From the Fleece of Fact to Narrative Yarns: a Computational Model of Composition

Pablo Gervás

Universidad Complutense de Madrid
pgervas@sip.ucm.es

Abstract

From a given observable set of events, a large number of stories may be composed, by deciding to select or omit specific events, by restricting attention to smaller subsets of the overall setting, by focusing on particular characters, or by narrating the chosen events in different order. This particular task of narrative composition is not covered by existing models of storytelling or cognitive accounts of the writing task. This paper presents a model of the task of narrative composition as a set of operations that need to be carried out to obtain a span of narrative text from a set of events that inspire the narration. To provide guidance in structuring the task, an analogy is drawn between the narrative composition task and that of manufacturing textile fibres, with corresponding concepts of heckling the original material into fibres, then twisting these fibres into richer and better yarns. The model explores a set of intermediate representations required to capture the structure that is progressively imposed on the material, and connects this content planning task with a classic pipeline for natural language generation. As an indicative case study, an initial implementation of the model is applied to a chess game understood as a formalised set of events susceptible of story-like interpretations. The relationships between this model and existing models from other fields (narratogical studies, cognitive accounts of writing, AI models of story generation, and natural language generation architectures) is discussed.

Keywords: Narrative generation, Theory of narrative, Representations, Natural language processing, Artificial intelligence, Cognitive science, Narratology.

1. Introduction

The task of composing a narrative based on a given set of events that have taken place has received little attention in terms of conceptual modelling. Efforts have been made to capture the structure of narratives as a finished product (by the narratology research community), to come up with a set of cognitive processes implied in the tasks of writing in general or of understanding narrative in particular (by the cognitive science community), to build models of how fictional plots are generated from scratch or of how discourse may be structured for a given plot (by the artificial intelligence community) and to construct functional architectures for generating text from conceptual data (by the natural language generation community). The task of putting together a narrative that conveys events that have already happened is related to all these aspects. It is also the kind of basic storytelling that people carry out in their everyday lives to communicate with one another, to convince, to inform, to remember the past, to interpret the present and to plan for the future.

Producing a model of this task is a significant challenge for various reasons. First, due to its complexity. The fact that it involves elements from various other tasks requires a certain familiarity with the different phenomena involved. The final product will need to have narrative structure, and the implied processes should at least be cognitively plausible. Second, because as a task it involves not just modelling the particular product resulting from the process (narrative structure), or of the processes themselves (narrative composition) but also requires a model of the input (the set of facts observed / remembered that constitute the source and the starting point for the composition). This element is different from the structures considered in cognitive accounts of narrative understanding (as it includes the events that may appear in the story but not necessarily the causal relations between them) and different from the set of events that are actually mentioned in existing stories (as the story will contain only a selection of all the events that actually happened, possibly filtered on the basis of a set of relevant causal relations postulated by the author). Such a representation of the input is implicit in the specification of the task, yet any attempt at computational modelling must start by representing it explicitly, as it will significantly influence the rest of the process.

The present paper attempts to address these challenges from an engineering point of view. The task is considered in terms of how its input and its expected output might be represented, and what processes might lead from one to the other. To manage the complexity of the problem, an analogy is established between the process of narrative composition and that of converting naturally occurring materials into lengths of line (whether yarns, thread or rope). This analogy is based on the similarities between the inputs (in both cases volumes of material of heterogeneous quality with no obvious unique linear presentation) and the outputs (for both a linear reorganization of selections of the same material, systematically structured to optimise certain desired qualities of the final result) of these processes. The analogy is discussed, and then further explored to provide a set of candidate subtasks for narrative composition that constitute a possible computational model for narrative composition. This analogy is taken in this paper to devise a computational model of the composition task which is then tested in a software implementation operating over a very simple representation of meaningful events involving a set of characters that interact over time in an elementary represented space. Chess provides a finite set of characters (pieces), a schematical representation of space (the board) and time (progressive turns), and a very restricted set of possible actions. Yet it also allows very elementary interpretations of game situations in terms of human concepts such as danger, threat, conflict, death, survival, victory or defeat, which can be seen as interesting building blocks for story construction.

The paper reviews existing models related to the narrative composition task, outlines and discusses the analogy between narrative composition and the spinning of textile yarns, describes the proposed computational model, presents the case study for narration of chess games and finishes with discussion and conclusions.

2. Previous Work

A number of models of related tasks and elements arising from different fields of research are reviewed in this section to provide background material for the discussion. Due to the breadth of fields under consideration, exhaustive review in any one of them is beyond the scope of the paper. An attempt has been made in each case to gather here the set of elementary concepts in each field that are relevant for the understanding of the arguments in the paper.

2.1. Narratology

According to many theorists, narrative has two components: what is told (what narrative is: its content, consisting of events, actions, time and location), and the way it is told (how the narrative is told: arrangement, emphasis / de-emphasis, magnification / diminution, of any of the elements of the content). These have been named different ways by different researchers, story and discourse, *histoire* and *discours*, *fabula* and *sujzet*. There are alternative analyses that postulate different subdivions. Even between theories that agree on having just two levels of analysis there seem to be many subtleties that cast doubt on whether the same thing is meant by the different words. This presents a serious obstacle for researchers from the computational field trying to address the treatment of stories in any form. In order to avoid ambiguity, we will restrict our analysis here to three levels of conceptual representation of a story, and refer to these as the *story* (the complete set of what could be told, organised in chronological order of occurrence), the *plot* (what has been chosen to tell, organised in the order in which it is to be told) and the *narrative* (the actual way of telling it).

Narratologists, who specialize in the study of narrative, consider the concept of *focalization* (Genette, 1980) as the way in which a narrator restricts what he is telling about a particular scene to what might have been perceived by someone present in that scene. This may be one of the characters if the scene is told in the first person, or the narrator himself as if he had been present (if the story is told in the third person). This has an interesting implication in the fact that, through focalization, narrative discourse (and thereby the structure of stories) is influenced by the perception of space: events that take place simultaneously in different locations that cannot be perceived at the same time (this may be different cities but also different neighbouring rooms separated by a wall) usually require different narrative threads.

2.2. Cognitive Accounts of Writing

Flower and Hayes (Flower and Hayes, 1981) define a cognitive model of writing in terms of three basic process: planning, translating these ideas into text, and reviewing the result with a view to improving it. These three processes are said to operate interactively, guided by a monitor that activates one or the other as needed. The planning process involves generating ideas, but also setting goals that can later be taken into account by al lthe other processes. The translating process involves putting ideas into words, and implies dealing with the restrictions and resources presented by the language to be employed. The reviewing process involves evaluating the text produced so far and revising it in accordance to the result of the evaluation. Flower and Hayes' model is oriented towards models of communicative composition (such writing essays or functional texts), and it has little to say about narrative in particular. Nevertheless, a computational model of narrative would be better if it can be understood in terms compatible with this cognitive model. An important feature to be considered is that the complete model is framed by what Flower and Hayes consider "the rhetorical problem", constituted by the rhetorical situation, the audience and the writers goals.

2.3. Cognitive Accounts of Narrative Comprehension

Although this paper is concerned with modelling the process of narrative composition, it is indirectly affected by models of narrative comprehension in as much as the results of composition must be suitable for comprehension. Narrative comprehension involves progresive enrichment of the mental representation of a text beyond its surface form by adding information obtained via inference, until a situation model (representation of the fragment of the world that the story is about) is constructed (van Dijk and Kintsch, 1983). A very relevant reference in this field is the work of (Trabasso et al., 1989), who postulate comprehension as the construction of a causal network by the provision by the user of causal relations between the different events of a story. This network representation determines the overall unity and coherence of the story.

2.4. Story Telling

Storytelling efforts in AI have focused on two different tasks: that of building fictional plots from scratch and that of structuring appropriate discourse for conveying a given plot

The importance of causal relations in narrative comprehension has led to AI models of plot generation that rely heavily on the concept of planning. Many existing storytelling systems feature a planning component of some kind, whether as a main module or as an auxiliary one. TALE-SPIN (Meehan, 1977), AUTHOR (Dehn, 1981), UNIVERSE (Lebowitz, 1983), MINSTREL (Turner, 1993) and Fabulist (Riedl and Young, 2010), all include some representation of goals and/or causality, though each of them uses it differently in the task of generating stories. An important insight resulting from this work (originally formulated by (Dehn, 1981) but later taken up by others) was the distinction between goals of the characters in the story or goals of the author.

A less frequently modelled aspect but also very relevant is emotion, which clearly plays a heavy role in the appreciation of narrative. The MEXICA storytelling system (Pérez y Pérez, 1999) takes into account emotional links and tensions between the characters as means for driving and evaluating ongoing stories. The system evaluates the quality of a partial draft for a story in terms of the the rising and falling shape of the arc of emotional tensions that can be computed from this information.

With respect to the task of building an appropriate discourse for rendering a given plot, significant efforts have been carried out in the domain of cinematic visual discourse. These include work on use of flashback and foreshadowing to produce surprise (Bae and Young, 2008) and automatic generation of camera placements over time to define a visual discourse that best fits the plot to be rendered (Jhala and Young, 2010). Both of these efforts rely on a planning based approach to narrative, with plots represented as plans.

2.5. Natural Language Generation

The general process of text generation takes place in several stages, during which the conceptual input is progressively refined by adding information that will shape the final text (Reiter and Dale, 2000). During the initial stages the concepts and messages that will appear in the final content are decided (*content determination*) and these messages are organised into a specific order and structure (*discourse planning*), and particular ways of describing each concept where it appears in the discourse plan are selected (*referring expression generation*). This results in a version of the discourse plan where the contents, the structure of the discourse, and the level of detail of each concept are already fixed. The *lexicalization* stage that follows decides which specific words and phrases should be chosen to express the domain concepts and relations which appear in the messages. A final stage of *surface realization* assembles all the relevant pieces into linguistically and typographically correct text. These tasks can be grouped into three separate sets: *content planning*, involving the first two, *sentence planning*, involving the second two, and surface realization. An additional task of *aggregation* is considered, that involves merging structurally or conceptually related information into more compact representations ("Tom fled. Bill fled." to "Tom and Bill fled." or "the boy and the girl" to "the children"). Aggregation may take place at different levels of the pipeline depending on its nature.

3. Narrative Composition from an Engineering Point of View

The type of narrative that we want to address in this paper involves a linear sequential discourse where only a single event can be told at any given point. Yet reality is not like that. Events to be reported may have happened simultaneously in physically separated locations, and constitute more of a cloud than a linear sequence, a volume characterised by 4 dimensional space time coordinates. Composing a narrative for such an input involves drawing a number of linear pathways through that volume, and then combining these linear pathways (or a selection thereof) together into a single linear discourse. This type of linear pathway is sometimes referred to as a *narrative thread*. The analogy between narrative and the production of textiles is pervasive. A narrative of real or fictitious adventures is sometimes referred to as a *yarn*, which is formally a continuous strand of twisted threads as used in weaving or knitting. To spin is to draw out and twist (fibres) into thread but also to relate a tale, or to provide an interpretation of something in a way meant to sway public opinion.

Before we can take this analogy further we need to consider how thread is produced. An initial material (wool, cotton, hemp, sisal...) is selected to provide the starting fibres. These materials occur as volumes of fleece, cottonseed, or stems or leaves of plants. A process is applied to transform these volumes into sets of linear fibres. This is known in various ways but we will refer to it as *heckling*. These fibres are then spun into longer *yarns*, which are continuous lengths of interlocked fibres, suitable for use in the production of textiles, sewing, crocheting, knitting, weaving, embroidery and ropemaking. Several yarns may be twisted together to form a *strand*. Strands may themselves be further twisted together to produce rope.

The elementary description of the problem of narrative composition given above matches the task of converting a volume of fleece into strands of wool. An original mass with no clear linear structure is first processed into a set of linear fibres, some of which are later combined together into more complex elements (yarns, strands, rope) which are still linear but collectively exhibit properties not present in individual fibres (perform better than individual fibres under certain parameters). The assumption underlying this analogy is that the reality (as a set of facts) from which the narrative is to be drawn plays the role of the initial volume of fleece.

As a research tool, this analogy poses the following questions:

1. what is the narrative equivalent for a fibre, and what is the corresponding process of heckling
2. what is the narrative equivalent for a yarn, and what is the corresponding process of spinning
3. what is the narrative equivalent for a strand, and what is the corresponding process of twisting

Not all of these are equally important, but trying to answer them all provides rich picture of narrative composition as a computationally modelable task. To bound the problem at the other end, let us assume that the narrative equivalents of a rope would be large-scale narrative works such as novels, capable of bearing heavy loads of expectation of narrative quality. As both the product and the specification of the requirements it should fulfill are beyond the scope of a paper such as this, we will restrict our present endeavour to the level of yarns, understood as spans of narrative reduced in size and complexity yet already exhibiting the overall formal properties that we desire of narratives (multi-threaded and involving a certain complexity in terms of chronological and spatial relative differences between threads). Thus they would be good candidates to represent small everyday narratives and the elementary abilities of narrative composition. Nevertheless, it should be kept in mind that these yarns or narrative spans would be susceptible of further integration into larger narrative elements.

4. A Computational Model of Narrative Composition

The task of heckling can be related to the identification of appropriate focalization decisions for conveying a given material. Focalization, understood as the decision of which character the narration should follow, and how much of the environment around him at each point should be conveyed to the reader of the narrative, heckles the perception of reality into individual fibres (one for each possible focalizer character) that are linear and sequential in nature. For each character involved in the set of events to be conveyed, a possible focalization fibre can be drawn. In contrast with the physical fibres of textiles, different elements of the material (locations, objects, characters, events...) may feature simultaneously in more than one fibre. This difference is not considered problematic, and it will allow the model to represent important features of narrative, such as the possibility of including multiple perpectives of a given event.

However, it introduces the need for a narrative-specific fibre selection subtask at this stage. As fibres will have redundant information, and very different coverage of the set of events to be reported, an important challenge during narrative composition will be to select the most promising set of fibres from a given inspirational set of events.

The tasks of spinning and twisting are conceptually very similar, involving as they do very similar operations (combining thin linear segments in to thicker ones). At this point the analogy with textiles stops being useful. The structure of textile yarns or strands is significantly different from the structure of narrative in terms of its threads. Textile yarns actually develop in three dimensional space, with fibres spiralling around one another along the length of the yarn. Narrative threads must be combined into a linear sequence of discourse with each fibre taking centre stage of the discourse for a while, then stopping and leaving room for another to take its place, which may later stop and allow for a return to the initial fibre. Thus the combination of narrative fibres needs to take a different form that takes into account this difference in structure.

As a first approximation, two basic operations are considered:

Grafting two separate narrative fibres converge at given point in the narrative, and from then only a single fibre exists

Splicing two (or more) separate narrative fibres are combined into a single discourse sequence by snipping each one of them into smaller fragments of similar relative size and interleaving them in an appropriate order to form a single discourse that switches back and forth between them, covering their whole length

Grafting is a simpler operation than splicing. Grafting occurs when fibres for two different focalizers converge at a point in the discourse and one of them does not continue beyond that (that particular focalization does not appear any further in the discourse). This may occur because the focalizer character for one of the fibres need not appear anymore in the narrative to be told (either because he/she is dead or no longer plays a significant role in the story or because their view on the story is already covered by other focalizations considered more desireable). This type of structure occurs when a secondary narrative thread is inserted into a primary one to provide explanation for certain aspects of it (how someone else happened to be there, why particular characters react in certain ways...).

Splicing is the fundamental operation for handling multi-threaded narratives. It involves the following operations:

1. identifying for each fibre a set of potential break points where the narrative flow may be interrupted to switch to a different thread

2. identifying pairs of origin-target break points in different fibres such that abandoning one fibre at the origin break point and retaking the other at the target break point results in a desireable discourse

The identification of potential breakpoints must rely on various criteria that are either specific to the domain being narrated, or determined by various advanced effects such as suspense or narrative tension. This issue will not be addressed generically here but only for a specific case study below.

The pairing of origin-target breakpoints is also too specific or too complex to be addressed generically here. However, it is important to point out that the relative displacement (in terms of time and space) between the moment and the location associated with the two breakpoints plays a very important role. Other parameters that may need to be considered are the relative size of fibre fragments that results from snipping at a given breakpoint, and conceptual relations that occur between the elements (characters, events, locations, time moments...) described in each fibre around the breakpoint.

This introduces an additional aspect that needs to be modelled. Whenever two fibres have been spliced together, some kind of contextualization must be added at the target breakpoint, to provide the reader with an idea of how the fibre begin retaken relates to the fibre just abandoned. Expressions of the kind “Two days earlier, at headquarters in London,...” play this role, usually occurring at the beginning of new paragraphs or new chapters dealing with new focalizations.

4.1. Fibres and Heckling

An *event* is something that happens with a potential for being relevant to a story. Events occur at a specific location at a given moment of time. They may have preconditions and postconditions. In order to have a generic representation, each event is considered to have an associated *event description* that allows both a descriptive and a narrative component. The descriptive component of an event description contains predicates that describe relevant preconditions. The narrative component of an event describes the action of the event itself, but it may also contain additional narrative predicates describing the effect of the action. For instance, if someone moves from one place to another, elements at the original location may be left behind, and elements in the final location appear. These are not considered separate events but included in the description of the moving event.

A *fibre* is a sequence of events that either involve or are seen by a given character. It represents a focalized perception of the world. The extent of information that is included in the events for a given fibre is defined by the range of perception that is being considered and by the presence of any obstacles to perception in the surrounding environment. It may also be affected by the direction in which the focalizer is facing, or where he is focusing his attention. As these further refinements require advanced authorial decisions, we decide that the standard representation should include all the information within perceptive range, and leave the decision of whether to mention it to a later stage of composition.

The task of *heckling* involves establishing the range of perception, tracking the set of all possible characters involved in the events to be narrated, and for each character constructing a fibre representation that includes descriptions of all the event that the character initiates, suffers or perceives. These descriptions will appear in the fibre in chronological order.

4.2. Yarns and Twisting

A *yarn* would constitute the most elementary type of narrative. It should hold a sequence of fragments of fibre, possibly coming from different fibres. The event descriptions within a given fibre fragment will appear in relative chronological order, but transition to a different fibre fragment in the sequence may involve a shift in chronology. This allows for the representation of phenomena such as flashbacks or flashforwards. Transition to a different fibre fragment may also imply a change in focalization. Together with the change in chronology, this allows for the representation of alternating narration as a single linear discourse of simultaneous narrative threads.

A yarn is obtained by combining together fibres. This can be done in several ways:

- grafting a secondary fibre onto a primary one
- splicing together two (or more) fibres of equal importance
- dividing a fibre in a number of fragments and recombining them in a different order

The criteria to be applied in each case would need further study. This study should take into consideration existing work on narrative and rhetorical effects, as well as domain dependent information and issues such as purpose of the narration or author goals (and in general, the rhetorical problem as described by Flower and Hayes).

Nevertheless, some basic indications have been provided above. Further considerations are made in section 5. for the particular case of narratives told about a game of chess.

4.3. From Yarns to Spans of Narrative Text

To obtain yarns involves composition at a reasonably abstract conceptual level. This may be said to correspond to the planning process of Flower and Hayes' model, or to a content planning task in terms of the classic natural language generation pipeline model. However, the translation process of the Flower and Hayes' model (from these structural plans for the narrative to actual text) involves a further set of elementary operations. Some of these operations[1] are described here.

4.3.1. Contextualization

Once a specific structure for the narrative has been decided upon in terms of yarns, an important task is to establish appropriate contextualizations after each transition into a new fibre fragment. When changes occur at a certain point in the discourse to the location, the narrative time, or the focalization, these changes need to be indicated in some ways to avoid confusion. In the given representation, these changes occur only at the transition between fibre fragments included in a yarn. When the structure of a yarn becomes fixed, it must be traversed in the correct sequence, adding any required contextualizations at the appropriate places. These contextualizations may take different forms:

- inclusion of temporal expressions or temporal discourse markers to indicate changes in chronology (usually with respect to the narrative time holding at the end of the previously narrated fragment)
- inclusion of spatial expressions to indicate changes in location
- relying on combinations of the above and possibly additional discourse features to allow the reader to infer changes in focalization

As the last bullet point indicates, these mechanisms can be fairly complex, and a detailed analysis is beyond the scope of this paper.

4.3.2. Setting Narrative Parameters

The groupings of event descriptions within a yarn provided by fibre fragments constitute very good candidates for the assignment of narrative parameters such as which person to use in narration or verb tense. Because transitions between fibre fragments will signal changes in focalization or chronology, spans of discourse corresponding to single fibre fragments are likely to have similar overall values for person and tense. In the case of tense, the same information used for contextualization (relative shift in chronology from the previous fibre fragment to the current one) will play a significant role in establishing correct values for tense.

They are also likely to be significant in determining layout information (for instance, forcing the introduction of paragraph or even chapter breaks to match transition into new fibre fragments).

[1]The set of operations addressed here is not intended to be exhaustive. Many more would probably arise if the task were addressed in detail.

4.3.3. Filtering and Contextualizing Descriptive Information

An additional operation that may be considered at this stage is filtering the information available in the event descriptions in the yarn according to the desired authorial decisions. Once the author has decided whether the narration is to take place in the first or the third person (or the second), it may be relevant to flag part of the descriptive information contained in the yarn for specific events, so that it does not get realized. Such information would for instance be things that the focalizer character cannot actually see or has not noticed even though they are within perceptual range. This would correspond to a content determination task in terms of the classic natural language generation pipeline.

Also, the actual form that spatial descriptions are going to take may need to be refined, changing from an absolute perspective with respect to the overall space to referential expressions relative to the position and/or the orientation of the focalizer. This would correspond to a referring expression generation task in terms of the classic natural language generation pipeline.

4.3.4. Realization

Further operations at a more basic level of natural language generation would be required.

If a fluent natural text is desired, a stage of aggregation should be considered, for instance to pack together descriptive statements of similar structure, or replace enumerations of relative position of several individuals with generic description of the position of the whole set using an appropriate collective noun.

Additional stages of referring expression generation may be required to replace some of the fully specified references with pronominal references. These pronominal references (and their placement) must be adequately constructed to ensure the resulting discourse remains understandable and no unnecessary ambiguity is introduced.

A stage of lexical choice can be introduced, to explore the possibilities of using different lexical terms for recurring occurrences of the same concept. Depending on the desired style for the resulting narrative, this can be a strong requirement (for more literary styles) or an encumbrance (for texts that value referential precision more highly than aesthetic value). Again, the rhetorical problem and/or constraints on the writing task should be considered here to provide decision criteria.

5. A Case Study: Narratives from a Chess Game

To provide a preliminary benchmark for the various intuitions outlined in the rest of the paper the simplest approximation to a case study that could be conceived is described in this section. This is done by considering a chess game as a very simple model of a formalised set of events susceptible of story-like interpretations. Chess provides a finite set of characters (pieces), a schematical representation of space (the board) and time (progressive turns), and a very restricted set of possible actions. Operating on simple representations of a chess game in algebraic notation, exploratory solutions for the tasks of content selection and content planning are explored based on a fitness function that aims to reflect some of the qualities that humans may value on a discourse representation of a story.

A basic software implementation has been written that reads a description of a chess game in algebraic notation (see Table 1) and builds for it the kind of representations that are described above. The aim of this exercise is to consider broadly what particular domain dependent criteria may be applicable at each of the decision points outlined for the generic case in the previous sections.

Each individual chess piece taking part in the game is considered a character. Perception range is defined as the small space of 3 x 3 squares of the boad that constitutes that immediate surroundings of each piece at any given moment.

Events are triggered by pieces moves. Whenever a piece moves, this constitutes an event for the piece itself, for any other piece captured during the move, and for any other piece that sees either the full move, the start of the move or the conclusion of the move.

Fibres for each of the pieces are built by collecting event descriptions for those moves that they are involved in or they see. The same event may get described differently in different fibres depending on the extent to which the corresponding focalizer is involved in it.

Initial trials with this set up for simple realization of single fibres uncovered a number of conceptual problems. Because chess was used as a context, and no specific measures had been taken, the resulting narratives are not understood as being focalised as they are intended. Chess as a concept tends to invoke unconsciously focalization over the complete board.

Two main problems were found. First, there was no way of identifying from the rendering of the story fragment for each fibre who the story "was about" (the focalizer). Second, focalization was not clear, because the text gave no indication of what was seen at each stage. So the reader, identifying the whole thing as a description from a chess game, assumed focalization to be over the whole board.

To address the first problem, an introductory sentence was added to each fragment presenting this problem, stating who the focalizer is and where she is ("The black queen was four squares north of the centre of the board."). Also, descriptions of events in which the focalizer does not actually take part are translated in indirect form ("The black queen saw the white left bishop appearing ahead.").

To address the second problem, a brief description of what can be seen is also added in the cases where focalization has changed or was previously unknown, to establish the actual range of perception.

Additionally, chronological information in terms of moves or turns in the game was translated into a more intuitive temporal framework by considering that each move corresponds to a day, and appropriate temporal expressions are generated based on that premise. In this way, temporal expression become "two weeks earlier" rather than "14 moves earlier". This significantly reduces the perceived awkwardness.

Spatial information throughout was reformulated either in terms of cardinal points of the compass ("north", "south",...) or relative to the focalizer ("right" "left"

1. e4 c5	16. Bxe2 Be6
2. Nf3 d6	17. Rfd1 Rfd8
3. d4 cxd4	18. Bc5 Rd5
4. Nxd4 Nf6	19. b4 a5
5. Nc3 g6	20. Bf3 Rxd1+
6. Be2 Bg7	21. Rxd1 e4
7. Be3 O-O	22. Bxe4 Bxc3
8. O-O Nc6	23. Bxc6 Rc8
9. h3 d5	24. b5 Bxa2
10. exd5 Nxd5	25. Bd4 Bb4
11. Nxd5 Qxd5	26. Be5 Be6
12. Bf3 Qc4	27. b6 Rxc6
13. Nxc6 bxc6	28. b7 Rb6
14. c3 e5	29. Rd8+
15. Qe2 Qxe2	1-0

Table 1: Algebraic notation for an example chess game

"ahead" "behind"...). The first option was used for describing piece movements, for describing the initial position of pieces at the beginning of the story, or for describing the relative position of the location where a new fibre fragment starts with respect to the location where the previous fibre fragment had ended.

As the case study is intended primarily for basic trial of the intuitions, only the simplest composition operation has been implemented. This constitutes the grafting of the life fibre of a captured piece into the fibre of the capturing piece. This was carried out using the following criteria:

- split the fibre for the capturing piece at the point of the attack
- start the yarn with the first of the resulting fibre fragments
- then add the full fibre for the captured piece (from beginning to attack, but omitting the actual death)
- then add the second fibre fragment obtained from the fibre for the capturing piece (starting with a recapitulation of the attach followed by the death of the captured piece)

The process of dealing with the redundant information present in both fibres, and distributing it appropriately over the resulting fibre fragments (before and after the break, and between the fibres with different focalizers) has for the time being been resolved empirically and would be in need of further study.

An excerpt of an example rendering of the narrative span for a yarn obtained by grafting the fibres for two chess pieces is given below. This constitutes the story of the confrontation between the black and white queens. The yarn starts by narrating the life of the black queen from the beginning of the game, and follows her closely until the beginning of her attack on the white queen. It then tracks back to tell the story of the white queen from the beginning of the game to that same point. Then it describes the outcome of the attack. It finishes by telling how the story of the black queen ends (she does not live long to enjoy her triumph).

> The black queen was four squares north of the centre of the board. The third black pawn was to the right. (...) The black queen saw the third black pawn leaving to the right. (...) Three days later, the black queen moved southeast. The third white pawn remained behind. (..) The black queen saw the white queen appearing ahead. The black queen attacked the white queen.
>
> A month earlier three squares northwest, the white queen was three squares south of the centre of the board. (...) The white queen saw the black queen arriving. The black queen attacked the white queen.
>
> The white queen died. The black queen saw the white right bishop arriving. The white right bishop attacked the black queen. The black queen died.

6. Discussion

In appraising the proposed model it is important to consider that it is based on a number of hypotheses as to what the starting data on which it operates are. Although efforts have been made to start from the simplest possible representation of input data describing the set of events on which the final narrative must be based, it is possible that different initial assumptions might have led to different characteristics in the model.

Along these lines, it may be worth considering that any author facing the task of narrative composition would likely be operating from a memory of the set of events in question. This memory would mediate the task in two different ways. First, it may be remembered in incomplete or incorrect form. This would result in a narrative not matching the inspiring events through no conscious decision of the author or no explicit operation in the composition process. Second, the set of events as remembered may be interpreted by the author in different ways during the composition process. This task of interpretation would affect many of the decision points described in the model, mainly through the effect of the rhetorical problem and the constraints on the writing tasks, but possibly in further ways that have not been contemplated.

A number of related efforts exist to automatically derive narratives from sport games (Allen et al., 2010; Lareau et al., 2011; Bouayad-Agha et al., 2011). These efforts operate on input data in the form of statistics on a given game, and produce texts in the manner of newspaper articles covering similar games. These efforts, for instance, are heavily mediated by the set of data they start from, which arises from purpose specific abstraction and filtering of the set of real world events, driven by the purpose of the desired story.

From the point of view of narratological theory, the model as described captures a number of important features, such as concepts of focalization and chronology, and relates them closely to the input data, the decision criteria, and the computational processes being modelled. It also provides representation for issues such as person, tense and narrative time which are important characteristics of narrative as studied by narratologists.

The set of events before processing might be considered to represent the narratological concept of *story* as defined in section 2.1.. Another possibility might be that the story be the set of fibres obtained after heckling, or the subset of those fibres that get selected for inclusion in the final result. On this issue, the model presented in this paper has the merit of uncovering the degree of vagueness in existing description of these concepts in narratology. The concepts of fabula or story plot as a plan used by AI storytelling systems constitute a different more elaborate representation than the input considered in this paper. To obtain this refined representation it would be necessary to enrich the input with the causal inferences that link together all the events and then to select a subgraph of the resulting causal network to be used as driving plot, ommitting those events in the graph that are not included in that subgraph. This process will be considered in further work.

The structure proposed for yarns as it stands is very plain, as a result of the general aim of solving each problem with the simplest possible mechanism. However, as soon as more complex problems are addressed with the same framework, richer representation structures are bound to be needed, for instance to address the role of causal relations in the selection process. Existing narratological theories and causal models will surely be of use in extending and refining the model.

This paper addresses the task of obtaining a discourse representation from the representation of the input to a narrative composition process based on real life events. Although a representation that captured details of causal relations between events, along the lines of the causal network model (Trabasso et al., 1989), may play a significant role in this process at a deeper level, the work described here constitutes a first approximation that tackles relatively simple structural issues dealing with time and space, not causality. Yet it is clear that further work should address the role of causality in the various decision processes identified. This may present a significant obstacle, as difficult pragmantic inferences will have to be made to interpret the causal structure underlying observed events. This task is far from trivial.

Existing work on AI models of storytelling using planning approaches may be of great assistance in that endeavour. Planning features prominently both in the storytelling literature and in Flower and Hayes' account of writing. The planning applied in most storytelling systems involves drawing causal networks that connect events to one another, in order to provide a guiding line through a story or to suggest additional events that might be added to it. Similar solutions might be applied in order to identify potentially interesting connections between events in different fibres. For instance when events in one fibre can be causally related to events in another. Such connections could provide a very strong basis for ways of structuring the fibres into high quality yarns. The concept of author goals as postulated by Dehn relates very closely to the generic descriptions of the rhetorical problem and the constraints on the writing tasks discussed by cognitive accounts of writing. Both of these issues deserve attention in further work on this model.

Modifications of temporal ordering as considered in (Bae and Young, 2008) could inform the task of dividing an fibre into a number of fragments and recombining them in a different order. Additional options of introducing temporally displaced (possibly partial) versions of a given fragment (as well as or instead of the original fragment as told in its chronological position) must be considered as further work on the model.

The use of the tension arc for a story in the MEXICA system to evaluate the quality of a partial draft points the way towards very plausible and effective criteria both for selecting fibres during the planning stage and for evaluating yarns during the reviewing stage. These possibilities should be explored in further work, once a sufficiently rich representation is found, capable of representing emotion. As it is, the current prototype based on chess games applies very simple criteria for fibre selection, based on identifying fibres with the highest number of piece captures, under the assumption that captures constitute more emotionally charged events than other moves. The criteria for selecting break points, and for redistributing information around graft points also take into account intuititive criteria to sustain and maximise tension.

From a cognitive point of view, the set of operations postulated for the task of narrative composition aligns reasonably well with the processes described by Flower and Hayes. In terms of Flower and Hayes' model, heckling the original material into fibres, fibre selection and twisting fibres into yarns would constitute specific operations of the planning process. Contextualization, setting narrative parameters and realization would constitute operations of the translation process. The model for narrative composition as described in the present paper constitutes a very simple initial description of the task as a one-pass attempt. In more refined versions, the task should be addressed in a cyclic way, involving additional processes of evaluation and revision (the reviewing process of Flower and Hayes' model, currently not represented in the model), and allowing interaction between the various processes as controlled by a monitor.

Finally, a brief note on the use of chess games as indicative case study. Chess games present the advantage of having very clear temporal and spatial constraints, and constituting at heart a sketchy representation of one of the most dramatic settings for human experience: war. In that sense, it provides a good ground for simple experiments, and it is not intended as a contribution but as an illustrative example of the operation of the model of sufficient simplificity to be describable within the size restrictions of a paper such as this. Three aspects are identified as problematic with the chess domain. First, adequate representation of opportunities and threats in a chess game involves some serious representational challenges (Collins et al., 1991). Although significant progress may be made on this point, the effort invested is unlikely to lead to compelling narratives or narratives that bring insight on the task of narrative composition. Second, its extreme simplicity requires a number of additional operations (conversion of chronology to a framework in terms of days, or the introduction of compass points for specifying spatial directions) that may be introducing noise into the experiment. Third, the chronological structure of a chess game is in truth purely sequential, in contrast with the sets of events that would be considered when narrating from real life events. This has not been considered a serious obstacle in as much as focalization breaks up the set of events into separate views corresponding to different fibres, and the numbering or moves in the game provides a good indication of relative chronology of events within any given fibre and across fibres. Yet it also introduces unnecessary complexity when computing chronological allignment between paired break points, for instance, or in determining when jumps in time between successive event descriptions in a fibre should be flagged in the discourse. For these reasons, it is considered advisable to explore further investigation of the suitability of the model when applied to case studies in other domains that are richer in terms of their representation of time and space and that may lead to more compelling narratives with a stronger human interest. The search for alternative domains is made difficult by the need to obtain for them reliable records of all events, both relevant and irrelevant to the story that may be told. Only if this condition is satisfied can the corresponding problem be considered equivalent to the human task that we want to model.

7. Conclusions

The model presented in this paper constitutes a first approximation to a computational model of the task of narrative composition. It draws upon an analogy with textile manufacturing well-based on popular culture. This analogy has provided a break down into subtasks that has lead to interesting insights in terms of specific knowledge-based operations that need to be carried out during composition. These operations relate reasonably well with structural features of narrative as described in literary studies, such as focalization and chronology. They can also be correlated to the set of processes described in cognitive accounts of writing. Finally, they allign and integrate well with generally accepted task divisions for natural language generation. Additionally, a preliminary implementation over a simple indicative case study based on narrating chess games has shown the feasilibity of the approach in practical terms, as well as uncovering a number of elementary issues that arise from particular properties of the chosen domain. Overall, it seems fair to assume that the model might constitute a good starting point for further work both in terms of refining the model and extending the implementation to more complex case studies in other domains.

As specific research lines deserving attention, it is worth listing investigation into the role of causality in the various decisions processes, criteria for selecting interesting fibres to use in a composition, definition of procedures for trimming selected fibres to retain only the parts of them that justify their inclusion, and implementation of solutions for summarizing fragments of fibres that are uninteresting but relevant to the overall structure of the yarn. The model and the implementation should be extended to address in more detail the challenge of identifying appropriate breakpoint at which to splice fibres together, and the task of splicing together more than two fibres.

8. Acknowledgements

The author thanks the anonymous reviewers that contributed important insights to the final version of the paper. This research is partially funded by research grant (TIN2009-14659-C03-01) from the Spanish government.

9. References

N. D. Allen, J. R. Templon, P.S. McNally, L. Birnbaum, and K.Hammond. 2010. Statsmonkey: A data-driven sports narrative writer. In *Computational Models of Narrative: AAAI Fall Symposium 2010*.

B.-C. Bae and R. M. Young. 2008. A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In *Proc. ICIDS 2008*.

N. Bouayad-Agha, G.Casamayor, and L. Wanner. 2011. Content selection from an ontology-based knowledge base for the generation of football summaries. In *Proc. ENLG 2011*, pages 72–81.

G. Collins, L. Birnbaum, B. Krulwich, and M. Freed. 1991. Plan debugging in an intentional system. In *Proc. IJCAI 1991*, pages 353–359.

N. Dehn. 1981. Story generation after tale-spin. In *Proc. IJCAI 1981*, pages 16–18.

L. Flower and J.R. Hayes. 1981. A cognitive process theory of writing. *College Composition and Communication*, 32(4):365–387.

G. Genette. 1980. *Narrative discourse : an essay in method*. Cornell University Press.

A. Jhala and R. M. Young. 2010. Cinematic visual discourse: Representation, generation, and evaluation. *IEEE Trans. on Comp. Int. and AI in Games*, 2(2):69–81.

F. Lareau, M. Dras, and R. Dale. 2011. Detecting interesting event sequences for sports reporting. In *Proc. ENLG 2011*, pages 200–205.

M. Lebowitz. 1983. Story-telling as planning and learning. In *Proc. IJCAI 1983*, volume 1.

James R. Meehan. 1977. Tale-spin, an interactive program that writes stories. In *Proc. IJCAI 1977*, pages 91–98.

R. Pérez y Pérez. 1999. *MEXICA: A Computer Model of Creativity in Writing*. Ph.D. thesis, The University of Sussex.

E. Reiter and R. Dale. 2000. *Building Natural Language Generation Systems*. Cambridge University Press.

M. Riedl and Michael Young. 2010. Narrative planning: Balancing plot and character. *J. Artif. Intell. Res. (JAIR)*, 39:217–268.

T. Trabasso, P. vand den Broek, and S.Y. Suh. 1989. Logical necessity and transitivity of causal relations in stories. *Discourse Processes*, 12:1–25.

Scott R. Turner. 1993. *Minstrel: a computer model of creativity and storytelling*. Ph.D. thesis, University of California at Los Angeles, Los Angeles, CA, USA.

T. A. van Dijk and W. Kintsch. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.