

Chapter 3: Probability

Yonghyun Kwon

Department of Mathematics, Korea Military Academy

Slides developed by Mine Çetinkaya-Rundel of OpenIntro.

The slides may be copied, edited, and/or shared via the CC BY-SA license.

Some images may be included under fair use guidelines (educational purposes).

Defining probability

Random processes

- A *random process* is a situation in which we know what outcomes could happen, but we don't know which particular outcome will happen.
- Examples: coin tosses, die rolls, whether the stock market goes up or down tomorrow, etc.

MP3 Players > Stories > iTunes: Just how random is random?

iTunes: Just how random is random?

By David Braue on 08 March 2007

- | | |
|---|---|
| <ul style="list-style-type: none">• Introduction• Say You, Say What? | <ul style="list-style-type: none">• A role for labels?• The new random |
|---|---|

Think that song has appeared in your playlists just a few too many times? David Braue puts the randomness of Apple's song shuffling to the test -- and finds some surprising results.

Quick -- think of a number between one and 20. Now think of another one, and another, and another.

Starting to repeat yourself? No surprise: in practice, many series of random numbers are far less random than you would think.

Computers have the same problem. Although all systems are able to pick random numbers, the method they use is often tied to specific other numbers -- for example, the time -- that means you could get a very similar series of 'random' numbers in different situations.

This tendency manifests itself in many ways. For anyone who uses their iPod heavily, you've probably noticed that your supposedly random 'shuffling' iPod seems to be particularly fond of the Bee Gees, Melissa Etheridge or Pavarotti. Look at a random playlist that iTunes generates for you, and you're likely to notice several songs from one or two artists, while other artists go completely unrepresented.



<http://www.cnet.com.au/>

itunes-just-how-random-is-random-339274094.htm



- There are several possible interpretations of probability but they (almost) completely agree on the mathematical rules probability must follow.
 - $P(A)$ = Probability of event A
 - $0 \leq P(A) \leq 1$
- The *probability* of an outcome is the proportion of times the outcome would occur if we observed the random process an infinite number of times.




Which of the following events would you be most surprised by?

- (a) exactly 3 heads in 10 coin flips
- (b) exactly 3 heads in 100 coin flips
- (c) exactly 3 heads in 1000 coin flips



Law of large numbers

Law of large numbers states that as more observations are collected, the proportion of occurrences with a particular outcome, $\hat{p}_n$, converges to the probability of that outcome, p .



Disjoint and non-disjoint outcomes

Disjoint (mutually exclusive) outcomes: Cannot happen at the same time.

- The outcome of a single coin toss cannot be a head and a tail.
- A student both cannot fail and pass a class.
- A single card drawn from a deck cannot be an ace and a queen.

Non-disjoint outcomes: Can happen at the same time.

- A student can get an A in Stats and A in Econ in the same semester.



Union of non-disjoint events

What is the probability of drawing a jack or a red card from a well shuffled full deck?




Figure from <http://www.milefoot.com/math/discrete/counting/cardfreq.htm>.


General addition rule

$$P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$$


Note: For disjoint events $P(A \text{ and } B) = 0$, so the above formula simplifies to $P(A \text{ or } B) = P(A) + P(B)$.

Note: Union of two events A and B : $A \text{ or } B = A \cup B$.

Intersection of two events A and B : $A \text{ and } B = A \cap B = AB$.



$A \cup B$



$A \cap B$

Sample space and complements

Sample space(S) is the collection of all possible outcomes of a trial.

- A couple has one kid, what is the sample space for the gender of this kid? $S = \{M, F\}$
- A couple has two kids, what is the sample space for the gender of these kids? $S =$

The *Complement* of an event, A^c , is the event that A does not occur.

- A couple has one kid. If we know that the kid is not a boy, what is gender of this kid? $A = \{M\} \rightarrow A^c = \{F\}$
- A couple has two kids, if we know that they are not both girls, what are the possible gender combinations for these kids?




The complement of set A within the sample space S

Complement

A and A^c are mathematically related as:

$$P(A) + P(A^c) = 1, \quad \text{i.e. } P(A) = 1 - P(A^c).$$



Two processes are *independent* if knowing the outcome of one provides no useful information about the outcome of the other.

- The outcome of one coin toss does not affect the next. → Coin tosses are
- The outcome of drawing a card (*without replacement*) affects the next draw. → Card draws (*without replacement*) are



Product rule for independent events

$$P(A \text{ and } B) = P(A) \times P(B)$$

Or more generally, $P(A_1 \text{ and } \cdots \text{ and } A_k) = P(A_1) \times \cdots \times P(A_k)$

You toss a coin twice, what is the probability of getting two tails in a row?

$$P(\text{T on the first toss}) \times P(\text{T on the second toss}) =$$



Disjoint vs. complementary

Do the sum of probabilities of two disjoint events always add up to 1?

Do the sum of probabilities of two complementary events always add up to 1?

Probability of “at least one”

To find the probability of at least one of something, calculate the probability of none and then subtract that result from 1.

$$P(\text{at least one}) = 1 - P(\text{none})$$



Practice

Roughly 20% of undergraduates at a university are vegetarian or vegan. What is the probability that, among a random sample of 3 undergraduates, at least one is vegetarian or vegan?

(a) $1 - 0.2 \times 3$

(b) $1 - 0.2^3$

(c) 0.8^3

(d) $1 - 0.8 \times 3$

(e) $1 - 0.8^3$



Conditional probability

Photo Classification Results

A classifier was tested on 1822 photos from a photo-sharing website to determine whether they depict fashion. Each photo was classified by:

- `mach_learn`: ML classifier output (`pred_fashion` or `pred_not`).
- `truth`: Human-verified label (`fashion` or `not`).

		truth		Total
		<code>fashion(A)</code>	<code>not(A^c)</code>	
<code>mach_learn</code>	<code>pred_fashion(B)</code>	197	22	219
	<code>pred_not(B^c)</code>	112	1491	1603
Total		309	1513	1822



Marginal probability

What is the probability that the ML classifier identified the photo as being about fashion?

		truth		Total
		fashion(A)	not(A^c)	
mach_learn	pred_fashion(B)	197	22	219
	pred_not(B^c)	112	1491	1603
Total		309	1513	1822

(Probability that mach_learn is pred_fashion)

=



Joint probability

What is the probability that the ML classifier predicted the photo as being about fashion and the photo is truly about fashion?

		truth		Total
		fashion(A)	not(A^c)	
mach_learn	pred_fashion(B)	197	22	219
	pred_not(B^c)	112	1491	1603
Total		309	1513	1822

(Probability that mach_learn is pred_fashion and truth is fashion)

=



Conditional probability

Conditional probability

The conditional probability of the outcome of interest A given condition B is calculated as

$$P(A|B) = \frac{P(A \text{ and } B)}{P(B)} \quad \begin{array}{l} \leftarrow \text{joint probability} \\ \leftarrow \text{marginal probability} \end{array}$$

	fashion(A)	not(A^c)	Total
pred_fashion(B)	197	22	219
pred_not(B^c)	112	1491	1603
Total	309	1513	1822

(Probability that truth is fashion given mach_learn is pred_fashion)

=



Conditional probability (cont.)

If we knew the ML classifier predicted the photo was about fashion, what is the probability that the photo is actually about fashion?

	fashion(A)	not(A ^c)	Total
pred_fashion(B)	197	22	219
pred_not(B ^c)	112	1491	1603
Total	309	1513	1822

$$P(\text{fashion} \mid \text{pred_fashion}) = \frac{197}{219} \approx 0.900$$

$$P(\text{fashion} \mid \text{pred_not}) = \frac{112}{1603} \approx 0.070$$



Conditional probability (cont.)

What is the probability that the ML prediction was correct if the photo was about fashion?

	fashion(A)	not(A ^c)	Total
pred_fashion(B)	197	22	219
pred_not(B ^c)	112	1491	1603
Total	309	1513	1822

$$P(\text{pred_fashion} \mid \text{fashion}) = \frac{197}{309} \approx 0.638$$

$$P(\text{pred_fashion} \mid \text{not fashion}) = \frac{22}{1513} \approx 0.015$$



General multiplication rule

- Earlier we saw that if two events are independent, their joint probability is simply the product of their probabilities.
- If the events are not believed to be independent, the joint probability is calculated slightly differently.

General multiplication rule

If A and B represent two outcomes or events, then

$$P(A \text{ and } B) = P(A|B) \times P(B) = P(B|A) \times P(A)$$

- Note that this formula is simply the conditional probability formula, rearranged.



Independence and conditional probabilities

Consider the following (hypothetical) distribution of gender and major of students in an introductory statistics class:

	social science	non-social science	total
female	30	20	50
male	30	20	50
total	60	40	100

- The probability that a randomly selected student is a social science major is
- The probability that a randomly selected student is a social science major given that they are female is

Does major of students in this class depend on their gender?



Independence and conditional probabilities (cont.)

Generically, if $P(A|B) = P(A)$ then the events A and B are said to be independent.


- Conceptually: Giving B doesn't tell us anything about A .
- Mathematically: We know that if events A and B are independent, $P(A \text{ and } B) = P(A) \times P(B)$. Then,

$$P(A|B) =$$



Inverting probabilities

In Canada, 0.35% of women over 40 develop breast cancer each year. A mammogram, a common screening test, has a false negative rate of 11% and a false positive rate of 7%. If a randomly selected woman over 40 tests positive(event B), what is the probability that she actually has breast cancer(event A)?



$$P(A|B)$$

=



Bayes' Theorem

The conditional probability formula we have seen so far is a special case of the *Bayes' Theorem*, which is applicable even when events have more than just two outcomes.


Bayes' Theorem:

Let $A_1, \dots, A_n$ and B be events where A_i are disjoint, $\cup_{i=1}^n A_i = S$, and $P(A_i) > 0$ for all $i = 1, \dots, k$. Then

$$P(A_i|B) = \frac{P(B|A_i)P(A_i)}{P(B|A_1)P(A_1) + P(B|A_2)P(A_2) + \dots + P(B|A_k)P(A_k)}.$$



Application activity: Inverting probabilities



$$P(\text{inf}|\text{pos}) = \frac{P(\text{inf and pos})}{P(\text{pos})} = \frac{0.099}{0.03 + 0.099 + 0.105} \approx 0.423$$

