Lec29

Saturday, April 11, 2015 2:57 PM

- Deterministic SSP: Principle of Optimality:

$$J_k(i) = \min_{j=1,2,\dots,N} \{ a_{ij} + J_{k+1}(j) \}$$

Viterbi algorithm

Tuesday, March 03, 2015 11:45 AM

- Even random systems sometimes lead to a deterministic shortest path problem

Hidden Markov Models

- An underlying Markov chain with some finite state space and given state-transition probability P_{ij}
- However, we cannot directly observe the state of the Markov chain
- Instead, whenever a transition occurs, we obtain an observation that relates to the transition

  $r(z; i, j) \triangleq$ probability to observe z if the underlying transition is from state $i \to j$

  - z only depends on (i, j), but not other transitions.
- Based on the observations, we want to estimate the sequence of underlying states & transitions.

Many Applications of HMM

- Speech recognition:
  - state = "phonemes": a piece of elementary speech unit
    - e.g. /t/, /i:/ /k/

- observation = "sound recording"
- Why "Markov chain": phonemes in human speech is not independent. (e.g. certain sound is more likely to follow another sound in a language).
    - /ti:/ vs /tk/
    - Can be "trained" to yield the p_{ij}.

- Convolutional coding & decoding (skip)
    - In convolutional code, the output codeword y_k depends both on the source symbol w_k, but also the "state" x_k of the convolutional coder, e.g.

$$y_k = c x_{k-1} + d w_k$$

    - The state x_k has some transition probabilities

$$x_k = A x_{k-1} + b w_k$$

    - The received signal z_k is a noisy version of y_k

$$p(z_k | y_k)$$

    - Based on $z_1, z_2, \cdots z_k$, we would like to estimate $w_1, w_2, \cdots w_k$.
    - Why Markov chain? (Why not set $y_k = w_k$?)
        - The constraint that $y_1, \cdots, y_k$ must be generated by a valid sequence of state transitions helps to recover isolated errors.

State Estimation

- Given $Z_N = \{z_1, z_2, \cdots, z_N\}$, we would like to estimate $X_N = \{x_1, x_2, \cdots, x_N\}$

- Let
$$p\{X_N | Z_N\} = P\{\text{underlying state transition is } X_N \text{ given that observed sequence is } Z_N\}$$

  - It makes sense to "estimate" X_N as the one that maximizes $p\{X_N | Z_N\}$

- Now,
$$p\{X_N | Z_N\} = \frac{P\{X_N, Z_N\}}{P\{Z_N\}}$$
fixed for all X_N

  - It is sufficient to maximize $P\{X_N, Z_N\}$

- This turns out to be a deterministic shortest path problem.

$$p\{X_N, Z_N\} = P\{x_0, x_1, z_1, x_2, z_2, \cdots, x_N, z_N\}$$
$$= \pi_{x_0} P\{x_1 z_1, x_2, z_2 \cdots x_N z_N | x_0\}$$
$$= \pi_{x_0} P\{x_1, z_1 | x_0\} \cdot P\{x_2, z_2 \cdots x_N z_N | x_1\}$$
$$\cdots$$
$$= \pi_{x_0} \cdot p_{x_0 x_1} r(z_1; x_0, x_1) \cdot p_{x_1 x_2} r(z_2; x_1, x_2)$$

$$\cdots \; p_{x_{N-1},x_N}\, r(z_N; x_{N-1}, x_N))$$

- Thus, maximizing $P(X_N \mid Z_N)$ is equivalent to minimizing

$$-\ln \lambda_{x_0} - \sum_{k=1}^{N} \ln\left(p_{x_{k-1}x_k} \cdot r(z_k; x_{k-1} x_k)\right)$$

Trellis diagram

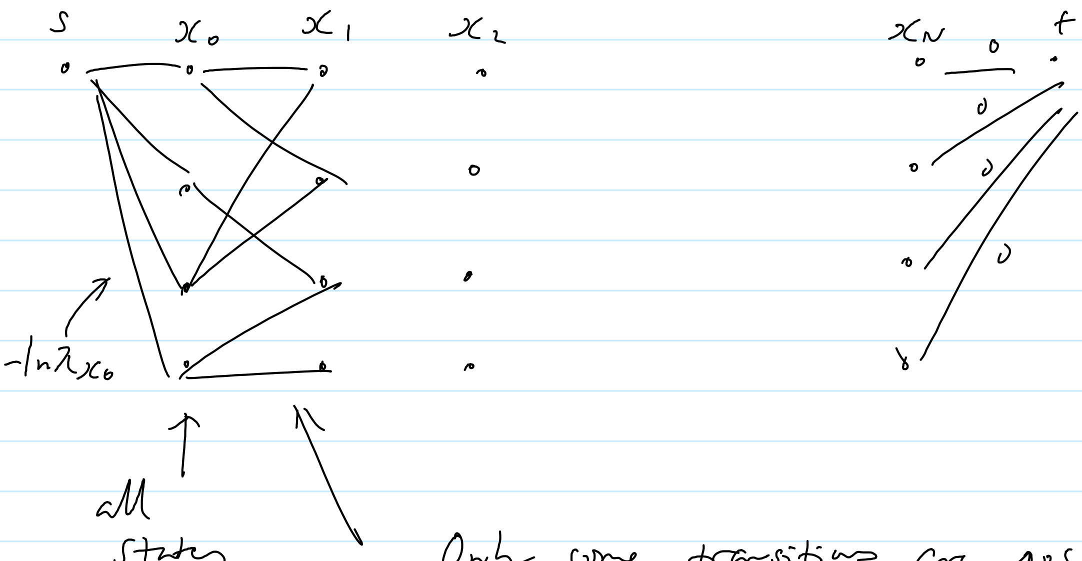

Only some transitions are possible ($p_{ij} > 0$)

the "cost" of each transition is

$$-\ln(p_{x_{k-1}x_k} \cdot r(z_k; x_{k-1} x_k))$$

Backward DP vs Forward DP

- We can now use backward DP to find the min-cost path from $s \to t$
  - Start from $J_N(x_N) \to J_{N-1}(x_{N-1}) \to \cdots$

- This backward DP would require us to know all the observations first before we can even run the first iteration.
    - It would be more desirable to run the iterations as each piece of observations comes in
- Instead, view S as the destination and t as the source.
    - Start from $J_N(x_0)$: cost from $S \to x_0$
    - Then $J_{N-1}(x_1)$, ...
    - This is called "forward DP"
- Another advantage of forward DP is that, if all shortest path currently known pass through a single node, then the state sequence up to this node can be determined without waiting for future observations

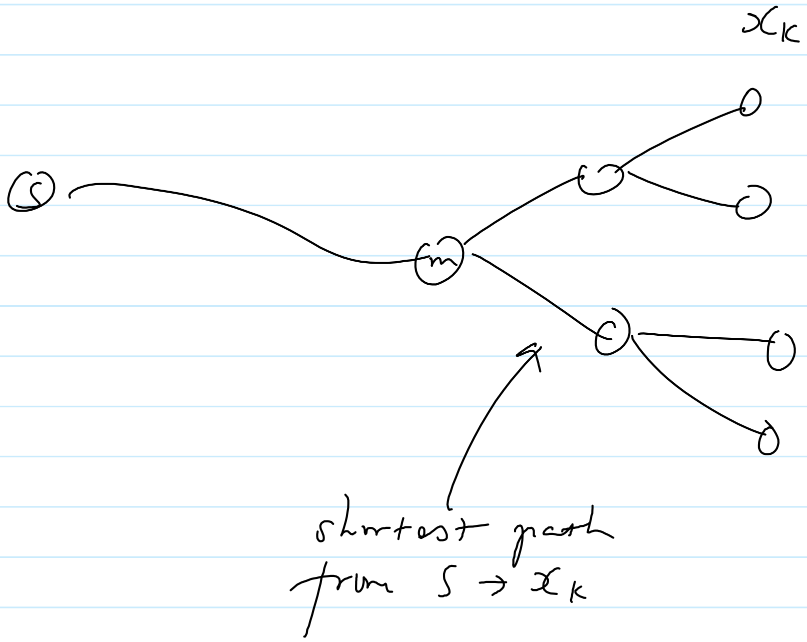

- Known as the Viterbi algorithm

Open-loop versus closed loop

Tuesday, March 03, 2015 11:04 AM

- In the above example, the entire sequence of optimal decisions can be computed in advance
- This is true for "deterministic" DP.
  - We can optimize directly over the outcome of the policy: $u_0, u_1, \cdots, u_{N-1}$
  - Open-loop: the decisions at all stages are computed in advance.

- However, if the system has uncertainty, then such an open-loop approach will be inadequate
  - Can you decide where to park in advance?
  - The optimal decision should naturally be adjusted based on the observations at previous stages.

    $\longrightarrow$ Closed loop (with feedback)

- In this case, we can not just optimize over the outcome of the decision policies. Rather, we need to optimize on the policy itself.
  - A policy μ_k describes how the action u_k depends on the state x_k
    - There is a u_k for every x_k!

Let us now illustrate the principle of optimality with the policy space.

① Start with

$$J_N(x_N) = g_N(x_N)$$ ← terminal cost.

② At $k = N-1$, consider the tail problem for x_{N-1}

$$J_{N-1}(x_{N-1}) = \min_{u \in U_{N-1}(x_{N-1})} E_{w_{N-1}}\left[g_{N-1}(x_{N-1}, u, w_{N-1}) + J_N\big(f_{N-1}(x_{N-1}, u, w_{N-1})\big)\right]$$

produce the random x_N

③ In general

$$J_k(x_k) = \min_{u_k \in U_k(x_k)} E_{w_k}\left[g_k(x_k, u_k, w_k) + J_{k+1}\big(f_k(x_k, u_k, w_k)\big)\right] \quad (*)$$

④ The "final" cost $J_0(x_0)$, should be equal to the optimal $J^*(x_0)$.

Further, the optimal policy is

$$\pi^* = \{\mu_0^*, \mu_1^*, \cdots \mu_{N-1}^*\}$$

where each $\mu_k^*(x_k)$ takes as input/feedback the current state x_k, and produces the action that minimizes the right-hand-side of $(*)$.

Proof by induction - skip

Tuesday, March 03, 2015 11:21 AM

Hypothesis : $J_k(x_k)$ is equal to $J_k^*(x_k)$, which is defined as the optimal cost of the tail sub-problem that starts at time k at state x_k

- True for $k = N$ (Trivial)

$$J_N(x_N) = g(x_N) = J_N^*(x_N)$$

- Assume that the Hypothesis holds for $k+1$

$$J_{k+1}(x_{k+1}) = J_{k+1}^*(x_{k+1}) \quad \text{for all } x_{k+1}$$

We want to show that it implies

$$J_k(x_k) = J_k^*(x_k)$$

Let $\pi_{k+1} = \{\mu_{k+1}, \mu_{k+2}, \cdots, \mu_{N-1}\}$

$\pi_k = \{\mu_k, \pi_{k+1}\}$

Note that

$$J_k^*(x_k) = \min_{\{\mu_k, \pi_{k+1}\}} \underset{w_k, \cdots w_{N-1}}{E} \Big\{ g_k(x_k, \mu_k(x_k), w_k) + g_N(x_N) + \sum_{i=k+1}^{N-1} g_i(x_i, \mu_i(x_i), w_i) \Big\}$$

$$= \min_{\mu_k} \min_{\pi_{k+1}} E_{w_k} \Big\{ \cancel{\underset{w_{k+1}, \cdots, w_{N-1}}{E}} \big[g_k(x_k, \mu_k(x_k), w_k) \big] + \underset{w_{k+1}, \cdots w_{N-1}}{E} \Big[g_N(x_N) + \sum_{i=k+1}^{N-1} g_i(x_i, \mu_i(x_i), w_i) \Big| w_k \Big] \Big\}$$

$$= \min_{\mu_k} E_{w_k} \Big\{ g_k(x_k, \mu_k(x_k), w_k) + \min \underset{w_{k+1}, \cdots w_{N-1}}{E} \Big[g_N(x_N) + \sum_{i=k+1}^{N-1} g_i(x_i, \mu_i(x_i), w_i) \Big| w_k \Big] \Big\}$$

$\vdots$

$$+ \underbrace{\min_{\pi_{k+1}} \mathbb{E}_{\omega_{k+1},\dots,\omega_{N-1}} \left[g_N(x_N) + \sum_{i=k+1}^{N-1} g_i(x_i, \mu_i(x_i), \omega_i) \middle| \omega_k \right]}_{= J^*_{k+1}(f_k(x_k, \mu_k(x_k), \omega_k))} \Bigg\}$$

$$||$$

$$= \min_{\mu_k} \mathbb{E}_{\omega_k} \Big\{ g_k(x_k, \mu_k(x_k), \omega_k) + J_{k+1}(f_k(x_k, \mu_k(x_k), \omega_k)) \Big\}$$

$$= J_k(x_k)$$

Two stage problem

Tuesday, March 17, 2015 2:29 PM

- Let us first study a simpler class of DP problems that only have two stages

Ex 1) "Newsvendor" problem

- In each morning, a newsboy needs to decide how many copies of newspaper to buy from the publisher
- Each copy of the newspaper costs him c. If he can sell it, he will earn $p > c$.
- However, the number of copies that he can sell, D, is unknown when he buys the papers.
- Unsold papers cannot be returned.
- If he buys too few $\longrightarrow$ lost sales

  If he buys too many $\longrightarrow$ wasting money.
- How much should he buy?

- Clearly, this is a two-stage problem

  Stage 1: buy newspaper

  Stage 2: sell newspapers
- Let q = # of copies bought from publisher
  - (q, D) is the state for the 2nd stage

- The decision at the 2nd stage is trivial
  - If $q \le D$, sell q copies
  - If $q > D$, sell D copies
  - $J_2(q, D) = -p \min(q, D)$ (negative cost means revenue)
- Go back to the 1st stage

$$J_1 = \min_q \; cq + E_D[J_2(q, D)]$$

$$\Leftrightarrow \quad \max_q \; p E[\min(q, D)] - cq$$

Solution:

- Let $\lambda(q) = p E[\min(q, D)] - cq$
  - This is concave in q
  - $$\frac{d\lambda}{dq} = p \cdot E[1_{\{q \le D\}}] - c$$
    $$= p \cdot P\{q \le D\} - c = 0$$
    $$\Rightarrow \quad P\{q \le D\} = \frac{c}{p}$$
    $$q = F_D^{-1}\left(1 - \frac{c}{p}\right)$$

Ex 2) Wind Energy integration.

Ref: R. Rajagopal, E. Bitar, F. Wu and P. Varaiya
"Risk-Limiting Dispatch of Wind Power".
in 2012 American Control Conference,
Montreal, Canada, June 2012.

- Future wind supply is uncertain.
- Suppose that a dispatcher (think of it as the utility company) wants to utilize wind energy to serve demand.
    - It needs to buy additional energy if wind power is not enough
- Usually, energy is cheaper to buy in advance (in the so called "day-ahead" market), versus buying it immediately ("real time")
- Let C_d be the price of buying energy day-ahead,

  and C_r be the price of buying energy in real-time
    - Assume $C_r > C_d$, otherwise one would always wait until real-time.
    - Assume wind is of zero cost.
- Let x be the amount of energy bought day ahead.

  Let W be the actual wind output
  Let L be the actual load
    - $D = L - W$ is the net load.

- If $x \geq D$, no real-time purchase is needed
  If $x < D$, need to buy $D-x$ in real-time.
- How should we determine x?

- Again, this is a two-stage problem:
  - Stage 1: day-ahead
  - Stage 2: real-time
- (x, D) is the state for the 2nd stage
  - $J_2(x, D) = C_r [D-x]^+$
- Go back to stage 1
  - $J_1 = \min_x C_d x + E[J_2(x, D)]$

$$= \min_x C_d x + E\left(C_r [D-x]^+ \right)$$

Solution:

- Let $\lambda(x) = C_d x + E\left(C_r [D-x]^+ \right)$

$$\frac{d\lambda}{dx} = C_d - C_r E\left[\mathbb{1}_{\{D \geq x\}} \right]$$

$$= C_d - C_r P\{D \geq x\} = 0$$

$$\Rightarrow \quad P\{D \geq x\} = \frac{C_d}{C_r}$$

$$x = F_D^{-1}\left(1 - \frac{C_d}{C_r}\right)$$

Generalizations

- More than 2 stages?
  - Multiple intermediate markets
- Future prices may be uncertain?
- More decisions at the various stages
  - Buy an "option"
    - Pay for capacity in advance, then pay for the actual energy
  - Sell back the energy
  - Storage.