10-418 / 10-618 Machine Learning for Structured Data

Machine Learning Department
School of Computer Science
Carnegie Mellon University

Structured Prediction as Search

Matt Gormley
Lecture 3
Sep. 4, 2019

Q&A

Q: Can we have the handwritten notes from lectures?

A: Okay fine…

https://1drv.ms/u/s!Aqk9RupCw3gqhUENhKlkfjqrG5Tv?e=3wGw1o

…but just be warned that lots of education research suggests that taking your own notes is the best way to learn!

Reminders

- Recitation: PyTorch
  - Fri, Sep. 06, same time/location as lecture
- Homework 1: DAgger for seq2seq
  - Out: Mon, Sep. 09 (+/- a day)
  - Due: Mon, Sep. 23 at 11:59pm

LEARNING TO SEARCH

Learning to Search

Whiteboard:

- Problem Setting
- Ex: POS Tagging
- Other Solutions:
  - Completely Independent Predictions
  - Sharing Parameters / Multi-task Learning
  - Graphical Models
- Today's Solution: Structured Prediction to Search
  - Search spaces
  - Cost functions
  - Policies

Learning to Search

***Whiteboard*:**

- Scoring functions for "Learning to Search"
- Learning to Search: a meta-algorithm
- Algorithm #1: Traditional Supervised Imitation Learning
- Algorithm #2: DAgger

LEARNING TO SEARCH: EMPIRICAL RESULTS

Dagger for Mario Tux Cart

Video from Stéphane Ross (https://www.youtube.com/watch?v=V0onpNnWzSU)

Experiments: Vowpal Wabbit L2S

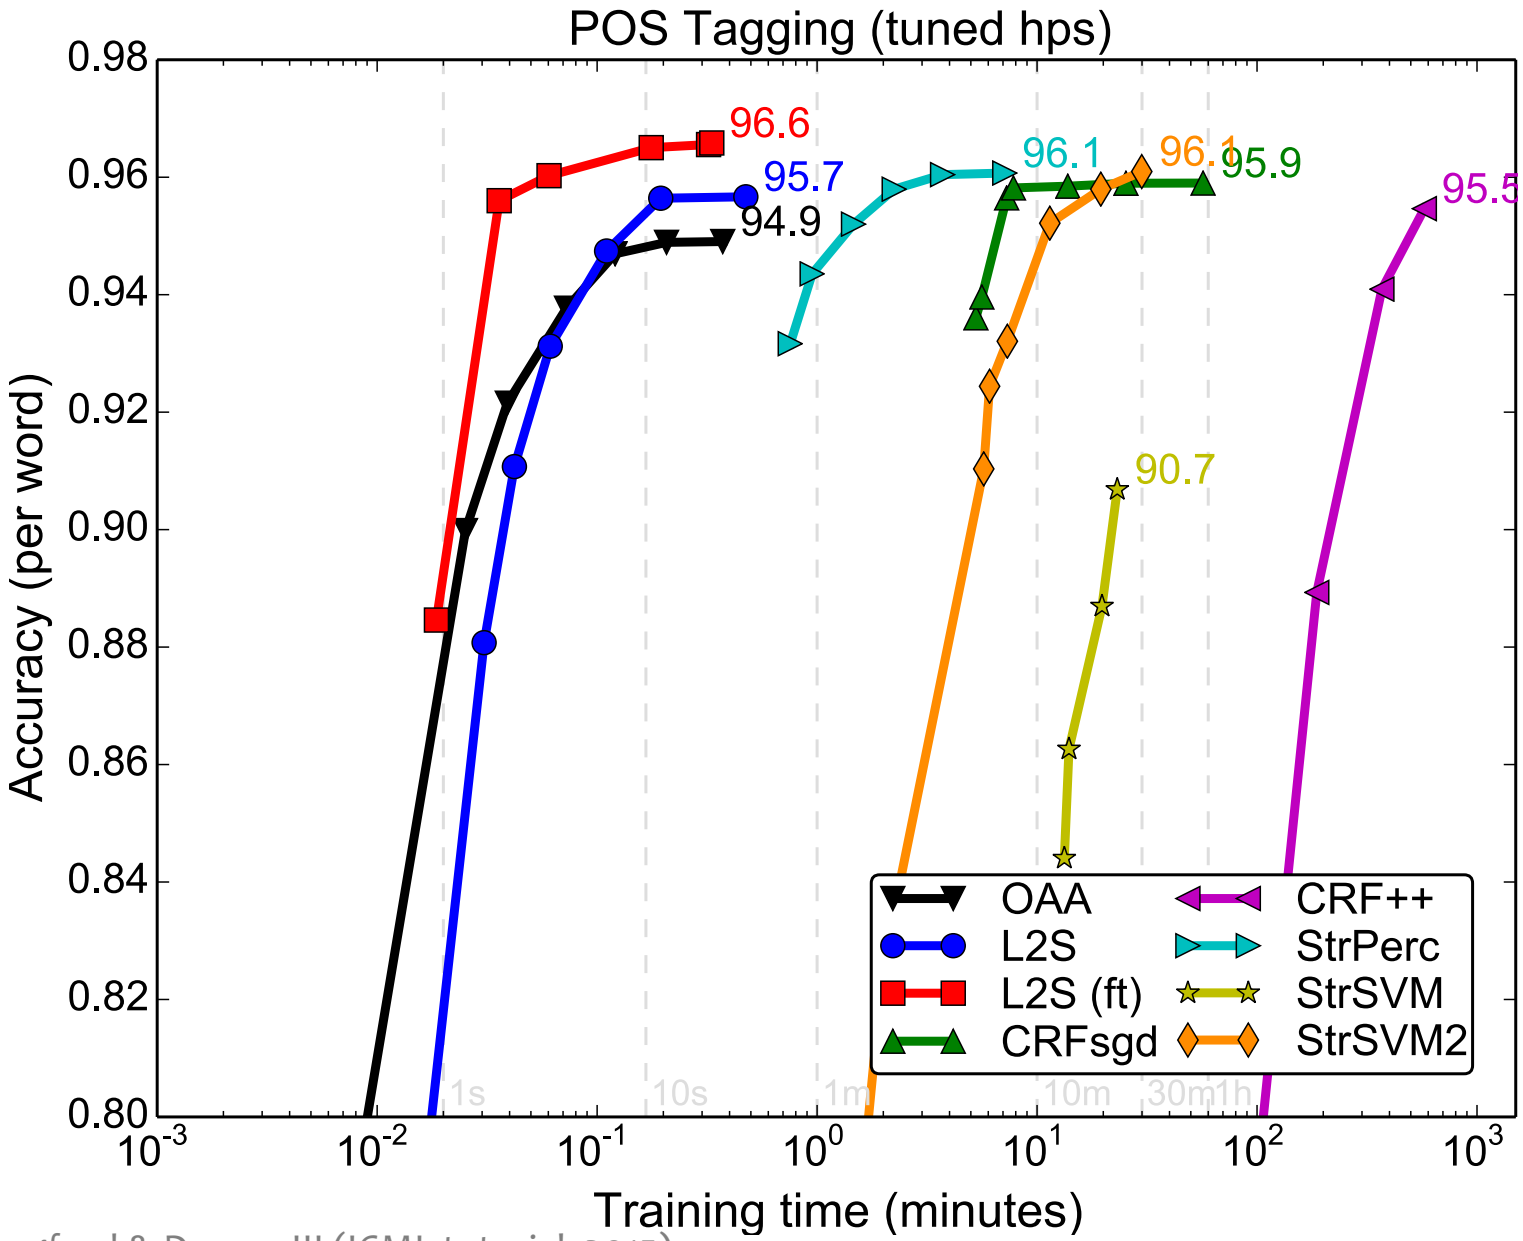

Figure from Langford & Daume III (ICML tutorial, 2015)

Experiments: Vowpal Wabbit L2S

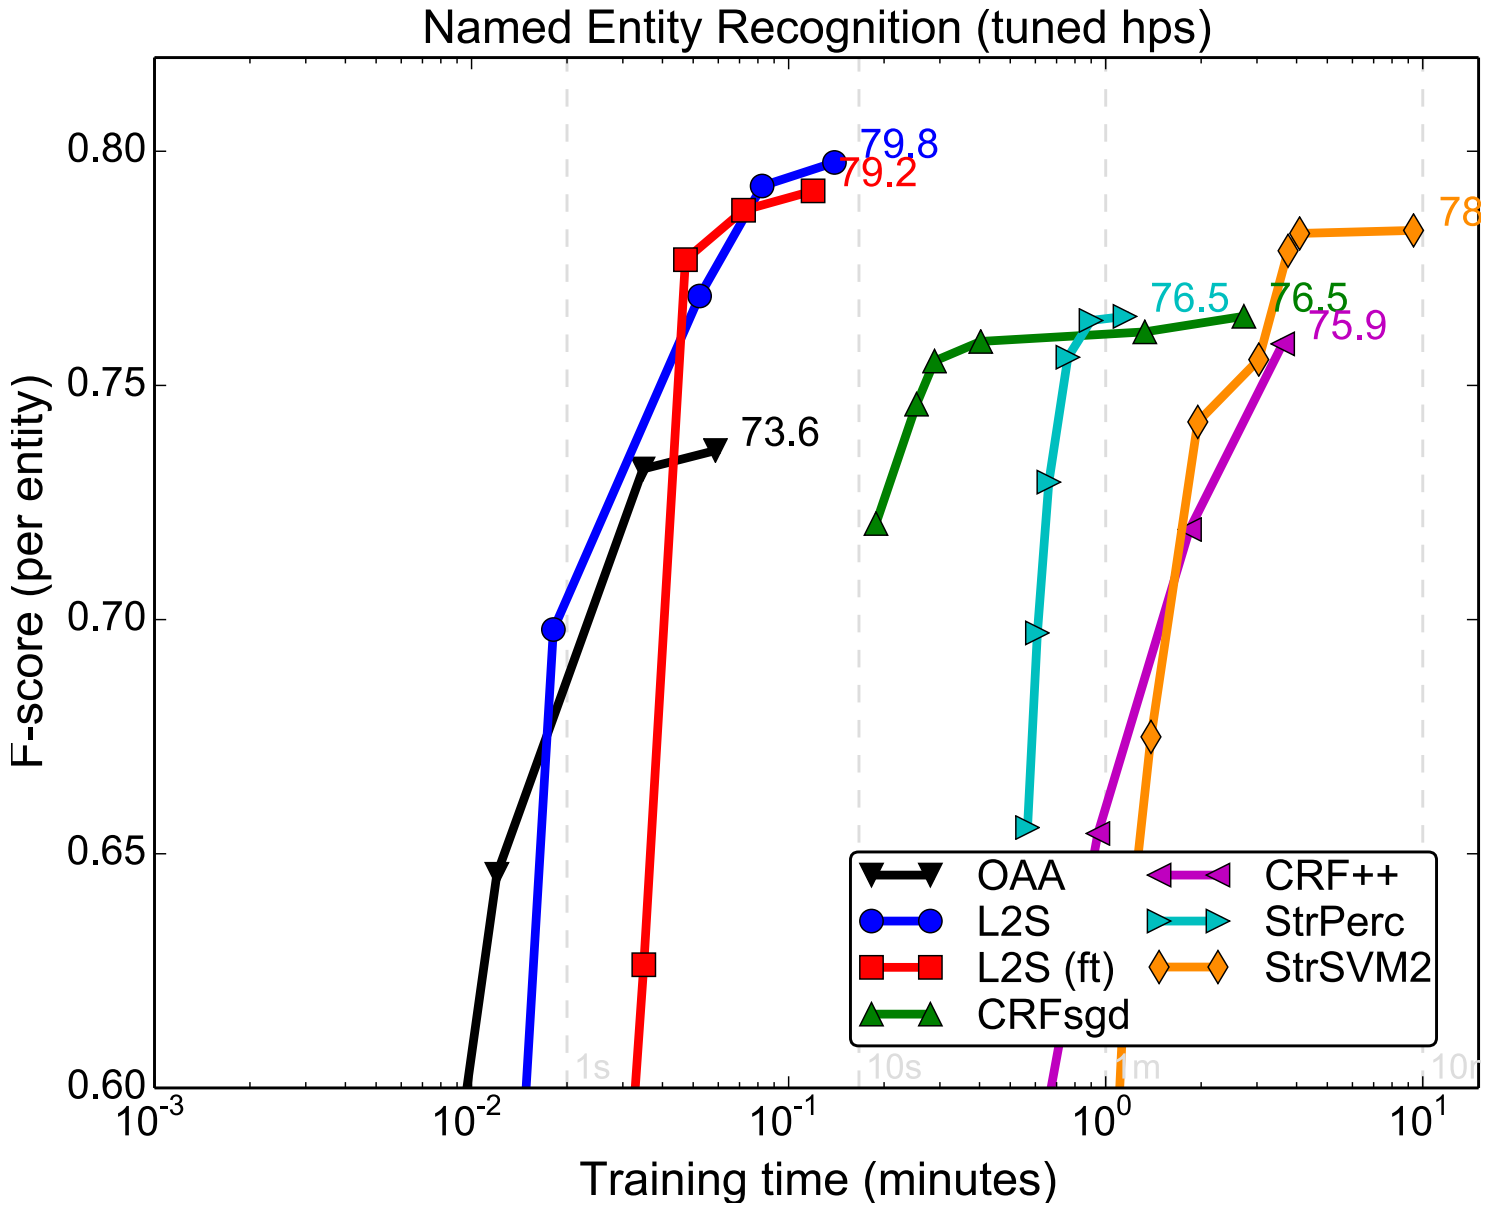

Figure from Langford & Daume III (ICML tutorial, 2015)

Experiments: Vowpal Wabbit L2S

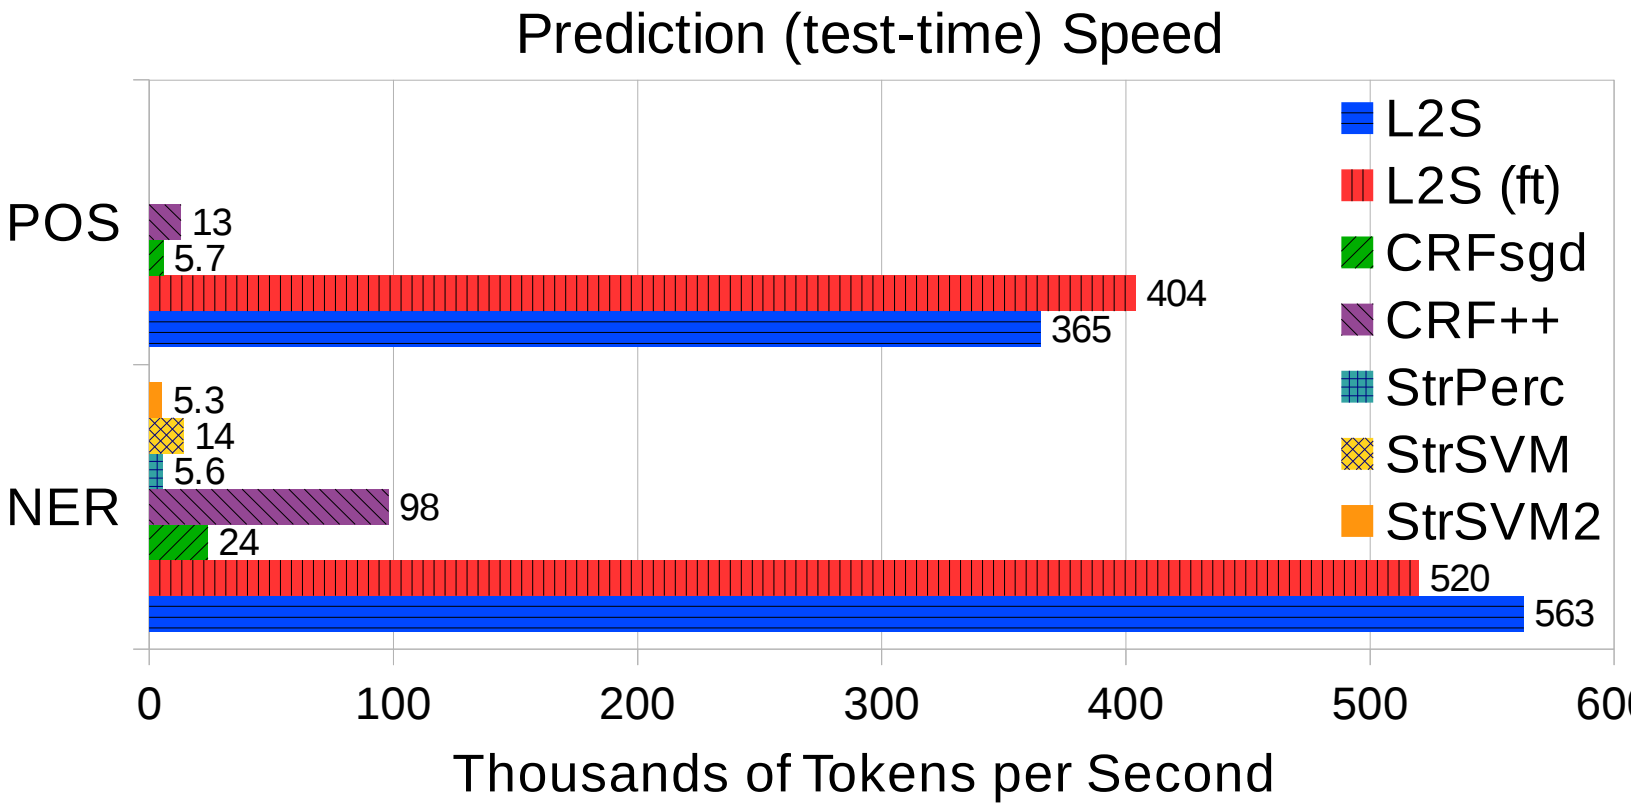

Figure from Langford & Daume III (ICML tutorial, 2015)

Learning 2 Search

Some key challenges:

- performance depends heavily on search order, but have to pick this by hand
- reference policy is critical, but what if it's too difficult to design one
- not always easy to make efficient on a GPU

Adapted from Langford & Daume III (ICML tutorial, 2015)

Learning Objectives

Structured Prediction as Search

You should be able to…

1. Reduce a structured prediction problem to a search problem
2. Contrast imitation learning with reinforcement learning
3. Implement Dagger, a learning to search algorithm
4. Explain the reduction of structured prediction to no-regret online learning
5. Contrast various learning2search algorithms based on their properties