Bayesian Updating: Discrete Priors: 18.05 Spring 2014

Jeremy Orloff and Jonathan Bloom

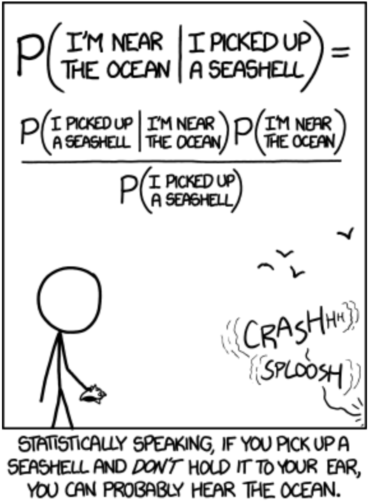

http://xkcd.com/1236/

Learning from experience

Which treatment would you choose?

1. Treatment 1: cured 100% of patients in a trial.
2. Treatment 2: cured 95% of patients in a trial.
3. Treatment 3: cured 90% of patients in a trial.

Learning from experience

Which treatment would you choose?

1. Treatment 1: cured 100% of patients in a trial.
2. Treatment 2: cured 95% of patients in a trial.
3. Treatment 3: cured 90% of patients in a trial.

Which treatment would you choose?

1. Treatment 1: cured 3 out of 3 patients in a trial.
2. Treatment 2: cured 19 out of 20 patients treated in a trial.
3. Standard treatment: cured 90000 out of 100000 patients in clinical practice.

Which die is it?

- Jon has a bag containing dice of two types: 4-sided and 20-sided.
- Suppose he picks a die at random and rolls it.
- Based on what Jon rolled which type would you guess he picked?

Which die is it?

- Jon has a bag containing dice of two types: 4-sided and 20-sided.
- Suppose he picks a die at random and rolls it.
- Based on what Jon rolled which type would you guess he picked?

Suppose you find out that the bag contained one 4-sided die and one 20-sided die. Does this change your guess?

Which die is it?

- Jon has a bag containing dice of two types: 4-sided and 20-sided.
- Suppose he picks a die at random and rolls it.
- Based on what Jon rolled which type would you guess he picked?

Suppose you find out that the bag contained one 4-sided die and one 20-sided die. Does this change your guess?

Suppose you find out that the bag contained one 4-sided die and 100 20-sided dice. Does this change your guess?

Learn from experience: coins

Three types of coins:

- Type A coins are fair, with probability .5 of heads
- Type B coins are bent, with probability .6 of heads
- Type C coins are bent, with probability .9 of heads

A box contains 4 coins:
2 of type A
1 of type B
1 of type C.

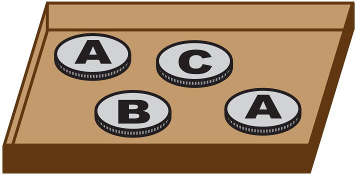

I pick one at random flip it and get heads.

1. What was the prior (before flipping) probability the coin was of each type?

2. What is the posterior (after flipping) probability for each type?

3. What was learned by flipping the coin?

Tabular solution

$\mathcal{H}$ = hypothesis: the coin is of type A (or B or C)

$\mathcal{D}$ = data: I flipped once and got heads

hypothesis	prior	likelihood	unnormalized posterior	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})P(\mathcal{H})$	$P(\mathcal{H}\|\mathcal{D})$
A	.5	.5	.25	.4
B	.25	.6	.15	.24
C	.25	.9	.225	.36
total	1		.625	1

Total probability: $P(\mathcal{D})$ = sum of unnorm. post. column = .625

Bayes theorem: $P(\mathcal{H}|\mathcal{D}) = \dfrac{P(\mathcal{D}|\mathcal{H})P(\mathcal{H})}{P(\mathcal{D})}$

Board Question: Dice

Five dice: 4-sided, 6-sided, 8-sided, 12-sided, 20-sided.

Jon picks one at random and, without showing it to you, rolls it and reports a 13.

1. Make a table and compute the posterior probabilities that the chosen die is each of the five dice.

2. Same question if he rolls a 5.

3. Same question if he rolls a 9.

(Keep the tables for 5 and 9 handy! Do not erase!)

Tabular solution

$\mathcal{D}$ = 'rolled a 13'

hypothesis	prior	likelihood	unnormalized posterior	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}\vert\mathcal{H})$	$P(\mathcal{D}\vert\mathcal{H})P(\mathcal{H})$	$P(\mathcal{H}\vert\mathcal{D})$
$\mathcal{H}_4$	$1/5$	0	0	0
$\mathcal{H}_6$	$1/5$	0	0	0
$\mathcal{H}_8$	$1/5$	0	0	0
$\mathcal{H}_{12}$	$1/5$	0	0	0
$\mathcal{H}_{20}$	$1/5$	$1/20$	$1/100$	1
total	1		$1/100$	1

Tabular solution

$\mathcal{D}$ = 'rolled a 5'

hypothesis	prior	likelihood	unnormalized posterior	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})P(\mathcal{H})$	$P(\mathcal{H}\|\mathcal{D})$
$\mathcal{H}_4$	$1/5$	0	0	0
$\mathcal{H}_6$	$1/5$	$1/6$	$1/30$	0.392
$\mathcal{H}_8$	$1/5$	$1/8$	$1/40$	0.294
$\mathcal{H}_{12}$	$1/5$	$1/12$	$1/60$	0.196
$\mathcal{H}_{20}$	$1/5$	$1/20$	$1/100$	0.118
total	1		.085	1

Tabular solution

$\mathcal{D}$ = 'rolled a 9'

hypothesis	prior	likelihood	unnormalized posterior	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})P(\mathcal{H})$	$P(\mathcal{H}\|\mathcal{D})$
$\mathcal{H}_4$	$1/5$	0	0	0
$\mathcal{H}_6$	$1/5$	0	0	0
$\mathcal{H}_8$	$1/5$	0	0	0
$\mathcal{H}_{12}$	$1/5$	$1/12$	$1/60$	0.625
$\mathcal{H}_{20}$	$1/5$	$1/20$	$1/100$	0.375
total	1		.0267	1

Iterated Updates

Suppose Jon rolled a 5 and then a 9.

Update in two steps:

- First for the 5
- Then update the update for the 9.

Tabular solution

$\mathcal{D}_1$ = 'rolled a 5'
$\mathcal{D}_2$ = 'rolled a 9'
Unnorm. posterior$_1$ = likelihood$_1 \times$ prior.
Unnorm. posterior$_2$ = likelihood$_2 \times$ unnorm. posterior$_1$

hyp.	prior	likel. 1	unnorm. post. 1	likel. 2	unnorm. post. 2	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}_1\|\mathcal{H})$	$***$	$P(\mathcal{D}_2\|\mathcal{H})$	$***$	$P(\mathcal{H}\|\mathcal{D}_1, \mathcal{D}_2)$
$\mathcal{H}_4$	1/5	0	0	0	0	0
$\mathcal{H}_6$	1/5	1/6	1/30	0	0	0
$\mathcal{H}_8$	1/5	1/8	1/40	0	0	0
$\mathcal{H}_{12}$	1/5	1/12	1/60	1/12	1/720	0.735
$\mathcal{H}_{20}$	1/5	1/20	1/100	1/20	1/2000	0.265
total	1				0.0019	1

Board Question

Suppose Jon rolled a 9 and then a 5.

1. Do the Bayesian update in two steps:
   First update for the 9.
   Then update the update for the 5.

2. Do the Bayesian update in one step
   The data is $\mathcal{D} =$ '9 followed by 5'

Tabular solution: two steps

$\mathcal{D}_1$ = 'rolled a 9'
$\mathcal{D}_2$ = 'rolled a 5'
Unnorm. $\text{posterior}_1 = \text{likelihood}_1 \times$ prior.
Unnorm. $\text{posterior}_2 = \text{likelihood}_2 \times$ unnorm. posterior_1

hyp.	prior	likel. 1	unnorm. post. 1	likel. 2	unnorm. post. 2	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}_1\|\mathcal{H})$	$***$	$P(\mathcal{D}_2\|\mathcal{H})$	$***$	$P(\mathcal{H}\|\mathcal{D}_1, \mathcal{D}_2)$
$\mathcal{H}_4$	1/5	0	0	0	0	0
$\mathcal{H}_6$	1/5	0	0	1/6	0	0
$\mathcal{H}_8$	1/5	0	0	1/8	0	0
$\mathcal{H}_{12}$	1/5	1/12	1/60	1/12	1/720	0.735
$\mathcal{H}_{20}$	1/5	1/20	1/100	1/20	1/2000	0.265
total	1				0.0019	1

Tabular solution: one step

$\mathcal{D}$ = 'rolled a 9 then a 5'

hypothesis	prior	likelihood	unnormalized posterior	posterior
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})$	$P(\mathcal{D}\|\mathcal{H})P(\mathcal{H})$	$P(\mathcal{H}\|\mathcal{D})$
$\mathcal{H}_4$	$1/5$	0	0	0
$\mathcal{H}_6$	$1/5$	0	0	0
$\mathcal{H}_8$	$1/5$	0	0	0
$\mathcal{H}_{12}$	$1/5$	$1/144$	$1/720$	0.735
$\mathcal{H}_{20}$	$1/5$	$1/400$	$1/2000$	0.265
total	1		0.0019	1

Board Question: probabilistic prediction

Along with finding posterior probabilities of hypotheses. We might want to make posterior predictions about the next roll.

With the same setup as before let:
$\mathcal{D}_1 =$ result of first roll
$\mathcal{D}_2 =$ result of second roll

(a) Find $P(\mathcal{D}_1 = 5)$.

(b) Find $P(\mathcal{D}_2 = 4 | \mathcal{D}_1 = 5)$.

Solution

$\mathcal{D}_1$ = 'rolled a 5'

$\mathcal{D}_2$ = 'rolled a 4'

hyp.	prior	likel. 1	unnorm. post. 1	post. 1	likel. 2	post. 1 × likel. 2
$\mathcal{H}$	$P(\mathcal{H})$	$P(\mathcal{D}_1\|\mathcal{H})$	$***$	$P(\mathcal{H}\|\mathcal{D}_1)$	$P(\mathcal{D}_2\|\mathcal{H},\mathcal{D}_1)$	$P(\mathcal{D}_2\|\mathcal{H},\mathcal{D}_1)P(\mathcal{H}\|\mathcal{D}_1)$
$\mathcal{H}_4$	1/5	0	0	0	$*$	0
$\mathcal{H}_6$	1/5	1/6	1/30	.392	1/6	$.392 \cdot 1/6$
$\mathcal{H}_8$	1/5	1/8	1/40	.294	1/8	$.294 \cdot 1/40$
$\mathcal{H}_{12}$	1/5	1/12	1/60	.196	1/12	$.196 \cdot 1/12$
$\mathcal{H}_{20}$	1/5	1/20	1/100	.118	1/20	$.118 \cdot 1/20$
total	1		.085	1		0.124

The law of total probability tells us $P(\mathcal{D}_1)$ is the sum of the unnormalized posterior 1 column in the table: $\boxed{P(\mathcal{D}_1) = .085}$.

The law of total probability tells us $P(\mathcal{D}_2|\mathcal{D}_1)$ is the sum of the last column in the table: $\boxed{P(\mathcal{D}_2|\mathcal{D}_1) = 0.124}$