The Narrative Braid: A Model for Tackling The Narrative Paradox in Adaptive Documentaries

Charlie Hargood
Electronics and Computer Science
University of Southampton
Southampton, United Kingdom
cah07r@ecs.soton.ac.uk

Michael O. Jewell
Electronics and Computer Science
University of Southampton
Southampton, United Kingdom
moj@ecs.soton.ac.uk

David E. Millard
Electronics and Computer Science
University of Southampton
Southampton, United Kingdom
dem@ecs.soton.ac.uk

ABSTRACT

The Narrative Paradox is a theory that describes interaction and narrative cohesion as being in tension, and asserts that the structure of a narrative is disrupted by user adaptivity, leading to possible incoherence as the system accounts for interaction. We propose an approach that may reduce this disruption. Specifically, we propose to model a narrative as a collection of threads, woven together into the final discourse as a narrative braid. By separately maintaining logical coherence within a thread and thematic coherence between threads we believe it is possible to introduce interactivity while maintaining a strong narrative structure. We discuss how this may be applied to adaptive documentaries which, with a wide base of recorded material and diverse plot threads, provide a rich medium for initial experimentation in this area.

Categories and Subject Descriptors

H.1 [**Models and Principles**]: General

General Terms

Standardization, Human Factors, Experimentation

Keywords

Narrative, Adaptive Narrative, Interactive Documentaries

1. INTRODUCTION

Documentaries, and other observational productions such as news reports, often record much more footage than is actually used. Their observational nature makes it difficult to plan shots, so as much footage is captured as possible and this is later refined at the editing stage. This presents an opportunity: there are large libraries of footage currently locked away that might be used to create new, or adaptive, documentaries.

NHT'12 @ HT'12, June 25, 2012, Milwaukee, US.

Author-centric narrative generation systems (such as ArtEquAKT [25]) have sought to utilise multimedia collections like these to generate narratives, often using predefined structures and either generated or annotated metadata [21]. Such a generator may be able to utilise the collections of footage to generate or adapt personalised documentaries by allowing the user to state their interest in viewed scenes. Based on this feedback the documentary could then regenerate and adapt the remainder of its content to include more of what the viewer was interested in.

However, a significant challenge in building interactive narratives of this form is that as the level of interactivity increases it becomes more difficult to maintain the shape of the narrative, Aylett and Louchart refer to this as the 'narrative paradox' [1] [18]. In this paper we propose a model to manage the narrative paradox in the context of interactive documentaries.

Documentaries are typically comprised of a selection of interweaved plot 'threads' based on different aspects connected with the subject matter. As such, the generation process would have to select and weave threads, and remove and replace key scenes which might affect the understanding of subsequent elements within a thread. Our proposed model works by separating logical and thematic cohesion within the narrative, focusing on maintaining logical cohesion within a thread, and thematic cohesion between different threads. In this way we can respond to user choices by adjusting the themes (and thus the choices around weaving threads) while being more rigorous about maintaining logical cohesion within each thread.

We call our approach the Narrative Braid. We propose using the OntoMedia [17] and Stories [15] ontologies to model logical relations within each thread, while thematic cohesion will build upon the Thematic Model Builder [9] and Thematic Illustrator [8].

2. BACKGROUND

2.1 Thematics and Narrative Cohesion

In his work *Thematics* [24] Tomashevsky asserts that 'To be coherent, a verbal structure must have a unifying theme', and explained the importance of theme both to giving a narrative direction and binding its component elements together. Tomashevsky's structuralist view was that themes within narrative could be deconstructed into *themes* (broad ideas such as 'politics' or 'drama') and *motifs* (more atomic elements directly related to the narrative such as 'the helpful

beast' or 'the thespian'). A motif is the smallest atomic thematic element and refers to an individual element within the narrative which connotes in some way the theme. Themes may always be deconstructed into other themes or motifs whereas a motif may not be deconstructed. This definition allows for us to see how literal denotations (motifs) based on the content of a section of narrative may be connected to less tangible concepts (themes) within the narrative subtext. It is these unifying concepts, themes, to which Tomashevsky attributes (at least in part) narrative cohesion.

Subsequent explorations in narrative cohesion demonstrate it to be a complicated function or potentially many different features, of which theme is one part. Studies performed by linguists such as those by Hudson [13], analysing narrative cohesion within children's stories, concentrate on narrative cohesion from the point of view of logical sense, such as causal connections between parts of the narrative. It is also possible that consistent connection to an identifiable genre has a part to play in ensuring narrative cohesion. This is demonstrated within the Coh-Metrix project [7], which produced a system capable of measuring narrative cohesion automatically according to a defined set of metrics which were based on causal, linguistic, cohesion, much like Hudson, but also genre. In our own work [10] in the area we have found evidence and literature to suggest narrative cohesion is connected to five key variables: logical sense, theme, genre, narrator, and style. However subsequent explorations [8] of this have also shown statistical correlations that suggest these variables may be interconnected: particularly theme, genre, and logical sense.

2.2 Narrative Systems

In order to perform the analysis of a piece of documentary footage, information about the individual scenes is required. This can be achieved by annotating the media, whether manually or automatically, with a standard vocabulary. OntoMedia, or Ontology for Multimedia, was developed to support this markup of heterogeneous media content. Its design was influenced by existing standards, including CIDOC CRM [6] and ABC [16].

At its core, OntoMedia defines events and entities. The former provide temporal information: events have occurrences on a timeline, and these can be described as preceding or following other occurrences. The latter provides extensions such as character, item, and space. These allow for the description of the characters involved in an event, any physical or abstract items, and the locations in which they occur. For example, in the context of a nature documentary there may be penguins (characters) on an ice floe (item) in Antarctica (location). Entities can also have traits - whether perceived by other characters or actual. These could include, in this case, gender and species. It is also possible to generate OntoMedia-compliant annotations from Text Encoding Initiative (TEI) performance texts, such as screenplays [14].

The Stories Ontology [15] was derived from the OntoMedia ontologies, and builds upon existing ontologies where possible (FOAF [3] for people, the Event Ontology [20] for events). It is purposefully light: a Story is defined as having a subject and a set of Events, and these Events contain information about the people, places, and things involved, as well as the time at which they occurred on a timeline. Interpretations can be built of these Events, and facts can be asserted as part of these. The Stories Ontology has been applied to content at the BBC [22], and has been used to compare the events within an episode of the popular science fiction series Doctor Who with text-chat conversations between viewers of that episode.

Narrative generation covers a range of narrative systems concerned with actually automatically creating narrative based on a set of requirements. Riedl and Young [21] classify narrative generation systems in three categories: character-centric, author-centric, and story-centric. The first two categories, character-centric and author-centric, are the most common by a significant margin.

Character-centric systems attempt to simulate the contents of a narrative and report on what occurs. This normally involves intelligent agents playing the roles of different characters using a set of rules that govern their personality, goals, and how they can interact with the environment and each other. Examples of such a system include the earliest example of narrative generation, TaleSpin [19], which handled quite limited characters and textual presentation, but also more sophisticated contemporary works such as those by Cavazza [5] with complicated models of character goals and desires. Author-centric narrative generation instead tries to simulate the authorial process of creation, often by filling either predefined or generated structures with relevant narrative material, much like filling a template. This material is often taken from the web, such as in ArtEquAKT [25] which generates artist biographies from web resources, or predefined collections, such as in GRIOT [11] where authors prepare materials for generated interactive poetry.

Riedl and Young [21] explain that arguably the best approach is a compromise between character and author centric generation which takes advantage of the character believability of one and the plot coherence of the other. They present their own solution in the form of ACONF [21] which utilises author-centric plot structures of scenes linked together by a planning system which has defined roles that are populated by character-centric expert systems. Other compromise approaches include the Virtual Storyteller [23] which largely adopts a traditional character-centric approach of character agents interacting with an environment but also includes a 'director' agent. The director agent is motivated by rules governing what creates an interesting and cohesive story and vets the actions of character agents to ensure this is maintained.

Interactive narratives, generated or otherwise, suffer from the problem that, as interactivity increases, the author loses control of the shape of the eventual narrative. This is what Aylett and Louchart [1] [18] refer to as the 'narrative paradox': allowing your audience to interact and change the course of a story reduces its cohesion. This argument can be reduced to the notion that narratives are all a collaboration between author and reader, and that even without interactivity the reader still affects the final resulting narrative through their own interpretation. This idea was explored by Calvino as the 'Labyrinth Challenge' [4] where it is explained that the author can only hope to define the maze, not the route the reader takes, and that while some readers will read to follow the plot and get to the end others will read to find out something in particular. Aylett and Louchart have explored solutions to this problem with regards to interactivity, with some success, including attempting to model the relationship pencil and paper roleplaying game 'Games Masters' have with interactivity in narrative [2], where player interac-

tion must always be incorporated into a coherent narrative. Harrell explains in his work on 'agency play'[12] that the key to solving this issue lies in finding a balance between *user agency* (interactive actions taken by user through the protagonist) and *system agency* (actions of protagonist forced by the system to ensure plot cohesion). Harrell explains that cohesion problems with interaction can be overcome by user constraints based on resulting effects to the narrative's cohesion, similar to how compromise narrative generation often vets character-centric agent actions with author-centric director agents.

3. WEAVING THE BRAID

In the following three sections we describe an approach for the generation of narrative for adaptive documentaries. This approach works towards a technique where documentaries involving multiple subplots might be adapted on the fly using additional recorded material based on user preference. The concepts behind the technique are detailed, a means to preserve coherence in the produced narratives is discussed, and finally we suggest a possible data model for the four layers of our conceptual model.

3.1 Conceptual Model

Our approach defines four concepts: from finest granularity to largest, at the narrative level these are the natom and nolecule, and at the discourse level the thread and the braid. A natom, or Narrative Atom [8], describes a segment of narrative from which a story is composed. These are the finest granularity of media available to a given system, so could be a paragraph of text, a scene from a film, or a photograph. These natoms may be rich with features, both machine-discoverable or authored, that may suggest motifs. In turn, these motifs can connote broader themes.

A discourse is therefore eventually derived from many connected natoms, but to build this the system must understand the logical relationships between natoms. A nolecule (Narrative Molecule) is thus defined as a directed (potentially cyclic) graph of natoms that are causally connected (see Figure 1). As such each narrative may be made up of a set of noleclues, that represent each subplot of the narrative, and in presentation that narrative would cut between these different subplots where appropriate. For example, in an adaptive wildlife documentary, there may be a nolecule that captures the potential story for a particular type of animal made up off all the natoms relevant to that animal and mapping the viable causal paths through that material.

There may be many possible causal traversals from one natom to another within a nolecule, and each chain of traversals is a possible 'thread'. As applied to our wildlife example, a given thread may cover the birth of an animal, but omit the process of finding a breeding ground. The thread represents the presented pathway through the nolecule for that particular subplot, a 'directors' cut of that material.

Given multiple threads, a narrative 'braid' can be woven. The 'braid' represents a continuous chain of natoms from cutting together parts of the component threads. The causal ordering within each thread is preserved - for example, given thread A with '123' and thread B with '789', a possible braid could be '172389'. Each of these threads is causally coherent, but they may have very different dominant themes (as represented diagrammatically in Figure 2). This requires great attention to how the threads are cut together so that

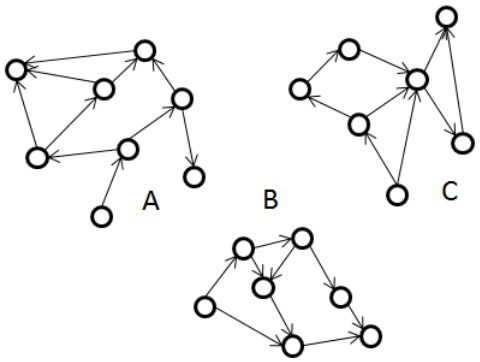

Figure 1: From the causal constraints for each narrative element a network of potential paths through each subplot is emergent.

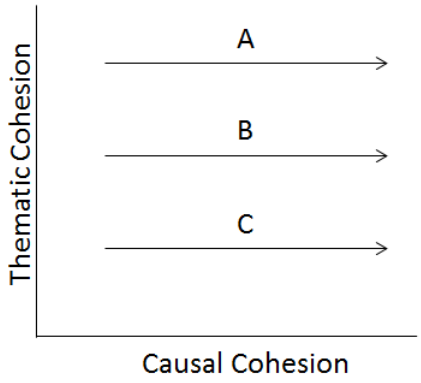

Figure 2: Three plot lines maintain causal coherence but must be woven together for thematic coherence.

a thematically coherent 'braid' may be woven together.

3.2 Thematic Coherence

The paths that narrative threads may take through each subplot network are based upon causal connections and constraints (e.g. 'scene A must occur before scene B'). In an adaptive documentary which of these paths are taken may be based on audience preference (e.g. if the viewer is interested in a particular animal the system might pick paths that show more of that animal). However, this may lead to thematically noisy stories where adjustments are made to meet user desires at the cost of overall narrative cohesion. Also causal connections are less useful when cutting from one plot thread to another where there might be no causality connecting the two threads. When these cuts should be made, and what this decision should be based on, is a significant challenge.

We suggest that the individual plot threads should be blended together through thematic coherence to form the resulting narrative braid. To do this the system would have to determine the current dominant theme of the narrative and select elements that emphasise this theme to ensure stronger cohesion. This could be done utilising a thematic model with definitions for individual themes, such as that explored by us in previous work [9]. The thematic model would con-

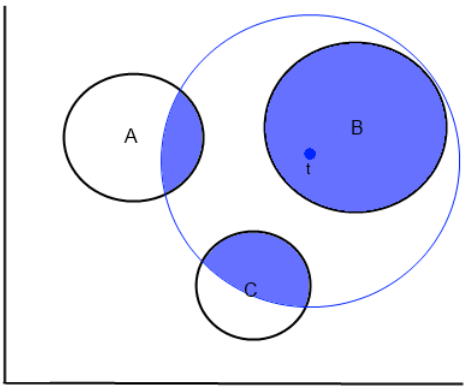

Figure 3: Radius of items from three plot lines relevant to dominant theme t

sult its thematic definitions, and currently selected narrative atoms, to determine the dominant theme by presence of relevant motifs. It could then be used to ensure subsequent elements are selected based on their thematic relevance to further strengthen the theme and increase thematic cohesion.

This can also be used to influence when a cut between threads is made: when the thematic relevance of possible connections in a plot line drops below a defined threshold the system will look to make a cut and shift to a new plot line with a node that is thematically strong and maintains cohesion. It is important to note that strong storytelling does not always require strong cohesion, and breaking cohesion can be done deliberately to good effect, but modelling such a choice is outside the scope of this particular piece of work where we will assume a strong cohesion is preferential.

It is likely that different plot threads will have different core themes, but may perhaps share individual motifs. The dominant theme for the narrative braid at any given time will likely be that of the plot thread that has thus far received the most exposure (in an adaptive documentary this could likely be due to audience preference for the content of that plot line). As such, thematically relevant elements will be selected from the majority of the plot line connected with the dominant theme, and the parts of other plot lines closest to this theme.

Figure 3 represents this diagrammatically. Each circle represents the range of natoms from threads A, B, and C in a multi-dimensional thematic feature space, with the dot representing the dominant theme t. The areas highlighted blue are those with a thematic relevance measured above an acceptable threshold and as such represent thematically relevant items. By keeping natom selection within this radius (even if we jump between threads) we strengthen the core theme and therefore improve thematic cohesion. Previous work has shown that an improvement in thematic cohesion is correlated to an improvement in other cohesion metrics [8], and as such this thematic emphasis may improve the narrative cohesion of the resulting narrative braid as a whole.

3.3 Data Model

The data will be modelling using the OntoMedia or Stories ontologies. The natoms themselves can be represented as Events, with the involved entities attached. In the case of a scene from a documentary, information about the content may also be linked to the natom. For example, the dominant colour, the amount of motion, or the textural content of the video.

The links in a nolecule may also be specified: OntoMedia allows for 'precedes' and 'follows' properties. This is a temporal link, however, so it may be necessary to construct a causal relationship. It would also be appropriate to model the 'bond' between two natoms. This may be a multidimensional property including shared attributes, and spatiotemporal information. These properties would be shared by the threads generated from the nolecular structures.

Finally, the braid would have similar bonding properties, though there may be extra attributes due to the natoms originating from different nolecules. There would also be a representation of the importance of elements while the braid is being manipulated or viewed: this could be attached to the natoms, and then propagated through the higher-level models to influence the bonding properties and thread selection process.

4. PROPOSED APPLICATION: ADAPTIVE DOCUMENTARIES

Documentaries have a large amount of available footage, much of which may be edited out of a final cut due to time constraints. As such, this is an ideal fit for our approach: a 'regular' documentary is simply a braid consisting of threads that have been manually selected to tell the story effectively within a limited time. If, however, a nolecule is formed for the various threads, it may be possible to dynamically alter the documentary at viewing time.

As a first step, the various videos would have to be annotated into natoms. This could be achieved part-automatically, using shot cut detection to segment the video. These natoms would then be arranged into nolecule structures. Given two plot threads, for example in our wildlife documentary a thread about penguins and a thread about seals, there would be two nolecules. It might be that one scene includes both penguins and seals, and therefore is shared between two nolecules (this is therefore quite likely to appear in a thread). The connection-building process should preserve the causal links within each thread, for example ensuring that 'searching for a breeding site' precedes 'a penguin is born', and as such these links form the potential paths for threads through each nolecule.

The process would then switch to a user-controlled application. Here a viewer would watch a documentary braid, constructed from threads selected to be narratively consistent. The threads would be causally consistent and cut together in a way that preserves thematic cohesion. As the narrative is presented a viewer would have the ability to up-vote scenes they enjoy or are interested in. By up-voting a given natom, causal links would be reweighted in the nolecules in favour of themes or content in the up-voted natom. This would cause the threads to be reformed and braided together based on this new preference, and causing subsequent scenes to alter from what was previously planned. The causal links would remain consistent to preserve the coherence of the documentary, but any new themes and attributes of the promoted natoms might be more prominent in the final narrative. So for example, if a viewer found the breeding sites of penguins interesting, and

thus up-voted the scene, the thread conveying this part of the story might be elongated, but still arranged within the final braid to make sense. In addition themes associated with that natom (for example, of 'birth' and 'family') would become more dominant, making it more likely that natoms with those themes would be selected when splicing in other threads.

5. CONCLUSIONS AND FUTURE WORK

This paper proposes an approach that allows for user interaction with a multi-layered adaptive documentary in such a way that narrative cohesion is kept throughout the viewing. It models the narrative as multiple threads, woven together to create a braid, and manages the narrative paradox by separating logical cohesion within threads from thematic cohesion across the whole braid.

The Narrative Paradox suggests that interactivity and authorial design are at odds with each other. We hope that with careful modelling of causal and thematic links this limitation can be overcome.

6. REFERENCES

[1] R. Aylett. Narrative in virtual environments-towards emergent narrative. Technical report, Working notes of the Narrative Intelligence Symposium, 1999.

[2] R. Aylett, S. Louchart, A. Tychsen, M. Hitchens, R. Figueiredo, and C. D. Mata. Managing emergent character-based narrative. In *Proceedings of the 2nd international conference on INtelligent TEchnologies for interactive enterTAINment*, INTETAIN '08, pages 5:1–5:8, ICST, Brussels, Belgium, Belgium, 2007. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[3] D. Brickley and L. Miller. FOAF Vocabulary Specification 0.97. Namespace document, January 2010.

[4] I. Calvino. *Una pietra sopra*, chapter La sfida al labirinto, pages 82–97. 1980.

[5] M. Cavazza, F. Charles, and S. J. Mead. Agents interaction in virtual storytelling. In *Intelligent Virtual Agents, Springer Lecture Notes in Artificial Intelligence 2190*, pages 156–170. Springer-Verlag, 2001.

[6] M. Doerr. The CIDOC CRM - an ontological approach to semantic interoperability of metadata. *AI Magazine*, 24:2003, 2003.

[7] A. Graesser, D. McNamara, M. Louwerse, and Z. Cai. Coh-metrix: Analysis of text on cohesion and language. *Behavior Research Methods*, 36:193–202, 2004.

[8] C. Hargood. *Semiotic Term Expansion as the Basis for Thematic Models in Narrative Systems.* PhD thesis, University of Southampton, 2011.

[9] C. Hargood, D. Millard, and M. Weal. A semiotic approach for the generation of themed photo narratives. In *HT '10: Proceedings of the 21st ACM conference on Hypertext and hypermedia*, pages 19–28, New York, NY, USA, 2010. ACM.

[10] C. Hargood, D. Millard, and M. Weal. Measuring narrative cohesion: A five variables approach. In *Narrative and Hypertext at Hypertext '11*, 2011.

[11] D. Harrell. Shades of computational evocation and meaning: The griot system and improvisational poetry generation. In *Proceedings, Sixth Digital Arts and Culture Conference*, pages 133–143. Citeseer, 2005.

[12] D. Harrell and J. Zhu. Agency play: Dimensions of agency for interactive narrative design. In *Proceedings of the AAAI 2009 Spring Symposium on Narrative Intelligence II. Menlo Park: AAAI Press*, pages 44–52, 2009.

[13] J. A. Hudson and L. Shapiro. *Developing Narrative Structure*, chapter From Knowing to Telling: The Development of Children's Scripts, Stories, and Personal Narratives, pages 89–109. American Psychological Association, 1991.

[14] M. O. Jewell. Semantic screenplays: Preparing TEI for linked data. In *Digital Humanities*, 2010. Paper presented as part of Scanning Between the Lines: The Search for the Semantic Story.

[15] M. O. Jewell and C. J. Hooper. Linking locations: Storytelling with pervasive technology. In *Narrative and Hypertext Workshop at Hypertext'11*, June 2011.

[16] C. Lagoze and J. Hunter. The ABC ontology and model. In *Dublin Core Conference*, pages 160–176, 2001.

[17] K. F. Lawrence, M. O. Jewell, m. c. schraefel, and A. Prugel-Bennett. Annotation of heterogenous media using OntoMedia. In *First International Workshop on Semantic Web Annotations for Multimedia (SWAMM)*, May 2006.

[18] S. Louchart and R. Aylett. Solving the narrative paradox in ves lessons from rpgs. In *Intelligent virtual agents: 4th International Workshop*, 2003.

[19] J. R. Meehan. Tale-spin, an interactive program that writes stories. In *In Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 91–98, 1977.

[20] Y. Raimond and S. Abdallah. The Event Ontology, 2007. http://motools.sf.net/event.

[21] M. O. Riedl and R. M. Young. Character-focused narrative generation for execution in virtual worlds. In *Proceedings of the International Conference on Virtual Storytelling*, pages 47–56, 2003.

[22] P. Rissen. Re-imagining the creative and cultural output of the BBC with the Semantic Web. In *Digital Humanities*, 2010.

[23] M. Theune, S. Faas, A. Nijholt, and D. Heylen. The virtual storyteller: Story creation by intelligent agents. In *TIDSE 2003: Technologies for Interactive Digital Storytelling and Entertainment*, 2003.

[24] B. Tomashevsky. *Russian Formalist Criticism: Four Essays*, chapter Thematics, pages 66–68. University of Nebraska Press, 1965.

[25] M. Weal, H. Alani, S. Kim, P. Lewis, D. Millard, P. Sinclair, D. D. Roure, and N. Shadbolt. Ontologies as facilitators for repurposing web documents. *International Journal of Human-Computer Studies*, 65:537–562, 2007.