

Probabilistic Temporal Reasoning Study Guide (AIMA 14)

Artificial Intelligence

1 Probabilistic Temporal Reasoning

1. You are creating discrete-time model for a system that changes approximately every minute. What is a good value for the time interval, Δ , between time slices?

Solution: Any value less than 1 minute would work. You usually don't want the system to change more than once since the last measurement. Consider the case of a sinusoidal function where the time slices and Δ are set such that each time you measure, you get 0. You would get an incorrect view of the system.

2. What is the general equation for the Markov assumption in a first-order Markov transition model over the variables \mathbf{X}

Solution: $Pr(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = Pr(\mathbf{X}_t | \mathbf{X}_{t-1})$

3. What is the general equation for for the Markov assumption in a second-order Markov transition model over the variables \mathbf{X}

Solution: $Pr(\mathbf{X}_t | \mathbf{X}_{0:t-1}) = Pr(\mathbf{X}_t | \mathbf{X}_{t-2}, \mathbf{X}_{t-1})$

4. Describe the four basic inference tasks in probabilistic temporal models.

Solution:

- Filtering, a.k.a., state estimation is the task of computing the belief state $Pr(\mathbf{X}_t | \mathbf{e}_{1:t})$ – the posterior distribution over the most recent state given all the evidence to date.
- Prediction is the task of computing the posterior distribution over the future state, given all evidence to date: $Pr(\mathbf{X}_{t+k} | \mathbf{e}_{1:t})$ for some $k > 0$.
- Smoothing is the task of computing the posterior distribution over a past state, given all evidence up to the present: $Pr(\mathbf{X}_k | \mathbf{e}_{1:t})$ for some k such that $0 \leq k < t$.
- Most likely explanation: Given a sequence of observations, find the sequence of states that is most likely to have generated those observations: $\operatorname{argmax}_{\mathbf{x}_{1:t}} Pr(\mathbf{x}_{1:t} | \mathbf{e}_{1:t})$.

5. Give an example of each kind of basic probabilistic temporal inference task.

Solution:

- Filtering: Umbrella example: compute probability of rain today given all umbrella observations so far.
- Prediction: Umbrella example: compute probability of rain three days from now, given all the observations to date.
- Smoothing: Umbrella example: compute the probability that it rained last Wednesday, given all the observations of the umbrella carrier made up to today.
- Most likely explanation: Umbrella example: if umbrella appears on each of the first three days and is absent on the fourth, then the most likely explanation is that it rained on the first three days and did not rain on the fourth. On the other hand, perhaps the director forgot an umbrella on a rainy day, or took an umbrella on a sunny day out of caution. Other examples: speech recognition.