

A Psychology-Inspired Approach to Automated Narrative Text Comprehension

Irene-Anna Diakidoy

University of Cyprus
eddiak@ucy.ac.cy

Antonis Kakas

University of Cyprus
antonis@cs.ucy.ac.cy

Loizos Michael

Open University of Cyprus
loizos@ouc.ac.cy

Rob Miller

University College London
rsm@ucl.ac.uk

Abstract

We report on an ongoing research program to develop a formal framework for automated narrative text comprehension, bringing together know-how from research in Artificial Intelligence and the Psychology of Reading and Comprehension. It uses argumentation to capture appropriate solutions to the frame, ramification, and qualification problems, and their generalizations as required for text comprehension. In this first part of the study we concentrate on the central problem of integration of the explicit information from the text narrative with the reader's implicit commonsense world knowledge, and the associated tasks of elaboration and revision.

Introduction

Text comprehension has long been identified as a key test for Artificial Intelligence (AI), with a central position in many forms of the Turing Test, and enormous benefits in human-computer interaction. The rise of computing over the Internet, where so much data is in the form of textual information, has given even greater importance to this topic. We report on a research program aiming to learn from the extensive text comprehension work in Psychology, and draw guidelines for developing frameworks for automated narrative text comprehension and, in particular, *story comprehension* (SC).

Our research program brings together know-how from research in Psychology and AI — in particular, our understanding of Reasoning about Actions and Change and Argumentation — to provide a formal representational and computational framework for SC that can be empirically evaluated and iteratively developed. This empirical evaluation, which forms an important part of the program, is based on the following methodology: (i) set up a corpus of stories with questions to test different aspects of story comprehension; (ii) harness the world knowledge used by human readers for comprehension; (iii) use this world knowledge in our framework and automated system, and compare its comprehension behaviour with that of the human readers.

The Psychology of Story Comprehension

Comprehending text entails the construction of a mental representation of the information contained in the text. How-

ever, no text specifies clearly and completely all the implications of text ideas or the relations between them. Therefore, comprehension depends on the ability to mentally represent the text-given information and to generate *bridging and elaborative inferences* that connect and elaborate text ideas resulting in a mental or *comprehension model* of the story. Inference generation is necessary in order to comprehend any text as a whole, i.e., as a single network of interconnected propositions instead of as a series of isolated sentences, and to appreciate the suspense and surprise that characterize narrative texts or stories, in particular (Brewer and Lichtenstein 1982; McNamara and Magliano 2009).

Although inference generation is based on the activation of background *world knowledge*, the process is constrained by text information. Concepts encountered in the text activate related conceptual knowledge in the readers' long-term memory (Kintsch 1988). Nevertheless, at any given point in the process, only a small subset of all the possible knowledge-based inferences remain activated and become part of the mental representation: those that connect and elaborate text information in a way that contributes to the *coherence* of the mental model (Rapp and den Broek 2005; McNamara and Magliano 2009). Inference generation is a task-oriented process that follows the principle of *cognitive economy* enforced by a limited-resource cognitive system.

However, the results of this coherence-driven selection mechanism can easily exceed the limited working memory capacity of the human cognitive system. Therefore, coherence on a more global level is achieved through *higher-level integration* processes that operate to create macro-propositions that generalize or subsume a number of text-encountered concepts and the inferences that connected them. In the process, previously selected information that maintains few connections to other information is dropped from the mental model. This results in a more consolidated network of propositions that serves as the new anchor for processing subsequent text information (Kintsch 1998).

Comprehension also requires an iterative general *revision mechanism* of the readers' mental model. The feelings of suspense and surprise that stories aim to create are achieved through discontinuities or changes (in settings, motivations, actions, or consequences) that are not predictable or are wrongly predictable solely on the basis of the mental model created so far. Knowledge about the structure and the func-

tion of stories leads readers to expect discontinuities and to use them as triggers to revise their mental model (Zwaan 1994). A change in time or setting in the text may serve as a clue for revising parts of the mental model while other parts remain and are integrated with subsequent text information.

The interaction of bottom-up and top-down processes for coherence carries the possibility of *different* but equally *legitimate or successful comprehension* outcomes, due to the qualitative and quantitative differences in the conceptual and mental state knowledge of different readers. Comprehension is successful if these differences are primarily in elaboration.

Scope of the Present Paper

We concentrate on the development of an appropriate Reasoning about Actions and Change and Default Reasoning framework for representing narratives and the *world knowledge* needed for the central comprehension process of *synthesizing and elaborating* the explicit text information with new inferences, and *revising* them in the presence of new narrative information. Our working hypothesis is that higher level features of comprehension, such as coherence and cognitive economy, can be tackled on top of the framework we develop. We are also assuming as solved the issue of correctly parsing the natural language of the text into some information-equivalent structured (e.g., logical) form of the story narrative, without discounting the importance of this latter problem, nor the possibility of the need to be tackled in conjunction with the problems on which we are focusing.

An excerpt from one of the stories used for the initial evaluation of our approach, is used as our running example.

Story excerpt: [...] *He aimed at the first turkey, and pulled the trigger. After a moment's thought, he opened his shotgun and saw there were no bullets in the shotgun's chamber. He loaded his shotgun, aimed at the turkey and pulled the trigger again. Undisturbed, the bird nearby continued to chirp and build its nest. Papa Joe was very confused. Would this be the first time that his shotgun had let him down?*

KRR for Story Comprehension

The close link between human commonsense reasoning, such as that used for SC, and argumentation has been recently re-enforced by new psychological evidence (Mercier and Sperber 2011) suggesting that human reasoning is in its general form inherently argumentative. In our proposed approach of KRR for SC, the reasoning to construct a comprehension model and its qualification and revision at all levels as the story unfolds, will be captured through a uniform acceptability requirement on the arguments that support the conclusions in the model. We use methods and results from Argumentation Theory in AI, e.g., (Dung 1995; Modgil and Prakken 2012), and its links to the area of Reasoning about Action and Change (RAC) with Default Reasoning on the static properties of domains (see (van Harmelen, Lifschitz, and Porter 2008) for an overview).

In particular, we use a typical RAC language of Fluents, Actions, and Times, to represent the narrative $\mathcal{N}$ of a story, and also the world knowledge $\mathcal{W}$ used for the story's comprehension. The exact time-points are mostly inconsequential, and are meant to stand for the abstract scenes in a story.

The narrative of our story excerpt is represented as a sequence of timed-observations (*pj* is short for "Papa Joe"):

OBS(<i>alive</i> (<i>turkey</i>), 1),	OBS(<i>aim</i> (<i>pj</i> , <i>turkey</i>), 1),
OBS(<i>pull_trigger</i> (<i>pj</i>), 1),	OBS($\neg$ <i>gun_loaded</i> , 4),
OBS(<i>load_gun</i> , 5),	OBS(<i>pull_trigger</i> (<i>pj</i>), 6),
OBS(<i>chirp</i> (<i>bird</i>), 10),	OBS(<i>nearby</i> (<i>bird</i>), 10).

World knowledge is represented as a collection of unit-arguments, in the form of simple associations between concepts in the language. This stems from a key observation in Psychology that typically all world knowledge, and irrespective of type, is inherently default. It is not in the form of an elaborate formal theory of detailed definitions of concepts, but rather is better regarded as a collection of relatively loose semantic associations between concepts, reflecting typical rather than absolute information. Thus, knowledge need not be fully qualified at the representation level, since it can be qualified via the reasoning process by the relative strength of other (conflicting) pieces of knowledge. This addresses the *endogenous qualification problem*, while the implicit priority of information in the narrative over any unit-argument also addresses the *exogenous qualification problem*.

The world knowledge needed to make sense of our story excerpt is given below, with the priorities $p1 \succ c1$, $p2 \succ c1$:

$c1 : \text{cau}(\text{fired_at}(pj, X), \{\text{aim}(pj, X), \text{pull_trigger}(pj)\})$
$c2 : \text{cau}(\neg \text{alive}(X), \{\text{fired_at}(pj, X), \text{alive}(X)\})$
$c3 : \text{cau}(\text{noise}, \{\text{fired_at}(pj, X)\})$
$c4 : \text{cau}(\neg \text{chirp}(\text{bird}), \{\text{noise}, \text{nearby}(\text{bird})\})$
$c5 : \text{cau}(\text{gun_loaded}, \{\text{load_gun}\})$
$p1 : \text{pro}(\neg \text{fired_at}(pj, X), \{\neg \text{gun_loaded}\})$
$p2 : \text{pro}(\neg \text{fired_at}(pj, X), \{\neg \text{noise}\})$ (story-specific)

Each unit-argument is of the general form $\text{arg}(H, B)$, with a literal (a fluent / action or its negation) H for its head, and a set of literals B for its body. Unit-arguments are able to capture information pertaining to how properties are caused to come about ($\text{cau}(H, B)$), how they relate to each other ($\text{pro}(H, B)$), and how they persist ($\text{per}(H, \{H\})$); unit-arguments of the latter type are implicitly assumed to be present for each literal H . Conflicts between unit-arguments are resolved based on the priority relation $\succ$ that holds between them. In general, the priority relation $\succ$ includes (i) $\text{cau}(H, B_1) \succ \text{per}(\neg H, B_2)$; (ii) $\text{per}(H, B_1) \succ \text{pro}(\neg H, B_2)$; and (iii) any given story-specific priorities.

The priority amongst these basic units of knowledge gives a form of non-monotonic reasoning (NMR) for deriving new properties that hold in the story. In particular, priorities in (i) address the *frame problem*, ensuring that properties cease to persist when there is causal evidence to the contrary. When we need to reason with defeasible property information, such as default rules about the normal state of the world in which a story takes place, we are also faced with a *generalized frame problem*, where "a state of the world persists irrespective of the existence of general state laws". Hence, if we are told that the world is in fact in some exceptional state that violates a general (default) property, this will continue to be the case in the future, until we learn of (or derive) some causal information that returns the world into its normal state. The solution to this generalized frame problem is

captured succinctly by inclusion of the priorities in (ii).

Story-specific priorities account for the proper qualification of information in the context of a specific story. In our example story, for instance, we are told (outside the excerpt) that the gun had always made a loud noise when it fired. This is captured by including the story-specific priority $p2 \succ c1$.

In accordance to psychological evidence, the use of unit-arguments as loose associations allows for the “*fast thinking*” and easy elaboration with new inferences critical in narrative comprehension. On the other hand, easy elaboration and the extreme form of qualification that it needs are mitigated by the requirement that elaborative inferences need to be *grounded* on the narrative and *sceptical* in nature.

Argumentation Semantics

Given a narrative $\mathcal{N}$ and a world knowledge $\mathcal{W}$, we formulate an argumentation framework $\langle \mathcal{A}, \mathcal{T}, \mathcal{D} \rangle$ by defining the set of arguments $\mathcal{A}$, the attacking relation $\mathcal{T} \subseteq \mathcal{A} \times \mathcal{A}$, and the defending relation $\mathcal{D} \subseteq \mathcal{A} \times \mathcal{A}$, as in the sequel.

Arguments in $\mathcal{A}$ are collections of tuples, each comprising a unit-argument, the time-point at which it is applied, and the inference that follows from its application. The introduction of time-points is necessitated as each given unit-argument can be possibly applied multiple times, with each application qualified by other conflicting knowledge in different ways.

The inference that follows is of the form (X, T) , where X is a fluent / action literal, and T the time-point at which it is inferred to hold. We allow a unit-argument $\arg(H, B)$ to be used in the usual forward direction to draw (H, T) as inference (with T the time-point at which the unit-argument’s head is applied), given that the premises B hold (at the appropriate time-point). In addition, we allow a unit-argument to be used in a backward direction (capturing, in this manner, reasoning by *reductio ad absurdum*) to draw $(\neg X, T)$ as inference (with T the time-point at which the unit-argument’s body is applied), given that $X \in B$ and that the premises $B \setminus \{X\}$ and $\neg H$ hold (at the appropriate time-points).

Despite the use of defeasible world knowledge, the framework includes *reasoning by contradiction*. The psychological debate on whether humans reason by contradiction, e.g., contraposition (Rips 1994; Johnson-Laird and Yang 2008) notwithstanding, it is still natural for a formal argumentation framework to capture such indirect reasoning; e.g., (Kakas and Mancarella 2013; Kakas, Toni, and Mancarella 2013). A main consequence is that it gives a form of *backward persistence*, e.g., from an observation to support (but not necessarily conclude, due to a possible qualification) that the observed property holds also at previous time-points. Note also that the separation of the inference type (e.g., forward and backward) is known to be significant in preference-based argumentation (Modgil and Prakken 2012), and necessitates a careful treatment of the definition of attacks and defences.

To reflect the suggestion by Psychology that inferences drawn by readers are strongly tied to the story, we require that the activation conditions of argument tuples must be eventually traced back to the explicit information in the narrative of the story representation. Thus, each argument is restricted so that the premises of each tuple it contains are supported either directly by the narrative itself, or by inferences

of other tuples in the argument in a non-cyclic manner.

The second key suggestion by Psychology that inferences are sceptical is reflected by allowing attacks to be generated “easily”, and insisting that defences are “harder” to generate. Therefore, in case of “equally strong” arguments (leading to a non-deterministic split), there is an attack, but no defense.

Roughly, an argument A_1 attacks an argument A_2 , and thus $(A_1, A_2) \in \mathcal{T}$, if they draw conflicting conclusions. The use of contrapositive reasoning for backward inference also means that it is possible to have arguments that support conclusions that are not contrary to each other, but whose unit-arguments have conflicting heads. For instance, in our running example we can use the causal unit-argument $c1$ in A_1 to forward derive $fired_at(pj, X)$ and the property unit-argument $p1$ in A_2 to backward derive gun_loaded from $fired_at(pj, X)$; despite that the derived conclusions of A_1 and A_2 are not in conflict, the unit-arguments used have conflicting heads. Although not all such *indirect* conflicts are important, a certain subset does need to be accounted for.

Assuming that $(A_1, A_2) \in \mathcal{T}$, an argument A_2 defends an argument A_1 , and thus $(A_2, A_1) \in \mathcal{D}$, if the former undercuts the latter, but the latter does not undercut the former. Undercutting means that a unit-argument in A_2 has a higher priority than a unit-argument in A_1 , capturing, in effect, a weakness in the chain of reasoning that A_1 used to attack A_2 ; cf. (Modgil and Prakken 2012). Again, care needs to be taken in dealing with indirect conflicts, as a result of the *de-feasible nature* of the world knowledge and the allowance of reasoning by contradiction on such defeasible information.

Finally, an argument Δ is acceptable if it does not have a conflict with the narrative or itself, and for every argument A that attacks Δ , Δ defends A . An acceptable argument corresponds to a comprehension model of the story $\mathcal{N}$ under the world knowledge $\mathcal{W}$. A comprehension model can be tested, as is often done in Psychology, through a series of multiple-choice questions with answers of the form “ X holds at T ”. Such an answer is accepted if (X, T) is entailed by the comprehension model, it is rejected if $(\neg X, T)$ is entailed by the comprehension model, and it is possible otherwise.

Evaluation through Empirical Studies

We have implemented the KRR framework as a Prolog system able to read a story narrative as a sequence of segments, compute a comprehension model, answer questions, and repeat the process, revising its conclusions as more parts of the story narrative become available. The process runs in time polynomial in the size of the world knowledge and the size of the story narrative, and is sound w.r.t. the framework.

In the first part of the evaluation of our approach we carried a *psychological study* to ascertain the world knowledge that is activated to successfully comprehend stories such as our example story on the basis of data obtained from human readers. We developed a set of inferential questions to follow the reading of pre-specified story segments. These assessed the extent to which readers connected, explained, and elaborated key story elements. Readers were instructed to answer each question and to justify their answers using a “think-aloud” method of answering questions while reading, in order to reveal the world knowledge that they had used.

The qualitative data from the readers was pooled together and analysed as to the frequencies of the types of responses in conjunction with the information given in justifications and think-aloud protocols. Considering those readers that demonstrated successful comprehension according to psychological criteria, our system was able to identify the most popular answers to questions, and also to recognize questions for which no single answer was (sceptically) accepted, following the variability demonstrated by human readers.

Related Work and Conclusions

Automated story understanding has been an ongoing endeavor of AI for more than forty years (Mueller 2002; 2013). Logic-related approaches have largely proceeded under the assumption that standard logical reasoning techniques can subsequently be applied; e.g., satisfiability (Mueller 2003) or planning (Niehaus and Young 2009). To our knowledge there has been very little work relating story comprehension with computational argumentation, an exception being (Bex and Verheij 2010), in which a case is made for combining narrative and argumentation techniques in the context of legal reasoning. Argumentation for reasoning about actions and change, on which our formal framework builds, has been studied in (Vo and Foo 2005; Michael and Kakas 2009).

To complete a fully automated approach to SC we continue drawing lessons from Psychology to address further the computational aspects of cognitive economy and coherence, by applying “computational heuristics” on top of the semantic framework we have developed. We are also investigating the systematic extraction or acquisition of commonsense world knowledge using lexical databases (Miller 1995; Baker, Fillmore, and Lowe 1998; Palmer, Gildea, and Kingsbury 2005), knowledge base archives (Lenat 1995), or automated extraction of commonsense knowledge from text (Michael and Valiant 2008; Michael 2009; 2011; 2013).

We envisage that the strong inter-disciplinary nature of our work can provide a concrete and important test-bed for evaluating the development of frameworks in AI while at the same time offering valuable feedback for Psychology.

References

- Baker, C. F.; Fillmore, C. J.; and Lowe, J. B. 1998. The Berkeley FrameNet Project. In *ACL'98*.
- Bex, F. J., and Verheij, B. 2010. Story Schemes for Argumentation about the Facts of a Crime. In *CMN'10*.
- Brewer, W., and Lichtenstein, E. 1982. Stories are to Entertain: A Structural-Affect Theory of Stories. *Journal of Pragmatics* 6:473–486.
- Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *AIJ* 77(2):321–358.
- Johnson-Laird, P. N., and Yang, Y. 2008. *Mental Logic, Mental Models, and Simulations of Human Deductive Reasoning*. The Cambridge Handbook of Computational Psychology. 339–358.
- Kakas, A., and Mancarella, P. 2013. On the Semantics of Abstract Argumentation. *Logic Computation* 23:991–1015.
- Kakas, A.; Toni, F.; and Mancarella, P. 2013. Argumentation for Propositional Logic and Nonmonotonic Reasoning. In *Commonsense'13*.
- Kintsch, W. 1988. The Role of Knowledge in Discourse Comprehension: A Construction-Integration Model. *Psychological Review* 95:163–182.
- Kintsch, W. 1998. *Comprehension: A Paradigm of Cognition*. NY: Cambridge University Press.
- Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *CACM* 38(11):32–38.
- McNamara, D. S., and Magliano, J. 2009. Toward a Comprehensive Model of Comprehension. *The Psychology of Learning and Motivation* 51:297–384.
- Mercier, H., and Sperber, D. 2011. Why Do Humans Reason? Arguments for an Argumentative Theory. *Behavioral and Brain Sciences* 34(2):57–74.
- Michael, L., and Kakas, A. C. 2009. Knowledge Qualification through Argumentation. In *LPNMR'09*.
- Michael, L., and Valiant, L. G. 2008. A First Experimental Demonstration of Massive Knowledge Infusion. In *KR'08*.
- Michael, L. 2009. Reading Between the Lines. In *IJCAI'09*.
- Michael, L. 2011. Causal Learnability. In *IJCAI'11*.
- Michael, L. 2013. Machines with Websense. In *Commonsense'13*.
- Miller, G. A. 1995. WordNet: A Lexical Database for English. *CACM* 38(11):39–41.
- Modgil, S., and Prakken, H. 2012. A General Account of Argumentation with Preferences. *AIJ* 195:361–397.
- Mueller, E. T. 2002. Story Understanding. In Nadel, L., ed., *Encyclopedia of Cognitive Science*, volume 4, 238–246. London: Macmillan Reference.
- Mueller, E. T. 2003. Story Understanding through Multi-Representation Model Construction. In *HLT-NAACL 2003 Workshop on Text Meaning*, 46–53.
- Mueller, E. T. 2013. Story Understanding Resources. <http://xenia.media.mit.edu/~mueller/storyund/stories.html>. Accessed February 28, 2013.
- Niehaus, J., and Young, R. M. 2009. A Computational Model of Inferencing in Narrative. In *INT'09*.
- Palmer, M.; Gildea, D.; and Kingsbury, P. 2005. The Proposition Bank: An Annotated Corpus of Semantic Roles. *Computational Linguistics* 31(1):71–106.
- Rapp, D., and den Broek, P. V. 2005. Dynamic Text Comprehension: An Integrative View of Reading. *Current Directions in Psychological Science* 14:297–384.
- Rips, L. 1994. *The Psychology of Proof*. MIT Press.
- van Harmelen, F.; Lifschitz, V.; and Porter, B. 2008. *Handbook of Knowledge Representation*. Elsevier Science.
- Vo, Q. B., and Foo, N. Y. 2005. Reasoning about Action: An Argumentation-Theoretic Approach. *JAIR* 24:465–518.
- Zwaan, R. A. 1994. Effect of Genre Expectations on Text Comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20:920–933.