

IRF proof

- Each step maximizes log-likelihood over ψ_c alone
- coordinate ascent



$$l(\theta|D) = \sum_x \sum_{k_c} m(k_c) \log \psi_c(k_c) - N \log Z$$

function of ψ_1, \dots, ψ_c

write single table entry



$$\frac{\partial}{\partial \psi_c(k_c)} = \frac{m(k_c)}{\psi_c(k_c)} - \frac{N}{Z} \frac{\partial}{\partial \psi_c(k_c)} \left[\sum_x \prod_D \psi_D(\tilde{x}_D) \right]$$

only hits terms in sum where table entry active

$$= \frac{m(k_c)}{\psi_c(k_c)} - \frac{N}{Z} \sum_x \delta(\tilde{x}_c, k_c) \frac{\partial}{\partial \psi_c(k_c)} \left[\prod_D \psi_D(\tilde{x}_D) \right]$$

locked at old ψ

$$\frac{\partial}{\partial \psi_c(k_c)} (\psi_c^{(t+1)}(k_c)) = \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} - \frac{N}{Z^{(t+1)}} \sum_x \delta(\tilde{x}_c, k_c) \prod_{D \neq c} \psi_D^{(t)}(\tilde{x}_D)$$

Z const

$$= \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} - \frac{N}{Z^{(t+1)}} \sum_x \delta(\tilde{x}_c, k_c) \prod_{D \neq c} \psi_D^{(t)}(\tilde{x}_D)$$

$$= \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} - \frac{N}{\psi_c^{(t+1)}(k_c)} \sum_x \delta(\tilde{x}_c, k_c) \frac{1}{Z^{(t+1)}} \prod_{D \neq c} \psi_D^{(t)}(\tilde{x}_D)$$

$$= \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} - \frac{N}{\psi_c^{(t+1)}(k_c)} \sum_x \delta(\tilde{x}_c, k_c) p(x)$$

$$= \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} - \frac{N}{\psi_c^{(t)}(k_c)} p(x)$$

$$= \frac{m(k_c)}{\psi_c^{(t+1)}(k_c)} \frac{p(x)}{p(x)} - \frac{N p(x)}{\psi_c^{(t)}(k_c)} = \frac{N p(x)}{\psi_c^{(t+1)}(k_c)} - \frac{N p(x)}{\psi_c^{(t)}(k_c)} = 0$$

M-step

example



$\psi_c(x) =$

summary: have efficient algo's for AIBLC, $p(y|x)$ & Max. Likelihood

Junction Tree Algorithm: - general inference tool for graphical mod.


- Workhorse, inference opdf is main tool

- Node elimination algo must be repeated for each query
redundant work gets discarded & repeated.


instead look closely at cliques we are eliminating
connect them & their separators.

- Tree: unique path between vertices, root node.

from node elimination



- cliques
- {x2, x5, x6}
 - {2, 3, 5}
 - {2, 4}
 - {1, 2, 3}
 - {1, 2}
 - {1, 3}




- junction tree property: each node \neq appears connected (in a path) to its other instantiations through the tree & all subtrees.

on unique path from V to W, all nodes have V \cap W.

- more generally: undirected graphs not on maximal cliques:

$$p(x) = \frac{1}{Z} \prod_{c \in C} \psi_c(x_c)$$



- Want to generally compute $p(x_E | x_C)$ or $p(x_E | x_C)$
- Efficiently infer marginals & cond's. \uparrow
 (over a set of nodes (clique for now))
- for a single clique

$p(x_C, \bar{x}_E)$ then set $p(\bar{x}_E) = \sum_{x_C} p(x_C, \bar{x}_E)$


then $p(x_C | \bar{x}_E) = \frac{p(x_C, \bar{x}_E)}{\sum_{x_C} p(x_C, \bar{x}_E)}$



$$p(x) = p(x_A) p(x_B | x_A) p(x_C | x_B)$$

easily can get $p(A, B)$ by $p(x_A) p(x_B | x_A)$
 but how to get $p(B, C)$?

First $p(x) = \prod_i p(x_i | x_{pa(i)}) = \frac{1}{Z} \prod_c \psi_c(x_c)$



- Review:
- Assignment 1, 2 coming
 - Assignment 3 online, projects online
 - Max. Likelihood for directed, undirected
 - decomposable, I.P.F.

Junction Trees Continued: want local marginals everywhere easy

$$\prod_i p(x_i | X_{A_i}) = \frac{1}{Z} \prod_c \psi_c \equiv \text{O} \xrightarrow{A} \text{O} \xrightarrow{B} \text{O} \xrightarrow{C}$$

Can write $\psi_{AB} = p(x_A) p(x_B | x_A) = p(x_A, x_B)$

$\psi_{BC} = p(x_C | x_B) \neq p(x_B, x_C)$


how to get? ↗

- Would like all cliques
- More flexible notation needed:

$$p(x) = \frac{1}{Z} \frac{\prod_c \psi_c(x_c)}{\prod_s \phi_s(x_s)} \leftarrow \begin{array}{l} \text{cliques} \\ \text{separators (between all clique-paths)} \end{array}$$

Z=1 in practice
put this in a $\psi_s(x_s)$ over null set


- Separators are function of intersection of neighbouring cliques
- Superset, if all $\psi_s(x_s) := 1$ get previous.
- But doesn't span anything new since any separator is subset of clique



$$p(x) = \frac{\psi_{AB} p(A, B) p(B, C) \psi_{BC}}{p(B) \phi_B}$$


now have nice marginals upstairs can just read off.

note division by zero can occur, but numerator will also be 0 so treat $\frac{0}{0} \equiv 0$
since $p(B) = \sum_{x_C} p(B, C)$



How to get the product over marginals form from $p(x_i | x_{\setminus i})$?

Junction Tree Algorithm Message Passing between cliques \Rightarrow consistency To get marginals with $p(x)$ staying consistent



if ψ_V is valid $p(A,B)$ can stop

message from V to W

$$\psi_W^* = \frac{\phi_S^* \psi_W}{\phi_S}$$

$$\psi_V^* = \psi_V$$

$\therefore p(x)$ unchanged

$$\frac{\psi_V^* \psi_W^*}{\phi_S^* \phi_S^*} = \frac{\psi_V \psi_W}{\phi_S}$$

if ψ_W is valid $p(B,C)$ can stop

message from W to V

$$\phi_S^{**} = \sum_{B,C} \psi_W^*$$

$$\psi_V^{**} = \frac{\phi_S^{**} \psi_V^*}{\phi_S^*}$$

$$\psi_W^{**} = \psi_W^*$$


$\therefore p(x)$ unchanged AND

$$\sum_{B,C} \psi_V^{**} = \sum_{B,C} \frac{\phi_S^{**}}{\phi_S^*} \psi_V^* = \frac{\phi_S^{**}}{\phi_S^*} \sum_{B,C} \psi_V^* = \phi_S^{**} = \sum_{A,B} \psi_V^*$$

Can pop these out now

\Rightarrow full clique margs V & W agree with marginal on S

mini-example 2



with evidence $X=1$, Init as above.

$$\phi_Y^* = \sum_x p(x,y) \delta(x=1) = p(x=1,y)$$

$$\psi_{Y,Z}^* = \frac{p(x=1,y)}{1} p(z|y) = p(x=1,y,z)$$

and $\psi_{X,Y}^* = p(x=1,y)$ as before end.

to get conds: $p(y,z|x=1) = \frac{\psi_{Y,Z}^*}{\psi_{X,Y}^*}$

- Message Passing Makes ^{active} marginals with proper interactions
- On big junction tree, can keep iterating messages but inefficient

- Send message only after hearing from all neighbours

J.T.A. - No need to iterate mindlessly

- initialize:
 - pick a root
 - $\phi_s = 1$ for all cliques
 - $\psi_c(x_c) = p(x_c | X_{\partial c}) \quad \forall c$

$$p(x) = \prod p(x_i | X_{\partial i}) = \prod \psi_c(x_c)$$

$$= \frac{\prod \psi_c(x_c)}{\prod_S \psi_S(x_S)} \leftarrow \text{denom} = 1$$

- update cliques & separators on a tree structure recursively

Collect Evidence (node)

for each child of node

{ update (node, collect Evidence (child)) };


return (node);

Distribute Evidence (node)


for each child of node

{ update (child, node);

distribute evidence (child) };



collect




distribute


TRIANGULATION To get a valid Junction Tree \rightarrow triangulate
 Undir. Node Elimination Algo \rightarrow Triangulates Graph.

Triangulation (like Moralization) adds links \rightarrow more complex graph, fewer independencies
 Such that: No 4 or more node cycle within graph where some nodes do not \perp


i.e. a 4-cycle



5-cycle



3-cycle






OK \checkmark


2-cycle \checkmark



1-cycle




example:





not a junction tree

triangulate:



or



- optimal triangulation (adds fewest links) is NP
- but any triangulation (can be bad heuristic) is P
- want to keep max clique size small
- also need resulting clique graph to be junction tree
 - a) tree (no loops)
 - b) satisfy junction tree property: cliques between path of γ_A & γ_B contain $A \cap B$ nodes

- When junction tree algo completes (induction proof)

$$\psi_C = p(x_C, \bar{x}_E)$$

$$\phi_S = p(x_S, \bar{x}_E)$$

← all cliques are marginals.

- Hugin (Junction Tree Algorithm) → Jensen's book

online

offline
(once for a given graph)

- Moralize (poly complex)

- Introduce evidence $x_E \rightarrow \bar{x}_E$ (poly)

- Triangulate (poly or NP for optimal)

- Construct Junction Tree (poly, Kruskal)

online

- Propagate Probabilities (poly in # of cliques, expo in clique size)

$$\phi_S^* = \sum_{\forall S} \psi_V$$




$$\psi_w^* = \frac{\phi_S^* \psi_w}{\phi_S}$$

(for different evidence queries etc.)

- Without Proof: Triangulated Graph \iff decomposable

junction tree \iff maximum weight spanning tree

largest cardinality of separator nodes



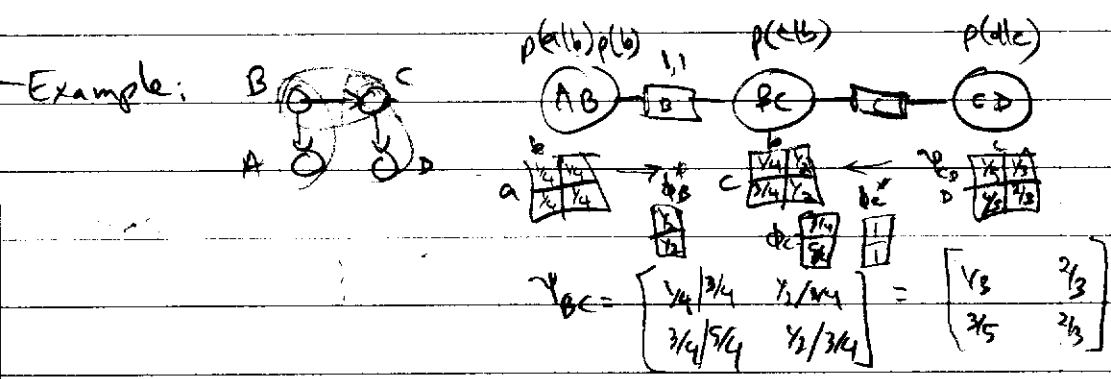
not a junction tree

$$|\text{separators}| = |D| + |CD| = 3$$


$$|\text{separators}| = |D| + |CD| = 4$$

- Algo to find junction tree: max card. of separators → Kruskal

- Kruskal: Max Span Tree:
 - 1) Begin with no edges between cliques
 - 2) Add edge with max separator card, & make sure we don't make a loop
 - 3) Repeat 2 until graph connected (all cliques are connected via a path)



with evidence
i.e. $X_C = 0$



- Can compute fast marginals
- Can also compute fast argmax


$$P(X_H, \bar{X}_E) = P(X_1, \dots, X_n, \bar{X}_{n+1}, \dots, \bar{X}_N)$$

← look all \bar{X}_E find X_H that maximizes $P()$

i.e. what is most likely state of patient if he has fever & headache?

$$\equiv \operatorname{argmax}_x P(x) = \max_x P(x) \overset{\max}{P}(x_2|x_1) \overset{\max}{P}(x_3|x_2)$$

$$\max_x P(x_1|x_2) \max_{x_2} P(x_3|x_2) \max_{x_3} P(x_4|x_3)$$



- flu
- fever
- sinus inf.
- temperature
- sinus swell
- headache