

Combinatorics: The Fine Art of Counting

Lecture Notes Week 3 – Counting Sets

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 $A / (B * C)$.

Set Theory and Boolean Logic

When working with sets in a particular problem, we will place the sets within a universe U which is simply the set of all objects currently under consideration, e.g. all the students in the school, all 5-digit numbers, all 8 letter strings, etc... Often U will be a finite set, but it may also be a geometric object or region, e.g. a circle or a cube.

Given some particular set of elements in U , call it A , we can define a function P_A which simply assigns to each element of U the value 1 if that element is in A , and 0 otherwise. We can interpret 1 as meaning *TRUE* and 0 as meaning *FALSE*. A function whose output is always 0 or 1 is called a **Boolean function**. The function P_A is the membership function of the set A and defines A in the sense that $A = \{x: x \in U, P_A(x) = 1\}$. $P_A(x)$ is equivalent to the statement “ x is a member of A ” and can be thought of as representing the “property” possessed by members of A . For example, if U is a set of integers and A is the subset of U which consists of multiples of 3, P_A can be thought of as representing the property of divisibility by 3.

Given any Boolean function P defined on U , we can use it to define a subset of U as above. There is a one-to-one correspondence between Boolean functions on U and subsets of U . If we fix a particular ordering on $U = \{x_1, x_2, x_3, \dots\}$ and then consider the sequence $P(x_1), P(x_2), P(x_2), \dots$ this will simply be a sequence of bits (1s and 0s). Conversely, any sequence of bits can be used to define a Boolean function on U .

Thus there is a one-to-one correspondence between each of the three classes of objects we have been considering: bit sequences, Boolean functions on U , and subsets of U . If U is a finite set of size n , there are 2^n sequences of n bits, 2^n distinct Boolean functions on U and 2^n subsets of U .

For any subset A of U , we can define the complement of A (denoted in these notes as A^c) as the set of elements in U which are not in A (if $A=U$, then $A^c=\emptyset$). Note that this set depends on U as well as A .

Given two sets A and B , the union of A and B (denoted $A \cup B$) is the set of elements which are in either A or B (or both), and the intersection of A and B (denoted $A \cap B$) is the set of elements which are in both A and B . Note the definition of these sets does not depend on U .

Each of the standard operations of Boolean logic (“and”, “or”, “not”, “if-then”, “if and only if” etc...) has an analogous operation in set theory which can be specified in terms of the complement, union, and intersection operations. See the handout titled “The 16 things you can say about A and B ” for Venn Diagrams of the different sets that can be defined in terms of two sets A and B .

Using Venn Diagrams to Count Sets

Venn diagrams are a useful starting point for organizing many types of problems. Perhaps the simplest case is where we need to correct for over-counting by excluding undesired elements. For any two sets A and B , the number of elements which are in A but not in B (denoted $A-B$) can be computed by subtracting $A \cap B$ from A , i.e. $|A-B| = |A| - |A \cap B|$.

Example: Let A be the set of poker hands which have all five cards the same suit and let B be the set of poker hands which contain a consecutive sequence in rank order (e.g. 3,4,5,6,7 not necessarily the same suit). The set of hands which are called "straight flushes" is $A \cap B$. $|A \cap B| = 4 \cdot 10$, since there are four possible suits and ten possible starting values for a straight (Ace through 10). The set of hands which are considered "flushes" is $A - B$ or, equivalently, $A - A \cap B$. Since $|A| = 4 \cdot \binom{13}{5}$ (we choose 5 out of 13 cards in one of 4 suits), the number of flushes is $|A-B| = |A| - |A \cap B| = 4 \cdot \binom{13}{5} - 40 = \mathbf{5,108}$. The number of "straights" is $|B-A| = |B| - |A \cap B| = 10 \cdot 4^5 - 40 = \mathbf{10,200}$. Thus a straight is nearly twice as likely as a flush.

Principle of Inclusion and Exclusion (PIE)

The principle of inclusion and exclusion (PIE) gives us a way to relate the size of the union of two or more sets to the sizes of the sets and their intersections. For two sets this is simply the following: $|A \cup B| = |A| + |B| - |A \cap B|$

This relationship can be easily seen by examining the Venn diagram. If we count up all the elements in A and all the elements in B , we will count each element of $A \cap B$ twice, so we subtract by the size of $A \cap B$ to correct for our over-counting. We first "include" everything, and then "exclude" the things we over-included.

Note that the formula above is equivalent to: $|A \cap B| = |A| + |B| - |A \cup B|$

The principle of inclusion/exclusion is useful whenever we know or can easily count three of the sets in the expression above in order to compute the size of the fourth. Usually we are trying to find either $A \cup B$ or $A \cap B$. Some examples:

Sample Problem #0: There are 100 students in a school, 50 don't play sports, 70 don't play music, and 40 don't do either. How many students do both?

Solution: Let U be the set of all students in the school, A be the subset who play sports and let B be the subset in the school band. We are given that $|U| = 100$, $|A^c| = 50$, $|B^c| = 30$, and $|(A \cup B)^c| = 40$. We want to compute $A \cap B$. Counting complements gives $|A| = 100 - 50 = 50$, $|B| = 100 - 30 = 70$, and $|A \cup B| = 100 - 40 = 60$. Applying the principle of inclusion/exclusion (PIE) we have: $|A \cap B| = |A| + |B| - |A \cup B| = 50 + 70 - 60 = \mathbf{20}$.

The problem above is such a simple example, that it is tempting to underestimate the value of using PIE to solve problems involving just two sets. The following problems give some more interesting cases.

Sample Problem #1: A Venn diagram depicts two circles of radius 1 with their centers a distance 1 apart. What is the area of the union of the two circles?

Solution: For a set S which is a region in the plane, let $|S|$ denote the area of the set. If A and B represent the two circles, then $|A \cup B| = |A| + |B| - |A \cap B|$ is the area we wish to find. Let $C = A \cap B$. Note that the two points in C which are on the boundaries of A and B are both distance one from the centers of A and B . Thus C contains two equilateral triangles and is the union of two sectors X and Y each with area $\pi/3$, since the central angle of each sector is 120° .

The intersection of X and Y is just the sum of two equilateral triangles with side length 1, or $2 \cdot \sqrt{3}/4 = \sqrt{3}/2$. Applying PIE gives $|C| = |XUY| = |X| + |Y| - |X \cap Y|$, so $|C| = \pi/3 + \pi/3 - \sqrt{3}/2$. Applying PIE a second time, $|A \cup B| = |A| + |B| - |A \cap B| = \pi + \pi - (2\pi/3 - \sqrt{3}/2) = \mathbf{4/3\pi + \sqrt{3}/2}$.

Sample Problem #2: How many 3 digit numbers contain at least one 2 and one 3?

Solution: Let A be the set of 3 digit numbers containing a 2 and let B be the set of numbers containing a 3. Both these sets can be counted by counting the complement within the universe U of all 3 digit numbers. $|A| = |U| - |A^c| = 900 - 8 \cdot 9^2 = 252$, and |B| is the same size. The set AUB can counted similarly: $|A \cup B| = |U| - |(A \cup B)^c| = 900 - 7 \cdot 8^2 = 452$. Applying PIE we obtain: $|A \cap B| = |A| + |B| - |A \cup B| = 252 + 252 - 452 = \mathbf{52}$. Note that this problem is essentially the same as problem #0 above.

We could have partitioned problem #2 into cases by looking at patterns of 2s and 3s, but there are a number of cases (first digit not 2 or 3, second or third digit not 2 or 3, or 223 or 332) and this quickly becomes unmanageable for larger numbers.

Generalizing problem #2 we see there are $2 \cdot (9 \cdot 10^{n-1} - 8 \cdot 9^{n-1}) - (9 \cdot 10^{n-1} - 7 \cdot 8^{n-1})$ n-digit numbers which contain both the digits 2 and 3, and similarly the number of n-letter words which contain both the letters A and B is $2 \cdot (26^n - 25^n) - (26^n - 24^n) = 26^n + 24^n - 2 \cdot 25^n$.

Sample Problem #3: How many integers between 10 and 100 (inclusive) are not divisible by 2 or 3?

Solution: Let U be the set of integers in the inclusive range 10 to 100. Let S_n denote the subset of numbers in U divisible by n. The most important step in a problem like this is to carefully avoid OBOEs (“Off-By-One-Errors”). $|U| = (100-10)+1 = 91$, $|S_2| = (100-1)/2 + 1 = 46$, and $|S_3| = (99-12)/3 + 1 = 30$. The set of numbers which are divisible by both 2 and 3 is the set of numbers divisible by their least common multiple, which is 6, so $S_2 \cap S_3 = S_6$, and $|S_6| = (96-12)/6 + 1 = 15$. Applying PIE we have $|S_2 \cup S_3| = |S_2| + |S_3| - |S_2 \cap S_3| = 46 + 30 - 15 = 61$. Counting the complement gives $|(S_2 \cup S_3)^c| = 91 - 61 = \mathbf{30}$, which is the desired result.

Caution: when counting the numbers divisible by both 2 and 3 we used their LCM, which in this case was simply their product. In general however, if the numbers are not relatively prime, using the product is incorrect, use the LCM.

Generalizing PIE to More Than Two Sets

Venn diagrams can be applied to combinations of three or more sets (although circles can't be used to depict all the possible intersections of 4 or more sets - see homework problem). The general formula for counting the union of three sets is:

$$|A \cup B \cup C| = |A| + |B| + |C| - (|A \cap B| + |B \cap C| + |C \cap A|) + |A \cap B \cap C|$$

This can be visualized as a the union of 3 overlapping circles being a flower composed of three leaves (circles), minus the three petals (intersections of two circles), plus the center (intersection of all three circles). To see why this formula is correct, note that when we add up the sizes of A, B, and C, we are over-counting elements in the intersection of two or more of these sets. We then subtract all elements contained in the intersection of two sets, but this results in under-counting elements contained in all three sets – these elements were counted three times (once for each set), then subtracted three times (once for each pair-wise intersection) - so we add them back in again.

The formula above generalizes to arbitrary numbers of sets – simply alternate adding intersections of odd numbers of sets and subtracting intersections of even numbers of sets.

This quickly becomes unwieldy however, as the number of terms becomes very large as n gets bigger, since there are 2^n distinct intersections possible among n sets.

In practice however, most of these intersections are either empty or have the same size and it is more efficient to simply analyze the over-counting and under-counting within the context of the particular problem. The principle of inclusion/exclusion is the key idea to remember, not the formulas.

Sample Problem #4: How many positive integers ≤ 2001 are multiples of 3 or 4 but not 5? (AMC12 2001 #12)

Solution: Let U be the set of positive integers in the inclusive range 1 to 2001, and let S_n denote the integers in U which are multiples of n . We want to compute $S_3 \cup S_4 - S_5$ which is equal to $S_3 \cup S_4 - ((S_3 \cup S_4) \cap S_5)$. After carefully looking at a Venn diagram, we can see that the set $(S_3 \cup S_4) \cap S_5$ equal to $(S_3 \cap S_5) + (S_4 \cap S_5) - (S_3 \cap S_4 \cap S_5)$ by applying PIE to the appropriate region of the Venn diagram (it is of course possible to prove this algebraically, but drawing a picture is strongly recommended). Now we just need to count all these various sets, remembering that if m is the LCM of j and k , then $S_j \cap S_k = S_m$. We obtain:

1. $|S_3| = 667$, $|S_4| = 500$, and $|S_3 \cap S_4| = |S_{12}| = 166$
2. $|S_3 \cup S_4| = |S_3| + |S_4| - |S_3 \cap S_4| = 667 + 500 - 166 = 1001$
3. $|S_3 \cap S_5| = |S_{15}| = 133$, $|S_4 \cap S_5| = |S_{20}| = 100$, $|S_3 \cap S_4 \cap S_5| = |S_{60}| = 33$
4. $|(S_3 \cup S_4) \cap S_5| = |S_3 \cap S_5| + |S_4 \cap S_5| - |S_3 \cap S_4 \cap S_5| = 133 + 100 - 33 = 200$
5. $|S_3 \cup S_4 - ((S_3 \cup S_4) \cap S_5)| = |S_3 \cup S_4| - |(S_3 \cup S_4) \cap S_5| = 1001 - 200 = \mathbf{801}$

Sample Problem #5: A hexagonal number H_n is defined as the number of points in a regular hexagonal lattice of points with n points on each of the six exterior sides. The first four hexagonal numbers are 1, 7, 19, and 37. What is H_n in terms of n ?

Solution: This problem can be solved using summations, but a simpler approach is to work geometrically and apply PIE to count points in overlapping regions. A hexagonal lattice can be split into easily counted (possibly overlapping) sets in a number of ways:

1. Three rhombuses of $n \times n$ points with adjacent pairs overlapping in a line of length n and all three overlapping in the center point. Each rhombus has n^2 points, the lines of overlap each have n points, and the center consists of 1 point. Applying PIE yields $H_n = 3n^2 - 3n + 1$.
2. Six equilateral triangles each containing $T_n = n(n+1)/2$ points with adjacent pairs intersecting in six lines of length n and all six overlapping in the center point. Applying PIE we have $H_n = 6 \cdot T_n - 6n + 1 = 3n^2 + 3n - 6n + 1 = 3n^2 - 3n + 1$.
3. Looking a little more closely at the situation, if we use slightly smaller triangles we can partition the hexagon into 6 non-overlapping triangles with $n-1$ points on a side plus the center point by simply subtracting the one of the lines from each triangle in the example above. This gives $H_n = 6 \cdot T_{n-1} + 1 = 6n(n-1)/2 + 1 = 3n^2 - 3n + 1$.

Note that in the second case we had six overlapping sets, but other than intersections of adjacent pairs, all of the other intersections involved just the center point. We could have computed the possible combinations of all the various other intersections (using binomial coefficients) and applied the complete formula for six sets, but it is far simpler to simply realize that when adding in the 6 triangles and subtracting 6 lines, we have counted the center point $6-6=0$ times so we need to add it back in.

How many ways can we count by choosing two?

1. **Summing Consecutive Integers:** What is $1+2+3+\dots+100$?

The value of this sum is well known. Legend has it was figured out by the young Gauss (along with many other bored math students no doubt) when his teacher told him to add

up all the numbers from 1 to 100. Gauss was so bored he decided to do the problem twice, once in increasing order, then again in decreasing order. After writing down both sums in two very long rows he noticed that he had 100 columns each with two numbers whose sum was 101, so he multiplied and divided by 2. In general we have $1+2+3+\dots+n = \mathbf{n(n+1)/2}$.

2. **Triangular Lattice:** Consider a triangular lattice with n points on each side. How many points does this lattice contain?

If we orient the triangle so that one of the sides is at the “bottom” and count horizontal rows, we find there are n rows of points and if we sum the rows starting from the top we get $1 + 2 + 3 + \dots + n = \mathbf{n(n+1)/2}$.

An alternative approach to counting the number of points in a triangular lattice is to apply the idea used in problem #5 in reverse. Consider a square $n \times n$ grid of points. It contains two triangles with n points on a side which overlap in a line of n points along the diagonal. Note that we are not concerned with the spacing between the points, just the number of points, so for our purposes the triangle formed along the diagonal of the square is the same as triangular lattice above: it has n points on each side and it can be oriented into n rows with $1, 2, 3, \dots, n$ points in each row. If T_n is the number of points in a triangular lattice with n points on each side, we can apply PIE to obtain the equation $n^2 = 2T_n - n$, and solving for T_n we get $T_n = (n^2 + n)/2 = \mathbf{n(n+1)/2}$ as above.

3. **Complete Graph:** How many edges are in the complete graph K_{n+1} on $n+1$ vertices?

K_{n+1} is a regular graph with degree n . If we pick a vertex to remove and remove all the n edges incident to that vertex we will get the complete graph K_n on the n remaining vertices. Repeating this process we will eventually remove all the edges in the graph, and if we add up all the edges removed we get the familiar sum $n + (n-1) + \dots + 1 + 0 = \mathbf{n(n+1)/2}$. Alternatively, we could have counted by simply picking (unordered) pairs of vertices in the graph. There are $\binom{n+1}{2} = (n+1)n/2!$ such pairs. Counting the edges in K_{n+1} in two different ways illuminates the combinatorial connection between the sum of the integers from 1 to n and the binomial coefficient $\binom{n+1}{2}$.

4. **Handshakes:** How many distinct handshakes are possible among $n+1$ people?

An equivalent approach to counting $\binom{n+1}{2}$ is to imagine a room with $n+1$ people entering one by one. As each person enters he or she shakes hands with everyone in the room. Counting handshakes we get $0+1+2+\dots+n = \mathbf{n(n+1)/2}$, or alternatively we could count how many distinct pairs of people we could choose to shake hands, getting $\binom{n+1}{2}$. In both cases we are counting all possible handshakes among $n+1$ people.

5. **Intersecting Lines:** What is the maximum number of intersection points possible among n lines in the plane?

The maximum is achieved when every line intersects every other (i.e. none are parallel), in which case every pair of lines intersect in a distinct point (this type of configuration is called **general position**). Since there are $\binom{n}{2}$ distinct pairs of lines, there are $\mathbf{\binom{n}{2}}$ intersection points.

6. **Ordered Partitions of n :** How many ordered triples (x,y,z) of non-negative numbers satisfy $x+y+z = n$?

Call this number $F(n)$, there are several ways to compute $F(n)$:

- a. Given any triple (x,y,z) which sum to $n-1$, the triple $(x+1,y,z)$ sums to n and all such triples are distinct and have the first element non-zero. Conversely, any triple which sums to n and has a non-zero first element corresponds to a triple which sums to $n-1$ – simply subtract 1 from the first element. Additionally, there

are $n+1$ distinct triples which have the first element 0, since we can choose y to be any number in the range 0 to n (inclusive), and then z is simply $n-y$. Thus we have the recurrence $F(n) = F(n-1) + (n+1)$. Noting that $F(0) = 1$ we find that $F(n) = (n+1) + n + (n-1) + \dots + 1 = (n+2)(n+1)/2 = \mathbf{(n+2 \ 2)}$.

- b. Alternatively, consider a sequence of n 1's. Imagine dropping two separators into the sequence one at a time. There are $n+1$ positions the first separator could land in, and $n+2$ positions for the second once the first is in place. Now sum up the 1's in 3 groups according to the placement of the separators (note that one or more groups could be empty) to get a triple (x,y,z) . Since the separators are indistinguishable, the order we drop them in doesn't matter and there are $(n+1)(n+2)/2 = \mathbf{(n+2 \ 2)}$ possible triples we will obtain. Note that every triple can be obtained in this way so we have computed $F(n)$.
- c. An equivalent approach is to imagine a row with $n+2$ open spaces. Pick two spaces and insert "+" signs, then count the size of the three (possibly empty) gaps on either side of and in between the plus signs and write these numbers in the corresponding positions to get an expression which sums to n . There are $\mathbf{(n+2 \ 2)}$ ways to choose the position of the "+" signs.
- d. If we change the problem to insist that all three numbers be non-zero, we can simply subtract 1 from each number to convert it to the original problem with n replaced by $n-3$, getting $(n-1 \ 2)$ positive ordered triples which sum to n .

7. **Bouquets of Flowers:** How many different bouquets of n roses can be made from three colors of roses (red, white, and pink)?

Assuming we don't distinguish the arrangement of flowers within the bouquet, just their color combinations, we can count bouquets in two ways:

- a. Write down the number of flowers of each color in a particular order, say red, white, then pink. This will result in a distinct ordered triple of three non-negative integers which sum to n . We have shown that there are $\mathbf{(n+2 \ 2)}$ such triples.
- b. Place twelve "blank" flowers in a row and drop two separators among them dividing the flowers into three groups, some of which may be empty. Paint all the flowers left of both separators red, all the flowers between the two separators pink, and all the flowers right of both separators white. As above, there are $(n+1)(n+2)/2 = \mathbf{(n+2 \ 2)}$ ways to do this given that the separators are indistinguishable.
- c. If we wish to add the constraint that there must be at least one flower of each color in every bouquet, we can simply start with one flower of each color and then make a bouquet of $n-3$ flowers without restriction as above, getting $(n-1 \ 2)$ different bouquets.

8. **Block Walking:** If we are walking in a city with streets running in a rectangular grid and our destination is n blocks east and 2 blocks north, how many different routes can we take assuming that we walk as short a distance as possible?

Every shortest route is exactly $n+2$ blocks long and only contains segments running north or east. Assuming we start at an intersection, there are $n+2$ places where we have the opportunity to choose between walking north or east, and we can only choose to walk north 2 times. We can describe any such route by a string of $n+2$ N's and E's which contains exactly 2 N's. Thus there are $\mathbf{(n+2 \ 2)}$ possible routes.