Semantic Parsing with Combinatory Categorial Grammars

Yoav Artzi, **Nicholas FitzGerald** and **Luke Zettlemoyer**
Computer Science & Engineering
University of Washington
Seattle, WA 98195
{yoav,nfitz,lsz}@cs.washington.edu

1 Abstract

Semantic parsers map natural language sentences to formal representations of their underlying meaning. Building accurate semantic parsers without prohibitive engineering costs is a long-standing, open research problem.

The tutorial will describe general principles for building semantic parsers. The presentation will be divided into two main parts: learning and modeling. In the learning part, we will describe a unified approach for learning Combinatory Categorial Grammar (CCG) semantic parsers, that induces both a CCG lexicon and the parameters of a parsing model. The approach learns from data with labeled meaning representations, as well as from more easily gathered weak supervision. It also enables *grounded* learning where the semantic parser is used in an interactive environment, for example to read and execute instructions. The modeling section will include best practices for grammar design and choice of semantic representation. We will motivate our use of lambda calculus as a language for building and representing meaning with examples from several domains.

The ideas we will discuss are widely applicable. The semantic modeling approach, while implemented in lambda calculus, could be applied to many other formal languages. Similarly, the algorithms for inducing CCG focus on tasks that are formalism independent, learning the meaning of words and estimating parsing parameters. No prior knowledge of CCG is required. The tutorial will be backed by implementation and experiments in the University of Washington Semantic Parsing Framework (UW SPF).[1]

[1] http://yoavartzi.com/spf

2 Outline

1. Introduction to CCG and lambda calculus
2. Learning
   (a) A unified learning algorithm
   (b) Learning with supervised data
      i. Lexical induction with templates
      ii. Unification-based learning
   (c) Weakly supervised learning without labeled meaning representations
3. Modeling
   (a) Questions for database queries
   (b) Plurality and determiner resolution in grounded applications
   (c) Event semantics and imperatives in instructional language

3 Instructors

Yoav Artzi is a Ph.D. candidate in the Computer Science & Engineering department at the University of Washington. His research studies semantic meaning and machine learning methods for natural language understanding within situated interactive systems. He is currently supported by the 2014 Microsoft Research PhD Fellowship, and is a recipient of the 2012 Yahoo KSC award.

Nicholas FitzGerald is a Ph.D. student at the University of Washington. His research interests are grounded natural language understanding and generation. He is a recipient of an Intel Science and Technology Center Fellowship and an NSERC Postgraduate Scholarship.

Luke Zettlemoyer is an Assistant Professor in the Computer Science & Engineering department at the University of Washington. His research interests are in the intersections of natural language processing, machine learning and decision making under uncertainty. Honors include best paper awards at UAI 2005 and ACL 2009, selection to the DARPA CSSG, and an NSF CAREER Award.

References

Y. Artzi and L.S. Zettlemoyer. 2011. Bootstrapping semantic parsers from conversations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

Y. Artzi and L.S. Zettlemoyer. 2013a. UW SPF: The University of Washington Semantic Parsing Framework.

Y. Artzi and L.S. Zettlemoyer. 2013b. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics*, 1(1):49–62.

SRK Branavan, H. Chen, L.S. Zettlemoyer, and R. Barzilay. 2009. Reinforcement learning for mapping instructions to actions. In *Proceedings of the Joint Conference of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*.

SRK Branavan, L.S. Zettlemoyer, and R. Barzilay. 2010. Reading between the lines: learning to map high-level instructions to commands. In *Proceedings of the Conference of the Association for Computational Linguistics*.

Q. Cai and A. Yates. 2013a. Large-scale semantic parsing via schema matching and lexicon extension. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

Q. Cai and A. Yates. 2013b. Semantic parsing freebase: Towards open-domain semantic parsing. In *Joint Conference on Lexical and Computational Semantics: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*.

B. Carpenter. 1997. *Type-Logical Semantics*. The MIT Press.

D.L. Chen and R.J. Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *Proceedings of the National Conference on Artificial Intelligence*.

D.L. Chen. 2012. Fast online lexicon learning for grounded language acquisition. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.

A. Church. 1932. A set of postulates for the foundation of logic. *The Annals of Mathematics*, 33:346–366.

A. Church. 1940. A formulation of the simple theory of types. *The journal of symbolic logic*, 5:56–68.

S. Clark and J. R. Curran. 2007a. Perceptron training for a wide-coverage lexicalized-grammar parser. In *Proceedings of the Workshop on Deep Linguistic Processing*.

S. Clark and J. R. Curran. 2007b. Wide-coverage efficient statistical parsing with CCG and log-linear models. *Computational Linguistics*, 33(4):493–552.

J. Clarke, D. Goldwasser, M. Chang, and D. Roth. 2010. Driving semantic parsing from the world's response. In *Proceedings of the Conference on Computational Natural Language Learning*.

M. Collins. 2002. Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

D. A Dahl, M. Bates, M. Brown, W. Fisher, K. Hunicke-Smith, D. Pallett, C. Pao, A. Rudnicky, and E. Shriberg. 1994. Expanding the scope of the ATIS task: The ATIS-3 corpus. In *Proceedings of the workshop on Human Language Technology*.

D. Davidson. 1967. The logical form of action sentences. *Essays on actions and events*, pages 105–148.

D. Davidson. 1969. The individuation of events. In *Essays in honor of Carl G. Hempel*, pages 216–234. Springer.

N. FitzGerald, A. Artzi, and L.S. Zettlemoyer. 2013. Learning Distributions over Logical Forms for Referring Expression Generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

M. Granroth-Wilding and M. Steedman. 2012. *Statistical parsing for harmonic analysis of jazz chord sequences*. Ann Arbor, MI: MPublishing, University of Michigan Library.

A. K. Joshi, K. V. Shanker, and D. Weir. 1990. The convergence of mildly context-sensitive grammar formalisms. Technical report, Department of Computer & Information Science, University of Pennsylvania.

J. Kim and R.J. Mooney. 2012. Unsupervised pcfg induction for grounded language learning with highly ambiguous supervision. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

J. Krishnamurthy and T. Kollar. 2013. Jointly learning to parse and perceive: Connecting natural language to the physical world. *Transactions of the Association for Computational Linguistics*, 1(1):193–206.

N. Kushman and R. Barzilay. 2013. Using semantic unification to generate regular expressions from natural language. In *Proceedings of the Human Language Technology Conference of the North American Association for Computational Linguistics*.

T. Kwiatkowski, L.S. Zettlemoyer, S. Goldwater, and M. Steedman. 2010. Inducing probabilistic CCG grammars from logical form with higher-order unification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

T. Kwiatkowski, L.S. Zettlemoyer, S. Goldwater, and M. Steedman. 2011. Lexical Generalization in CCG Grammar Induction for Semantic Parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

T. Kwiatkowski, E. Choi, Y. Artzi, and L.S. Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.

T. Lei, F. Long, R. Barzilay, and M. Rinard. 2013. From natural language specifications to program input parsers. In *Proceedings of the the annual meeting of the Association for Computational Linguistics*.

P. Liang, A. Bouchard-Côté, D. Klein, and B. Taskar. 2006. An end-to-end discriminative approach to machine translation. In *Proceedings of the Conference of the Association of Computational Linguistics*.

P. Liang, M.I. Jordan, and D. Klein. 2011. Learning dependency-based compositional semantics. In *Proceedings of the Conference of the Association for Computational Linguistics*.

M. MacMahon, B. Stankiewics, and B. Kuipers. 2006. Walk the talk: Connecting language, knowledge, action in route instructions. In *Proceedings of the National Conference on Artificial Intelligence*.

M. MacMahon. 2007. *Following Natural Language Route Instructions*. Ph.D. thesis, University of Texas at Austin.

C. Maienborn, K. Von Heusinger, and P. Portner. 2011. *Semantics: An international handbook of natural language and meaning*. Walter de Gruyter.

C. Matuszek, D. Fox, and K. Koscher. 2010. Following directions using statistical machine translation. In *Proceedings of the international conference on Human-robot interaction*.

C. Matuszek, N. FitzGerald, L.S. Zettlemoyer, L. Bo, and D. Fox. 2012. A joint model of language and perception for grounded attribute learning. *Proceedings of the International Conference on Machine Learning*.

S. Miller, R. Bobrow, R. Ingria, and R. Schwartz. 1994. Hidden understanding models of natural language. In *Proceedings of the Conference of the Association of Computational Linguistics*.

S. Miller, D. Stallard, R. Bobrow, and R. Schwartz. 1996. A fully statistical approach to natural language interfaces. In *Proceedings of the Conference of the Association of Computational Linguistics*.

T. Parsons. 1990. *Events in the Semantics of English*. The MIT Press.

G. Scontras, P. Graff, and N. D. Goodman. 2012. Comparing pluralities. *Cognition*, 123(1):190–197.

N. Singh-Miller and M. Collins. 2007. Trigger-based language modeling using a loss-sensitive perceptron algorithm. In *IEEE International Conference on Acoustics, Speech and Signal Processing*.

M. Steedman and J. Baldridge. 2003. Combinatory categorial grammar. *Non-Transformational Syntax*, pages 181–224.

M. Steedman. 1996. *Surface Structure and Interpretation*. The MIT Press.

M. Steedman. 2000. *The Syntactic Process*. The MIT Press.

M. Steedman. 2011. *Taking Scope*. The MIT Press.

L.R. Tang and R.J. Mooney. 2001. Using multiple clause constructors in inductive logic programming for semantic parsing. In *Proceedings of the European Conference on Machine Learning*.

S. Tellex, T. Kollar, S. Dickerson, M. Walter, A.G. Banerjee, S. Teller, and N. Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the National Conference on Artificial Intelligence*.

S. Tellex, P. Thaker, J. Joseph, and N. Roy. 2013. Learning perceptually grounded word meanings from unaligned parallel data. *Machine Learning*, pages 1–17.

J.M. Zelle and R.J. Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Proceedings of the National Conference on Artificial Intelligence*.

L.S. Zettlemoyer and M. Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.

L.S. Zettlemoyer and M. Collins. 2007. Online learning of relaxed CCG grammars for parsing to logical form. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*.

L.S. Zettlemoyer and M. Collins. 2009. Learning context-dependent mappings from sentences to logical form. In *Proceedings of the Joint Conference of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing*.