The Tale of Peter Rabbit: A Case-study in Story-sense Reasoning

Malcolm Ryan
Nicholas Hannah
School of Computer Science and Engineering
UNSW, Australia
malcolmr,nhannah@cse.unsw.edu.edu

Joshua Lobb
School of English
UNSW, Australia
j.lobb@unsw.edu.edu

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".

1. INTRODUCTION
Stories form a common part of our everyday lives as human beings. Not only do they form a large portion of our entertainment, in television, movies and books, but they also allow us to construct meaning from our experiences. 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 artificial intelligence representations of narrative would prove a useful project, building 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.

An AI representation of narrative could also be beneficial in human-computer interaction applications. Conversation is full of narratives (both complete and implied). A computer which was able to understand the narratives with which it is presented and express itself in a similar fashion could be a great improvement over current attempts at conversational interfaces.

Interactive Entertainment 2007, Melbourne Victoria Australia

All of 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, there are a number of patterns that regularly recur. 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 (?, ?, ?, ?). A single problem, such as the Surprise Birthday Present 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 couple of challenges to drive story-sense reasoning. Taking *The Tale of Peter Rabbit* by Beatrix Potter (?, ?) as our source, we present a number of short extracts from the text, which we shall proceed to analyse from the perspectives of narrative theory and AI. Where we can, we point to existing work in AI which addresses some of the problems raised, or else provide our own thoughts on how these problems might be approached. Our plan is to challenge some of the shortcomings of existing work and to inspire directions for future innovation.

The paper is laid out as follows: We will first lay out the foundations for what we see as the necessary components of an AI model, representing the world at the levels of action, characters and plot. We then describe four short scenes from Peter Rabbit and enumerate the different requirements they place on our models. Finally we conclude with our observations on what we have already achieved and what challenges lie ahead for the narrative AI community.

2. BACKGROUND

Before we can determine what our representation is going to be, we must first ask what we are going to use it for. What do we mean when we say that the AI "understands" a narrative? The following criteria seem reasonable requirements:

- The ability to *explain* a narrative, answering questions such as "Why 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 (?, ?) and Prince (?, ?), 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 (?, ?), 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 (?, ?, ?).

The interaction between this kind of 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 (?, ?) 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, desires 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 the Knowledge Representation and Reasoning community.

2.1 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. Depending on the genre, action in fictional narratives is not always naturalistic (e.g, in the real world, rabbits don't wear shoes) but in a well-constructed narrative it should at least be consistent with its own internal logic, and have a reasonable amount of overlap with our commonsense understanding.

There are many approaches to modeling physical systems, some of which, such as planning operators, are being put to good use for narrative generation (?, ?, ?). However these are often not expressive enough to cover the wide range of physical phenomena that occur even in seemingly simple scenarios (?, ?), such as cracking an egg into a bowl (?, ?) or, as we shall see below, getting stuck in a gooseberry net.

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* (?, ?, ?) and the *event calculus* (?, ?, ?). 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). These languages and associated programming and theorem proving tools (?, ?, ?) have been used for challenging applications such as high-level robot control and perception (?, ?, ?).

The commonsense reasoning community has constructed models of several different physical phenomena in languages such as these, including: topological and metric spaces (?, ?, ?), and time, concurrency and processes (?, ?, ?).

2.2 The Character Model

Characters are the driving forces of narrative. The drama of a narrative arises from the characters' desires and goals, how they act upon them, and whether those goals are frustrated or conflicted (?, ?). As James puts it, *what is character but the determination of incident? What is incident but the illustration of character?* 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*. As we will see, this will involve modeling their beliefs, goals, plans, emotions and other character traits which influence their action.

These qualities are the subject of research in a number of connected areas of AI: knowledge representation, intelligent agents and multi-agent systems (?, ?, ?), and naive psychology (?, ?). There exists a large number of well developed logical formalisations (?, ?), programming and specification languages (?, ?, ?), and software development frameworks (?, ?), which allow a system designer to appeal directly to mental attitudes.

Several narrative theorists have worked to distil the motives of fictional characters into character types' (?, ?, ?): an adaptation of these approaches might prove a manageable way of quantifying character intention.

2.3 The Plot Model

While a narrative can be told by simply describing the interactions of a group of characters in a world, most readers would complain that, as narrative, these sequences are not very interesting.[1] We expect stories to progress, that is, to have some kind of recognisable dramatic structure. Narrative theorists, particularly the Russian Formalists and Structuralists of the 20th century (?, ?, ?), have striven to understand what these structures are, and catalogue their usage.[2]

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 genuinely original narratives without requiring large amounts of human intervention, more general models of drama are needed. Some efforts that 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 (?, ?); the various works on generating folk tales, often based on the work of Propp (?, ?); limited work on suspense (?, ?).

3. THE CHALLENGES

We have taken four short passages from the Peter Rabbit text for examination. Each passage 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.

3.1 Scenario 1: 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.'*

Here 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 she 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.

Earlier we saw that 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 not only appear to have such mental attitudes but also to act based on the attitudes of others. Blumberg (?, ?) notes that this meta-level reasoning is essential to uphold the "illusion of life" (?, ?), needed to keep an audience engaged in a narrative. We beleive that this is a good example of how narrative modeling can provide new motivations and problems to extend mainstream agents research.

[1]Fans of Big Brother may disagree.
[2]Cavazza (?, ?) provides a concise summary for AI practitioners new to this field.

A long standing practical approach to defining and establishing relationships (with associated commitments etc.) involves the use of protocols, e.g. Contract-Net (?, ?). 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 (?, ?). 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.

FIXME: This isn't there.

AI approaches to modeling communication are disscussed below, when Peter has a narrow escape from an angry farmer.

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. For a rabbit, this is clearly a bad thing. 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 reasonably be stretched (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 (?, ?). 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 patten 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, which is primarily told from his perspective. Mrs Rabbit is a lesser character, we do not follow her actions in as much detail. Her role is to be the Mentor, and her primary task is to take part in this scene. She will also reenter 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 conception of roles, as a collection of abstract objectives decoupled from any specific actor, fits very closely with recent work in open agent societies (?, ?). Here the stated aim is to allow heterogenous participants to take up and enact roles in order to play a part in some structured interaction, for example, an auction (?, ?). Apart from carrying out their necessary function in modeling narrative, roles may be a useful software engineering tool. Role specifications and artificial actors could be combined in different ways and so reused for different scenarios.

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 (?, ?). Mrs Rabbit's warning establishes a micro-narrative, which contains the basic story elements of Peter's later exploits (the escape from mundane tasks, the exploits in Mr McGregor's garden, the 'accident' in the gooseberry net). 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 (?, ?, ?). This work has been used in story-generation as a source of inspiration (?, ?, ?) and also as a means of generating foreshadowing (?, ?).

3.2 Scenario 2: On the run

> *Mr. McGregor was on his hands and knees planting out young cabbages, but he jumped up and ran after Peter, waving a rake and calling out "Stop thief!"*
>
> *Peter was most dreadfully frightened; he rushed all over the garden, for he had forgotten the way back to the gate.*
>
> *He lost one shoe among the cabbages, and the other amongst the potatoes. After losing them, he ran on four legs and went faster so that I think he might have got away altogether if he had not unfortunately run into a gooseberry net and got caught by the large buttons on his jacket.*

The fears of the previous scenario now begin to be realised. Mr McGregor enters the story in the role of Antagonist[3] (see (?, ?)). His goals are in direct conflict with those of the Protagonist. In this case, he wants to protect the vegetables that Peter wants to eat. As a result, he wants to catch (and possibly kill) Peter, who now wants to get away. In a simple narrative, it is expected that the antagonist will almost achieve his desire, but will be foiled by the protagonist in the end.

To model this situation we need to recognise that an adversarial game is now being played between Peter and the farmer. Each one is acting based on the expectation of what the other will do. The drama lies in the uncertainty of who will succeed, and how. For the character model, this means that we need to represent reasoning about conflict, possibly as simultaneous move games. For the drama model, we need

[3] Propp refers to this role as the Villain, and to the Protagonist as the Hero, although in general these roles may not coincide.

to represent an uncertain situation as one with many different foreseeable outcomes, some good and some bad. The tension increases as the good outcomes become fewer and less likely.

The situation also affects character motives and therefore action. Peter is afraid. His fear and the urgency of the situation prevent him from thinking clearly. He acts less rationally and more reactively. If this were a game-tree search, we might say that he doesn't look very far ahead to decide what to do. He also forgets important information, like his sense of direction. In modeling agents, we are accustomed to making them as rational as possible, but in this instance we need to model an agent's irrationality. What is it reasonable for a character to forget? What kinds of mistakes are realistic under these circumstances? It is often their mistakes that make computers seem so alien, as they do not fail in the same kinds of ways as people.

Peter's final mistake of getting caught in the gooseberry net is not really his fault. It is not a normal effect of the action of running. Partly it can be explained by his confusion, but mostly it is just an accident. From the point of view of the action model, it is a very unlikely event, but it is clearly not random. It has been deliberately added by author for dramatic purposes.

This is a problem for our models. Stories need to contain unlikely events; accidents, coincidences or other fortuitous and infelicitous occurrences are the meat of many a tale. Therefore they need to exist as possibilities in our model of the world, but explicitly enumerating every possibility is not realistically possible. It is difficult to see how this might be resolved. Perhaps they could be modeled as emergent effects of a more general model at a lower level of abstraction. If this general model is sufficiently elaboration tolerant then it would allow for various outcomes (?, ?).

3.3 Scenario 3: Hiding

> *He rushed into the tool-shed and jumped into a can.*
>
> *It would have been a beautiful thing to hide in, if it had not had so much water in it. Mr. McGregor was quite sure that Peter was somewhere in the tool-shed, perhaps hidden underneath a flower-pot.*
>
> *He began to turn them over carefully, looking under each.*
>
> *Presently Peter sneezed "Kertyschoo!"*
>
> *Mr. McGregor was after him in no time, and tried to put his foot upon Peter, who jumped out of a window, upsetting three plants.*

This scenario introduces some important issues in reasoning about knowledge. In order to choose a place to hide, Peter needs to be able to reason about what Mr McGregor does and doesn't know. While Mr McGregor is outside the shed, he cannot see Peter and does not know what he is doing. Peter expects him to come into the shed, so he moves into the watering can, where the farmer will not be able to see him. Because this movement was made while the farmer was not looking, and because there are many possible places that Peter could have hidden, the farmer will not know where Peter is.

When the farmer enters the shed, he knows Peter is somewhere inside, but he doesn't know where. So he builds a plan to find out, by deliberately checking each hiding place in turn.

So our models of action and character overlap at this point. The physical structure of the world dictates the knowledge available to agents in the world. Peter is small, so he can fit in a watering can. The can is enclosed and opaque so that Peter cannot be seen from the outside. For the same reason, the pots are also possible hiding places. Mr McGregor knows this, so he looks inside each one to check.

Sight isn't the only way of locating someone. Unfortunately for Peter, sound also provides information and it obeys different rules. A loud sound, such as a sneeze, can be heard through a watering can, and can give your position away. To model these ideas we need to formalise quite complex rules of visibility and audibility, and how they affect character's knowledge.

Sensing, communication and their effect on knowledge are well studied (?, ?). However, the models are only able to make inferences based on the content of the communicative act. They do not take into account anything to do with the act itself, or the properties of the channel. Both of these things are crucial in this case: by being informed of 'sneeze' Mr. McGregor knows that Peter is close, and the fact that the channel was auditory rather than visual means that Mr. McGregor doesn't know Peter's exact location.

The dramatic character of this scene is one of suspense. No longer in the fast pace of the chase, here the narrative slows down. Peter is in a situation of impending danger, but is unable to do anything but wait. The reader must also wait, knowing the danger but enduring the delay while Mr McGregor examines each flowerpot. This combination of slowly encroaching danger, with no apparent escape that characterises a suspenseful situation. In narrative terms, this waiting is often characterised by a plotted 'stretch': the imagined slowing down of time for Peter is represented by a lengthening of the text, which slows down reading time (?, ?, ?).

The tension is resolved suddenly when Peter sneezes. He is found, but narrowly escapes by jumping through the window. Implicit in the idea of a narrow escape is the counterfactual: what if is happened otherwise? Counterfactuals like this appear implicitly and explicitly throughout the text. A narrative isn't just about the events that occur, but also about the events that might have occurred but didn't, as we will see more starkly in the next scenario.

For a comprehensive treatment of AI approaches to representing counterfactuals, see (?, ?).

3.4 Scenario 4: The white cat

> *Presently he came to a pond where Mr. McGregor filled his water-cans. A white cat was staring at some gold-fish; she sat very, very still,*

but now and then the tip of her tail twitched as if it were alive. Peter thought it best to go away without speaking to her.

He had heard about cats from his cousin, little Benjamin Bunny.

This seemingly irrelevant moment in the story exemplifies a prevalent, but almost unaccountable component of narrative construction. In any narrative, there is always a story element that has no function in the plot. The inclusion of the white cat here does not progress the plot: rather, in some ways, it obstructs it. Some narrative theorists have conjectured that these moments amplify the verisimilitude of the fictional work: that is, provide a sense of the believability of the story world, by providing superfluous details, just like real space (?, ?). The inclusion of such asides must temper our commitment to plot coherence and dramatic structure in our models. We must allow for the inclusion of irrelevant details.

Others have proffered that such a moment represents Socrates' notion of the 'aporia': a conundrum in the narrative, that creates more questions than it answers (?, ?). The 'explanation' of Peter's motives causes more uncertainty: what happened to Benjamin Bunny? In narrative terms, this creates a more intense production of the 'forbidden' scenario outlined above. The reader surmises several narratives (Benjamin is attacked by/deliberately taunts/seeks help from a cat), each with its own narrative outcome (Benjamin escapes/is injured/is killed). This amplifies the suspense of the situation: the reader is presented with the possibility of many more narrative choices (what if Peter didn't know Benjamin's story, and did speak to the cat? What if he continued to be naughty and spoke to the cat anyway? What if there are other dangerous animals?). The identification of such function is essential for the construction of AI narrative models.

This explosion of hypothetical narratives could be modeled by constructing many possible paths at the dramatic and character level, without presenting them in full physical detail. But this brings up a broader question of whether any part of a narrative sequence that is not going to be shown needs to be created in the first place. If it doesn't, this would seem to challenge the current practice of generating story independently of discourse because, in fact, discourse can drive story generation.

4. 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.[4]

[4]For readers who wish to know how Peter's story ends, a longer paper including a scene-by-scene analysis of the entire tale is forthcoming.