Information Cascades

CMPS 17

Theresa Migler-VonDollen

Information cascades

Suppose you are visiting a city you have never visited before and you want to go to dinner.
You yelp in the neighborhood that you are in and see a restaurant the has great reviews.
BUT the restaurant is almost empty.
AND the restaurant next door is packed, though it isn't listed on yelp.
What should you do?

Information cascades

Majority red?, majority blue?
"Herding" experiment.
10 volunteers please.

Information cascades

Suppose the first two people in the experiment drew blue.
Even if the third person drew red, it would be better to guess blue.
And similarly for every person after the third.
How can we formalize this?
Bayes' Law!

Bayes' law

- Thomas Bayes (1701 - 1761) was an English statistician, philosopher and Presbyterian minister.

- He didn't actually publish his work, Richard Price did.
- The law (theorem) relates current to prior belief.

Bayes' law

Definition

Given any event A, denote its probability of occurring by $Pr[A]$.

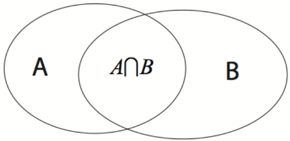

In this diagram, $A \cap B$ denotes that both event A and event B occur.

What is the probability that A occurs given that some other event B occurs?

Definition

The *conditional probability* of A given B, $Pr[A|B]$, is the probability of A given that B has occurred.

Bayes' law

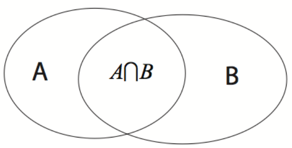

$Pr[A|B] = \frac{Pr[A \cap B]}{Pr[B]}$

Symmetrically, $Pr[B|A] = \frac{Pr[B \cap A]}{Pr[A]} = \frac{Pr[A \cap B]}{Pr[A]}$

$Pr[A|B] \times Pr[B] = Pr[A \cap B] = Pr[B|A] \times Pr[A]$.

Bayes' law

$Pr[A|B] = \frac{Pr[A] \times Pr[B|A]}{Pr[B]}$

Bayes' law - example

- Suppose you have been tested positive for a disease; what is the probability that you actually have the disease?
  - It depends on the accuracy and sensitivity of the test
  - And on the background (prior) probability of the disease.
- Let $Pr[\text{Test=Positive} \mid \text{Disease=True}] = 0.95$
- So $Pr[\text{Test=Negative} \mid \text{Disease=True}] = 0.05$
- Let $Pr[\text{Test=Positive} \mid \text{Disease=False}] = 0.05$
- Suppose that the disease is rare, $Pr[\text{Disease=True}] = 0.01$

What is the probability of having the disease given that you tested positive?
What is $Pr[\text{Disease =True} \mid \text{Test = Positive}]$?

Bayes' law - example

$$Pr[\text{Disease} = \text{True} \mid \text{Test} = \text{Positive}] =$$

$$\frac{Pr[\text{Disease} = \text{True}] \times Pr[\text{Test} = \text{Positive} \mid \text{Disease} = \text{True}]}{Pr[\text{Test} = \text{Positive}]} =$$

$$\frac{0.01 \times 0.95}{Pr[\text{Test} = \text{Positive}]}$$

$$Pr[\text{Test} = \text{Positive}] =$$

$$Pr[\text{Test} = \text{Positive} \mid \text{Disease} = \text{True}] \times Pr[\text{Disease} = \text{True}] +$$

$$Pr[\text{Test} = \text{Positive} \mid \text{Disease} = \text{False}] \times Pr[\text{Disease} = \text{False}]$$

$$= 0.95 \times 0.01 + 0.05 \times 0.99 = 0.0590$$

Bayes' law - example

$$Pr[\text{Disease} = \text{True} \mid \text{Test} = \text{Positive}]$$
$$= \frac{Pr[\text{Disease} = \text{True}] \times Pr[\text{Test} = \text{Positive} \mid \text{Disease} = \text{True}]}{Pr[\text{Test} = \text{Positive}]}$$
$$= \frac{0.95 \times 0.01}{0.0590}$$
$$= 0.161$$

The probability of having the disease, given that you tested positive, is approximately 16%.

Bayes' law - herding experiment

A student should choose blue if:

Pr[majority-blue | what she has seen or heard] $> .5$.

In the beginning:

Pr[majority-blue] $=$ Pr[majority-red] $= .5$.

Also,

Pr[blue|majority-blue] $=$ Pr[red|majority-red] $= .67$.

Bayes' law - herding experiment

Suppose the first student draws a blue. What is the probability that the majority is blue?:

$$Pr[\text{majority-blue} \mid \text{blue}] = \frac{Pr[\text{majority-blue}] \times Pr[\text{blue} \mid \text{majority-blue}]}{Pr[\text{blue}]}$$

and

$$\begin{aligned}
&Pr[\text{blue}] = \\
&Pr[\text{majority-blue}] \times Pr[\text{blue} \mid \text{majority-blue}] + \\
&Pr[\text{majority-red}] \times Pr[\text{blue} \mid \text{majority-red}] \\
&= .5 \times .67 + .5 \times .33 = .5
\end{aligned}$$

Bayes' law - herding experiment

If the first student draws blue, the likelihood that the majority is blue is 67%:

$$Pr[\text{majority-blue} \mid \text{blue}] =$$
$$\frac{Pr[\text{majority-blue}] \times Pr[\text{blue} \mid \text{majority-blue}]}{Pr[\text{blue}]} =$$
$$\frac{.5 \times .67}{.5} = .67$$

So, if the first student draws blue, she should guess blue. Similarly for the second student.

Bayes' law - herding experiment

Now suppose that the third student draws red. What is the probability that the majority is blue?

$$Pr[\text{majority-blue} \mid \text{blue, blue, red}] = \frac{Pr[\text{majority-blue}] \times Pr[\text{blue, blue, red} \mid \text{majority-blue}]}{Pr[\text{blue, blue, red}]}$$

$$Pr[\text{blue, blue, red} \mid \text{majority-blue}] = .67 \times .67 \times .33 \approx .15$$

Bayes' law - herding experiment

$$Pr[\text{blue, blue, red}] =$$
$$Pr[\text{ majority-blue}] \times Pr[\text{blue, blue, red} \mid \text{majority-blue}] +$$
$$Pr[\text{ majority-red}] \times Pr[\text{blue, blue, red} \mid \text{majority-red}]$$
$$= .5 \times .67 \times .67 \times .33 + .5 \times .33 \times .33 \times .67 \approx .11$$

Bayes' law - herding experiment

Therefore, when the third student draws a red, the probability of the majority being blue is 67%

$$Pr[\text{majority-blue} \mid \text{blue, blue, red}] =$$
$$\frac{Pr[\text{majority-blue}] \times Pr[\text{blue, blue, red} \mid \text{majority-blue}]}{Pr[\text{blue, blue, red}]} \approx$$
$$\frac{.5 \times .15}{.11} \approx .67$$

General cascade model

There is a group of people sequentially making decisions about an option.
We say that the two choices are to *accept* or *reject* an option.

Cascade model

- *States of the World*: At the very beginning, we assume that the world is randomly placed in one of two possible *states*:
  - A state where the option is a good idea, G.
  - A state where the option is a bad idea, B.
- Suppose that the world is placed in state G with probability p and is placed in state B with probability $1 - p$.
  - $Pr[G] = p$
  - $Pr[B] = 1 - p$

General cascade model

There is a group of people sequentially making decisions about an option.
We say that the two choices are to *accept* or *reject* an option.

Cascade model

- *Payoffs*: Each individual receives a payoff based on her decision to accept or reject the option.
  - If the individual chooses to reject the option then her payoff is 0.
  - If the individual chooses to accept the option:
    - If the option is a good idea, then the payoff obtained from accepting is positive, $v_g > 0$.
    - If the option is a bad idea, then the payoff obtained from accepting is negative, $v_b < 0$.

General cascade model

There is a group of people sequentially making decisions about an option.
We say that the two choices are to *accept* or *reject* an option.

Cascade model

- *Signals*: We need to model the effect of private information.
  - Before any decisions are made, each individual gets a *private signal* that provides information about whether accepting or rejecting is a good idea.
  - There are two possible signals:
    - A *high signal*, H, which suggests that accepting is a good idea.
    - A *low signal*, L, which suggests that accepting is a bad idea.
  - $Pr[H|G] = q > 1/2$ and $Pr[L|G] = 1 - q < 1/2$.
  - Similarly, $Pr[L|B] = q$ and $Pr[H|B] = 1 - q$.

Urn experiment in terms of cascade model

Cascade model

- *States of the World*: majority-blue or majority-red
- *Signals*: Accepting = guessing majority-blue.
  This is good, G, if the urn really has majority-blue and a bad idea, B, otherwise.
  The prior probability of accepting being a good idea is $p = 1/2$.
  The signal is high if the drawn ball is blue.
  $Pr[H|G] = q = 2/3$.

Cascade model

Individual decisions

- Suppose that a person gets the high signal.
- This shifts the expected payoff from $v_g Pr[G] + v_b Pr[B] = 0$ to $v_g Pr[G|H] + v_b Pr[B|H]$
- We use Bayes Law to calculate $Pr[G|H]$.

$$
\begin{aligned}
Pr[G|H] &= \frac{Pr[G] \times Pr[H|G]}{Pr[H]} \\
&= \frac{Pr[G] \times Pr[H|G]}{Pr[G] \times Pr[H|G] + Pr[B] \times Pr[H|B]} \\
&= \frac{pq}{pq + (1-p)(1-q)} \\
&> p
\end{aligned}
$$

Cascade model

Multiple signals

Suppose that there is a sequence, S, of a high signals and b low signals.

1. $Pr[G|S] > Pr[G]$ when $a > b$.
2. $Pr[G|S] < Pr[G]$ when $a < b$.
3. $Pr[G|S] = Pr[G]$ when $a = b$.

$$
\begin{aligned}
Pr[G|S] &= \frac{Pr[G] \times Pr[S|G]}{Pr[S]} \\
Pr[S|G] &= q^a(1-q)^b \\
Pr[S] &= Pr[G] \times Pr[S|G] + Pr[B] \times Pr[S|B] \\
&= pq^a(1-q)^b + (1-p)(1-q)^a q^b
\end{aligned}
$$

Cascade model

Multiple signals

Suppose that there is a sequence, S, of a high signals and b low signals.

1. $Pr[G|S] > Pr[G]$ when $a > b$.
2. $Pr[G|S] < Pr[G]$ when $a < b$.
3. $Pr[G|S] = Pr[G]$ when $a = b$.

$$
\begin{aligned}
Pr[G|S] &= \frac{Pr[G] \times Pr[S|G]}{Pr[S]} \\
&= \frac{pq^a(1-q)^b}{pq^a(1-q)^b + (1-p)(1-q)^a q^b}
\end{aligned}
$$

How does this compare to $Pr[G]$ or p?
Replace the second term in the denominator with $(1-p)q^a(1-q)^b$ and compare.

Sequential decision making

- Person 1 follows her own signal.
- Person 2 follows her own signal.
  - If Person 2's signal differs from Person 1's, Person 2 is indifferent and chooses her own signal.
  - If Person 1's signal is the same as Person 2's, Person 2 chooses that signal.
- Person 3 follows the two signals before if they are the same, or her own if they differ.
- .....
- Person N?

Sequential decision making

For person N:

As long as the number from acceptances differs from the number of rejections by at most one, each person in sequence is simply following their own signal.

But once the number of acceptances differs from the number of rejections by two or more, a cascade takes over, and everyone simply follows the majority decision forever.

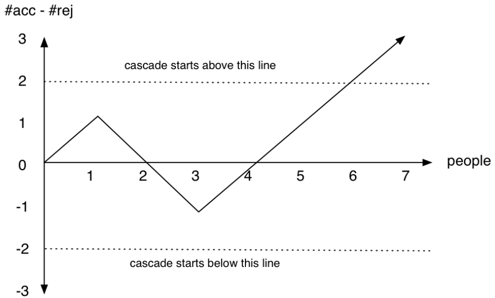

Lessons from cascades

- Cascades can be wrong.
- Cascades can be based on very little information.
- Cascades are fragile.