

Markov Logic in Natural Language Processing: Theory, Algorithms, and Applications

Hoifung Poon, University of Washington

Natural languages are characterized by rich relational structures and tight integration with world knowledge. As the field of NLP/CL moves towards more complex and challenging tasks, there has been increasing interest in applying joint inference to leverage such relations and prior knowledge. Recent work in statistical relational learning (a.k.a. structured prediction) has shown that joint inference can not only substantially improve predictive accuracy, but also enable effective learning with little or no labeled information. Markov logic is the unifying framework for statistical relational learning, and has spawned a series of successful NLP applications, ranging from information extraction to unsupervised semantic parsing. In this tutorial, I will introduce Markov logic to the NLP community and survey existing NLP applications. The target audience of the tutorial is all NLP researchers, students and practitioners. The audience will gain the ability to efficiently develop state-of-the-art solutions to NLP problems using Markov logic and the Alchemy open-source software.

1. Structure

The tutorial will be structured as follows:

Part One: Markov Logic

In the first part I will motivate statistical relational learning (SRL) for NLP problems, and introduce Markov logic as the unifying framework. I will present state-of-the-art learning and inference algorithms in Markov logic, and give an overview of the Alchemy open-source software package. The duration of this part will be approximately one hour and half.

Part Two: NLP Applications: Supervised Learning

In the second part I will describe how to use Markov logic and Alchemy to develop state-of-the-art solutions very efficiently for a variety of NLP problems, including: logistic regression, text and hypertext classification, vector-space and link-based information retrieval, entity resolution, information integration, hidden Markov models, Bayesian networks, information extraction, semantic role labeling, and biomedical text mining. This part will also cover practical tips on using Markov logic and Alchemy — the kind of information that is rarely found in research papers, but is key to

developing successful applications. This part will focus on supervised learning and the duration will be approximately an hour.

Part Three: NLP Applications: Unsupervised Learning

In the third and final part I will introduce the emerging direction for statistical relation learning that leverages prior knowledge and relational structures to enable effective learning with little or no labeled data. As examples I will present recent work in applying Markov logic to unsupervised coreference resolution and unsupervised semantic parsing. I will also briefly touch on the exciting prospect of machine reading from the Web. The duration will be about half an hour.

2. Instructor

Hoifung Poon

Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195, USA
hoifung@cs.washington.edu

Hoifung Poon is a fifth-year Ph.D. student at the University of Washington, working with Pedro Domingos. His main research interest lies in advancing machine learning methods to handle both complexity and uncertainty, and in applying them to solving challenging NLP problems with little or no labeled data. His most recent work developed unsupervised learning methods for a number of NLP problems ranging from morphological segmentation to semantic parsing, and received the Best Paper Awards in NAACL-09 and EMNLP-09. At Washington, he helped design course materials for the first offering of the undergraduate machine learning course, and gave guest lectures in both undergraduate and graduate machine learning classes. His prior experience includes teaching undergraduate math classes in West Virginia University, for which he was awarded the Outstanding Graduate Teaching Assistant by the University.