MAKEBELIEVE: Using Commonsense Knowledge to Generate Stories

Hugo Liu, Push Singh

MIT Media Laboratory
20 Ames Street 320D
Cambridge, MA, 02139 USA
{hugo, push}@media.mit.edu

Introduction

This paper introduces MAKEBELIEVE, an interactive story generation agent that uses commonsense knowledge to generate short fictional texts from an initial seed story step supplied by the user. A subset of commonsense describing causality, such as the sentence "a consequence of drinking alcohol is intoxication," is selected from the ontology of the Open Mind Commonsense Knowledge Base (Singh, 2002). Binary causal relations are extracted from these sentences and stored as crude trans-frames (Minsky, 1988). By performing fuzzy, creativity-driven inference over these frames, creative "causal chains" are produced for use in story generation. The current system has mostly local pair-wise constraints between steps in the story, though global constraints such as narrative structure are being added.

Our motivation for this project stems from two questions: Can a large-scale knowledge base of commonsense benefit the artificial intelligence community by supplying new knowledge and methods to tackle difficult AI problems? And given that full commonsense reasoning has yet to mature, can we demonstrate any success with a more "fail-soft" approach? We picked story generation because we feel it is a classic AI problem that can be approached with a creative use of commonsense knowledge. Compared to problem solving or question answering, story generation is a "softer" problem where there is no wrong solution per se, and that solution is evaluated subjectively.

This paper is organized as follows. In the first section we frame our approach in the context of previous work in story generation. In the second section, we present the system's representation of commonsense and techniques for fuzzy, creativity-driven inference. The third section presents an overview of the story generation architecture of MAKEBELIEVE. We conclude with an evaluation of the system followed by some discussion.

Previous Work

Previous work on story generation has generally taken one of two approaches: structuralist, and transformationalist. Structuralists such as Klein (1973, 1975) use real-world story structures such as canned story sequences and story grammars to generate stories. In contrast, transformationalists believe that story-telling expertise can be encoded by rules (Dreizin et al, 1978), or narrative goals (Dehn, 1981), which are applied to story elements such as setting and characters. The best example of this approach is TALE-SPIN (Meehan, 1977), which treats story generation as being analogous to problem solving. A story TALE-SPIN produces is essentially a description of the steps taken in the course of solving one or more problems.

MAKEBELIEVE inherits from both the structuralist and transformationalist traditions. First, MAKEBELIEVE is essentially a transformational story generator in its assumption that story generation is the result of simulation guided by rules manifested as local and global constraints. On the other hand, the commonsense facts about causality used in our simulation exhibit the property of being real-world story structures preferred by structuralists. The chief advantage of this is that the causal relationship between two events in a piece of commonsense is not bounded by any small set of simulation rules, but are instead bounded by the much greater variety of events and causal relations describable by commonsense knowledge.

Techniques for Using Commonsense

There are two large-scale knowledge bases of commonsense that we are aware of: Lenat's CYC (1995) and Open Mind Commonsense (OMCS). CYC contains over a million hand-crafted assertions, expressed in formal logic while OMCS has over 400,000 semi-structured English sentences, gathered through a web community of collaborators. Our current implementation uses Open Mind but CYC will be considered in the future. Sentences in OMCS are semi-structured, due to the use of sentence templates in the acquisition of knowledge, so it is relatively easy to extract relations and arguments. Compared with a logic representation, there is more semantic ambiguity associated with English sentences, but our fail-soft approach to commonsense inference is rather tolerant to ambiguity.

From the OMCS ontology, we selected a subset of 9,000 sentences that describe causation, such as the following:

- A consequence of bringing in a verdict is that the defendant is nervous

- Something that might happen when you act in a play is you forget your lines
- A consequence of eating in a fast food restaurant may be constipation

Before inference can be performed, we normalize the English sentences into a consistent form, for which we chose crude trans-frames, each with a before (cause) and after (effect) event, further decomposed into verb-object form with the help of a constituent structure parser. An example of sentence and its corresponding frame follows:

"The effect of keeping things orderly and tidy is living a better life." → {VERB: "keep" OBJS: "thing" MANNER: ("orderly", "tidy") EFFECT: "living a better life"}

Fuzzy, creativity-driven inference. Once we have a repository of trans-frames, we perform inference by trying to match the EFFECT of one frame to the CAUSE of some other frame. The heuristic for fuzzy matching is a scoring function that assigns points based on how closely the verb, object, and manner of two events are related through lexical semantics. We used WordNet nymic relations (Fellbaum, 1998) to measure semantic proximity between non-verbs, and Levin's verb classes (1993) in a similar way for verbs. While not perfect, using lexical semantics to connect related, but not identical ideas overcomes some of the brittleness associated with precise inference, and also has the effect of lending creativity to the storyline.

MAKEBELIEVE Architecture

For brevity, our system's architecture for story generation can be summarized into the following processing steps:

1. The user enters the first sentence of the story.

2. The sentence is parsed into verb-object form and fuzzy inference matches this initial event to the CAUSE slot of some trans-frame in the repository.

3. The EFFECT slot of the same trans-frame is parsed and inference continues, generating a whole chain of events.

4. After each step of inference, elements of the current story step are modified by analogous or synonymous elements taken from lexical semantic resources. This is to make sure that story steps deviate somewhat from the sometimes too logical causality that is characteristic of commonsense knowledge.

4. A global manager evaluates the chain of events to make sure it is free of cycles and contradictions, and if necessary it can backtrack to explore other storylines. Other global constraints such as narrative structure will be added here.

5. In cases where the inference chain is completely stuck, users may be asked to enter the next line in the story.

6. Frames of the inference chain and their corresponding sentences are used to generate English sentences, with the main character from the seed sentence being inserted. The structure and syntax of sentences are kept very simple in our current implementation.

An example. For brevity we present only one short story, which is typical of stories generated by MAKEBELIEVE.

John became very lazy at work. John lost his job. John decided to get drunk. He started to commit crimes. John went to prison. He experienced bruises. John cried. He looked at himself differently.

Conclusion

We have built MAKEBELIEVE, an interactive story generation agent that can generate short fictional texts of 5 to 20 lines when the user supplies the first line of the story. Our fail-soft approach to story generation represents a hybrid approach inheriting from both the structuralist and transformationalist traditions. It also incorporates a novel knowledge source, commonsense, which unlike other story knowledge bases, is not specifically purposed for story telling. Using a subset of knowledge in Open Mind, which describes causation, MAKEBELIEVE performs fuzzy and creative inference to generate casual chains, which become the basis for a storyline.

What sorts of limitations were encountered? The ambiguity inherent in any natural language representation makes it difficult to resolve the bindings of agents to actions when more than one agent is involved. For example, in this sentence from OMCS, "the effect of kicking someone is pain", we do not know enough to bind "pain" to the kicker or the kicked. This ambiguity precludes our current system from being able to tell multiple character stories.

A preliminary evaluation of MAKEBELIEVE was completed to serve as a baseline for future studies. 18 users were asked to judge the creativity, quality, and coherence of several five-line stories, generated as they interacted with the agent. On average, users scored the stories 10 out of a possible 15 points.

It was a pleasant surprise that despite being assembled out of commonsense knowledge, stories turned out to be much more interesting and dramatic from the user's perspective than we might have imagined. Furthermore, even though we did not add plot devices to the system such as motifs, climax, tension, etc., many users in our evaluation nonetheless felt that these devices were present in the generated stories. Where the generated story leaves off, the imagination of the reader seems ready to pick up.

Selected References

Klein, S. et al. 1973. *Automatic Novel Writing: A Status Report.* Wisconsin University

Meehan, J.R. 1977. TALE-SPIN, An interactive program that writes stories in *Proceedings of IJCAI-77*, pp. 91-98

Singh, P. (2002). The public acquisition of commonsense knowledge. *In Proceedings of AAAI Spring Symposium: Acquiring (and Using) Linguistic (and World) Knowledge for Information Access.* Palo Alto, CA, AAAI.