


Exploring Combinatorics: Inclusion–exclusion*

Like the textbook (section 3.7), we introduce inclusion–exclusion with a silly problem.

Problem 1. *A group of students attends three kinds of classes: art classes, biology classes, and chemistry classes¹. It is known that in this group 11 take an art class, 8 take biology, and 4 take chemistry. Some students attend more than one class. Four students take both art and biology, two take both art and chemistry, two take biology and chemistry, and one person attends all three classes. How many students are there taking at least one class?*

A natural impulse is to add the number of students in each of class, i.e., estimate the number of students by $|A| + |B| + |C|$, where A , B and C are the sets of the students in the art, biology and chemistry classes respectively. Each person who takes just one class is counted exactly once by this sum, but those who take several classes are counted more than once (*overcounted*). The Venn diagram below illustrates the issue:




Count multiplicity for $|A| + |B| + |C|$

The integers in the regions of the Venn diagram record the number of times $|A| + |B| + |C|$ counts students that belong to the region. For example, the 1 in the bottom left signifies that students taking only art are counted once, whereas 3 means that the student taking all three classes is counted 3 times.

To correct the overcount it is natural to subtract the number of those students that are counted twice. We thus obtain the expression $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C|$ and the following Venn diagram:

*These notes are available from the course webpage, and directly from http://www.borisbukh.org/CombinatoricsSpring13/notes_inclusion_exclusion.pdf

¹Of course all the students take math classes.



Count multiplicity for $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C|$

Hence, $|A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C|$ *undercounts* the number of students, and we need to correct for the student taking all the subjects. We thus arrive to the following formula:

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

The process by which we arrived at the formula is called inclusion–exclusion, and the resulting formula is generally known as the *inclusion–exclusion principle* (for three sets). Plugging the values into the formula, we find that there are $11 + 8 + 4 - 4 - 2 - 2 + 1 = 18$ students in all.

The inclusion–exclusion principle is not restricted to counting elements of sets. For instance, if A , B and C are three regions in the plane, of (finite) areas $\text{area}(A)$, $\text{area}(B)$, $\text{area}(C)$ respectively, then we can find the area of $A \cup B \cup C$ in the similar manner: overestimate by $\text{area}(A) + \text{area}(B) + \text{area}(C)$, then underestimate by $\text{area}(A) + \text{area}(B) + \text{area}(C) - \text{area}(A \cup B) - \text{area}(A \cap C) - \text{area}(B \cap C)$, and finally getting it right,

$$\begin{aligned} \text{area}(A \cup B \cup C) &= \text{area}(A) + \text{area}(B) + \text{area}(C) - \text{area}(A \cap B) - \text{area}(A \cap C) \\ &\quad - \text{area}(B \cap C) + \text{area}(A \cap B \cap C). \end{aligned} \tag{1}$$

Remark. An astute reader will notice that similar formulas exist for length, volume and some other quantities. For this and other similar reasons, a popular notation is $|A|$ for the area (or volume, or length, depending on the dimension) of A . Area, volume, length and the number of elements are examples of *measures*, and the inclusion–exclusion principle holds for all measures.

Area of a spherical triangle

As an application, we can derive a handy formula for the area of a spherical triangle. Unlike a triangle in the plane, the angles of a triangle on the sphere do not add up to 180° . The angles of a spherical triangle in fact determine the triangle, much like the ordinary triangle is determined by its sidelengths.

In the plane a triangle has three sides, which are line segments. A line segment is the shortest path between its endpoints. On the sphere, the shortest paths between points are arcs of great circles. Similarly, in the plane a triangle is an intersection of three

halfplanes, which are sets bounded by a line. Since great circles bound halfspheres, this motivates the following definition.

Definition. A spherical triangle is an intersection of three halfspheres. We allow only triangles of positive area.

Consider a sphere S and let A, B, C be three halfspheres whose intersection is non-empty, and let $T = A \cap B \cap C$ be the spherical triangle that they determine. We rewrite (1) as


$$\text{area}(T) = \text{area}(A \cup B \cup C) - \text{area}(A) - \text{area}(B) - \text{area}(C) + \text{area}(A \cap B) + \text{area}(A \cap C) + \text{area}(B \cap C).$$

We shall evaluate the terms on the right of the equation in order.

The term $\text{area}(A \cup B \cup C)$: Since halfspheres are pretty big, and their union is bigger still, it is more convenient to think of a smaller set $S \setminus (A \cup B \cup C)$ rather than $A \cup B \cup C$. A point does not belong to halfsphere A if and only if the antipodal point² belongs to A . Similarly, for halfspheres B and C . Thus, $S \setminus (A \cup B \cup C)$ consists of points that are antipodal to $A \cap B \cap C$. Hence, $\text{area}(A \cup B \cup C) = \text{area}(S) - \text{area}(T)$.

The terms $\text{area}(A)$, $\text{area}(B)$, and $\text{area}(C)$: These are all equal to $\frac{1}{2} \text{area}(S)$ since A, B and C are halfspheres.

The terms $\text{area}(A \cup B)$, $\text{area}(A \cup C)$ and $\text{area}(B \cup C)$: Each of these terms measures the area of a *lune* (also known as a 2-gon) formed by bounding great circles of respective halfspheres.



A lune with angle α

A lune is determined by the angle at which the two halfspheres meet. We claim that the area of a lune is proportional to its angle. More accurately, our claim is that if the lune's angle is α , then its area is $\text{area}(\text{lune}) = \frac{\alpha}{360^\circ} \text{area}(S)$. It is easiest to see this in the case α divides 360° . If $\alpha = \frac{1}{q} 360^\circ$ for an integer q , then successively rotating the lune by α degrees, we can cover the sphere by q rotated copies of the lune, which implies that $\text{area}(\text{lune}) = \text{area}(S)/q$. Similarly, if $\alpha = \frac{p}{q} 360^\circ$ for some rational number $\frac{p}{q}$, then by successively rotating the lune the angle $\frac{1}{q} 360^\circ$ we will obtain a collection of q copies of the original lune that cover each points of the sphere p times. Thus $p \text{area}(S) = q \text{area}(\text{lune})$. The case of an irrational angle α follows by taking a limit³.

²If $x \in S$, then the antipodal point to x is the other point of the sphere that lies on the ray from x directed to the sphere's center. If the sphere is centered at the origin, the antipode of x is $-x$.

³Formally, this is cheating because we have neither defined surface area on a sphere, nor shown that it behaves properly when we take the limits.

Conclusion: Putting all of these together, we obtain that if angles in a spherical triangle are α , β , and γ , then

$$\text{area}(T) = (\text{area}(S) - \text{area}(T)) - 3 \cdot \frac{1}{2} \text{area}(S) + \text{area}(S) \frac{\alpha + \beta + \gamma}{360^\circ}$$

or equivalently,

$$\text{area}(T) = \frac{1}{2} \text{area}(S) \frac{\alpha + \beta + \gamma - 180^\circ}{360^\circ}.$$

If we express the angles in radians, and recall that the area of a sphere of radius r is $4\pi r^2$, then we obtain a very neat formula

$$\text{area}(T) = r^2(\alpha + \beta + \gamma - \pi).$$

The quantity $\alpha + \beta + \gamma - \pi$ is known as *spherical excess*.

The inclusion–exclusion principle in general

In this section we generalize what we did for three sets to any number of sets. Suppose $A_1, A_2, \dots, A_n$ are n finite sets, we know the sizes of $|A_i|$, $|A_i \cap A_j|$, $|A_i \cap A_j \cap A_k|$, etc, for all $i, j, k, \dots \in [n]$, and we wish to compute the size of $A_1 \cup A_2 \cup \dots \cup A_n$. As above, the natural first step is to estimate the size of the union by $|A_1| + |A_2| + \dots + |A_n|$. As above, this is an overestimate. For example, all the elements of $\bigcup_i A_i$ that belong to two of the A_i 's are counted twice. Thus, it is natural to estimate $|A|$ by

$$\sum_i |A_i| - \sum_{i < j} |A_i \cap A_j|. \quad (2)$$

The new sum counts elements that occur only in one or two A_i correctly, but it undercounts all the other elements. In particular, the elements that occur in three of the A_i 's are counted with total multiplicity 0 in (2), namely 3 times in $\sum |A_i|$, and -3 times in $\sum_{i,j} |A_i \cap A_j|$. Thus, we consider

$$\sum_i |A_i| - \sum_{i < j} |A_i \cap A_j| + \sum_{i < j < k} |A_i \cap A_j \cap A_k|.$$

Whereas the elements that occur in 1, 2 or 3 sets A_i are counted correctly by this sum, all the others are not. By this point, the pattern should be obvious, and one is willing to hazard a guess that $|A|$ is equal to a certain sum of sums. The main difficulty now is not so much to prove this conjecture, but to find a sane notation in which we can express it. Our choice of notation goes a step beyond of what our textbook calls “devilish” (page 100).

For a non-empty set $I \subset [n]$ we define

$$A_I \stackrel{\text{def}}{=} \bigcap_{i \in I} A_i.$$

For example, $A_{\{2,5,6\}}$ is a shorthand for $A_2 \cap A_5 \cap A_6$.

We can then write the inclusion–exclusion principle as

$$\left| \bigcup_{i \in [n]} A_i \right| = \sum_{\substack{I \subseteq [n] \\ I \neq \emptyset}} (-1)^{|I|+1} |A_I|. \quad (3)$$

We can easily convert our previous back-and-forth reasoning into a formal proof of (3). Let x be an arbitrary element of $A_1 \cup A_2 \cup \dots \cup A_n$. Suppose x belongs to some k sets among $A_1, A_2, \dots, A_n$. Without loss of generality, x belongs to all of $A_1, A_2, \dots, A_k$, but not to any of $A_{k+1}, A_{k+2}, \dots, A_n$. The element x is counted precisely once by the left side of (3). The contribution of x to the right side of (3) is

$$\sum_{\substack{I \subseteq [n] \\ I \neq \emptyset \\ x \in A_I}} (-1)^{|I|+1} = \sum_{\substack{I \subseteq [k] \\ I \neq \emptyset}} (-1)^{|I|+1} = 1 + \sum_{I \subseteq [k]} (-1)^{|I|+1}.$$

Since there are as many sets of odd size as there are sets of the even size (first lecture; also Proposition 3.1.3 in the textbook), it follows that $\sum_{I \subseteq [k]} (-1)^{|I|+1} = 0$, and hence x is counted exactly once by the right side (3). Since x is arbitrary, (3) is proven.

Probability that two numbers have no common factors

As an application of the general inclusion–exclusion principle we consider the following mind-boggling question:

How likely are two random natural numbers to have a common factor?

As the question stands, it makes no sense. If we pick a random natural number so that each natural number is equally likely, how likely are we to pick 7? Certainly the probability of picking 7 is zero, for there are infinitely many other numbers to pick from. However, in that case the probability of picking any other number must be zero too, and so the probability of picking anything at all must be zero! We have reached the state of confusion that is frighteningly familiar to anybody who has ever done mathematics: we are unsure of what words mean, and definitions seem elusive. To bring clarity we change the question:

Let n be a large natural number, and consider the n^2 pairs of numbers (a, b) where $a \in [n]$ and $b \in [n]$. How many of these pairs have no common factor?

This question is much better than the question that we started with. We replaced the imprecise notion of a “random natural number” by a concrete “random number from $[n]$ ”. We have a question; let’s find the answer!

We start by naming the set whose elements we want to count. We call it G , or in symbols

$$G \stackrel{\text{def}}{=} \{(a, b) \in [n]^2 : \text{there is no } d \text{ such that } d \mid a \text{ and } d \mid b\}.$$

We will apply the inclusion–exclusion principle to the *complement* of G . This step is similar to the way one counts derangements (Section 3.8 of the textbook). So, we let

$$B \stackrel{\text{def}}{=} [n]^2 \setminus G = \{(a, b) \in [n]^2 : \text{there is a } d \text{ such that } d \mid a \text{ and } d \mid b\}.$$

(For the curious: the letters G and B stand for “good” and “bad”. There are no shades of gray in mathematics.)

For a prime number p we define

$$A_p \stackrel{\text{def}}{=} \{(a, b) \in [n] : p \mid a \text{ and } p \mid b\}. \quad (4)$$

We then have $B = \bigcup_{p \leq n} A_p$ where the union is over all the primes p that do not exceed n . Therefore we are in the position to apply the inclusion–exclusion principle to compute $|B|$. We let P be the set of all the primes not exceeding n . In our application of the inclusion–exclusion principle the sets are indexed by the elements of P rather than by the elements of $[n]$, and so we obtain the formula

$$|B| = \sum_{\substack{P' \subseteq P \\ P' \neq \emptyset}} (-1)^{|P'|+1} |A_{P'}|.$$

Computation of $|A_P|$ is scarier than it is hard. If we extend the definition (4) to allow non-prime subscripts,

$$A_d \stackrel{\text{def}}{=} \{(a, b) \in [n] : d \mid a \text{ and } d \mid b\} \quad \text{for any integer } d,$$

and observe that a number is divisible by every prime in P if and only if it is divisible by $\prod_{p \in P} p$ (fundamental theorem of arithmetic), then we arrive at a very simple relation

$$A_P = A_{\prod_{p \in P} p}.$$

Hence, our next task is to compute the size of A_d where d is some number. A pair in A_d consists of two multiples of d not exceeding n . As there are $\lfloor \frac{n}{d} \rfloor$ multiples of d not exceeding n , it follows that $|A_d| = \lfloor \frac{n}{d} \rfloor^2$. Plugging this into the formula for $|B|$ we obtain

$$|B| = \sum_{k \geq 1} \sum_{\substack{d \text{ is a product} \\ \text{of some } k \text{ distinct} \\ \text{primes from } P}} (-1)^{k+1} \left\lfloor \frac{n}{d} \right\rfloor^2. \quad (5)$$

There are two simplifications that we can effect. First, we may restrict the summation to $d \leq n$, for $\lfloor \frac{n}{d} \rfloor$ vanishes for larger values of d . Second, we can introduce a notation which will take care of tracking the pesky minus signs for us:

$$\mu(d) \stackrel{\text{def}}{=} \begin{cases} (-1)^k & \text{if } d = p_1 p_2 \cdots p_k \text{ is a product of } k \text{ distinct primes,} \\ 0 & \text{if } p^2 \mid d \text{ for some prime } p. \end{cases}$$

This definition of the function⁴ μ is clever; by defining $\mu(d)$ to be zero when d is not a product of distinct primes, we can extend the summation to all values of d . Note that $\mu(1) = 1$ since 1 is a product of zero distinct primes; this will come handy later. These two simplifications together transform (5) into

$$|B| = \sum_{2 \leq d \leq n} -\mu(d) \left\lfloor \frac{n}{d} \right\rfloor^2.$$

⁴The function μ is known as the Möbius function. Call it by that name when you meet it next time.

The sum above looks pretty, save for the minus sign and the missing term $d = 1$. Replacing $|B|$ by $|G|$ fixes these blemishes, and also shifts our attention from B (whose role is auxiliary) to G (which is our primary interest):

$$|G| = n^2 - |B| = n^2 + \sum_{2 \leq d \leq n} \mu(d) \left\lfloor \frac{n}{d} \right\rfloor^2 = \sum_{1 \leq d \leq n} \mu(d) \left\lfloor \frac{n}{d} \right\rfloor^2.$$

If we could factor out the n^2 from the sum, we could consider ourselves done. Only the floor signs stand in our way. We focus on them next. If x is any real number, then $x = \lfloor x \rfloor + \{x\}$ where $\{x\}$ is the *fractional part* of x , better known as the “stuff after the decimal point”, as in $\{217.446\} = 0.446$. We compute

$$\lfloor x \rfloor^2 - x^2 = (x - \{x\})^2 - x^2 = -2x\{x\} + \{x\}^2 = O(x)$$

since $\{x\} \leq 1$. Thus,

$$\begin{aligned} |G| &= \sum_{d \leq n} \mu(d) \left(\frac{n^2}{d^2} + O\left(\frac{n}{d}\right) \right) \\ &= n^2 \sum_{d \leq n} \frac{\mu(d)^2}{d^2} + O\left(n \sum_{d \leq n} \frac{|\mu(d)|}{d}\right). \end{aligned}$$

Because $|\mu(d)| \leq 1$ for all d , the sum in the big-oh term is small,

$$\sum_{d \leq n} \frac{|\mu(d)|}{d} \leq \sum_{d \leq n} \frac{1}{d} = H_n = O(\log n).$$

As $n \rightarrow \infty$ the sum $\sum_{d \leq n} \frac{\mu(d)}{d^2}$ converges to some constant $\mathcal{M} \stackrel{\text{def}}{=} \sum_{d=1}^{\infty} \frac{\mu(d)}{d^2}$. One can even show that $\left| \mathcal{M} - \sum_{d \leq n} \frac{\mu(d)}{d^2} \right| \leq 1/n$. Amazingly, it turns out that $\mathcal{M} = 6/\pi^2$! Putting all together we announce the answer

$$|G| = \frac{6}{\pi^2} \cdot n^2 + O(n \log n).$$

We have arrived at the answer:

The probability that two random natural numbers are coprime is $6/\pi^2$.

Computing $\mathcal{M}$

In this section we derive the amazing identity $\mathcal{M} = 6/\pi^2$ and show that the difference $\mathcal{M} - \sum_{d \leq n} \frac{\mu(d)}{d^2}$ is indeed bounded by $1/n$. We prove the latter claim first:

$$\left| \mathcal{M} - \sum_{d \leq n} \frac{\mu(d)}{d^2} \right| = \left| \sum_{d > n} \frac{\mu(d)}{d^2} \right| \leq \sum_{d > n} \frac{1}{d^2} = \int_n^{\infty} \frac{dx}{\lceil x \rceil^2} \leq \int_n^{\infty} \frac{dx}{x^2} = \frac{1}{n}.$$

Remarkably the infinite sum that defines $\mathcal{M}$ can be rewritten as an infinite product. Let $p_1, p_2, p_3, \dots$ be all the prime numbers. Then we have the identity

$$\mathcal{M} = \sum_{d=1}^{\infty} \frac{\mu(d)}{d^2} = \left(1 - \frac{1}{p_1^2}\right) \left(1 - \frac{1}{p_2^2}\right) \left(1 - \frac{1}{p_3^2}\right) \cdots$$

The identity holds because when we expand the product⁵ we obtain all possible terms of the form $(-1)^k / p_{i_1}^2 p_{i_2}^2 \cdots p_{i_k}^2$ where $p_{i_1}, p_{i_2}, \dots, p_{i_k}$ are distinct primes. These are precisely the non-zero terms in $\sum_d \mu(d)/d^2$.

A similar expansion of $1/\mathcal{M}$ yields a sum without the μ function:

$$\begin{aligned} 1/\mathcal{M} &= \left(\frac{1}{1 - p_1^{-2}}\right) \left(\frac{1}{1 - p_2^{-2}}\right) \cdots = \left(1 + \frac{1}{p_1^2} + \frac{1}{p_1^4} + \cdots\right) \left(1 + \frac{1}{p_2^2} + \frac{1}{p_2^4} + \cdots\right) \cdots \\ &= 1 + \frac{1}{2^2} + \frac{1}{3^2} + \frac{1}{4^2} + \frac{1}{5^2} + \cdots \end{aligned}$$

The last sum has a long history. Its evaluation was a famous problem until 1735, when Euler finally succeeded to show that it is equal to $\pi^2/6$. You can read a version of the proof in section 12.7 of the textbook.

⁵We cheat again by ignoring convergence of the infinite product, and legality of formal expansion of the product. It can be justified by considering truncated products, and appealing to absolute convergence of the series $\sum \mu(d)/d^2$.