

Book Reviews

Semantic Role Labeling

Martha Palmer,^{*} Daniel Gildea,[‡] and Nianwen Xue[¶]

(*University of Colorado, Boulder; ‡University of Rochester; ¶Brandeis University)

Morgan & Claypool (Synthesis Lectures on Human Language Technologies, edited by Graeme Hirst, volume 6), 2010, ix+91 pp; paperbound, ISBN 978-1-59829-831-4, \$30.00 or by subscription

Reviewed by

Alessandro Moschitti

University of Trento

A basic aim of computational linguistics (CL) is the study and design of computational models of natural language semantics. Although frequency-based approaches—for example, distributional semantics—provide effective and concrete solutions for natural language applications, they still fail to fully reconcile the field with the theoretical-linguistic soul. In contrast, semantic role labeling (SRL), a recent new area of CL, aims to automatically provide (shallow) semantic layers using modern linguistic theories of semantic roles, also exploitable by language applications. The centrality and importance of such theories in CL has promoted the development of a rather large body of work on SRL; its many aspects and research directions make it difficult to survey the field.

Palmer, Gildea, and Xue's book provides an excellent description of such work, detailing all its main concepts and practical aspects. The authors accurately illustrate all important ingredients to acquire a global and precise view of the field, namely, (i) the theoretical framework, ranging from linking theory to theta roles, Levin's classes and frame semantics; (ii) computational models based on syntactic representations derived from diverse parsing paradigms; (iii) several resources in different languages; (iv) many machine learning approaches and strategies; and (v) portability to other languages and domains.

This book is mainly directed to practitioners who want to contribute to SRL or who want to simply use its technology in natural language applications. As an "Ariadne's ball of thread," this book will guide the reader through the conceptual SRL labyrinth, saving months of work needed to understand theory and practice of this exciting research field. The book is divided into four content chapters.

Chapter 1. A sound and natural way to approach SRL, independently of one's own background and interest, is to start from an overall understanding of the theory of semantic roles. This phase is not painless for the practitioner as there is no general agreement on the theory of semantic roles. Several concepts must be acquired, including theta roles, proto-roles, Levin's classes, and frame semantics. As a consequence, to understand the potential of the different formalisms, entire books (e.g., Jackendoff 1990; Levin 1993) should be examined. This step, although interesting, probably would slow down the work of the practitioners and may frustrate those not having the required linguistic background. By contrast, this chapter is an optimal compromise as it simply summarizes all the important aspects of semantic role theory, also providing concrete examples. This allows for saving much of the time needed to recombine and interpret the different results obtained from different theories.

Chapter 2. One interesting aspect of SRL is the practical instantiation of linguistic theories in some corpora. A nice feature of this chapter is the ability to quickly focus on practical aspects by presenting the most useful data sets. The authors' technical description is minimal but precise and highlights the differences between resources and their annotation from theoretical and practical viewpoints: The chapter provides what is needed to understand the format and meaning of the data so that writing code for using it becomes straightforward. Advanced topics such as linking different resources to boost the accuracy of automatic SRL systems are also presented. One drawback of the chapter, which is a consequence of the limited available space, is the lack of a description of other important resources (e.g., NomBank [Meyers et al. 2004], and resources in languages other than English.

Chapter 3. For a computational linguist, this is the most fascinating chapter. It presents the typical computational models used to design automatic SRL systems by illustrating the most effective pipeline architectures. These are typically composed of different modules performing different tasks; the filtering, identification, classification, and joint inference stages are described in detail by proposing features and models that have proved to be effective during several years of research. Additionally, the chapter discusses important aspects of SRL systems such as (i) the impact of syntactic information along with different parsing paradigms, (ii) models combining different types of syntactic representation, and (iii) models integrating syntactic and semantic parsers in a single approach. These descriptions constitute key information, which allows for enhancing accurate basic systems to state-of-the-art shallow semantic parsers. The chapter concludes with valuable information for testing the quality of an SRL system, that is, a description of the most-used accuracy measures for different kinds of parsing paradigms. Typical aspects that impact system performance such as domain variability, combinations of different resources, and the use of unsupervised approaches are also illustrated. However, this chapter may still be considered incomplete as some architectures exploiting advanced machine learning techniques, for example, kernel methods (Moschitti 2004), are not reported.

Chapter 4. After all the important topics of English SRL have been presented, this chapter is dedicated to discussing the problem of extending semantic parser models to other languages (e.g., Chinese). After a brief presentation of resources for other languages, which serves only the purpose of showing their availability, the chapter focuses on interesting topics such as semantic role projection and alignment. The former allows for automatically generating labeled data using annotation on one language and parallel corpora, whereas the latter aligns both annotated corpora and can be exploited for machine translation. The second main topic of the chapter is an interesting discussion of the adaptation of system architectures when they are applied to languages different than English. Although the focus is only on the English–Chinese pair, the derived guidelines are generally applicable to other languages. Finally, the presence of nominal predicates in the Chinese PropBank is used for introducing the important topic of nominal SRL. Although this explanation along with other related resources (e.g., NomBank) could have received a larger space, it is enough to complete all the most relevant topics of the field.

What is missing? This is a sound and complete book on SRL. It can also work as an initial manual of SRL systems as it provides indications for designing state-of-the-art parsers.

However, its small size suggests that it could not have been comprehensive, given the very large body of work in shallow semantic parsing. Consequently, there are some missing or not completely described topics. One of them is the practical use of semantic parses for concrete tasks or real-world applications. That is, the question “Now that we have such a nice shallow semantic representation, how do we use it for a concrete (e.g., commercial) task?” remains unanswered. It is not easy to respond to this question as, at the moment, no industrial company is using SRL (or getting from it a resounding success). However, a discussion of previous work that has successfully exploited SRL—for example, for question answering (Moschitti et al. 2007; Shen and Lapata 2007; Surdeanu, Ciaramita, and Zaragoza 2008), for sentiment analysis (Johansson and Moschitti 2010), and for cross-document coreference resolution (Ponzetto and Strube 2006)—could have been attempted.

Another potentially interesting chapter would have been a survey of machine learning approaches. Although the book wisely presents a well-assessed and restricted set of techniques, the prolific SRL research has developed many other interesting methods, for example, in CoNLL (Carreras and Màrquez 2005; Surdeanu et al. 2008). Additionally, more evidence on the accuracy and speed achievable by the different SRL models on different corpora and tasks would have been useful for practitioners to estimate the expected performance of new systems in new application domains.

Finally, a description of the available resources, classified with respect to the semantic role theory, the language, and the genre would have been very nice. It would have given a clear picture of the spread of SRL in the natural language processing or related fields, for example, semantic Web or data mining.

Despite these points, this is a unique book on SRL that condenses essential knowledge and years of research in linguistics, computational linguistics, and machine learning for SRL into a small number of pages. It allows the reader to save time in getting acquainted with SRL research. It is also useful in helping the design of one’s own system and can serve as a starting point for conducting advanced research in SRL. In conclusion, this book is indispensable for researchers who are approaching SRL.

References

- Carreras, Xavier and Lluís Màrquez. 2005. Introduction to the CoNLL-2005 shared task: Semantic role labeling. In *Proceedings of the Ninth Conference on Computational Natural Language Learning (CoNLL 2005)*, pages 152–164, Ann Arbor, MI.
- Jackendoff, Ray S. 1990. *Semantic Structures*. The MIT Press, Cambridge, MA.
- Johansson, Richard and Alessandro Moschitti. 2010. Syntactic and semantic structure for opinion expression detection. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning (CoNLL 2010)*, pages 67–76, Uppsala.
- Levin, Beth. 1993. *English Verb Classes and Alternations*. University of Chicago Press, Chicago, IL.
- Meyers, A., R. Reeves, C. Macleod, R. Szekeley, V. Zielinska, B. Young, and R. Grishman. 2004. The NomBank project: An interim report. In *Proceedings of the HLT-NAACL 2004 Workshop: Frontiers in Corpus Annotation*, pages 24–31, Boston, MA.
- Moschitti, Alessandro. 2004. A study on convolution kernels for shallow semantic parsing. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL 2004)*, pages 335–342, Barcelona.
- Moschitti, Alessandro, Silvia Quarteroni, Roberto Basili, and Suresh Manandhar. 2007. Exploiting syntactic and shallow semantic kernels for question/answer classification. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics (ACL 2007)*, pages 776–783, Prague.
- Ponzetto, Simone Paolo and Michael Strube. 2006. Exploiting semantic role labeling, WordNet and Wikipedia for coreference resolution. In *Proceedings of the 2006 Human Language Technology Conference*

of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 2006), pages 192–199, Brooklyn, NY.

- Shen, Dan and Mirella Lapata. 2007. Using semantic roles to improve question answering. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2007)*, pages 12–21, Prague.
- Surdeanu, Mihai, Massimiliano Ciaramita, and Hugo Zaragoza. 2008. Learning to

rank answers on large online QA collections. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL 2008)*, pages 719–727, Columbus, OH.

- Surdeanu, Mihai, Richard Johansson, Adam Meyers, Lluís Màrquez, and Joakim Nivre. 2008. The CoNLL 2008 shared task on joint parsing of syntactic and semantic dependencies. In *Proceedings of the 12th Conference on Computational Natural Language Learning (CoNLL 2008)*, pages 159–177, Manchester.

This book review was edited by Pierre Isabelle.

Alessandro Moschitti is a professor of the Computer Science and Information Engineering Department of Trento University, Italy. He has designed many diverse SRL systems: three PropBank systems for English, one system for Arabic PropBank, two FrameNet parsers for English (one of which was the most accurate system in the SRL challenge of SENSEVAL 2004), and one parser for Italian FrameNet. Currently, his main research direction in SRL is the design of models for exploiting shallow semantics in practical applications, for example, question answering and opinion mining. Moschitti's address is DISI, University of Trento, Via Sommarive 14, 38123 POVO (TN) - Italy; e-mail: moschitti@disi.unitn.it.