Experiment design

Bandit problems and Markov decision processes

Christos Dimitrakakis

Chalmers

November 6, 2018

The reinforcement learning problem

Learning how to act optimally in an unknown world through interaction and reinforcement.

- Optimal behaviour implicitly defined through rewards.
- Learning about an unknown world.
- Interactive data collection.

Sequential problems

- Observation x_t.
- Decision a_t.
- Steps $t = 1, \ldots, T$.

General utility function

Utility $U(x_1, x_2, \ldots, x_T, a_1, \ldots, a_T)$

Linear utility function

$$U(x_1, x_2, \ldots, x_T, a_1, \ldots, a_T) = \sum_{t=1}^{T} \rho(x_t, a_t).$$

This is the standard reinforcement learning setting

Sequential problems: full observation

Example 1

- n meteorological stations $\{\mu_i \mid i = 1, \ldots, n\}$
- The i-th station predicts rain $P_{\mu_i}(y_t \mid y_1, \ldots, y_{t-1})$.

- Observation x_t: the predictions of all stations.
- Decision a_t.
- Steps $t = 1, \ldots, T$.

Linear utility function

Reward function is $\rho(x_t, a_t) = \mathbb{I}\{x_t = a_t\}$ simply rewarding correct predictions with utility being

$$U(y_1, y_2, \ldots, y_T, a_1, \ldots, a_T) = \sum_{t=1}^{T} \rho(y_t, a_t),$$

the total number of correct predictions.

The n meteorologists problem is simple, as:

- You always see their predictions, as well as the weather, no matter whether you bike or take the tram (full information)
- Your actions do not influence their predictions (independence events)

In the remainder, we'll see two settings where decisions are made with either partial information or in a dynamical system. Both of these settings can be formalised with Markov decision processes.

Experimental design and Markov decision processes

The following problems

- Shortest path problems.
- Optimal stopping problems.
- Reinforcement learning problems.
- Experiment design (clinical trial) problems
- Advertising.

can be all formalised as Markov decision processes.

Applications

- Robotics.
- Economics.
- Automatic control.
- Resource allocation

Bandit problems

Bandit problems

Applications

- Efficient optimisation.

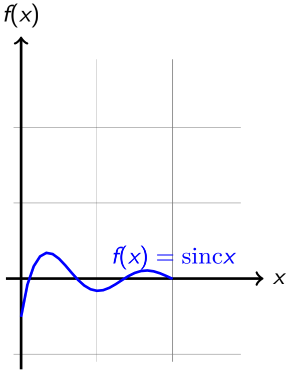

Bandit problems

Applications

- Efficient optimisation.
- Online advertising.

Bandit problems

Applications

- Efficient optimisation.
- Online advertising.
- Clinical trials.

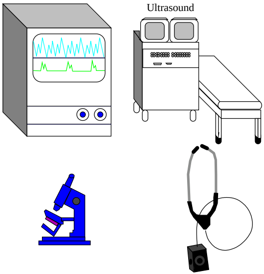

Bandit problems

Applications

- Efficient optimisation.
- Online advertising.
- Clinical trials.
- ROBOT SCIENTIST.

The stochastic n-armed bandit problem

Actions and rewards

- A set of actions $\mathcal{A} = \{1, \ldots, n\}$.
- Each action gives you a random reward with distribution $\mathbb{P}(r_t \mid a_t = i)$.
- The expected reward of the i-th arm is $\rho_i \triangleq \mathbb{E}(r_t \mid a_t = i)$.

Interaction at time t

1. You choose an action $a_t \in \mathcal{A}$.
2. You observe a random reward r_t drawn from the i-th arm.

The utility is the sum of the rewards obtained

$$U \triangleq \sum_t r_t.$$

We must maximise the expected utility, without knowing the values ρ_i.

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^{\pi}(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

Exercise 1

Why should our action depend on the complete history?

A The next reward depends on all the actions we have taken.

B We don't know which arm gives the highest reward.

C The next reward depends on all the previous rewards.

D The next reward depends on the complete history.

E No idea.

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^\pi(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

Example 3 (The expected utility of a uniformly random policy)

If $\mathbb{P}^\pi(a_{t+1} \mid \cdot) = 1/n$ for all t, then

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^\pi(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

Example 3 (The expected utility of a uniformly random policy)

If $\mathbb{P}^\pi(a_{t+1} \mid \cdot) = 1/n$ for all t, then

$$\mathbb{E}^\pi U = \mathbb{E}^\pi \left(\sum_{t=1}^T r_t \right) = \sum_{t=1}^T \mathbb{E}^\pi r_t = \sum_{t=1}^T \sum_{i=1}^n \frac{1}{n} \rho_i = \frac{T}{n} \sum_{i=1}^n \rho_i$$

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^\pi(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

The expected utility of a general policy

$$\mathbb{E}^\pi U = \mathbb{E}^\pi \left(\sum_{t=1}^{T} r_t \right)$$

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^\pi(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

The expected utility of a general policy

$$\mathbb{E}^\pi U = \mathbb{E}^\pi \left(\sum_{t=1}^T r_t \right) = \sum_{t=1}^T \mathbb{E}^\pi (r_t) \tag{1.1}$$

Policy

Definition 2 (Policies)

A policy π is an algorithm for taking actions given the observed history $h_t \triangleq a_1, r_1, \ldots, a_t, r_t$

$$\mathbb{P}^\pi(a_{t+1} \mid h_t)$$

is the probability of the next action a_{t+1}.

The expected utility of a general policy

$$\mathbb{E}^\pi U = \mathbb{E}^\pi \left(\sum_{t=1}^{T} r_t \right) = \sum_{t=1}^{T} \mathbb{E}^\pi(r_t) \tag{1.1}$$

$$= \sum_{t=1}^{T} \sum_{a_t \in \mathcal{A}} \mathbb{E}(r_t \mid a_t) \sum_{h_{t-1}} \mathbb{P}^\pi(a_t \mid h_{t-1}) \mathbb{P}^\pi(h_{t-1})$$

Bernoulli bandits

Decision-theoretic approach

- Assume $r_t \mid a_t = i \sim P_{\theta_i}$, with $\theta_i \in \Theta$.
- Define prior belief ξ_1 on Θ.
- For each step t, select action a_t to maximise

$$\mathbb{E}_{\xi_t}(U_t \mid a_t) = \mathbb{E}_{\xi_t}\left(\sum_{k=1}^{T-t} \gamma^k r_{t+k} \;\middle|\; a_t\right)$$

- Obtain reward r_t.
- Calculate the next belief

$$\xi_{t+1} = \xi_t(\cdot \mid a_t, r_t)$$

How can we implement this?

A simple heuristic for the unknown reward case

Say you keep a running average of the reward obtained by each arm

$$\hat{\theta}_{t,i} = R_{t,i}/n_{t,i}$$

- $n_{t,i}$ the number of times you played arm i
- $R_{t,i}$ the total reward received from i.

Whenever you play $a_t = i$:

$$R_{t+1,i} = R_{t,i} + r_t, \qquad n_{t+1,i} = n_{t,i} + 1.$$

Greedy policy:

$$a_t = \arg\max_i \hat{\theta}_{t,i}.$$

What should the initial values $n_{0,i}$, $R_{0,i}$ be?

Bayesian inference on Bernoulli bandits

- Likelihood: $\mathbb{P}_\theta(r_t = 1) = \theta$.
- Prior: $\xi(\theta) \propto \theta^{\alpha-1}(1-\theta)^{\beta-1}$ (i.e. $\mathcal{B}eta(\alpha, \beta)$).

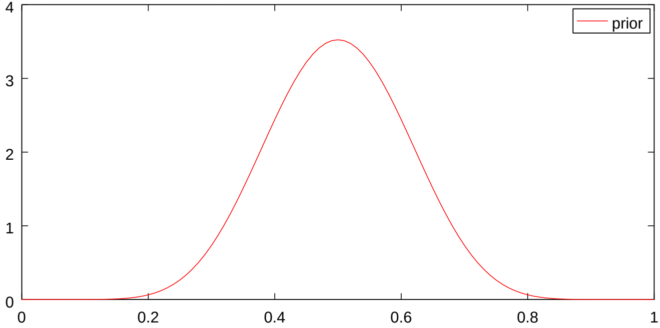

Figure: Prior belief ξ about the mean reward θ.

Bayesian inference on Bernoulli bandits

For a sequence $r = r_1, \ldots, r_n$, $\Rightarrow P_\theta(r) \propto \theta_i^{\#1(\mathrm{r})}(1 - \theta_i)^{\#0(\mathrm{r})}$

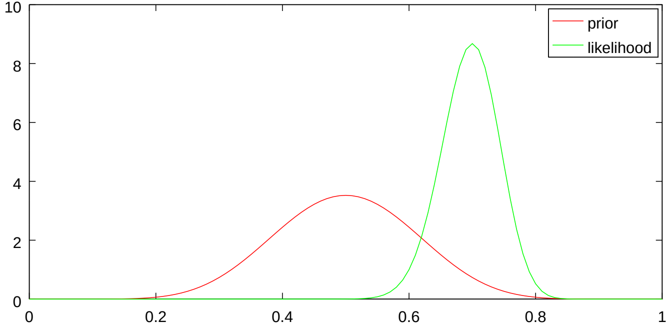

Figure: Prior belief ξ about θ and likelihood of θ for 100 plays with 70 1s.

Bayesian inference on Bernoulli bandits

Posterior: $\mathcal{Beta}(\alpha + \#1(\mathrm{r}), \beta + \#0(\mathrm{r}))$.

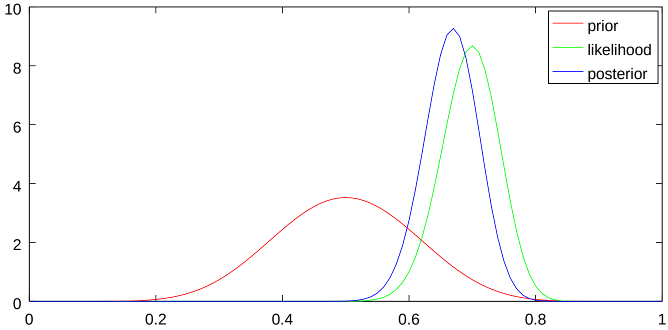

Figure: Prior belief $\xi(\theta)$ about θ, likelihood of θ for the data r, and posterior belief $\xi(\theta \mid r)$

Bernoulli example.

Consider n Bernoulli distributions with unknown parameters θ_i ($i = 1, \ldots, n$) such that

$$r_t \mid a_t = i \sim \mathcal{B}ernoulli(\theta_i), \qquad \mathbb{E}(r_t \mid a_t = i) = \theta_i. \tag{1.2}$$

Our belief for each parameter θ_i is $\mathcal{B}eta(\alpha_i, \beta_i)$, with density $f(\theta \mid \alpha_i, \beta_i)$ so that

$$\xi(\theta_1, \ldots, \theta_n) = \prod_{i=1}^{n} f(\theta_i \mid \alpha_i, \beta_i). \qquad \text{(a priori independent)}$$

$$N_{t,i} \triangleq \sum_{k=1}^{t} \mathbb{I}\left\{a_k = i\right\}, \qquad \hat{r}_{t,i} \triangleq \frac{1}{N_{t,i}} \sum_{k=1}^{t} r_t \, \mathbb{I}\left\{a_k = i\right\}$$

Then, the posterior distribution for the parameter of arm i is

$$\xi_t = \mathcal{B}eta(\alpha_i^t, \beta_i^t), \qquad \alpha_i^t = \alpha_i + N_{t,i}\hat{r}_{t,i}\ , \ \ \beta_i^t = \beta_i N_{t,i}(1 - \hat{r}_{t,i})).$$

Since $r_t \in \{0, 1\}$ there are $O((2n)^T)$ possible belief states for a T-step bandit problem.

Belief states

- The state of the decision-theoretic bandit problem is the state of our belief.
- A sufficient statistic is the number of plays and total rewards.
- Our belief state ξ_t is described by the priors α, β and the vectors

$$N_t = (N_{t,1}, \ldots, N_{t,i}) \tag{1.3}$$

$$\hat{r}_t = (\hat{r}_{t,1}, \ldots, \hat{r}_{t,i}). \tag{1.4}$$

- The next-state probabilities are defined as:

$$\mathbb{P}_{\xi_t}(r_t = 1 \mid a_t = i) = \frac{\alpha_i^t}{\alpha_i^t + \beta_i^t}$$

  as ξ_{t+1} is a deterministic function of ξ_t, r_t and a_t
- Optimising this results in a Markov decision process.

Markov process

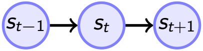

Definition 3 (Markov Process – or Markov Chain)

The sequence $\{s_t \mid t = 1, \dots\}$ of random variables $s_t : \Theta \to \mathcal{S}$ is a Markov process if

$$\mathbb{P}(s_{t+1} \mid s_t, \dots, s_1) = \mathbb{P}(s_{t+1} \mid s_t). \tag{1.5}$$

- s_t is state of the Markov process at time t.
- $\mathbb{P}(s_{t+1} \mid s_t)$ is the transition kernel of the process.

The state of an algorithm

Observe that the R, n vectors of our greedy bandit algorithm form a Markov process. They also summarise our belief about which arm is the best.

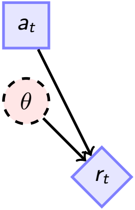

Figure: The basic bandit MDP. The decision maker selects a_t, while the parameter θ of the process is hidden. It then obtains reward r_t. The process repeats for $t = 1, \ldots, T$.

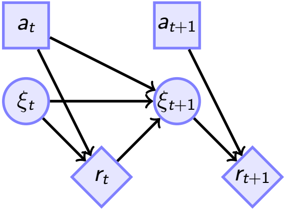

Figure: The decision-theoretic bandit MDP. While θ is not known, at each time step t we maintain a belief ξ_t on Θ. The reward distribution is then defined through our belief.

Backwards induction (Dynamic programming)

for $n = 1, 2, \ldots$ and $s \in \mathcal{S}$ **do**

$$\mathbb{E}(U_t \mid \xi_t) = \max_{a_t \in \mathcal{A}} \mathbb{E}(r_t \mid \xi_t, a_t) + \gamma \sum_{\xi_{t+1}} \mathbb{P}(\xi_{t+1} \mid \xi_t, a_t)\, \mathbb{E}(U_{t+1} \mid \xi_{t+1})$$

end for

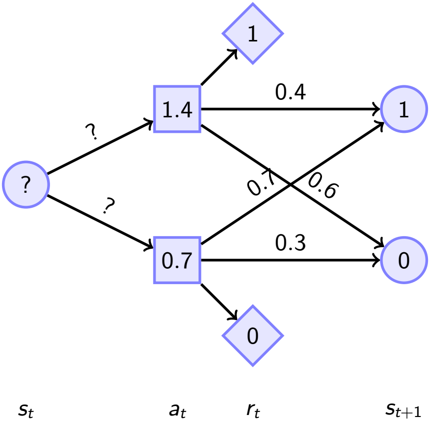

Exercise 1

What is the value $v_t(s_t)$ of the first state?

- A 1.4
- B 1.05
- C 1.0
- D 0.7
- E 0

Backwards induction (Dynamic programming)

for $n = 1, 2, \ldots$ and $s \in \mathcal{S}$ **do**

$$\mathbb{E}(U_t \mid \xi_t) = \max_{a_t \in \mathcal{A}} \mathbb{E}(r_t \mid \xi_t, a_t) + \gamma \sum_{\xi_{t+1}} \mathbb{P}(\xi_{t+1} \mid \xi_t, a_t) \, \mathbb{E}(U_{t+1} \mid \xi_{t+1})$$

end for

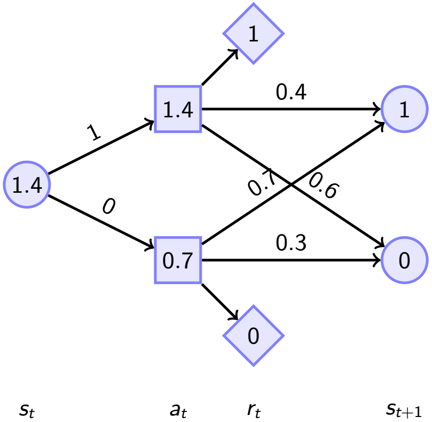

Exercise 1

What is the value $v_t(s_t)$ of the first state?

- A 1.4
- B 1.05
- C 1.0
- D 0.7
- E 0

Heuristic algorithms for the n-armed bandit problem

Algorithm 1 UCB1

Input $\mathcal{A}$

$\hat{\theta}_{0,i} = 1, \forall i$

for $t = 1, \ldots$ **do**

$a_t = \arg\max_{i \in \mathcal{A}} \left\{ \hat{\theta}_{t-1,i} + \sqrt{\frac{2 \ln t}{N_{t-1,i}}} \right\}$

$r_t \sim P_\theta(r \mid a_t)$ `// play action and get reward // update model`

$N_{t,a_t} = N_{t-1,a_t} + 1$

$\hat{\theta}_{t,a_t} = [N_{t-1,a_t} \theta_{t-1,a_t} + r_t] / N_{t,a_t}$

$\forall i \neq a_t, \; N_{t,i} = N_{t-1,i}, \; \hat{\theta}_{t,i} = \hat{\theta}_{t-1,i}$

end for

Algorithm 2 Thompson sampling

Input $\mathcal{A}, \xi_0$

for $t = 1, \ldots$ **do**

$\hat{\theta} \sim \xi_{t-1}(\theta)$

$a_t \in \arg\max_a \mathbb{E}_{\hat{\theta}}[r_t \mid a_t = a]$.

$r_t \sim P_\theta(r \mid a_t)$ `// play action and get reward // update model`

$\xi_t(\theta) = \xi_{t-1}(\theta \mid a_t, r_t)$.

end for

Example 4 (Clinical trials)

Consider an example where we have some information x_t about an individual patient t, and we wish to administer a treatment a_t. For whichever treatment we administer, we can observe an outcome y_t. Our goal is to maximise expected utility.

Definition 5 (The contextual bandit problem.)

At time t,

- We observe $x_t \in \mathcal{X}$.
- We play $a_t \in \mathcal{A}$.
- We obtain $r_t \in \mathbb{R}$ with $r_t \mid a_t = a, x_t = x \sim P_\theta(r \mid a, x)$.

Example 6 (The linear bandit problem)

- $\mathcal{A} = [n]$, $\mathcal{X} = \mathbb{R}^k$, $\theta = (\theta_1, \ldots, \theta_n)$, $\theta_i \in \mathbb{R}^k$, $r \in \mathbb{R}$.
- $r \sim \mathcal{N}(\theta_a^\top x), 1)$

Example 7 (A clinical trial example)

- $\mathcal{A} = [n]$, $\mathcal{X} = \mathbb{R}^k$, $\theta = (\theta_1, \ldots, \theta_n)$, $\theta_i \in \mathbb{R}^k$, $y \in \{0, 1\}$.
- $y \sim \mathcal{Bernoulli}(1/(1 + exp[-(\theta_a^\top x)^2])$.
- $r = U(a, y)$.