

Discrete Mathematics lecture notes 9-2

November 1, 2013

26. The Inclusion-Exclusion Principle

Let's finish counting-for-the-sake-of-counting portion of the course¹ with a powerful method of computing the size of finite union of finite sets. We begin with a case we've already seen:

Proposition 1. For A, B any sets, $|A \cup B| = |A| + |B| - |A \cap B|$.

Proof. Set $C := A \cap B$, $A' := A - C$, and $B' := B - C$. Clearly $A = A' \amalg C$ and $B = B' \amalg C$ and $A \cup B = A' \amalg B' \amalg C$, so $|A| = |A'| + |C|$, $|B| = |B'| + |C|$, and $|A \cup B| = |A'| + |B'| + |C|$. We then have $|A| + |B| = |A'| + |B'| + 2|C| = |A \cup B| + |C|$, which is just a restatement of the desired result. \square

Let's actually think about the proof of Proposition 1 for a moment: Suppose we wanted to calculate the size of $A \cup B$ by counting the elements in A and adding that number to the size of B . This is a good first approximation, except we have *overcounted*: Elements that lie in $A \cap B$ have been counted twice by this procedure. So we need to correct our first guess by subtracting $|A \cap B|$ to undo our doublecounting. The happy observation is, in this case at least, this one correction is all that is needed.

Proposition 2. For A, B, C any sets, $|A \cup B \cup C| = |A| + |B| + |C| - |A \cap B| - |A \cap C| - |B \cap C| + |A \cap B \cap C|$.

Proof. Exercise for you. Try to construct a proof based on the same idea of the proof of Proposition 1, say by considering $|A \cap B \cap C|$, $|(A \cap B) - (A \cap B \cap C)|$, etc., then writing $A \cup B \cup C$ as a disjoint union of $A \cap B \cap C$ with $(A \cap B) - (A \cap B \cap C)$, etc. \square

Consider the following informal argument:

- Our first guess for the size of $A \cup B \cup C$ is just to add the sizes of A , B , and C individually. This first guess is an overestimate, in that any element of $A \cap B$ is counted once for being in A and a second time for being in B , and similarly for a point in $A \cap C$ or $B \cap C$.
- Therefore we make our second guess: Take the sum of the sizes of A , B , and C , and from that subtract the sizes of the pairwise intersections $A \cap B$, $A \cap C$, and $B \cap C$. If you follow the counting process, you'll see that we've now counted each element that is *only* A exactly once (and similarly for elements only in B or C); while if $x \in A \cap B$ but $x \notin C$, we've counted x twice and then subtracted once. It looks like our second guess might actually be correct, until we consider points in $A \cap B \cap C$: Such a point is counted three times (in the first sum of the sizes of A , B , and C), but then subtracted three times again (since it lies in each pairwise intersection). If we stopped with subtracting the sizes of the pairwise intersections of subsets, we would miss the threefold intersection $A \cap B \cap C$.
- So we need to correct our correction by adding back elements that have been unjustly cancelled, hence the final term $+|A \cap B \cap C|$.

One could try to generalize this result, and the proof technique along with it. Of course, we'll have to avail ourselves of the rigor of mathematical formalism if we want our arguments to be anything more than vigorous handwaving.

Theorem 3 (Inclusion-Exclusion). Let A_1, A_2, \dots, A_n be a set of n finite sets.² Then

$$\left| \bigcup_{i=1}^n A_i \right| = \left(\sum_{1 \leq i \leq n} |A_i| \right) - \left(\sum_{1 \leq i < j \leq n} |A_i \cap A_j| \right) + \left(\sum_{1 \leq i < j < k \leq n} |A_i \cap A_j \cap A_k| \right) - \dots + (-1)^{n-1} \left| \bigcap_{i=1}^n A_i \right|.$$

¹Well... maybe the counting-for-the-sake-of-counting-with-no-extra-machinery section.

²One can formulate infinite generalizations of this theorem; we will not.

We'll give two techniques for proving this Theorem. The first is an inductive argument, which requires the following technical

Lemma 4. *Let $\{X_\alpha\}_{\alpha \in A}$ be a set of sets and Y some other set.³ Then*

$$Y \cap \left(\bigcup_{\alpha \in A} X_\alpha \right) = \bigcup_{\alpha \in A} (Y \cap X_\alpha).$$

In other words, intersections distribute over unions.⁴

Proof. Two sets are equal if and only if they have the same elements. If $z \in Y \cap \left(\bigcup_{\alpha \in A} X_\alpha \right)$, then $z \in Y$ and $z \in \bigcup_{\alpha \in A} X_\alpha$, so there is some particular $\alpha_0 \in A$ such that $z \in X_{\alpha_0}$. But then $z \in Y \cap X_{\alpha_0}$, so $z \in \bigcup_{\alpha \in A} (Y \cap X_\alpha)$.
Going the other direction is a matter of reading the above argument backwards. \square

Inclusion-Exclusion proof, using induction. Let $\mathcal{P}(n)$ be the statement “For any set of n finite sets A_1, \dots, A_n , the size of the union $\left| \bigcup_{i=1}^n A_i \right|$ is computed by the conclusion of the Theorem. Clearly $\mathcal{P}(0)$, $\mathcal{P}(1)$, and $\mathcal{P}(2)$ are all true.⁵

Suppose now that $\mathcal{P}(k)$ is true, and fix a set $\{A_i\}_{i=1}^{k+1}$ of $k+1$ finite set. We write

$$U := \bigcup_{i=1}^{k+1} A_i = B \cup A_{k+1} \quad \text{for } B := \left(\bigcup_{i=1}^k A_i \right).$$

Since we know $\mathcal{P}(2)$ is true, we have

$$|U| = |B| + |A_{k+1}| - |A_{k+1} \cap B|.$$

B is the union of k finite sets, so in principle we can compute its size by IE and the inductive hypothesis. Similarly, A_{k+1} is a term that shows us in the statement of IE, so let's focus for a moment on the last term:

$$A_{k+1} \cap B = A_{k+1} \cap \left(\bigcup_{i=1}^k A_i \right) = \bigcup_{i=1}^k (A_{k+1} \cap A_i)$$

where the last equality follows by the preceding technical Lemma. Writing $C_i := A_{k+1} \cap A_i$, we see that we can compute $|A_{k+1} \cap B|$ as $\left| \bigcup_{i=1}^k C_i \right|$, which is a union of k finite sets, so our inductive hypothesis applies.

We now write out explicitly the inductive hypotheses:

$$|B| = \left| \bigcup_{i=1}^k A_i \right| = \left(\sum_{1 \leq i \leq k} |A_i| \right) - \left(\sum_{1 \leq i < j \leq k} |A_i \cap A_j| \right) - + \dots + (-1)^{k-1} \left| \bigcap_{i=1}^k A_i \right|$$

and

$$|A_{k+1} \cap B| = \left| \bigcup_{i=1}^k C_i \right| = \left(\sum_{1 \leq i \leq k} |C_i| \right) - \left(\sum_{1 \leq i < j \leq k} |C_i \cap C_j| \right) - + \dots + (-1)^{k-1} \left| \bigcap_{i=1}^k C_i \right|.$$

Substituting in $C_i := A_{k+1} \cap A_i$, we get

$$|A_{k+1} \cap B| = \left(\sum_{1 \leq i \leq k} |A_{k+1} \cap A_i| \right) - \left(\sum_{1 \leq i < j \leq k} |A_{k+1} \cap A_i \cap A_j| \right) - + \dots + (-1)^{k-1} \left| \bigcap_{i=1}^k A_{k+1} \cap A_i \right|.$$

³Note there are no finiteness conditions here.

⁴The word “distribute” is meant suggestively: Can you draw an analogy with the operations of multiplication and addition?

⁵By, in reverse order, Proposition 1; obviousness; and vacuity.

Note two things: First, a generic term in the second (and all future) expression here should look like $A_{k+1} \cap A_i \cap A_{k+1} \cap A_j$, but we've simplified by only writing one intersection with A_{k+1} . Second, the *sign* of these terms appears to be off from what we'd expect in the conclusion of IE: The terms with an even number of intersected sets have a + sign in front, while those with an odd number intersected have a -. Luckily, we're saved by the fact that $|A_{k+1} \cap B|$ appears as something being *subtracted* in our calculate $|U| = |B| + |A_{k+1}| - |A_{k+1} \cap B|$. Thus the sign of the terms involving an intersection with A_{k+1} are what we'd expect.

At this point it is simply a matter of plugging in our computations of $|B|$ and $|A_{k+1} \cap B|$ into this formula to see that the whole union U 's size is computed by IE. \square

This proof is fairly typical of inductive proofs in that it is quite literally making the obvious choice one step at a time. There was no flash of insight needed; instead we just needed to write down what we know from before, what we want to show now, and how these are related to each other. Just as it is fairly straightforward, it is also almost entirely unenlightening.⁶ Let's see if we can take a different viewpoint that will give some insight about what's going on. The tradeoff for this proof, again in a fairly typical exchange, is that it may seem far less motivated at first blush.

Inclusion-Exclusion proof, using combinatorics. Combinatorial proofs are all about obtaining relationships between *numbers* by considering *sets*. One common technique that we've already encountered is to introduce two different partitions on a set, then count the total size of the set in the two corresponding different ways; we know the end results must be the same number, but the *interpretations* of each method can vary dramatically. Here, we'll try the same thing: From the finite set of finite sets whose union we're interested in, we'll construct a large (but finite) list of numbers. When we partition these numbers in one way, add within the partition, then add the results, we'll get the size of the union; when we partition the number a different way, sum within that different partition, and then sum the results, we'll get the expression of IE. This is just writing one number as a sum in two different ways, so we'll conclude that the size of the union is given by the IE expression.

Explicitly: Fix $\{A_i\}_{i=1}^n$ a set of n finite sets. We construct a table with rows indexed by the elements of $U := \bigcup_{i=1}^n A_i$ and whose rows are indexed by the intersections of the A_i . In the $\left(x, \bigcap_{j=1}^k A_{i_j}\right)$ -position⁷, put a 0 if $x \notin \bigcap_{j=1}^k A_{i_j}$, put a 1 if $x \in \bigcap_{j=1}^k A_{i_j}$ and k is odd, and a -1 if $x \in \bigcap_{j=1}^k A_{i_j}$ and k is even.

For example, let $A = \{x, y, z\}$, $B = \{y, w\}$, and $C = \{y, z, w, u\}$. The corresponding array looks like (omitting 0s)

	A	B	C	$A \cap B$	$A \cap C$	$B \cap C$	$A \cap B \cap C$	row sum
x	1							1
y	1	1	1	-1	-1	-1	1	1
z	1		1		-1			1
w		1	1			-1		1
u			1					1
column sum	3	2	4	-1	-2	-2	1	

Let's make two observations about this example:

- The sum of entries in row is equal to 1.
- The sum of the entries in each column is equal to the size of the set indexing that column, up to a sign. Moreover, the sign on a given column is positive if it corresponds to an intersection of an odd

⁶Not really. See if you can find connections with the proof of Pascal's Lemma in the last part of the proof that I glossed over.

⁷I.e., the position whose row is indexed by x and whose column is indexed by the intersection of the k set $A_{i_1} \cap A_{i_2} \cap \dots \cap A_{i_k}$.

number of sets and negative if an even number. You might recognize these as the terms and signs that appear in the conclusion of IE.

As there are, by construction, as many rows as there are elements of $A \cup B \cup C$, the first observation indicates that the sum of *all* the entries in the array, computed by first summing the rows and then summing the row sums, is $|A \cup B \cup C|$. On the other hand if we compute the sum of all the array entries by grouping along columns and then summing the column sums, we get exactly the alternating sum of the sizes of the intersection. Therefore, to prove IE in general, we just have to prove that both of our observations generally obtain.

The second point is more immediate: For a given column, there is one nonzero entry for each element of the corresponding set. Moreover, by construction of the table, all of these nonzero entries will have the same sign, which is also the one required by the conclusion of IE. Therefore summing by columns first gives the right hand side of the IE equality.

The first takes a little more work: Fix some element $x \in U$, and suppose that x is in exactly ℓ of the $\{A_i\}$. In the first block of the x th row, corresponding to the single-intersection of the $\{A_i\}$ (i.e., the A_i themselves), there must therefore be exactly ℓ nonzero entries, all of which are 1. In the next block, corresponding to the double intersection, the only nonzero entries will correspond to pairs of subsets which both contain x . By definition of the binomial coefficients, there are $\binom{\ell}{2}$ of these, and each has a negative sign. The sum in the third block is similarly $+\binom{\ell}{3}$, in the fourth $-\binom{\ell}{4}$, etc. We conclude that the sum of the entries in the x th row is

$$S := \sum_{i=1}^{\ell} \binom{\ell}{i} (-1)^{i-1},$$

corresponding to our choice of partition of the row.

The binomial theorem tells us that

$$\left((-1) + 1\right)^{\ell} = \sum_{j=0}^{\ell} \binom{\ell}{j} (-1)^j (1)^{\ell-j} = \sum_{j=0}^{\ell} \binom{\ell}{j} (-1)^j = 1 - S,$$

where the last equality corresponds to splitting off the $j = 0$ case to give the 1, and then observing that the signs in S and the binomial expression are opposite. On the other hand $(-1) + 1 = 0$, so we've just shown that $0 = 1 - S$ or $S = 1$. This completes the proof. □