

Deep Natural Language Understanding of News Text

Jaya Shree,¹ Emily Liu,² Andrew S. Gordon,¹ and Jerry R. Hobbs¹

¹University of Southern California, Los Angeles CA USA

²Duke University, Durham, NC USA

shree@usc.edu, emily.f.liu@duke.edu, gordon@ict.usc.edu, hobbs@isi.edu

Abstract

Early proposals for the deep understanding of natural language text advocated an approach of “interpretation as abduction,” where the meaning of a text was derived as an explanation that logically entailed the input words, given a knowledge base of lexical and commonsense axioms. While most subsequent NLP research has instead pursued statistical and data-driven methods, the approach of interpretation as abduction has seen steady advancements in both theory and software implementations. In this paper, we summarize advances in deriving the logical form of the text, encoding commonsense knowledge, and technologies for scalable abductive reasoning. We then explore the application of these advancements to the deep understanding of a paragraph of news text, where the subtle meaning of words and phrases are resolved by backward chaining on a knowledge base of 80 hand-authored axioms.

Introduction

Typical natural language applications today do an excellent job of performing relatively shallow tasks, such as determining whether a text expresses a predominantly positive or negative sentiment, or doing a fairly direct translation of a string from one language to another. But when people read a text, they construct a much richer model of it than is evident in the output of these applications. In the project described in this paper, we have attempted to explicate all the inferences that people draw in comprehending one 4-sentence, 75-word paragraph from business news, encode the necessary knowledge in first-order logic, and then use an abductive theorem-prover to identify the correct interpretation for the entire paragraph. The knowledge base we constructed for this task consisted of only those axioms needed for the target interpretation, but were written in a general

style that did not cater to the requirements of this specific text. Our goal was to explore the scope of the axioms that were required, and to determine whether we could derive the correct interpretation of the whole paragraph using recent advances in incremental abductive reasoning.

The paragraph we used for this exploration was as follows:

Uber’s innovations reflect the changing ways companies are managing workers amid the rise of the freelance-based “gig economy.” Its drivers are officially independent business owners rather than traditional employees with set schedules. This allows Uber to minimize labor costs, but means it cannot compel drivers to show up at a specific place and time. And this lack of control can wreak havoc on a service whose goal is to seamlessly transport passengers whenever and wherever they want.

Among the problems this text poses are the following:

1. What are the relations between Uber and “companies”, and between “innovations” and “changing ways”, as indicated by the verb “reflect”? What does “reflect” mean here?
2. What causal information is provided by the preposition “amid”?
3. What is the relation between gigs and the economy in “gig economy” and how does that relate to “freelance-based”?
4. What are the relations among “workers”, “drivers”, “employees”, and “labor”?

5. What are the relations among “managing workers”, “independent”, “set schedules”, “cannot compel”, and “lack of control”?
6. What are the relations among “schedules”, “a specific place and time” and “whenever and wherever”?
7. Can we automatically recognize the discourse structure of this paragraph? That is, can we verify the contrast relation between the two clauses of Sentence 3, the causal relation between Sentence 2 and Sentence 3, and the causal relation between Sentences 2-3 and Sentence 4?
8. In the first sentence of the next paragraph there is the referring expression “this fundamental problem”. Can we resolve this to Uber’s lack of control of its workers? Why is the lack of control a problem?

In this project we were able to address all these problems and enhance the abductive theorem-prover to the point where the proof graph it produced correctly encoded the answers to all these questions. Obviously scaling up will require much more knowledge and more ways of dealing with the combinatorial explosion that this will trigger. But this small-scale exploration has already led to solutions to significant problems and points the way toward larger-scale experiments.

Interpretation as Abduction

Hobbs et al. (1993) presented an approach to language interpretation that rooted in the logical reasoning approach of abduction, or inference to the best explanation. The approach, *interpretation as abduction*, provided an integrated account of syntax, semantics, and pragmatics as a type of search problem, where the aim was to find a set of assumptions that would logically entail the observable words of a text, given a knowledge base of linguistic and commonsense axioms. Among the worked-out examples provided in this paper, the interpretation of a short sentence (“The Boston office called.”) is disambiguated by assuming unmentioned entities and relations that connect the words to our commonsense understanding of the world (a person in the office located in Boston made the call). Given sufficiently rich knowledge bases, logical abduction produces many candidate solutions, necessitating a means of favoring

some interpretations over others. Here Hobbs et al. describe a scheme of *weighted abduction*, implemented in the TACITUS system, where literals in knowledge base axioms are annotated with numerical weights that transfer and scale costs associated with the input text to terms in the solutions, where the least-cost set of assumptions becomes the preferred interpretation.

Although hugely influential at the time, the proposal of Hobbs et al. (1993) left many challenges that needed to be overcome in order to apply the approach to unconstrained natural language text, including:

1. What commonsense knowledge is necessary, and how should it be obtained?
2. How should knowledge base axioms be annotated with weights?
3. What algorithm would allow logical abduction to scale to large documents and knowledge bases?

In the decades since the publication of this early work, research in natural language processing has followed a radically different course, beginning first with the data-driven approaches of statistical NLP, leading then to contemporary deep learning approaches that treat syntax, semantics, and pragmatics as implicit, latent encodings of neural network activations. Despite these changing trends, the intervening years has seen enormous progress on addressing the three challenges listed above.

Parsing the Logical Form of the Text

Hobbs et al.’s (1993) proposal viewed the problem of syntax as the conversion of a sequence of words into the *logical form of the text*, where individual morphemes in the text were reified as literals whose arguments encoded their syntactic relationship to other elements, following Hobbs (1985). In their original proposal, this form was abductively derived via a knowledge base of syntactic axioms. However, the emergence of high-accuracy statistical parsers makes this a less than optimal approach to syntactic analysis.

Beginning first with systems that generated the logical form of the text from constituency parsers (Rathod, 2005), recent interpretation pipelines have opted to generate these forms using the output of English Slot Grammar parsers, Combinatory Categorical Grammar parsers, or syntactic

dependency parsers, e.g., (Ovchinnikova et al., 2014b,a).

In the current exploration we first used the Boxer system (Bos, 2008) to parse the text and translate it into logical form. This achieved about 80% accuracy, where the two types of mistakes were the usual attachment ambiguities and misalignments between Boxer’s output and the logical representation we required. Fixing the latter would have been tedious and unenlightening, and the former was one class of problems we expected to solve with inferencing. So for the rest of this exploration we began with a manually constructed logical form for the text. This enabled us to focus on the less well-understood issues around semantics and pragmatics.

Encodings of Commonsense Knowledge

Spurred by promising results on open-domain question-answering (Moldovan et al., 2003), recent interpretation pipelines have relied on broad-coverage knowledge bases of axioms derived from lexical resources, e.g., using the relations and the glosses in WordNet (Fellbaum, 1998). Ovchinnikova (2012) typifies this approach, where axioms automatically derived from lexical resources are used in abductive reasoning applied to the tasks of recognizing textual entailment, semantic role labeling, and the interpretation of noun dependencies.

Complementing these automatically-derived approaches has been continued progress on the large-scale manual formalization of commonsense knowledge, most notably in the area of commonsense psychology (Gordon and Hobbs, 2017). While the hand-authoring of commonsense domain theories affords a certain level of precision that is not readily obtained using automatic methods, it requires an additional set of so-called *lexical axioms* to bridge the semantic gap between words and the literals used these theories. Montazeri (2014) demonstrates how many of these lexical axioms can be semi-automatically derived by annotating smaller sets of words from large-scale lexical resources.

Probability-ordered Abduction

A frequent critique of Hobbs et al.’s (1993) proposal for weighted abduction was that the weights assigned to knowledge base literals seem arbitrary, lacking a connection to more commonly used numerics such as probability. Indeed, Ovchinnikova

et al. (2013) showed that the weights used in weighted abduction cannot be interpreted as probabilities. These concerns have led several researchers to pursue different abductive reasoning frameworks that are more solidly grounded in probability theory, e.g., Blythe et al.’s. (2011) and Kate and Mooney’s (2009) implementations of abduction using Markov Logic Networks.

A more recent advance has been Gordon’s (2016) Etcetera Abduction, which builds on earlier work by Poole (1991) on estimating the probability of solutions in Horn-clause abduction. Gordon noted that these solutions could be expressed as conjunctions of so-called *etcetera literals* that reified the uncertainty in a defeasible axiom, following Hobbs et al.’s (1993) variant of McCarthy’s (1986) \neg *abnormal* literal, and showed that their probabilities could be readily interpreted as prior and conditional probabilities.

Scalable Abductive Reasoning

Implementations of abductive reasoning must carefully manage the combinatorial search process in order to process long passages of text with sufficient depth, as naive implementations will fail to scale when presented with more than a handful of words. In recent years, several researchers have explored the application of optimized solvers to abductive reasoning, aiming to find solutions for increasingly larger input texts. Inoue and Inui (2013) describe an approach that formulates a weighted abduction problem as a set of linear equations that can be passed to a contemporary integer linear programming solver. Kazeto et al. (2015) further this approach by pre-estimating the relatedness between predicates, and implement their solution in a robust software library called Phillip¹. Inoue and Gordon (2017) pursue a similar approach within the framework of Etcetera Abduction.

Although the use of optimized solvers allows for substantially longer input, the combinatorial nature of the search problem ultimately limits the scalability of these approaches. An alternative approach to scalable Etcetera Abduction is pursued in Gordon (2018), in which arbitrarily long input sequences are interpreted incrementally, using the best interpretations of previous segments as contexts for the interpretation of the current input window. Given finite window sizes and a finite

¹<https://github.com/kazeto/phillip>

beam of running hypotheses, incremental Etcetera Abduction can fail to find the overall best (most-probable) solution, particularly when supporting evidence appears over long distances in the input stream. However, Gordon demonstrated that even with modest window and beam sizes, the available implementation² can find near-optimal solutions for interpretation problems with several dozen input literals.

Interpretation of a Paragraph of News Text

To explore the application of contemporary incremental abductive reasoning engines to the interpretation of naturalistic texts, we attempted to use Gordon’s (2018) implementation of incremental Etcetera Abduction to interpret a passage of news text. We chose a paragraph from the New York Times article “How Uber Uses Psychological Tricks to Push Its Drivers’ Buttons” by Noam Scheiber, which appeared online on April 2, 2017,³ (presented in the introduction of this paper). The passage explains how the company has undertaken an extraordinary experiment in behavioral science to subtly entice an independent workforce to maximize its growth. It starts with explaining how Uber has changed the ways of managing workforce by making them feel more independent, and then it explains the contrast between how its strategy is good in minimizing labor cost but also bad because it can no longer compel its drivers. Finally, it explains how this lack of control is damaging services provided by Uber, which is the fundamental problem they are trying to solve. Our overall conception of the coherence structure of this passage is depicted in Figure 1.

Our aim in this exploration was to determine if we could hand-author a set of first-order axioms (definite clauses) such that the deep meaning of this passage could be automatically recovered following the “interpretation as abduction” approach.

Logical Form of the Text

We began our exploration by applying a contemporary CCG parser (Bos, 2008) to generate the logical form of the text. After some preliminary work with the resulting output, we judged that the automatically-generated logical form of the text

²<https://github.com/asgordon/EtcAbductionPy>

³<https://www.nytimes.com/interactive/2017/04/02/technology/uber-drivers-psychological-tricks.html>

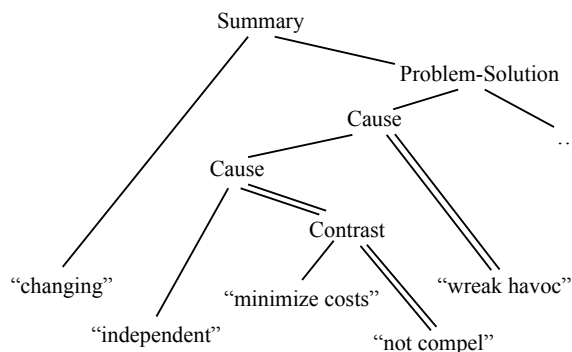


Figure 1: Overall coherence structure of the text

contained too many errors to serve as the starting point for our current investigation. For this reason, we instead opted to hand-author the logical form of the text for each of the four sentences in the passage. The first sentence, “Uber’s innovations reflect the changing ways companies are managing workers amid the rise of the freelance-based gig economy,” was encoded as follows:

```

(uber u) (poss u x13)
(innovation x13 u x12 x14)
(reflect x13 w11)
(changeIn x15 w11)
(prog x15) (way' e11 w11 e12)
(plural w11 s11)
(company u) (plural u s12)
(manage' e12 u w12) (prog e12)
(pres e12) (workFor w12 u)
(amid w11 r11) (rise r11 x11)
(of r11 x11) (freelance x16)
(base x11 x16) (gig t11 w12 u)
(nn t11 x11) (economy x11)
  
```

Knowledge Base and Interpretations

To derive the target interpretation for this text passage, we hand-authored a set of 80 axioms (definite clauses) consisting of only those needed as part of the abductive proof structure. While required for this particular text, these axioms were written in a general style so as to better assess the feasibility of the approach on general textual input. Each axiom was assigned an arbitrary probability, reified as an etcetera literal as required by Etcetera Abduction, which will become more relevant as the size of the knowledge base increases. The following are two examples of 80 axioms authored during the course of this exploration.

```

;; failure of u to control is bad for u
(if (and (cannot' e35 e37)
  
```

```

(control' e37 u d)
(etcl.badFor 0.9 e35 u)
(badFor e35 u)

;; meaning of allow
(if (and (causallyInvolved e31 e32 e33)
        (etcl.allow' 0.9 e31 e32 e33))
    (allow' e31 e32 e33))

```

The following items, depicted in Figure 2, illustrate the kind of knowledge we encoded and the subtleties of text meaning we are capturing in the proof graphs that represent the interpretations.

Meaning of “reflect”

The word “reflect” in this text has deep semantics that expresses why innovation done by Uber is a reflection of changing ways of managing workers. In general terms, an event reflects another event when the former causes knowing the later (Figure 2a).

Meaning of “amid”

To understand the causal force of the preposition “amid” we see first a change in managing as one instance of a change in the economy, since managing is one task in producing goods and services, which is the sort of activity that economies are made up of. Second, we see a change in a whole defeasibly causing change in its parts. As a by-product of explaining “amid” in this way, we resolve the attachment of “amid” to “changing ways” rather than “reflect”, “managing” or “workers” (Figure 2b).

Meaning of “rather than”

“Rather than” indicates a contrast, so the interpretation should say what that contrast is. Owners contrast with employees in that the former are not managed by any company, and employees that are managed with some schedule set by the company they work for (Figure 2c).

Contrast between clauses “minimize cost but cannot compel”

Minimized labor cost is good for Uber. However, the inability to compel drivers to show up at specific schedule is not good for them. This contrast between things that are good and bad for Uber explains the presence of the “but” in the sentence (Figure 2d).

Meaning of “This... means”

The coherence relation between sentence 2 and sentence 3 is the predicate-argument relation where “means” is the predicate and the argument is the drivers’ independence in sentence 2. The implicational meaning of “mean” is justified by the implicational relation between independence and lack of control.

Causal relationship between sentences 2, 3, and 4

The occurrence of “wreaking havoc” is due to Uber’s lack of control (sentence 4), because they cannot compel drivers (sentence 3, which in turn is caused by drivers being independent, sentence 2). Therefore, there is a causal coherence relationship between “wreaking havoc” and “cannot control” (Figure 2f).

“The fundamental problem”

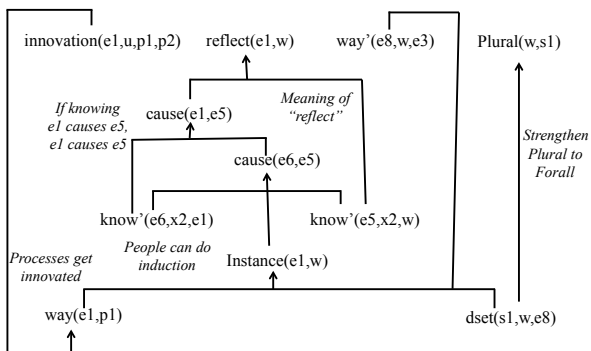
The first sentence of the paragraph that follows this one contains the referring expression “this fundamental problem”. We are able to resolve its referent as follows: the lack of control causes damage to the transportation services provided by Uber, which is bad for Uber because it is not able to achieve its goal. And hence, this damage to the service is the fundamental problem for Uber (Figure 2f).

Figure 2 shows how the axioms in our knowledge base resolve each of the challenges listed above, where each of the six graphs are subgraphs of the final interpretation derived for this paragraph of news text.

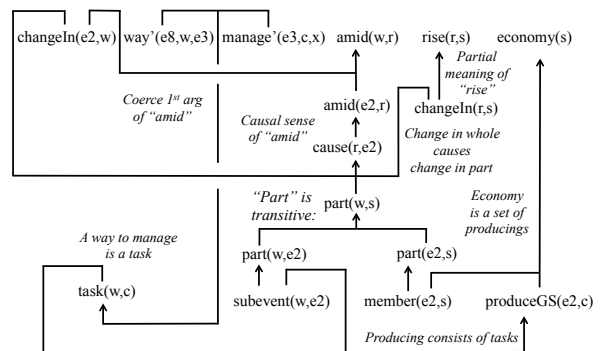
Over-unification of Literals

In using incremental Etcetera Abduction to derive the target interpretation for this paragraph, we encountered problems that required changes to the available implementation of the reasoning algorithm.


The principal problem we faced was the over-merging of assumptions. In the abductive framework most coreference problems are resolved by inferring implicit redundancies from different parts of the text. That is, entities are identified with each other because they share a property. The difficulty is that if not carefully controlled, this process can identify entities that are not coreferential. We implemented two heuristics that virtually eliminated this problem.



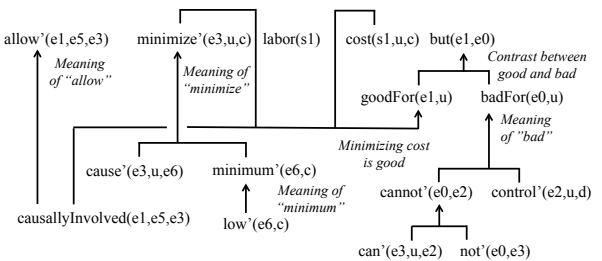
a) innovations **reflect** the changing ways



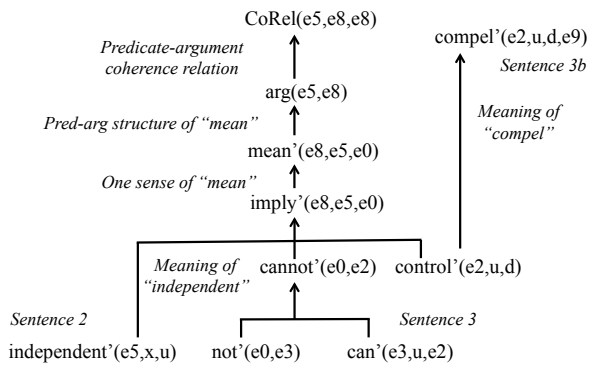
b) changing ways...managing...amid the rise of the...economy



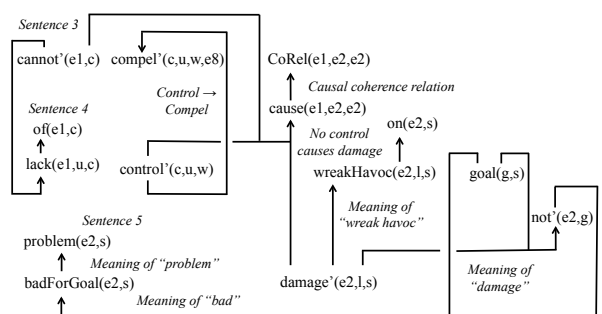
c) independent business owners **rather than** traditional employees with set schedules



d) allows Uber to **minimize labor costs, but means**



e) independent business owners...**This means** that it cannot compel



f) cannot compel drivers...**wreak havoc** on a service whose goal...solve this fundamental problem...

Figure 2: Six examples of coherence relationships in a paragraph of news text

The first heuristic was this: Suppose we know $animal(x)$, $animal(y)$, $dog(x)$ and $cat(y)$. Should we identify x and y on the basis of their both being animals? Obviously not, because they have other, contradictory properties—dog and cat. The general rule schema underlying this heuristic is

$$p(\dots, x, \dots) \wedge q(\dots, x, \dots) \rightarrow \perp$$

or equivalently,

$$p(\dots, x, \dots) \wedge q(\dots, y, \dots) \rightarrow x \neq y$$

We implemented this class of constraints efficiently in terms of bit matrices.

An example of the second class of constraints is that something cannot be a part of itself. That is, you can't have $part(x, x)$. The general rule schema for this heuristic is

$$p(\dots, x, \dots, x, \dots) \rightarrow \perp$$

or equivalently,

$$p(\dots, x, \dots, y, \dots) \rightarrow x \neq y$$

The two heuristics together blocked every illegitimate case of merging in our data, while letting the correct ones through. As a side-effect of this, it greatly speeded up processing.

Depth, Window, and Beam Parameters

At the beginning of our efforts, we attempted to use Gordon's original implementation of Etcetera Abduction (Gordon, 2016) to interpret each sentence individually, but found that the size of the input was too great, leading to a combinatorial explosion in the search space. We subsequently used the implementation of incremental Etcetera Abduction (Gordon, 2018), treating the entire paragraph as a single input sequence, ignoring sentence boundaries. To achieve our target interpretations, we modified the software to prevent existentially quantified variables from being turned into constants after each increment of the interpretation process.

Using our hand-authored knowledge base of 80 axioms, we were able to achieve our target interpretation of this text using the modified version of incremental Etcetera Abduction. We found that this interpretation could be found using a modest window parameter of only four literals, and a

very small beam of two running hypotheses. However, a large depth parameter of seven backward-chaining steps was required given our formalization of the requisite semantic knowledge, which is substantially larger than required for previous non-linguistic interpretation problems (Gordon, 2016, 2018). The final abductive proof graph consisted of 32 assumptions of prior probabilities and 71 assumptions of conditional probabilities in order to logically entail the logical form of this passage of text given our knowledge base.

Conclusions

In the end we were able to run the entire 75-word paragraph and produce the correct interpretation (proof graph). No incorrect identifications of entities were made. This interpretation included the correct coherence structure and the correct resolution of the definite noun phrase “this fundamental problem” to the lack of control referenced several places in the paragraph. Our success in achieving the target interpretation using incremental logical abduction demonstrates that recent technological advances constitute real progress toward practical implementations.

As encouraging as this result is, there are several obvious questions. First, how will it do on the next paragraph, and the next? How large a knowledge base will be needed before previously unseen texts can be processed and understood? How should that knowledge base be constructed? Given that large knowledge base, can the combinatorial explosion be contained for realistically long and complex texts? What further techniques will be required beyond the incremental processing we employed here?

This work does offer insight into how a knowledge base can be devised to best address subtleties that are prevalent in real-world text. About a third of the axioms we encoded were essentially lexical knowledge, of the sort that standard lexical resources can be expected to provide, such as the implicational sense of “means” and the relation between “workers” and “labor”. Another third of the axioms were rules we had already encoded in core commonsense theories (Gordon and Hobbs, 2017), such as the transitivity of “part” in the theory of composite entities, and the axiom in the theory of knowledge management that says people can do induction, i.e., draw general conclusions from specific instances. On the other hand,

we also had to encode axioms to coerce or shuffle arguments around, such as coercing from a causal relation to the effect, in order to get the predicate-argument relations right. These rules had a very ad hoc feel, and it would be good to develop a more general approach to this class of problems.

Our analysis of this one paragraph also helped gauge the utility of current technologies for identifying the logical form of the text, where we see the need for further improvement. Likewise, our efforts identified a number of problems with the available implementation of incremental Etcetera Abduction, which we addressed by making modifications to this software. Our expectation is that we would have faced these problems regardless of the passage selected for our analysis, and that future analyses of a similar sort would uncover additional problems to address. In this respect, we see a path forward in this line of research that analyzes different and longer passages of text, uncovering and solving new technical problems and further characterizing the scope of the knowledge engineering requirements. As the software architecture becomes more robust and the knowledge base becomes well-understood, efforts can be increasingly directed toward the automatic acquisition of the axioms required for the deep understanding of arbitrary news text.

The results of this exploration have been good enough to encourage us to continue this line of research, but at best, it has so far given us only a partial proof of possibility of abductive interpretation of complex real-world discourse.

Acknowledgments

This work is supported by Contract W911NF-15-1-0543 with the US Defense Advanced Research Projects Agency (DARPA). This research was supported by the Office of Naval Research, grant N00014-16-1-2435.

References

James Blythe, Jerry R. Hobbs, Pedro Domingos, Rohit J. Kate, and Raymond J. Mooney. 2011. Implementing weighted abduction in markov logic. In *Proceedings of the Ninth International Conference on Computational Semantics, IWCS '11*, pages 55–64, Stroudsburg, PA, USA. Association for Computational Linguistics.

Johan Bos. 2008. Wide-coverage semantic analysis with boxer. In *Proceedings of the 2008 Conference*

on Semantics in Text Processing, pages 277–286. Association for Computational Linguistics.

Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.

Andrew S. Gordon. 2016. Commonsense interpretation of triangle behavior. In *Thirtieth AAAI Conference on Artificial Intelligence*, pages 3719–3725, Palo Alto, CA. AAAI Press.

Andrew S. Gordon. 2018. Interpretation of the Heider-Simmel film using incremental etcetera abduction. *Advances in Cognitive Systems*, 6:1–16.

Andrew S. Gordon and Jerry R. Hobbs. 2017. *A formal theory of commonsense psychology: How people think people think*. Cambridge University Press, Cambridge, UK.

Jerry R. Hobbs. 1985. Ontological promiscuity. In *Proceedings of the 23rd Annual Meeting on Association for Computational Linguistics*, pages 60–69. Association for Computational Linguistics.

Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63(1-2):69–142.

Naoya Inoue and Andrew S. Gordon. 2017. A scalable weighted Max-SAT implementation of propositional etcetera abduction. In *Proceedings of the 30th International Conference of the Florida AI Society*, pages 62–67, Palo Alto, CA. AAAI Press.

Naoya Inoue and Kentaro Inui. 2013. Ilp-based inference for cost-based abduction on first-order predicate logic. *Journal of Natural Language Processing*, 20(5):629–656.

Rohit J. Kate and Raymond J. Mooney. 2009. Probabilistic abduction using markov logic networks. In *Proceedings of the IJCAI-09 Workshop on Plan, Activity, and Intent Recognition (PAIR-09)*.

John C. McCarthy. 1986. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28:89–116.

Dan Moldovan, Christine Clark, Sanda Harabagiu, and Steve Maiorano. 2003. Cogex: A logic prover for question answering. In *Proceedings of the North American Chapter of the Association for Computational Linguistics*.

Niloofar Montazeri. 2014. *Building a Knowledgebase for Deep Lexical Semantics*. Ph.D. thesis, University of Southern California.

Ekaterina Ovchinnikova. 2012. *Integration of World Knowledge for Natural Language Understanding*. Atlantis Press.

Ekaterina Ovchinnikova, Andrew S. Gordon, and Jerry R. Hobbs. 2013. Abduction for discourse interpretation: A probabilistic framework. In *Proceedings of the Joint Symposium on Semantic Processing*, pages 42–50.

Ekaterina Ovchinnikova, Ross Israel, Suzanne Wertheim, Vladimir Zaytsev, Niloofar Montazeri, and Jerry Hobbs. 2014a. Abductive inference for interpretation of metaphors. In *Proceedings of the Second Workshop on Metaphor in NLP*, pages 33–41.

Ekaterina Ovchinnikova, Niloofar Montazeri, Theodore Alexandrov, Jerry R. Hobbs, Michael C. McCord, and Rutu Mulkar-Mehta. 2014b. Abductive reasoning with a large knowledge base for discourse processing. In H. Hunt, J Bos, and S. Pulman, editors, *Computing Meaning*, volume 4, pages 107–127. Springer.

David Poole. 1991. Representing bayesian networks within probabilistic horn abduction. In *Proceedings of the Seventh Conference on Uncertainty in AI*, pages 271–278.

Nishit Rathod. 2005. LFToolkit.
<https://www.isi.edu/hobbs/LFToolkit/>.

Kazeto Yamamoto, Naoya Inoue, Kentaro Inui, Yuki Arase, and Juníchi Tsugii. 2015. Boosting the efficiency of first-order abductive reasoning using pre-estimated relatedness between predicates. *International Journal of Machine Learning and Computing*, 5(2):114–120.