Computational Cognitive Science 2026/27

Week 7

Frank Keller Chris Lucas

Tutorial 5: Active learning

This tutorial provides some hands-on exposure to concepts from the active learning lecture.

Imagine you come across a curiosity shop that has four devices for sale at £100 each. Each device can automatically create a delicious cake when one pulls its lever, but the devices vary in their reliability:

- one device never works,
- one device works exactly 20 percent of the time,
- one works 30 percent of the time, and
- one works 75 percent of the time.

We don't know which device is which, but the shopkeeper will let us use the devices a few times to choose one.

Our task is to find out which machine is which.

Let's take an information-theoretic approach to this active learning problem.

Here are some helper functions.

```
plogp <- function(p) {
  # bits, not nats
  if(p==0) 0 else p*log2(p) # Consider why we can't just use p*log2(p).
}

# entropy for a categorical random variable
cat_entropy <- function(catv) {
  -sum(sapply(catv,plogp))
}
```

Let's define our hypothesis space as all of the different cake probabilities the machines before us might have from lef to right. Fore example `c(0,.2,.3,.75)` means the machines increase in their probabilities from left to right, and `c(.75,.3,.2,0)` means they decrease from left to right.

Question 1 How does the code below define that hypothesis space?

```
require(combinat) # From googling "all list permutations in R"
```

```
## Loading required package: combinat
```

```
##
## Attaching package: 'combinat'
```

```
## The following object is masked from 'package:utils':
##
##     combn
```

```
probs <- c(0,.2,.3,.75)
hyps <- permn(probs)
```

Question 2 What is our entropy, assuming we know nothing about the order of the machines?

```
prior <- c(0) # Replace this with a prior over all hypotheses
cat_entropy(prior)
```

```
## [1] 0
```

Question 3 Suppose we turn the lever on the first machine and get nothing. What probability should we now assign our different hypotheses? Edit the code below.

```
update_probs <- function(prior,choice,outcome) {
  # prior (vector): one prior probability per hypothesis
  # choice (int): index of the machine (1-4) on which the lever was turned
  # outcome (int): whether we get cake (1) or not (0)

  likelihood <- function(this_hyp,this_choice,this_outcome) {
    1.0 # Replace this 1.0 with an actual likelihood
  }

  likes <- sapply(hyps,function(h) {likelihood(h,choice,outcome)})  # likelihoods
  unnp <- likes*prior # unnormalized posterior
  z <- sum(unnp) # normalization constant
  unnp/z
}

post <- update_probs(prior,choice=1,outcome=0)
print(post)
```

```
##  [1] NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN NaN
## [20] NaN NaN NaN NaN NaN
```

Question 4 What is our new entropy? Did we learn much? Why or why not?

Question 5 What is the expected probability getting a cake if we choose machine 2 next?

```
# First we want the probability of each outcome in each hyp
# Then we take the weighted sum over hypothesis probabilities (Bayesian hypothesis averaging)
expected_outcome <- function(i,ph) {
  # ph (vector): posterior, one prior probability per hypothesis
  # i (int): machine index
  0 # Replace this
}
print(expected_outcome(2,post))
```

```
## [1] 0
```

Question 6 What is the expected entropy if we choose machine 2 next?

```
expected_entropy <- function(choice,p_hyps) {
  pwin <- expected_outcome(choice,p_hyps)
  0 # Replace this
```

```
}
print(expected_entropy(2,post))
```

```
## [1] 0
```

Question 7 Suppose we do choose machine 2 and it yields a cake. What is our current entropy and what are the expected entropies after each of our next possible choices?

Question 8 What does this imply we should do next?

Question 9 Is this at all counter-intuitive? What might we do differently if we are only interested in knowing which is the most reliable machine?