


CS188 Spring 2014 Section 10: ML

Pacman and Mrs. Pacman have been searching for each other in the Maze. Mrs. Pacman has been pregnant with a baby, and just this morning she has given birth to Pacbaby (Congratulations, Pacmans!).

Because Pacbaby was born before Pacman and Mrs. Pacman reunited in the maze, he has never met his father. Naturally, Mrs. Pacman wants to teach Pacbaby to recognize his father, using a set of Polaroids of Pacman. She also has several pictures of ghosts to use as negative examples.

Because the polaroids are black and white, and were taken from strange angles, Mrs. Pacman has decided to teach Pacbaby to identify Pacman based on more salient features: the presence of a bowtie (b), hat (h), or mustache (m).



The following table summarizes the content of the Polaroids. Each binary feature is represented as 1 (meaning the feature is present) or 0 (meaning it is absent). The subject y of the photo is encoded as +1 for Pacman or -1 for ghost.

(m)	(b)	(h)	Subject (y)
0	0	0	+1
1	0	0	+1
1	1	0	+1
0	1	1	+1
1	0	1	-1
1	1	1	-1

1 Naive Pacbaby

Suppose Pacbaby has a Naive Bayes based brain.

- Write the Naive Bayes classification rule for this problem (i.e. write a formula which given a data point $x = (m, b, h)$ returns the most likely subject y). Write the formula in terms of conditional and prior probabilities. Be explicit about which parameters are involved, but you do not need to estimate them yet.
- Assuming no smoothing, give estimates for the parameters of the classification rule based on the Polaroids.

	$y = +1$	$y = -1$
$P(y)$		
P	$y = +1$	$y = -1$
$P(m = 1 y)$		
$P(b = 1 y)$		
$P(h = 1 y)$		

c) Suppose a character comes by wearing a hat but without a mustache or bowtie. What would happen if Pacbaby had to guess the identity of the character?

d) Suppose now that Pacbaby performs Laplace smoothing with strength $k = 1$ (on both the prior and class-conditional parameters). Re-estimate the parameters. Now how will Pacbaby classify this new character with the hat and without a mustache or bowtie?

	$y = +1$	$y = -1$
$P(y)$		


P	$y = +1$	$y = -1$
$P(m = 1 y)$		
$P(b = 1 y)$		
$P(h = 1 y)$		

2 Perceptron Pacbaby

Suppose Pacbaby has a perceptron based brain, meaning his “internal classifier” learns through perceptron updates and is limited to learning linear classification rules. Further, suppose Pacbaby “augments” each example it sees with a bias feature that is always equal to 1 (this allows Pacbaby to learn decision rules with boundaries that do not pass through the origin).

(m)	(b)	(h)	Subject (y)
0	0	0	+1
1	0	0	+1
1	1	0	+1
0	1	1	+1
1	0	1	-1
1	1	1	-1

- a) Will Pacbaby be able to learn a rule that makes no mistakes on the set of Polaroids? In other words, is the training set linearly separable? (Plot the training data)



- b) Suppose there was another Polaroid of a character without a mustache or a hat, but who was wearing a bowtie. If this Polaroid was of Pacman, would the data be linearly separable? What if it contained a ghost?

- c) Suppose we start with the training weights $[-1, 1, -1, -1]$, and wish to train a perceptron on the above data. Perform two updates of the Perceptron algorithm, processing the training data in the order they appear. The first 3 weights correspond to the features $(m), (b), (h)$, respectively. The last weight corresponds to the bias feature. If a training example has weight exactly 0, classify it as +1.

	w_1	w_2	w_3	w_4
Initial weights	-1	1	-1	-1
Training: (0,0,0,1) → +1				
Training: (1,0,0,1) → +1				