Bayes Rule

Intuitive Description of Bayes

Initial belief plus new evidence = new and improved belief.

$$P(H|D) = \frac{P(H)*P(D|H)}{P(D)}$$

Bayes Rule Components

Hypotheses – beliefs you have about the world

Probabilities – strength of belief that each hypothesis is true

Priors – assumptions about the world (before you see data)

Likelihoods – your assumptions about how the data you see was generated by hypotheses

Steps in computing Bayes

1. Compute the prior for each hypothesis, P(H)

2. Compute the likelihood of the data under each hypothesis, P(D|H)

3. Multiply prior times likelihood to get P(H) P(D | H)

4. Re-normalize these over all hypotheses (e.g. scale each P(H) P(D|H) value so that together they sum to 1; equivalently, divide each by the sum of all of them).

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.

Test A is 98 percent reliable.

- e.g., 98 out of 100 people who have cancer will test positive, and 98 out of 100 who are healthy will test negative.

Your test is positive. How probable is that you have cancer?

Modified from: http://blogs.scientificamerican.com/cross-check/bayes-s-theorem-what-s-the-big-deal/

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.
Test A is 98 percent reliable.
Your test is positive. How probable is that you have cancer?

1. Compute the prior for each hypothesis, P(H)
2. Compute the likelihood of the data under each hypothesis, P(D|H)
3. Multiply prior times likelihood to get P(H) P(D | H)
4. Re-normalize these over all hypotheses

Hypothesis	P(H)	P(D|H)	P(H)P(D|H)	P(H|D)
You have Cancer				
You don’t have Cancer				

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.
Test A is 98 percent reliable.
Your test is positive. How probable is that you have cancer?

1. Compute the prior for each hypothesis, P(H)
2. Compute the likelihood of the data under each hypothesis, P(D|H)
3. Multiply prior times likelihood to get P(H) P(D | H)
4. Re-normalize these over all hypotheses

Hypothesis	P(H)	P(D|H)	P(H)P(D|H)	P(H|D)
You have Cancer	0.01			
You don’t have Cancer	0.99			

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.
Test A is 98 percent reliable.
Your test is positive. How probable is that you have cancer?

1. Compute the prior for each hypothesis, P(H)
2. Compute the likelihood of the data under each hypothesis, P(D|H)
3. Multiply prior times likelihood to get P(H) P(D | H)
4. Re-normalize these over all hypotheses

Hypothesis	P(H)	P(D|H)	P(H)P(D|H)	P(H|D)
You have Cancer	0.01	0.98		
You don't have Cancer	0.99	0.02		

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.
Test A is 98 percent reliable.
Your test is positive. How probable is that you have cancer?

1. Compute the prior for each hypothesis, P(H)
2. Compute the likelihood of the data under each hypothesis, P(D|H)
3. Multiply prior times likelihood to get P(H) P(D | H)
4. Re-normalize these over all hypotheses

Hypothesis	P(H)	P(D|H)	P(H)P(D|H)	P(H|D)
You have Cancer	0.01	0.98	0.0098	
You don’t have Cancer	0.99	0.02	0.0198	
			P(D) = .0296	

Plug-in Example:

Cancer A is estimated to occur in one percent of people your age.
Test A is 98 percent reliable.
Your test is positive. How probable is that you have cancer?

1. Compute the prior for each hypothesis, P(H)
2. Compute the likelihood of the data under each hypothesis, P(D|H)
3. Multiply prior times likelihood to get P(H) P(D | H)
4. Re-normalize these over all hypotheses

Hypothesis	P(H)	P(D|H)	P(H)P(D|H)	P(H|D)
You have Cancer	0.01	0.98	0.0098	0.0098/0.0296 = 0.33
You don't have Cancer	0.99	0.02	0.0198	0.0198/0.0296 = 0.66
			P(D) = .0296	

More Examples

1. Your friend pulls out a coin and flips it 4 times, only to get Tails each time. What is the probability that the coin is:
   a. 1 head-1tail
   b. 2 tail
   c. weighted 70H-30T

   1A. Assume a uniform prior.
   1B. Assume a 50-25-25 prior.

2. Your friend pulls out a coin and flips it to get heads. Your posterior probability for each hypothesis is:
   a. H1 (1h1t) = .7
   b. H2 (2h) = .3

What is your prior for each hypothesis?

3A. There is a box that makes noises when it is given certain numbers. You are shown that the box responds to 2 and 6. Assuming that you have a uniform prior over the hypotheses and using a size principle likelihood, what is your posterior probability for the following hypotheses after observing these 2 events?

a. H1 = even numbers {2, 4, 6, …, 100}
b. H2 = all numbers {1, 2, 3, …, 100}
c. H3 = powers of 2 {1, 2, 4, 8, 16, 32, 64}
d. H4 = only numbers 1-10

3B. Pretend that instead, you saw that the box responds to 1 and 6. Assuming that there is a uniform prior over the hypotheses and a size principle likelihood, what is your posterior probability for the following hypotheses after observing these 2 events?