Combinatorics: The Fine Art of Counting

Lecture Notes – Counting 101

Note – to improve the readability of these lecture notes, we will assume that multiplication takes precedence over division, i.e. A / B*C always means A / (B*C).

Vocabulary Notes: For the purposes of these lecture notes, we will assume the following terminology. Words that are considered synonyms may be used interchangeably.

- The words *distinct*, *different*, and *distinguishable*, are all synonyms.
- The words *ordering, arrangement,* and *permutation* are all synonyms.
- The words *outcome* and *possibility* are synonyms.
- The words *string* and *sequence* are synonyms that refer to an ordered list of elements. The elements in a string need not be distinct. Any non-empty string of letters will be considered a *word*.
- The words *separate*, *non-overlapping*, and *disjoint* are all synonyms that refer to two sets with no elements in common, i.e. whose intersection is the empty set.
- A binary digit or *bit* is an element of {0, 1}. A decimal digit is an element of {0,1,2,3,4,5,6,7,8,9}. A letter is an element of {A, B, C,…,Z}.

Set Notes:

- The elements of a set are distinct by definition. A set always contains exactly one of each element (objects which can contain multiple copies of elements are called multi-sets).
- The union of two sets A U B is the set of elements in either (possibly both) sets. {1,2} U {2,3} = {1,2,3}. The union of two sets is the smallest set which contains both sets as subsets.
- The intersection of two sets A ∩ B is the set of elements they have in common. {1,2} ∩ {2,3} = {2}. The intersection of disjoint sets is the empty set, denoted $\emptyset$.
- |A| denotes the number of elements in A.
- A *partition* of a set if a collection of disjoint subsets whose union equals the entire set.

0th law of counting: Add exclusive options. When choosing one item from two disjoint sets, A and B, there are |A| + |B| items to choose from, or |A| + |B| possible outcomes (Note that |A| denotes the number of elements in the set A). This generalizes to an arbitrary number of sets and is the basis of counting by cases (see below).

Example: If you are choosing among doing homework for one of your 3 classes versus going to see one of 4 movies this evening, you have 7 options to choose from.

1st law of counting: Multiply successive choices. If you are first choosing one option from set A and then making a second choice from set B (A and B could be the same, overlap, or be completely different), the total number of possible outcomes for the sequence of choices is |A| * |B|. This generalizes to an arbitrary number of successive choices.

Bear in mind that an "outcome" is an ordered list of the results of each choice - this makes a difference when A and B are not necessarily disjoint. If A and B are both the set of decimal digits, the outcome 25 is not the same as 52 and both outcomes will be counted when you multiply |A| * |B|. In a situation where the problem dictates that the order of the choices doesn't matter, you must then apply the 2^{nd} law of counting to correct for over-counting outcomes that you want to consider identical. Some simple examples are listed below:

- 3 shirts, 4 pairs of pants, and 2 pairs of shoes give 3*4*2 = 24 possible outfits to wear
- 10 coin flips result in 2^{10} possible outcomes (this is still true even if the coin is biased – this may change the *likelihood* of an outcome, but not how many there are).
- 3 rolls of a six-sided die gives 6^3 possible outcomes
- 3 billion base pairs of DNA give $4^{3,000,000,000}$ possible genomes
- 10^6 possible strings of six decimal digits, but only $9*10^5$ six-digit decimal numbers since the first digit cannot be zero.
- 10^3*26^3 possible license plates with three decimal digits followed by three letters.
- 26^5 sequences of five letters. $26^4 + 26^3 + 26^2 + 26$ words with less than 5 letters.
- 2^n sequences of n bits, and 2^n subsets of a set with n elements in it – when constructing a subset, we make a binary "choice" for each element. Thus any sequence of n bits corresponds to a subset.

Note that the 1^{st} law of counting still holds even if choosing an option from A changes the number of options available in successive choices (i.e. the particular element chosen from set A changes the set B in some way), provided the change is *uniform*. If we are picking distinct names out of a hat to form a list of people, we have fewer names left in the hat after the first choice, but the number of choices remaining is the same regardless of which name was chosen first. But be careful, if some choices from A result in different sized B's, then the first law no longer applies and we must analyze particular cases. More examples:

- There are 12*11*10*9 = 11,880 ways that four books out of a bag of 12 books can be placed on a shelf.
- 5*4*3*2*1 = 5! = 120 permutations of the string of letters ABCDE.
- 26*25*24*23*22 = 7,893,600 five letter words with distinct letters.
- 9*9*8*7*6 = 27,216 five digit numbers with distinct digits (note the first digit can't be zero, the second can be any digit but the first).
- $26*25^n$ words of length n that don't contain a double letter – the first letter can be any letter, the second can't equal the first, the third can't equal the second, etc…
- $26^{n/2}$ palindromes of length n (n even) and $26^{(n+1)/2}$ palindromes of length n (n odd) – note that the first n/2 or n/2+1 letters of a palindrome can be anything, but there is only one choice for each remaining letter.
- 9^3 five digit palindrome decimal numbers with no adjacent digits the same.
- 4*10 = 40 four digit palindrome decimal numbers which are even (the first digit cannot be zero and the last digit must be even and equal to the first digit so there are only four choices. The second digit can be any decimal digit, and the third digit must be the same as the second).
- 8^2*7^2 = 3,136 ways a black and a white rook can be placed on an 8x8 chessboard so that neither threatens the other. The first rook may be placed anywhere. Wherever it is placed it threatens or occupies all the positions in one column and one row, leaving exactly 7^2 positions unthreatened. This generalizes to $8^2*7^2*6^2*$… for up to 8 distinct

rooks. If the rooks were not distinct we would need to apply the 2nd law of counting and divide by k!, where k is the number of rooks.

Note that the example of peaceful rook placement does *not* apply to bishops because the number of positions threatened by a bishop depends on the position the bishop occupies (e.g. a bishop in the corner only threatens one diagonal line, while a bishop in the center threatens two). Peaceful bishop placement is a much harder problem because of this difference.

2nd law of counting (Shepard's Law): One way to count sheep is to count the legs and divide by four (provided each sheep has exactly four legs). In mathematical terms, over-counting by a uniform multiplicative factor can be corrected by dividing by this factor.

This situation most often occurs when there are multiple outcomes which share a common feature that we wish to regard as equivalent (such a group is called an **equivalence class**). For example when counting edges in a graph by summing over the degree of each vertex, we are really counting vertex-edge pairs. We want to consider two vertex-edge pairs equivalent if they have the same edge. Since each edge has two vertices, there will always be two vertex-edge pairs per edge, so we have over-counted by a factor of two.

To apply the 2nd law of counting we need to know two things. First we must be sure that we are over-counting in a uniform manner (all sheep must have the same number of legs). The size of each equivalence class must be the same. The second is to compute what this size actually is (how many legs does each sheep have). These two steps were trivial in the vertex-edge example, but in general verifying that the over-counting is uniform may require a bit of thought, and then computing size of the equivalence classes may require further counting in its own right.

(Note – this law can be generalized to handle mutant sheep, as long as the average number of legs per sheep is unchanged. One three-legged sheep plus one five-legged sheep = two four-legged sheep. We will see examples of this in later lectures).

The simplest examples of the 2nd law of counting are where we want to ignore order:

- There are 12*11*10*9 / 4*3*2*1 = 495 ways to choose 4 books out of 12 arranged on a shelf and put them in a bag, since there are 12*11*10*9 ways to pick the books, but once they are in the bag we don't care about the order in which we picked them. For any set of 4 books we choose, there are 4*3*2*1 different ways we could have taken the same books off the shelf. Contrast this with putting 4 books from a bag of 12 onto the shelf.
- 10*9*8 / 3*2*1 = 120 possible committees of three people that can be formed from a group of ten people.
- 52*51*50*49*48 / 5*4*3*2*1 = 2,598,960 possible five card poker hands.
- 36*35*34*33*32*31 / 6*5*4*3*2*1 possible winning lottery tickets if there are 6 balls chosen from a set of 36 and the order of the numbers on the ticket does not matter (or is fixed, e.g. automatically put into increasing order).

In all of these cases, the equivalence classes were simply permutations of the chosen elements. Since we were always choosing the same number of elements, the number of permutations was always the same (step one of applying the 2nd law), and the number of permutations of k distinct objects is simply k! (step two).

Choosing subsets is such a common operation in combinatorics that there is a standard notation for it which will be denoted in these lecture notes by (n k), read "n choose k", which stands for the number of ways you can choose k distinct elements from a set of n elements.

It is normally written vertically not horizontally. We will eventually switch over to using this notation exclusively, but for now we will often write out the multiplication and division.

- (10*9*8 / 3*2*1) * (7*6*5 / 3*2*1) = 10*9*8*7*6*5 / (3*2*1 * 3*2*1) = 4200 ways of choosing two disjoint committees of three people from a group of ten people.
- (10*9*8 / 3*2*1) * (7*6*5 / 3*2*1) * (4*3*2*1 / 4*3*2*1) = 10! / 3! 3! 4! = 4200 ways of dividing ten people up into two groups of three and one group of four – this is identical to the example above since the third group of four is simply those not in a committee.
- n! / k!(n-k)! to divide n people into two groups of size k and n-k. Note that this is equivalent to choosing a group of k people, since the other group of n-k people are simply those not chosen, thus (n k) = n! / k!(n-k)!

Rearranging sequences of elements which are not necessarily distinct:

- The string ABC can be permuted in 3! = 6 ways: ABC, ACB, BAC, BCA, CAB, CBA
- The string ABA can be permuted in 3! / 2! = 3 distinct ways: ABA, AAB, BAA. We can see this by considering the string A_1BA_2 which can be permuted in 3! = 6 ways and then grouping into equivalence classes strings which only differ in the way A_1 and A_2 are ordered within the string – these strings will be indistinguishable if we consider A_1 and A_1 to be identical. There are always 2! = 2 ways to order A_1 and A_2 so our equivalence classes are all the same size and we can apply the 2nd law of counting.
- The string AAA can be permuted in 3! / 3! = 1 distinct way by exactly the same argument. We are now grouping all the permutations of $A_1A_2A_3$ into a single equivalence class of size 3! = 6.

We can generalize the above argument to permutations of arbitrary strings containing possibly repeated letters. Thus the string ABCCC can be permuted in 5!/3! = 20 distinct ways, since we can permute $ABC_1C_2C_3$ in 5! ways and then group these permutations into equivalence classes by regarding to strings as equivalent if they only differ in the order of C_1, C_2, and C_3 which can be ordered in 3! different ways. The string AABCC can be permuted in 5!/2!2! different ways – given a permutation of $A_1A_2BC_1C_2$ changing the order of A_1 and A_2 or of C_1 and C_2 (or both) will result in a permutation in the same equivalence class, and there are always 2!*2! ways we can do this (pick an ordering of A_1 and A_2 and then pick an ordering of C_1 and C_2).

Mississippi Rule: To count the number of distinct permutations of a string of letters where not all of the letters are necessarily distinct, we count the number of permutations as if the letters were distinct and then divide by the number of equivalent permutations that can be made by permuting identical letters.

- 5! / 2! = 60 distinct permutations of HAPPY.
- 4! / 2!2! = 6 distinct permutations of NOON.
- 7! / 3!3! = 140 distinct permutations of BANNANA.
- 11! / 4!4!2! = 34,650 distinct permutations of MISSISSIPPI.
- 6! / 2!2!2! = 90 distinct stacks of 6 red, green, and blue blocks with two of each color.

Strings with repeated letters can be used to partition sets of objects into groups. To split up a group of 10 people into two groups of 3 and one group of 4 we can simply fix the order of the people and then permute the word AAABBBCCCC to assign them to groups A, B, and C. By the Mississippi rule, there are 10! / 3!3!4! ways to do this, which agrees with our earlier calculation.

The second class of examples where we can apply the 2^{nd} law of counting involve situations where our equivalence classes correspond to symmetric configurations.

There are n! ways to arrange n people in a line, but only n! / n = (n-1)! distinct ways to arrange them in a circle if we don't distinguish between rotationally symmetric arrangements. To see this note that given any ordering of n people around a circle, there are always exactly n distinct rotations which we want to regard as equivalent.

Note however that once we can distinguish any particular position in the circle, the rotational symmetry is lost and there are n! distinct arrangements. Thus while there are 8! / 8 = 7! distinct ways eight people can sit around a table, if one of the chairs is broken (and we choose to take note of this fact) then there are 8! distinct ways eight people can sit – pick one person to sit in the broken chair, then a person to sit to their right, and so on.

Sample problem 1: How many ways can eight people sit in a circle if Alice and Bob must sit next to each other?

Solution 1a: Treat Alice and Bob as a single unit and then seat 7 "people" in a circle. There are two ways to combine Alice and Bob (Alice on the left or Bob on the left) and 7!/7 = 6! ways to seat 7 in a circle, giving a total of 2*6! = 1440.

Solution 1b: Seat Alice anywhere. Because of the rotational symmetry, we can't distinguish which seat she is in so we haven't really made a choice (any circular arrangement could be rotated to put Alice anywhere without changing the arrangement). Now sit Bob in one of the two seats adjacent to Alice and pick the remaining six people one at a time to fill in the remaining seats, giving 2*6*5*4*3*2*1 choices or 2*6! = 1440.

Note that in the second solution, once Alice was seated, the symmetry was broken and we could distinguish the seats by their position relative to Alice.

The importance of looking carefully at equivalence classes when applying the 2^{nd} law of counting is highlighted in the following problem:

Sample problem 2: How many distinct ways can we arrange four A's and two B's around a circle?

Bogus Solution: As above, there are 6!/4!2! = 15 distinct permutations of AAAABB. There are 6 rotations of each of these permutations we want to regard as identical, so there must be 15/6 = 2.5 ways. Oops, that can't be right, it's not even an integer…what went wrong?

Real Solution: If we think about placing n A's and two B's around a circle, we can classify the arrangements by looking at the gaps between the two B's. The B's partition the A's into two runs, and the shortest distance between the B's completely determines the arrangement for we can rotate any arrangement with the same shortest distance line up the B's, and all the other positions are A's so they must match. Thus the number of distinct arrangements is simply the number of possible shortest gap sizes. For four A's the possible shortest gap sizes are 0, 1, and 2 for a total of 3 distinct arrangements.

Where did we go wrong in our bogus solution? The problem lies in not checking our application of the 2^{nd} law of counting carefully. For some permutations of AAAABB, the 6 rotations are not always distinct, for example rotating BAABAA by 3 positions gives the same string. As a result, this equivalence class only has 3 strings in it (BAABAA, ABAABA, and AABAAB), while the other two equivalence classes have 6 strings each. Our sheep don't all have the same number of legs!

In many situations the easiest way to count something is to count a larger set which contains the set of objects we wish to count, and then subtract out what we don't want. We place our desired set within some "universe" of objects we can count, then compute the size of the set's complement within this universe (i.e. count what we don't want), then subtract.

For example, the number of 5 letter sequences containing the letter A is equal to the number of all possible 5 letter sequences (our universe) minus the number of 5 letter sequences that *don't* contain the letter A. Thus there are $26^5 - 25^5$ = 2,115,171 five letter sequences which contain the letter A. We could try to count this set directly, but then there are a lot of cases to consider – counting sequences containing exactly one A is pretty straight-forward ($5 * 26^4$), but then we need to think about what happens if there is more than one A. We will learn how to break this down and count all the different possibilities, but it is still much easier to solve this simple problem by counting the complement. Some more examples:

- $26^5 - 24^5$ five letter sequences that contain an A or a Z, since there are 24^5 that don't (this one would be even harder to count directly).
- $26^5 - 21^5$ five letter sequences that contain a vowel, since there are 21^5 five letter sequences of consonants.
- $26^5 - 26*25^5$ five letter sequences that contain a double letter since there are $26*25^5$ five letter sequences with all adjacent letters distinct.
- $6^4 - 5^4$ way of rolling at least one six in four rolls of a die. The probability of getting at least one six in four rolls is thus $(6^4 - 5^4) / 6^4$ or about 52%.
- $26^3 - 21^3$ five letter palindromes that contain a vowel, since there are 26^3 five letter palindromes and 21^3 of these do not contain a vowel.

Sample problem 3: How many different mixed-gender committees of 3 people can be chosen from a group of 5 men and 5 women?

Solution: There are (10 3) was of selecting a committee of three from a group of 5+5=10 people. There are (5 3) ways of choosing an all male committee from the 5 men and (5 3) was of choosing an all female committee from the 5 women. Thus there are (10 3) – 2* (5 3) or 120 – 2*20 = 80 committees which have at least one man and one woman.

Sample problem 4: How many ways can 8 people sit around a circular table if Alice and Bob won't sit next to each other?

Solution 4a: Simply subtract the number of ways the 8 people can be seated with Alice and Bob adjacent which we computed in problem 1 from the total number of ways 8 people can sit around a circular table without any restriction to get 7! – 2*6! = (7-2)*6! = 5*6!

Solution 4b: Seat Alice anywhere (as in problem 1, we don't count this as a choice since we can't distinguish her seat). Now sit Bob in one of the five seats not adjacent to Alice and pick the remaining six people one at a time to fill in the remaining seats, giving 5*6*5*4*3*2*1 choices or 5*6! as above.

The second solution above is an example of constructive counting. To count the number of possibilities we construct one in some particular way and count the number of options we had to choose from when doing it. To use this method we must be sure of three things:

1) Our method can be used to construct any element of the set we are counting.
2) The same number of options were available during any particular construction.
3) We will never construct the same object in two different ways, i.e. by making different choices during our construction.

Sample problem 5: How many ways can four boys and four girls be seated around a circular table alternating boy/girl?

Solution: Seat 4 girls in alternating seats in a circle of 4 – there are 3! ways to do this. Now pick a boy to fill in each gap between two girls – there are 4*3*2*1=4! ways to do this giving a total of 3!*4! = 144 arrangements.

To see that this satisfies our three requirements, we can reason backwards from a given boy/girl arrangement. We can certainly construct any such arrangement by simply using the arrangement to guide our choices so (1) is satisfied and (2) is clearly also satisfied. To see that (3) is satisfied (this is usually the only tricky part), note that if the arrangement of the girls we chose varies, then certainly the boy/girl arrangement will vary, and given a particular arrangement of girls, each of the gaps between them is distinguished by the pair of girls on either side. If a different boy is placed in any gap, a different boy/girl arrangement will result.

Sample problem 6: How many distinct ways can the faces of a cube be painted with 6 different colors? Two paintings are considered identical if the cube can be oriented so that the cubes look the same.

Solution 6a: Any painting of the cube can be oriented so that a particular color, say red, is on the bottom, so we can assume that we start with a cube that has five blank faces with the bottom painted red. Pick one of the 5 colors remaining to paint the top (5 options), and arrange the 4 remaining colors in a circle and paint the sides accordingly – there are 3! ways to do this. Thus there are a total of 5*3! = 30 ways to paint the cube.

We now check that our construction satisfies (1)-(3) above, It is easy to see that (1) and (2) are satisfied. For (3) note that if we choose a different color to put opposite red, we will get a distinctly different painting of the faces, and that the arrangement of colors about the other 4 faces corresponds in a 1-1 fashion with a circle of 4 colors so different circles will result in distinguishable paintings of the faces.

Solution 6b: If we fix the orientation of the cube and number its faces, we will get a permutation of the six colors for any particular painting. There are 6! permutations of six colors. We want to regard as equivalent permutations which correspond to different orientations of the same cube. A cube can be uniquely oriented by picking one of the 6 faces to be on the bottom, and then rotating a particular one of the 4 side faces to be in front, thus there are 24 distinct orientations and our equivalence classes all have 24 elements. Applying the 2^{nd} law of counting, there are 6! / 24 = 5*3! = 30 distinct ways to paint the cube.

The last method of counting we will look at in this lecture is counting by cases. This is often the method of last resort, but it many situations it is the easiest approach to use and may allow us to quickly break down the problem into sub-problems we can easily solve using methods we already know. The key requirement is that our cases must **partition** the set of objects we are counting, i.e.:

1) The cases must not overlap, otherwise we will some things twice (although we will see methods for correcting for this when we talk about inclusion/exclusion)
2) The cases must cover all possibilities, otherwise we may miss something.

Once we have partitioned the problem into cases, we count each case and then sum over all the different cases as in the 0^{th} law of counting. Note that the cases need not contain the same number of elements. Some may, and these we can group and multiple to add them up.

Sample problem 7: How many ways can 8 people stand in a line if Alice and Bob refuse to stand next to each other?

Solution: We will partition the problem according to where Alice stands.

- If Alice is either first or last in line (2 cases), then there are 6 spots where Bob can stand and not be next to Alice. The other 6 people can fill in the remaining spots in 6! ways, giving 6*6! lineups for each of these 2 cases.
- If Alice is not first or last then she must be in one of the six middle positions (6 cases) and there are only 5 spots Bob can stand in. The remaining spots can be filled in 6! ways giving 5*6! lineups for each of these 6 cases.
- The total number of possible lineups is thus 2*6*6! + 6*5*6! = 7*6*6! = 6*7! = 30240.

In this problem it was easy to see that our case analysis partitioned the possibilities. Alice has to stand somewhere and she can't stand in two places at once.

Sample problem 8: Each faces of a six-sided die is randomly painted red or blue with equal probability. What is the probability the resulting die contains a mono-chromatic ring of four faces (i.e. the die can be oriented so that all the vertical faces are the same color)?

Solution: First recall that computing a uniform probability simply means computing the ratio of two numbers: How many ways can the event we are interested in occur, divided by the total number of possible things that can happen. The denominator is clearly $2^6 = 64$, since there are six sides, each of which can be painted in two ways. To compute the numerator, we will partition the problem based on the number of side painted black.

- If 0 or 6 sides are painted black (2 cases) then there is always a mono-chromatic ring. There is only 1 way this can happen in each case.
- If 1 or 5 sides are painted black (2 cases) then there is always a mono-chromatic ring. There are 6 ways this can happen in each case.
- If 3 sides are painted black (1 case) then there is no mono-chromatic ring because there aren't four sides the same color.
- If 2 sides are painted black (1 case) then there is a mono-chromatic ring if and only if two opposing sides are painted black. There are 3 pairs of opposing sides that can be painted black.
- If 4 sides are painted black (1 case) then there is a mono-chromatic ring if and only if two opposing sides are painted white. As above there are 3 ways this can happen.
- Adding up the cases we get 2*1 + 2*6 + 2*3 = 20, so the probability is 20/64 = 5/16.

As above it is easy to see that our cases partitioned the possibilities. In the problem below this is not so clear.

Sample problem 9: Given a 10 point equilateral triangular lattice of points distance one apart, how many distinct equilateral triangles can be formed with the vertices of the lattice?

Solution: We will partition the problem based on the side lengths of the triangles. Assume the lattice is oriented with 4 points horizontal at the base

- Side length 1: there are 3+2+1 triangles pointing up and 2+1 triangles pointing down or a total of 9 for this case.
- Side length 2: there are 2+1 triangles pointing up for a total of 3 in this case.
- Side length 3: just one triangle for this case.
- Side length $3^{1/2}$: there are two triangles for this case, both tilted.
- Thus the total would appear to be 9+3+1+2 = 15.

This problem is an example of a difficult case analysis because it is not entirely clear that we have covered all the possible cases, nor is it obvious that we have counted every possibility in each case. More work is required to prove that our solution is correct.

Another situation where case analysis is difficult is when the most obvious way to split the problem up into cases results in overlapping cases. An example of this is computing the binary strings of length 3 which contain at least one 1. We know this is just $2^3 - 1 = 7$, but suppose we wanted to argue that we can construct a string with at least one 1 in three different ways, put a 1 in the first, second, or third position, and then fill the other two positions arbitrarily. This argument yields three cases, each with 4 possibilities for a total of 12 which is of course impossible. The problem is that the cases overlap because some strings have more than one 1. We could correct for this overlap, but it requires more careful analysis.

The last problem we will look at is an example of counting by cases which uses recursion, a concept we will be exploring in much more detail in the weeks ahead.

Sample problem 10: How many strings with n bits contain no adjacent 1's?

Solution: Let F(n) be the number of strings of n bits with no adjacent 1's. F(1) = 2 since there is only 1 bit, and F(2) = 4-1 = 3 since there is only one 2-bit string with adjacent 1's. For general n, it is tempting to try to analyze this by partitioning on the number of 1's in the string. This can be done, but it quickly gets very complicated. A much simpler approach is to think about constructing such a string for n > 2. The first bit is either a 0 or a 1. If it is a 0, then the remaining string can be any string of n-1 bits that doesn't contain a 1. There are F(n-1) such strings. If it is a 1, then the second bit must be 0, and the remaining string can be any string of n-2 bits, and there are F(n-2) of these. Thus we have partitioned the problem into two cases based on the first bit of the string, and we get F(n-1) strings in one case and F(n-2) strings in the other case. This leads to the following recurrence relationship for F(n):

$$F(n) = F(n-1) + F(n-2) \text{ for } n > 2 \qquad F(1) = 2,\ F(2) = 3$$

This is a very well known recurrence relationship (F(n) is the n+2nd Fibonacci number) which we probably could have guessed if we had just computed F(n) for a few more small n. The power of this technique is that it would have worked just as well if we had been analyzing strings of ternary digits with no adjacent 2's, where the numbers would have looked much more mysterious. In this situation we would get the recurrence:

$$G(n) = 2(G(n-1) + G(n-2)) \text{ for } n > 2 \qquad G(1) = 3,\ G(2) = 8$$

The next few terms are G(3) = 22, G(4) = 60, G(5) = 164. This recurrence relationship is a second order linear recurrence which can be easily solved, just like the Fibonacci sequence, using techniques we will learn later in the course.

Combinatorics: The Fine Art of Counting
Summer 2007