

Hands-on Learning to Search for Structured Prediction

Hal Daumé III¹, John Langford², Kai-Wei Chang³, He He¹, Sudha Rao¹

¹ University of Maryland, College Park

² Microsoft Research, New York

³ University of Illinois, Urbana-Champaign

me@hal3.name jl@hunch.net kchang10@illinois.edu
hhe@cs.umd.edu raosudha@cs.umd.edu

1 Introduction

Many problems in natural language processing involve building outputs that are structured. The predominant approach to structured prediction is “global models” (such as conditional random fields), which have the advantage of clean underlying semantics at the cost of computational burdens and extreme difficulty in implementation. An alternative strategy is the “learning to search” (L2S) paradigm, in which the structured prediction task is cast as a sequential decision making process.

One can then devise training-time algorithms that learn to make near optimal collective decisions. This paradigm has been gaining increasing traction over the past five years: most notably in dependency parsing (e.g., MaltParser, ClearNLP, etc.), but also much more broadly in less “sequential” tasks like entity/relation classification and even graph prediction problems found in social network analysis and computer vision.

This tutorial has precisely one goal: an attendee should leave the tutorial with hands on experience writing small programs to perform structured prediction for a variety of tasks, like sequence labeling, dependency parsing and, time-permitting, more.

2 Format

This tutorial is unique (to our knowledge) among ACL tutorials in this regard: half of the time spent will be in the style of a “flipped classroom” in which attendees get hands on experience writing structured predictors on their own or in small groups. All course materials (software, exercises, hints, solu-

tions, etc., will be made available at prior to the event so that students can download the required data ahead of time; we will also bring copies on USB in case there is a problem with the internet).

3 Contents

The first half of the tutorial will be mostly “lecture” style, in which we will cover the basics of how learning to search works for structured prediction. The goal is to provide enough background information that students can understand how to write and debug their own predictors, but the emphasis will not be on how to build new machine learning algorithms. This will also include a brief tutorial on the basics of Vowpal Wabbit, to the extent necessary to understand its structured prediction interface. The second half of the tutorial will focus on hands-on exploration of structured prediction using the Vowpal Wabbit python “learning to search” interface; a preliminary python notebook explaining the interface can be viewed at <http://tinyurl.com/pyvwsearch>; an elaborated version of this notebook will serve as the backbone for the “hands on” part of the tutorial, paired with exercises.