Bridging Text and Knowledge with Frames

Srini Narayanan
Google Zurich / Brandschenkestrasse 110, 8002 Zurich, Switzerland
snarayanan0@gmail.com

Abstract

FrameNet is the best currently operational version of Chuck Fillmore's Frame Semantics. As FrameNet has evolved over the years, we have been building a series of increasingly ambitious prototype systems that exploit FrameNet as a semantic resource. Results from this work point to frames as a natural representation for applications that require linking textual meaning to world knowledge.

1 Introduction

Frame Semantics (Fillmore, 1976) defines the meaning of a word with respect to the conceptual structure (Frame) that it evokes. The promise of Frame Semantics is that it is a principled method to connect language analysis with concepts and knowledge. This paper summarizes over a decade of research at Berkeley[1] on computational models bridging text and inference using Frame Semantics. We will start with a brief description of the lexical resource, FrameNet[2], designed with the explicit goal to capturing insights and findings from Frame Semantics in an on-line lexicon. We then describe computational models that exploit the semantic information in FrameNet for a variety of NLP tasks.

2 FrameNet

The Berkeley FrameNet project (Fillmore, Johnson, & Petruck, 2003) is building a lexicon based on the theory of Frame Semantics. In FrameNet, the meanings of lexical items (lexical units (LU)) are defined with respect to larger structured representations, called Frames. Individual lexical units evoke specific frames and establish a binding pattern to specific slots or roles (frame elements (FE)) within the frame. FrameNet describes the underlying frames for different lexical units, examines sentences related to the frames using a very large corpus, and records (annotates) the ways in which information from the associated frames are expressed in these sentences. The result is a database that contains a set of frames (related through hierarchy and composition), a set of frame elements for each frame, and a set of frame annotated sentences that covers the different patterns of usage for lexical units in the frame. Collin Baker's paper in this conference has more details on the FrameNet project including the current state of the resource which is now available in multiple languages. This paper summarizes the results of applying FrameNet in a variety of NLP applications.

2.1 FrameNet data as seed patterns for Information Extraction

While FrameNet frames and FE tags are meaningful to human interpreters, they are not suitable for direct use in NLP applications. One early project explored using the FrameNet annotated dataset to automatically compile patterns and a lexicon for Information Extraction (IE) (Mohit and Narayanan, 2003). A distinguishing feature that made FrameNet attractive for this purpose was its explicit mandate to cover all the valence patterns for a target word, not just the frequent ones. Thus, FrameNet annotations and valence alternations were designed to capture the long tail for every target lexeme. We hypothesized that using a highly precise set of patterns and a lexicon automatically compiled from the FrameNet frame relational database and annotations should result good performance for the task. To increase coverage, we extended the frame lexicon with WordNet synsets. As a first test, we culled a set of news stories from Yahoo News Service with topics re-

[1] All the work described was done while the author was at the University of California, Berkeley and the International Computer Science Institute (ICSI) / 1947 Center Street, Berkeley CA 94704.

[2] http://framenet.icsi.berkeley.edu

Proceedings of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929–2014), pages 22–25,
Baltimore, Maryland USA, June 27, 2014.

lated to the topic of *crime*. We also compiled a set of IE patterns and lexicon from several crime related frames (such as Arrest, Detain, Arraign and Verdict.) We were able to achieve an average precision of 76.5% and an average recall to 66% for the stories in this domain. However, the relatively sparse and uneven domain coverage of FrameNet and the absence of high quality parsers and named entity annotators (used for building expressive and general patterns) at the time made the pilot task difficult to repeat in an open domain setting. While the coverage of FrameNet is still an issue, the enormous gains made in the quality and amount of parsed and named entity annotated data could make this early work attractive again where FrameNet can be used as a high precision seed pattern generator in a semi-supervised IE setting.

3 From Frames to Inference

A fundamental aspect of Frame Semantics, one that directly connected the linguistic insights of Chuck Fillmore to the early work in AI by Schank, Abelson, Minsky, and others was the idea that Frames were central to how inferences were packaged. In this view, framing provided preferential access to specific *expected* inferences. These inferences were said to be *in the frame*. Schankian scripts (such as the famous restaurant script) (Schank and Abelson, 1977) are a good example of such inferential packaging in terms of expected sequences of events, participants, and outcomes. In addition to providing such general inferences, Chuck Fillmore observed that *linguistic framing* also provided a way to delineate multiple perspectives on an event (including foregrounding, backgrounding, and participant perspective). An example can be found in the perspective difference provided by the lexical items *sell*, *buy*, or *pay*, which all evoke the commercial transaction frame.

(Chang, Narayanan, & Petruck, 2002), built a computational formalism that captured structural frame relationships among participants in a dynamic scenario. This representation was used to describe the internal structure and relationships between FrameNet frames in terms of parameters for active event simulations for inference. We applied our formalism to the *commerce domain* and showed how it provides a flexible means of handling linguistic perspective and other challenges of semantic representation. While this work was able to computationally model subtle inferential effects in perspective and framing choice, it remains a proof of concept demonstration and there was a need to do an automatic translation to an inference formalism which would enable us to use more robust reasoners (the trade-off was of course that these off the shelf reasoners produced shallower inferences).

(Scheffczyk, Baker, & Narayanan, 2010) automatically translated a crucial portion of FrameNet to the description logic based web ontology language OWL, and showed how state-of-the-art description logic reasoners can make inferences over FrameNet-annotated sentences. Thus, annotated text becomes available to the Semantic Web and FrameNet itself can be linked to other ontologies. While our OWL translation is limited to facts included in FrameNet, links to ontologies make world knowledge available to reasoning about natural language. Therefore, are linked FrameNet to the Suggested Upper Merged Ontology (SUMO). This ground work gives a clear motivation for the design of further ontology bindings and defines the baseline for measuring their benefits.

Fillmore's further insights into the connections between textual inference and world knowledge led us to ask the question of how a linguistic analysis of a written document can contribute to identifying, tracking and populating the eventualities that are presented in the document, either directly or indirectly, and representing degrees of belief concerning them. This work, reported in (Fillmore, Narayanan, & Baker, 2006), attempts to clarify the boundary between on the one hand the information that can be derived on the basis of linguistic knowledge alone (composed of lexical meanings and the meanings of grammatical constructions) and on the other hand, reasoning based on beliefs about the source of a document, world knowledge, and common sense. In particular, we show that the kind of information produced by FrameNet can have a special role in contributing to text understanding, starting from the basic facts of the combinatorial properties of frame-bearing words (verbs, nouns, adjectives and prepositions) and arriving at the means of recognizing the anaphoric properties of specific unexpressed event participants. Framenet defines a new layer of anaphora resolution and text cohesion based on the annotations of the different types of null instantiations (Definite Null Instantiation (DNI), Indefinite Null Instantiation (INI), and Constructional Null

Instantiation (CNI)). A full exploitation of these linguistic signals in a coreference resolver is still pending.

4 Frame Semantics in Question Answering

As FrameNet matured, we started asking if it could be used for semantically based question answering for questions that went beyond factoids and required deeper semantic information. (Narayanan and Harabagiu, 2004; Sinha and Narayanan, 2005; Sinha, 2008) report on a prototype question answering system that attempted to answer questions related to causality, event structure, and temporality in specific domains. The project on Question Answering (QA) was a joint effort with Sanda Harabagiu's group at UT Dallas.

The QA work was based on the fact that events, while independent of language themselves, are frequently discussed in natural language, yielding copious data in that form. To reason about complex events requires an interface from event models to data sources. We sought to exploit semantic frames as an intermediate structure and interface between event descriptions in natural language and event models that produce inferences to answer questions. In the course of this project, we came up with the basic framework and algorithms combining a variety of NLP techniques including Parsing, Topic Modeling, Named Entity Recognition, and Semantic Role Labeling with deep event structure inference in multiple domains. The frame structure in language provides a bi-directional mapping from language to event models, enabling us to link information found in text about an event of interest to models that represent that event. The proof of concept system used frame parsed input with a set of hand built domain ontologies for specific domains. The system was able to answer domain questions involving causal, diagnostic, and hypothetical reasoning. While the results clearly showed the utility of FrameNet as a resource supporting deep semantic inference, it also delineated the necessity and role of domain specific ontologies and inference required to realize an end-to-end system using FrameNet.

5 Frames, Constructions and Grammar

Yet another of Fillmore's profound insights was the observation that every unit of grammar is most effectively described as a pairing between form and meaning, aka a construction. Constructions exist at lexical (and sub-lexical) levels as well as at larger granularities (phrases, discourse) playing a crucial role in the compositionally of language. This proposal, entitled construction grammar, has gained considerable empirical support in large part due to the investigations of Fillmore, his colleagues and students, reported in a series of beautiful papers on the grammatical and compositional properties of constructions.[3]

Research on construction grammar has played a fundamental role within our Berkeley interdisciplinary project, NTL[4], which is attempting to build cognitively plausible computational models of language acquisition and use. Specifically, NTL research has resulted in the grammar formalism called Embodied Construction Grammar (ECG), where the meaning pole of a construction is expressed in terms of bindings between bodily schemas (also called Image Schemas) and frames. ECG allows constraints of all kinds (phonological, syntactic, semantic, etc.) in a unification based probabilistic framework, where the best fitting interpretation in context is selected as the analysis of the input. ECG is formally defined and computationally implemented, and has been used for linguistic analysis, in models of language comprehension, and for cognitive models of language learning[5].

6 Frame Semantics and Metaphor

FrameNet has long held the goal of including information about metaphorical usage in language. The most recent project on Frame Semantics is the ICSI MetaNet project, where the goal is to build a system that extracts linguistic manifestations of metaphor (words and phrases that are based on metaphor) from text and interprets them automatically in four different languages.

The MetaNet project, is a multi-lingual, multi-university, multi-disciplinary effort that incorporates FN methodology as well as corpus and machine learning techniques, deep cognitive linguistics, and behavioral and imaging experiments.

MetaNet models metaphor as a mapping between two different frames. Such mappings

[3] `http://www.constructiongrammar.org/` is a currently active resource on the topic with contributions from a variety of international scholars.

[4] `http://ntl.icsi.berkeley.edu/ntl`

[5] `http://ecgweb.pbworks.com/w/page/15044343/FrontPage`

project information from a source frame to a target frame. The information projected is partial and can include the frame, its slots, and filler constraints. An initial repository of mappings that draws on FrameNet frames as sources and targets of the mappings is used as base information by a system that extracts additional metaphors using machine learning. The system uses what it has learned about the relationships between the frame elements of conceptual metaphors to find new metaphors in text. The MetaNet Wiki[6] is a database of such mappings, drawing on FrameNet's inventory of Frames. The mappings currently exist in four different languages. FrameNet frames and mappings constrain the search for new metaphors, and the discovery of new metaphors by a corpus based machine learning algorithm both a) provides empirical support for the existing frames and mappings and b) more importantly potentially extends the set by identifying gaps and inconsistencies in coverage. This interaction facilitates an iterative design process in MetaNet which is empirically driven and usage based, just as Fillmore would have insisted.[7]

7 Conclusion and Future Directions

Frame semantics in general and FrameNet in particular show considerable promise for use in deep semantic analysis. FrameNet frames are intended to capture crucial generalizations not available in other lexical resources. Various prototype systems have clearly demonstrated the potential of FrameNet to make a qualitative difference in semantic NLP. There remain two crucial gaps that have to be bridged. One is the issue of coverage. The second is the lack of a formal representation covering the more subtle inferential aspects of FrameNet. Progress is being made of both fronts as is evidenced in some of the papers in this workshop. If successful, the next few years should see an increasing use and acceptance of FrameNet as a crucial semantic resource bridging language analysis with inference. This will lead to scalable versions of the systems described in this paper, but will also give rise to new applications. One particularly intriguing area of research is the use of frames for cross-modal semantic representation bridging text, speech, and vision.

[6] http://metaphor.icsi.berkeley.edu

[7] Even at 83, Chuck's brilliant attention to detail and infectious enthusiasm fundamentally shaped the early MetaNet project on a day-to-day basis, till his illness sadly made direct participation impossible after 2012.

Acknowledgments

All the projects described here were collaborations with the FrameNet and NTL groups at ICSI and Berkeley. This line of inquiry will continue to be dedicated to and guided by Chuck Fillmore's invaluable insights, vision, and body of work.

References

Nancy Chang, Srini Narayanan, and Miriam R. L. Petruck. 2002. Putting Frames in Perspective. *Proceedings of 19th International Conference on Computational Linguistics. Taipei*: COLING, 2002.

Charles Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, 280, 1,20-32, 1976.

Charles Fillmore, Christopher R. Johnson, and Miriam R. L. Petruck. 2003. Background to FrameNet. *International Journal of Lexicography*. 16.3 (2003): 235-250.

Charles Fillmore, Srini Narayanan, Collin Baker. 2006. What Linguistics can contribute to Event Extraction. AAAI Workshop on Event Extraction, AAAI 2006, Boston.

Behrang Mohit and Srini Narayanan. 2003. *Semantic Extraction from Wide-Coverage Lexical Resources*, HLT-NAACL, Short Papers, 2003.

Srini Narayanan and Sanda Harabagiu 2004. Question Answering based on Semantic Structures. *International Conference on Computational Linguistics* (COLING 2004), Geneva, Switzerland, August 22-29, 2004.

Roger Schank and Robert P. Abelson. 1977 *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*, Erlbaum, 1977. ISBN 0-470-99033-3.

Jan Scheffczyk, Collin F. Baker, Srini Narayanan. 2010. Reasoning over Natural Language Text by Means of FrameNet and Ontologies. *Ontology and the Lexicon: A Natural Language Processing Perspective*. Eds. Chu-ren Huang, et al. Studies in Natural Language Processing. Cambridge, UK: Cambridge University Press, 2010. 53-71.

Steve Sinha and Srini Narayanan. 2005. Model Based Answer Selection. Textual Inference in Question Answering Workshop, AAAI 2005, Pittsburgh, July 2005.

Steven Sinha. 2008. *Answering Questions about Complex Events* PhD Dissertation, Department of EECS, Berkeley, CA, 2008.