Bayes' Theorem (7.3)

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

- P(b) is the *prior* probability of b.
- P(b | a) is the *posterior* probability, after taking the evidence a into account.
- P(a | b) is the *likelihood* of the evidence, given the hypothesis.
- P(a) is the *prior* probability of the evidence
  - used as a normalizing constant

Why is this useful? Consider a medical diagnosis. Diagnostic evidence P(disease | symptom) is often hard to get. *<u>But it's what you really want.</u>* Causal evidence P(symptom | disease) is often easier to get. P(disease) is easy to get.

Diagnosing a Rare Disease

- Meningitis is rare: P(m) = 1/50000
- Meningitis causes stiff neck: P(s|m) = 0.5
- Stiff neck is not so rare: P(s) = 1/20
- You have a stiff neck. What is P(m|s)?

The point being that we can quickly figure out if a stiff neck is a good reason to guess someone has Meningitis.

Evaluating Public Policy

- Terrorists are rare: $P(t) = 1/10^6$
- Suppose we have a very accurate test:
  - Accuracy: $P(pos \mid t) = 0.99$
  - Specificity: $P(neg \mid \neg t) = 0.99$

	pos	neg	
t			
$\neg t$			
			300,000,000

So we want to know P(t|pos). P(pos|t) is known as is P(t). But how do we compute P(pos)?

$P(pos) = P(pos \mid t)P(t) + P(pos \mid \neg t)P(\neg t)$ ← Does that seem reasonable? If so, we get:

$$P(t \mid pos) = \frac{P(pos \mid t)P(t)}{P(pos \mid t)P(t) + P(pos \mid \neg t)P(\neg t)}$$

That form is the one we find in the text.

BAYES' THEOREM Suppose that E and F are events from a sample space S such that $p(E) \neq 0$ and $p(F) \neq 0$. Then

$$p(F \mid E) = \frac{p(E \mid F)p(F)}{p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F})}.$$

Applications

There are a huge number of applications, mainly in artificial intelligence, related to Bayes' theorem. One you interact with most every day is called Bayesian filtering, and it's (mostly) what is used to keep your inbox spam free.

Say you have a set of B messages known to be spam and a set of G messages known to not be spam (Google, for example, gets a good sense of this when you label things as spam…). We could then search for words (or addresses, or whatever) that tend to occur in B but are less common in G.

> Suppose that we have found that the word "Rolex" occurs in 250 of 2000 messages known to be spam and in 5 of 1000 messages known not to be spam. Estimate the probability that an incoming message containing the word "Rolex" is spam, assuming that it is equally likely that an incoming message is spam or not spam. If our threshold for rejecting a message as spam is 0.9, will we reject such messages? (Example 3, page 473 though we'll approach it a bit differently.)

We want to find the probability that a word with "Rolex" in it is spam. We don't know what percent of all messages are spam, so let's just assume 50% for now. For p(spam | Rolex) we get

$$P(s \mid r) = \frac{P(r \mid s)P(s)}{P(r \mid s)P(s) + P(r \mid \neg s)P(\neg s)} = \frac{P(r \mid s)}{P(r \mid s) + P(r \mid \neg s)}$$ (That last is true if P(s)=.5).

Of course, filtering based on one word isn't a very good idea—you get lots of false positives. Instead you use lots of different filters all together. *Example 4 on page 474* does a nice job of giving an example of combining filters.

Monty Hall

What is the probability of winning with strategy 1? What would you have had to choose if strategy 1 is to work? If strategy 2?

This question was actually the subject of a big debate where serious mathematicians got the answer wrong.

Monty Hall Problem

- Monty: secretly choose one door to be the car (the other two are "empty").
- Contestant: tell Monty which door you choose.
- Monty: reveal an "empty" door, and offer to let contestant switch doors.

- Contestant: decline the offer (not switch doors) (strategy 1)
- Contestant: accept the offer (switch doors) (strategy 2)

- Monty: reveal the prize!

Note: In the original problem, the contestant can freely choose to decline or accept without committing to a fixed strategy. However, we will consider two fixed policies to analyze the problem.

Monty Hall (using Bayes Theorem)

Consider 3 discrete random variables all taking on values in {1,2,3}

- C: the number of the door hiding the **C**ar
- S: the number of the door **S**elected by player
- H: the number of the door opened by **H**ost

$$P(C=2|H=3,S=1) = \frac{P(H=3|C=2,S=1)P(C=2|S=1)}{\sum_{i=1}^{3} P(H=3|C=i,S=1)P(C=i|S=1)}$$
$$= \frac{1\times\frac{1}{3}}{\frac{1}{2}\times\frac{1}{3}+1\times\frac{1}{3}+0\times\frac{1}{3}} = \frac{2}{3};$$

$$P(C=1|S=1) = P(C=2|S=1) = P(C=3|S=1) = \tfrac{1}{3}.$$
$$P(H=3|C=1,S=1) = \tfrac{1}{2},$$
$$P(H=3|C=2,S=1) = 1,$$
$$P(H=3|C=3,S=1) = 0.$$

Prob. of winning by switching

Of course, we could also use Bayes' Theorem to work this out. I'm showing that mainly as a review, but also as a bit more complex of an example.

Recurrence relations (Chapter 8)

Consider the following problem:

> A young pair of rabbits (one of each sex) is placed on an island. A pair of rabbits does not breed until they are 2 months old. After they are 2 months old, each pair of rabbits produces another pair each month, as shown in Figure 1. Find a recurrence relation for the number of pairs of rabbits on the island after n months, assuming that no rabbits ever die. (Example 1,p 502)

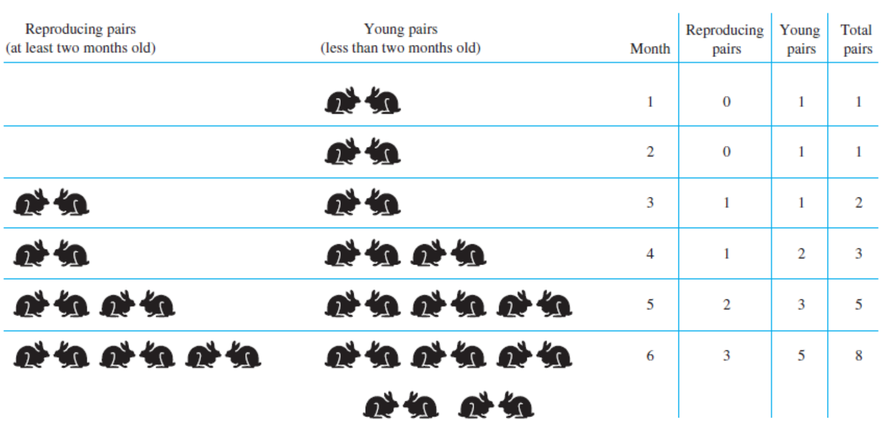

Month	Reproducing pairs	Young pairs	Total pairs
1	0	1	1
2	0	1	1
3	1	1	2
4	1	2	3
5	2	3	5
6	3	5	8

Notice that f_1=1 and f_2=1. Beyond that $f_n=f_{n-1}+f_{n-2}$. How many pairs of rabbits do we expect to have on month 5?

The above is a form of *recurrence relation* (which we looked at back in the first week of classes).

Modeling problems using recurrence relations.

It turns out that recurrence relations are a common way of looking at the complexity of algorithms. Let's look at one of the most common examples of an algorithm whose complexity can be modeled as a recurrence relation: the Towers of Hanoi.

Towers of Hanoi

Goal: move all discs from peg A to peg C

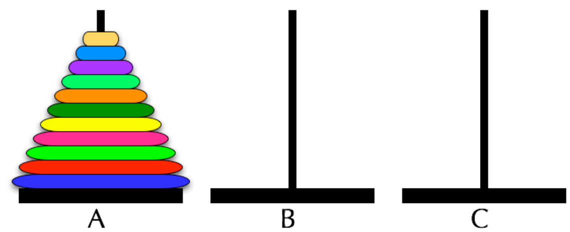

Rules:

1. only one disc can be moved at a time
2. a disc cannot be put on a smaller disc

The algorithm for doing this isn't initially all that intuitive. Notice that at *some* point you need to move the bottom disc to C. At that point, what where *must* the rest of the discs be?

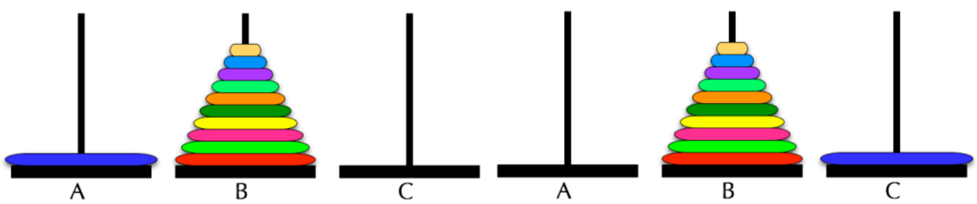

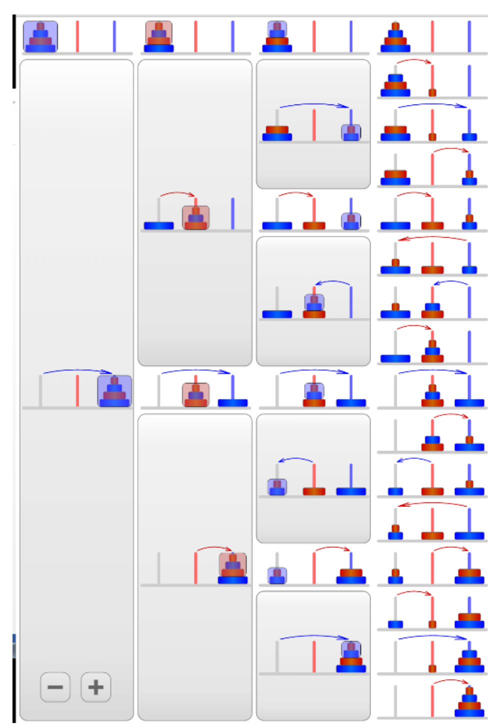

What we notice is that we must get all the other discs to B so that the largest disc can move from A to C (can you provide an argument to that effect?) And what we've basically done is reduce the problem: we now need to move n-1 discs from A to B, then move the largest disc to C and then move the n-1 discs from B to C.

Basically, we need to do the following[i]

ToH(n,A,B,C) : *move n discs from A to C using B if necessary.*
  If n=1, move the disc from A to C and stop
  ToH(n-1,A,C,B) ← from A to B using C
  Move disc n from A to C
  ToH(n-1,B,A,C) ← from B to C using A

---

[i] http://upload.wikimedia.org/wikipedia/commons/2/20/Tower_of_Hanoi_recursion_SMIL.svg is pretty cool. Click around on it.

So what is the complexity of this algorithm? Pretty clearly, in order to solve this for "n" discs, we need to solve it twice for "n-1" and do one extra move.

Define T(n) = number of moves when there are n discs.

- T(1) = 1
- T(n) = 2T(n-1) + 1

How do we find a closed form expression for T(n)? With an educated guess we could try to prove it by induction!

What is T(1)? T(2)? T(3)? Any guesses on T(n)?

Let's prove it.

---

*Warning, the following is going to be taught mostly as formula "plug and chug". The text occasionally has proofs of things, but generally not very convincing ones and it the proofs are often (literally) left as an exercise for the reader. We aren't going to worry about **why** too much here, just what.*

Now, how would we go about finding a closed form for something like the Fibonacci sequence? Or even finding a solution for the Towers of Hanoi without having to guess first?

Solving Linear Recurrence Relations (8.2)

A *linear homogeneous recurrence relation of degree k with constant coefficients* is a recurrence relation of the form

$$a_n = c_1 a_{n-1} + c_2 a_{n-2} + \cdots + c_k a_{n-k},$$

where $c_1, c_2, \ldots, c_k$ are real numbers, and $c_k \neq 0$.

- The recurrence is *linear* because the all the "a_n" terms are just the terms (not raised to some power nor are they part of some function). So $a_n=2a_{n-1}$ is linear but $a_n=2(a_{n-1})^2$ is not.
- It is *homogeneous* because *all* terms are multiples of some previous value of a_n. So $a_n=2a_{n-1}$ is homogeneous, but $a_n=2a_{n-1}+1$ is not.
- It is of *degree* k because a_n is expressed in terms of the last k terms of the sequence.
- And it has *constant coefficients* because all the c terms are constants (not a function of n).

OK, we've got a few definitions out of the way. It turns out that linear homogeneous recurrence relations typically have a solution of the form $a_n=r^n$ where r is some constant. Replacing a_n with r^n we get:

$$r^n = c_1 r^{n-1} + c_2 r^{n-2} + \cdots + c_k r^{n-k}$$

Now, if we divide both sides of this equation by r^{n-k} and move the things around a bit, we get:

$$r^k - c_1 r^{k-1} - c_2 r^{k-2} - \cdots - c_{k-1} r - c_k = 0.$$

We call the above equation the ***characteristic equation***.

Examples

Consider the following recurrence relations and indicate if they are linear, homogeneous and have constant coefficients. If they do, find the characteristic equation associated with the recursion.

1. $a_n=a_{n-1}+a_{n-2}$
2. $a_n=a_{n-1}+a_{n-2}+2$
3. $a_n=2a_{n-1}+a_{n-2}$
4. $a_n=a_{n-2}+a_{n-3}$
5. $a_n=na_{n-1}+a_{n-2}$
6. $a_n=4(a_{n-1})-a_{n-2}$

Working with LH RR with CC and degree 2

Let c_1 and c_2 be real numbers. Suppose that $r^2 - c_1 r - c_2 = 0$ has two distinct roots r_1 and r_2. Then the sequence $\{a_n\}$ is a solution of the recurrence relation $a_n = c_1 a_{n-1} + c_2 a_{n-2}$ if and only if $a_n = \alpha_1 r_1^n + \alpha_2 r_2^n$ for $n = 0, 1, 2, \ldots$, where α_1 and α_2 are constants.

That feels like it jumps out of nowhere. Let's focus on what it says, time allowing we will try to show it makes sense (or at least a bit of sense). The proof is on pages 515 and 516 of the text.

Consider the Fibonacci sequence defined by $a_n=a_{n-1}+a_{n-2}$. This is saying we need to find the roots of the characteristic equation and then the solution for this relation is of the form $a_n = \alpha_1 r_1^n + \alpha_2 r_2^n$, where r_1 and r_2 are those roots. Notice that α is just an arbitrary constant (that we'd have to figure out…)

(Mostly) Worked Example (p516 example 3)

Let's look at $a_n=a_{n-1} + 2a_{n-2}$ where $a_0=2$ and $a_1=7$

- Is this of the correct form? YES
- What is the characteristic equation? $r^2-r-2=0$
- What are the roots of the characteristic equation? $r=2$ and $r=-1$

From the above, we get $a_n = \alpha_1 2^n + \alpha_2(-1)^n$. Now we just need to solve for the two alphas. We know that $a_0=2$ and $a_1=7$. So we get two equations and two unknowns.

$$a_0 = 2 = \alpha_1 + \alpha_2,$$
$$a_1 = 7 = \alpha_1 \cdot 2 + \alpha_2 \cdot (-1)$$

So what is α_1 and α_2? What is the final solution?

You do this one (example 4 page 517)

Find the closed form solution for the Fibonacci numbers. Use $f(n)=f(n-1)+f(n-2)$ where $f(0)=0$ and $f(1)=1$.

- Is this of the correct form?
- What is the characteristic equation?
- What are the roots of the characteristic equation?

Repeated roots

Let c_1 and c_2 be real numbers with $c_2 \neq 0$. Suppose that $r^2 - c_1 r - c_2 = 0$ has only one root r_0. A sequence $\{a_n\}$ is a solution of the recurrence relation $a_n = c_1 a_{n-1} + c_2 a_{n-2}$ if and only if $a_n = \alpha_1 r_0^n + \alpha_2 n r_0^n$, for $n = 0, 1, 2, \ldots$, where α_1 and α_2 are constants.

Let's do $a_n=6a_{n-1} -9a_{n-2}$ where $a_0=1$ and $a_1=6$

- What is the characteristic equation?
- What are the roots of the characteristic equation?

General form for arbitrary degree

In this class, we will only be working with relations of this type of degree 2. But there is a more general result.

Let $c_1, c_2, \ldots, c_k$ be real numbers. Suppose that the characteristic equation

$$r^k - c_1 r^{k-1} - \cdots - c_k = 0$$

has t distinct roots $r_1, r_2, \ldots, r_t$ with multiplicities $m_1, m_2, \ldots, m_t$, respectively, so that $m_i \geq 1$ for $i = 1, 2, \ldots, t$ and $m_1 + m_2 + \cdots + m_t = k$. Then a sequence $\{a_n\}$ is a solution of the recurrence relation

$$a_n = c_1 a_{n-1} + c_2 a_{n-2} + \cdots + c_k a_{n-k}$$

if and only if

$$\begin{aligned} a_n = {} & (\alpha_{1,0} + \alpha_{1,1} n + \cdots + \alpha_{1,m_1-1} n^{m_1-1}) r_1^n \\ & + (\alpha_{2,0} + \alpha_{2,1} n + \cdots + \alpha_{2,m_2-1} n^{m_2-1}) r_2^n \\ & + \cdots + (\alpha_{t,0} + \alpha_{t,1} n + \cdots + \alpha_{t,m_t-1} n^{m_t-1}) r_t^n \end{aligned}$$

for $n = 0, 1, 2, \ldots$, where $\alpha_{i,j}$ are constants for $1 \leq i \leq t$ and $0 \leq j \leq m_i - 1$.

Again, we aren't going to worry about this general case, but I want you aware it exists.