CS354

Nathan Sprague

September 27, 2016

Probability Notation

- Probability Functions/Distributions:
  - $P(A)$ is a function that maps from all possible values of A to the probability of the corresponding event.
  - Examples:
    - $P(A = true) = .9$
      $P(A = false) = .1$
    - $P(B = red) = .8$
      $P(B = blue) = .1$
      $P(B = green) = .1$
- $P(A)$ is also referred to as a prior probability.

Sample Spaces and Joint Probability Distributions

- Sample space is the set of all possible outcomes.
- The full joint probability distribution assigns a probability to each element of the sample space:
  - S - Squished, U - Under falling Piano

S	U	$P(S, U)$
T	T	.008
T	F	.002
F	T	.001
F	F	.989

Conditional Probability

- $P(A \mid B)$ Expresses the probability of assignments to A given assignments to B.
  - $P(SQUISHED = true) = .01$
  - $P(SQUISHED = true \mid UNDER\_PIANO = true) \approx .89$

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}$$

Bayes Rule

$$P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$$

Very handy for updating our beliefs on the basis of evidence.

Bayes Rule Example

- Robot is in a simple four room maze, rooms are labeled a-d.
- Initially, we think he is most likely to be in the left half, $P(X = a) = .4$, $P(X = b) = .4$, ...

a	b	c	d
.4	.4	.1	.1

Bayes Rule Example

- Robot has a sensor designed to tell him what room he is in.
- Sensor is not perfect: only 80% likely to report he is in the correct room. 20% of the time the sensor is off by one. (Errors at the edge wrap around.)
- Distribution of sensor readings when robot is in *a*:

a	b	c	d
.8	.1	0	.1

- In probability notation, where X is the position and Z is sensor reading.
  - $P(Z = a \mid X = a) = .8$
    $P(Z = b \mid X = a) = .1$
    $P(Z = c \mid X = a) = 0$
    $P(Z = d \mid X = a) = .1$

Bayes Rule Example

- Given that we have a sensor model, Baye's rule enables us to update our prior beliefs based on sensor input:

$$P(X \mid Z) = \frac{P(Z \mid X)P(X)}{P(Z)}$$

Bayes Rule Example

- Let's calculate $P(X = a \mid Z = b)$

$$P(X = a \mid Z = b) = \frac{P(Z = b \mid X = a)P(X = a)}{P(Z = b)}$$

- $P(Z = b \mid X = a) = .1$ (From our sensor model)
- $P(X = a) = .4$ (Our prior)
- $P(Z = b)$ (??)

Bayes Rule Example

To calculate $P(Z = b)$, we can use the total probability theorem:

$$P(Z) = \sum_{i}^{N} P(X = x_i) P(Z \mid X = x_i)$$

We can also treat $P(Z)$ as an unknown constant,

$$P(X \mid Z) = \eta P(Z \mid X) P(X)$$

and set it to whatever value makes $P(X \mid Z)$ sum to 1. The two approaches are equivalent.

Bayes Rule Example

Back to work...

$$P(X = a \mid Z = b) = \frac{P(Z = b \mid X = a)P(X = a)}{P(Z = b)}$$

$$= \eta \times .1 \times .4 = .04\eta$$

Similarly:

$$P(X = b \mid Z = b) = \eta \times .8 \times .4 = .32\eta$$

$$P(X = c \mid Z = b) = \eta \times .1 \times .1 = .01\eta$$

$$P(X = d \mid Z = b) = \eta \times 0 \times .1 = 0$$

Bayes Rule Example

Therefore, after our sensor reading, the updated distribution over possible robot locations is:

a	b	c	d
$.04\eta$	$.32\eta$	$.01\eta$	0

We know the robot is *somewhere*, so we know that:

$$.04\eta + .32\eta + .01\eta = 1$$

$$\eta = \frac{1}{.04 + .32 + .01} = 1/.37 \approx 2.70$$

Bayes Rule Example

Finally, we have an updated belief about the robot location:

a	b	c	d
.108	.865	.027	0

We may use this as our new prior, and incorporate additional sensor readings.