
Introduction

ANNIE ZAENEN, CLEO CONDORAVDI, VALERIA DE PAIVA

The introduction of the RTE (Recognizing Textual Entailment) paradigm (ref) in 2004 constituted a turning point in computational work on inference in natural language. That paradigm sees the task of determining the relation between two texts as one where assuming the truth of the first text, T, the thesis, leads to concluding to the (likely) truth or falsity of the other text, called H, the hypothesis. Later variants extended the relation also to contradictions. It has focused work on textual inference on tasks that are on the one hand feasible and on the other hand based on real texts. In this volume we present some of the work that arose from this conceptualization of the task, mostly but not only focussing on methods that involve logical formalizations. The volume is based on three workshops that we organized in 2011 and 2012: the *CSLI Workshop on Natural Logic, Proof Theory, and Computational Semantics* on April 8 and 9, 2011 at Stanford,¹ *Semantics for textual inference* on July 9/10, 2011 at the University of Colorado at Boulder,² and the CSLI workshop on *Semantic Representations for Textual Inference* on March 9 and 10, 2012 at Stanford.³

The first paper in this volume, “The BIUTEE Research Platform for Transformation-based Textual Entailment Recognition” by Asher Stern and Ido Dagan, is from the lab that introduced the current textual inference paradigm, RTE. The paper is an interim report on BIUTEE, an open source and open architecture platform that allows users to develop components for textual inference systems. The system consists of a preprocessing module that does parsing, named en-

¹<http://www.stanford.edu/~icard/logic&language/>

²<http://www.stanford.edu/~cleoc/Sem-Text-Inf/>

³<http://semanticrepresentation.stanford.edu>

tity recognition and coreference resolution and several modules that do inference recognition. The architecture allows the user to insert different lexical/knowledge resources (e.g. DIRT, WordNet, Wikipedia) and to add his/her own modules. The inference recognition engine itself is transformation-based, constructing derivations that go from text to hypothesis via rewrite rules. This procedure that can be seen as an instance of a proof, with the difference from traditional formal proofs being that the procedure allows for conclusions to likely true (false) instead of aiming for certainty. The system presented in this paper has meanwhile been integrated in an wider platform, EXCITEMENT, that allows more flexibility.⁴ The hope is that ultimately the community interested in textual inference will have a platform comparable to that provided to the (statistical) translation community by MOSES.⁵

There are two other papers describing full systems. They concentrate on building systems that allow true textual entailments and are not geared to allowing conclusions that are only *likely* to be drawn by native speakers.

The first paper “Is there a place for logic in recognizing textual entailment” by Johan Bos describes Nutcracker, a system that integrates Combinatory Categorical Grammar, Discourse Representation Theory (translated into first-order formulas through the reification of modality) and first order theorem proving. It argues that the main problem for such systems is the acquisition of the relevant *background knowledge* and shows how axioms deriving synonyms and hyponyms relations can be automatically derived from WordNet information. It proposes axioms for verbs of ‘saying’ that allow one to conclude that a reported event holds in the real world assuming one trusts the source and axioms learned from the RTE sets themselves. Several different theorem provers can be plugged into the system as well as different model builders that search for models up to a specified domain size. The analysis terminate either with (i) a proof, (ii) a finite counter model of size n , or (iii) neither. The system has, as expected, high precision but low recall as the knowledge acquisition problem is only partially solved.

Another take on textual inference is provided in the paper by Lenhart Schubert “NLog-like Inference and Commonsense Reasoning”. The paper describes the key properties of the Epilog system, an implementation of Episodic Logic, whose language is meant to provide a target representation for natural language. Epilog performs Natural Logic (refs) kind of inferences, but it goes beyond them in two ways: it

⁴<http://www.excitement-project.eu>

⁵<http://www.statmt.org/moses/>

can perform goal-directed and forward inferences, not just inferences with known premises and conclusions, as well as inferences based on lexical knowledge and language-independent world knowledge. Moving beyond narrow textual inference to common sense inference, essential to natural language understanding, requires addressing the problem of the “knowledge acquisition bottleneck”. The paper describes various under way and potential efforts for gathering knowledge from texts. These are based on the idea that one can recover some of the background knowledge assumed in one text from what is stated in another. In addition to lexical knowledge, necessary for inferences based on meaning, and world/common sense knowledge, necessary for general reasoning and true understanding, it recognizes semantic pattern knowledge. This knowledge comprises general ‘factoids’ that guide parsing and interpretation. Such factoids can also be used to acquire world knowledge via factoid strengthening and factoid sharpening as defined in the contribution.

The RTE challenge stressed the importance of working on real text and to take into account all the phenomena that contribute to making an inference valid or not. In this it is contrasted with earlier approaches that would concentrate on getting inferences involving specific semantic phenomena, e.g. quantifiers (Cooper et al. 1996, 1994) but after several years, the need to decompose the TE task into basic phenomena and the way these basic phenomena interact. Two papers address that issue.

Elena Cabrio and Bernardo Magnini in “Decomposing Semantic Inferences” look at the problem from an empirical angle. They analyze a TE data set looking at the nature of the inference (deductive, inductive, adductive) and the linguistic phenomena involved (e.g. synonymy, coreference, negation, active/passive alternation). Their results show that a huge amount of background information is required to approach the TE task. They also decompose the inferences of the thesis-hypothesis TH pairs into smaller atomic inference pairs consisting of a linguistic and an inference pattern. They conclude that “the polarity of most of the phenomena is not predictable for the logical judgments” and point out the consequences for attempts to learn from the annotated RTE data sets.

Assaf Toledo et al., in “Towards a Semantic Model for Textual Entailment Annotation”, develop a theoretical model of entailment recognition. The main idea consists in providing an interpreted lexicon, which specifies the semantic types and denotations of content and function words, and which ultimately serves as a target canonical representation for constructions in Text and Hypothesis sentences. After “binding to” an interpreted lexicon, an inferential relation can

be proven between T and H using predicate calculus and lambda calculus reduction, or disproven by the construction of a countermodel. Starting from the assumption that the model can incrementally incorporate increasingly complex phenomena, the authors concentrate on three prevalent inferential phenomena in the RTE data bases: intersective, restrictive, and appositive modification. At the same time, they acknowledge that interaction between phenomena might significantly complicate scaling up their model. The contrast between intersective vs. restrictive modification provides a nice illustration of how expressions with the same syntactic structure can have radically different inferential properties and how “binding to” an interpreted lexicon can model the difference.

Four of the remaining papers focus on logic more traditionally construed.

Alex Djalali’s “Synthetic Logic” reminds us that derivability plays as central a role in NL semantics as that of entailment and that it makes sense to develop a more proof-based approach to the logic of Natural Language, where one tries to capture the sorts of inferences speakers make in practice. Djalali considers MacCartney and Manning’s model of Natural Logic (MacCartney and Manning 2009; MacCartney 2009) as a kind of generalized transitive reasoning system and makes it, in the process, much easier to understand as a system of logic than the original system. His proof rules in Gentzen-style sequent calculus are divided into M-rules (which explain the composition of MacCartney and Manning relations) and D-rules, which correspond to structural properties of the MacCartney relations themselves. Djalali’s soundness and completeness proofs are crisp and enlightening, despite, like MacCartney and Manning, dealing only with a fragment of the algorithm developed for the implemented system NatLog.

The paper by Icard and Moss summarizes classic as well as more recent work on monotonicity reasoning in natural language. They first offer an informal overview of work on the Monotonicity Calculus, beginning with van Benthem and Snchez-Valencia in the late 1980s, and continuing on to the present day, including extensions, variations, and applications. Alongside examples from natural language, they also present analogous examples from elementary algebra, illustrating the fact that the Monotonicity Calculus makes sense as a more general system for reasoning about monotone and antitone functions over (pre)ordered sets. Following a discussion of current logical, computational, and psychological work on monotonicity in natural language, they develop a fully explicit Monotonicity Calculus using markings on simple types, with a well-defined language, semantics, and proof system, and culmi-

nating in an overview of soundness and completeness results, pointing to recent and forthcoming work by the authors.

The paper by Ian Pratt-Hartman, “The Relational Syllogistic Revisited” is part of the tradition of extending the original syllogistic calculus. In previous work Moss and Pratt-Hartman (refs) introduced the relational syllogistic, an extension of the language of classical syllogisms in which predicates are allowed to feature transitive verbs with quantified objects. They showed that this relational syllogistic does not admit a finite set of rules whose associated direct derivation relation is sound and complete. Thus for the relational syllogistic, indirect reasoning, in the form of *reduction ad absurdum* is essential. Pratt-Hartmann’s paper in this volume presents a modest extension of the relational syllogistic language which is sound and complete, as desired for direct proofs. This shows that the impossibility of providing a finite rule-set for the relational syllogistic can be overcome by a modest increase in expressive power. The proof is quite complicated. Still one important conclusion from the existence of a sound and complete proof system defined by a finite set of syllogism-like rules such as the ones here is that adding relations (as transitive verbs) to a basic syllogistic logic does not represent a logical ‘boundary’ with respect to the expressiveness of fragments of natural language. From the previous result of Moss and Pratt-Hartmann one could get the (wrong) impression that syllogistic extensions could not be provided for transitive verbs while keeping the system sound and complete. The system RE introduced shows that this is not the case, soundness and completeness are within reach.

In “Intensions as Computable Functions”, Shalom Lappin deals with an long standing problem of intensional logic, proposing a type theoretical solution. Classical intensional semantic representation languages, like Montague’s Intensional Logic do not accommodate fine-grained intensionality well. In the traditional work intensional identity is reduced to equivalence of denotation across possible worlds and logically equivalent expressions are semantically indistinguishable. Thus not only all mathematical truths are the same, but also the denotations of belief statements are all logically equivalent. Lappin’s paper shows that terms in the type theory PTCT (Property Theory with Curry Typing) proposed by Fox and Lappin (to appear) constitute an alternative intensional semantic representation framework. PTCT uses two notions of equality: intensional identity and extensional equivalence, and while intensional identity implies extensional equivalence, the converse is not true. Their fine-grained notions allow PTCT to prove the equivalence of mathematical truths, while allowing the non-equivalence of all belief statements. Here, Lappin proposes to characterize the distinction

between intensional identity and provable equivalence *computationally* by invoking the contrast between operational and denotational semantics in programming language. Since the terms of PTCT are lambda-expressions that encode computable functions and since Lappin has identified these with the intensions of words and phrases in natural language, given the distinction between denotational and operational meaning, he can interpret the non-identity of terms in the representation language as an operational difference in the functions that these terms express. In other words if the terms compute the same result set through different sets of procedures, they are different. This approach factors modality and possible worlds out of the specification of intensions.

While the series of conferences this volume is based on had several talks on machine learning approaches to semantics, none of the authors could find the time to write up their work in a way to that would have fitted it. But “Frege in Space”, the contribution of Marco Baroni, Raffaella Bernardi and Roberto Zamparelli, fills the gap with an extensive discussion of distributional semantics and its relation to traditional compositional semantics. One of the problems of classical approaches to semantics is that most lexical items are unanalyzed (‘prime semantics’). Distributional semantics proposes a way to handle this through a ‘the meaning of a word, is the company it keeps’ approach. The approach has shown to deliver interesting results for noun and noun adjective combinations. It is less clear how to extend it to argument taking predicates and how to handle compositionality. “Frege in Space” proposes an ambitious program to do so and shows the way for a synthesis between both approaches.

We would like to dedicate the volume to our colleagues of the now defunct PARC NLTT group, especially Danny Bobrow, Dick Crouch, Lauri Karttunen, Ron Kaplan, Martin Kay, Tracy King and John Maxwell. They kindled our interest in the problems of the relation between logic, computation and natural language understanding that the volume aims to be a contribution to.

References

- Cooper, R., R. Crouch, J. van Eijck, C. Fox, J. van Genabith, J. Jaspars, H. Kamp, D. Milward, M. Pinkal, M. Poesio, and S. Pulman. 1996. Using the framework. Technical Report LRE 62-051 D-16, The FraCaS Consortium.
- Cooper, R., R. Crouch, J. van Eijck, C. Fox, J. van Genabith, J. Jaspars, H. Kamp, M. Pinkal, M. Poesio, S. Pulman, and E. Vestre. 1994. Describing the approaches.

- Fox, C. and S. Lappin. to appear. *Formal Foundations of Intensional Semantics*. Blackwell.
- MacCartney, B. 2009. *Natural Language Inference*. Ph.D. thesis, Stanford University.
- MacCartney, Bill and Christopher Manning. 2009. An extended model of natural logic. In *The Eighth International Conference on Computational Semantics (IWCS-8)*.