

Bayesian Probability: Formally updating your beliefs given observations

On your first day of class, what was the probability of you doing well in this course?

Assign a probability to each grade according to your prior belief (*remember $\Sigma(p_i) = 1$, and try not to rule out any possibility entirely – ie. don't assign 0 probability to any grades – absolute certainty rules out learning!*)

P(F)____ P(D)____ P(C) ____ P(B) ____ P(A) ____ : $P(\text{Course Grade}) \leftarrow$ Prior probability

You've received your grade back from your first assignment. Now you've made an observation! You can now update your beliefs about your course grade using Bayes formula:

$$P(\text{Course Grade} | \text{Marks So Far}) = \frac{P(\text{Marks So Far} | \text{Course Grade}) P(\text{Course Grade})}{P(\text{Marks So Far})}$$

Your observation enters through what's called a *likelihood function*:

$$P(\text{Marks So Far} | \text{Final Grade}) \leftarrow \text{Likelihood}$$

The model we have of grades is geometric:
(likelihood drops off with the square of the distance between grades)

$$P(\text{Marks So Far} | \text{Course Grade}) \propto P_i \left(\frac{1}{(c_i - m)^2 + 1} \right)$$

$$i \in \{A, B, C, D, F\}$$

We have everything we need to compute the *posterior probability* of our course mark given our observation (up to a normalizing constant):

$$P(\text{Course Grade} | \text{Marks So Far}) \leftarrow \text{Posterior}$$

formula $P(\text{Course Grade} | \text{Marks So Far}) \propto P(\text{Marks So Far} | \text{Course Grade}) P(\text{Course Grade})$

Prior Probability	Likelihood	Posterior (normalize by dividing by the sum)
P(A)=	$1/(\{\text{difference between observed mark and this grade}\}^2 + 1)$ =	Likelihood*Prior/Sum =
P(B)=		
P(C)=		
P(D)=		
P(F)=		
	Sum =	