

Towards Representing What Readers of Fiction Believe

Toryn Q. Klassen and Hector J. Levesque and Sheila A. McIlraith
{**toryn**,hector,sheila}@cs.toronto.edu

Department of Computer Science
University of Toronto

November 6, 2017

Story understanding is a long-standing AI problem.

See e.g.

- *Toward A Model Of Children's Story Comprehension* by Charniak (1972)
- "An Example for Natural Language Understanding and the AI Problems It Raises" by McCarthy (1990)

Story understanding is a long-standing AI problem.

See e.g.

- *Toward A Model Of Children's Story Comprehension* by Charniak (1972)
- “An Example for Natural Language Understanding and the AI Problems It Raises” by McCarthy (1990)
- “Representing and Reasoning about Time Travel Narratives” by Morgenstern (2014)
- “Story Comprehension Through Argumentation” by Diakidoy, Kakas, Michael, and Miller (2014)
- “One Hundred Challenge Problems for Logical Formalizations of Commonsense Psychology” by Maslan, Roemmele, and Gordon (2015)

What does a reader, with given initial beliefs, come to believe after reading (part of) a fictional story?

Outline

1. Review of Friedman and Halpern (1999)
2. Literary logic
3. Applying literary logic to problems within story understanding
 - 3.1 carrying over real world knowledge into fiction
 - 3.2 how the reader expects information to be presented over the course of reading

Friedman and Halpern (1999) developed a logic for describing beliefs over time in dynamic systems.

- It's a **temporal** logic; at any time the accessible worlds are those that are consistent with the observations so far.
- Reasoning is **non-monotonic**; agents assign plausibilities to worlds, and believe what is true in the most plausible accessible worlds.

Literary logic is based on Friedman and Halpern's logic.

Notable differences are that

- Literary logic is **first order**.
- Instead of timelines defined by infinite sequences of observations, we have **finite timelines** defined by the finite sequence of sentences of a story.
- We use **more temporal operators** (the complete set from Lichtenstein et al. (1985)).

Furthermore,

- The way we define plausibility is a special case of theirs.
- We have two types of predicates, **real** and **imaginary**.

In literary logic:

- The reader considers a set of possible worlds, ordered by their plausibility.
 - We will be using special “**abnormality**” predicates in the reader’s **knowledge base** to define this ordering.
- Each world includes a **discourse** and specifies truth values for **real** and **imaginary** predicates.
 1. The **discourse** is a sequence of first order sentences, which can be thought of as standing in for the sentences of a natural language story.
 2. The **real** predicates describe things in the reader’s world.
 3. The **imaginary** predicates describe affairs within the world of the story.

Real predicates describe the real world, and imaginary predicates the world of the story.

Real predicates may, for example, be about the physical world, as with $Rabbit(x)$, or “literary” properties like the genre of the story being read, e.g. $FantasyGenre$.

Imaginary predicates may share names with real predicates, e.g. $Rabbit(x)$.

If α is a first-order sentence using only imaginary predicates, then

$$\mathbf{I}\alpha \quad (\text{“}\alpha \text{ is imagined”})$$

is a literary logic sentence.

Literary logic describes the beliefs of a reader.

There is a belief operator, \mathbf{B}_ψ , parameterized by a sentence ψ (which we will call the “knowledge base” or KB).

The agent described by \mathbf{B}_ψ initially considers possible exactly those worlds where ψ is true (again, the possible worlds are ordered by plausibility).

We will use prioritized abnormality predicates to define the plausibility ordering on worlds.

- Suppose that we have a finite set of distinguished **abnormality** predicates, each with an associated **priority**.
- Roughly, the more plausible worlds will be those with fewer abnormal objects (where higher priorities count for more).
- See the paper for details; the ordering is a prioritized version of the preference relation from **cardinality-based** circumscription (Liberatore and Schaerf, 1997; Moinard, 2000).

We will use prioritized abnormality predicates to define the plausibility ordering on worlds.

We will write

$$\Vdash \phi$$

if ϕ is true using this abnormality-based plausibility ordering.

Example

$$\Vdash \mathbf{B}_{(P \supset Ab)} \neg P$$

Literary logic's temporal operators describe time for the reader.

The '►' operator

► ϕ means that ϕ is true at the final time (when the entire discourse has been read).

In contrast, time within the story is not described with any special operators.

So, for example, if we had the story

John picked up a block.

Then he put it back down.

we might formalize it as the discourse

```
<  
  John(#1)  $\wedge$  Block(#2)  $\wedge$  Pickup(#1, #2, #3),  
  Precedes(#3, #4)  $\wedge$  Putdown(#1, #2, #4),  
  End  
>
```

Outline

1. Review of Friedman and Halpern (1999)
2. Literary logic
3. Applying literary logic to problems within story understanding
 - 3.1 carrying over real world knowledge into fiction
 - 3.2 how the reader expects information to be presented over the course of reading

World knowledge is important for understanding stories, but in fiction, it does not always apply.

- For example, in reality, (most) animals don't talk, but there are stories in which they do.
- Often, differences between the world of the story and the real world follow genre conventions.

Nonetheless, often world knowledge can be carried over.

The factual premisses [...] may carry over into the fiction, not because there is anything explicit in the fiction to make them true, but rather because there is nothing to make them false.

– Lewis (1978, p. 42)

For example, we may assume that humans are mortal in a story, because they are in reality.

How to model carry-over has not been much explored in AI (though see (Rapaport and Shapiro, 1995; Moorman and Ram, 1994)).

A first, syntactic, approximation to carry-over can be achieved by adding sentences to the KB.

- Suppose you have the KB $\psi = \forall \vec{x}(\phi(\vec{x})) \wedge \dots$
- Suppose $\mathbf{I}(\phi(\vec{x}))$ is also a formula
 - i.e., ϕ does not use modal operators or real predicate symbols for which there are not imaginary counterparts
- Then you could automatically generate a **defeasible imaginary copy** of ϕ ,

$$\forall \vec{x}(Ab(\vec{x}) \vee \mathbf{I}(\phi(\vec{x}))),$$

where Ab is some abnormality predicate not used in ψ .

- This new sentence could be conjoined with ψ .

Genre conventions may override carry-over.

For example, we may believe that dragons breathe fire even though no real animal does (Lewis, 1978).

That fictional dragons breathe fire can be expressed in literary logic.

$$Ab \vee \mathbf{I}(\forall x(Dragon(x) \supset BreathesFire(x))).$$

Here Ab would represent the abnormality of a story about dragons which didn't breathe fire.

- Ab would need sufficiently high priority so that this sentence would overrule knowledge about animals not breathing fire.
- A reader with this sentence in its KB could still regard real dragons that breathe fire as (even) less plausible than real dragons that don't.

Outline

1. Review of Friedman and Halpern (1999)
2. Literary logic
3. Applying literary logic to problems within story understanding
 - 3.1 carrying over real world knowledge into fiction
 - 3.2 how the reader expects information to be presented over the course of reading

Literary logic models the beliefs of a reader over time.

This allows for encoding knowledge about the organization of information presentation in stories.

Reading a mystery story, a reader may expect to (eventually) find out who is guilty.

- Suppose we have an imaginary predicate Guilty, with the intended interpretation that $\text{Guilty}(x)$ means that x is guilty.
- Suppose that ψ is the KB of a reader, and we want to inform this reader that they should expect to find out who is guilty.
- How should we modify ψ to accomplish this?

Reading a mystery story, a reader may expect to (eventually) find out who is guilty.

An extended KB can be written

$$\psi' = \psi \wedge \exists x(\blacktriangleright \mathbf{B}_{\psi} \mathbf{I} \text{Guilty}(x))$$

Under some conditions on ψ , a reader with KB ψ' believes they will find out who's guilty:

$$\|\sim \mathbf{B}_{\psi'} \exists x(\blacktriangleright \mathbf{B}_{\psi'} \mathbf{I} \text{Guilty}(x))$$

- See Proposition 1 in the paper for details.
- In the paper, we also consider how to tell a reader that the author is going to try to mislead them (see Proposition 2).

Conclusion

We've introduced literary logic, and applied it to

1. carry-over and genre knowledge, and
2. expectations about discourse organization.

In future work, we will

- further investigate carry-over
- construct fully worked-out examples with complete stories

We encourage investigating how other formal approaches developed for modelling belief over time, or non-monotonic reasoning, could similarly be useful in these issues.

References

- Eugene Charniak. Toward a model of children's story comprehension, 1972. MIT AI Laboratory Technical Report 266.
- Irene-Anna Diakidoy, Antonis Kakas, Loizos Michael, and Rob Miller. Story comprehension through argumentation. In *Computational Models of Argument - Proceedings of COMMA 2014*, pages 31–42, 2014.
- Nir Friedman and Joseph Y. Halpern. Modeling belief in dynamic systems, part II: Revision and update. *Journal of Artificial Intelligence Research (JAIR)*, 10:117–167, 1999.
- David Lewis. Truth in fiction. *American Philosophical Quarterly*, 15(1):37–46, 1978.
- Paolo Liberatore and Marco Schaerf. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence*, 93(1):261–296, 1997.
- Orna Lichtenstein, Amir Pnueli, and Lenore D. Zuck. The glory of the past. In *Proceedings of the Conference on Logic of Programs*, pages 196–218, 1985.
- Nicole Maslan, Melissa Roemmele, and Andrew S. Gordon. One hundred challenge problems for logical formalizations of commonsense psychology. In *Commonsense 2015*, pages 107–113, 2015.
- John McCarthy. An example for natural language understanding and the AI problems it raises. In Vladimir Lifschitz, editor, *Formalizing Common Sense: Papers by John McCarthy*, pages 70–76. Ablex Publishing Corporation, 1990.
- Yves Moinard. Note about cardinality-based circumscription. *Artificial Intelligence*, 119(1):259 – 273, 2000.
- Kenneth Moorman and Ashwin Ram. A model of creative understanding. In *AAAI 1994*, pages 74–79, 1994.
- Leora Morgenstern. Representing and reasoning about time travel narratives: Foundational concepts. In *KR 2014*, pages 642–645, 2014.
- William J. Rapaport and Stuart C. Shapiro. Cognition and fiction. In Judith F. Duchan, Gail A. Bruder, and Lynne E. Hewitt, editors, *Deixis in Narrative: A Cognitive Science Perspective*, pages 107–128. Lawrence Erlbaum Associates, Inc., 1995.