Computational Cognitive Science (2018–2019)

School of Informatics, University of Edinburgh
Original exercises by Frank Keller, with modifications by Chris Lucas

Tutorial 3: Bayesian Estimation; Concept Learning

Please work through this tutorial sheet on your own time as much as possible before arriving in tutorial. We encourage you to work together and discuss your methods and solutions.

1 Exercises in Bayesian Estimation

Question 1: I have three fair dice in a jar: one with 6 sides, one with 8 sides, and one with 12 sides. I select one at random from the jar, roll it once, and tell you that the number is greater than four.

1. If you had to place a bet on which die I picked, which would you bet on? Why?
2. If I roll the **same** die again, what's the probability that the outcome will be greater than four?

Question 2: In an experiment on face recognition, subjects are presented with images of people they know, and asked to identify them. The images are presented for a very short period of time so that subjects may not have time to see the details of the entire face, but are likely to get a general impression of things like hair color and style, overall shape, skin color, etc. In this question we will consider how to formulate the face recognition problem as a probabilistic inference model.

1. What is the hypothesis space in this problem? Is it continuous or discrete? Finite or infinite?
2. What constitutes the observed data y and what kinds of values can it take on?
3. Write down an equation that expresses the inference problem that the subjects must solve to identify each face. Describe what each term in the equation represents.
4. What factors might influence the prior in this situation?
5. Suppose one group of subjects sees clear images, such as the one on the left below, and another group sees noisy images, such as the one on the right below. Which term(s) in your equation will be different for the noisy group compared to the clear group?

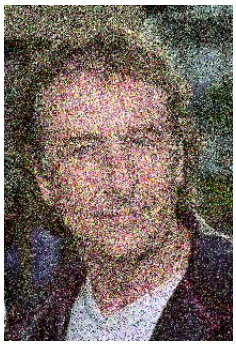

6. What does the model predict about subjects' performance with noisy images compared to clear images? Rather than working with the full scenario above, you can simplify by supposing the experiment has images of only two people, $h_1 =$ Eric Idle and $h_2 =$ John Cleese.

How does image noisiness affect the model's probability of inferring that the image is of Eric Idle, rather than John Cleese? (Hint: since you are considering only two hypotheses, think about the *posterior odds*. How will this quantity differ between the case where the image is noisy and the case where it is clear?) Do you need to make any further assumptions in order to make predictions about subject behavior?

2 A Bayesian Model of Concept Learning

Question 3: In this question we will consider a simplified version of the Tenenbaum, 2000 model of generalization where there are only four hypotheses under consideration, each of which has equal prior probability:

$h_1 = \{\text{odd numbers}\}$, $h_2 = \{\text{even numbers}\}$, $h_3 = \{\text{multiples of 5}\}$, $h_4 = \{\text{multiples of 10}\}$

Given the set of examples $X = \{10, 40\}$ from the target concept, first compute the posterior probability of each hypothesis under the model. Then, determine the model's predicted probability that each of the following new data points is also part of the same concept: 2, 3, 5, and 20. Now explain, with reference to the terms in the model, why the model's predictions tend to become more 'rule-like' as more examples are seen (e.g., if 50 is shown as a third example from the concept).

References

Tenenbaum, J. B. (2000). Rules and similarity in concept learning. In *Advances in neural information processing systems* (pp. 59–65).