Inference and Natural Language Semantics

Christof Monz[1,2] Maarten de Rijke[2]

[1] IMS, Universität Stuttgart
Azenbergstr. 12, 70174 Stuttgart, Germany
christof@ims.uni-stuttgart.de
[2] ILLC, University of Amsterdam
Pl. Muidergracht 24, 1018 TV Amsterdam, The Netherlands
{christof, mdr}@wins.uva.nl

1 Introduction

One of the concluding slogans of the FraCaS project on Frameworks for Computational Semantics is that '[t]here can be no semantics without logic' [CFvGG98]. We take this to mean that formalisms for semantic representation should be developed hand-in-hand with inference methods for performing reasoning tasks with representations and algorithms for representation construction.

Clearly, to be usable in the first place, representation formalisms need to come equipped with construction methods, and this explains the need for algorithmic tools. But what about the need for inference methods? At least types of three reasons can be identified. For *cognitive* purposes one may want to test the truth conditions of a representation against (a model of) speakers' intuitions—this amounts to a model checking or theorem proving task. Also, the whole issue of what it is to understand a discourse may be phrased as a model generation task. *Computationally*, we need various reasoning tasks and AI-heuristics to help resolve quantifier scope ambiguity, or to resolve anaphoric relations in information extraction and natural language queries. And last, but not least, the very construction of semantic representations may require inference tools to be used in checking for consistency and informativity. At the end of the day, the main purpose of a semantic representation is that we can *do* something with it, both algorithmically and in terms of inference tasks.

The present times are exciting ones for anyone with an interest in inference for natural language semantics. On the one hand, there is work in semantics that has little or no attention for inferential aspects. This is certainly the case for a lot of work in dynamic semantics and underspecified representation, and in the recent *Handbook of Logic and Language* [vBtM97] inferential methods for semantic representations are largely absent, despite the fact that a substantial part of the book is devoted to representational matters.

At the same time, there is a growing body of work aimed at developing inference methods and tools for natural language semantics, fed by a growing realization that these are 'the heart of the enterprise' [BB98, page viii]. This is manifested not only by various research initiatives (see below), but also by the fact that a number of textbooks and monographs on natural language semantics and its inferential and algorithmic aspects are in preparation [BB98, CFvGG98], and by a recent initiative to set up a special interest group on Computational Semantics (see http://www.coli.uni-sb.de/~patrick/SIGICS.html for details).

In this note we survey some of the ongoing work on inference and natural language semantics; we identify commonalities, as well as possibilities and the main logical challenges we are confronted with in the field.

2 Putting Semantics to Work

2.1 Lines of Attack

It has often been claimed that classical reasoning based on first-order logic (FOL) is not appropriate as an inference method for natural language semantics. We are pragmatic in this matter: try to stick to existing formats and tools and see how far they get you, and only if they fail, one should develop novel formats and tools. Traditional inference tools (such as theorem provers and model builders) are reaching new levels of sophistication, and they are now widely and easily available. Blackburn and Bos [BB98] show that the 'conservative' strategy of using first-order tools can actually achieve a lot. In particular, they use first-order theorem proving techniques for implementing Van der Sandt's approach to presupposition.

Although one may want to stick to first-order based tools as much as possible, for reasons of efficiency, or simply to get 'natural representations' it may pay to move away from the traditional first-order realm. Such a move may be particularly appropriate in two of the areas that currently pose the biggest challenges for computational semantics: *ambiguity* and *dynamics* [CFvGG98, Chapter 8]. Let us consider some samples of deductive approaches in each of these two areas.

Reasoning with Ambiguity

While the problem of ambiguity and underspecification has recently enjoyed a considerable increase in attention from computational linguists, computer scientists, the focus has mostly been on semantic aspects, and 'reasoning with ambiguous sentences is still in its infancy' [vDP96]. Lexical ambiguities can be represented pretty straightforwardly by putting the different readings into a disjunction, doing so for quantificational ambiguities ignores additional structure that may help in the disambiguation process. To exploit this structural information and to avoid the state explosion problem, one needs a concise representation of the possible readings, and the proofs themselves need to be incremental in the sense that at any stage we have a 'partial' proof that can easily be extended to cope with novel information. [MdR98c] achieve this by adopting Bos' hole semantics in the style of [Rey93, Rey96], and by developing a tableaux calculus for quantificational ambiguity that interleaves disambiguation and proper deductive reasoning while preserving non-redundant representations as much as possible.

But there is more to reasoning with quantificational ambiguity than just developing a calculus for it. In the presence of multiple readings of premises and conclusions, fundamental logical notions such as *entailment* receive new dimensions. Should *all* possible readings of the conclusion follow (in the traditional sense) from all possible readings of the premises for the ambiguous conclusion to really qualify as a consequence of an ambiguous premise? Basic research in this direction has been carried out by a number of people [vD96, Rey95, vEJ96, Jas97]. Ultimately, the aim here is to obtain insights into the development and implementation of theorem provers for underspecified representations.

Deduction for Dynamics

A number of calculi have recently been proposed for reasoning with dynamic semantics. [KR96] and [RG94] present natural deduction calculi for Discourse Representation Theory. And in the area of Dynamic Predicate Logic (and its many variations), [Vel97] presents some ground-tableau calculi and [vE98a] a sequent calculus, but none of these calculi are designed to be efficiently implemented. [MdR98a, MdR98b] focus on efficiency, but the resolution method that is used there calls for a clausal form transformation which destroys the structural information of the input. As a consequence, the system is unable to reflect the incremental feature of the semantics. Furthermore, none of these calculi are adapted to incremental consistency checking as the syntactic information of 'actively quantified variables' is lost during a proof. Van Eijck's calculus for 'Dynamic Predicate Logic without variables' is incremental for the price of losing the destructive assignments brought about by repeated quantifications of the same variable in dynamic predicate logic.

However, [Ker98] presents an incremental free-variable tableau calculus for Dynamic Predicate Logic where unification takes care of the choice in substitutions and information on 'actively quantified variables' are recorded alongside a proof, thus enabling an efficient implementation of substitutions and incremental consistency checking.

2.2 Lessons Learned

The brief sketches of recent work on inference and natural language semantics given above show a number of things. First, all traditional computational reasoning tasks (theorem proving, model checking, model generation) are needed, but often in novel settings and as part of larger logical architectures. Second, there are novel logical concerns both at a fundamental and at an architectural level. The former is illustrated by the proliferation of notions of entailment and by the need for incremental, structure preserving proof procedures. As to the latter, to move forward we need to develop methods for integrating specialized inference engines, possibly operating on different kinds of information, with other computational tools such as statistical packages, parsers, and various interfaces. We propose to use combinations of small specialized modules rather than large baroque systems. Of course, similar strategies in design and architecture have gained considerable attention in both computer science [Boo91], and in other areas of applied logic and automated reasoning [BS96].

3 Further Directions and Challenges

The findings of the previous section are supported by a number of further and novel developments in more applied areas adjacent to natural language semantics. We will restrict ourselves to three examples.

First, in syntactic analysis, partial or underspecified approaches to parsing are becoming increasingly popular [Abn96]. Just like underspecified representations in semantics, a partial parse fully processes certain phrases, but leaves some ambiguities such as modifier attachment underspecified. Given this similarity, it is natural to ask whether underspecified semantics can somehow be combined with partial parsing. An ongoing project at ILLC studies to which extent one can, for instance, use semantic information into account to resolve syntactic ambiguities. Note that combinations of underspecified representation and packed syntactic trees (parse forests) have been considered before [Sch96, Dör97], but no methods for using semantic information to resolve syntactic ambiguities are reported there.

Second, assuming that underspecified representations can usefully be combined with partial parsing, we may be able to improve methods in Information Extraction (IE). Common approaches to IE suffer from the fact that they either give only a very shallow analysis of text documents, as in approaches using word vectors, or that they are domain dependent, as in the case of template filling. More general techniques using some kind of logical representation could circumvent these disadvantages. Now, IE techniques provide the right data structures, but to access the information one needs the right retrieval algorithms. Logic-based Information Retrieval (IR) has been around, at least theoretically, since the mid 1980's [vR86]. An ongoing project at ILLC investigates to which extent underspecified reasoning and representation can be used for IR. We do not believe that these techniques can compete with IR methods for very large data collections, where logic-based techniques seem to be intractable, but we are confident about substantial quality improvements for smaller domains.

Finally, and coming from a completely different direction, there is work on the use of dynamic semantics to explain the meaning of programs in hybrid programming languages such as Alma-0 [ABPS] that combine the imperative and declarative programming paradigms. [vE98b] shows how dynamic predicate logic provides an adequate semantics for a non-trivial fragment of Alma-0, and how inference tools for dynamic predicate logic become verification tools for the hybrid programming language.

4 Conclusion

In this note we have identified the main concerns of doing inference for natural language semantics. We have illustrated these concerns by means of samples from ongoing research initiatives, and, in addition, we have listed what we take to be some of the main challenges and most promising research directions in the area.

Acknowledgments. Maarten de Rijke was supported by the Spinoza project 'Logic in Action' at ILLC, the University of Amsterdam.

References

[Abn96] S. Abney. Tagging and partial parsing. In *Corpus-Based Methods in Language and Speech*. Kluwer Academic Publishers, 1996.

[ABPS] K.R Apt, J. Brunekreef, V. Partington, and A. Schaerf. Alma-0: An imperative language that supports declarative programming. *ACM Toplas*, 1998. In press.

[BB98] P. Blackburn and J. Bos. *Representation and Inference for Natural Language*. Available from `http://www.coli.uni-sb.de/~bos/comsem/`, 1998.

[Boo91] G. Booch. *Object-Oriented Design*. Benjamin/Cummings, 1991.

[BS96] F. Baader and K. Schulz, editors. *Frontiers of Combining Systems*, Applied Logic Series. Kluwer Academic Publishers, 1996.

[CFvGG98] D. Crouch, C. Fox, J. van Genabith, and S. Glasbey, editors. *The FraCaS Book*. Unpublished manuscript, 1998.

[Dör97] J. Dörre. Efficient construction of underspecified semantics under massive ambiguity. In *Proc. of the 35th Annual Meeting of the ACL*, 1997.

[Jas97] J. Jaspars. Minimal logics for reasoning with ambiguous expressions. CLAUS-report 94, Universität des Saarlandes, 1997.

[Ker98] G. Kerdiles. Theorem proving in DPL. Manuscript, ILLC, Amsterdam, 1998.

[KR96] H. Kamp and U. Reyle. A calculus for first order Discourse Representation Structures. *Journal of Logic, Language and Information*, 5:297–348, 1996.

[MdR98a] C. Monz and M. de Rijke. Labeled resolution for discourse semantics. In *Proc. LD'98*. Dept. of Computer Science, Univ. of Freiburg, 1998.

[MdR98b] C. Monz and M. de Rijke. A resolution calculus for dynamic semantics. In J. Dix et al, editors, *Proc. JELIA'98*, LNAI 1489, pages 184–198. Springer, 1998.

[MdR98c] C. Monz and M. de Rijke. A tableaux calculus for ambiguous quantification. In H. de Swart, editor, *Proc. Tableaux-98*, LNAI 1397, pages 232–246. Springer, 1998.

[Rey93] U. Reyle. Dealing with ambiguities by underspecification: Construction, representation, and deduction. *Journal of Semantics*, 10(2):123–179, 1993.

[Rey95] U. Reyle. On reasoning with ambiguities. In *Proceedings of the 6th Meeting of the EACL*, pages 1–8, Dublin, 1995.

[Rey96] U. Reyle. Co-indexing labelled DRSs to represent and reason with ambiguities. In [vDP96].

[RG94] U. Reyle and D. Gabbay. Direct deductive computation on Discourse Representation Structures. *Linguistics and Philosophy*, 17(4):343–390, 1994.

[Sch96] M. Schiehlen. Semantic construction from parse forests. In *Proc. COLING'96*, 1996.

[vBtM97] J. van Benthem and A. ter Meulen, editors. *Handbook of Logic and Language*. North-Holland, 1997.

[vD96] K. van Deemter. Towards a logic of ambiguous expressions. In [vDP96].

[vDP96] K. van Deemter and S. Peters, editors. *Semantic Ambiguity and Underspecification*. CSLI Publications, 1996.

[vE98a] J. van Eijck. Axiomatizing dynamic logics for anaphora, 1998. CWI, Amsterdam.

[vE98b] J. van Eijck. Programming with dynamic logic. CWI, Amsterdam, 1998.

[vEJ96] J. van Eijck and J. Jaspars. Ambiguity and reasoning. CWI, Amsterdam, 1996.

[Vel97] F. Veltman. Dynamic predicate logic, 1997. ILLC, Amsterdam.

[vR86] K. van Rijsbergen. A non-classical logic for information retrieval. *The Computer Journal*, 29:481–485, 1986.