

Introduction to Semantic Role Labelling

Behrang QasemiZadeh

zadeh@phil.hhu.de

Computational Linguistics Department, HHU

October 2018–January 2019; Draft .1

Preface

This material provides instructions and a summary of discussions for the Semantic Role Labelling course offered at the Winter Semester of the Philosophy Faculty of Heinrich-Heine-Universität Düsseldorf.

Preface: A simplified motivation/introduction I

Before we start, let's see what are our motivations in simple, nonacademic language, using everyday speech, and simple vocabulary.

Preface: A simplified motivation/introduction I



Preface: A simplified motivation/introduction I



- (1)
- The boy broke the window with a ball.
 - The boy hit the window with a ball.

Preface: A simplified motivation/introduction I



- (2)
- The window broke.
 - *The window hit.**

Preface: A simplified motivation/introduction I

WHY the two verb 'hit' and 'break' (and many other pairs of verbs) show different behavior, and why these differences are the way they are?! (Levin and Hovav, 2005)

And, **How** to assert that below although both 'boy' and 'window' are the subject of the sentence, still they are somehow different?

- (3) a. The boy broke the window
- b. The window broke.

Preface: A simplified motivation/introduction I

Given sentences such as

- (4) a. The window broke.
- b. John broke the window.
- c. The ball broke the window.

How to write a computer program that can, correctly, answer questions such as:

- ▶ **Who** broke the window?

('ergative' verbs are just one example of complexities we may face).

Preface: A simplified motivation/introduction I

In general, how to provide computers with a “semantic representation” of natural language sentences so that they can be used for applications such as *question answering*, *information extraction*, *machine translation* etc.? Or, how to build computational models that are capable of showing a level of “natural language understanding”?

How to help computers to avoid mistakes such as:

- ▶ **USER:** *Siri, I am bleeding, call me an ambulance?*
- ▶ **SIRI:** *From now on, I'll call you 'An Ambulance'.
OK?*

Semantic Role Labelling: A simplified motivation/introduction I

The relationship between verbs and their arguments is a widely debated topic in linguistics:¹

From linguistics perspective: How the meaning/semantics of a verb can determine the morphosyntactic realization of its arguments (e.g., when dativ, when akkusativ)? Contrariwise, from computational perspectives (for practical applications): How to determine the meaning of a verb from the morphosyntactic realization of its arguments?

¹The study is naturally extended to other part-of-speech categories, such as noun, adjectives, etc. but our basic focus will be on event verbs and their nominalizations.

Semantic Role Labeling: A simplified motivation/introduction I

As you could infer, the topic is important from several perspectives:

In Linguistics, for instance, there are several theories that suggests a close relationship between the meaning of a verb and its syntactic structure (remember the break example); perhaps, we could find patterns in syntactic usages of verbs that are specific to (or, determined by) the meaning of the verbs.

In turn, this can be used to bridge between lexical semantics (analysis of word meanings with respect to their senses, implications, presupposition, etc.) and syntax-oriented research in computational linguistics (e.g., parsing).

Semantic Role Labeling: A simplified motivation/introduction II

In a more general way, we can assume that:

1. Language is an instrument that we use to convey information regarding events and situations.
2. These events/situations are the fruit of our conceptualization from things we experience/understand/synthesize/... (or, simply, thinking!).
3. In turn, we assume the existence of a thing called **event structure**, i.e., we hypothesis that events (things that we have conceptualized and articulated using language) are complex entities consist of several elements; and, that there exist relations between these elements.

Semantic Role Labeling: A simplified motivation/introduction III

4. Not only that, we (people doing research on semantic role labelling – people who follow subjects concerning Linguistic/Philosophical Semantics, in a shallow but empirical manner in favour of brute pragmatism) believe that:
 - ▶ event structure and the complex relationships between its parts can be captured and represented in a (semi)-formal manner, e.g., using set theory, with the help of logical statements and predicate calculus, as well as formal grammars designed for this purpose.
 - ▶ **NOTE:** Here the word **semi** is used because we can use statistical/distributional models to capture event structure (i.e., instead of a symbolic formal language, we can use a mathematical model).

Semantic Role Labeling, three slides history! I



Semantic role labelling, as a topic in computational linguistics, owes its existence to the seminal work of Charles Fillmore, **Case Grammar (1968)** and **Frame Semantics (1976)**: simply put, a grammar that Fillmore developed for the analysis the structure of sentences with respect to relationships between "deep case" (read it semantic roles, we will be discussed soon).

Semantic Role Labeling, three slides history! II

With the help of remarkable advances in the beginning of the new millennium, i.e., with respect to the computational power and performance, statistical machine learning, and perhaps most importantly the availability of hand annotated linguistic resources suitable for finding semantic roles (which we fully describe later), Gildea and Jurafsky (2002) published a paper titled "Automatic Labeling of Semantic Roles".

Semantic Role Labeling, three slides history! III

The paper gained an immense amount of attraction right away such that we can consider it as a landmark for (even the birth of) “semantic role labelling” as a topic in computational linguistics. But also, in reality, Gildea and Jurafsky (2002) work was a computerized/commercialized application of Fillmore’s legacies: Case Grammar, Frame Semantics, and FrameNet.

Since then, a lot of work has been done and the field has had an enormous, yet insufficient, progress. In our course, we cover the basics required to go through this large body of research.

Our plan, administrative matters, etc.

Overall Plan I

This is an introductory–multidisciplinary course about **Semantic Role Labelling** and the content is designed with respect to the fact that attendees are coming from different scientific backgrounds and levels (Bachelor or Master degree in CS, CL, Ling., IS.).

The topic of this course is semantic relations and predicate-argument structure for natural language understanding (practical – application oriented).

Overall Plan II

We intend to cover the following major topics:

1. Introduction to Semantic Role and its Key Concept
2. Lexical Resources for Semantic Role Labeling
3. Computational Methods for Semantic Role Labelling

Item 1 includes topics Linguistics. Item 3 embraces topics from Machine Learning. Item 2 stands somewhere in between.

Overall Plan III

Our intended agenda, in more details:

- ▶ Introduction to Semantic Role and its Key Concept
 - ▶ Arguments/Adjuncts
 - ▶ Proto-Roles
 - ▶ Frame Semantics
 - ▶ Levin's Verb Classes and Alternation patterns
 - ▶ Typical types of semantic roles
- ▶ Lexical Resources for Semantic Role Labeling
 - ▶ FrameNet
 - ▶ VerbNet
 - ▶ PropBank
- ▶ Computational Methods of Semantic Role Labelling

Overall Plan IV

- ▶ Using Supervised Machine Learning for Semantic Role Labelling
 - ▶ A Gentle Introduction
 - ▶ Feature Analysis
 - ▶ Off-the-shelf Semantic Role Labelling Systems
 - ▶ Evaluation methods
 - ▶ Evaluation Campaigns
- ▶ Unsupervised Methods
 - ▶ A Gentle Introduction
 - ▶ Evaluation Campaigns

Overall Plan V

During the course, depending on their background and interest, participants can deepen their knowledge regarding theoretical or/and computational aspects of semantic role labelling through practical experiences (annotation and cross-theory comparison of semantic roles, building their own computational model, e.g. training semantic role parsers, etc.).

Goals of the course I

To emphasize on the obvious: The ultimate goal is to **learn**.

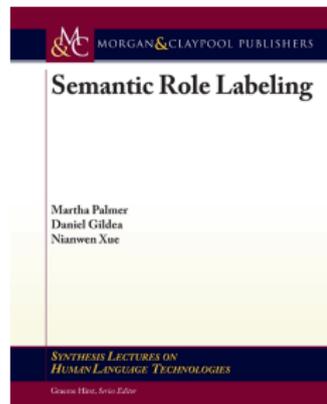
We will try to gain knowledge from two sides of a spectrum: Linguistic Semantics and Machine Learning, for the purpose of Semantic Role Labelling, in which we try to maintain a balance between theory and practice.

- ▶ Introducing the terminology and basic concepts in Semantic Role Labeling.
- ▶ Introducing essential resources and methods available for semantic role labeling.

Goals of the course II

- ▶ We will touch the tip of an iceberg; the aim is to equip you with the basic knowledge to be able to study a lot of work which will be left in this course.

Our main text book:



Palmer, M., Gildea, D., and Xue, N. (2010).

Semantic Role Labeling.

Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers

<https://doi.org/10.2200/S00239ED1V01Y200912HLT006>

Additional references will be provided for each of the topics.

Homework and grading I

Participants are asked to do homework every other week; homeworks are essential for the **BN**.

A wide range of topics will be available as homework: critical essays, annotation practices, etc. In some cases, the result must be presented and discussed with other attendees. Although attending the course is not mandatory, it may be required for choosing a homework and to deliver it.

For **AP**, there will be the usual written exam from the content of the book and what we discuss in the class, however, you are welcome to do projects (for both Linguistics enthusiasts and Computer coding enthusiasts).

I kindly request you ... I

If I am not clear, e.g., you feel that I keep things behind a door, please let me know.

Bibliography I

- Gildea, D. and Jurafsky, D. (2002). Automatic labeling of semantic roles. *Comput. Linguist.*, 28(3):245–288.
- Levin, B. and Hovav, M. R. (2005). *Argument realization*. Cambridge University Press.
- Palmer, M., Gildea, D., and Xue, N. (2010). *Semantic Role Labeling*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.