Probabilistic Programming

Lecture #3: Markov Chains

Joost-Pieter Katoen

Software Modeling and Verification Chair | RWTH AACHEN UNIVERSITY

RWTH Lecture Series on Probabilistic Programming 2018

Overview

Overview

Probability distribution

Probability distribution

A (discrete) probability distribution on countable set X is a function $\mu : X \to [0,1] \subseteq \mathbb{R}$ such that $\sum_{x \in X} \mu(x) = 1$.

The set $\{ x \mid \mu(x) > 0 \}$ is the support set of probability distribution μ.

Let $Dist(X)$ denote the set of all probability measures on X.

Andrei Andrejewitsch Markow

Markov chains

Markov chain

A Markov chain (MC) D is a triple $(\Sigma, \sigma_I, \mathbf{P})$ with:

- Σ being a countable set of states
- $\sigma_I \in \Sigma$ the initial state, and
- $\mathbf{P} : \Sigma \to Dist(\Sigma)$ the transition probability function

where $Dist(\Sigma)$ is a (discrete) probability measure on Σ.

A state $\sigma \in \Sigma$ for which $\mathbf{P}(\sigma, \sigma) = 1$ is called absorbing.

Transition probability matrix

For MC D with finite state space Σ, function $\mathbf{P}$ is called the *transition probability matrix* of D.

Properties:

1. $\mathbf{P}$ is a (right) *stochastic* matrix, i.e., it is a square matrix, all its elements are in $[0, 1]$, and each row sum equals one.
2. $\mathbf{P}$ has an eigenvalue of one, and all its eigenvalues are at most one.
3. For all $n \in \mathbb{N}$, $\mathbf{P}^n$ is a stochastic matrix.

Paths

Paths

Path $\pi = \sigma_0 \, \sigma_1 \, \ldots$ is a *path* through MC D whenever $\mathbf{P}(\sigma_i, \sigma_{i+1}) > 0$ for all natural i.

Let $Paths(D)$ denotes the set of paths in D that start in its initial state σ_I.

Example

Cylinder sets

Cylinder set

The *cylinder set* of finite path $\hat{\pi} = \sigma_0 \, \sigma_1 \ldots \sigma_n$ in MC D is defined by:

$$Cyl(\hat{\pi}) \;=\; \left\{ \pi \in Paths(D) \mid \hat{\pi} \text{ is a prefix of } \pi \right\}$$

The cylinder set spanned by finite path $\hat{\pi}$ consists of all infinite paths that have prefix $\hat{\pi}$.

Probability measure on sets of infinite paths

Probability measure

Pr is the unique *probability distribution* defined on cylinder sets by:

$$Pr\big(Cyl(\sigma_0 \ldots \sigma_n)\big) \;=\; \prod_{0 \leq i < n} \mathbf{P}(\sigma_i, \sigma_{i+1})$$

for $n > 0$ and $\mathbf{P}(\sigma_0) = 1$ iff $\sigma_0 = \sigma_I$.

By standard results in probability theory, Pr is a distribution on all sets of infinite paths that are countable (disjoint) unions and/or complements of cylinder sets.

Reachability

Reachability

Let MC D with countable state space Σ and $G \subseteq \Sigma$ the set of *goal* states. The event *eventually reaching G* is defined by:

$$\Diamond G \;=\; \{ \pi \in Paths(D) \mid \exists i \in \mathbb{N}.\, \pi[i] \in G \}$$

where $\pi[i] = \sigma_{i+1}$ for $\pi = \sigma_0 \, \sigma_1 \ldots$.

The event $\Diamond G$ is measurable, i.e., the probability $Pr(\Diamond G)$ is well defined.

Reachability probabilities: Knuth's die

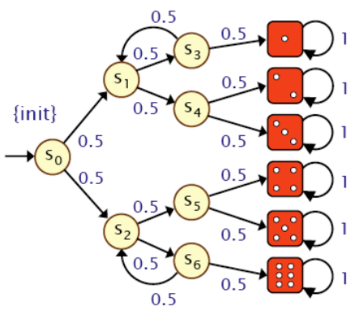

- Consider the event $\Diamond 4$
- We have:

$$Pr(\Diamond 4) = \sum_{s_0 \ldots s_n \in (\Sigma \setminus 4^*)4} \mathbf{P}(s_0 \ldots s_n)$$

- This yields:
$\mathbf{P}(s_0 s_2 s_5 4) + \mathbf{P}(s_0 s_2 s_6 s_2 s_5 4) + \ldots\ldots$
- Or: $\sum_{k=0}^{\infty} \mathbf{P}(s_0 s_2 (s_6 s_2)^k s_5 4)$
- Or: $\frac{1}{8} \cdot \sum_{k=0}^{\infty} (\frac{1}{4})^k$
- Geometric series: $\frac{1}{8} \cdot \frac{1}{1 - \frac{1}{4}} = \frac{1}{8} \cdot \frac{4}{3} = \frac{1}{6}$

For finite state spaces, reachability probabilities can be obtained algorithmically.

Reachability probabilities

Problem statement

Let D be an MC with finite state space Σ, $\sigma \in \Sigma$, and $G \subseteq \Sigma$.

Aim: determine $Pr(\sigma \models \Diamond G) = Pr_\sigma(\Diamond G) = Pr\{\, \pi \in Paths(D_\sigma) \mid \pi \in \Diamond G \,\}$

where D_σ is the MC D with initial state σ.

Characterisation of reachability probabilities

Let variable $x_\sigma = Pr(\sigma \models \Diamond G)$ for any state σ be defined by:

- if $\sigma \notin Pre^*(G)$, then $x_\sigma = 0$
- if $\sigma \in G$, then $x_\sigma = 1$
- otherwise:

$$x_\sigma = \underbrace{\sum_{\tau \in \Sigma \setminus G} \mathbf{P}(\sigma, \tau) \cdot x_\tau}_{\text{reach } G \text{ via } \tau \in \Sigma \setminus G} + \underbrace{\sum_{\gamma \in G} \mathbf{P}(\sigma, \gamma)}_{\text{reach } G \text{ in one step}}$$

$Pre^*(G)$ is the set of states in Σ from which G is reachable, i.e., $\{\, \sigma \in \Sigma \mid Pr(\sigma \models \Diamond G) > 0 \,\}$.

Reachability probabilities: Knuth-Yao's die

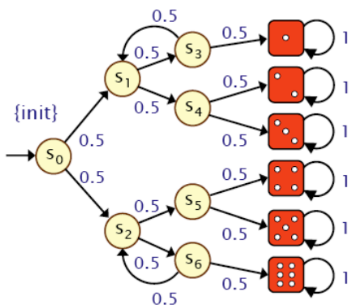

- Consider the event $\Diamond 4$
- The previous characterisation yields:

$x_1 = x_2 = x_3 = x_5 = x_6 = 0$ and $x_4 = 1$

$x_{s_1} = x_{s_3} = x_{s_4} = 0$

$x_{s_0} = \frac{1}{2} x_{s_1} + \frac{1}{2} x_{s_2}$

$x_{s_2} = \frac{1}{2} x_{s_5} + \frac{1}{2} x_{s_6}$

$x_{s_5} = \frac{1}{2} x_5 + \frac{1}{2} x_4$

$x_{s_6} = \frac{1}{2} x_{s_2} + \frac{1}{2} x_6$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Linear equation system

- Let $\Sigma_? = Pre^*(G) \setminus G$, the states that can reach G by > 0 steps
- $\mathbf{A} = \big(\mathbf{P}(\sigma,\tau)\big)_{\sigma,\tau\in\Sigma_?}$, the transition probabilities in $\Sigma_?$
- $\mathbf{b} = \big(b_\sigma\big)_{\sigma\in\Sigma_?}$, the probs to reach G in 1 step, i.e., $b_\sigma = \sum_{\gamma\in G} \mathbf{P}(\sigma,\gamma)$

Theorem

The vector $\mathbf{x} = (x_\sigma)_{\sigma\in\Sigma_?}$ with $x_\sigma = Pr(\sigma \models \Diamond G)$ is the unique solution of the linear equation system:

$$\mathbf{x} \;=\; \mathbf{A}\cdot\mathbf{x} \;+\; \mathbf{b} \quad \text{or, equivalently} \quad (\mathbf{I}-\mathbf{A})\cdot\mathbf{x} \;=\; \mathbf{b}$$

where $\mathbf{I}$ is the identity matrix of cardinality $|\Sigma_?|\cdot|\Sigma_?|$.

Reachability probabilities: Knuth-Yao's die

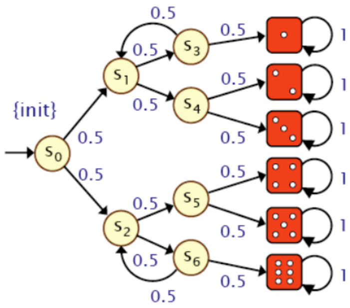

- Consider the event $\Diamond 4$
- $\Sigma_? = \{ s_0, s_2, s_5, s_6 \}$

$$\begin{pmatrix} 1 & -\frac{1}{2} & 0 & 0 \\ 0 & 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & 0 & 1 & 0 \\ 0 & -\frac{1}{2} & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} x_{s_0} \\ x_{s_2} \\ x_{s_5} \\ x_{s_6} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \frac{1}{2} \\ 0 \end{pmatrix}$$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Computing reachability probabilities

Polynomial complexity

Reachability probabilities in finite MCs can be computed in polynomial time.

Overview

First visit probabilities

First visit probabilities

For states $\sigma, \tau \in \Sigma$, let

$$f^{(n)}_{\sigma,\tau} = Pr\{ \text{ first visit to } \tau \text{ after exactly } n \text{ steps from } \sigma \text{ } \}$$

(This differs from the probability to move from τ to σ in n steps.)
We have:

$$\mathbf{P}^n(\sigma, \tau) \;=\; \sum_{\ell=1}^{n} f^{(\ell)}_{\sigma,\tau} \cdot \mathbf{P}^{n-\ell}(\tau, \tau)$$

The probability to reach τ from state σ equals:

$$Pr(\sigma \models \Diamond \tau) \;=\; f_{\sigma,\tau} \;=\; \sum_{n=1}^{\infty} f^{(n)}_{\sigma,\tau}$$

Return probabilities

Return probabilities

For state $\sigma \in \Sigma$, let

$$f^{(n)}_{\sigma} = Pr\{ \text{ first return to } \sigma \text{ after exactly } n \text{ steps } \}$$

We have:

$$f^{(n)}_{\sigma} \;=\; f^{(n)}_{\sigma,\sigma} \;=\; Pr\{\text{first visit to } \sigma \text{ after } n \text{ steps from } \sigma\}.$$

The return probability to state σ equals: $Pr(\sigma \models \Diamond \sigma) \;=\; f_\sigma \;=\; \sum_{n=1}^{\infty} f^{(n)}_{\sigma}$.

Transient and recurrent states

The return probability to σ equals: $Pr(\sigma \models \Diamond \sigma) \;=\; f_\sigma \;=\; \sum_{n=1}^{\infty} f^{(n)}_{\sigma}$.

Transient and recurrent states

State σ is called *recurrent* if $f_\sigma = 1$, i.e., with probability one (aka: almost surely) the MC returns to σ.

State σ is called *transient* otherwise, i.e., if $f_\sigma < 1$. With a positive probability, the MC does not return to a transient state.

Example on the black board.

Null and positive recurrence

Let σ be a recurrent state, i.e., $Pr(\sigma \models \Diamond \sigma) = f_\sigma = 1$.

Mean recurrence time

The mean recurrence time of recurrent state σ equals

$$m_\sigma \;=\; \sum_{n=1}^{\infty} n \cdot f^{(n)}_{\sigma}$$

This is the expected number of steps between two successive visits to σ.

Null and positive recurrent states

State σ is called positive recurrent whenever $m_\sigma < \infty$. Otherwise, state σ is called null recurrent; then $m_\sigma = \infty$.

Example on the black board.

Null and positive recurrence in finite MC

1. Every state in a finite MC is either positive recurrent or transient.
2. At least one state in a finite MC is positive recurrent.
3. A finite MC has no null recurrent states.

Foster's theorem

A countable Markov chain is "non-dissipative"
if almost every infinite path eventually enters
— and remains in — positive recurrent states.

Foster's theorem

A sufficient condition for being non-dissipative is:

$$\sum_{j \geq 0} j \cdot \mathbf{P}(i, j) \;\leq\; i \qquad \text{for all states } i$$

F. Gordon Foster	
Born	24 February 1921 Belfast, United Kingdom
Died	20 December 2010 (aged 89) Dublin, Ireland
Nationality	Irish
Known for	Foster's theorem
Scientific career	
Doctoral advisor	David George Kendall

Frederic Gordon Foster
Markoff chains with an enumerable number of states
and a class of cascade processes
1951

Periodicity and ergodicity

Periodic state

A state σ is called *periodic* if for all n it holds:

$$f_{\sigma}^{(n)} > 0 \quad \text{implies} \quad n = k \cdot d \quad \text{where period } d > 1 \text{and} k \in \mathbb{N}.$$

A state is aperiodic otherwise.

A state is ergodic if it is positive recurrent and aperiodic.
An MC is ergodic if all its states are ergodic.

Example on the black board.

Connected states are of the same "type"

Let σ and τ be mutually reachable from each other. Then:

σ is transient	iff	τ is transient
σ is null-recurrent	iff	τ is null-recurrent
σ is positive recurrent	iff	τ is positive recurrent
σ has period d	iff	τ has period d

Irreducibility

Irreducible

A MC is irreducible if it is strongly connected, i.e., all states are mutually reachable.

Markov's theorem

A finite, irreducible MC is positive recurrent, and if aperiodic then:

$$\mathbf{P}^{\infty} \; = \; \lim_{n \to \infty} \mathbf{P}^{n} = \begin{pmatrix} v \\ \cdot \\ \cdot \\ v \end{pmatrix} \quad \text{where} \quad v = \left(\frac{1}{m_1}, \ldots, \frac{1}{m_k} \right)$$

where $k = |\Sigma|$.

Stationary distribution

Stationary distribution

A probability vector $\mathbf{x}$ satisfying $\mathbf{x} \; = \; \mathbf{x} \cdot \mathbf{P}$ is called a stationary distribution of MC D.

$$x_\sigma \; = \; \sum_{\tau \in \Sigma} x_\tau \cdot \mathbf{P}(\tau, \sigma) \quad \text{iff} \quad \underbrace{x_\sigma \cdot (1 - \mathbf{P}(\sigma, \sigma))}_{\text{outflow of } \sigma} \; = \; \underbrace{\sum_{\tau \neq \sigma} x_\tau \cdot \mathbf{P}(\tau, \sigma)}_{\text{inflow of } \sigma}$$

An irreducible, positive recurrent MC has a unique stationary distribution satisfying $x_\sigma = \frac{1}{m_\sigma}$ for every state σ.

Limiting distribution

Ergodic stochastic matrix

Stochastic matrix $\mathbf{P}$ is called *ergodic* if:

$$\mathbf{P}^{\infty} \; = \; \lim_{n \to \infty} \mathbf{P}^{n} \quad \text{exists and has identical rows}$$

Limiting distribution

If $\mathbf{P}$ is ergodic, then each row of $\mathbf{P}^{\infty}$ equals the limiting distribution.

Limiting = stationary distribution

For ergodic (aka: aperiodic and positive recurrent) MCs, the stationary and limiting distribution are equal.

Overview

Rewards

To reason about resource usage in MCs: use rewards.

MC with rewards

A reward MC is a pair (D, r) with D an MC with state space Σ and $r : \Sigma \to \mathbb{R}$ a function assigning a real reward to each state.

The reward $r(\sigma)$ stands for the reward earned on leaving state σ.

Cumulative reward for reachability

Let $\pi = \sigma_0 \ldots \sigma_n$ be a finite path in (D, r) and $G \subseteq \Sigma$ a set of target states with $\pi \in \Diamond G$. The cumulative reward along π until reaching G is:

$$r_G(\pi) = r(\sigma_0) + \ldots + r(\sigma_{k-1}) \text{ where } \sigma_i \notin G \text{ for all } i < k \text{ and } \sigma_k \in G.$$

If $\pi \notin \Diamond G$, then $r_G(\pi) = 0$.

Expected reward reachability

Expected reward for reachability

The expected reward until reaching $G \subseteq \Sigma$ from $\sigma \in \Sigma$ is:

$$\mathrm{ER}(\sigma, \Diamond G) = \sum_{\pi \models \Diamond G} Pr(\widehat{\pi}) \cdot r_G(\widehat{\pi})$$

where $\widehat{\pi} = \sigma_0 \ldots \sigma_k$ is the shortest prefix of π such that $\sigma_k \in G$ and $\sigma_0 = \sigma$.

Conditional expected reward

Let $\mathrm{ER}(\sigma, \Diamond G \mid \neg \Diamond F)$ be the conditional expected reward until reaching G under the condition that no states in $F \subseteq \Sigma$ are visited.

Expected rewards in finite Markov chains

Polynomial complexity

Expected rewards in finite MCs can be computed in polynomial time.