UNIVERSITY OF SOUTHAMPTON
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

A mini-thesis submitted for transfer from MPhil to PhD

Supervisor: Dr David Millard
Advisor: Dr Mark Weal
Examiner: Prof Paul Lewis

Exploring the Importance of Themes in Narrative Systems

by Charlie Hargood

July 23, 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

A mini-thesis submitted for transfer from MPhil to PhD

by Charlie Hargood

The problem of presenting information surrounds the need to make information relevant, personalised and engaging. As a prevalent form of presenting information narratives can be used to make information engaging and by their nature are relevant to their audience. Were we able to automaticly generate narratives for presenting information we could also personalise it, creating custom personalised narratives that selected relevant information to present and did so through an engaging experience for the audience. However existing narrative systems can either fail to generate quality narratives or sacrifice their autonomy to do so. This document presents a machine understandable thematic model for expressing themes in narratives. This model can be used in a thematic system to give a thematic subtext to the presentation of information, enriching generated narratives, and improving the thematic relevance of information. A prototype using this model in an experiment with photo montages as simple narratives demonstrated that use of the model can successfully cause results to connote desired themes and improve their relevance to titles with thematic content over simpler keyword methods. There are a variety of ways such a system could be integrated with modern narrative systems in order to enrich their results without sacrificing autonomy and the development of a thematic presenter with thematic analysis could be used to elaborate thematic content in a narrative improving its thematic cohesion and relevance.

Contents

Chapter 1

Introduction

Narratives or stories are a traditional human form of communicating information represented in all cultures. The personal focus of narrator and audience means the information is specifically positioned to be relevant to the audience as well as engaging in order to best fulfill that narratives objectives whether it be to communicate information, entertain, or both. This potentially makes narratives a powerful representation of information that could be personalised as well as relevant and engaging that could be used both as a means to tackle information overload [7] as well as make presentations of information entertaining. However narratives are difficult to make individually personal and it is infeasible to human author a narrative for every user request. Narrative Generation presents a potential solution to this constraint, potentially allowing custom narratives to be generated automatically. If a system were able to author a narrative from a set of requests and conditions then relevant information could be presented in an engaging fashion and the system could personalise it to the user at hand. However there is more to generating engaging narratives then just creating a plot.

Themes give a richness to the narrative they inhabit by giving it subtext and communicating subtleties to the audience beyond the forthright message of the story. Ensuring a narrative has thematic objectives as well as plot ones and thematic coherence makes the story more elaborate and engaging and ensures the narrative coherently flows together. Thematics is an area of narratology that seeks to study themes within narrative from a structuralist approach, Tomashevskys work [37] in this area worked towards identifying the important structures in narrative that build towards a thematic subtext.

This mini thesis describes a machine understandable thematic model based on Tomashevskys structuralist work in thematics. It describes the thematic content of narrative in terms of the key components of themes, motifs, features and natoms (narrative atoms). Using an instance of this model it is possible to define the required components for a narrative to successfully connote a desired theme, it is also possible to use it to analyse the themes already present within a narrative. It is our intention to demonstrate that

this thematic model may be used in order to improve representations of information by making them more relevant and by making narrative representations more engaging by improving the thematic subtext.

Quality generated narratives could be used in a range of use cases ranging from a computer game that writes the plot as it goes along, adapting the plot to the players actions and decisions; to a personalised news reader that always knows what you are interested in and what you already know. Narrative generation could also be used in a more traditional search context - representing search results in an engaging narrative context rather then an abstract list.

Some projects are focused on understanding narrative as a representation of information such as OntoMedia [21][22] which seeks to build a shared vocabulary for annotating narrative, where as adaptive hypermedia systems often seek to create a personalised experience for the user either through user modeling such as in projects like AHA! [10] or by allowing the user to personalise their own experience through their choices such as in HyperCafe [35]. Similar systems also use narrative techniques in order to make presented information more engaging such as the Topia project [3] which uses the ideas of sequencing and emphasis to create rudimentary narratives of search results in order to better engage the user.

Full narrative generation systems seek to create the content itself, rather then represent existing information in a more engaging or personalised way. Various systems have implemented different aspects of narrative generation with some requiring prewritten characters and settings such as AConf [33] and others generating everything such as Talespin [28]. These systems are often composed of a variety of smaller systems each generating a component part to represent the multi layered task that is narrative generation. Resolving the plot of the narrative is very different to presenting it as a desired piece of media.

The theory of narrative generation is often based upon structuralist work in narratology where literary theory is used to deconstruct narrative into its component elements and devices. These identified devices are then often what is used to generate the desired narratives in narrative generation however the specific approach used varies from system to system with some taking a very structuralist author centric approach dealing with the narrative to be generated as if it were a story to be written where as others take a more character centric approach seeking to simulate the events the narrative describes.

Existing narrative generation techniques experience mixed success, while many successfully generate stories they can seem bland and very straight forward. They can become very formulaic being heavily reliant on the defined structures for generation or in the case of more simulated methods of generation with emergent narratives they can be comprised of a simple uninspired list of character actions. It is possible that narrative generations focus of creating a story and plot has lead it to overlook other essential parts

of the authorial process, such as themes. This motivation lead to the work conducted on the creation of a thematic model, which could potentially be used to improve the results of narrative generation.

A prototype called the TMB (Thematic Model Builder) has been developed which uses the thematic model to calculate the thematic quality of candidate natoms using a variety of thematic metrics. This has been used to build themed photo montages from Flickr[1]. Using this we conducted an experiment into the effectiveness of the thematic model by assessing the relevance of montages it built for particular queries. The results of this experiment are presented within this mini-thesis and using these, the quality of both the model and the method of calculating thematic quality are examined as well as the beginnings of an investigation into a formalised method for building instances of the model.

Finally we consider the future of this research. The development of a formalised process for authoring instances of the model, the building of a new prototype that is capable of using the model for thematic analysis and emphasis, and how such a prototype will be evaluated to show the effect of the thematic model on stories (generated or otherwise) and whether a better thematic cohesion is possible using the model.

[1] http://www.flickr.com

Chapter 2

Background Research

2.1 Narratology

Although narratology, as a study of literature, is mostly focused on the analysis of narrative it provides a detailed insight into how narratives are built.

One approach to narratology, structuralism, deconstructs narrative and aims to learn about the components from which a story is built and how they are connected and contrasted against each other. As this defines tangible objects within a narrative that can be modeled there is much narrative generation can use from structuralism as it can seek to generate the structures that structuralists have defined. Most structuralist theories asserts that a narrative is composed of any series of human experiences [27], and may be deconstructed into a story and a discourse [6] where the story (or fabula) represents a chronology of all the information to be communicated and the discourse (or sjuzhet) represents what parts of the story are told and how those parts are presented (shown in Figure 2.1). Other layered approaches to narratives exist as well such as Bal [4] which similarly makes a division between the sum of all events (Fabula) and the story that is told but divides the selection of story elements and its presentation (collectively the discourse according to Barthes) into seperate layers of story and narrative.

The story element is constructed by the experiences that make up the subject of the narrative. In a virtual collection of resources the story represents the collection of experiences represented as resources. The discourse however represents what parts of the story are told (the story selection) and how it is told (the story presentation); if the collection is the story then the result of narrative generation (telling the story) is the discourse.

The discourse is the result of a multitude of different mechanics including how the story is presented, what medium is used, the style, the genre, and the themes of the narrative. The study of thematics approaches themes with a structuralist method of deconstruction

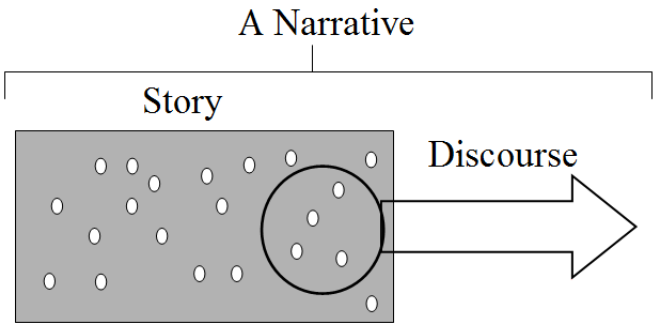

FIGURE 2.1: A narrative can be deconstructed into story and discourse

and attempts to identify the narrative elements that communicate themes. Themes are a subtle subtext of the narrative communicating an objectivity aside from the literal events in the plot and add a personel influence to the story, they could be considered a part of the authorial voice [9] of the narrative.

Tomashevsky deconstructs thematic elements 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') [37]. He describes a structure of themes being built out of sub-themes and motifs. 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.

2.1.1 Semiotics

Semiotics or semiology is the study of signs and how we extract meaning from them. Saussure wrote that all signs are built of two parts[34], a signifier (the physical signal from the sign such as the appearance of an apple) and a signified (the denotation of that sign such as the concept of 'apple-ness' or 'fruit'). Barthes made a distinction between denotative signs (signifiers that lead directly to their signified, such as a word having a literal meaning) and connotative signs (signifiers that lead indirectly to some contextual or culturally important signified, such as the red light implying Stop to a driver)[5]. Barthes goes on to point out that should a sign connote something then the signifier of such a sign would itself be built out of a denotative sign (a picture of a red light denotes a red light, red light connotes Stop). In such a way we can draw contextual cultural concepts from static basic objects that in a particular context have a greater meaning.

2.2 Narrative Generation

Narrative generation has a variety of applications in systems that deal with different information, as a narrative can be any collection of human experience it is not limited to written prose but to any representation of human experience, indeed the importance of non textual narratives highlighted in [15] shows that increasing multimedia communication can benefit from a narrative context. Some systems use narrative as a lens through which to view a larger collection, for example PhotoCopia [38] which presents narrative photo montages. While many systems seek to generate full narratives for entertainment such as the virtual storyteller [36] and AConf [33] some systems use narrative generation to add additional meaning to information by representing it as a narrative using narratological devices like sequencing, emphasis and omission such as in Topia [3] and some adaptive hypermedia systems like the Hyperdoc [30] and AHA! [10] where narrative generation allows for the adaptive presentation of information based on initial objectives that is fundamental to adaptive hypermedia [11].

As a process narrative generation can be broken down into three stages; story, plot, and presentation generation. Depending on the project in question these stages can be consolidated together or separated, (for example, in the virtual storyteller, presentation generation is broken down in narration and presentation [36]). The majority of narrative generation projects deal with the creation of the narrative elements (story generation); resolution of the sequence of events that comprise the narrative and selection of narrative elements to be exposed and building of relationships between these elements (plot generation); and presentation of the narrative through a chosen medium (presentation generation). Figure 2.2 illustrates this process.

According to Riedl and Young [33] narrative systems take either a character or author centric approach depending on whether the system seeks to model the characters within the story, the authorial process itself, or whether the system is a compromise of both approaches. [33] also identifies a third approach in the form of story centric approaches these are however less common and due to their more linguistic focus are less relevant to this research.

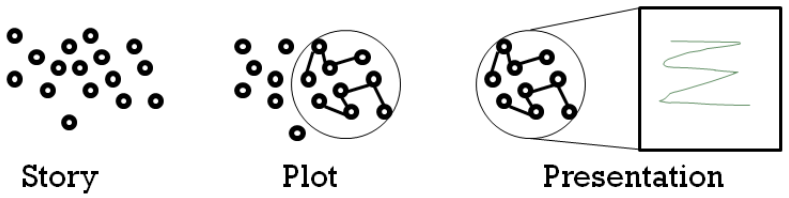

FIGURE 2.2: Narrative generation can be broken down into three stages

2.2.1 Character Centric

Character centric narrative generation revolves around the perspective of modeling the behavior and goals of the characters of a story. With the characters successfully simulated they are released to pursue their goals and their actions are exposed, the idea being that stories are everywhere and an engaging narrative will naturally emerge from the actions of a set of well-motivated characters.

Character centric narrative generation systems often use agent technology to suitably simulate the characters and their behaviors with a purpose built agent taking the part of each character such as in work by Cavazza [14] and in the Facade system [25] (Facade is not entirely character centric, but its approach is very similar). Sometimes the intelligence is much more simplistic and a reasoning system will handle the goals and behavior of all characters, such as in TaleSpin [28]. However, these systems lack the power to generate varied narratives and although short simple stories are generated the lack of in-depth modeling of individual characters behavior removes personalized variety from their actions.

Automatic generation of story elements is rare in character centric narrative generation. This is because elegantly written characters with sophisticated behavior are key to narratives being successfully emergent from the generated result and at present the only way to ensure this is to build the characters by hand. Some story elements are generated by using character archetypes with cliche behavior such as with the supporting characters in work by Cavazza [14] but it is rare to find this for key characters.

Plot generation in character centric generation is therefore a direct result of the characters behavior as dictated by the agents playing them or the intelligence modeling all of the characters. The actions they take to achieve their goals builds the relationships between story elements and the sequence of events that makes a plot. Presentation generation is not specifically tied to the character centric approach but the focus on entities and modeling their actions make character centric approaches ideal for presentation in game engines (for example AConf [33] used the UT engine through the mimesis project [40]). Although the presentation of character centric systems still sometimes uses text as a medium of choice either using sentence templates such as in talespin [28] or generated text using natural language processing.

The main weakness of character centric narrative generation is its reliance on an engaging narrative successfully emerging from the exposition of the characters actions. Often these systems generate bland stories that merely report on a series of uninteresting actions. These stories are thus often sensible and varied but lack narrative richness or interesting plot.

2.2.2 Author Centric

Author Centric narrative generation seeks to model the authorial process itself rather then the content of the narrative. The systems seek to model the process by creating rule based systems or narrative grammars that use well defined structures that are typical of the desired genre of narrative in order to generate stories, it is typical of earlier work in the area such as that by Lebowitz [23] or Hovy [20] to seek the relationships between elements that make up the discourse of a story in order to build up structural grammars.

Author centric narrative generation also lends itself better to the representation of existing knowledge as narrative as its story elements are not necessarily the narrative devices such as characters and objects but the devices the author needs to construct a story. Systems such as ArtEquAKT [39] create narratives out of a variety of resources and media from the internet and for this project story generation is the compilation of these resources. ArtEquaAKT automaticly retrieved relevant story elements off to web [2] and using structures and grammars for its desired genre (Biographies) and the link server AuldLinky [29] generated working narratives for artists. The same could be said for narrative influenced hypertext systems such as Topia [3] which generated simplistic narratives from search results by using sequencing, emphasis, and omission at the presentation level in order to create an effective narrative.

Some systems do model the contents of the narrative to be generated as part of story generation but still remain author centric. Universe [23] builds stories around a set of author goals and constructs a structure for a story to satisfy these but does so using the actions of characters modeled from cliche archetypes and a finite set of actions. In other author centric systems the story structure itself is not explicitly generated, but emerges from the selection of a predefined set of story elements, such as in Card Shark [8].

Plot generation in these systems is a case of applying the rules of the system for the desired genre, utilizing the grammar with the available resources, or filling a story template with appropriate resources. Presentation for author centric systems is often text based, either using templates such as Universe[23] or ArtEquAKT [39] or simply exposing the elements in sequence such as in Card Shark [8].

Author centric systems tend to be highly specialized for one particular type of narrative, making them inflexible and also often not with a view to generic narrative generation. The stories are seldom varied as they all follow a similar authoring process with the same rules and/or grammars and as such can generate engaging but not often varied narratives.

2.2.3 Compromise Approaches

Many narrative generation systems often seek a compromise between these two approaches in order to counteract the weakness of using one approach or another. Some systems such as Facade[25] and Universe[23] will only make slight compromises, such as the ideal story drama curve approach in Facade or the choice to model characters in Universe, but others make much larger steps towards marrying the two approaches.

The virtual storyteller [36] at first seems to be a character centric approach that uses agents to model the behavior of its characters, the difference arises with the addition of an extra director agent. The director agent has a set of rules about what makes an engaging story, much like an author centric approach, and uses these rules to influence the narrative by vetting character actions, influencing them by giving them new goals, and creating story events to channel the emergent narrative into being more engaging.

AConf[33] models each 'actor' as an expert system seeking to achieve its goals, giving it characteristics of a character centric approach, but it is fundamentally much more author centric as its process of plot generation centers around the structure of the narrative building it as a network of events using text and story planners.

The presentation generation for these systems also vary. In the virtual storyteller [36] the director agent directly communicates with a narrator and presenter to generate text using sentence templates whereas AConf [33] uses its character's modeling and plot planning using Longbow [41] to interact with a system called mimesis [40] which uses the UT game engine to present the narratives.

These systems experienced mixed success with both reporting the generation of successful narratives. However both suffered from similar problems to character centric approaches, while the addition of measures to ensure the narratives structure is engaging does have a positive effect the engaging narrative can at times still fail to emerge from the result and the systems can be reliant of stories that are heavily predefined at the request stage rather then being entirely generated.

Chapter 3

The Thematic Model

3.1 The Model

Authors use themes to communicate a subtext within a narrative. This subtext may be an agenda or simply an emphasis of a particular part of the narrative or even simply an emphasis of the authors own style. This subtext gives a narrative direction beyond merely communicating a chronology leading to deeper narratives and giving an authorial voice to stories. In previous work [17] we proposed a thematic underpinning to narrative generation in the form of a thematic model that described how themes are constructed within a narrative.

Existing systems work in thematics often looks at themes as a classification of the content itself in a piece of information, performing keyword extraction in order to classify a document into a particular topic such as in [24]. These IR projects often seek to show that thematic classification can show a marked improvement over traditional keyword search but approach the concept of themes from the perspective of the core content of a document rather then its subtext. Projects such as [32] bring thematic extraction closer towards discourse theory investigating human readers observations of what makes a documents theme but are still tethered to the idea of themes as topical classifications of content. This model seeks to represent the thematic subtext of a narrative and the concepts that influence the narratives content rather then classifications of the content itself.

To do this we go back to Tomashevsky's structurist work on thematics[37]. Features within the narrative denote Motifs and from these Themes can be identified. We assume a situation where a story is compiled with many small segments of narrative that are structured together, in this case the selection of these small atomic segments and their content are key to communicating a theme. We use the term Narrative-Atoms or *Natoms* to describe these segments; small atomic pieces of narrative that cannot be further broken down, for example a single photo or paragraph. The content of these natoms is rich with

information, however only some of it visible to a machine (such as generated meta data and authored tags), we call these visible computable elements *Features*. Natoms contain any number of features which may or may not work towards connoting a theme in a story. Features can each denote a motif, a basic thematic object that has connotations within the story, for example the feature *cake* denotes the motif of *food*. These motifs in turn connote broader themes in the context in which they are presented, for example *food* in the context of a gathering may connote *feasting*. These themes, when combined with other themes or motifs could in turn be used to further connote other themes, for example feasting might connote celebration. This forms the foundation of our thematic model of a narrative:

- Natoms contain tagged features
- Features denote motif's
- Themes are connoted by other themes and motifs

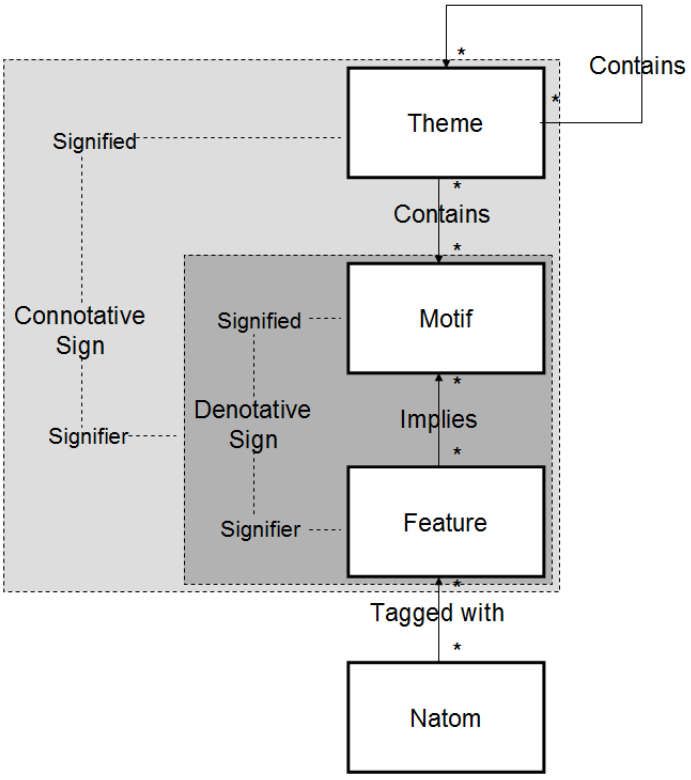

FIGURE 3.1: The Thematic Model

The model is shown in Figure 3.1, which also shows how the parts of the model map to Barthes' idea of denotative signs as the signifiers for connotative signs. Features denote Motifs because motifs are directly associated with the feature (normally as a generalized version of it). Themes are broader concepts communicated over the entirety

of the narrative, typically by numerous motifs. By their nature they cannot be denoted as they rely on some cultural context which cannot be contained within a natom, as such a theme is a connotation of the motifs, and by extension the features, within the narrative. This model is but one part of a narrative generation system, it contains no rules for the presentation of elements or the narrative structure. However it can be used to select natoms to be used within a discourse. As such we could use themes constructed from this model to influence the story selection in author or character centric systems to give them a thematic subtext.

When a narrative is formed a part of the story is selected and then presented as a discourse[6]. We can consider virtual collections of resources as our story, and should we want to create a discourse to tell a story of Tuesday it would select all the natoms (photos, blog entries, etc.) of that day. Using an appropriately populated thematic model we could examine the features of those natoms in order to identify motifs and thus potential themes. Natoms that connote these popular themes could then be selected or emphasized to create a final discourse that felt more purposeful. If the virtual collections were very large we could set out to look for natoms that supported particular themes, for example, by using public photo collections to create a discourse (a photo montage) with the themes of family, winter and Christmas.

Because features could be tagged in any way for such a system to work every motif object would need a list of features that could denote the motif. In turn theme object will also require some way of knowing what motif's are suitable for them, however in this case it is less simple as themes are contextual things not simply denoted. It seems likely that a theme should be described as having core thematic elements that are required for a theme to be communicated, such as a wedding theme requiring a bride motif, as well as optional thematic elements that exaggerate or promote the theme but are not essential (such as a religious theme). Themes would need to keep a set of required and optional thematic elements (both motif's and sub themes). The power of the thematic approach will be proportional to the quantity and richness of these feature-motif and motif-theme connections.

3.2 An Example

Figure 3.2 shows a simple example of how a collection of natoms connotes a theme in the terms of the model, in this case a passage of text[1], and two photographs connoting the theme of winter. The features presented are present within the given natoms, it is feasible that the natoms would be tagged with them or that they might be automatically extracted from them. These features literally denote the motifs of snow, cold, and warm clothing. As snow demonstrates many different features might denote the device of snow

[1]text from William Shakespears Blow, Blow, Thou Winter Wind

but in this case thematically they serve the same effect. Finally in the context of each other these motifs connote the concept and theme of winter.

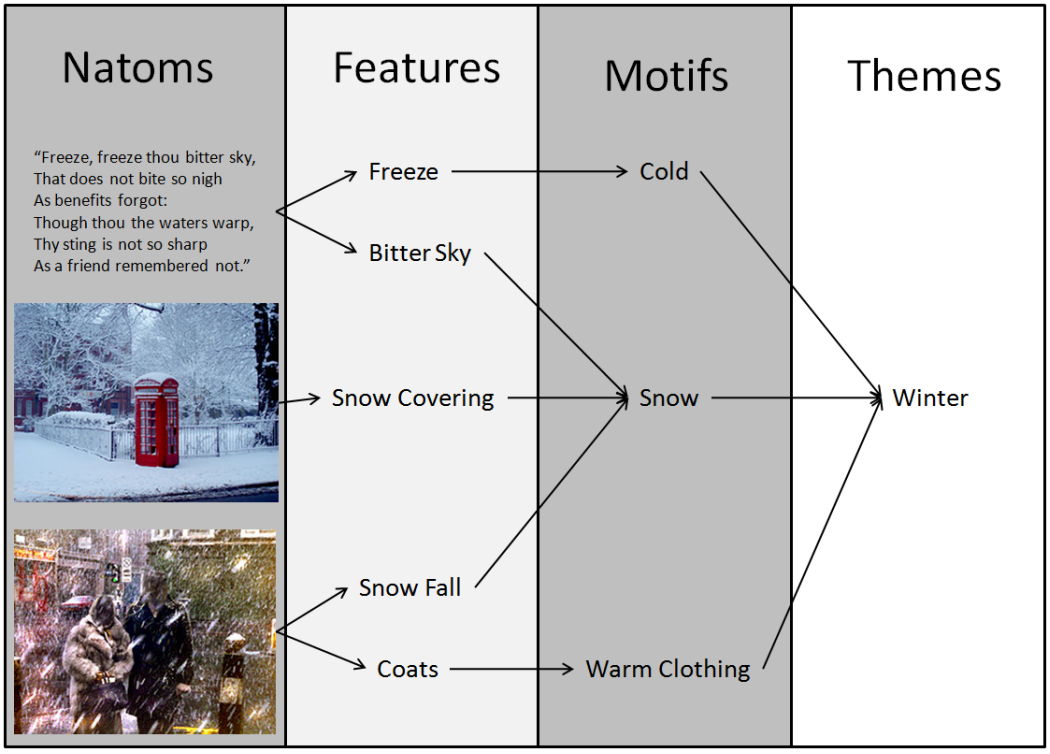

FIGURE 3.2: A Worked Example

3.3 Authoring Themes

Instances of the current model have to be authored manually. While automatic authoring of themes could be possible potentially through a combination of term expansion and thematic analysis of texts it is not within the scope of this research. In order to formalise the process of authoring themes and ensure similarity between themes authored by different authors the process of writing themes has been broken down into a set of rules that if strictly followed should lead to an acceptable instance of the model.

The initial set of rules was an outcome of analysis of the thought process of creating the initial instances of the themes. In order to create a formal set of rules these early rules would have to go through a process of being used by different authors and iterative expert reviews in order to identify weakness' in the rules that lead to insufficient or inaccurately authored themes. This investigation would culminate in an evaluation that assess whether themes authored by multiple authors are capable of successfully connoting the desired themes and whether they are suitably similar.

Currently this process has reached a single expert review of the initial rules. The following rules were given to an author who had no prior knowledge of the process of authoring

themes and they were asked to create an instances of the model for the theme of danger. The subject was allowed to ask clarification on the wording of rules but was given no advice on authoring of the theme itself.

- List Connotations: List all concepts, objects, and words that to you connote the idea of the desired theme. List everything that you associate with it in anyway and to you helps build the idea of the theme in your head.
- Divide Tangible Objects and Concepts: Divide the listed connotations into those that are anchored to specific objects and devices that could be included in a narrative element and those that are broader concepts connoted by many things and less tangible. These broader concepts become Themes.
- Group motifs: Group similar tangible objects together. Consider the relationship the object has with the desired theme and group together objects that belong to the same narrative device. For example, in the theme of picnic chicken sandwich and scotch egg all serve the same purpose of denoting food. These grouped together objects become your themes motifs.
- Iteratively write the contents for sub-themes and motifs: For each theme and motif repeat step 1. For motifs this will be slightly different, as you are not considering a desired high level concept but a much more tangible object you will be listing denotations not connotations, for example list every specific object that might exist in a narrative element that would lead to denoting this concept. Be careful only to list things that directly denote the motif, not associated words, these are the motifs features. For themes the process is identical as step 1 was before. Repeat this step until all sub themes and motifs have been iterated through and written.
- Identify associated themes and motifs: Check the components of every sub-theme and motif of the desired theme, and in turn every sub-theme and motif of each of those. Ensure the entire contents of a sub-theme or motif is relevant to the parent theme and in turn connotes the parent theme. A sub-theme (or motif) that contains elements irrelevant to the parent theme becomes an associated theme and is removed from the model.

The review resulted in an instance for danger that highlighted similar important elements to that authored by the expert, the exact results shown in figures 3.3 and 3.4. However the process did highlight several weakness' in the initial rule set that need to be addressed. The existing rule set asks the author to list all concepts that connot the desired theme where this can seem infinite, without any guidelines as to a sensible number of concepts the author can potentially not do a broad enough selection to encompass all motifs of a theme or go too far and list concepts with only a tangential relevance that slow down the process of authoring when they are only removed anyway in stage

five after being expanded. The semiotic terminology in the existing rules also needs explanation and there should be guidance on the grouping of motifs in step three.

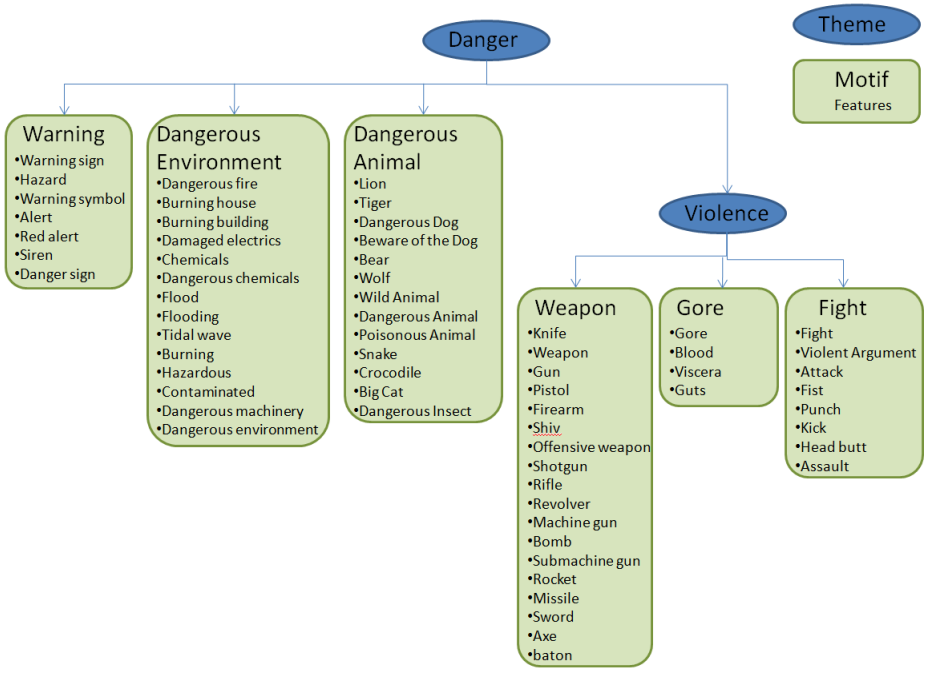

FIGURE 3.3: Original instance for theme of danger

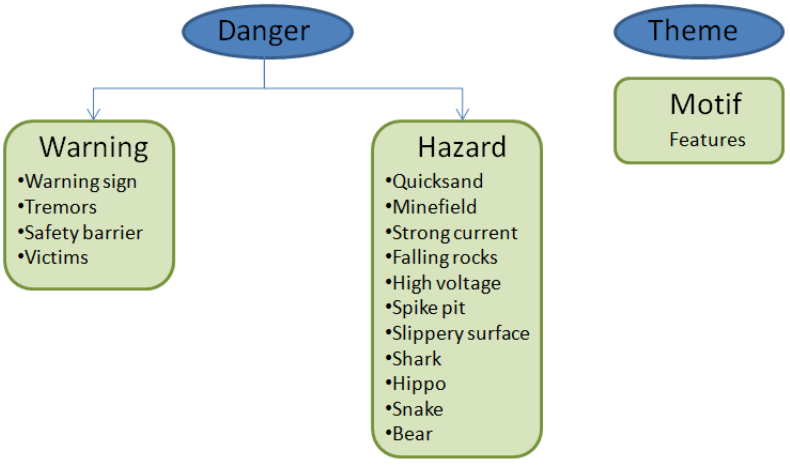

FIGURE 3.4: Instance for theme of data by inexperienced author following rules

Chapter 4

Thematic System Experiment

4.1 The TMB

In order to evaluate the effectiveness of the model a prototype system was built that utilised an instance of the model. The prototype uses the model to select images from Flickr that have strong relevance to particular themes. The prototype went under the working name of the Thematic Model Builder (TMB).

This instance of the model was built in xml and four themes were modeled and expanded (all sub themes and motifs were modeled as well): winter, spring, celebration, and family. The process of defining an instance of the model for particular themes is a complex and subjective one [17]. We explored a systematic method for building themes based on semiotics. Initially we identify what *connotes* that theme, these connotative signs will make up the themes sub themes and motifs. However, these signs become sub-themes only if all of the aspects of their concept in turn connote the theme being built, otherwise the sign should become a seperate theme in its own right. Thematic objects anchored to a particular device within the narrative become motifs which have their features defined by likely tags that *dennote* the object.

The prototype itself was written in java with a simple JSP front end. For the purposes of this prototype and evaluating the model, Flickr was chosen as a source of natoms. As a folksonomy its items have rich semantic annotations in meta data [1] that make the features in each image apparent and it has a large freely available body of resources. The library of images (the fabula) was generated by making a keyword search of Flickr on the desired subject and storing the top n images (where n is the desired size).

The system then followed an algorithm of measuring the thematic quality of each natom in the fabula. A number of metrics were developed as candidates for measuring thematic quality, including coverage(number of relevant elements present), focus(percentage of a natoms features relevant to the request), and match(percentage of a given components

features or motifs relevant to a given natom). Each of these metrics can be measures at the thematic, motif, or component(top level thematic elements of the desired theme) level. Of these two were used for measuring thematic quality in the prototype, the system returns the natoms with the highest scores according to these two metrics:

- *Component coverage*: the proportion of high-level sub-themes or motifs that a natom has features for - this is useful for measuring how strongly a natom matches the desired theme. (for example, winter expands several high-level sub-theme and motifs including christmas, snow and cold. A natom matching just one of these has less coverage than one that matches many)
- *Thematic coverage*: the proportion of desired themes that a natom has features for - this is useful for searches with multiple themes

The TMB Prototype allows us to compare the effectiveness of selecting photos according to their theme with the process of selecting photos based directly on their tags.

Using the thematic model to use motifs to expand a list of features unified under a single theme to retrieve relevant images for a title makes the TMB similar in concept to keyword unification image retrieval projects such as [42] or traditional term expansion such as [12] [16], leading us to expect that it will be more successful then traditional keyword search. The current human authoring process required to make the instances of the thematic model means the expansion process is not automated as in [12]. However because all expansion in the thematic model is relevant specifically to the core theme rather then only associated terms expanded are unlikely to experience the query drift from desired concepts identified in [31] enforcing stronger thematic cohesion without having to resort to term weighting as in [13].

4.2 Evaluation Plan

For the evaluation it was important to measure what advantage there was in using a thematic system for natom selection over a keyword search system, but we also wanted to see whether themes emerged more strongly from in the more narrative context of a group of natoms; in this case a montage of images than with individuals. The experiment also sought to refine the process by which thematic quality is calculated, to do this correlations between natom relevance and where the natom measure with regards to the other metrics would be observed to find if some metrics often correlated with images successfully connoting desired themes. From this we can identify three experiment objectives.

- Evaluate the effectiveness of the TMB in selecting images connoting desired themes in comparison to simple keyword search

- Evaluate the effectiveness of the TMB in the narrative context of generating montages
- Observe any correlations between image relevance and different thematic metrics

The evaluation asked participants to rate images individually and in sets according to how they matched a given subject and theme (for example, 'London in Winter'). The images and sets were generated in four different ways:

- *TMB*: Using the TMB and Flickr API to search by subject and select by component coverage
- *Flickr*: Using Flickr to search by subject and theme, filtered by relevance
- *BaseL(ow)*: Selecting images from Flickr at random
- *BashH(igh)*: Using Flickr to search tags by subject and filter manually

In this way we hoped to compare the performance of the TMB with keyword search on Flickr, and place both of these methods in context by comparing them to random and hand-picked samples. For each test the user would be presented with two titles and under each the images for the test (depending on the test either individually or in groups) and asked to rate them 1-5 on their relevance to the title. To ensure the data was representative we chose titles composed of contrasting themes and fabulas (such as taking a fabula built on the content of 'factory' and theming it with the theme of 'family') as well as well regular or complimentary theme and fabula pairings. We also included titles that included more then one theme in seperate tests.

In order to make the evaluation fair we presented the single image text first (so participants would not already have associated them with a group). The images on the single image test were also randomly shuffled and for the group tests we randomised the order in which sets appeared. We also added a restriction on image groups that no more than one image would be allowed per Flickr author - this is because image sets published by a single author are often taken as a part of a set and have naturally flow and would artificially seem to be stronger montages. Finally users were only allowed to take the evaluation once, a unique evaluation link for each user was given out per email address.

Each test contained two titles composed of different subjects and themes from the four the TMB prototype was able to use, in each test one title paired the theme with a complementing fabula, the other title paired the theme with a contrasting fabula to observe performance under different conditions. The titles chosen for single themes were London in Winter, Celebration and Earthquake, Spring Picnic, and Family Factory and for multiple themes My Family in New York at Winter, and Celebrating the New House in Spring.

Our pilot study was performed with 22 users. While this is a relatively low number of people it still gave us a large amount of data, as each user was asked to rate 40 images and 4 groups for each of the 4 sources. This resulted in 880 data points for single images and 88 for groups, enough for early indications of quantitive significance to emerge (which we measured with a t test). The pilot study was important to test the stability of the evaluation and also find if there were any significant improvements that needed to be made to the test. On completion of the pilot we found the test to be stable and only minor rewording of instructions to make the test easier to understand. The results of the pilot study are present in a published paper [19].

4.3 Results and Analysis

The full evaluation achieved 107 test subjects and shows some significant results. The mean rating of natoms from the TMB is higher then that for a keyword search (Flickr) in both single and group images. Figure 4.1 and Tables 4.1 and 4.2 show the data and t-tests for single images. Figure 4.2 and Tables 4.3 and 4.4 show the data and t-tests for grouped images. The hypothesis that the TMB selects natoms more relevant to the title then a keyword search is true with less then 0.0005 percent probability of error for both group and single images.

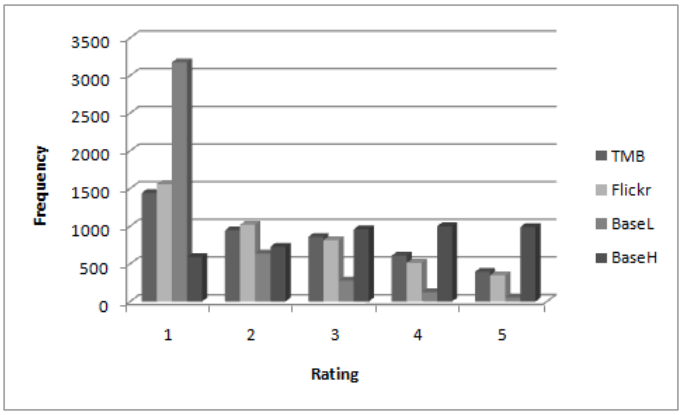

FIGURE 4.1: Single Image Rating Frequency

Set	1	2	3	4	5	Total
TMB	1437	944	857	609	393	4240
Flickr	1558	1019	812	513	346	4248
BaseL	3176	638	277	122	53	4266
BaseH	588	725	957	998	988	4256

TABLE 4.1: Single Images Rating Frequency

At first glance the difference between the TMB and Flickr only appears to be slight however it must be seen in the context of the difference in results between a best case scenario (human selection: BaseH) and a worst case scenario (random selection: BaseL).

Set	Mean	SD	Variance
TMB	2.428	1.329	1.766
Flickr	2.310	1.295	1.678
BaseL	1.414	0.832	0.693
BaseH	3.252	1.350	1.822

t=4.227, df=8486, p=0.0005

TABLE 4.2: Single Images Rating Satistics

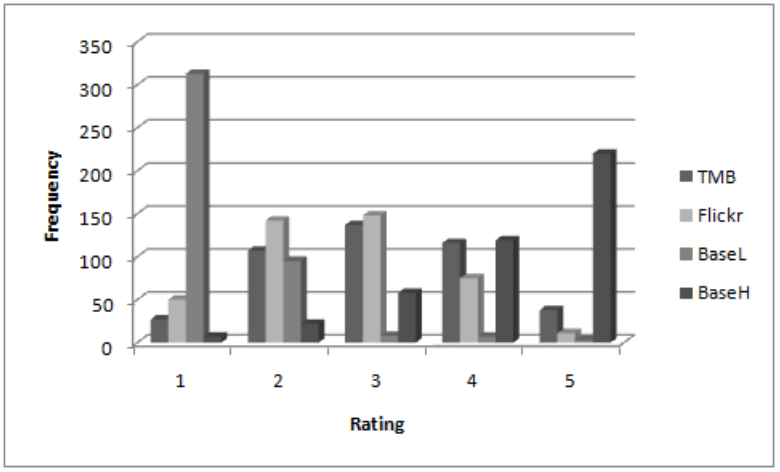

FIGURE 4.2: Grouped Image Rating Frequency

Set	1	2	3	4	5	Total
TMB	27	106	135	116	37	421
Flickr	50	141	147	73	11	422
BaseL	311	93	8	7	3	422
BaseH	7	22	57	119	217	422

TABLE 4.3: Grouped Images Rating Frequency

Set	Mean	SD	Variance
TMB	3.071	1.064	1.133
Flickr	2.654	0.983	0.967
BaseL	1.336	0.668	0.447
BaseH	4.225	0.979	0.958

t=5.902, df=841, p=0.0005

TABLE 4.4: Grouped Images Rating Satistics

Figures 4.3 and 4.4 show the relevant means and standard deviations in a way that they can be compared. These ranges are rather smaller than we might expect, and in this context the improvement given by TMB is rather more impressive.

As expected the results also show that the TMB proves better in a montage context with significance where it can build themes over a group of natoms, a t-test shows this hypothesis to be true with only a 0.0005 percent probability of error. In addition the data shown in table 4.5 reveals that while both a keyword search and TMB improved

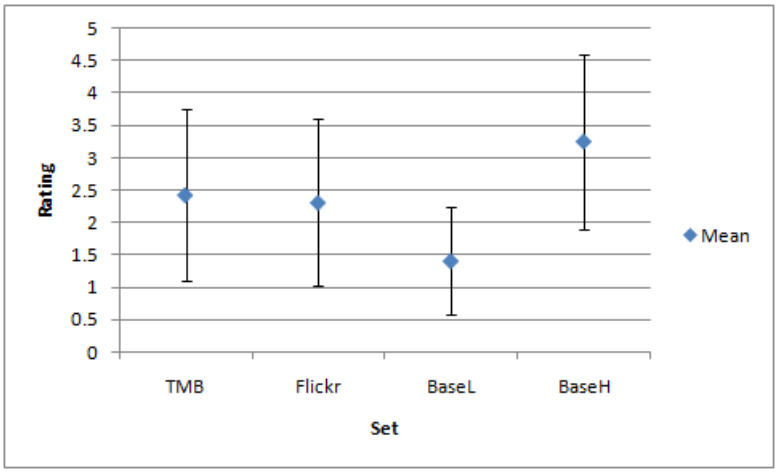

FIGURE 4.3: Single Image Mean and Std. Dev.

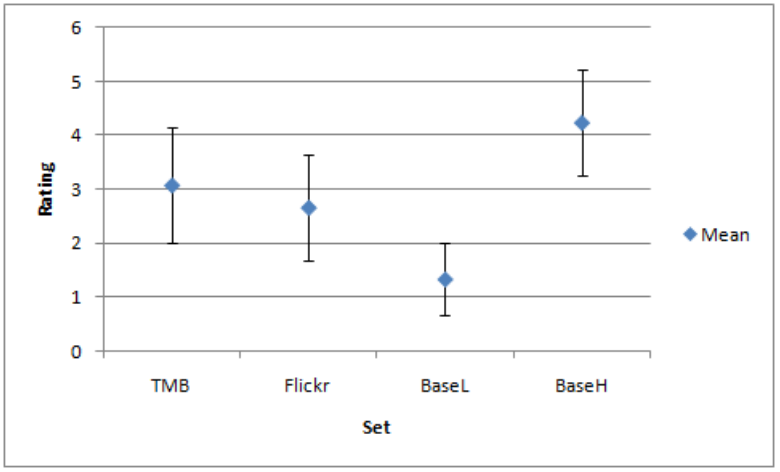

FIGURE 4.4: Grouped Image Mean and Std. Dev.

when their natoms were presented as a group the TMBs improvement was much more significant, the hypothesis that the TMBs improvement was greater then the improvement of a keyword search in a group context is shown with this data to be true according to a t-test with less then 0.0005 percent probability of error.

Set	Mean	SD	Variance
TMB	0.650	1.420	2.018
Flickr	0.356	1.447	2.095

t=9.328, df=8486, p=0.0005

TABLE 4.5: Grouped Images Improvement Statistics

As explained in the previous section we also wanted to observe how the TMB performed under a range of situations so deliberately included titles that had contradictory theme fabula pairings as well as titles which included multiple themes. Tables 4.6 and 4.7 and figures 4.5 and 4.6 show the contrast of results results for single and grouped images respectively between titles with multiple themes and those with just one theme where as tables 4.8 and 4.9 and figures 4.7 and 4.8 show the contrast of results for single and

grouped images between titles with contradictory theme fabula pairings and regular pairings.

Set	Mean	SD	Variance
TMB Single Theme	2.456	1.399	1.958
Flickr Single Theme	2.496	1.337	1.790
BaseL Single Theme	1.404	0.842	0.710
BaseH Single Theme	3.235	1.427	2.037
TMB Multiple Theme	2.399	1.253	1.571
Flickr Multiple Theme	2.122	1.223	1.496
BaseL Multiple Theme	1.425	0.822	0.676
BaseH Multiple Theme	3.268	1.267	1.606

TABLE 4.6: Single Images Single/Multiple Themes in Title Contrast Statistics

Set	Mean	SD	Variance
TMB Single Theme	2.981	1.135	1.288
Flickr Single Theme	2.849	0.971	0.943
BaseL Single Theme	1.292	0.659	0.435
BaseH Single Theme	4.037	1.109	1.230
TMB Multiple Theme	3.164	0.983	0.968
Flickr Multiple Theme	2.471	0.962	0.926
BaseL Multiple Theme	1.383	0.673	0.453
BaseH Multiple Theme	4.415	0.787	0.619

TABLE 4.7: Grouped Images Single/Multiple Themes in Title Contrast Statistics

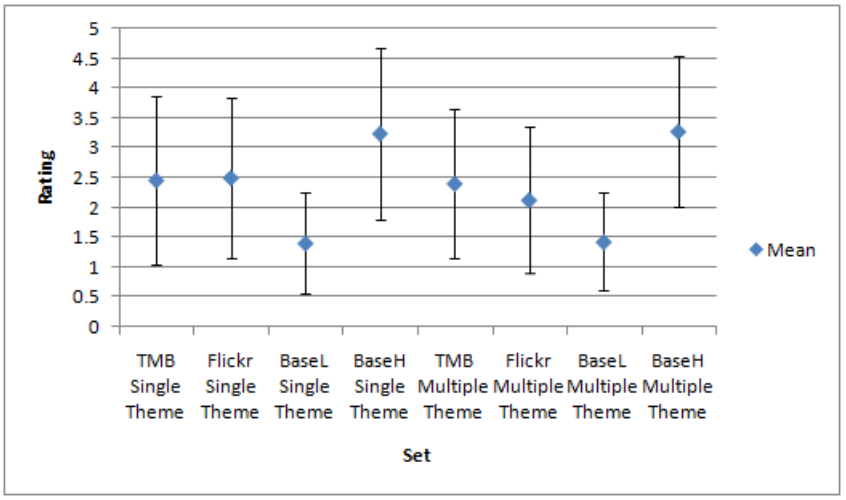

FIGURE 4.5: Single Images Single/Multiple Themes in Title Contrast

The results show that in both single and grouped images Flickrs keyword search always performed worse in titles with multiple themes where as the TMB only performed worse on single images and that TMB performed better then Flickr in all cases, multiple themes included. The results also show that in both single and grouped images Flickr performs better then TMB for contradictory theme fabula pairings and TMB performs better then Flickr for regular pairings. What this tells us is that the TMB is better at consolidating

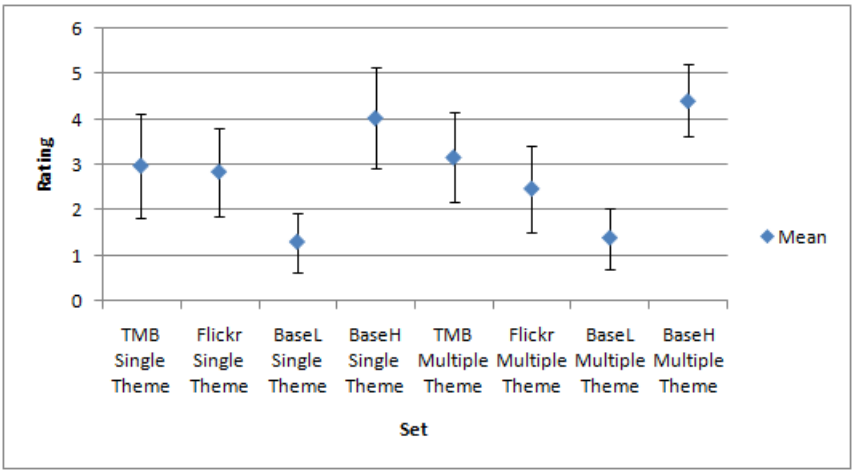

FIGURE 4.6: Grouped Images Single/Multiple Themes in Title Contrast

Set	Mean	SD	Variance
TMB Regular Theme Fabula Pairing	2.563	1.336	1.785
Flickr Regular Theme Fabula Pairing	2.258	1.282	1.644
BaseL Regular Theme Fabula Pairing	1.409	0.822	0.675
BaseH Regular Theme Fabula Pairing	3.546	1.265	1.601
TMB Contradictory Theme Fabula Pairing	2.028	1.224	1.499
Flickr Contradictory Theme Fabula Pairing	2.464	1.322	1.749
BaseL Contradictory Theme Fabula Pairing	1.429	0.864	0.746
BaseH Contradictory Theme Fabula Pairing	2.370	1.204	1.451

TABLE 4.8: Single Images Contradictory/Regular Theme Fabula Pairing in Title Contrast Statistics

Set	Mean	SD	Variance
TMB Regular Theme Fabula Pairing	3.332	0.988	0.977
Flickr Regular Theme Fabula Pairing	2.578	0.966	0.934
BaseL Regular Theme Fabula Pairing	1.337	0.656	0.431
BaseH Regular Theme Fabula Pairing	4.534	0.737	0.544
TMB Contradictory Theme Fabula Pairing	2.292	0.894	0.799
Flickr Contradictory Theme Fabula Pairing	2.905	1.001	1.001
BaseL Contradictory Theme Fabula Pairing	1.339	0.702	0.493
BaseH Contradictory Theme Fabula Pairing	3.301	1.034	1.069

TABLE 4.9: Grouped Images Contradictory/Regular Theme Fabula Pairing in Title Contrast Statistics

a wide variety of themes then simple keyword search, this can be attributed to the way the TMB consolidates a thematic request into a single list of relevant features, while keyword search might look for each theme separately the TMB will scores natoms with both themes present highly. That the TMB performed worse for contradictory pairings is also no surprise, because of the way this experiment has been designed with the TMB scoring a finite fabula its possible that having built the Fabula for a specific content that it would not contain any features to use in connoting a contradictory theme where

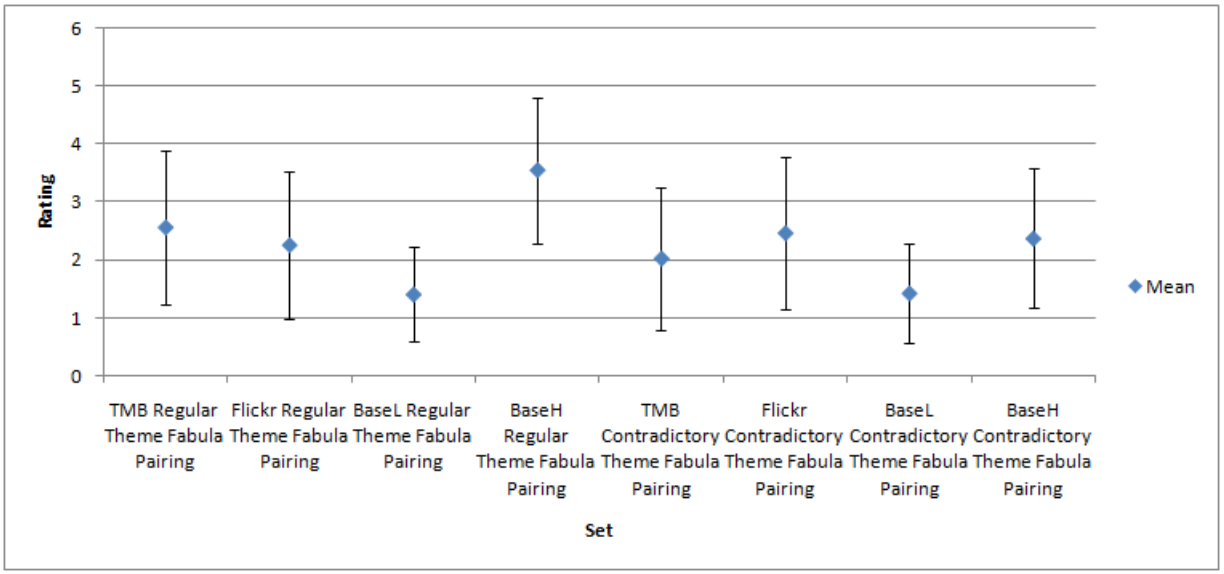

FIGURE 4.7: Single Images Contradictory/Regular Theme Fabula Pairing in Title Contrast

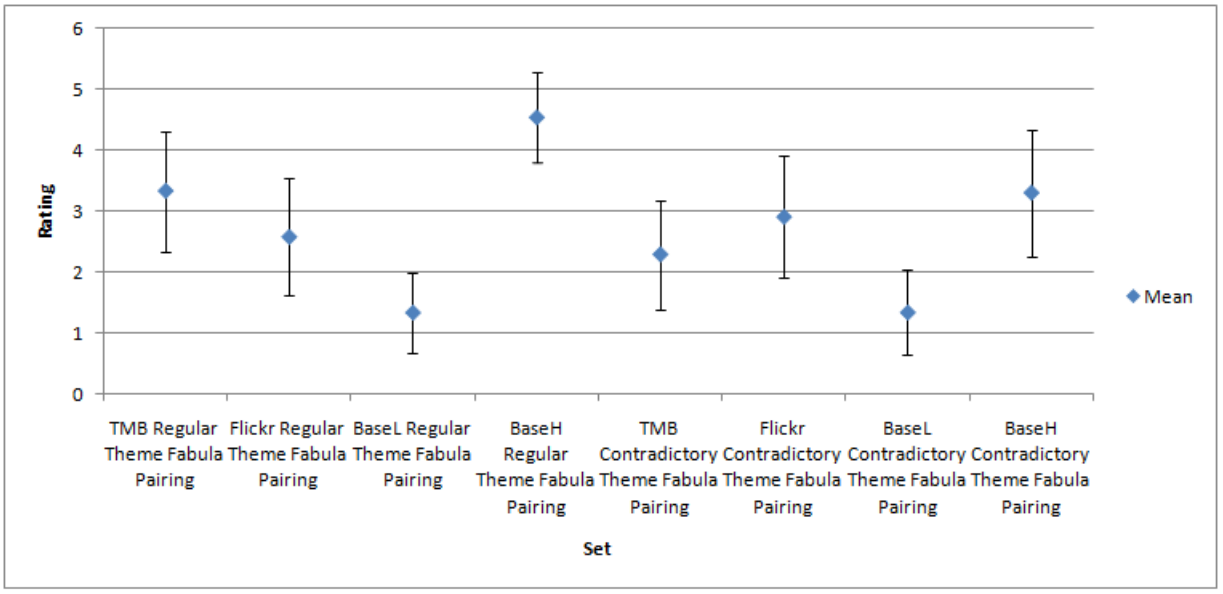

FIGURE 4.8: Grouped Images Contradictory/Regular Theme Fabula Pairing in Title Contrast

as the keyword search retrieves each word separately from a much wider pool of images and as such can pick images for each idea.

We also observed correlations between the ratings of different images and how these images rated on the different metrics for thematic quality we were measuring for them. This was so we could observe any strong correlations between high scoring images and particular metrics so that we might use these metrics to refine our method of calculating thematic quality which is currently very simple. However we observed no strong correlations between the scoring of natoms and the other metrics. However this could be attributed to the design of our evaluation, by using only the images that had been

scored and selected using one metric these gave us relatively small samples for which to look for correlations. Many of the individual metrics themselves (such as match and focus) were not designed to be used by themselves and are calculated on the relationship between a specific natom and thematic element. As such we were looking for correlations in samples of 10 and as such could not expect to find strong correlations. Future experiments seeking to refine the method of calculating thematic quality should seek a direct comparison between methods that use different metrics to avoid this problem, rather then using one and looking for correlations with others.

These results offer encouraging observations. The TMB seems to be performing better then a keyword search with some significance and further more it seems the TMB is very strong within a group context, this could lead us to believe it could perform similarly strongly within a narrative context. It is also encouraging to see the TMB is able to consolidate multiple themes as well although it is an accepted constraint that such a system performs less strongly with contradictory theme fabula content pairings. Although the third evaluation objective was not fulfilled due to a failing in the experiment design the process has made clear that a future experiment should seek a more direct comparison in order to refine the way thematic quality is measured.

Chapter 5

Investigation into a Thematic Integration

The thematic model has the potential to improve narrative generation by enriching its results with themes. However, as was explored in [18], how it should be integrated with existing methods of narrative generation is unclear. Also, which approach of narrative generation does a thematic solution best lend itself to?

Referring back to the division of narrative generation illustrated in figure 2.2, we can explore the possibility of a thematic systems involvement at different levels of narrative generation. Themes are intangible concepts, a subtext rather then a core focus of the narrative, and for this reason it seems at first that narrative generation would benefit from thematic involvement at the presentation level. Here themes could be connoted by emphasis given through the presentation to the features within the narrative that denote motifs which in turn connote the desired themes. At this level a thematic subtext would become present through elaboration on the presentation of the plot. However this is a process that could potentially fail if there were no relevant features present in the narrative that could be elaborated upon to help connote the desired theme, the system might find that at the presentation level a thematic system might only be able to offer from a subset of themes.

At the story level of narrative generation a thematic systems involvement would be in some ways the opposite of its involvement at a presentation level. Instead of offering elaboration on existing narrative features at the story level a thematic system would generate additional narrative elements based on a shopping list of required motifs for the desired themes. This way themes would become apparent through the presence of certain story elements that connoted the desired themes. This could potentially fail however if the systems plot generation did not make use of the thematic story elements or they were not properly exposed possibly leading to absence of key motifs. Also, such an approach could damage the generated narrative more then help it potentially

flooding the system with elements irrelevant to the plot. For some systems as well story generation integration is not always an option, at least on a fully autonomous level, with many systems generating plot out of pre written and defined story elements. These semi automatic approaches to story generation would require a very different approach perhaps with thematic guidance on the creation of these elements as supposed to influencing the automatic generation process in others.

At the plot generation level thematics could play a role in the story selection of narrative elements as well as the way relationships build. The first part of this would be similar to how the TMB prototype builds photo montages in that a list of desired motifs would be compiled and this would be used to thematicly score potential story elements and influence their selection and inclusion in the plot. Also, the relationships between story elements and actions of elements could in turn be factored in as features that denote motifs, as such potential actions at the plot generation of the story could be thematicly scored influencing what occurs. For example a story in which violence is a desired theme might see the protagonist kill the antagonist rather then banish or imprison them. However, like inclusion at the story level it is possible that heavy thematic involvement could damage the plot itself, making its involvement a dangerous balancing act, potentially forcing plot actions that damage the quality of the narrative. Furthermore like involvement at the presentation level, a lack of complete control over the story elements could potentially restrict available themes.

The question of approach is a very different one to that of level of involvement, in that feasibility becomes an important issue as well as potential benefit. The current basic implementation of the thematic model in the TMB is quite different to most forms of narrative generation. For a successful integration we would need to closely inspect what would need to change.

Character centric approaches are perhaps the most different from the current implementation using the thematic model. Rather then the natoms of narrative segments that our model describes, character centric approaches start by simulating the content of the narrative itself, modeling characters, locations, entities, and events. However, through the process of plot and presentation, character centric approaches still go on to generate natoms that contain features and by extension denote motifs. An integration would have to seek to ensure that certain features were planted in order for the themes to become apparent in the finished narrative. To do this involvement at the story level seems obvious as this is where the elements present within the narrative are generated. However, as a more semi automatic approach with predefined elements is more common then automatic generation at this level in character centric approaches it could be difficult to integrate a thematic approach with the prewritten characters. At the plot and presentation levels an integration seems more possible, potential character actions and story events can be thematicly scored to influence actions taken to be conducive with desired themes and then presented in a way that emphasises the relevant thematic

content. Character centric generations frequent use of game engines means that integration at the presentation level may be easier where knowledge of the entities present in a particular scene is much more exact then in natural language. However as already discussed a reliance on integration at these levels potentially limits the available themes.

Author centric approaches are more similar in process to the current implementation of the thematic approach in that they're heavily based on structures and largely concerned with the authoring process rather then modeling the content of the narrative. The story generation process for some is about composing a pool from large collections of potential natoms, often from the Web, based on their relevance to required parts of the narrative structure. This is very similar to the way the TMB currently puts together selections for montage, and it is easy to see that with author centric projects that work this way thematic integration would be a relatively simple process of scoring potential segments to generate. At the plot level, integration could be a similar process to the integration that would be used with character centric approaches, in that elements selected for exposure could be chosen on their thematic qualities rather then only narrative ones. However, for rule based systems of plot generation thematic rules would need to be written for the system. The feasibility of this would need to be added on a system by system basis. At a presentation level natural language generation poses difficulties for thematic integration as a full lexicon for desired features would need to be developed and integrated with the system. Forcing it to uses a small subset of words might make the language clumsy and its important to remember that the thematic model was created with the theory of structuring a narrative in mind where as the structure of individual pieces of language is very different. However, for those systems that use templating or selected pre authored text presentation, using thematics becomes more feasible where narrative techniques such as emphasis (spatially or visually) can be used to highlight relevant segments to help connote a theme.

The possibilities apparent from this investigation into how a thematic integration with a narrative system could be attempted are summarised in the table in figure 5.1. Decisions and selections made in generation may be influenced thematically by making the objective of the decision thematic as well as for plot objectives. Further thematic integration can be achieved through emphasis at the presentation or plot levels and other presentation choices such as style may have an influence that could be worked in favor of desired themes.

		Narrative			
		Story		Discourse	
				Plot	Presentation
		Semi Auto	Auto		
Generation Approach	***Character Centric***	*Uses Prewritten Characters.* Could be supported by thematic writing tools.	*Goals and Rules are generated.* Generation process influenced thematically.	*Characters and Elements Selected. Rules applied. Actions chosen.* Thematically influenced selection of characters and elements as well as actions chosen.	*Generation of media. Revealed narrative chosen.* Emphasis of relevant elements and relationships. Choice of style.
	Compromise	*Uses Prewritten Characters and purpose built Director Agent or an author rule set.* Could be supported by thematic writing tools and have a thematically influenced author/director.	*Both characters and author goals and rules generated.* Goal and rule generation thematically influenced.	*Characters and Elements Selected. Rules applied. Actions and Events chosen. All Influenced and vetted by Director Agent or Author rule set.* Thematically influenced selections of elements and actions. Thematic vetting by director/author.	
	Author Centric	*Material generated or collected.* Thematically influence which resources collected/generated.		*Populating/Building Structure.* Influence selection and positioning of resources. Influence structure built.	

FIGURE 5.1: Summary of potential integrations of a thematic system with narrative generation

Chapter 6

Conclusions and Future Work

6.1 Conclusions

Themes are an important part of narratives, they allow an author to give a cohesive thematic subtext to a story that ties its elements together and engage the audience with the underlying concepts and messages of the narrative beyond the individual actions and events that comprise the plot. By introducing themes to generated narratives we are able to portray concepts and subtle story elements beyond what the basic construction of the plot allows, leading to a richer story experience. Similar themes can be used to insure a thematic cohesion in other representations of information along with other narrative devices leading to more engaging and relevant search results.

In this report a thematic model has been presented, instances of which could be used as representations of various themes which a system could use to imbed a thematic subtext in presentations of information. Instances of the model are currently built manually out of the concepts of themes, motifs, and features. Work is underway to formalise a method for authoring themes and the early stages of expert review are leading to a set of emerging rules based on semiotics.

The effectiveness of an instance of this model was analysed by creating a prototype capable of generating themed photo montages. An evaluation experiment aimed to ascertain if the results of the prototype accurately represented the desired theme and if the results connoted the desired themes any better then results ascertained by including the theme names in a simple keyword search. The experiment also aimed to find correlations between different thematic metric and the quality of montages so that a more refined metrics for thematic quality could be developed.

The results show that collections generated using a system utilising the model are more relevant to their given requests then those using just keyword search by a significant degree. The results also show that the thematic system also faired better in the narrative

context of a montage. This is demonstrated not only by the thematic system receiving higher evaluations in the more narrative context of a montage but also in that the improvement from presenting images for a set as a montage rather then individually that was experienced by both methods was much greater for the thematic system, highlighting the benefit of such an approach in a narrative environment. The results also show that the thematic system was better able to consolidate titles which contained multiple themes maintaining a stronger thematic cohesion which it likely owes to the process of identifying common features in multiple theme requests. The results did however highlight that the system showed weaker performance under titles with contrasting theme fabula pairings, this is due to the way the system will only return items in its fabula which were selected for its fabulas subject not the titles thematic element. This lead to the system being unable to return a thematicly relevant montage. Also no significant correlations were found in the results between natom quality and different thematic metrics due to limitations in the experiment itself, further experimentation is necessary in order to measure the effectiveness of different metrics which would ideally look at a more direct comparison rather then correlations. However the experiments did fulfill its other objectives demonstrating thats the thematic model could be used to ensure montages successfully connoted desired themes, were more relevant then those generated from a simple keyword search, and that it performed well in a narrative context, therefore a thematic system could be used the improve the thematic cohesion in narrative systems.

Having shown that the thematic model could be accurately used to attach a thematic subtext to a narrative an investigation was made into how a thematic system could be integrated with a narrative system. It was assessed that a thematic system could be integrated with many different styles of narrative generation and at different levels but that each integration would have different limitations and advantages. Integration at the story and plot levels while potentially effective comes with several disadvantages in that as the narrative is still being formed its possible thematic influence at this level could lead to the introduction of inappropriate story elements that damage the plot, also the semi automatic approach of many systems would make thematic influence in system where for example pre written characters are present difficult.

Integration at the presentation level avoids these problems because the plot is already in place and thematic integration in the presentation of the narrative as media should be possible in all systems. Thematic representation also ensures that the thematic influence will keep its subtlety as it centers around the elaboration of existing thematic elements rather then the introduction of new ones. However, integration at the presentation level is limited by existing elements within the narrative, much like the problems experienced by the TMB with contrasting theme fabula pairings it would be impossible to create a thematic underpinning for a theme in a narrative that contained no relevant features. As a result a thematic integration at the presentation level would require thematic analysis as well as generation in order to detect which themes exist in potential within

the narrative.

Having observed the success of the thematic model at generating montages that successfully connote desired themes and investigating potential solutions for integrating the thematic model with a narrative system there are clear advantages to making a thematic integration with the presentation level of a narrative system the next step. This would take the form of a thematic representer that could take a narrative and as part of its presentation elaborate desired themes within the narrative. Tying this in with existing work it could take the form of an automatic illustrator that strengthened the thematic cohesion through relevant images but emphasis of the desired thematic features could be achieved through a variety of techniques.

6.2 Hypothesis

I have resolved the following research hypothesis:

A thematic model that identifies the themes in a piece of information and then emphasises them can make search results more relevant within a specified context and improve the thematic cohesion of narratives.

This can be broken down into specific questions:

1. Can the thematic model be used to identify themes within a piece of information?
2. Can the thematic model be used to elaborate desired themes?
3. Can this make search results more relevant?
4. Can this improve the thematic cohesion of narratives?

Existing contributions in the form of the experiment using the TMB have demonstrated positive answers to items two and three. Where a prototype using the model was shown to be evaluated as presenting montages that were more relevant to specific titles then those created with a simple keyword search. These montages we built from fabulas on a specific topic containing a variety of themes, the desired ones of which were elaborated through their selection to the montage. Although it is also possible that the thematic model might allow elaboration in alot of other ways rather then just selection and further experiments may be able to demonstrate this.

With regards to item one existing experiments have worked on the assumption a given theme is present within a fabula and poor results ensue if it is not, no work as of yet has been completed on actually identifying present themes within a piece. A thematic

analyser that could identify the themes of a narrative, perhaps using co-occurrence keyword extraction such as in [26] could demonstrate this ability if the keywords extracted matched those detected by human analysis or expert review to a sensible degree.

To answer item four requires we show that narratives, such as the results of narrative generation, can have a thematic subtext improved or added to them which might work towards combating the failings in many methods of narrative generation as well as enriching other narrative representations of information. To demonstrate this requires a prototype that can take short narratives which have their existing themes analysed and then desired ones elaborated which then could be used as the subject of an experiment which demonstrated the elaborated stories as having stronger identifiable themes in those desired and demonstrate that these had a positive effect on the narrative.

6.3 Future Work

The following plan aims to reach an evaluation of my hypothesis within the remaining 14 months. Figure 6.1 presents a gantt chart for the following plan.

6.3.1 Plan

6.3.1.1 Work Package 1: Formalise authoring process and develop new metrics

Objectives

- Develop new improved candidate methods for measuring thematic quality
- Complete formalising method for authoring themes
- Evaluate method for authoring themes and new measures of thematic quality in small experiment

Description: This package represents the completion of work around the thematic model itself. Continued expert reviews will allow a forming of a formalised set of rules for authoring themes, this set of rules can then be evaluated by allowing test subjects to author specific themes to create new instances of the model which can in turn be evaluated in an experiment similar to the previous one where themes from multiple authors are compared in order to measure the consistency of the authoring method. Such an experiment could also be used as an opportunity to compare the effect of different methods of measuring thematic quality.
Deliverables: Journal article on thematic model and formalised and evaluated method

of authoring instances of the thematic model.
Estimated Time: 4 Months

6.3.1.2 Work Package 2: Develop Thematic Analyser

Objectives

- Develop a keyword extraction tool
- Develop a system capable of identifying themes within a narrative from extracted keywords

Description: This package is key to item 1 of the objectives of my research hypothesis. Although it has been shown that the thematic model can be used to generate thematic content no experimentation has yet been done on using it in reverse for thematic analysis. In theory this should be possible using keyword extraction such as presented in [26] where extracted keywords could be treated as features and used to build the motifs and themes of a given narrative. The problem with this approach is the instance of the model would need to be broad enough in order to detect all the possible themes in a given narrative. This large instance of the model will be a side effect of work package 1 where test subjects will have authored a large number of themes using the new formalised rules for evaluation. Having developed keyword extraction and integrated it with thematic scoring from earlier prototypes a new thematic analyser prototype should be able to accurately identify themes in a given narrative and its effectiveness may be evaluated by comparing its detected themes of a specific narrative to those listed by human review.
Deliverables: Prototype thematic analyser and a conference paper on thematic analysis.
Estimated Time: 3 Months

6.3.1.3 Work Package 3: Develop and Evaluate Thematic Presenter

Objectives

- Integrate thematic analyser and earlier work on scoring thematic quality
- Develop a new prototype capable of illustrating narratives to elaborate desired identified themes
- Add other thematic methods of elaboration such as emphasis

- Evaluation experiment that demonstrates that the presenter, using the thematic model, can improve the thematic cohesion of given narratives

Description: In order to measure the thematic models effect on narratives and narrative generation an integration at the presentation level will be pursued. Not only does an integration at this level provide perhaps the most significant effect based on our investigation but should also be relevant to all approaches of narrative generation as well as other narrative representations of information. Developing the final thematic illustrator will largely be a process of bringing together earlier prototypes. The thematic analyser combined with the TMB will allow the system to work as a kind of automatic illustrator attaching images to a narrative relevant to desired themes in order to elaborate them. Further narratological devices can be used to enhance this elaboration by using emphasis on parts of the narrative relevant to specific themes and sequencing of illustrations in order to have the most profound effect. Having developed such a prototype it may be evaluated by conducting a user study that compares plain narratives and those represented using the prototype. Test subjects will be asked to identify the themes in each and the strength of the presentation of these themes, it is our prediction that narratives represented will have a more focused set of themes, with greater thematic cohesion, and significantly stronger presentation of the desired themes.
Deliverables: Completed thematic presenter prototype and journal paper on effectiveness of thematic presentation.
Estimated Time: 6 Months

6.3.1.4 Work Package 4: Write up Final Thesis

Objectives

- Complete PhD with write up of thesis

Description: Having completed the experiments a thesis will be written to detail my findings on the advantages of the thematic model.
Deliverables: Thesis.
Estimated Time: 5 Months

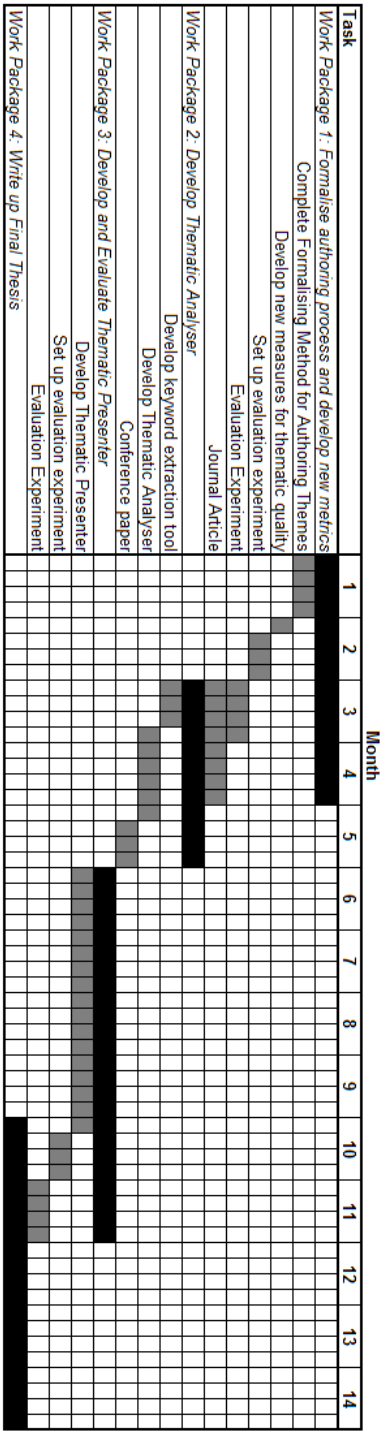

FIGURE 6.1: Gantt chart

Bibliography

[1] H. Al-Khalifa and H. Davis. Folksonomies versus automatic keyword extraction: An empirical study. *IADIS International Journal On Computer Science And Information Systems (IJCSIS)*, 1:132–143, 2006.

[2] H. Alani, S. Kim, D. Millard, M. Weal, W. Hall, P. Lewis, and N. Shadbolt. Automatic ontology-based knowledge extraction and tailored biography generation from the web. *IEEE Intelligent Systems*, 18:14–21, 2003.

[3] M. Alberink, L. Rutledge, and M. Veenstra. Sequence and emphasis in automated domain-independent discourse generation. In *Information systems*, pages 1–10, 2003.

[4] M. Bal. *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, January 1998.

[5] R. Barthes. *Mythologies*. Editions du Seuil, 1957.

[6] R. Barthes and L. Duisit. An introduction to the structural analysis of narrative. *New Literary History*, 6:237–272, 1975.

[7] H. Berghel. Cyberspace 2000: dealing with information overload. *Commun. ACM*, 40(2):19–24, 1997.

[8] M. Bernstein. Card shark and thespis: exotic tools for hypertext narrative. In *Proceedings of the twelfth ACM conference on Hypertext and Hypermedia*, 2001.

[9] W. Booth. *The Rhetoric of Fiction*, chapter Types of Narration, pages 69–74. University of Chicago Press, 1974.

[10] P. D. Bra, A. Aerts, B. Berden, B. de Lange, B. Rousseau, T. Santic, D. Smits, and N. Stash. Aha! the adaptive hypermedia architecture. In *Proceedings of the fourteenth ACM conference on Hypertext and hypermedia*, pages 81–84, 2003.

[11] P. D. Bra, P. Brusilovsky, and G.-J. Houben. Adaptive hypermedia: from systems to framework. In *ACM Computing Surveys (CSUR)*, volume 31, 1999.

[12] C. Buckley. Automatic query expansion using smart : Trec 3. In *In Proceedings of The third Text REtrieval Conference (TREC-3)*, pages 69–80, 1995.

[13] C. Carpineto, R. de Mori, G. Romano, and B. Bigi. An information-theoretic approach to automatic query expansion. *ACM Trans. Inf. Syst.*, 19(1):1–27, 2001.

[14] 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.

[15] D. M. Frohlich, D. Rachovides, K. Riga, R. Bhat, M. Frank, E. Edirisinghe, D. Wickramanayaka, M. Jones, and W. Harwood. Storybank: mobile digital storytelling in a development context. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 1761–1770, New York, NY, USA, 2009. ACM.

[16] G. Fu, C. B. Jones, and A. I. Abdelmoty. Ontology-based spatial query expansion in information retrieval. In *In Lecture Notes in Computer Science, Volume 3761, On the Move to Meaningful Internet Systems: ODBASE 2005*, pages 1466–1482, 2005.

[17] C. Hargood, D. Millard, and M. Weal. A thematic approach to emerging narrative structure. In *Web Science at Hypertext08*, 2008.

[18] C. Hargood, D. Millard, and M. Weal. Investigating a thematic approach to narrative generation. In *DAH at Hypertext 09*, 2009.

[19] C. Hargood, D. Millard, and M. Weal. Using a thematic model to enrich photo montages. In *Proceedings of Hypertext 09*, 2009.

[20] E. Hovy. Automated discourse generation using discourse structure relations. *Artificial Inteligence*, 63:341–385, 1993.

[21] M. Jewell, F. Lawrence, and M. Tufield. Ontomedia: An ontology for the representation of hetrogeneous media. In *Multimedia Information Retrieval Workshop*, 2005.

[22] F. Lawrence, M. Tuffield, M. Jewell, A. Prugal-bennett, D. Millard, M. Nixon, monica schraefel, and N. Shadbolt. Ontomedia - creating an ontology for marking up the contents of heterogeneous media. 2005.

[23] M. Lebowitz. Planning stories. In *Program of the Ninth Annual Conference of the Cognitive Science Society*, volume 9, pages 234–241, Seattle, Washington, 1987. Cognitive Science Scoiety Conference.

[24] N. Maria and M. J. Silva. Theme-based retrieval of web news. In *Selected papers from the Third International Workshop WebDB 2000 on The World Wide Web and Databases*, pages 26–37, London, UK, 2001. Springer-Verlag.

[25] M. Mateas and A. Stern. Faade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference*, 2003.

[26] Y. Matsuo and M. Ishizuka. Keyword extraction from a single document using word co-occurrence statistical information, 2003.

[27] M. McQuillan. *The Narrative Reader*. Routledge, London, 2000.

[28] 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.

[29] D. Michaelides, D. Millard, M. Weal, and D. D. Roure. Auld leaky: A contextual open hypermedia link server. In *Revised Papers from the nternational Workshops OHS-7, SC-3, and AH-3 on Hypermedia: Openness, Structural Awareness, and Adaptivity*, volume 2266, pages 59–70, 2001.

[30] D. Millard, C. Bailey, T. Brody, D. Dupplaw, W. Hall, S. Harris, K. Page, G. Power, and M. Weal. Hyperdoc: And adaptive narrative system for dynamic multimedia presentations. Technical report, ECS, University of Southampton, 2003.

[31] M. Mitra, A. Singhal, and C. Buckley. Improving automatic query expansion. In *SIGIR '98: Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 206–214, New York, NY, USA, 1998. ACM.

[32] F. Paradis and C. Berrut. Experiments with theme extraction in explanatory texts. In *In: Second International Conference on Conceptions of Library and Information Science, CoLIS*, pages 433–437, 1996.

[33] 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.

[34] F. Saussure, C. Bally, A. Sechehaye, and A. Riedlinger. *Course in General Linguistics*. McGraw-Hill, 1966.

[35] N. Sawhney, D. Balcom, and I. Smith. Hypercafe: narrative and aesthetic properties of hypervideo. In *Proceedings of the the seventh ACM conference on Hypertext*, pages 1–10, 1996.

[36] 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.

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

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

[39] 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.

[40] R. Young. Story and discourse: A bipartite model of narrative generation in virtual worlds. *Virtual Reality*, 2003.

[41] R. M. Young, M. E. Pollack, and J. D. Moore. Decomposition and causality in partial-order planning, 1994.

[42] X. S. Zhou and T. S. Huang. Unifying keywords and visual contents in image retrieval. *IEEE MultiMedia*, 9(2):23–33, 2002.