HMMs

CPS 1/296

Ronald Parr

HMM Applications

- Target tracking
- Patient/factory monitoring
- Speech recognition
- Robot mapping and localization

Back to Atomic Events

- We began talking about probabilities from the perspective of atomic events
- An atomic event is an assignment to every random variable in the domain
- For n random variables, there are 2^n possible atomic events
- State variables return later

States

- When reasoning about time, we often call atomic events states
- States, like atomic events, form a mutually exclusive and jointly exhaustive partition of the space of possible events
- We can describe how a system behaves with a state-transition diagram

State Transition Diagram

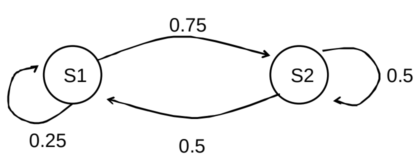

P(S2|S1)=0.75
P(S1|S1)=0.25
P(S2|S2)=0.50
P(S1|S2)=0.50

Don't confuse states with state variables!
Don't confuse states with state variables!
Don't confuse states with state variables!

State Transition Diagrams

- Make a lot of assumptions
  - Transition probabilities don't change over time (*stationarity*)
  - The event space does not change over time
  - Probability distribution over next states depends only on the current state (*Markov assumption*)
  - Time moves in uniform, discrete increments

The Markov Assumption

- Let S_t be a random variable for the state at time t
- $P(S_t|S_{t-1},\ldots,S_0) = P(S_t|S_{t-1})$
- (Use subscripts for time; S0 is different from S_0)
- Markov is special kind of conditional independence
- Future is independent of past given current state

Markov Models

- A system with states that obey the Markov assumption is called a *Markov Model*
- A sequence of states resulting from such a model is called a *Markov Chain*
- The mathematical properties of Markov chains are studied heavily in mathematics, statistics, computer science, electrical engineering, etc.

What's The Big Deal?

- A system that obeys the Markov property can be described succinctly with a transition matrix, where the i,jth entry of the matrix is P(Sj|Si)
- The Markov property ensures that we can maintain this succinct description over a potentially infinite time sequence
- Properties of the system can be analyzed in terms of properties of the transition matrix
  - Steady-state probabilities
  - Convergence rate, etc.

Observations

- Introduce E_t for the observation at time t
- Observations are like evidence
- Define the probability distribution over observations as function of current state: P(E|S)
- Assume observations are conditionally independent of other variables given current state
- Assume observation probabilities are stationary

Applications

- Monitoring/Filtering
  - S is the current status of the patient/factory
  - E is the current measurement
- Prediction
  - S is the current/future position of an object
  - E are our past observations
  - Project S into the future

Applications

- Smoothing/hindsight
  - Update view of the past based upon future
  - Diagnosis: Factory exploded at time t=20, what happened at t=5 to cause this?
- Most likely explanation
  - What is the most likely sequence of events (from start to finish) to explain what we have seen?

Viterbi Path

From definition of HMM and conditional indendpence:

$$P(S_0E_0 \ldots S_tE_t) = P(S_0)P(E_0 \mid S_0)\prod_{i=1}^{t} P(S_i \mid S_{i-1})P(E_i \mid S_i)$$

Suppose we want max probability sequence of states:

$$\max_{S_0 \ldots S_t} P(S_0E_0 \ldots S_tE_t) = \max_{S_0 \ldots S_t} P(S_0)P(E_0 \mid S_0)\prod_{i=1}^{t} P(S_i \mid S_{i-1})P(E_i \mid S_i)$$

$$= \max_{S_1 \ldots S_t} \prod_{i=1}^{t} P(S_i \mid S_{i-1})P(E_i \mid S_i) \max_{S_0} P(S_1 \mid S_0)P(S_0)P(E_0 \mid S_0)$$

$$= \max_{S_2 \ldots S_t} \prod_{i=2}^{t} P(S_i \mid S_{i-1})P(E_i \mid S_i) \max_{S_1} P(S_2 \mid S_1)P(E_1 \mid S_1) \max_{S_0} P(S_1 \mid S_0)P(S_0)P(E_0 \mid S_0)$$

Keep distributing max over product!

Viterbi Path vs. Stereo DP

- At first seem like very different things
- Both are actually the SAME DP algorithm
- Stereo DP:
  - Can be viewed as shortest path through disparity space
  - Uses fact that addition is associative
- Viterbi path:
  - Computes lowest cost path through HMM
  - Uses fact that multiplication is associative
- Logs vs. probabilities
  - Maximizing product of probs equivalent to maximizing sum of log probabilities
  - Converting HMM probabilities to log probabilities makes viterbi classical shortest path problem

Deriving Tracking Equations Our Main Tool

$$P(A \wedge B) = P(B \wedge A)$$

$$P(A \mid B)P(B) = P(B \mid A)P(A)$$

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$$

Extending Bayes Rule

$$P(A \mid BC) = \frac{P(B \mid AC)P(A \mid C)}{P(B \mid C)}$$

How to think about this: The C is like “extra” evidence.
This forces us into one corner of the event space.
Given that we are in this corner, everything behaves the same.

Monitoring

We want: $P(S_t \mid e_t \ldots e_0)$

$$P(S_t \mid e_t \ldots e_0) = \frac{P(e_t \mid S_t, e_{t-1} \ldots e_0)P(S_t \mid e_{t-1} \ldots e_0)}{P(e_t \mid e_{t-1} \ldots e_0)}$$

$$= \alpha P(e_t \mid S_t e_{t-1} \ldots e_0)P(S_t \mid e_{t-1} \ldots e_0)$$

$$= \alpha P(e_t \mid S_t)P(S_t \mid e_{t-1} \ldots e_0)$$

$$= \alpha P(e_t \mid S_t)\sum_{S_{t-1}} P(S_t \mid S_{t-1})P(S_{t-1} \mid e_{t-1} \ldots e_0)$$

Recursive

Hindsight

$$P(S_k \mid e_t \ldots e_0) = \alpha P(e_t \ldots e_{k+1} \mid S_k, e_k \ldots e_0)P(S_k \mid e_k \ldots e_0)$$

$$= \alpha P(e_t \ldots e_{k+1} \mid S_k)P(S_k \mid e_k \ldots e_0)$$ Monitoring!

$$P(e_t \ldots e_{k+1} \mid S_k) = \sum_{S_{k+1}} P(e_t \ldots e_{k+1} \mid S_k S_{k+1})P(S_{k+1} \mid S_k)$$

$$= \sum_{S_{k+1}} P(e_t \ldots e_{k+1} \mid S_{k+1})P(S_{k+1} \mid S_k)$$

$$= \sum_{S_{k+1}} P(e_{k+1} \mid S_{k+1})P(e_t \ldots e_{k+2} \mid S_{k+1})P(S_{k+1} \mid S_k)$$

Recursive

Hindsight Summary

- Forward: Compute k state distribution given
  - Forward distribution up to k
  - Observations up to k
  - Equivalent to monitoring up to k
  - Equivalent to eliminating variables <k
- Backward: Compute conditional evidence distribution after k
  - Work backward from t to k
  - Equivalent to to eliminating variables >k
- Smoothed state distribution is proportional to product of forward and backward components

Checkpoint

- Done: Forward Monitoring and Backward Smoothing
- Monitoring is recursive from the past to the present
- Backward smoothing requires two recursive passes
- Called the forward-backward algorithm
  - Independently discovered many times throughout history
  - Was classified for many years by US Govt.

What's Left?

- We have seen that filtering and smoothing can be done efficiently, so what's the catch?
- We're still working at the level of atomic events
- There are too many atomic events!
- Not all systems are discrete
- Need:
  - Ways of dealing with continuous variables
  - Ways of dealing with very large state spaces

Continuous Variables

- How do we represent a probability distribution over a continuous variable?
  - Probability density function
  - Summations become integrals
- Very messy except for some special cases:
  - Distribution over variable X at time t+1 is a multivariate normal with a mean that is a linear function of the variables at the previous time step
  - This is a linear-Gaussian model

Inference in Linear Gaussian Models

- Filtering and smoothing integrals have closed form solution
- Elegant solution known as the Kalman filter
  - Used for tracking projectiles (radar)
  - State is modeled as a set of linear equations
    - S=vt
    - V=at
  - What about pilot controls?

Related Topics

- Continuous time
  - Need to model system using differential equations
- Non-stationarity
  - What if the model changes over time?
  - This touches on learning
- What about controlling the system w/actions?
  - Markov decision processes