1

Computing

Exercise 1.1. Draw a binary tree with the minimum possible depth to:

a. Distinguish among the numbers $0, 1, 2, \ldots, 15$.

Solution. There are sixteen (2^4) different numbers, so the minimum depth binary tree has four levels. There are many different trees that could work, so long as each decision divides the remaining possible numbers in two equal-sized sets. The easiest way to construct it is to start with the 3-bit tree in Figure 1.1 which distinguishes 0–7. We make two copies of this, but add 8 to each number in the second copy, and add an extra binary question at the top to decide which subtree to use.

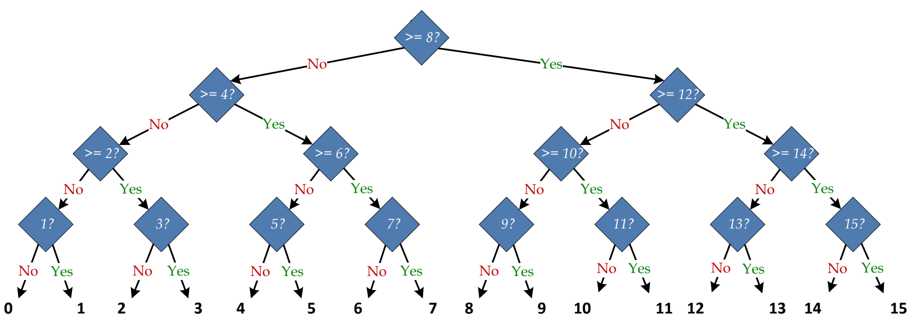

b. Distinguish among the 12 months of the year.

Solution. Since there are 12 months, we need $\log_2 12 \approx 3.58$ bits to distinguish them. This means we need a tree with 4 levels, but on some paths only three questions are needed. One solution would be to use 0–11 from the previous solution. A more natural solution might divide the year into quarters.

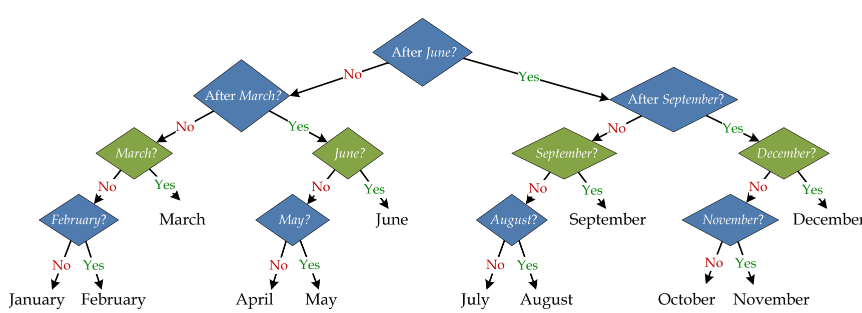

Note that the answers for the green boxes do not provide a full bit of information since the "No" answer leads to two leaf nodes (e.g., January and February), but the "Yes" answer only leads to one leaf node (e.g., March). If each month is equally likely, the answer should be "No" two thirds of the time.

Exercise 1.2. How many bits are needed:

a. To uniquely identify any currently living human?

Solution. According to the U. S. Census Bureau, the world population (on July 4, 2011) is 6.95 Billion people. See http://www.census.gov/main/www/popclock.html for an updated count. To identify each person, we need

$$\lceil \log_2 6,950,000,000 \rceil = 33 \text{ bits}$$

(the notation $\lceil x \rceil$ means the "ceiling" of x which is the least integer larger than x)k. Since $2^3 3 = 8,589,934,592$, 33 bits should be enough to uniquely identify every living person for at least the next decade.

b. To uniquely identify any human who ever lived?

Solution. This is much tougher, and requires defining a *human*. The best estimate I can find come from Carl Haub's article for the Population Reference Bureau, originally written in 1995 and updated in 2002 (http://www.prb.org/Articles/2002/HowManyPeopleHaveEverLivedonEarth.aspx). He estimated that modern humans emerged about 50,000 years ago, and that the total number of humans born up to 2002 was 106 Billion. The world birth rate is approximately 130 million per year, so a current estimate would be perhaps 1-2 Billion more. So, 36 bits is certainly not enough ($2^3 6 = 68,719,476,736$), but 37 bits should be enough for a long time ($2^3 7 = 137,438,953,472$).

c. To identify any location on Earth within one square centimeter?

Solution. The total surface area of the Earth is $510,072,000\text{km}^2$. One kilometer is $100 * 1000 = 100,000$ centimeters, so one square kilometer is $(100 * 1000)^2 = 10,000,000,000$ square centimeters. So, the total surface area of the Earth is $5,100,720,000,000,000,000$ square centimeters. The number of bits needed to represent this is $\log_2 5,100,720,000,000,000,000 \approx 62.145$ so 63 bits is enough to uniquely identify every square centimeter of the Earth's surface. (For comparison, Internet addresses are 128 bits. This is enough for each square centimeter of the Earth to have over 66 quadrillion IP addresses!)

d. To uniquely identify any atom in the observable universe?

Solution. The number of atoms in the universe is estimated to be 10^{80} (see http://en.wikipedia.org/wiki/Observable_universe for a good explanation of how this number was derived). To determine the number of bits needed to uniquely identify every atom, we need to convert from decimal to binary. It requires $\log_2 10 \approx 3.32$ bits to represent each decimal digit, so $10^{80} \approx 2^{265.7}$. Hence, 266 bits is enough to uniquely identify every atom in the universe.

Exercise 1.3. The examples all use binary questions for which there are two possible answers. Suppose instead of basing our decisions on bits, we based it on *trits* where one trit can distinguish between three equally likely values. For each trit, we can ask a ternary question (a question with three possible answers).

a. How many trits are needed to distinguish among eight possible values? (A convincing answer would show a ternary tree with the questions and answers for each node, and argue why it is not possible to distinguish all the values with a tree of lesser depth.)

Solution. Two trits are required since $log_3 8 \approx 1.893$. Here is one such tree:

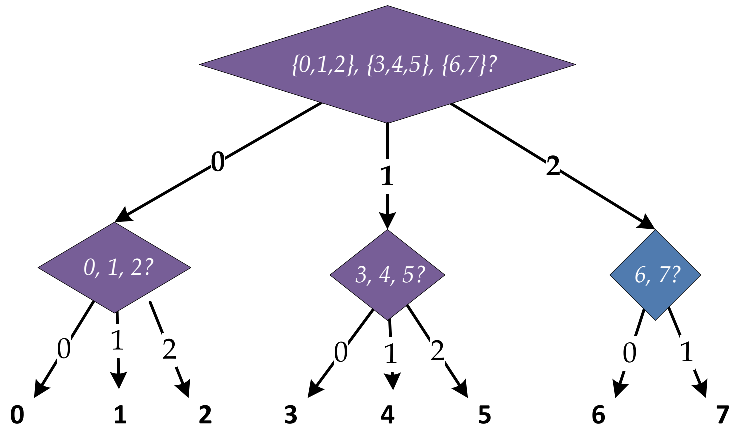

All of the decision nodes except for the bottom rightmost one use three possible outputs. Hence, we could distinguish one additional value using 2 trits (which makes sense since $3^2 = 9$).

b. [⋆] Devise a general formula for converting between bits and trits. How many trits does it require to describe b bits of information?

Solution. To convert between logarithm bases, we use the formula

$$\log_b x = \frac{\log_a x}{\log_a b}.$$

So, for a n-bit value x,

$$\log_3 x = \frac{n}{log_2 3} \approx 0.63093n$$

Thus, to represent an n-bit value requires up to $\lceil 0.63093n \rceil$ trits.