COMP 150PP: Programming With Mature Systems

Due Tuesday, November 8, 2016 at noon

October 26, 2016

Goals

Now that we've had a chance to do our own language designs and write our own inference algorithms, it's time to compare that experience with the experience of using a mature system that infers posterior distributions using statistical methods. In principle, these systems ought to be able to solve dice problems quickly, but not necessarily exactly.

Your assignment is to *pick a mature probabilistic programming system* and use it to *solve one problem from each group of problems below*. (The problems are all problems you've seen before.)

This assignment has two purposes:

- It will give you information about an alternative system that you can compare with the system you designed and built yourself. We will compare experiences and code in class on Wednesday, November 9.
- It will give you sufficient experience with the system so you can decide if you want to use it for your project.

The assignment is offered on the usual terms: you may work by yourself or with any others you wish; and the deadline is soft—although I do expect you do have done enough to make an in-class discussion possible. *Send me your solutions* by noon on Tuesday, November 8 so we can have Code Show and Tell on Wednesday, November 9.

Languages and systems to consider

The systems named below are hyperlinked. If your PDF reader does not understand hyperlinks, you can also find links on the course home page, on the syllabus, and on the HTML version of this handout at http://www.cs.tufts.edu/comp/150PP/handouts/1026outside-problems2c.html.

Most likely to be easy to use

The three languages WebPPL, Church, Anglican have the best tutorials.

- WebPPL is the implementation language of the electronic book The Design and Implementation of Probabilistic Programming Languages.
- Church is a Scheme-like language with probabilistic semantics. I have my issues with the computational model, but it is the foundation of many other projects. An extensive tutorial can be found in the electronic book Probabilistic Models of Cognition.
- Anglican is inspired by Church and integrated with Clojure. It compiles to relatively efficient code for the Java Virtual Machine. There is a highly polished tutorial.

Possibly more powerful

The three languages Hakaru, Probabilistic C, and Wolfe each have something more interesting to offer than the basic model of distributions over "evaluation histories," but they are probably more difficult to learn. To start, depending on which system you pick, you had better be comfortable using Haskell, C, or Scala.

- Hakaru is a probabilistic language embedded in Haskell. Its distinctive feature is that it can do exact inference by symbolic analysis. We will read about this inference in class.
- Probabilistic C is intended as a compilation target for higher-level probabilistic languages. It is surprising how much can be accomplished by adding just a couple of primitives to C—and by exploiting parallelism.
- Wolfe is a probabilistic language embedded in Scala. It has a particularly clean way of talking about probabilistic domains.

Other possibilities

A longer list, with other commentary, can be found at `probabilistic-programming.org`.

The evaluation task

Using the probabilistic-programming system of your choice, in teams of your choice, write code that answers some of the questions below. Answer at least one question from each group. On November 9, come to class prepared to show code and to analyze and explain your experience.

The questions

These are mostly the same questions you've seen before.

Simple distribution and probability questions

A. What distribution of integers results from throwing a single d6? What about a d12?

B. Given that you have a d6 and a d12, what's the probability of throwing a total of 11?

C. You're given a `P D` (distribution over dice) representing the ability to draw one die at random from a bag. You draw two dice. (Assume you are drawing dice "with replacement", so that you can draw two by doing the exact same probabilistic draw twice.)

- What's the probability of drawing a d6 and a d12?
- What's the probability that exactly one of the two dice is a d20?

D. You draw two dice, throw them, and total the numbers. What's the ("joint") probability of drawing a d6 and a d12 and throwing 11?

E. You draw two dice and throw them. The numbers total 11. Given that you know the total of 11, what's the ("conditional") probability that you drew a d6 and a d12?

Dice and coins

F. You throw a d6 to get number N, then toss that number of coins. What's the probability of observing exactly three heads?

G. You throw a d6 to get number N, then toss that number of coins. What's the probability of observing at least three heads?

H. You throw a d6 to get number N, then toss that number of coins. The coins land in a row. What's the probability that the row contains three *consecutive* heads?

I. I throw a d6 to get number N, then toss that number of coins. I don't tell you what N is, but I do tell you that exactly three of the coins came up heads. What's the probability that $N = 3$?

J. I throw a d6 to get number N, then toss that number of coins. I tell you that $N > 4$. What's the probability that exactly three of the coins come up heads?

K. I throw a d6 to get number N, then toss that number of coins. I don't tell you what N is, but I do tell you that exactly three of the coins came up heads. What's the probability that $N = 3$?

L. I throw a d6 to get number N, then toss that number of coins. I don't tell you what N is, but I do tell you that exactly *one* of the coins came up heads. What's the most likely value of~N?

M. I throw a d6 to get number N, then toss that number of coins. I don't tell you what N is, but I do tell you that exactly three of the coins came up heads. What's the posterior distribution of~N?

Tally-sheet questions

In all these questions, you throw a tally sheet (draw two dice, throw 30 times, mark left column for totals less than 16, the middle column for totals of exactly 16, and the right column for totals of more than 16).

N. What's the probability that you have drawn two d12's and put 3 marks in the right column of the tally sheet?

O. What's the probability that all marks are in the left column?

P. On average, how many marks can you expect to put in the right column of the tally sheet?

Gambling questions

Probability was invented by gamblers for gamblers. Let me describe the *tally-sheet game*:

> The tally-sheet game has two players: the *thrower* and the *guesser*.
>
> - The thrower creates a tally sheet by drawing two dice and throwing them 30 times using the procedure you know.
> - The guesser sees the tally sheet but not the dice and has three chances to try to guess the dice.
>
> It's not gambling unless you play for money. Here are the payouts:
>
> - If the guesser gets the right dice on the first try, the guesser wins one dollar ($1.00) from the thrower.
> - If the guesser gets the right dice on the second try, the guesser wins fifty cents ($0.50) from the thrower.
> - If the guesser gets the right dice on the third try, the guesser wins a quarter ($0.25) from the thrower.
> - If the guesser fails to get the right dice in three tries, the guesser loses and pays a dime ($0.10) to the thrower.

Write code to answer the following questions about the tally-sheet game:[1]

Q. Who is likely to win money: the guesser or the thrower?

R. You have your choice of playing guesser or thrower for a series of 1000 games. You choose whichever is more likely to win. How much do you expect to win?

S. You're playing thrower, and I ask if you'd be willing to change the rules so that instead of throwing 30 times, you throw 50 times. Assuming we're going to play 1000 games, what's the minimum amount you should insist on charging me for this privilege?

T. As guesser, I hate it that I can never tell the difference between d6+d8 and d6+d6. I want the option to change the rules so that the middle column corresponds to 13, not 16. (The left and right columns change correspondingly.) Assuming I'm going to play 1000 games as guesser, what's the most I should be willing to pay for the privilege of exercising this option?

U. Suppose that in addition to the original dice, I add 12 d4's (four-sided dice) that I kept out of the original experiment. Now putting the middle at 16 is even less helpful. Where should the split be to maximize the expected return for the guesser? (Assume the guesser always picks the most likely guess in the posterior distribution, which is the *maximum a posteriori* choice.)

The slips problem

V. Somebody gives Norman a hat containing five slips of paper, numbered 1 to 5 respectively. Norman draws a slip from the hat. The number on the slip is called n. Norman then repeats the following procedure ten times:

- Take n dice from the bowl, throw them, report the total t, then put the dice back in the bowl.

The totals reported are 21, 15, 34, 12, 18, 38, 46, 13, 24, and 27. The question is, *what is the number on the slip Norman drew?* (That is, what is n?) We will call this problem the *slip problem*.

Here are the initial contents of the bowl:

Number of sides	Number of those dice in the bowl
d4	—
d6	9
d8	9
d10	—
d12	14
d20	14

[1] In all these questions, please assume that the cost of throwing dice and the cost of computing answers are both zero. In other words, the questions should consider only the payouts from the game; all actions in the real world are free.