

This lecture introduces several natural language inference problems and discusses them briefly in order to introduce the general topic we will study in this course. We then present the topics that will be covered in the course and the plan for the class and final project.

1 Motivation

The Spelling poem and some more specific disambiguation problems exemplified provide a preliminary “feeling” to what we all know. Attempts to write programs that solve natural language problems even at a fairly low and well defined manner (before we even discuss what does it mean to “understand” the spelling poem) need to address a wide variety of questions. There is no question today that a Learning need to have a principle role in any attempt to make progress on these questions.

Learning has multiple purposes in this process, from knowledge acquisition and its role in the integration of various knowledge sources to ensure robust behavior.

Why Study Natural Language: Natural Language is the main channel of Communication and the main channel through which we acquire knowledge. Both these aspects are important both for the study of language from a cognitive point of view and from an engineering perspective.

A grand application to keep in mind that embodied all aspects and difficulties of natural language is that of human computer interaction. In particular, the problems of language understanding and generation as well as knowledge acquisition can be studied from the HCI point of view. Specific tasks could include: question answering, summarization, translation.

Learning Perspective: Why study learning in *natural language*? Natural Language is a great example and for studying and developing theories in learning. Any non trivial task in natural language involves the study of both learning and inference. The tasks involved, e.g., language comprehension, are large scale phenomena in terms of both knowledge and computation. Making progress in this area requires an integrated approach; there is a need to solve many problems in learning, inference and knowledge representation. And, it is not possible to “cheat”: no toy problems are of any practical interest, so to be relevant we need to address real issues.

Natural Language Perspective: Why study *learning* in NLP? There isn’t any significant aspect of language that can be studied without giving learning a principle role. As a consequence, the role of learning in the study of natural language is significant both from the engineering point of view as it is from the cognitive point of view. Learning is important in language acquisition, language change, language variation, language generation and language comprehension.

2 This course:

This course is not on statistical NLP (the more popular name of the topics we discuss in this class). We are going to discuss and gain a lot from the fact that language satisfies statistical regularities; but, the goal is not to explore these statistical regularities (not only, at least) but rather to gain from their existence in order to make progress towards language understanding. In doing that, *knowledge representation* is as important. We therefore prefer to call it *learning* rather than statistics.

There are many topics to cover in a course of this type. It is a very active field of research, many ideas are floating around, most of them will not stay. I will attempt to cover some of the main ideas and techniques. In particular, we will make an attempt to abstract away from specific works and see the main paradigms. In addition, we will exemplify the techniques with some real problems. The plan is to choose problems that will lead to the final project.

3 Handouts:

- Tentative Plan for the course.
- Students survey
- Steve Abney's paper on Statistics and Linguistics.
- Deep Read Paper (Mitre group) as background for the final project.