22nd International Conference on Computational Linguistics

Advanced Dynamic Programming in Computational Linguistics: Theory, Algorithms and Applications Tutorial notes

Liang Huang

Department of Computer and Information Science
University of Pennsylvania

Order copies of this and other Coling proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
`acl@aclweb.org`

Design by Chimney Design, Brighton, UK
Production and manufacture by One Digital, Brighton, UK

Outline

Dynamic Programming (DP) is an important class of algorithms widely used in many areas of speech and language processing. It provides efficient solutions to seemingly intractable inference over exponentially-large spaces by sharing overlapping subproblems. Notable instances of DP include the well-known Viterbi and Forward-Backward Algorithms for finite-state models, and the CKY Algorithm for context-free parsing and syntax-based machine translation. The Dijkstra and A* Algorithms, although less obvious, can also be viewed as DP instances. In fact, almost all inference algorithms in the NLP/CL literature involve some sort of DP.

With this overwhelming popularity, a unified view of various DP algorithms would not only provide NLP researchers a better understanding but help them design new DP algorithms in practice. This tutorial, therefore, surveys two such theoretical frameworks: the semiring framework in the context of finite-state methods, and the hypergraph framework in the context of parsing and machine translation. Under each of these two paradigms, we review two most important types of DP algorithms, namely the Viterbi-style topological algorithms and the Dijkstra-style best-first algorithms. Wherever relevant, we will discuss typical instances of these algorithms in practice, which include applications in tagging and chunking, word alignment, phrase-based translation, syntactic parsing, and syntax-based translation.

Structure

The tutorial will be structured as follows:

1. Dynamic Programming on Lattices/Graphs under the Semiring Framework

   (a) Motivations and Examples
   (b) Semirings
   (c) Viterbi Algorithm
   (d) Dijkstra and A* Algorithms
   (e) Comparison between Viterbi and Dijkstra/A* Algorithms

2. Dynamic Programming on Packed Forests under the Hypergraph Framework

   (a) Hypergraphs and Related Formalisms
   (b) Examples in Parsing and Machine Translation
   (c) Generalized Viterbi Algorithm; CKY Parsing
   (d) Knuth and A* Algorithms

3. Extensions: Non-Optimization Problems and k-best Inference

   (a) Forward-backward and Inside-Outside Algorithms
   (b) k-best Inference in Graphs and Hypergraphs

Prerequisites

None!

Instructor

Liang Huang// Department of Computer & Information Science (CIS)// 3330 Walnut Street// Levine Hall// University of Pennsylvania// Philadelphia, PA 19104, USA// Email: lhuang3 at cis.upenn.edu

Liang Huang is finishing his PhD study at the University of Pennsylvania, co-supervised by Aravind Joshi and Kevin Knight (USC/ISI). He is mainly interested in the theoretical aspects of computational linguistics, in particular, efficient algorithms in parsing and machine translation, generic dynamic programming, and formal properties of synchronous grammars. His thesis work develops a series of "forest-based methods" that have been applied to many problems in NLP including k-best parsing, forest rescoring and reranking, and forest-based translation. This work received the outstanding paper award at ACL 2008 and a best paper award nomination at ACL 2007. He was also a recipient of the University Prize for Excellence in Teaching at Penn, and had designed and lectured one of the first Python Programming courses for CS undergraduates in the US.