, and in particular existing theories on how juxtaposition and montage of images can be used for effect [22].

5 Conclusions and Future Work

In this paper we have argued 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. In the future work section we have identified a system, similar to ArtEquAKT insofar as they both take a "translation to narrative" approach, highlighting the need for accessible demonstrations of SW technologies.

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

We have grounded this 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.

In our own work we have explored the ontological modeling of the Fabula, and have 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.

6 Acknowledgments

This work is supported under the AKT IRC, which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N157-64/01. Mike Jewell and Faith Lawrence are also acknowledged for their contribution to the OntoMedia project.

References

1. Tuffield, M.M., Millard, D.E., Shadbolt, N.R.: 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, ACM Hypertext 2005 (2005)
2. Campbell, J.: *The Hero with a Thousand Faces*. Fontana Press, London:UK (1949)
3. Murray, J.: *Hamlet on the Holodeck*. MIT University Press, Cambridge:Mass (1998)
4. Bal, M.: *Introduction to the theory of Narrative*. 2nd edn. University of Toronto Press, Toronto:Canada (1997)
5. Blair, D., Mayer, T.: Tools for interactive virtual cinema. In Trappl, R., Petta, P., eds.: *Creating Personalities for Synthetic Actors: Towards Autonomous Personality Agents*, Springer (1997)
6. Mateas, M., Sengers, P.: An introduction to narrative intelligence. In: *Narrative Intelligence Symposium AAAI 1999 Fall Symposium Series, AAAI (1999)*
7. m. c. schraefel, Karam, M., Zhao, S.: mSpace: interaction design for user-determined, adaptable domain exploration in hypermedia. In Bra, P.D., ed.: *Proceedings of AH 2003: Workshop on Adaptive Hypermedia and Adaptive Web Based Systems.*, Nottingham, UK, Springer (2003) 217–235

8. Huynh, D., Karger, D., Quan, D.: Haystack: A platform for creating, organising, and visualising information using rdf. In: Semantic Web Workshop, The Eleventh World Wide Web Conference 2002. (2002)
9. Rutledge, L., van Ossenbruggen, J., Hardman, L.: Making RDF representable = global and local semantic web browsing. In: In Proceedings of the 14th International World Wide Web Conference (WWW05), Chiba, Japan (2005) 199–206
10. Millard, D.E., Gibbons, N.E., D. T. Michselides, M.J.W.: Mind the semantic gap. In: In Proceedings of ACM Hypertext 2005, SIGWEB (2005) 54–63
11. Alani, H., Kim, S., Millard, D.E., Weal, M.J., Lewis, P.H., Hall, W., Shadbolt, N.R.: Automatic ontology-based knowledge extraction from web documents. *IEE Intelligent Systems* **18** (2003) 14–21
12. Gemmell, J., Lueder, R., Bell, G.: The MyLifeBits lifetime store. In: Proceedings of ACM SIGMM 2003 Workshop on Experiential Telepresence (ETP 2003), Berkeley, CA (2003)
13. Bilasco, I.M., Gensel, J., Villanova-Oliver, M.: Stamp: A model for generating adaptable multimedia presentations. *Int. Journal Multimedia and Applications* **25** (2005) 361–375
14. Matheas, M., Stern, A.: Facade, an experiment in building fully realised interactive drama. In: Games Design Track, Games Developers Conference (2003)
15. Szilas, N.: Structural models for interactive drama. In: The 2nd International Conference on Computational Semiotics for Games and New Media, Augsburg, Germany (2002)
16. Bocconi, S., Nack, F., Hardman, L.: Using rhetorical annotations for generating video documentaries. In: Proceedings of the IEEE International Conference on Multimedia and Expo, (ICME) (2005)
17. Propp, V.: Morphology of a Folktale. 2nd edn. University of Texas Press, Texas:USA (1968) Translated by L. Scott and L. Wagner.
18. Bailey, P.: Searching for storiness: Story generation from a reader’s perspective. In: Narrative Intelligence Symposium AAAI 1999 Fall Symposium Series, AAAI (1999)
19. Shneiderman, B.: Designing information-abundant web sites: Issues and recommendations (1997)
20. Falkovych, K.: Creating a semantic based discourse model for hypermedia presentations: (un)discovered problems. In: Proceedings of 1st Workshop on Narrative, Musical, Cinematic and Gaming Hyperstructure ACM Hypertext 05, Salzburg, Austria, ACM Hypertext 2005 (2005)
21. Toulmin, S., Rieke, R., Janik, A.: Introduction to Reasoning. 2nd edn. MacMillan Publishing Company, New York (1984)
22. McCloud, S.: Understanding Comics, The Invisible Art. HarperPerennial, NY:USA (1994)
23. Jewell, M.O., Lawrence, F., Tuffield, M.M., Nixon, M.S., Prügel-Bennett, A., Shadbolt, N.R., Millard, D.E., m c schraefel: OntoMedia: An ontology for the representation of heterogeneous media. In: Proceedings of MultiMedia Information Retrieval (MMIR) Workshop at SIGIR, SIGIR (2005)