The Tale of Peter Rabbit: A Case Study in Story-Sense Reasoning

Malcolm Ryan and **Nicholas Hannah**
{malcolmr,nhannah}@cse.unsw.edu.au
School of Computer Science and Engineering
University of New South Wales

Joshua Lobb
j.lobb@unsw.edu.au
School of English
University of New South Wales

Abstract

The telling and understanding of stories is a universal part of human experience. If we could reproduce even part of the process inside a computer, it could expand the possibilities for human-computer interaction enormously. We argue that in order to do so, we need to model narrative at three levels of abstraction, in terms of physics, characters and plot. Taking four scenes from the children's story *The Tale of Peter Rabbit*, we describe some of the challenges they present for modeling this kind of "story-sense reasoning".

Introduction

Stories are a common part of our everyday lives as humans. Not only do they form a large portion of our everyday entertainment, but they also allow us to construct meaning from our lives. We are quick to compile narratives from the muddle of our experiences, and enjoy relating them to one another. For both of these reasons, building AI representations of narrative would prove a useful project, giving a more accurately 'human' model of comprehension.

In the field of entertainment, story AI could help writers build more ambitious works of art, in much the same way that encoding the rules of perspective and lighting into 3D graphics has allowed artists to make works of unprecedented complexity. Much is made of the possibility of "interactive narratives", but the sheer complexity of such works makes it seem unlikely that they will ever fulfill the ideal without a great deal of the story being constructed and manipulated automatically. For this to work, we need computers which can understand the author's intention for the dramatic outcomes of the work, so as to be able to incorporate the reader's actions appropriately.

This requires some ability to recognise what a narrative is, and to distinguish a well-structured narrative from a random sequence of events. Narrative theorists have studied this question for well over a century, and while there is no universal narrative structure, common patterns exist. For several decades, AI researchers have tried to encode this knowledge into some kind of computational form, for narrative understanding and generation. This turns out to be a difficult task. Narratives operate on multiple levels. Much of the reasoning behind them is based on a commonsense understanding of how the everyday world operates. Encoding this kind of knowledge has long been recognised as a major challenge in AI.

In addition to this, narratives follow particular patterns that distinguish them from arbitrary sequences of events. This "story-sense reasoning" includes such concepts as suspense, climax, resolution, etc. It also includes particular approaches to the modeling of the physical world and of interactions between characters.

There is a tradition among the commonsense reasoning community of posing "challenge problems", brief scenarios which encapsulate significant concepts from commonsense knowledge (Morgenstern & Miller 2007). A single problem might require reasoning about knowledge, planning, and space and proximity. The goal of building these challenges is to provide a circumscribed but non-toy problem to drive the creation of reusable theories of commonsense reasoning and to test the limits of existing theories.

Following this tradition, in this paper we offer a challenge to drive story-sense reasoning. Taking *The Tale of Peter Rabbit* by Beatrix Potter (1902) as our source, we present a scene which offers challenges in a number of different areas of AI. Our plan is to challenge some of the short-comings of existing work and to inspire directions for future innovation.

Background

Before we describe the scene, we must first decide on our goals. What do we mean when we say that an AI "understands" a narrative? The following criteria seem reasonable:

- The ability to *explain* a narrative, answering questions such as "Why did Peter go into the garden?"
- The ability to *speculate* about the narrative, answering questions such as "What is going to happen next?"
- The ability to *generate* narratives which follow the standard rules of narrative.

Based on the work of literary narrative theorists such as Barthes (1982) and Prince (1982), we adopt the view that these questions can be approached by dividing the general concepts of story into different narrative functions. We are not referring to the often-quoted division between story and discourse (Genette 1980), but a more basic division between *story* and *plot*. A story comprises a set of circumstances in a fictional world (the physical dynamics); a plot is the construction of these dynamics into a sequence of cause and effect (Forster 1927; Shklovsky 1965).

The interaction between story and plot is made manifest in the author's representation of character, which transforms *actions* into *motives*. Narrative production, then, is dependent on three things: the physical laws of the story world, the deliberate intentions of the characters, and the plotting of the author.

So to reason about narrative, we need three different levels of representation:

- An *action model* which describes the physical dynamics of the world: actions, their effects, time, continuous processes, concurrency, etc.,
- A *character model* which describes the psychology and motives of the characters, their traits, goals, beliefs, emotions and so forth,
- A *plot model* which describes the actions and characters in terms of their dramatic purposes in the narrative.

Furthermore, since a good narrative must make sense at all three of these levels, and will contain events motivated by more than one of them, a representation scheme must be able to capture the balanced interaction between models. The importance of finding this balance is well documented.

For example, Riedl points out in his thesis (Riedl 2004) that a story must have plot coherence, that is, an outcome which is causally connected to the events of the story, and character believability, that is, the events of the story must be reasonably motivated by belief and goals of the characters that participate in it. We would go further and say that narrative, by definition, implies dramatic structure, which gives significance to characters and events by making them part of a recognisable plot. The advantage of the proposed division, which will be seen shortly, is that the given models correspond closely to well-established work within AI, especially within knowledge representation and reasoning.

The Action Model

Action is what fundamentally sets the narrative apart from other forms of writing. Time passes, characters act, things happen and the world changes. To understand narratives we must first be able to understand action.

There are many approaches to modeling physical systems, some of which, such as planning operators, are being put to good use for narrative generation (Charles *et al.* 2003; Young *et al.* 2004). However these are rarely expressive enough to cover the wide range of physical consequences that occur even in seemingly simple scenarios such as cracking an egg into a bowl (Shanahan 2004).

Here, we look towards research in the area of reasoning about action and change. Notably, two major logical formalisms have arisen from this work. The *situation calculus* (Levesque, Pirri, & Reiter 1998) and the *event calculus* (Sadri & Kowalski 1995). Both attempt to formalise our understanding of how things change over time and how agents' actions and other events effect those changes. They both support a number of different types of inference: e.g. temporal projection (deduction), planning (abduction) and diagnosis (explanation). The commonsense reasoning community has constructed models of several different physical phenomena in languages such as these, including: topological and metric spaces (Randell, Cui, & Cohn 1992; Davis 1995), and time, concurrency and processes (Shanahan 1990).

The Character Model

Characters are the driving forces of narrative. Drama arises from the characters' desires and goals, how they act upon them, and whether those goals are frustrated or conflicted (James 1957). So to understand narrative, there is a strong need to model the psychological impulses of its characters. We need to understand not just *how* they act but *why*.

These qualities are the subject of research in a number of connected areas of AI: intelligent agents and multi-agent systems, naive psychology, and knowledge representation and reasoning in general. There exist many logical formalism which deal directly to mental attitudes; too many to enumerate here, but see (Woolridge 2000) for a summary.

Several narrative theorists have worked to distil the motives of fictional characters into character types' (Forster 1927; Rimmon-Kenan 1983): an adaptation of these approaches might prove a manageable way of quantifying character intention.

The Plot Model

Narrative is more than just action and character: there is always an organising principle, a structure that manages the smaller incidents into a coherent whole. Narrative theorists, particularly the Russian Formalists and Structuralists of the 20th century (Shklovsky 1965; Barthes 1982), have striven to identify these structures, and catalogue their use.[1]

Work in this area of AI is less prolific than in modeling action and characters, as it is difficult to build models which capture subtle concepts such as drama, irony or suspense. Some practitioners resort to hand-labeling situations as more or less dramatic, but this is unsatisfying, as it requires the author to foresee every possible path the narrative may take and label it accordingly. To understand previously unseen narratives, and to generate original narratives without large amounts of human intervention, more general models of drama are needed. Some efforts have been made towards this, including: the work of Sgouros on generating Aristotilean-style plot, that is, a series of conflicts, an antagonistic climax, and ending with an unambiguous solution (1999); the various works on generating folk tales, often based on the work of Propp (Gervás *et al.* 2005); limited work on suspense (Cheong 2006).

The Challenge

We have taken a short passage from the Peter Rabbit text for examination. It illustrates reasoning at the levels of action, character and plot, and places particular requirements on each of these models. Where possible, we indicate work which addresses these requirements, but in many cases we recognise that they are still open problems.

[1] Cavazza (2006) provides a concise summary for AI practitioners new to this field.

Scenario: Forbidden

> *'Now, my dears,' said old Mrs. Rabbit one morning, 'you may go into the fields or down the lane, but don't go into Mr. McGregor's garden: you Father had an accident there; he was put in a pie by Mrs. McGregor.'*
>
> *'Now run along, and don't get into mischief. I am going out.'*

Already, on the second page of the story, we are thrown into the depths of great difficulty. Why does Mrs Rabbit forbid the children from entering the garden? There are many implicit assumptions required to make sense of this action:

- The children are able to go into the garden. Mrs Rabbit knows this. She would not forbid them from doing something they could not do.
- The children desire to enter the garden. Mrs Rabbit knows this. She would not forbid them from doing something she has no reason to suspect they might do.
- She believes that something bad might also happen to the children if they enter the garden.
- As their mother, she cares about them and wants to keep them from harm.
- She is going away, so will not be able to keep them safe.
- She believes that if she tells them not to go into the garden, then they will obey her (because she is their mother).

None of these facts are spelled out in the text, yet to make sense of the act of forbidding, we must understand them all. This is common sense: the "elder" must have a motivation to prevent the forbidden act, but is unable to do so directly, so a command is given. The elder must also have the authority to believe that this command will be obeyed. Furthermore, the elder would not bother forbidding something which was not likely to happen, being impossible or already undesirable to the one forbidden.

So at the level of character modeling, we need agents which can reason about each other's desires and abilities. We also need to establish relationships between characters (in this case, mother and children) which entail properties of authority and concern. Then we can define a communication act in the action model, which would also require suitable models of proximity and audibility.

There exists significant work on representing modalities such as desire within individual agents however, for the most part, this has not been extended to handle reasoning about the mentality of others. Unlike in agents research, we have the additional requirement that characters act based on the attitudes of others. Blumberg (2001) notes that this meta-level reasoning is essential to uphold the "illusion of life" (Thomas & Johnston 1995), needed to keep an audience engaged in a narrative. We believe that this is a good example of how narrative modeling can provide new motivations and problems to extend mainstream agents research.

A long standing practical approach to defining and establishing relationships (with associated commitments etc.) involves the use of protocols, e.g. Contract-Net (Smith 1980). Clearly this is not flexible enough for our purposes; capturing Peter's disobedience and Mrs. Rabbit's subsequent response would be extremely difficult. Perhaps a more promising approach is the use of an appropriate deontic logic (Dignum 1999). These logics have been used to describe how agents adopt, violate, or adhere to social norms. They can capture dependencies between agents without compromising autonomy.

We have elided Mrs Rabbit's reason for believing the garden is dangerous. This is also implicit. She knows the history of Mr Rabbit, who was caught and put in a pie. By analogy, she reasons that the same thing might occur to the children. Modeling analogy is complex. While it is easy to see that "Mr Rabbit" can sensibly be generalised to "any rabbit", deciding how far the generalisation can be taken (any animal? any living thing?) is a subtler problem.

As for the dramatic purpose of this event, it is instructive to read the scenario using Propp's Interdiction-Absentation-Violation functions (Propp 1968). In Propp's terminology, the Hero (Peter) is forbidden to do something by an elder (Mrs Rabbit). The elder proceeds to leave home, providing an opportunity for the Hero to disobey. This he does, and the results are predictably bad for him (but not irrecoverably so). This is a very common plot-pattern, especially in fairy tales. It is dramatic insofar as it gives the reader an expectation danger ahead, and the moral satisfaction of seeing a disobedient character meet his come-uppance.

To model this pattern we would need to be able to represent the idea of generic roles that characters play in the narrative. Peter is the Protagonist or Hero, the principle character in the narrative. Mrs Rabbit is constructed as the Mentor, and her primary task is to take part in this scene. She will also re-enter at the end of the narrative, to conclude the pattern. Given these roles, we can encode the pattern as a set of abstract goals and actions for the Mentor and the Protagonist, which we can match with the concrete goals and actions of Mrs Rabbit and Peter. This idea of roles, as a collection of abstract objectives decoupled from any specific actor, resembles recent work in open agent societies (Davidsson 2001).

The expectation of danger in the scene is heightened by the production of an "implied narrative", springing from, and then running parallel to Peter's subsequent actions (Eco 1989). Mrs Rabbit's warning establishes a micro-narrative, which contains the basic story elements of Peter's later exploits. Each time the narrative reaches these elements in Peter's story, there is an implicit link made to the "Mr. Rabbit" micro-narrative: the sense of danger is formed by the reader's knowledge of Mr Rabbit's fate.

To model this pattern we must again turn to analogical reasoning, but in this case the analogy also contains a deliberate contrast. While Mr Rabbit is caught, Peter escapes, and this difference is important. In a sense, these two stories are symmetrical – the outcomes are not just different, but opposites. When we create analogies in our stories, such opposites can be as a kind of similarity.

For work on analogical reason in AI, we refer the reader to the long study of case-based reasoning (Falkenhainer, Forbus, & Gentner 1989; Riesbeck & Schank 1989). This work has been used in story-generation as a source of inspiration (Perez & Sharples 2001; Gervás *et al.* 2005) and also as a means of generating foreshadowing (Turner 1994).

Conclusion

In this paper we have identified the importance of narrative modeling as a foundation of understanding and generating stories. These two tasks are not discrete. There is a continuum of activity from explanation, through speculation to the production of narrative. All of these activities require sound models in order to meaningfully understand the reasons behind events in the story world.

Both narrative theory and established AI practice suggest a natural hierarchy of models, explaining events in terms of physical causes, characters' motives, and their role in the plot. Applying this hierarchy, our critical analysis of Beatrix Potter's *The Tale of Peter Rabbit* has provided an illustration of the kinds of knowledge needed in each level. To understand even a simple story like this, we need rich commonsense models of action in the world, the psychology of characters and the dramatic structure of story. We present these scenes as challenges for what we call 'story-sense reasoning'. Solving these challenges will involve drawing together work from different areas of AI and narrative theory, as we have tried to indicate. We believe that this is an exciting field with many possible avenues for exploration.

References

Barthes, R. 1982. Introduction to the structural analysis of narratives. In Sontag, S., ed., *A Barthes Reader*. Hill & Wang.

Blumberg, B.; Tomlinson, B.; and Downie, M. 2001. Multiple conceptions of character-based interactive installations. In *CGI '01: Proc. of the Int. Conf. on Computer Graphics*, 5. Washington, DC, USA: IEEE Computer Society.

Cavazza, M., and Pizzi, D. 2006. Narratology for interactive storytelling: A critical introduction. In *Proc. of the 3rd Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE)*, 72–83.

Charles, F.; Lozano, M.; Mead, S.; Bisquerra, A.; and Cavazza, M. 2003. Planning Formalisms and Authoring in Interactive Storytelling. In *Proc. of the 1st Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE)*.

Cheong, Y.-G. 2006. A computational model of narrative generation for suspense. In *Proc. of the 21st AAAI Computational Aesthetic Workshop*.

Davidsson, P. 2001. Categories of artificial societies. In *Proc. of the 2nd Int. Workshop on Engineering Societies in the Agents World II (ESAW '01)*, 1–9. London, UK: Springer-Verlag.

Davis, E. 1995. A highly expressive language of spatial constraints. Technical Report TR1995-714, NYU.

Dignum, F. 1999. Autonomous agents with norms. *Artificial Intelligence and Law* 7(1):69–79.

Eco, U. 1989. *The Open Work*. Harvard University Press. trans. Anna Cancogni.

Falkenhainer, B.; Forbus, K.; and Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence* 41(1):1–63.

Forster, E. M. 1927. *Aspects of the Novel*. San Diego: Harcourt.

Genette, G. 1980. *Narrative Discourse: An Essay in Method*. Cornell University Press. trans. J. E. Levin.

Gervás, P.; Díaz-Agudo, B.; Peinado, F.; and Hervás, R. 2005. Story plot generation based on CBR. *Knowledge-Based Systems* 18(4-5):235–242.

James, H. 1957. *The House of Fiction: Essays on the Novel*. London: Rupert Hart-Davis.

Levesque, H.; Pirri, F.; and Reiter, R. 1998. Foundations for a calculus of situations. *Electronic Transactions of AI* 2(3–4):159–178.

Morgenstern, L., and Miller, R. 2007. Common sense problem page. http://www-formal.stanford.edu/leora/commonsense/.

Perez, R., and Sharples, M. 2001. MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)* 13(2):119–139.

Potter, B. 1902. *The Tale of Peter Rabbit*. F Warne & Co. http://www.gutenberg.org/etext/14304.

Prince, G. 1982. *Narratology: The Form and Functioning of Narrative*. Berlin: Walter de Gruyter and Co.

Propp, V. A. 1968. *Morphology of the Folktale*. Publications of the American Folklore Society. University of Texas Press.

Randell, D. A.; Cui, Z.; and Cohn, A. 1992. A spatial logic based on regions and connection. In *KR'92. Principles of Knowledge Representation and Reasoning.*

Riedl, M. O. 2004. *Narrative generation: balancing plot and character*. Ph.D. Dissertation, North Carolina State University.

Riesbeck, C., and Schank, R. 1989. *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates.

Rimmon-Kenan, S. 1983. *Narrative Fiction: Contemporary Poetics*. London: Methuen.

Sadri, F., and Kowalski, R. 1995. Variants of the event calculus. In *Proc. of the 12th Int. Conf. on Logic Programming*.

Sgouros, N. M. 1999. Dynamic generation, management and resolution of interactive plots. *Art. Intel.* 107(1):29–62.

Shanahan, M. 1990. Representing continuous change in the event calculus. In *Proc. of European Conference on AI*, 598–603.

Shanahan, M. 2004. An attempt to formalise a non-trivial benchmark problem in common sense reasoning. *Artificial Intelligence* 153(1-2):141–165.

Shklovsky, V. 1965. Art as technique. In *Russian Formalist Criticism: Four Essays*. Lincoln: University of Nebraska Press. trans. Lee T. Lemon and Marion J. Reis.

Smith, R. G. 1980. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29(12):1104–1113.

Thomas, F., and Johnston, O. 1995. *The Illusion of Life: Disney Animation*. Disney Editions, revised edition.

Turner, S. 1994. *The Creative Process: a computer model of storytelling and creativity*. Lawrence Erlbaum Associates.

Woolridge, M. 2000. *Reasoning about Rational Agents*. MIT Press.

Young, R.; Riedl, M.; Branly, M.; Jhala, A.; Martin, R.; and Saretto, C. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. *Journal of Game Development* 1(1):51–70.