Story Understanding... Calculemus!

Loizos Michael
Open University of Cyprus
loizos@ouc.ac.cy

Abstract

Stories lie at the heart of the human experience, and efforts to comprehend the latter can be aided by investigating how humans understand stories. We take a computational view of this problem, and suggest that a form of logic can be fruitfully used to represent and reason with stories. Instead of seeking to define what counts as story understanding, we investigate tasks typically associated with story understanding (e.g., question answering), and demonstrate how those can be formalized and tackled within the logic-based framework we propose.

Introduction

Stories are ubiquitous, as evidenced by their widespread use across human cultures (e.g., flood and creation myths), and the prominence of story reading and understanding as objectives across the educational ladder. It is, thus, natural and desirable to investigate how to build machines that understand stories, both as a means to understand humans themselves, but also as a way to improve human-machine interactions.

Efforts have been underway in the past few years to bring together scientists from Computer Science, Linguistics, Philosophy, and Psychology to study narrative from a computational perspective (Finlayson, Richards, and Winston 2010; Finlayson 2011). Although research in Computational Models of Narrative (ComMoN) is at its early stages, a community started forming, pooling insights and knowhow from a diverse set of research fields with the aim to pursue this goal. This work seeks to contribute in that direction by putting forward a particular logic-based model-theoretic computational framework for problems pertaining to story understanding.

Rather than adopting any one particular, and ultimately ad hoc, definition of what a story is, or what it means to understand one, we embrace the inclusive view of Altman (2008) that *"virtually any situation can be invested with [those] characteristics [necessary to] perform the narrational function"*. We investigate, instead, certain tasks that one may undertake in relation to stories: *(i)* identifying coherence, *(ii)* answering questions, *(iii)* summarizing, and *(iv)* comparing for similarity. Our emphasis is in making these tasks precise in the context of the proposed framework, and showing them to be amenable to a formal and computational treatment.

We acknowledge upfront that others may have a different view than ours on what these tasks mean, and what properties a solution to these tasks should have. Our intention here is to present plausible interpretations for the chosen tasks, and argue that the approach we follow in tackling them is sufficiently powerful and flexible to accommodate other nuances of the tasks that other scholars may wish to consider.

It is also not our intention to deal with issues particular to any modality for communicating stories (e.g., text, comics, video). Each modality carries its own deep questions, which are targeted by large research communities. Rather, we focus on aspects surrounding story understanding once stories are encoded in a formal language. We suggest that logic can usefully fulfill the role of this formal language, and that existing knowhow from research on logic-based knowledge representation and commonsense reasoning (Mueller 2006) can be brought to bear on the problem of story understanding.

Formal Representation

> *Alice and Bob meet one day around noon. Alice asks Bob if he has had lunch yet, and upon receiving an affirmative reply from Bob, Alice says: "Too bad. I was thinking of taking you out for lunch.". Bob proceeds to ask Alice whether herself has had lunch, and upon receiving a negative reply from Alice, Bob says: "Too bad. I was thinking of taking you out for coffee.".*

Although the story above is presented in text, this is done purely for the benefit of the readers of this work. Henceforth we shall assume that *every story is given in a formal representation that is, itself, the direct object of investigation*.

The Story Discourse

What would be an appropriate formal representation for our example story? At the very least, the representation should capture the discourse: the objects of interest and the relations that hold between them, the events that occur according to the story, and the scenes across which the story unfolds. We shall represent scenes by abstract states S_i, and use clauses to encode facts we are told hold or events we are told occur within the story. The story representation could then be

time(noon) `holds-at` S_1
meeting(alice,bob) `holds-at` S_1
ask(alice,bob,lunch) `occurs-at` S_1
reply(bob,yes) `occurs-at` S_2
say(alice,bob,think(alice,bob,lunch)) `occurs-at` S_3

ask(bob,alice,lunch) `occurs-at` S_4
reply(alice,no) `occurs-at` S_5
say(bob,alice,think(bob,alice,coffee)) `occurs-at` S_6

along with the information that scenes are ordered linearly in the story. This last piece of information is represented by using a partial ordering $\preceq^s$ and by writing that $S_i \preceq^s S_j$ for every $i < j$. We repeat that we do not consider whether this formal representation is in any sense "appropriate", but we take the representation as the object of interest itself; the formal representation *is* the story that we are considering.

The story employs grounded predicates (e.g., time(noon), reply(alice,no)), effectively drawn from a finite dictionary over which stories are defined (much like what is done when writing stories in any natural language). We denote this dictionary by the tuple $\langle \mathcal{F}, \mathcal{A} \rangle$, where $\mathcal{F}$ includes all statements used to name fluents / observations (e.g., time(noon)), and $\mathcal{A}$ includes all statements used to name actions / events (e.g., reply(alice,no)) in stories. The language is sufficiently expressive and natural to represent higher-order modalities that capture what the story characters believe, desire, plan, think, say, etc. (e.g., say(bob,alice,think(bob,alice,coffee))).

We, thus, define a ***discourse*** $\mathbf{C}$ as a triple $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ comprising a finite set $\mathcal{C}$ of clauses of the form A `occurs-at` S and L `holds-at` S, a finite set $\mathcal{S}$ of states, and an acyclic partial ordering $\preceq^s$ over $\mathcal{S}$, where $A \in \mathcal{A}$, L is a literal over $\mathcal{F}$, and $S \in \mathcal{S}$. These will be the objects of our investigation.

Reader's Knowledge

Theories of reading acknowledge that readers employ certain knowledge to understand stories (Anderson and Pearson 1984; Pressley et al. 1992). We take the view that this knowledge expresses rules, using which the reader draws inferences given a story. For our example story, a particular reader could employ, amongst other rules, the following:

`static:` lying(X,Y) $\rightarrow$ $\neg$honest(X,Y)
`static:` say(X,Y,S) $\rightarrow$ ($S \oplus$ lying(X,Y))
`static:` think(X,Y,Z) $\wedge$ $\neg$realizable(X,Y,Z) $\rightarrow$ lying(X,Y)
`static:` (think(X,Y,lunch) $\wedge$ answer(Y,lunch,yes)) $\rightarrow$ $\neg$realizable(X,Y,lunch)
`static:` (think(X,Y,coffee) $\wedge$ answer(Y,lunch,no)) $\rightarrow$ $\neg$realizable(X,Y,coffee)
ask(X,Y,Q) `causes` question(Y,Q)
(reply(Y,A) $\wedge$ question(Y,Q)) `causes` $\neg$question(Y,Q)
(reply(Y,A) $\wedge$ question(Y,Q)) `causes` answer(Y,Q,A)

As in the case of discourses, we do not concern ourselves here with the question of where such background knowledge comes from. Instead, our focus is on how such background knowledge is to be understood, and how a story is to be interpreted given a reader's background knowledge (whatever it might be). We shall return to this issue in the last section.

A reader's background knowledge encodes rules relating various fluents and actions to each other. These rules are of two main types: *(i)* Static rules stipulate constraints that hold between predicates at any given state. So, according to the second rule above: if X says S to Y, then either S is the case or X is lying to Y. *(ii)* Causal rules stipulate cause and effect relations between predicates across states. So, according to the last rule above: if Y replies with A, and Y had been asked Q, then A is caused to be Y's answer to Q. The second to last rule above captures the complimentary effect that once Y replies to a question Q that Y had been asked, then question Q is retracted. These rules can be taken to correspond to the reader's commonsense knowledge regarding conversations. As such, they do not purport to be infallible, but only sufficiently accurate for commonsense reasoning.

For succinctness we have employed a high-level language that allows predicates with variables. Its syntax can be made precise by treating, for instance, each variable as a type with a fixed domain. So, variable X could be set to range over the values $\{\text{alice}, \text{bob}\}$. The values of certain variables may involve other predicates, which themselves may employ variables. So, variable S could include the value think(X,Y,W) in its domain, which should be thought itself as standing for the set of all of its grounded instances, as determined by the values that X,Y,W can take. Due to space constraints we shall forego the precise formulation of this high-level language, and we shall appeal to the intuitive reading of these rules as a short-hand notation for the set of all their grounded instances (over the grounded predicates in $\langle \mathcal{F}, \mathcal{A} \rangle$).

We, thus, define a ***domain*** as a finite set $\mathcal{D}$ of clauses of the form `static:` φ and φ `causes` L, where L is a literal over $\mathcal{F}$, and φ is a formula over $\mathcal{A} \cup \mathcal{F}$.

Further Extensions

Other than general rules, a reader's knowledge may include facts or actions that occurred in real-life. Stories (especially factual ones) may rely on the reader having and using such knowledge. To account for such knowledge, we shall assume a ***time-line*** $\langle \mathcal{T}, \preceq^t \rangle$ comprising a countable set $\mathcal{T}$ of time-points, and a well-order relation $\preceq^t$ on it. We shall not explicitly reference $\langle \mathcal{T}, \preceq^t \rangle$ in the sequel. Constructs dealing with time will be implicitly taken to be defined over $\langle \mathcal{T}, \preceq^t \rangle$.

The ***domain*** definition can be extended to allow inclusion of clauses of the form A `occurs-at` T, L `holds-at` T, where $A \in \mathcal{A}$, L is a literal over $\mathcal{F}$, and $T \in \mathcal{T}$.

Assuming $\langle \mathcal{T}, \preceq^t \rangle$ represents dates, our example domain could include kill(brutus,caesar) `occurs-at` 15/03/44BC and $\neg$alive(caesar) `holds-at` 16/03/44BC. Stories read given this domain would need to be understood in the context of, and in temporal relation to, these time-specific facts and actions. Each possible embedding of a story in the time-line could lead to a different reading of the story, so each such embedding should be considered, as long as it respects the ordering that the story itself imposes on its discourse.

We, thus, define an ***embedding of*** a given discourse $\mathbf{C} = \langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ as a set $\mathcal{B}$ of clauses of the form A `occurs-at` T and L `holds-at` T that results by replacing each state in $\mathcal{S}$ with a time-point in $\mathcal{T}$, so that if time-points $T_1, T_2 \in \mathcal{T}$ replace states $S_1, S_2 \in \mathcal{S}$ with $S_1 \preceq^s S_2$, then $T_1 \preceq^t T_2$.

Other extensions are also possible. Discourses could be extended to include time-specific facts and actions, or even static and causal rules. Some stories (fables and fictional stories, in particular) communicate not only what happens or holds, but also the rules under which a reader is to make sense of them; e.g., "Anytime he would rub the magic lamp, a genie would appear.". The inclusion of such rules in fables is often prompted by phrases like "Once upon a time in a faraway land...", alerting the reader to suppress her domain knowledge in favor of that found in the story. Some of these story-specific rules, or other rules in the reader's knowledge, may be geared towards allowing the reader to infer the scene ordering, instead of assuming that it is given by the story.

Knowledge vs Beliefs

We have already alluded to the fact that not all of a reader's knowledge is strict and indisputable. Some of it may encode beliefs, biases, or maybe even guesses, which the reader may be willing to abandon while trying to make sense of a story.

One such belief is that humans usually do not lie — one can certainly understand why this rule should be treated as defeasible! To accommodate this belief, the reader may include `static:` ¬lying(X,Y) in the domain $\mathcal{D}$ that she employs. She may further specify that if only a subset of $\mathcal{D}$ were to be used — if using the entire $\mathcal{D}$ would preclude making sense of a story — then that subset should include all rules other than `static:` ¬lying(X,Y). In doing so, the reader indicates that this rule is weak compared to the others.

More generally, the reader may specify a preference $\preceq^d$ over subsets of $\mathcal{D}$ to indicate how to suppress rules if necessary. If certain "indisputable" rules (e.g., definitions) should never be suppressed, then this can be accommodated by not specifying a preference over subsets that exclude those rules.

We, thus, define a ***default domain*** $\mathbf{D}$ as a triple $\langle \mathcal{D}, \mathbb{D}, \preceq^d \rangle$ comprising a domain $\mathcal{D}$, a set $\mathbb{D} \subseteq 2^{\mathcal{D}}$ of subsets of $\mathcal{D}$, and a transitive preference relation $\preceq^d$ over $\mathbb{D}$.

One can informally verify that to make sense of our example story, one can respect all the rules in the example domain presented earlier, but cannot respect at the same time the rule `static:` ¬lying(X,Y). We examine rule satisfaction later.

Among the pieces of defeasible knowledge that a reader may have, of particular interest is one related to the celebrated Frame Problem: things change over time only if there is a reason (McCarthy and Hayes 1969). It is the case that a story may specify the actions that occur and change things. But a story may also simply present a new scene, and the reader is expected to accept that certain things have changed despite the lack of mention of an action to that effect.

Our example story, for instance, could proceed with "That night, Alice called Bob.", with the discourse expanded to include time(night) `holds-at` S_7, and the domain expanded to include `static:` ¬(time(noon) ∧ time(night)). Now, the story states that it was (and still is, since nothing happened to change it) noon and now it is night as well; this conflicts with the reader's knowledge. The only way to resolve the conflict is to suppress `static:` ¬(time(noon) ∧ time(night)), which is absurd. Clearly, the intention is not for both noon and night to hold, but for the former to change to the latter.

The issue is naturally resolved if, for every literal L over $\mathcal{F}$ and time-point T, we include exogenous(L) `causes` L and exogenous(L) `occurs-at` T in the domain, and we also extend the preference $\preceq^d$ so that *ceteris paribus* it gives preferences to subsets of the domain that *do not* include exogenous(L) `occurs-at` T. Since the latter is actually a collection of rules (one for each L and T), then $\preceq^d$ seeks to minimize their use, capturing the belief that exogenous actions should not be assumed unless if needed to understand a story; in that case the action causes L to come about. In our example, the subset of the domain to be used is that with exogenous(¬time(noon)) `occurs-at` T_1 and exogenous(time(night)) `occurs-at` T_2 for appropriate T_1, T_2.

Story Interpretation

Assume we are given the discourse of a story and the default domain that a reader employs to make sense of it. How do the two interact? How does a reader understand a story?

Degree of Coherence

An important aspect of story understanding — irrespectively of whether it occurs unconsciously or not in the mind of the reader — that one may attempt to formalize, is to examine a story for coherence. According to text linguistics and discourse analysis (de Beaugrande and Dressler 1981), coherence relates to the semantical meaningfulness of (communicative) text, such as the degree of abiding to conventions.

Such conventions in our framework are captured by the background knowledge of the reader. The degree of abiding to these conventions can be formalized as the extent to which the rules in the reader's default domain are satisfied.

We have already introduced a notion of *how much* of a default domain is satisfied, in terms of a preference relation $\preceq^d$ over its subsets. We have also discussed how a story discourse can be embedded in a time-line, so that it can be meaningfully considered along with a domain. It remains to define what it means for the rules in a domain to be *satisfied*.

Recall that the union of a domain with the embedding of a discourse results in a set of clauses that is itself a domain. It is this resulting domain that we seek to satisfy. Intuitively, its satisfaction amounts to ensuring that action occurrences and fact observations are respected, static constraints are satisfied (logic-theoretically), and causal change is brought about whenever its conditions are met (and only then). We adopt a model-theoretic treatment of this problem, following work in logic-based commonsense reasoning (Mueller 2006).

An ***assignment*** is a mapping M of each pair $X \in \mathcal{A} \cup \mathcal{F}$ and $T \in \mathcal{T}$ to a truth-value $M(X,T)$. The truth-assignment over $\mathcal{A} \cup \mathcal{F}$ that is induced by projecting / restricting the mapping M to a given time-point $T \in \mathcal{T}$ is denoted by $M(T)$.

Since language $\langle \mathcal{F}, \mathcal{A} \rangle$ uses only grounded predicates, it is effectively propositional. We shall use, then, the entailment operator $\models$ of Propositional Calculus as usual. Since $\mathcal{T}$ is countable and $\preceq^t$ is a well-ordering, we shall write $T+1$ to mean the time-point in $\mathcal{T}$ that follows T according to $\preceq^t$.

Definition 1 (Model of Domain)**.** *A **model of** a given domain $\mathcal{D}$ is an assignment M such that for each $A \in \mathcal{A}$, each literal L over $\mathcal{F}$, each formula φ over $\mathcal{A} \cup \mathcal{F}$, and each $T \in \mathcal{T}$, the following conditions hold:*

(i) $M(T) \models A$ *if and only if* A `occurs-at` $T \in \mathcal{D}$.

(ii) $M(T) \models L$ *if* L `holds-at` $T \in \mathcal{D}$.
$M(T) \models \varphi$ *if* `static:` $\varphi \in \mathcal{D}$.

(iii) $M(T+1) \models L$ *if* $M(T) \models \varphi$ *for some* φ *such that* φ `causes` $L \in \mathcal{D}$.

(iv) $M(T+1) \models L$ *if* $M(T) \models L$ *and* $M(T) \not\models \varphi$ *for every* φ *such that* φ `causes` $\neg L \in \mathcal{D}$.

A domain $\mathcal{D}$ *is* ***satisfiable*** *if there exists a model of* $\mathcal{D}$.

Coherence can then be defined. A discourse $\mathbf{C}$ is ***coherent with*** a given domain $\mathcal{D}$ if there exists an embedding $\mathcal{B}$ of $\mathbf{C}$, whose union with $\mathcal{D}$ results in a satisfiable domain. We shall call embedding $\mathcal{B}$ a ***witness to the coherence of*** $\mathbf{C}$ ***with*** $\mathcal{D}$.

In terms of our expanded example story, it can be verified that the story's discourse is not coherent with the domain $\mathcal{D}_1$ that includes `static:` $\neg$(time(noon) $\wedge$ time(night)), but it is coherent with any one of its subsets that suppresses (i.e., does not include) that particular rule. Furthermore, it can be verified that the story's discourse is coherent with the domain $\mathcal{D}_2 \supset \mathcal{D}_1$ that also includes exogenous(L) `causes` L and exogenous($\neg$time(noon)) `occurs-at` T_1 and exogenous(time(night)) `occurs-at` T_2 for appropriate T_1, T_2, as we have intuitively explained in the previous section.

Since there exist infinitely many embeddings in the general case (i.e., when $\mathcal{T}$ is infinite), it is not evident that checking for coherence is decidable. The next result shows that indeed it is, and offers an explicit prescription on which embeddings suffice to be considered to determine coherence.

Theorem 1 (Decidability of Coherence)**.** *Checking whether a discourse* $\mathbf{C}$ *is coherent with a domain* $\mathcal{D}$ *is decidable.*

Proof. Let T_0 be the largest time-point in $\mathcal{D}$. For each embedding $\mathcal{B}$ of $\mathbf{C}$, let $T_{\mathcal{B}}$ be the set of time-points in $\mathcal{B}$ after time-point T_0. It can be shown that: if there is an embedding $\mathcal{B}_1$ of $\mathbf{C}$ such that $\mathcal{D} \cup \mathcal{B}_1$ is satisfiable and $\max(T_{\mathcal{B}_1}) \geq T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}|$, then there is an embedding $\mathcal{B}_2$ of $\mathbf{C}$ such that $\mathcal{D} \cup \mathcal{B}_2$ is satisfiable and $\max(T_{\mathcal{B}_2}) < \max(T_{\mathcal{B}_1})$.

To check for coherence, it suffices to check for embeddings of $\mathbf{C}$ that include time-points only before $T_0 + (2^{|\mathcal{F}|} + 2) \cdot |\mathcal{C}|$, and whose union with $\mathcal{D}$ is satisfiable. Checking all the said embeddings can be done in finite time. $\square$

NB: Theorem 1 offers the basis to establish the computability of tasks considered later; i.e., the relevant notions are decidable. We shall not formally state and prove these corollaries.

Of course a reader does not specify a domain $\mathcal{D}$, but a default domain $\mathbf{D} = \langle \mathcal{D}, \mathbb{D}, \preceq^d \rangle$. Depending on which subset of $\mathcal{D}$ a discourse is coherent with, we can derive a qualitative measure of coherence of the discourse with $\mathbf{D}$. In particular, we can show that we can order any set of discourses in terms of their degree of coherence with a given default domain.

Drawing Inferences

A story might be coherent with multiple subsets of a reader's background knowledge. Intuitively, the reader uses those of the subsets that are most preferred; or, those subsets with which the story is maximally coherent. There is no need, for instance, to give up the belief that "if Alice and Bob talk to each other, then they speak the same language", if doing so is not necessary to make sense of our example story.

Since preference over the subsets of a default domain $\mathbf{D}$ is determined by a preference relation $\preceq^d$, we should, then, seek the maximal elements of $\preceq^d$, after restricting our attention to subsets with which a given discourse is coherent.

We, thus, define a discourse $\mathbf{C}$ to be ***maximally coherent*** with a given domain $\mathcal{D}_1$ ***under*** a default domain $\mathbf{D}$ if: *(i)* $\mathbf{C}$ is coherent with $\mathcal{D}_1 \in \mathbb{D}$; and *(ii)* for every $\mathcal{D}_2 \in \mathbb{D}$ with which $\mathbf{C}$ is coherent, if $\mathcal{D}_1 \preceq^d \mathcal{D}_2$ then $\mathcal{D}_2 \preceq^d \mathcal{D}_1$.

The existence of multiple subsets of $\mathcal{D}$ with which a discourse $\mathbf{C}$ is maximally coherent accounts for the fact that a story may have multiple readings, depending on which parts of background knowledge the reader may decide to suppress in favor of others. Even if some subset is fixed, a story may still have multiple readings depending on the embedding the reader considers in relation to time-specific facts and actions in her background knowledge. These combinations amount to the possible interpretations of a story by the reader.

Definition 2 (Interpretation)**.** *An* ***interpretation of*** *a given discourse* $\mathbf{C}$ ***under*** *a default domain* $\mathbf{D}$ *is a model of the domain that results from the union of a domain* $\mathcal{D}_1$ *and an embedding* $\mathcal{B}$ *of* $\mathbf{C}$*, such that* $\mathbf{C}$ *is maximally coherent with* $\mathcal{D}_1$ *under* $\mathbf{D}$*, and* $\mathcal{B}$ *is a witness to the coherence of* $\mathbf{C}$ *with* $\mathcal{D}_1$.

Each interpretation assigns a truth-value to every fluent and action in the language $\langle \mathcal{F}, \mathcal{A} \rangle$ for every time-point, in a way that satisfies both the story (i.e., its embedding $\mathcal{B}$) and the background knowledge of the reader to the extent possible (i.e., the subset $\mathcal{D}_1$ with which the story is maximally coherent). In a precise sense, then, each interpretation is a way to complete all information that is not explicit in the story, in a maximally plausible (according to the reader) manner.

Answering Questions

Question answering is a typical way to evaluate story understanding. We consider below questions that inquire whether a certain fact holds / action occurs in the context of a story. In our example story, such a question could be "Does Alice stand 10 meters away from Bob when she first speaks to him?". A correct answer is not explicitly provided by the story, but a reader is reasonably expected to offer one.

The example question above can be formalized as a clause distance(alice,bob,10m) `holds-at` S_1, where S_1 is the state (or scene) in which Alice first speaks to Bob according to the example story's discourse. As in the case of stories, we do not consider issues surrounding the extraction of this formal representation of the question from its textual form. Rather, the clause *is* the question that we are considering.

In general, a question is a clause, either X `holds-at` S or X `occurs-at` S.[1] We shall write $Q(X,S)$ to denote a question, where S belongs in the set $\mathcal{S}$ of states of the discourse $\langle \mathcal{C}, \mathcal{S}, \preceq^s \rangle$ of interest, and X is the fluent or action of our inquiry. Answering the question amounts to checking whether it follows from the interpretations of the story.

Recall that an interpretation refers to time-points. What time-point does the state S of $Q(X,S)$ correspond to? Put differently, at which time-point in an interpretation should one look to answer $Q(X,S)$? We reason thus: Each interpretation is associated with an embedding $\mathcal{B}$ of the discourse. The embedding implicitly maps each state S to a time-point $t(S)$. Mapping t gives the appropriate time-point for each S.

[1]Composite questions can be treated in an analogous manner.

Definition 3 (Question Answering)**.** *A question $Q(X, S)$ is **possibly true in** a given discourse $\mathbf{C}$ **under** a default domain $\mathbf{D}$ if there exists an interpretation M of $\mathbf{C}$ under $\mathbf{D}$ with $M(t(S)) \models X$, for the mapping t prescribed by M. A question $Q(X, S)$ is **certainly true in** $\mathbf{C}$ **under** $\mathbf{D}$ if it is possibly true in $\mathbf{C}$ under $\mathbf{D}$, but the question $Q(\neg X, S)$ is not.*

Our framework can deal with multiple-choice questions as well. If exactly one choice is certainly true, then this is the answer. Otherwise, none of the choices is by itself sufficiently plausible (in an absolute sense) in the story, so the reader seeks to find the most plausible (in a relative sense) among them. To do so, one extends the story with each of the choices in turn, and ranks the resulting versions of the story in terms of their degree of coherence. The answer is the choice included in the most coherent version of the story.

Encoding Expectations

When reading a story, a reader may have certain expectations. *(i)* Reader-independent: The story author (or critic, or narrator) might externally prescribe certain inferences that readers are expected to draw. So, if told that one is expected to infer that Alice and Bob are dishonest to each other in our example story, then this affects the way the story is read. *(ii)* Reader-specific but story-independent: Some reader might expect, even before reading, that stories have a certain structure. The expectation to have named characters, for instance, is met by our example story. *(iii)* Reader-specific and story-specific: While reading a story, a reader might infer what to subsequently expect (or even attempt to infer what the story author expects readers to infer). For instance, when reading Alice's response "Too bad. I was thinking of taking you out for lunch." in our example story, the reader may expect the story to clarify how this unresolved situation will conclude.

Whether expectations of the first type (e.g., authorial intentions) can exist has been at the center of a long debate by those that study the interpretation of language. One school of thought insists that an author's intent is the *only* way that a story should be understood (Michaels and Knapp 1982), or at least that it points to the *right* way among possibly many ways (Hirsch 1967). Another school of thought dismisses such a view as a fallacy (Wimsatt and Beardsley 1946), and decrees that the author of a story is irrelevant (Barthes 1967), and the interpretation of a story relies solely on the story itself and the reader's background knowledge and beliefs.

A solution to the full-fledged problem of dealing with expectations of all types is sufficiently involved so that it cannot be presented to a reasonable level of clarity herein. We shall restrict our attention to types *(i)* and *(ii)*, which can still encode very general and powerful expectations, while also representing both sides of the debate mentioned above.

Type *(ii)* expectations can be naturally included in the default domain $\mathbf{D}$ that represents a reader's background knowledge, with the preference $\preceq^d$ over rules in $\mathbf{D}$ applying also to expectations. Much like domain rules, some weak expectations may need to be suppressed to make sense of a story (perhaps even because the story author planned, through the choice of the discourse, for this to happen). In our example story, the expectation that the story clarifies at each instance whether it is day or night could be represented by including expect[time(day) $\texttt{holds-at}$ T] $\vee$ expect[time(night) $\texttt{holds-at}$ T] for each time-point T, and asking that $\preceq^d$ prefers the inclusion of these disjunctions.

Type *(i)* expectations (e.g., authorial intentions) should be thought as accompanying a story, even if not part of it per se. To account for their default nature (in a manner analogous to the treatment of type *(ii)* expectations), we associate with each story a new structure $\mathbf{E}$, defined as a triple $\langle \mathcal{E}, \mathbb{E}, \preceq^e \rangle$ comprising a set $\mathcal{E}$ of expectations, a set $\mathbb{E} \subseteq 2^{\mathcal{E}}$ of subsets of $\mathcal{E}$, and a transitive preference relation $\preceq^e$ over $\mathbb{E}$.

Each expectation is a propositional formula ε over terms expect[κ], where κ is any type of clause found in a discourse or domain, respectively, for type *(i)* or type *(ii)* expectations.

The expectations $\mathcal{E}$ for our example story could include

expect[¬honest(alice,bob) $\texttt{holds-at}$ S_3] $\wedge$
expect[¬honest(bob,alice) $\texttt{holds-at}$ S_6]
expect[¬friend(alice,bob) $\texttt{holds-at}$ S_1] $\wedge$
expect[¬friend(bob,alice) $\texttt{holds-at}$ S_1]
expect[$\texttt{static:}$ (meeting(alice,bob) $\wedge$ time(noon))
$\rightarrow$ think(alice,bob,lunch)]

with $\mathbb{E} \subseteq 2^{\mathcal{E}}$ comprising the subsets of $\mathcal{E}$ that include at least the first expectation, and with $\preceq^e$ being the subset relation. These choices would capture the expectation (e.g., authorial intention) that the reader infers that: Alice and Bob are being dishonest to each other at the time they speak; they are not friends at the start of the story; every time they meet and it is noon, Alice thinks of taking Bob out for lunch.[2] The former expectation would be considered (e.g., by the author) key to the story, whereas the latter two would be less important and the reader would not miss too much if she did not meet them.

To account for the type *(i)* expectations $\mathbf{E} = \langle \mathcal{E}, \mathbb{E}, \preceq^e \rangle$ associated with a story $\mathbf{C}$ and / or the type *(ii)* expectations in a default domain $\mathbf{D} = \langle \mathcal{D}, \mathbb{D}, \preceq^d \rangle$, we shall proceed as follows: First, we shall consider a single subset $\mathcal{E}_i \in \mathbb{E}$ of the expectations associated with $\mathbf{C}$, and define when these expectations are met. We shall also assume, for simplicity of the discussion, that all expectations in $\mathcal{E}_i$ and $\mathbf{D}$ are singletons expect[κ] rather than general propositional formulas.

Ignoring all the expectations in $\mathbf{D}$, consider an interpretation M of $\mathbf{C}$ under $\mathbf{D}$. By Definition 2, M specifies a domain $\mathcal{D}_1 \in \mathbb{D}$ and an embedding $\mathcal{B}$ of $\mathbf{C}$, whose union is used to compute M. Now, $\mathcal{D}_1$ might include certain (reader-specific) type *(ii)* expectations $\mathcal{E}_{ii}$ that must be met along with the (reader-independent) type *(i)* expectations in $\mathcal{E}_i$. Although the latter ones reference states instead of time-points, we can map those states into time-points using the mapping t prescribed by M; we write $\mathcal{E}_i^t$ to mean the resulting set. Thus, $\mathcal{E}_{ii} \cup \mathcal{E}_i^t$ is the set of all expectations to be met by M.

We continue to check whether M is a model of the domain $\mathcal{D}_1 \cup \mathcal{B} \cup \{\kappa \mid \text{expect}[\kappa] \in \mathcal{E}_{ii} \cup \mathcal{E}_i^t\}$. In other words, we check whether M, which already satisfies all clauses in $\mathcal{D}_1 \cup \mathcal{B}$ according to Definition 1, also satisfies all clauses in the expectations in $\mathcal{E}_{ii} \cup \mathcal{E}_i^t$. If this happens to be the case, then say that M ***meets the expectations*** $\mathcal{E}_i$ ***for*** $\mathbf{C}$; that is, M

[2]Note the exceptionally powerful expectation that the reader infers a rule about the story-world through some process of induction.

meets both the expectations $\mathcal{E}_i$ that are externally provided, and those $\mathcal{E}_{ii}$ that are present in the subset $\mathcal{D}_1$ that was used to interpret $\mathbf{C}$. If every interpretation meets the expectations $\mathcal{E}_i$ for $\mathbf{C}$, then say that $\mathbf{D}$ ***meets the expectations*** $\mathcal{E}_i$ ***for*** $\mathbf{C}$.

If it happens that $\mathbf{D}$ does not meet the expectations $\mathcal{E}_i$ for $\mathbf{C}$, for any $\mathcal{E}_i \in \mathbb{E}$ (i.e., $\mathbf{D}$ does not meet even the most important expectations), then this may prompt the reader to dismiss the story, or at least to seek to understand why this is so. Note that if one wishes not to deal with type *(i)* expectations at all, then this can be done by simply choosing $\mathcal{E}$ to be empty, and $\mathbb{E} = \{\emptyset\}$. In that case, our approach will make sure that only the reader-specific expectations must be met.

Intra-Story Reasoning

Certain story understanding tasks may involve multiple stories. We have already seen this phenomenon when comparing coherence degrees across stories to reply to multiple-choice questions. Below we consider two more tasks that fall in this category: story summarization and story similarity.

Story Summarization

Intuitively, a summary of a story is a second story that abstracts the former, but not too much. The theory of plot units (Lehnert 1981) takes "abstraction" to mean the identification of the plot units of the story, and "too much" to mean retaining most (or at least the pivotal) plot units in the summary.

Our view of "abstraction" and "too much" acknowledges the role of the reader's background knowledge in the summarization process, and the expectations that are present.

There are arguably many ways to define abstractions in the context of stories. We suggest two general and flexible criteria: a syntactic one and a semantic (reader-specific) one.

Definition 4 (Abstraction)**.** *A discourse* $\mathbf{C}_1 = \langle \mathcal{C}_1, \mathcal{S}_1, \preceq_1^s \rangle$ *is an* ***abstraction of*** *a given discourse* $\mathbf{C}_2 = \langle \mathcal{C}_2, \mathcal{S}_2, \preceq_2^s \rangle$ *under a default domain* $\mathbf{D}$ *if one of the following holds*

(i) ***syntactic abstraction****:* $\mathcal{C}_1 \subseteq \mathcal{C}_2$, $\mathcal{S}_1 \subseteq \mathcal{S}_2$, $\preceq_1^s \subseteq \preceq_2^s$*;*

(ii) ***semantic abstraction****: every interpretation of* $\mathbf{C}_2$ *under* $\mathbf{D}$ *is also an interpretation of* $\mathbf{C}_1$ *under* $\mathbf{D}$*;*

or if there is a discourse $\mathbf{C}_3$ *such that:* $\mathbf{C}_1$ *is an abstraction of* $\mathbf{C}_3$ *under* $\mathbf{D}$*, and* $\mathbf{C}_3$ *is an abstraction of* $\mathbf{C}_2$ *under* $\mathbf{D}$.

Syntactic abstraction amounts to omitting parts of a story, whereas semantic abstraction amounts to rewriting the story to make it less committing to what "happened"; i.e., to have more (or the same) possible interpretations. In reference to our example story, omitting time(noon) `holds-at` S_1 from the discourse is a syntactic abstraction, whereas replacing it with time(day) `holds-at` S_1 is a semantic one, since the resulting story admits more interpretations, assuming the default domain includes `static:` time(noon) $\rightarrow$ time(day).

The "too much" part in summarization is naturally accommodated by asking that the resulting story does not ignore (all of) the expectations associated with the original story.

We, thus, define a discourse $\mathbf{C}_1$ as a ***credulous summary of*** a given discourse $\mathbf{C}_2$ ***under*** a default domain $\mathbf{D}$ if: *(i)* $\mathbf{C}_1$ is an abstraction of $\mathbf{C}_2$ under $\mathbf{D}$; and *(ii)* $\mathbf{D}$ meets some expectations $\mathcal{E}_i \in \mathbb{E}$ for $\mathbf{C}_1$, according to structure $\langle \mathcal{E}, \mathbb{E}, \preceq^e \rangle$ associated with $\mathbf{C}_2$. If $\mathcal{E}_i = \mathcal{E}$, the summary is ***skeptical***.

Considering the particular expectations $\mathcal{E}_i$ that are met for $\mathbf{C}_1$, the more important (according to the preference $\preceq^e$) they are, the more skeptically $\mathbf{C}_1$ summarizes $\mathbf{C}_2$. This metric of evaluating summaries measures only the extent to which they preserve the important (according to the associated type *(i)* expectations) information of the original story. Although we shall not do so here, one can clearly introduce additional metrics (e.g., how concise a summary is, how well it meets the reader's type *(ii)* expectations) and seek summaries that balance the chosen metrics in some way.

One can verify that the next discourse is a credulous summary of the last discussed version of our example story:

$\neg$honest(alice,bob) `holds-at` S_1
$\neg$honest(bob,alice) `holds-at` S_2

Identifying Similarity

Structure-mapping theory (SMT) (Gentner 1983) views two situations as analogous if they share n-ary relations that hold across objects, but not (necessarily) unary relations that hold on individual objects. This same idea has been argued to account to some extent for the task of determining similarity between stories (Gentner, Rattermann, and Forbus 1993).

Our view ties similarity to having a common summary. Unlike SMT, this view makes explicit the subjectivity that the reader's background knowledge imposes on the notion of similarity. Equally importantly, our view of similarity is insensitive to the granularity of story representation — since summaries deal with excess granularity — an issue that has been argued to be problematic in SMT (Chalmers, French, and Hofstadter 1992) and other frameworks that view similarity as an isomorphism (Löwe 2010). The interplay of abstraction and similarity also relates to the views of Schank (1990), that abstractions of stories are performed unconsciously by humans for comprehension, and that stories one hears are mapped to those that one knows. This mapping, we suggest, happens when the stories are similar (in our sense), i.e., when a common abstraction / summary is found.

The proposed view of similarity is, we believe, plausible, and is offered as a hypothesis for further study and empirical psychological validation. In the context of this work we shall be content to make this particular hypothesis concrete.

We, thus, define discourses $\mathbf{C}_1$ and $\mathbf{C}_2$ to be ***credulously / skeptically similar under*** $\mathbf{D}$ if there exists a common credulous / skeptical summary $\mathbf{C}$ of $\mathbf{C}_1$ and $\mathbf{C}_2$ under $\mathbf{D}$. Intermediate degrees of similarity follow from degrees of how skeptical a summary is: the more skeptical the common summary of the stories (i.e., the more important the expectations of the stories that are met by the summary), the stronger their similarity. If no common (credulous) summary exists, then the stories are too different to be considered similar.

Beyond the direct comparison of similarity between two stories, we also consider a relative similarity test that seeks to determine which of two stories is more similar to a target story (see, e.g., (Falkenhainer, Forbus, and Gentner 1989)). As is natural given the approach we have followed so far, we suggest that more similar to the target story is the story that has a more skeptical common summary with the target story.

We, thus, define, discourse $\mathbf{C}_1$ to be ***weakly more similar***

than discourse $\mathbf{C}_2$ ***to*** discourse $\mathbf{C}_0$ ***under*** $\mathbf{D}$ if there exists a common summary $\mathbf{C}_{10}$ of $\mathbf{C}_1$ and $\mathbf{C}_0$ under $\mathbf{D}$, such that: for every common summary $\mathbf{C}_{20}$ of $\mathbf{C}_2$ and $\mathbf{C}_0$ under $\mathbf{D}$, $\mathbf{C}_{20}$ is not a more skeptical than $\mathbf{C}_{10}$ summary of $\mathbf{C}_0$ under $\mathbf{D}$. $\mathbf{C}_1$ is ***(strictly) more similar than*** $\mathbf{C}_2$ ***to*** $\mathbf{C}_0$ ***under*** $\mathbf{D}$ if $\mathbf{C}_1$ is weakly more similar than $\mathbf{C}_2$ but not vice versa.

Conclusions

Formally modeling story understanding may help offer: *(i)* hypotheses on how humans deal with stories, and *(ii)* a basis for developing machines to understand stories. We hope that this work touched upon a sufficient variety of issues and to a sufficient depth to aid in pushing forward both frontiers.

In the first direction, we have suggested concrete hypotheses, such as that a form of logic is appropriate for reasoning with stories, or that story similarity derives from stories having a common summary. The extent to which our hypotheses capture what humans do when reading stories is a matter of psychological validation, which we are actively pursuing.

A first study we have performed offers some evidence for the adequacy of logic. In our experiments we encoded stories from elementary school textbooks, and domains with commonsense knowledge aimed to answer simple true / false questions. The responses offered by the reasoning module were contrasted against those offered by elementary school students, yielding an agreement rate of around 90%; interestingly, in cases of disagreement it was typically the students, not the reasoning module, that offered a wrong answer.

In other recent work (Kypridemou and Michael 2013) we have investigated the "similarity as common summary" hypothesis. In our experiments we presented human participants with triples of stories, and asked them to score the similarity of the two first stories, and also the extent to which the third story was a good summary of each of the first two stories. These three scores were aggregated across a number of trials and under varying conditions. Analysis of the results revealed that our hypothesis is strongly supported.

In the second direction, our use of logic and the decidability results that can be derived for the various tasks, already offer a concrete basis for their mechanization. Of course, the ultimate goal in this direction is for machines to deal with stories provided to them not in formal logic, but in a modality used by humans, such as natural language text. A lot of work exists towards translating text to predicates (Collins 1999; de Marneffe, MacCartney, and Manning 2006; Punyakanok, Roth, and tau Yih 2008), for temporally anchoring or ordering statements in text (Mani, Schiffman, and Zhang 2003), and so on. Such works can be brought to bear to extract the discourse of stories in an automated manner.

Regarding the acquisition of appropriate default domains, one can envision two approaches: First, through the use of a crowdsourcing platform where humans contribute appropriate knowledge, by taking into account lessons learned from Cyc (Lenat 1995) and OpenMind (Stork 1999). Second, by capitalizing textual corpora (including text found on the Web) to extract factual statements (Etzioni et al. 2005; Carlson et al. 2010) or rules (Michael and Valiant 2008). Learning-theoretic frameworks offer the substrate to extract both static rules (Valiant 1984; Michael 2010) and causal rules (Michael 2011), as well as to extract preferences between such rules (Dimopoulos and Kakas 1995).

Rules extracted from the Web encode *websense*, a certain form of commonsense knowledge, which has been used, in particular, for drawing inferences that follow — in a precise sense (Michael 2008; 2009) — from pieces of text (Michael 2013). Adopting the use of such rules for story understanding would tackle the problem of identifying simultaneously the background knowledge needed to reason with stories in a formal language, and the knowledge needed to make sense of the natural language text in which stories are encoded.

The central role of background knowledge for story understanding is acknowledged also in some earlier work of the Computational Models of Narrative community (Verheij 2009; Mueller 2009). The latter work, in particular, offers a model-theoretic semantics to domains, like we do (cf. Definition 1). Unlike those works, we consider tasks other than question answering, and support general rules in stories, defeasible beliefs in domains, author and reader expectations.

In terms of developing concrete story understanding tools, argumentation offers a natural means to implement the preference relation among domains that our work assumes, and to offer a computationally efficient reasoning mechanism for the tasks considered herein. Such an approach is taken in recent work (Diakidoy et al. 2013), where argumentation is investigated as a substrate for psychologically valid narrative text comprehension for the task of question answering.

Our particular take on story understanding echoes Grice's Maxims (1991) on certain presumptions that readers rely on (Bach 2005): the Quality Maxim is echoed by the interpretation of stories according to a reader's held truth (background knowledge), the Relation and Quantity Maxims are echoed by the expectations of the reader, and the Manner Maxim is echoed by the non-ambiguity of the formal language. At the same time, purposeful ambiguity in the sense of a story having multiple interpretations (cf. pluralism) and priorities among them (Levinson 2003) is fully accommodated.

This work's aim was to tell a coherent story of how logic can plausibly offer a computational basis to certain tasks of relevance to story understanding. Extending the framework to apply to a wider list of such tasks, and to further deal with their intricacies, is part of another story waiting to be told!

References

Altman, R. 2008. *A Theory of Narrative*. Columbia University Press.

Anderson, R. C., and Pearson, P. D. 1984. *Handbook of Reading Research*. Longman. chapter A Schema-Theoretic View of Basic Processes in Reading Comprehension.

Bach, K. 2005. The Top 10 Misconceptions About Implicature. In Birner, B., and Ward, G., eds., *A Festschrift for Larry Horn*. John Benjamins.

Barthes, R. 1967. Death of the Author. *Aspen* 5–6.

Carlson, A.; Betteridge, J.; Kisiel, B.; Settles, B.; Hruschka Jr., E. R.; and Mitchell, T. M. 2010. Toward an Architecture for Never-Ending Language Learning. In *Proc. of 24th AAAI Conference on Artificial Intelligence (AAAI'10).*

Chalmers, D. J.; French, R. M.; and Hofstadter, D. R. 1992. High-Level Perception, Representation, and Analogy: A Critique of Artificial Intelligence Methodology. *Journal of Experimental and Theoretical Artificial Intelligence* 4(3):185–211.

Collins, M. 1999. *Head-Driven Statistical Models for Natural Language*. Ph.D. Dissertation, University of Pennsylvania, U.S.A.

de Beaugrande, R.-A., and Dressler, W. U. 1981. *Introduction to Text Linguistics*. Longman.

de Marneffe, M.-C.; MacCartney, B.; and Manning, C. D. 2006. Generating Typed Dependency Parses from Phrase Structure Parses. In *Proc. of 5th International Conference on Language Resources and Evaluation (LREC'06)*.

Diakidoy, I.-A.; Kakas, A.; Michael, L.; and Miller, R. 2013. Narrative Text Comprehension: From Psychology to AI. In *Proc. of 11th International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense'13)*.

Dimopoulos, Y., and Kakas, A. 1995. Learning Non-Monotonic Logic Programs: Learning Exceptions. In *Proc. of 8th European Conference on Machine Learning (ECML'95)*, 122–137.

Etzioni, O.; Cafarella, M.; Downey, D.; Popescu, A.-M.; Shaked, T.; Soderland, S.; Weld, D.; and Yates, A. 2005. Unsupervised Named-Entity Extraction from the Web: An Experimental Study. *Artificial Intelligence* 165(1):91–134.

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

Finlayson, M. A.; Richards, W.; and Winston, P. H. 2010. Computational Models of Narrative: Review of a Workshop. *AI Magazine* 31(2):97–100.

Finlayson, M. A. 2011. Reports of the AAAI 2010 Fall Symposia: Computational Models of Narrative. *AI Magazine* 32(1):96–97.

Gentner, D.; Rattermann, M. J.; and Forbus, K. D. 1993. The Roles of Similarity in Transfer: Separating Retrievability from Inferential Soundness. *Cognitive Psychology* 25:524–575.

Gentner, D. 1983. Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science* 7(2):155–170.

Grice, P. 1991. *Studies in the Way of Words*. Harvard University Press.

Hirsch, E. D. 1967. *Validity in Interpretation*. Yale University Press.

Kypridemou, E., and Michael, L. 2013. Narrative Similarity as Common Summary. In *Proc. of 4th Workshop on Computational Models of Narrative (ComMoN'13)*.

Lehnert, W. G. 1981. Plot Units and Narrative Summarization. *Cognitive Science* 4:293–331.

Lenat, D. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *Communications of the ACM* 38(11):33–38.

Levinson, J. 2003. *The Oxford Handbook of Aesthetics*. Oxford University Press. chapter 1, 3–24.

Löwe, B. 2010. Comparing Formal Frameworks of Narrative Structure. In *Proc. of 2nd Workshop on Computational Models of Narrative (ComMoN'10)*.

Mani, I.; Schiffman, B.; and Zhang, J. 2003. Inferring Temporal Ordering of Events in News. In *Proc. of Human Language Technology Conference (HLT-NAACL'03)*, 55–57.

McCarthy, J., and Hayes, P. 1969. Some Philosophical Problems from the Standpoint of Artificial Intelligence. *Machine Intelligence* 4:463–502.

Michael, L., and Valiant, L. G. 2008. A First Experimental Demonstration of Massive Knowledge Infusion. In *Proc. of 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, 378–389.

Michael, L. 2008. *Autodidactic Learning and Reasoning*. Ph.D. Dissertation, Harvard University, U.S.A.

Michael, L. 2009. Reading Between the Lines. In *Proc. of 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, 1525–1530.

Michael, L. 2010. Partial Observability and Learnability. *Artificial Intelligence* 174(11):639–669.

Michael, L. 2011. Causal Learnability. In *Proc. of 22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, 1014–1020.

Michael, L. 2013. Machines with Websense. In *Proc. of 11th International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense'13)*.

Michaels, W. B., and Knapp, S. 1982. Against Theory. *Critical Inquiry* 8(4):723–742.

Mueller, E. T. 2006. *Commonsense Reasoning*. Morgan Kaufmann.

Mueller, E. T. 2009. Story Understanding Through Model Finding. In *Proc. of 1st Workshop on Computational Models of Narrative (ComMoN'09)*.

Pressley, M.; Wood, E.; Woloshyn, V. E.; Martin, V.; King, A.; and Menke, D. 1992. Encouraging Mindful Use of Prior Knowledge: Attempting to Construct Explanatory Answers Facilitates Learning. *Educational Psychologist* 27(1):91–109.

Punyakanok, V.; Roth, D.; and tau Yih, W. 2008. The Importance of Syntactic Parsing and Inference in Semantic Role Labeling. *Computational Linguistics* 34(2):257–287.

Schank, R. 1990. *Tell Me A Story: A New Look at Real and Artificial Memory*. Scribners.

Stork, D. 1999. The Open Mind Initiative. *IEEE Expert Systems and Their Applications* 14(3):16–20.

Valiant, L. G. 1984. A Theory of the Learnable. *Communications of the ACM* 27(11):1134–1142.

Verheij, B. 2009. Argumentation Schemes, Stories and Legal Evidence: A Computational Perspective. In *Proc. of 1st Workshop on Computational Models of Narrative (ComMoN'09)*.

Wimsatt, W. K., and Beardsley, M. C. 1946. The Intentional Fallacy. *Sewanee Review* 54:468–488.