Lab: Probability, Testimony and Belief

CSI 3305: Introduction to Computational Thinking

December 13, 2010

1 Introduction

One major topic in philosophy is the evaluation of our beliefs to determine when and under what conditions a belief has a certain merit, like rationality. The rationality of our beliefs has a lot to do with evidence. Rational beliefs are proportioned to our evidence. Determining the weight of piece of evidence is a very complicated procedure, typically modeled by Bayesian reasoning. The Bayesian model can yield very surprising results, and there is strong evidence that we are not naturally very good at estimating accurately the weight of our evidence. Thinking about the rationality of our beliefs, at this point, is heavily computational, and often involves very large numbers. Usually, with a little instruction, students can work with the Bayesian model on pen and paper. But it is difficult, and, a bit of a distraction from the main philosophical point. In fact, students may well understand the nature of Bayesian reasoning better by turning to computational tools. The important point is not crunching the numbers, so to speak, but understanding how the numbers are interdependent.

2 Problem Statement

One good way to illustrate the issues surrounding the weight of evidence is to calculate the tipping point for believing something based on testimony of eyewitnesses (or of experts). No matter how unlikely an event may be (provided that the event is not impossible) there is a number of credible independent witnesses that would convince us that the event had occurred, and a point at which it would be positively irrational not to believe that the event had occurred. The computational problem is to fix that number, and to see how it is a function of the improbability of the event in question, of the reliability of the witnesses (or experts) and of their independence. Using Bayesian reasoning, and assigning plausible estimates (which they should be able to explain), students should be able to compute the number of eyewitness (or expert) testimony needed to convince them to believe something, no matter how improbable. Good examples can be drawn from any number of domains. One might ask, questions like these: How many qualified testifiers would it take to convince me that two students composed precisely the same 500 word essay? How many trusted reports would it take to convince me that a man survived a sky-diving accident in which his shoot failed to open? Or, since this is Baylor, how many reports would it take to convince me that a man had been raised from the dead?

For this lab, you will create a small program that calculates the probability of an event, using Bayes' theorem, given several parameters that you will adjust. You will answer the question: based on reasonable estimates of prior probabilities, how many eyewitnesses would be necessary to convince a rational person that Jesus had risen from the dead?

3 Tools

You will be using the Python programming language and interpreter for this lab. The script will be written using a text editor, and will be run using the command line.

4 Setup and Programming

1. Within a local folder, create a new file called bayes.py
2. Open bayes.py in the text editor of your choice (such as Notepad or Vim.)
3. Enter the following code in the file:

   ```
   from __future__ import division
   ```

   This code will allow Python to handle division in a more intuitive way, doing floating point division by default.

4. Next, enter the following code:

   ```
   """ Variables """
   witnesses = 10
   P_H = 1 / 100000000000
   P_E_H = 0.5
   P_NOT_H = 1 - P_H
   prob_of_single_witness = 0.1
   P_E_NOT_H = prob_of_single_witness**witnesses
   ```

   These are the parameters controlling our Bayesian probability estimate. The first is the number of eyewitnesses to an event, which we've estimated as ten. The second is the prior probability of our hypothesis holding. We've estimated it at one in a hundred billion, which is the same as saying that we assume one out of every hundred billion people rises from the dead. Notice that this estimate is pretty subjective, since we don't know beforehand what the probability of a resurrection occurring three days after death really is. Since we're uncertain, we choose a small probability, though we could have chosen a smaller (or larger) one.

   The next variable, `P_E_H`, is the probability of seeing the evidence given that our hypothesis holds. We've estimated this as 0.5. This is even harder to estimate, since we have no good way of estimating what the probability of the evidence is (multiple, independent eyewitness testimony) given that the hypothesis is true. So we choose a 50-50 probability.

   Next is the probability of the event NOT occurring, which should be one minus the probability that it does occur. The next variable, `prob_of_single_witness`, is the probability of a single witness seeing the events in question given that the hypothesis does not hold. For example, the probability of a person seeing a risen Jesus (a post-death sighting) assuming that he did not actually rise from the dead. We assume that a person has a ten percent chance of having such a mistaken encounter. The last variable, `P_E_NOT_H`, is the probability that of the evidence given that the hypothesis does not hold. It assumes that the eyewitnesses are independent and calculates the probability of several people having an eyewitness encounter given that the hypothesis does not actually hold. In other words, it is the probability of ten people seeing Jesus after his death, assuming that he was actually still in a tomb.

   These are the parameters you will need to adjust in order to answers the questions in the Questions section.

5. Next, type in the following:

```
""" Bayes' Theorem """
P_H_E = P_H * P_E_H / (P_H * P_E_H + P_NOT_H * P_E_NOT_H)
```

   This line calculates `P_H_E` which is the Bayesian probability that hypothesis H holds given the evidence E (eyewitness testimony).

6. Lastly, enter the following:

```
print "Probability that hypothesis H holds:", P_H_E
```

   This prints out our calculated probability.

7. Save the file and close it.

8. Open the command line (From the Windows menu bar: Start > Run... > cmd). Navigate to the folder where you created your file (use the following command: cd "C:\Folder\where file\is".)

9. Call the bayes.py file from the command prompt, using the following command and hitting enter:

```
python bayes.py
```

   You should see a probability displayed, between 0.0 and 1.0.

5 Questions

1. Given the parameters as defined in the Setup and Programming section, what is the probability of the hypothesis H holding? (Hint: it is the output of the program.)

2. The gospels and Paul's letters indicate that more than ten people had encounters with Jesus after his death. What happens to the probability of the hypothesis H if you change the number of eyewitnesses to fifteen instead of ten?

3. What is the "tipping point", i.e. the number of independent witnesses needed to make hypothesis H more probable than not given our model? In other words, what is the minimum number of witnesses necessary to make the probability exceed 50%?

4. The probability `P_H` has been estimated in a subjective way. Is it a reasonable estimate? If so, explain why. If not, give your own estimate for this value and justify why your estimate is more reasonable.

5. Is ten percent a reasonable estimate for `prob_of_single_witness`? In other words, is it reasonable to assume that for ever person that dies, ten percent of all people will have a sighting of them after death? Explain your reason for agreeing or disagreeing. Then choose a more reasonable estimate for this probability and explain why your number is more reasonable. Lastly, what is the new probability of hypothesis H given your new value for `prob_of_single_witness`?

6. Is the estimate for `P_E_H` a reasonable one? Why or why not? Remember, this is the probability that we'd have the eyewitnesses evidence of a risen Jesus if he rose from the dead. This probability estimate is somewhat subjective, so come up with your own value and justify why your estimate is reasonable. Then record what changing this value does to your overall probability.

7. Based on your experimentation with different values and estimates, roughly how many eyewitnesses would be necessary to make belief in Jesus' resurrection plausible?