

Intension, Attitude, and Tense Annotation in a High-Fidelity Semantic Representation

Gene Kim and Lenhart Schubert

University of Rochester

Department of Computer Science

{gkim21, schubert}@cs.rochester.edu

Abstract

This paper describes current efforts in developing an annotation schema and guidelines for sentences in Episodic Logic (EL). We focus on important distinctions for representing modality, attitudes, and tense and present an annotation schema that makes these distinctions. EL has proved competitive with other logical formulations in speed and inference-enablement, while expressing a wider array of natural language phenomena including intensional modification of predicates and sentences, propositional attitudes, and tense and aspect.

1 Introduction

Episodic Logic (EL) is a semantic representation and knowledge representation that extends FOL to more closely match the expressivity of natural languages. It echoes both the surface form of language, and more crucially, the semantic types that are found in all languages. Some semantic theorists view the fact that noun phrases denoting both concrete and abstract entities can appear as predicate arguments (*Aristotle*, *humanity*, *the fact that there is water on Mars*) as grounds for treating all noun phrases as being of higher types (e.g., second-order predicates). EL instead uses a small number of reification operators to map predicate and sentence intensions to individuals. As a result, quantification remains first-order (but allows quantified phrases such as *most people who smoke*, or *hardly any errors*). Another distinctive feature of EL is that it treats the relation between sentences and episodes (including events, situations, and processes) as a *characterizing relation*, written “***”. This coincides with the Davidsonian treatment of events as extra variables of pred-

icates, as long as we restrict ourselves to positive, atomic predications. But it also allows for logically complex characterizations of episodes, such as episodes of not eating anything all day, or of each superpower menacing the other with its nuclear arsenal (Schubert, 2000).

EL has been shown to be suitable for deductive inference, uncertain inference, and Natural-Logic-like inference (Morbini and Schubert, 2009; Schubert and Hwang, 2000; Schubert, 2014). Most recently, Kim and Schubert (2016) developed a system that generated EL verb gloss axioms from WordNet, which enabled inferences that were competitive with the state-of-the-art even with greater expressivity.

In a supplementary document for the above paper, Kim and Schubert present an illustration of EL appropriately handling the intensional predicate modifier *nearly*. The illustration uses the gloss for the second sense of *stumble*, which is *miss a step and fall or nearly fall* and shows that using EL as the representation enables inferences that are not possible using intersective predicate modification.

We are currently underway on an annotation project that is aimed at creating a corpus that can be used to train a reliable, general-purpose ULF (unscoped logical form) transducer. ULF is a preliminary, indexical EL representation with syntactic marking of residual scope ambiguity. If the project is successful, it would overcome the primary limitations of Kim and Schubert’s work: scalability and accuracy.

2 Project Overview

Kim and Schubert’s system relies in part on manually specified transduction rules that try to construct complete, interpretable sentences from WordNet verb glosses, which are in a stylized,

phrasal form. Often it is not enough to just expand a gloss into a sentence (understandable to a human reader) to enable reliable semantic parsing. The sentence must often be further transformed and broken into multiple, simpler sentences before somewhat reliable semantic parsing is possible. Even then, both the transduction rules and semantic parsing may introduce errors into the resulting definitional axioms(s). Kim and Schubert note that almost a third of the extracted axioms had come from EL formulas that were erroneously transduced from English. These were due to linguistic phenomena that did not show up in the development set or due to sheer sentence complexity. Such errors would become even more of a problem for noun glosses, which can contain quite complex descriptive material. A reliable, general-purpose, semantic parser would eliminate most of this labor and improve the project’s scalability. We expect that a statistical semantic parser trained on a large corpus would have better coverage of linguistic phenomena and function robustly for larger sentences.

We plan to annotate several thousand sentences from topically varied sources and have experimented so far with the Brown corpus, the Gigaword newswire corpus and *The Little Prince*. Annotating ULF has many advantages over directly annotating EL logical forms. ULF enables the separation of determining the semantic type structure from replacing indexial expressions and disambiguating quantifier scopes, word senses, and anaphora – tasks which in general require the context of the sentence to resolve. Since we are tackling a range of subtle semantic phenomena beyond those ordinarily considered, this decomposition is likely to achieve better results than a fell-swoop approach. An undisambiguated representation also has the advantage of adaptability to a wide range of tasks – a topic discussed in depth by Bender et al. (2015).

3 Semantic Handling of Intension and Attitudes in EL

This section briefly describes how the semantic interpretation of EL enables proper handling of intension and attitudes. For a fuller description of EL semantics please refer to (Schubert and Hwang, 2000).

3.1 Intensional Modifiers

EL semantic types distinguish *predicate* modifiers from *sentence* modifiers. Predicate modifiers are interpreted as mappings from predicate meanings to predicate meanings, where these are intensional functions based on possible episodes (whose maximal elements are possible worlds). This enables proper interpretation of non-intersective predicate modifiers such as *very*, *fairly*, and *big*, including intensional ones such as *nearly*, *fake*, and *resemble*. For example, EL can express the following fact:

```
(all x [[x (fake.a flower.n)] =>
  [(not [x flower.n]) and
    [x (resemble.v flower.n)]]])
```

Similarly, *intensional* sentence modifiers (e.g., *probably*, *according to Fox News*) map sentence intensions to sentence intensions, whereas *extensional* sentence modifiers (e.g., *in the forest*, *at dawn*) become simple predications about episodes upon “deindexing”.

3.2 Attitude Predicates

Attitude predicates such as *assert*, *believe*, and *assume* relate an individual to a proposition. Propositions are treated as abstract entities, namely, reified sentence intensions. Of course an attitude predication can be true without the proposition being true. Unlike some semantic representations, EL does not conflate propositions with episodes. Episodes are real (often physical) entities occupying time intervals, whereas propositions are informational entities. Propositions are formed from sentences using a *that* operator, since they are most commonly instantiated as that-clauses in English (e.g., Jim knows *that* there is water on Mars).

4 ULF Syntax

This section will act as a brief introduction to ULF syntax for understanding the examples presented. Atoms in ULF that correspond to lexical entries are followed by a suffix derived from the part of speech. Atoms without the suffix are special EL operators that correspond to particular morpho-syntactic phenomena; see the first visualization in Figure 1 for examples. ULF uses three different brackets: round brackets to indicate prefixed operators, square brackets for sentential formulas with infix predicates, and angle brackets for (prefixed) operators with ambiguous scope. The sec-

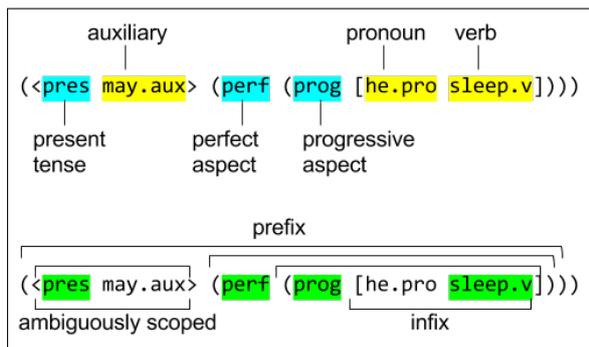


Figure 1: Visualization of ULF syntax for example sentence *He may have been sleeping*. Yellow shows atoms that are represent lexical entries, blue shows special EL operators, and green shows atoms that are acting as the operator in their clauses.

ond visualization in Figure 1 shows a labeling of this. Note that operators can themselves be complex expressions (e.g., `<pres may.aux>`).

5 Annotating Intension and Attitude in ULF

Annotation of modifiers in ULF requires distinguishing predicate modifiers from sentence modifiers, since these have different semantic types. If a modal auxiliary or modal adverb(ial) modifies a sentence without affecting what the sentence predicate says about the subject (e.g., *A major earthquake {may, could} occur; {perhaps, surprisingly, in my opinion} there is no life on Mars*), then it is a sentence-level modifier. If instead the modal auxiliary or adverb(ial) alters the property attributed to the subject, then it must be a predicate-level modifier (e.g., *The cadet must (i.e., is obligated to) obey; the skater {nearly, awkwardly} fell*).

This distinction can be quite subtle since it is dependent on both the lexical entry and the syntax. Consider the following sentences:

- (a) “Mary confidently spoke up”
- (b) “Mary undoubtedly spoke up”
- (c) “Koko is surprisingly intelligent”
- (d) “Surprisingly, Koko is intelligent”

In sentence (a) *confidently* is a predicate modifier whereas in sentence (b) *undoubtedly* is a sentence modifier. Clearly, this is entirely determined by the lexical entry since the syntax trees of the

two sentences are identical. Then compare sentence (c) and sentence (d). The only difference between them is the placement of the modifier *surprisingly*, which changes its semantic type.

Annotating attitudes merely requires recognizing when a sentence functions as a propositional argument (rather than, for instance, as an adverbial or relative clause), and using reifying operator *that* accordingly. The operator must be used even if *that* is elided in the surface text: *I’m sure (that) you’ve heard of him*. Since attitude predicates have the same type structure as extensional predicates, no additional annotation is necessary for ULF.¹

6 Annotating Aspect and Tense

Aspect is generally captured by the lexical entries in our annotations (e.g., *daily, used to*). However, we introduce *perf* and *prog* as operators for perfect and progressive aspect, since they are generated morpho-syntactically in English, via the auxiliaries *have* and *be* respectively. Semantically, aspect describes the way an event relates to time, so they are sentence modifiers in EL.

EL has two operators for tense – *past* and *pres* – for past and present. We treat the English modal auxiliary *will* as a present-tense verb operating at the sentence level and meaning *at a time after now*.² We regard tense as an unscoped operator in ULF (to be “raised” to the sentence level), and consequently it is simply annotated as operating on the verb that bears the tense inflection (this is always the first verb – the head verb – of a tensed verb phrase in English). Some examples:

- (a) “He is sleeping”
`(<pres prog> [he.pro sleep.v])`
- (b) “He has left Rome”
`(<pres perf> [he.pro (leave.v Rome.c)])`
- (c) “He had left Rome”
`(<past perf> [he.pro (leave.v Rome.c)])`
- (d) “He has been sleeping”
`(<pres perf> (prog [he.pro sleep.v]))`

¹Some clauses used as arguments denote episode types, e.g., *For Mary to be late is unusual*; we distinguish such cases but omit details here.

²Formal details of the treatment of tense and temporal adverbials in EL are given in (Hwang and Schubert, 1994).

(e) “He may have been sleeping”

```
(<pres may.aux>
  (perf (prog [he.pro sleep.v])))
```

Sentence (a) is a simple sentence where the tense is determined by the verb. Sentence (b), (c), and (d) show how *had* and *has* determine the tense of the sentence. Note that in all three cases the perfect auxiliary is followed by the past participle form of the verb. This is simply a syntactic requirement in English. Sentence (e) shows an example where the modal auxiliary determines the tense.

7 Remarks on Strategy

We have experimented with annotating randomly chosen examples from parsed and unparsed corpora such as the Brown corpus, the Gigaword newswire corpus and *The Little Prince*. This experimentation has led to an annotation strategy that starts with phrasal bracketing, followed by adding parts of speech (with manual correction of automatic tagging errors), followed by substituting type-suffixed lexical interpretations for words, followed by addition of any tacit reification and type-shifting operators. Here is a simple example:

```
(Mary (confidently (spoke up)) →
(Mary.nnp (confidently.rb
  (spoke.vbd up.prt))) →
[Mary.prp (confidently.adv-a
  (<past speak.up.v>)]).
```

Replacement of *confidently.adv-a* by *undoubtedly.adv-s* would cause subsequent automatic “raising” of the adverb to the sentence level.

Development of annotator tools, such as a possible role supplier for common words and access to the extant semantic parser, as well as evaluation of the described annotation strategy are underway. In parallel, ULF annotation methods of more linguistic phenomena are being developed. For these reasons, the annotation guidelines will not yet be publicly released. Also, since the phenomena described in this document cannot be annotated in isolation in our framework, there are no semantic category-specific preliminary annotations to report.

We expect the annotation effort to be successful because ULF is syntactically close to surface English and the annotator tools under development will simplify the annotation task. Similarly,

we expect machine translation methods such as Synchronous Tree Substitution Grammars (Eisner, 2003; Gildea, 2003) to be successful in automating this annotation because of the close syntactic correspondence to the surface form.

8 Generalization to Other Languages

In view of its English-like syntax, our annotation scheme it will not map directly to other languages. For example, Mandarin does not have grammatical tense markers, relying on lexical operators instead. This is in clear contrast with how our annotation schema marks tense on the verb. Of course, languages also differ in their vocabulary and surface operator-operand structure. Thus our corpus will not be cross-lingual.

However, the superficial tense operators of ULF are reduced to more fundamental constructs (predications about episodes) by deindexing, and in general conversion from ULF to ELF yields representations intended to be language-independent in terms of semantic types. The expressive devices employed in those representations, such as event reference, general quantification, reification, and modification are shared by all languages. Generalizing our work to other languages will require developing a ULF for the target language, close to its surface form, and methods of converting the ULF to ELF (in context). This is not a trivial task, but the resulting formulas will be type-coherent and capable of supporting inference.

9 Related Work

Previous efforts have been made toward training a transducer for broad coverage meaning representation of sentences, perhaps most prominently OntoNotes (Hovy et al., 2006) and AMR (Banarescu et al., 2013). These representations employed PropBank, WordNet, VerbNet, and FrameNet as semantic resources, but were not designed to be formally interpretable. Semantic types of nodes are not defined, there is no distinction between extension and intension (or between what is real and what is hypothetical), and thus there is no clear basis for inference. The representations also set aside some important linguistic phenomena, such as tense (hence, how events are temporally linked); and quantifiers are added in modifier-like fashion, much as if they were attributes of entities. DeepBank is a corpus of annotations in English Resource Semantics (ERS),

which is a canonicalized and grammar-constrained semantic representation (Flickinger et al., 2012). ERS handles a wide-array of linguistic phenomena, while allowing semantic underspecification by using minimal recursion semantics as its met-alanguage representation (Copestake et al., 2005). Although ERS is highly descriptive, it lacks machinery for generating general inferences from fully-resolved formulas.

10 Conclusions and Future Work

We have described how some semantically significant, often neglected phenomena of natural language can be captured in Episodic Logic. We outlined some requirements and methods for annotating a topically broad corpus with unscoped versions of EL, to be used as a basis for training a high-fidelity semantic parser for English. Because EL (and even more so, ULF) is close in form to the surface text, use of machine translation techniques should yield good performance for such a machine learning task. As noted earlier, we believe that a divide-and-conquer approach to resolving various sorts of residual indeterminacy in ULFs is likely to achieve better results than a fell-swoop approach, particularly since we are tackling a range of subtle semantic phenomena beyond those ordinarily considered. High-fidelity interpretations of NL into EL would greatly facilitate many NL applications, including knowledge extraction from lexical and encyclopedic sources, as well as text and dialogue understanding tasks.

11 Acknowledgements

This work was supported by a Sproull Graduate Fellowship from the University of Rochester, DARPA CwC subcontract W911NF-15-1-0542, and NSF EAGER grant IIS-1543758. We are also grateful to the anonymous reviewers for their comments.

References

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria, August. Association for Computational Linguistics.

Emily M. Bender, Dan Flickinger, Stephan Oepen, Woodley Packard, and Ann Copestake. 2015. Layers of interpretation: On grammar and compositionality. In *Proceedings of the 11th International Conference on Computational Semantics*, pages 239–249, London, UK, April. Association for Computational Linguistics.

Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A. Sag. 2005. Minimal Recursion Semantics: An introduction. *Research on Language and Computation*, 3(2):281–332.

Jason Eisner. 2003. Learning non-isomorphic tree mappings for machine translation. In *The Companion Volume to the Proceedings of 41st Annual Meeting of the Association for Computational Linguistics*, pages 205–208, Sapporo, Japan, July. Association for Computational Linguistics.

Dan Flickinger, Yi Zhang, and Valia Kordoni. 2012. Deepbank. a dynamically annotated treebank of the Wall Street Journal. In *Proceedings of the 11th International Workshop on Treebanks and Linguistic Theories*, pages 85–96.

Daniel Gildea. 2003. Loosely tree-based alignment for machine translation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 80–87, Sapporo, Japan, July. Association for Computational Linguistics.

Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. Ontonotes: The 90% solution. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA, June. Association for Computational Linguistics.

Chung Hee Hwang and Lenhart K. Schubert. 1994. Interpreting tense, aspect and time adverbials: A compositional, unified approach. In *Proceedings of the First International Conference on Temporal Logic, ICTL '94*, pages 238–264, London, UK, UK. Springer-Verlag.

Gene Kim and Lenhart Schubert. 2016. High-fidelity lexical axiom construction from verb glosses. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 34–44, Berlin, Germany, August. Association for Computational Linguistics.

Fabrizio Morbini and Lenhart Schubert. 2009. Evaluation of Epilog: A reasoner for Episodic Logic. In *Proceedings of the Ninth International Symposium on Logical Formalizations of Commonsense Reasoning*, Toronto, Canada, June.

Lenhart K. Schubert and Chung Hee Hwang. 2000. Episodic logic meets little red riding hood: A comprehensive natural representation for language understanding. In Lucja M. Iwańska and Stuart C. Shapiro, editors, *Natural Language Processing and*

Knowledge Representation, pages 111–174. MIT Press, Cambridge, MA, USA.

Lenhart K. Schubert. 2000. The situations we talk about. In Jack Minker, editor, *Logic-based Artificial Intelligence*, pages 407–439. Kluwer Academic Publishers, Norwell, MA, USA.

Lenhart Schubert. 2014. From treebank parses to Episodic Logic and commonsense inference. In *Proceedings of the ACL 2014 Workshop on Semantic Parsing*, pages 55–60, Baltimore, MD, June. Association for Computational Linguistics.