Informatics 2D: Reasoning and Agents

Alex Lascarides

School of Informatics
The University of Edinburgh

Lecture 26a: Time and Uncertainty: Stationary Processes and the Markov Assumption

Where are we?

So far...

- Completed our account of Bayesian Networks
- Dealt with methods for exact and approximate inference in BNs
- Enumeration, variable elimination, sampling, MCMC

Today ...

- **Time and uncertainty I**

Time and uncertainty

- So far we have only seen methods for describing uncertainty in static environments
- Every variable had a fixed value, we assumed that nothing changes during evidence collection or diagnosis
- Many practical domains involve uncertainty about processes that can be modelled with probabilistic methods
- Basic idea straightforward: imagine one BN model of the problem for every time step and reason about changes between them

States and observations

- Adopted approach similar to situation calculus: series of snapshots (**time slices**) will be used to describe process of change
- Snapshots consist of observable random variables $\mathbf{E}_t$ and non-observable ones $\mathbf{X}_t$
- For simplicity, we assume sets of (non)observable variables remain constant over time, but this is not necessary
- Observation at t will be $\mathbf{E}_t = \mathbf{e}_t$ for some set of values $\mathbf{e}_t$
- Assume that states start at $t = 0$ and evidence starts arriving at $t = 1$

States and observations

- Example: underground security guard wants to predict whether it is raining but only observes every morning whether director comes in carrying umbrella
- For each day, $\mathbf{E}_t$ contains variable U_t (whether the umbrella appears) and $\mathbf{X}_t$ contains state variable R_t (whether it's raining)
- Evidence $U_1, U_2, \ldots$, state variables $R_0, R_1, \ldots$
- Use notation $a : b$ to denote sequences of integers, e.g. $U_1, U_2, U_3 = U_{1:3}$

Stationary processes and the Markov assumption

- How do we specify dependencies among variables?
- Natural to arrange them in temporal order (causes usually precede effects)
- Problem: set of variables is unbounded (one for each time slice), so we would have to
  - specify unbounded number of conditional probability tables
  - specify an unbounded number of parents for each of these
- Solution to first problem: we assume that changes are caused by a **stationary process** – the laws that govern the process do not change themselves over time (not to be confused with "static")
- For example, $\mathbf{P}(U_t|Parents(U_t))$ does not depend on t

Stationary processes and the Markov assumption

- Solution to second problem: **Markov assumption** – the current state only depends on a finite history of previous states
- Such processes are called Markov processes or Markov chains
- Simplest form: **first-order Markov processes**, every state depends only on predecessor state
- We can write this as $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{0:t-1}) = \mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$
- This conditional distribution is called **transition model**
- Difference between first-order and second-order Markov processes:

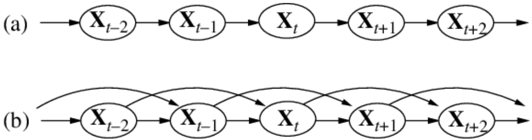

Stationary processes and the Markov assumption

- Assume that evidence variables are conditionally independent of other stuff given the current state:

$$\mathbf{P}(\mathbf{E}_t|\mathbf{X}_{0:t},\mathbf{E}_{0:t-1}) = \mathbf{P}(\mathbf{E}_t|\mathbf{X}_t)$$

- This is called the **sensor model** (**observation model**) of the system
- Notice direction of dependence: state causes evidence (but inference goes in other direction!)
- In umbrella world, rain causes umbrella to appear
- Finally, we need a prior distribution over initial states $\mathbf{P}(\mathbf{X}_0)$
- These three distributions give a specification of the complete JPD:

$$\mathbf{P}(\mathbf{X}_0,\mathbf{X}_1,\ldots,\mathbf{X}_t,\mathbf{E}_1,\ldots,\mathbf{E}_t) = \mathbf{P}(\mathbf{X}_0)\prod_{i=1}^{t}\mathbf{P}(\mathbf{X}_i|\mathbf{X}_{i-1})\mathbf{P}(\mathbf{E}_i|\mathbf{X}_i)$$

Umbrella world example

- Bayesian network structure and conditional distributions
- Transition model $P(Rain_t|Rain_{t-1})$, sensor model $P(Umbrella_t|Rain_t)$

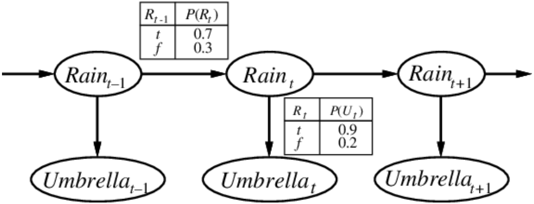

- Rain depends only on rainfall on previous day, whether this is reasonable depends on domain!

Stationary processes and the Markov assumption

- If Markov assumptions seems too simplistic for some domains (and hence, inaccurate), two measures can be taken
  - We can increase the order of the Markov process model
  - We can increase the set of state variables
- For example, add information about season, pressure or humidity
- But this will also increase prediction requirements (problem alleviated if we add new sensors)
- Example: dependency of predicting movement of robot on battery power level
  - add battery level sensor

Summary

Time and Uncertainty:

- In a dynamic environment, random variables change values over time
- There are (Latent) state variables and variables whose values are observed.
- Stationarity and Markov assumptions are important for obtaining a compact representation of an unbounded process
- They're also important for practical inference!
- Next time: **Time and Uncertainty: Inference**