

From Structured Prediction to Inverse Reinforcement Learning

Hal Daumé III

School of Computing, University of Utah
and UMIACS, University of Maryland
me@hal3.name

1 Introduction

Machine learning is all about making predictions; language is full of complex rich structure. Structured prediction marries these two. However, structured prediction isn't always enough: sometimes the world throws even more complex data at us, and we need reinforcement learning techniques. This tutorial is all about the *how* and the *why* of structured prediction and inverse reinforcement learning (aka inverse optimal control): participants should walk away comfortable that they could implement many structured prediction and IRL algorithms, and have a sense of which ones might work for which problems.

2 Content Overview

The first half of the tutorial will cover the “basics” of structured prediction: the structured perceptron and Magerman's incremental parsing algorithm. It will then build up to more advanced algorithms that are shockingly reminiscent of these simple approaches: maximum margin techniques and search-based structured prediction.

The second half of the tutorial will ask the question: what happens when our standard assumptions about our data are violated? This is what leads us into the world of reinforcement learning (the basics of which we'll cover) and then to inverse reinforcement learning and inverse optimal control.

Throughout the tutorial, we will see examples ranging from simple (part of speech tagging, named entity recognition, etc.) through complex (parsing, machine translation).

The tutorial does not assume attendees know anything about structured prediction or reinforcement learning (though it will hopefully be interesting even to those who know some!), but *does* assume some knowledge of simple machine learning (eg., binary classification).

3 Tutorial Outline

Part I: Structured prediction

- What is structured prediction?
- Refresher on binary classification
 - What does it mean to learn?
 - Linear models for classification
 - Batch versus stochastic optimization
- From perceptron to structured perceptron
 - Linear models for structured prediction
 - The “argmax” problem
 - From perceptron to margins
- Search-based structured prediction
 - Training classifiers to make parsing decisions
 - Search and generalizations

Part II: Inverse reinforcement learning

- Refresher on reinforcement learning
 - Markov decision processes
 - Q learning
- Inverse optimal control and A* search
 - Maximum margin planning
 - Learning to search
- Apprenticeship learning
- Open problems

References

See <http://www.cs.utah.edu/~suresh/mediawiki/index.php/MLRG/spring10>.