

Book Reviews

Representation and Management of Narrative Information: Theoretical Principles and Implementation

Gian Piero Zarri

(Politecnico di Milano)

Springer Verlag (Advanced Information and Knowledge Processing series, edited by Lakhmi Jain and Xindong Wu), 2009, x+301 pp; hardbound, ISBN 978-1-84800-077-3, \$99.00; e-book, ISBN 978-1-84800-078-0; DOI 10.1007/978-1-84800-078-0

Reviewed by

Frank Schilder

Thomson Reuters

Gian Piero Zarri's book summarizes more than a decade of his research on knowledge representation for narrative text. The centerpiece of Zarri's work is the Narrative Knowledge Representation Language (NKRL), which he describes and compares to other competing theories. In addition, he discusses how to model the meaning of narrative text by giving many real-world examples. NKRL provides three different components or capabilities: (a) a representation system, (b) inferencing, and (c) an implementation. It is implemented via a Java-based system that shows how a representational theory can be applied to narrative texts.

The book consists of five chapters and two appendices. Chapter 1 introduces the basic principles of NKRL. The chapter first defines the focus on nonfiction narratives by contrasting the domain with fictional narratives, for example, a novel. Zarri chooses n -ary predicates in order to represent events formally. He argues for a neo-Davidsonian knowledge representation following Schank (1980), Schubert (1976), and others, and at the same time he sets his approach apart from the knowledge representation proposals one can find in Semantic Web representation languages such as RDF and OWL. However, Zarri emphasizes that NKRL, despite its similarity to conceptual graphs (Sowa 1999), is more focused on practical applications. The chapter concludes by introducing so-called templates in an attempt to demonstrate the practical usefulness of NKRL.

Chapter 2 provides an in-depth description of NKRL. Four connected components are introduced:

- The definitional component provides a hierarchy of abstract concepts (e.g., artifact, company, activity) called HClass (hierarchy of classes).
- The descriptive component is a hierarchy of event types called HTemp (hierarchy of templates) commonly found in the domain of non-fiction narratives (e.g., moving an object, producing a task or activity).
- The factual component describes the concrete instantiation of an event. For example, the sentence *Berlex Laboratories have performed an evaluation of a given compound* would be represented as

[PRODUCE: [SUBJ BERLEX_LABORATORIES, OBJ: ASSESSMENT_1,
TOPIC: COMPOUND_27]]

- The enumerative component links the values of the event description from the factual component to unique instantiations of these participants of an object, such as in `COMPOUND_27`.

This chapter also contains a comparison of NKRL and other formalisms that deal with the representation of temporal information, such as TimeML (Pustejovsky et al. 2003) or Discourse Representation Theory (DRT; Kamp and Reyle 1993). A detailed description is given of how NKRL represents temporal information based on Allen's (1984) interval calculus and how NKRL approaches the problem of underspecified or coarse temporal information such as in *around December 25, 2005*.

Chapter 3 adds more information about the semantics and the ontologies in NKRL. A set of predefined conceptual structures is introduced and numerous examples taken from real-world narratives are provided. The definitional component introduced in Chapter 2 is fleshed out and the distinction between sortal and non-sortal concepts (e.g., `CHAIR` versus `GOLD`) is described. This hierarchy is quite similar to other so-called upper-level ontologies such as `CYC` or `SUMO` (Guha and Lenat 1991; Pease, Niles, and Li 2002), which are introduced to the reader in more detail in the beginning of this chapter. In addition to the conceptual hierarchy `HClass`, the descriptive component `HTemp` holds a set of often-used templates. Each template is described with a specific example that shows how different slots of the templates may be filled.

Chapter 4 covers in more detail how inferences can be drawn within the implemented system. The NKRL system provides several query tools for retrieving information from the knowledge base encoded in NKRL annotations, as described by the previous chapters. The query tools comprise querying by search patterns, unification / filtering operations, and indexing temporal information. The indexing of temporal information considers different levels of temporal information including a temporal perspective. The temporal perspective is used to represent information about when an event starts or ends in addition to whether it is observed by somebody.

Chapter 5 provides the author's conclusions, suggesting technological and theoretical enhancement to the current version of NKRL. Appendix A contains a detailed description of the NKRL software and Appendix B discusses the treatment of a particular linguistic phenomenon within NKRL: plural entities.

The book offers a unique combination of different tools for modeling narrative information. It contains valuable discussions of important questions such as whether n -ary predicates should be used. However, some of these discussions would have benefited from a more in-depth treatment. The comparison with TimeML (Pustejovsky et al. 2003), for example, only partly covers recent developments and does not mention software that utilizes TimeML. For example, the `TARSQI` toolset,¹ not mentioned by Zarri, allows the user to extract events and temporal expressions while temporal links are derived and consistency checks can be run via a constraint propagation component.

Zarri often compares and contrasts his work with the representation languages used for the Semantic Web, such as `RDF` and `OWL`. He rightly points out similarities while addressing shortcomings of the Semantic Web technology (e.g., restriction to ternary predicates). But he overlooks an important point: `RDF` and `OWL` were not created for the semantic representation of non-fictional narratives—the focus of this book. Halevy, Norvig, and Pereira (2009), for instance, point out that one needs to distinguish between

¹ <http://www.timeml.org/site/tarsqi/toolkit/index.html>.

approaches to the semantic interpretation problem of natural language and the representation of the Semantic Web. The former tries to address the question of how language can be formally represented while the latter focuses on the interoperability of semantic information expressed by Web pages (e.g., flight information provided by a travel agency). Because Zarri proposes an approach for solving the semantic interpretation problem, a clear distinction as to what problem is being addressed would be helpful in order to avoid giving the wrong impression that the proposed methods are suitable for the Semantic Web.

On a more detailed level, Zarri's introduction to sortal and non-sortal concepts does not cite work by Krifka (1992) and Dowty (1991). These authors presented theories on sortal event hierarchies similar to the one discussed by Zarri, and both theories discuss how the sortal quality of an object influences the event type (e.g., eating an apple vs. eating apple purée). Zarri's introduction to his theoretical framework could have been improved by a more in-depth treatment of event semantics, such as the theories introduced by Krifka, Dowty, and others.

On a more minor note, the readability of the book is slightly hampered by the overuse of italics and other fonts. Less would have been definitely more here. In addition, a list of abbreviations and a glossary of important terms would have been useful so that the reader could use the book as a quick reference, for example.

Though the working implementation of the system is intriguing, and the book provides many narrative texts exemplifying the expressiveness and capability of NKRL, an unaddressed issue is whether the implementation can be scaled. Unfortunately, there is no way for the reader to decide, because the book does not include a CD with a demo of the system, let alone the source code.

This book would be useful to researchers and practitioners in the field of modeling narrative information, but not to beginners in this field. Students, for example, who are interested in this area would need a more guided approach to the topic. Nevertheless, the book provides a solid theoretical foundation for representing information extracted from narrative text such as news messages; and I am pleased to see that the author undertook the effort of implementing his theory in an actual system that has the potential for many different and exciting practical applications.

References

- Allen, James. 1984. Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123–154.
- Dowty, David. 1991. Thematic proto-roles and argument selection. *Language*, 67:547–619.
- Guha, Ramanathan V. and Douglas B. Lenat. 1991. Cyc: A mid-term report. *Applied Artificial Intelligence*, 5(1):45–86.
- Halevy, Alon, Peter Norvig, and Fernando Pereira. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 24(2):8–12.
- Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer Academic.
- Krifka, Manfred. 1992. Thematic relations as links between nominal reference and temporal constitution. In Ivan Sag and Anna Szabolcsi, editors, *Lexical Matters*. CSLI Publications, Stanford, CA, pages 29–53.
- Pease, Adam, Ian Niles, and John Li. 2002. The suggested upper merged ontology: A large ontology for the Semantic Web and its applications. In *Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web*, Edmonton.
- Pustejovsky, James, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, and Graham Katz. 2003. TimeML: Robust specification of event and temporal expressions in text. In *Proceedings of the Fifth International Workshop on*

Computational Semantics (IWCS-5),
Tilburg.
Schank, Roger C. 1980. Language and
memory. *Cognitive Science*, 4(3):243–284.
Schubert, Lenhart K. 1976. Extending
the expressive power of semantic

networks. *Artificial Intelligence*,
7(2):163–198.
Sowa, John F. 1999. *Knowledge Representation:
Logical, Philosophical, and Computational
Foundations*. Brooks Cole Publishing Co.,
Pacific Grove, CA.

Frank Schilder is a senior research scientist at the Thomson Reuters Research & Development department. His research interests include temporal information extraction, summarization, and discourse processing. His address is Thomson Reuters R&D, 610 Opperman Drive, St. Paul, MN 55123, USA; e-mail: frank.schilder@thomsonreuters.com.

Dependency Parsing

Sandra Kübler, Ryan McDonald, and Joakim Nivre
(Indiana University, Google Research, and Uppsala and Växjö Universities)

Morgan & Claypool (Synthesis Lectures on Human Language Technologies, edited by Graeme Hirst, volume 2), 2009, xii+115 pp; paperbound, ISBN 978-1-59829-596-2, \$40.00; e-book, ISBN 978-1-59829-597-9, \$30.00 or by subscription

Reviewed by
John Carroll
University of Sussex

This book is a survey of the latest research in dependency parsing, describing some of the approaches currently being investigated by the research community, assessing their strengths and weaknesses, and exploring relationships between them. The book is one of the first in a series of short monographs entitled *Synthesis Lectures on Human Language Technologies*; other related series include speech and audio processing, and artificial intelligence and machine learning. A PDF version of the book can be purchased on-line for \$30.00.

Dependency parsing is currently a very active area of research, so this book is timely. Until recently, most research into inducing statistical parsers from syntactically annotated text has worked with phrase structure tree representations. However, over the last ten years or so a number of researchers have argued (Lin 1995; Briscoe and Carroll 2006) that dependency analyses—in which syntactic structure is represented by linking pairs of words by labeled dependency relations—have a number of important advantages over phrase structure trees. In particular:

- For evaluation of parsing accuracy (with respect to a gold standard), measuring correctness of dependency relations gives more reliable results than measuring similarity of phrase structure.
- For applications requiring automatic linguistic analysis, dependencies capture some aspects of predicate–argument structure in a more convenient form than phrase structure.

In addition, as annotated corpora for languages other than English have become more numerous, it has become evident that:

- For free word order languages, dependencies are better able to capture syntactic generalizations than phrase structure trees.

The CoNLL shared tasks from 2006 onwards have played an important role in raising interest in dependency parsing and supporting work in this area, by organizing the annotation and distribution of dependency corpora in several languages, providing the stimulus of friendly competition, and facilitating comparisons between the various techniques used by the systems that have been entered.

The book describes a range of dependency parsing models, presenting them in terms of a common framework consisting of three elements: a set of constraints defining the space of permissible dependency structures for a given sentence, a set of parameters (possibly empty), and a parsing algorithm. Two main classes of models are distinguished: data-driven, in which the parser is learned from a corpus of dependency structures; and grammar-based, in which the parser is directed by a formal grammar (and also possibly by weights controlling how alternative solutions are explored).

At only a little over 100 pages, the book does not have room to address issues such as how dependency grammar can be used to describe the linguistic facts about a given language, how particular syntactic constructions could be analyzed, or how hand-crafted dependency grammars are developed. The focus is on parsing algorithms, and the two main data-driven approaches of transition-based and graph-based parsing are described in some detail. The focus on the process of parsing itself also means that the book does not refer (even in a further-reading section) to some widely distributed parsing systems which use other types of grammar internally but which output dependencies, for example the C&C parser (Clark and Curran 2007). And although the CoNLL data sets have been very influential, there is no mention of the PARC 700 Dependency Bank (King et al. 2003), which has been used in comparative dependency-based evaluations of a wide range of parsers.

The authors state that they are aiming the book at graduate students and researchers in computer science, linguistics, and computational linguistics. The reader is assumed to have some background knowledge of linguistics and computer science, and although not necessary for an understanding of most of the material, at some points the reader is expected to be comfortable with proof by induction, analysis of the complexity of algorithms, and algorithms for directed graphs. Two of the chapters begin with a few pages of formal definitions, propositions, and proofs, with little in the way of motivating or illustrative examples, and here the less mathematically confident reader will have to persevere.

The publisher has produced the book very quickly: Indeed, the date of publication is only one month after the date at the end of the Preface. However, a side effect of this is a lack of polish: There are a number of typographical errors (and at least one incomplete sentence), occasional lapses into unidiomatic English, some instances of incorrect referencing style, and there is no index. Also, decisions on when references are given within the body of a chapter or postponed to end-of-chapter notes are inconsistent, some technical terms (e.g., ‘spurious ambiguity’) are not explained at the first point at which they are mentioned, and others (e.g., ‘gap degree’) are not explained at all.

However, given that these problems are mostly textual and in general do not impede the reader’s understanding, this book serves as a very useful and up to date survey of the burgeoning research area of dependency parsing.

References

- Briscoe, Ted and John Carroll. 2006. Evaluating the accuracy of an unlexicalized statistical parser on the PARC DepBank. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, Sydney, pages 41–48.
- Clark, Stephen and James R. Curran. 2007. Wide-coverage efficient statistical parsing with CCG and log-linear models. *Computational Linguistics*, 33(4):493–552.
- King, Tracy H., Richard Crouch, Stefan Riezler, Mary Dalrymple, and Ronald M. Kaplan. 2003. The PARC 700 Dependency Bank. In *Proceedings of the 4th International Workshop on Linguistically Interpreted Corpora*, Budapest, pages 1–8.
- Lin, Dekang. 1995. A dependency-based method for evaluating broad-coverage parsers. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Montreal, pages 1420–1425.

This book review was edited by Pierre Isabelle.

John Carroll is Professor of Computational Linguistics at the University of Sussex, working on parsing, tactical generation, opinion mining, and real-world applications of language technology. Carroll's address is School of Informatics, University of Sussex, Falmer, Brighton BN1 9QH, UK; e-mail: J.A.Carroll@sussex.ac.uk.