

Counting

Based on a handout by Mehran Sahami and Chris Piech

In-lecture: Section 1 (Sum Rule), Section 2 (Product Rule).

Pre-recorded (post-lecture): Section 3 onwards.

Questions at end of lecture (go over by yourself): Section 1.2 (Example 2: Datacenters), Section 2.2 (Example 4: Dice rolls), Lecture Notes 2 Example 4.

Although you may have thought you had a pretty good grasp on the notion of counting at the age of three, it turns out that you had to wait until now to learn how to really count. Aren't you glad you took this class now?! But seriously, below we present some properties related to counting which you may find helpful in the future.

1 Sum Rule

Sum Rule of Counting: If the outcome of an experiment can either be one of m outcomes or one of n outcomes, where none of the outcomes in the set of m outcomes is the same as the any of the outcomes in the set of n outcomes, then there are $m + n$ possible outcomes of the experiment.

Rewritten using set notation, the Sum Rule states that if the outcomes of an experiment can either be drawn from set A or set B , where $|A| = m$ and $|B| = n$, and $A \cap B = \emptyset$, then the number of outcomes of the experiment is $|A| + |B| = m + n$.

1.1 Example 1

Problem: You have the choice of either rolling a 6-sided die or drawing a card from a standard 52-card deck (but you have to do exactly one of them). How many possible outcomes do you have?

Solution: Since the outcomes of the die roll and the outcomes of the card draw are mutually exclusive (e.g., a roll of 1 can't possibly be considered an Ace of Spades, and vice versa), we can use the Sum Rule to get 58 ($= 6 + 52$) outcomes.

1.2 Example 2

Problem: You are running an on-line social networking application which has its distributed servers housed in two different data centers, one in San Francisco and the other in Boston. The San Francisco data center has 100 servers in it and the Boston data center has 50 servers in it. If a server request is sent to the application, how large is the set of servers it may get routed to?

Solution: Since the request can be sent to either of the two data centers and none of the machines in either data center are the same, the Sum Rule of Counting applies. Using this rule, we know that the request could potentially be routed to any of the 150 ($= 100 + 50$) servers.

2 Product Rule

Product Rule of Counting: If an experiment has two parts, where the first part can result in one of m outcomes and the second part can result in one of n outcomes regardless of the outcome of the first part, then the total number of outcomes for the experiment is mn .

Rewritten using set notation, the Product Rule states that if an experiment with two parts has an outcome from set A in the first part, where $|A| = m$, and an outcome from set B in the second part (regardless of the outcome of the first part), where $|B| = n$, then the total number of outcomes of the experiment is $|A||B| = mn$.

Note that the Product Rule for Counting is very similar to "the basic principle of counting" given in the Ross textbook.

2.1 Example 3

Problem: Flip two coins, one after the other. How many possible outcomes are there?

Solution: Define our outcome as a an ordering of two coin flips, where each coin flip has two outcomes ("Heads" or "Tails"). By the product rule, the total number of potential outcomes is therefore $4 (= 2 \times 2)$ outcomes. The possible outcomes can also be enumerated as: $\{ (H, H), (H, T), (T, H), (T, T) \}$, where the first entry of each tuple is the outcome of the first and second entries of each tuple are the outcomes of the first and second coin flips, respectively.

2.2 Example 4

Problem: Two 6-sided dice, with faces numbered 1 through 6, are rolled. How many possible outcomes of the roll are there?

Solution: Note that we are not concerned with the total value of the two dice, but rather the set of all explicit outcomes of the rolls. Since the first die¹ can come up with 6 possible values and the second die similarly can have 6 possible values (regardless of what appeared on the first die), the total number of potential outcomes is $36 (= 6 \times 6)$. These possible outcomes are explicitly listed below as a series of pairs, denoting the values rolled on the pair of dice:

(1, 1)	(1, 2)	(1, 3)	(1, 4)	(1, 5)	(1, 6)
(2, 1)	(2, 2)	(2, 3)	(2, 4)	(2, 5)	(2, 6)
(3, 1)	(3, 2)	(3, 3)	(3, 4)	(3, 5)	(3, 6)
(4, 1)	(4, 2)	(4, 3)	(4, 4)	(4, 5)	(4, 6)
(5, 1)	(5, 2)	(5, 3)	(5, 4)	(5, 5)	(5, 6)
(6, 1)	(6, 2)	(6, 3)	(6, 4)	(6, 5)	(6, 6)

2.3 Example 3

Problem: Consider a hash table with 100 buckets. Two arbitrary strings are independently hashed and added to the table. How many possible ways are there for the strings to be stored in the table?

¹"die" is the singular form of the word "dice" (which is the plural form).

Solution: Each string can be hashed to one of 100 buckets. Since the results of hashing the first string do not impact the hash of the second, there are $100 * 100 = 10,000$ ways that the two strings may be stored in the hash table.

3 The Inclusion-Exclusion Principle

Inclusion-Exclusion Principle: If the outcome of an experiment can either be drawn from set A or set B , and sets A and B may potentially overlap (i.e., it is not guaranteed that $A \cap B = \emptyset$), then the number of outcomes of the experiment is $|A \cup B| = |A| + |B| - |A \cap B|$.

Note that the Inclusion-Exclusion Principle generalizes the Sum Rule of Counting for arbitrary sets A and B . In the case where $A \cap B = \emptyset$, the Inclusion-Exclusion Principle gives the same result as the Sum Rule of Counting since $|\emptyset| = 0$.

3.1 Example 4

Problem: An 8-bit string (one byte) is sent over a network. The valid set of strings recognized by the receiver must either start with 01 or end with 10. How many such strings are there?

Solution: The potential bit strings that match the receiver's criteria can either be the 64 strings that start with 01 (since that last 6 bits are left unspecified, allowing for $2^6 = 64$ possibilities) or the 64 strings that end with 10 (since the first 6 bits are unspecified). Of course, these two sets overlap, since strings that start with 01 and end with 10 are in both sets. There are $2^4 = 16$ such strings (since the middle 4 bits can be arbitrary). Casting this description into corresponding set notation, we have: $|A| = 64$, $|B| = 64$, and $|A \cap B| = 16$, so by the Inclusion-Exclusion Principle, there are $64 + 64 - 16 = 112$ strings that match the specified receiver's criteria.

4 General Principle of Counting

General Principle of Counting: If an experiment has r parts such that part i has n_i outcomes for all $i = 1, \dots, r$, then the total number of outcomes for the experiment is $\prod_{i=1}^r n_i = n_1 \times n_2 \times \dots \times n_r$.

In the same way that the Inclusion-Exclusion Principle generalizes the Sum Rule of Counting, our next rule, the General Principle of Counting (also commonly known as the Fundamental Principle of Counting) generalizes the Product Rule of Counting.

4.1 Example 5

Problem: California license plates prior to 1982 had only 6-place license plates, where the first three places were uppercase letters A-Z, and the last three places were numeric 0-9. How many such 6-place license plates were possible pre-1982?

Solution: We can treat each of the positions of the license plate as a separate part of an overall six-part experiment. That is, the first three parts of the experiment each have 26 outcomes, corresponding to the letters A–Z, and the last three parts of the experiment each have 10 outcomes, corresponding to the digits 0–9. By the General Principle of Counting, we have $26 \times 26 \times 26 \times 10 \times 10 \times 10 = 17,576,000$ possible license plates.

Interestingly enough the current population of California is 39.5 million residents as of 2017, so this would not nearly be enough license plates such that each person can own one vehicle. Fortunately, in 1982, California changed to 7-place license plates by prepending a numeric digit, resulting in $10 \times 26 \times 26 \times 26 \times 10 \times 10 \times 10 = 175,760,000$ possible 7-place license plates. This is enough for each resident in California to own approximately 4.5 vehicles.

5 Floors and Ceilings: They’re Not Just For Buildings Anymore...

Floor and *ceiling* are two handy functions that we give below just for reference. Besides, their names sound so much neater than “rounding down” and “rounding up”, and they are well-defined on negative numbers too. Bonus.

Floor function: The floor function assigns to the real number x the largest integer that is less than or equal to x . The floor function applied to x is denoted $\lfloor x \rfloor$.

Ceiling function: The ceiling function assigns to the real number x the smallest integer that is greater than or equal to x . The ceiling function applied to x is denoted $\lceil x \rceil$.

5.1 Example 6

$$\begin{aligned} \lfloor 1/2 \rfloor &= 0 & \lfloor -1/2 \rfloor &= -1 & \lfloor 2.9 \rfloor &= 2 & \lfloor 8.0 \rfloor &= 8 \\ \lceil 1/2 \rceil &= 1 & \lceil -1/2 \rceil &= 0 & \lceil 2.9 \rceil &= 3 & \lceil 8.0 \rceil &= 8 \end{aligned}$$

6 The Pigeonhole Principle

Basic Pigeonhole Principle: For positive integers m and n , if m objects are placed in n buckets, where $m > n$, then at least one bucket must contain at least two objects.

In a more general form, this principle can be stated as:

General Pigeonhole Principle: For positive integers m and n , if m objects are placed in n buckets, then at least one bucket must contain at least $\lceil m/n \rceil$ objects.

Note that the generalized form does not require the constraint that $m > n$, since in the case where $m \leq n$, we have $\lceil m/n \rceil = 1$, and it trivially holds that at least one bucket will contain at least one object.

6.1 Example 7

Problem: Consider a hash table with 100 buckets. 950 strings are hashed and added to the table.

- a) Is it possible that a bucket in the table contains no entries?
- b) Is it guaranteed that at least one bucket in the table contains at least two entries?
- c) Is it guaranteed that at least one bucket in the table contains at least 10 entries?
- d) Is it guaranteed that at least one bucket in the table contains at least 11 entries?

Solution:

- a) Yes. As one example, it is possible (albeit very improbable) that all 950 strings get hashed to the same bucket (say bucket 0). In this case bucket 1 would have no entries.
- b) Yes. Since, 950 objects are placed in 100 buckets and $950 > 100$, by the Basic Pigeonhole Principle, it follows that at least one bucket must contain at least two entries.
- c) Yes. Since, 950 objects are placed in 100 buckets and $\lceil 950/100 \rceil = \lceil 9.5 \rceil = 10$, by the General Pigeonhole Principle, it follows that at least one bucket must contain at least 10 entries.
- d) No. As one example, consider the case where the first 50 bucket each contain 10 entries and the second 50 buckets each contain 9 entries. This accounts for all 950 entries ($50 * 10 + 50 * 9 = 950$), but there is no bucket that contains 11 entries in the hash table.

7 Bibliography

For additional information on counting, you can consult a good discrete mathematics or probability textbook. Some of the discussion above is based on the treatment in:

K. Rosen, *Discrete Mathematics and its Applications, 6th Ed.*, New York: McGraw-Hill, 2007.