Probabilistic Graphical Models (Concluded)

Lecture 20

Ganesh Ramakrishnan

CS337, Artificial Intelligence and Machine Learning

Recap: Inferencing in Graphical Models

- Exact inferencing
 Message passing
 Beam Search (exact with beams)
 A* search beams
 Junction tree algorithms Junction tree algorithms
 - Integer Linear Programming based inferencing

Exact Inferencing in Graphical Models

- Why hard ? We saw
- Message passing algorithm
 - Applicable to tree-structured directed and undirected graphical models – illustration on HMM Part of Speech Tagging
- Extra & Optional) Junction Tree algorithm
 - Applicable to arbitrary undirected graphs with cycles

Message Passing algorithm: Broad Idea

Li

x,

 $p(\pi_3|\pi_2) \geq p(\pi_1)$

- Example
- Find marginal for a particular node x_i

х,

 $p(x_i) = \sum_{x_1} \sum_{x_2} \dots \sum_{x_{i-1}} \sum_{x_{i+1}} \dots \sum_{x_L} p(x_1, x_2 \dots x_L)$ state nodes, cost is $O(M^L)$ ential in length of chain $p(x_2 | x_2)$. for M-state nodes, cost is $O(M^{L})$

x,

- exponential in length of chain
- but, we can exploit
 - the graphical structure (conditional independences)
 - 2. Avoidance of redundant computations through dynamic programming

[Exploiting that χ_i "blocks" info from $\chi_1 \ldots \chi_{i-1}$ To Message Passing: Broad idea $\chi_{i+1} \ldots \chi_L$]



CS337, Artificial Intelligence and Machine Learning



CS337, Artificial Intelligence and Machine Learning

Max Product Algorithm: Broad Idea

- Goal: find $x^{MAP} = \arg \max_{x} p(x)$ Recap Vitchi For POS
- Define $\phi(x_i) = \max_{x_1} \cdots \max_{x_{i-1}} \max_{x_{i+1}} \cdots \max_{x_L} p(x_1, \dots, x_L)$
- Then

$$x_i^{\mathsf{MAP}} = \arg\max_{x_i} \phi(x_i)$$

- Message passing algorithm with "sum" replaced by "max"
- Generalization of Viterbi algorithm for HMMs

Recap: Viterbi in HMM for POS For middle words, calculate the minimum score for all possible previous POS tags (we saw apphration of "beam" on Viterbi) natural best_score["2 NN"] = min(1:NN best_score["1 NN"] + -log P_r(NN|NN) + -log P_e(language | NN), best_score["1 JJ"] + -log P_r(NN|JJ) + -log P_r(language | NN), 1:JJ 2:JJ best_score["1 VB"] + -log P_T(NN|VB) + -log P_F(language | NN), best_score["1 LRB"] + -log P_r(NN|LRB) + -log P_r(language | NN), 1:VB 2:VB best_score["1 RRB"] + -log P_r(NN|RRB) + -log P_r(language | NN), 1:LRB 2:LRB best_score["2 JJ"] = min(best_score["1 NN"] + -log P_r(JJ|NN) + -log P_r(language | JJ), :RRB 2:RRE best_score["1 JJ"] + -log P_(JJ|JJ) + -log P_(language | JJ), best_score["1 VB"] + -log P_r(JJ|VB) + -log P_r(language | JJ), M

Summary of inference methods

	Chain (online)	Low treewidth	High treewidth
Discrete	BP = forwards (Viderbi	VarElim, Jtree,	Loopy BP, mean field,
Abitore	Boyen-Koller (ADF), beam search (BFS) A* (DFS)	recursive conditioning	structured variational, EP, graphcuts Gibbs
Gaussian	BP = Kalman filter	Jtree = sparse linear algebra	Loopy BP Gibbs
Other	EKF, UKF, moment matching (ADF) Particle filter	EP, EM, VB, NBP, Gibbs	EP, variational EM, VB, NBP, Gibbs

Exact Deterministic approximation Stochastic approximation

BP=belief propagation, EP = expectation propagation, ADF = assumed density filtering, EKF = extended Kalman filter, UKF = unscented Kalman filter, VarElim = variable elimination, Jtree= junction tree, EM = expectation maximization, VB = variational Bayes, NBP = non-parametric BP Slide Credit: Kevin Murphy

Search Algorithms: Further Illustrated

Lecture 20

Specifically culminating in A*

Ganesh Ramakrishnan

CS337, Artificial Intelligence and Machine Learning

Search Algorithms Illustrated



- Consider the problem of searching shortest path(s) from S (start) to G (goal).
- Refer: <u>https://moodle.iitb.ac.in/pluginfile.php/270556/mod_resource/content/1/searchingOnGraphs.pdf</u>

Depth First Search



Expand using stack



Beam Search

Suppose in addition to the standard search specification we also have a *heuristic*.

A heuristic function maps a state onto an estimate of the cost to the goal from that state.

Can you think of examples of heuristics?
Full for for for the short for the short

- Eucledian distance?
- Smallest of all outgoing edge weights at a node?
- Recap HMM: The highest emission probability a^{a} a^{b} a^{b} to a cost value.

Recap: A* Search: Optimistic Heuristic for HMM POS Taggiong

- Consider the words remaining
 - Use Optimistic Heuristic:
 - Upper bound on score (highest possible score)
 - Optimistic heuristic for tagging:
 - Best Emission Prob

natural	language	processing
log(P(natural NN)) = -2.4	$\log(P(lang. NN)) = -2.4$	log(P(proc. NN)) = -2.5
log(P(natural JJ)) = -2.0	log(P(lang. JJ)) = -3.0	log(P(proc. JJ)) = -3.4
log(P(natural VB)) = -3.1	log(P(lang. VB)) = -3.2	$\log(P(proc. VB)) = -1.5$
log(P(natural LRB)) = -7.0	log(P(lang. LRB)) = -7.9	log(P(proc. LRB)) = -6.9
$\log(P(natural RRB)) = -7.0$	$\log(P(lang. RRB)) = -7.9$	$\log(P(proc. RRB)) = -6.9$

$$H(1+) = -5.9$$
 $H(2+) = -3.9$ $H(3+) = -1.5$ $H(4+) = 0.0$

Euclidian Heuristic

Euclidian Heuristic

Steepest Ascent Hill Climbing

Steepest Ascent Hill Climbing

Init-PriQueue(PQ)

Insert-PriQueue(PQ,START,h(START))

while (PQ is not empty and PQ does not contain a goal state)

(s,h) := Pop-least(PQ)

foreach s' in succs(s)

if s' is not already in PQ and s' never previously been visited

Insert-PriQueue(PQ,s',h(s'))

N	number of states in the problem
В	the average branching factor (the average number of successors) (B>1)
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest cycle-free path from start to anywhere
Q	the average size of the priority queue

Algorithm		Comp lete	Optimal	Time	Space
BestFS	Best First Search	Y	N	O(min(N,B ^{LMAX}))	$O(min(N,B^{LMAX}))$

A few improvements to this algorithm can make things better

- Steepest Ascent Hill Climbing is clearly not guaranteed to find optimal S→A→B→C→G
- Obvious question: What can we do to avoid the stupid mistake? h(s) + c

A* - The Basic Idea

- Steepest Ascent Hill Climbing: When you expand a node n, take each successor n' and place it on PriQueue with priority h(n')
- A*: When you expand a node n, take each successor n' and place it on PriQueue with priority

g(n) + h(n)

 $(\underbrace{\text{Cost of getting to } n') + h(n')}_{\text{Let } g(n) = \text{Cost of getting to } n}$ (1) (2)

and then define...

(3)

A* Looking Non-Stupid

When should A* terminate?

Idea: As soon as it generates a goal state? Look at this example:

Correct A* termination rule:

A* Terminates Only When a Goal State Is Popped from the Priority Queue

A* revisiting nodes with longer paths are ignorable

Another question: What if A* revisits a state that was already expanded, and discovers a shorter path?

S h = 3Β *h* = *8* h = 7Α 1/2 h = 1 In this example a state that Need to update entries in Prionty Queue had been expanded gets G re-expanded. How and why?

A* revisiting states

What if A* visits a state that is already on the queue?

The A* Algorithm

= h(s) because

g(start) = 0

- Priority queue PQ begins empty.
- V (= set of previously visited (state,f,backpointer)-triples) begins empty.
- Put S into PQ and V with priority f(s) = g(s) + h(s)
- Is PQ empty?
 - > Yes? Sadly admit there's no solution
 - > No? Remove node with lowest f(n) from queue. Call it n.
 - > If *n* is a goal, stop and report success.
 - > "expand" *n* : For each n' in **successors**(n)....
 - Let f' = g(n') + h(n') = g(n) + cost(n,n') + h(n')

use sneaky trick to compute g(n)

- If n' not seen before, or n' previously expanded with f(n')>f', or n' currently in PQ with f(n')>f'
- Then Place/promote n' on priority queue with priority f' and update V to include (state=n', f', BackPtr=n).
- Else Ignore n'

Admissible Heuristics

- Write h*(n) = the true minimal cost to goal from n.
- A heuristic h is admissible if h(n) <= h*(n) for all states n.
- An admissible heuristic is guaranteed never to overestimate cost to goal.
- An admissible heuristic is optimistic.

A* with Admissible Heuristic Guarantees Optimal Path

State	h_1	h_2
S	5	4
A	3	2
B	6	6
C	2	1
D	3	3
G	0	0