A Fundamental Element for Narrative Parsing

Richard Doust
richard.doust@gmail.com

The Open University, UK

Abstract. Computational approaches to complex cultural objects such as narrative have tended to focus on the complete modelling problem, proposing high-level concepts that are manipulated to generate or analyse stories. For example, Tale-Spin [13] uses the characters' goals and Mexica [14] uses a tension curve to represent love, emotion and danger in order to drive the generation process.
In this paper, we put forward a different view based on the idea of a fundamental element that differentiates narrative from other forms of human cognition and that we call a narrative thread.
Furthermore, this paper supports the position that splitting phenomena such as narrative into constituent parts and researching each one separately, thereby postponing research into their precise interaction, is a pertinent and fertile approach. Our narrative model focusses specifically on suspense and relies on the assumption that the empathetic effects of a given narrative can be isolated from its plot-level description. This assumption has enabled us to build and get feedback from a model that should increase our understanding of the different roles played by event order and empathy in story-telling. We implemented our model computationally for some simple textual narratives and obtained promising evaluation results (N=46) (see [5] and [6]).

Introduction

If we have a pile of bricks, we know that if we put them together in a few standard ways, we can end up building a wall and maybe even a whole building. But what are the bricks of narrative? For sentence understanding, the bricks may be words or perhaps morphemes. For vision, there are edges and surfaces from which we derive the existence of objects. But, what is the right elemental object to capture the 'storiness' of narratives?

In this paper, we explore the potential explanatory power of one answer to that question, and show that our particular narrative brick can at least be used to build a wall, and perhaps soon, a whole building. Of course, like all scientific processes, the real world will give our model feedback on the accuracy of the analysis it proposes, pointing to faults and weaknesses and enabling new constraints and modifications to be imposed.

Additionally, our focus for now will be on one ingredient which we believe is constitutive of what makes a good story, namely: suspense. As discussed in [4], suspense is a pervasive narrative phenomenon that is associated with greater

enjoyment and emotional engagement. We thus ask: what media and domain-independent model of suspense can we develop that could be useful in a wide variety of situations for improving our techniques for creating new narratives?

We frame this article by generalising the development steps leading to our model with the aim of encouraging similar 'fundamental' formal and computational approaches aimed at understanding complex human cognition. The rest of this paper details these steps.

The first step we follow is to look for something similar to a syntax/semantics distinction. Secondly, we collect the constraints on a potential fundamental element that could generate the syntax observed. We then create an element that potentially satisfies them and with it, derive a variety of theoretical behaviours and effects which can be tested and modified or refuted. Once the model achieves a degree of explanatory power for some very specific human experiences, we can test it further and also use it generatively to create new cultural artefacts.

Step 1: Making a syntax/semantics distinction

1.1 Event segmentation

Before we can start our process, we must make the initial step of dividing up the 'stuff' we have into its basic parts. Stories and the real world have in common that human cognition divides them up into events. [18] suggests that this may reflect the existence of a general network for understanding event structure. Recent work by [12] has also found very strong convergence on event boundary judgments across film and textual media. Therefore, the first front-end process we consider is event segmentation. [21] and [22] propose an event indexing model which lists features such as space, time, causality, and we base our event segmentation on this model.

1.2 Narrative syntax

In 'Why anyone would read a story anyway', [9] categorised narrative interest into two kinds: the cognitive interest that arises from the story structure and the emotional interest that arises from the emotional context of the story. Consequently, the first step we propose is to separate the narrative content concerned with emotions and empathy for story characters from the simple sequences of interlinked events called a plot.

This distinction is, of course, analoguous to the syntax/semantics distinction prevalent in language understanding. To emphasise the parallels, we can say the following: syntax/plot is the collection of rules that govern how words/events are assembled into meaningful sentences/narrative sequences, whilst semantics/empathy is concerned with the meaning/importance of the words/events themselves.

Of course, in stories, creating empathy is essentially linked to the concept of a character. In our view, perhaps 70% of the sentences of a novel are concerned

primarily with creating and maintaining the basic empathetic link between the reader and the main characters. However, we postpone work on story semantics by assuming the existence of an 'empathy module' that produces a certain output. Our focus will be on the narrative parsing of the story plot.

Step 2: Creating a fundamental building block

We next postulate the psychological existence and computational utility of a fundamental element with which we can capture all the narrative-specific processing that is needed to generate and understand the 'story syntax', or plot. To create such an element, we first identify the available and derivable constraints such an element must satisfy.[1]

2.1 Collecting the constraints on the building block

Here we combine research in psychology on the requirements for knowledge structures used for sentence understanding (see for example [7]), research on the requirements for narrative understanding by [1] and some of the requirements from the GLAIVE project ([19]) to provide the basis for a definition of the useful size and complexity of a fundamental building block for narrative.

Starting point and end-point

– The element must have a clear starting point and end-point.

This is derived from Brewer and Lichtenstein's psychological theory of narrative understanding which suggests that three major discourse structures account for the enjoyment of a large number of stories: surprise, curiosity and suspense. Their approach is based on the existence of Initiating Events (IE) and the corresponding Outcome Events (OE) that are triggered by them.

To produce suspense, the IE and OE must be ordered chronologically. In addition, 'often additional discourse material is placed between the initiating event and the outcome event, to encourage the build-up of suspense' ([1, p. 17]).

Linearity of events

– The element must have a clear linear path to completion.

This is derived from the constraints on causal consequences grounded in psychological research such as [21], the constructionist and prediction-sustantiation models of narrative comprehension [7] and scripts [10, 16]. This psychogical research places the following conditions on ease of recall from long-term memory:

[1] A complete account of the derivation and motivation of the following constraints is given in [5].

– there is either a strongly *supporting* context: many different sources point to the same object.
– or a strongly *directive* context: very few alternative inferences are possible.

We interpret these conditions together as indicating that our element must consist of a *linear* series of events.

Consistency

– The element must be internally consistent: no event can contradict any other event in the same narrative thread.

These constraints are inspired by the GLAIVE project ([19]) which distinguishes between causal chains and intentional paths and imposes the following conditions:

– No event in a causal chain can negate the preconditions of another event in that chain.
– A character must consent to all steps in a intentional path and intends the final effect of the last step during all the preceding steps.

Interruptibility

– The element must be interruptible.

Much as events in the *same* fundamental element should not be in conflict, events in *different* elements must have some potential for conflict in order to produce the uncertainty necessary for suspense. We therefore need to model information about the following type of interaction between events: if Event E occurs in a story, then Event F can no longer occur in this story. This is analogous to syntactic rules that exclude certain category sequences in sentence grammar.

2.2 Creating the fundamental element

We now use all these constraints to deduce the fundamental element that will form the basis of our model of narrative. We call this element a *narrative thread*.

Combining constraints from 2.1 and 2.1, we extend the simple IE-OE pair to obtain the following:

$$IE \rightarrow Event_1 \rightarrow Event_2 \rightarrow Event_3 \rightarrow ... \rightarrow OE \tag{1}$$

This simple script-form has the necessary clear final result or outcome in the storyworld that we need for suspense and, according to the constructionist model, is also habitually and easily generated during narrative comprehension.

Using the constraints in 2.1, we can allow our narrative threads to be built by combining both causal chains and intentional paths whilst checking for internal consistency.

Finally, to create the necessary interruptibility mentioned in 2.1, we postulate the existence of a set of disallowing pairs (E, F) such that an event E in one narrative thread can disallow an event F in a different thread.

Here is an example of a narrative thread taken from [5]:

1. A wants to kill B
2. A plants a bomb in B's car
3. A checks that B gets in the car
4. A triggers a remote control device
5. The countdown starts on the remote control
6. The countdown starts in the car too
7. The countdown goes on for some time
8. The countdown reaches the end
9. The bomb explodes
10. B gets killed

Narrative threads thus model a reader's expectations about what might happen next in a given story. For a given storyworld, there will be multiple narrative threads in competition with each other. Their interaction models what goes on when a reader experiences stories in that storyworld.

Checking for the modularity between construction/function Narrative threads are simply lists of events that are likely to follow each other. In general, threads are informed by a variety of inferential and associative mechanisms: scripts, models of story characters involving beliefs, goals and desires, principles of naive physics and even previous narrative experiences[2].

Our claim is that it is not necessary to know exactly how different information sources were used to construct narrative threads for them to be used successfully to model narrative. The available storyworld information and inferences can always be *translated* into the linear structure that we call a narrative thread. We thus postulate a separation between the *inferential sources* that build narrative threads, from the *structure* of the ongoing suspense processes they trigger and maintain.

Checking for generality We must distinguish certain typical narrative *instantiations* from the fundamental narrative processes that underpin them. This is the token/type distinction. Thus, for example, the fairytale models developed in [15], should be seen rather as the socially calibrated and sedimented result of narrative play than as fundamental models of narrative creation in themselves. Moreover, these models should be *explainable* in terms of narrative threads.

[2] Although models of intention such as plans and goals structures may be key elements for many narrative threads, to use such structures as a fundamental element would exclude narratives based more on causal interactions with the physical world. We claim that it is possible to feel suspense about, say, an ice floe melting where no clear plans and goals are present.

Similarly, the structures used to *model* narrative should not be entangled with the specific *types* of narrative for which they create narrative effects. Just as with syntax, we should be able to create two narratives that have the same formal structure even though they exist in two completely different storyworlds.

Step 3: Creating a scaffold for new concept generation

We now create a theoretical scaffold for the story-telling phenomena upon which we can formalise a variety of narrative phenomena.

3.1 Modelling world knowledge

Firstly, to model a given storyworld, we derive the set of narrative threads that model the causal and intentional links it contains or suggests. We also create a set of disallowing event-pairs that define the interactions between the threads[3].

3.2 Modelling the input

The next step is to model the input to the storyworld. In this case, a story is just one particular sequence of events chosen from all possible events, that is, from all the events in the narrative threads used to model the storyworld. Telling a story, under this account, is then simply to evoke one by one this sequence of events.

3.3 Modelling their interaction

When an event E is told in the story, any narrative thread T that contains event E is activated. Intuitively, this means that the reader predicts that the events that follow E in thread T will be told in the story. As a story progresses, narrative threads are activated and de-activated. Also, some pairs of upcoming events in different active narrative threads may be in conflict with each other, that is, they may belong to the set of disallowing event-pairs.

Step 4: Defining new concepts

Using our basic building block and this theoretical story-telling scaffolding, we can now derive some structurally refined definitions of a number of narrative concepts. In our case, we make the claim that the following three fundamental mechanisms occur in stories to create suspense and are usefully defined as distinct narrative phenomena:

[3] For the moment, the narrative threads used in our computational model are derived by hand. However, our narrative thread is very similar to the narrative schemas developed by [2] and our disallowing relations are similar to the exclusion relations derived in [11], both of which can be harvested automatically. In future work, we aim to apply our method to such automatically harvested schemas and extend our model to different storyworlds and story variants.

– completion-based suspense
– conflict-based suspense
– revelatory suspense

4.1 Completion-based suspense

Translated into a context of goals and plans, completion-based suspense would correspond to the achievement of a goal. An example would be the captain of a football team walking slowly towards the tribune to receive and raise in the air the prize cup her team has just won. As the outcome event becomes ever more imminent the suspense increases; we can talk of *completion imminence*. There may be absolutely no expectation that the achievement of this goal will be interrupted and yet we experience a kind of suspense in such cases[4].

Completion imminence is a function of the number of events in the thread that remain to be told before it is complete. Figure 1 shows a thread with a completion imminence of 4 events at this stage in the story. We have tagged these events to show how the cup-raising episode could appear.

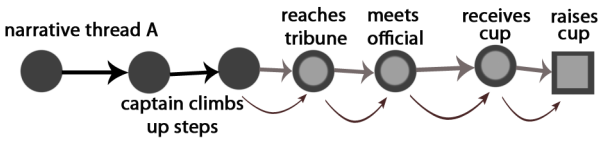

Fig. 1. Completion Imminence. Told events are black, untold grey.

4.2 Conflict-based suspense

This occurs when at least two active narrative threads contain upcoming incompatible events of which only one can actually occur. Furthermore, a big difference in story outcomes is expected depending on which event actually does occur. An example would be a chase between a policeman and a thief who is trying to cross a border. As the policeman gets closer and closer to catching the thief, the imminence of the catching event increases; we can talk of *interruption imminence*.

The conflict-based suspense of a thread A is related to the smallest number of events still to be told in some other thread B before an event can be told which interrupts A. To illustrate conflict-based suspense between two threads, in Figure 2, we show thread A which has an interruption imminence of 3 events due to thread B. We have tagged these events to show how the chase episode described could appear.

[4] Completion-based suspense may be linked to a phenomenon discovered by Lithuanian psychologist Bluma Zeigarnik who observed the effect of interruption on memory processing. [20] proposes that when a task is started, it creates a quasi-need for its completion such that when the task's progress is momentarily halted, the subject remains in a state of tension until it is completed.

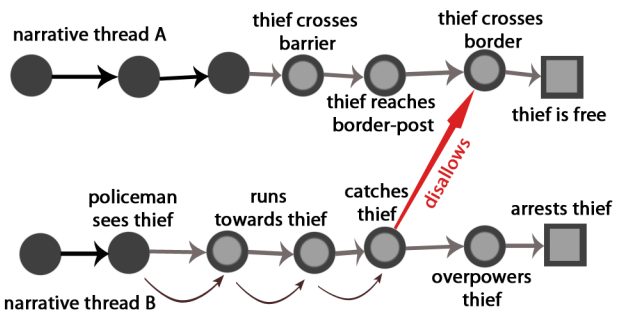

Fig. 2. Interruption Imminence

4.3 Revelatory suspense

Curiosity in Brewer and Lichtenstein's model of narrative is about the past; we wonder what *happened* before an event. If we see a man in sunglasses smiling as he sees another man some distance away getting into his car, we might imagine that it is because he had planned to meet his friend, or because he had planned some sinister plot. Such situations evoke what we call relevatory suspense which is linked to the disambiguation of a story event (smiling) that belongs to more than one narrative thread.

To model this, we imagine that in such situations a common event is present in two different narrative threads as illustrated in Figure 3. When the common event is told in the story, both threads become activated as candidates to explain the event's presence in the story. Subsequent story events may disallow one of the threads leaving one thread as the correct 'explanation' of the common event. This epistemological gap-filling process is suspenseful in itself: we know that every time we find out new information, it is likely to have a major effect on the set of active narrative threads.

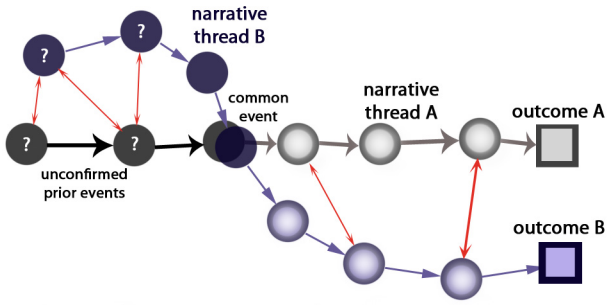

Fig. 3. Threads with a common event. Implicated prior events have a question mark. Bidirectional arrows show disallowing relations.

4.4 Modelling contextual factors

Relative importance At any one moment in a given story, a variety of suspenseful situations may be present. Relatively unimportant suspenseful situations may coexist with life-or-death situations. The importance of these situations will also often depend on the reader's emotional involvement with them and this could be low or high and positively or negatively valenced.

In our model, we presuppose the existence of a 'empathy module' that ascribes a relative emotional importance value to all the narrative threads. The value ascription of a narrative thread is related to the reader's appraisal of the state of the storyworld when the *last* event in the thread has been told. This use of a single value to encompass a multitude of factors enables us to take into account the modelling of emotions but at the same time keep our focus on the structure of the information flow and its relation to suspense.

Often, the importance of events is related to the fate of some story character. In this case, we can base the importance values on the i) the current level of sympathy (or antipathy) towards a character involved in an event, and ii) the perceived desirability (or undesirability) of the event in relation to that character.

Foregroundedness If a given narrative thread T is not evoked for several story steps during the telling of a story, we assume its effect on suspense will drop because it is less present in the reader's mind. At the same time, other narrative threads are of course active and competing for the reader's attention. Of course, as soon as a new event belonging to thread T is told, it comes again to the foreground and regains its full potential for suspense creation. Therefore, in addition to a measure of the relative importance of the narrative threads, we also need a measure for the degree of foregrounding of a thread.[5]

4.5 Defining a complete model of behaviour

With a suitable heuristic, we can now combine our measures for the three types of suspense: *completion-based*, *conflict-based* and *revelatory* with the additional contextual measures of *importance* and *foregroundedness* to obtain values for the suspense contributions of all active threads. We then combine these values to create an overall suspensefulness value for each story-step. We now have the apparatus needed to predict a step-by-step suspense profile for a given story.

Step 5: Testing the conceptual framework

To test this conceptual framework, we first implemented our narrative model computationally. We then derived storyworld-specific information in the form

[5] In our model, foregroundedness is roughly equivalent to recency of mention. Recency has been extensively researched in the psychological field as an important factor influencing memory (see for example [8]).

of threads and disallowing pairs of events for a story about an anti-Mafia judge driving back home one evening. We then created an online experimental set-up that used participants' self-reporting of perceived suspense levels whilst reading step by step through a story. We carried out a first study to calibrate some of the importance levels of our narrative threads. We then created a story variant, used our computational model to create a predicted suspense profile for this variant, and again used our online set-up to gather participants' self-reported suspense levels for this variant (N=46 for 31 steps).

Analysis of the results using a variety of statistical tests show that our narrative model predicts step-by-step fluctuations in suspense levels for short stories that have a *high correlation* with average self-reported human suspense judgements[6].

Our computational implementation is fully independent of the storyworld information. We thus claim to have created a model of suspense that could be used for any media or genre.

Step 6: Extending the model

We see this work as a signpost towards further development of models of narrative based on what we see as its fundamental ingredients. Further work should explain how higher-level narrative concepts, such as certain types of plot and central characters, naturally include the key ingredients of what is needed to build a successful story. We now examine other directions for further research.

6.1 Scene-switching: the power of meanwhile

Scene-switching can be defined as the alternation between different narrative view-points that show different sequences of events belonging to the same story. It is ubiquitous in suspenseful films. One possible explanation for its use is that scene-switching increases the length of time that the suspense generated by a particular narrative thread is present in the story. If we have two narrative threads A and B, and we show first A then switch to B, then as long as A does not get forgotten, its suspensefulness can continue to affect the reader or viewer, even as we are watching events in narrative thread B[7].

6.2 The narrative cycle

The inspiration for our model came from work on suspense, curiosity and surprise introduced in [1]. Future work should therefore extend the narrative thread model

[6] The Pearson Correlation Coefficient was 0.82 and Spearman's Rho Coefficient was 0.79. See [5] and [6] for the details of the experiment and the results.

[7] The power of scene-switching may also be linked to what is called the Zeigarnik effect ([20]), which suggests that sequences that are interrupted have more longer lasting cognitive effects and, for example, are remembered more precisely. Conversely, sequences that achieve closure can be quickly forgotten.

to the concepts of curiosity and surprise in order to complete the typology of story-building elements. Because situations where curiosity is evoked necessarily involve unconfirmed threads, such situations often evoke surprise.

Our narrative model can also model surprise. If thread A has many told events and thread B has no told events, then surprise will occurs if an event from thread B occurs in the story which suddenly disallows thread A.

In our view, the key moments of many narratives in film or text combine all three of these entertaining narrative effects in what we dub the 'narrative breathing cycle', illustrated in Figure 4 on page 11. In this cycle,

- first a suspenseful situation is interrupted by a surprising event
- this event sets up a new phase of revelatory or conflict-based suspense
- another surprising event occurs which interrupts this suspense phase

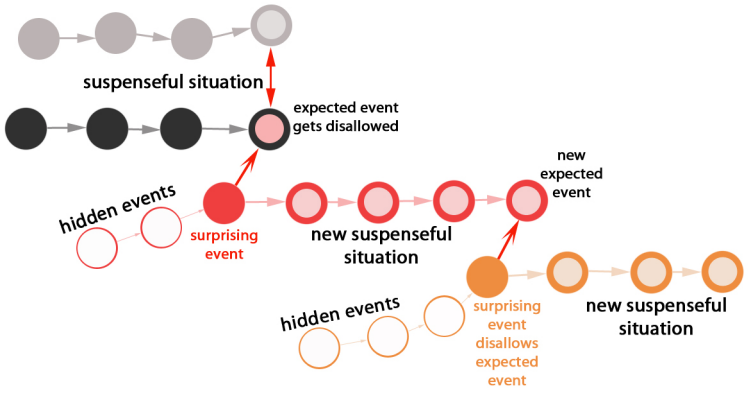

Fig. 4. The narrative breathing cycle. Filled-in circles represent told events. Each narrative thread has a different colour.

Of course, stories differ in the amounts of conflict-based and revelatory suspense they evoke and also in the length of time that suspense is maintained before a new surprising event occurs. We can say that stories have different *suspense and surprise profiles*.

6.3 A functional theory of narrative

Earlier work on a net-linguistic implementation of an Earley Parser (see [17]) influenced the development of our model. A clear analogy can be drawn between linguistic theories on sentence disambiguation and the model of suspense we propose. Narrative threads are analogous to lists of grammatical categories and the disambiguating process between narrative threads is analogous to syntactic

parsing. Revelatory suspense is analogous to words that seems to belong to two or more linguistic categories. Surprise could be seen as analogous to the effect of garden-path sentences.

In this light, our research raises the possibility of constructing a complete functional theory of narrative. Such a theory would postulate that all the story steps that an author produces must modify, disallow or activate at least one narrative thread and thus have some effect on the surprise, curiosity or suspense of the story at that point. Thus, just as we identify the function that each word plays in a sentence, we could identify the function that each story step plays in a story. We could derive a functional content-independent map of a story, much like the syntactic analysis of a sentence.

Conclusion

In our quest for a fundamental element underpinning narrative, we have been able to separate out suspense as an essential ingredient of narrative and also given a possible theoretical analysis of how it functions. We have also proposed a typology of suspense. This not only enables us to predict the suspense profiles of textual or filmic stories, but potentially also of other narrative-like artefacts such as music, advertising and humour. We think that it should be possible to use our model to generate suspenseful stories from a pre-defined storyworld *ex nihilo*.

In making these contributions, we have followed the position that splitting phenomena such as narrative into its constituent parts is an approach that allows a high degree of portability into a variety of domains. We believe that the increased creativity and inter-disciplinarity that this way of proceeding fosters is a good measure of its potential.

Furthermore, as readers of suspenseful novels know, suspense makes us focus. If the balance between skill and challenge is just right for a reader constantly attempting to understand and predict what will happen next in a suspenseful story, they may enter into the creativity-boosting state of Flow (see [3]). Suspense can be therefore be seen as having a direct effect on creativity.

Acknowledgments

Most of this research was undertaken as part of a Ph.D. program at the Open University, UK under the supervision of Richard Power and Paul Piwek.

References

1. Brewer, W., Lichtenstein, E.: Stories are to entertain: A structural-affect theory of stories. Journal of Pragmatics **6**(5-6) (1982) 473–486
2. Chambers, N., Jurafsky, D.: Unsupervised learning of narrative schemas and their participants. In: Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP. Volume 2., Association for Computational Linguistics (2009) 602–610

3. Csikszentmihalyi, M.: Flow: The psychology of optimal experience. Volume 41. HarperPerennial New York (1991)
4. Delatorre, P., Arfe, B., Gervás, P., Palomo-Duarte, M.: A component-based architecture for suspense modelling. Proceedings of AISB, 3rd International on Computational Creativity (CC2016) (2016) 32–39
5. Doust, R.: A domain-independent model of suspense in narrative. PhD thesis, The Open University (2015)
6. Doust, R., Piwek, P.: A model of suspense for narrative generation. In: International Natural Language Generation (INLG2017). (2017)
7. Graesser, A.C., Singer, M., Trabasso, T.: Constructing inferences during narrative text comprehension. Psychological review **101**(3) (1994) 371
8. Jones, M., Love, B.C., Maddox, W.T.: Recency effects as a window to generalization: separating decisional and perceptual sequential effects in category learning. Journal of Experimental Psychology: Learning, Memory, and Cognition **32**(2) (2006) 316–332
9. Kintsch, W.: Learning from text, levels of comprehension, or: Why anyone would read a story anyway. Poetics **9**(1) (1980) 87–98
10. Lebowitz, M.: Story-telling as planning and learning. Poetics **14**(6) (1985) 483–502
11. Li, B., Lee-Urban, S., Johnston, G., Riedl, M.: Story generation with crowdsourced plot graphs. In: AAAI. (2013)
12. Magliano, J., Kopp, K., McNerney, M.W., Radvansky, G.A., Zacks, J.M.: Aging and perceived event structure as a function of modality. Aging, Neuropsychology, and Cognition **19**(1-2) (2012) 264–282
13. Meehan, J.: Tale-spin, an interactive program that writes stories. In: Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI), Cambridge, MA, USA. Volume 1. (1977) 91–98
14. Pérez y Pérez, R., Sharples, M.: Mexica: A computer model of a cognitive account of creative writing. Journal of Experimental and Theoretical Artificial Intelligence **13**(2) (2001) 119–139
15. Propp, V.I.: Morphology of the Folktale. Volume 9. University of Texas Press (1968)
16. Schank, R.C., Abelson, R.P.: Scripts, plans, goals and understanding: an inquiry into human knowledge structures. Lawrence Erlbaum Associates Publishers, Hillsdale, NJ (1977)
17. Schnelle, H., Doust, R.: A Net-Linguistic "Earley" Parser. In: Connectionist Approaches to Language Processing. Routledge (1992) 170–205
18. Speer, N., Zacks, J., Reynolds, J.: Perceiving narrated events. In: Proceedings of the 26th Annual Meeting of the Cognitive Science Society, Chicago, IL. (2004)
19. Ware, S.G., Young, R.M.: Glaive: A state-space narrative planner supporting intentionality and conflict. In: Proceedings of the 10th International Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE2014, North Carolina State University, Raleigh, NC USA (2014) 80–86
20. Zeigarnik, B.: On finished and unfinished tasks. A source book of Gestalt psychology **1** (1938) 300–314
21. Zwaan, R.A., Langston, M.C., Graesser, A.C.: The construction of situation models in narrative comprehension: An event-indexing model. Psychological Science **6**(5) (1995) 292–297
22. Zwaan, R.A., Radvansky, G.A.: Situation models in language comprehension and memory. Psychological bulletin **123**(2) (1998) 162