

Fall 2008 – COMP 131 Midterm 2 Review¹

1 Bayes Nets

Let's say you go to the horse races. A shady character comes to offer you a free tip (a tip is a piece of information): he says Belle did not eat her breakfast. Assume that:

- The probability that a horse will win is dependent on the horses health and its speed.
- A horses health and its speed are independent.
- A healthy horse has a higher probability of eating breakfast than does a sick horse.
- Your informant is known to be accurate 80% of the time.

1. Draw a Bayesian network with 5 variables (T = you got this tip; B = Belle ate her breakfast; H = Belle is healthy; W = Belle will win; F = Belle is fast). The relationships between the variables should reflect the problem description.

2. What is $P(W)$ (in terms of values stored in the networks conditional probability tables)?

3. What is $P(W|T)$ (in terms of values stored in the network's conditional probability tables)?

¹Based on MIT's OpenCourseWare materials for 6.034 Artificial Intelligence by Profs. Kaelbling and Lozano-Perez, Berkeley's AI course by Prof. Russell, and Rochester University CSC 242 by Prof. Brown.

2 Gene sequencing

A gene is a sequence of nucleotide bases A,T,C and G. Each gene is made of exons – sequences of bases which get transcribed into proteins, and introns – sequences that are not transcribed. The single nucleotide where the switch from one to the other occurs is called a splice site. In 95% of the cases, the splice site is a G. It can also be an A, rarely, but never a C or a T. All bases are equally likely to appear in an exon, but introns are A and T-heavy (A and T are 4 times as likely as C and G, and equally likely as each other). We expect quite a few bases in an exon or an intron sequence. Say, we know that if one nucleotide is part of an exon (or an intron), the next one in the sequence has a 75% chance of being part of the same exon (or intron).

1. Draw a graphical representation of our knowledge in this domain. What are the random variables in your model? Include any relevant probability distributions.
2. What kind of model have you drawn? Be specific.
3. We have isolated a sequence CTGGA^{ACTC}, and we know that the first C comes from an exon, and the last C comes from an intron. Where is the most likely splice site between the two?

3 Markov Decision Processes

Suppose an agent inhabits a world with two states, s and $\neg s$. It can do exactly one of two actions, a and b . Action a does nothing and action b flips from one state to the other. The reward function is $R(s) = 3$, $R(\neg s) = 2$, and $\gamma = 0.5$. Complete the columns of the following table by performing policy iteration to find the optimal policy.

	π^0	U^{π^0}	π^1	U^{π^1}	π^2
s	a				
$\neg s$	a				

4 Reinforcement Learning

1. You are a temporal difference agent in a completely observable environment. Your initial estimate for the utility $U(s)$ of states you haven't visited (initially all) is 0. Your future reward discount $\gamma = 1.0$. Your learning rate $\alpha = 0.1$. Here is the history of one trial taking you from state 1 to an enjoyable state 2, to a mildly painful terminal state 3 (It reads down the columns, left to right).

step	0	1	2
state s	1	2	3
reward R	-1	5	-4
action a	go2	go3	
new state s'	2	3	

What is your estimate of the utilities $U(s)$ for all states s at the end of step 1? At the end of step 2?

2. You are an Adaptive Dynamic Programming agent in the same environment. You make two trials, the first being the same as above in part 1. Here is the second:

step	0	1	2	3	4
state s	1	2	2	2	3
reward R	-1	5	5	5	-4
action a	go2	go2	go3	go3	
new state s'	2	2	2	3	

What is your estimate of the transition function (or table) $T(s, a, s')$ at the end of the first trial? At

the end of the second trial?

3. How will you use the estimated values of $T(s, a, s')$?

5 More Reinforcement Learning




Figure 1: Simple MDP world.

The simple world in the figure has a start state S and a goal state G , with action s (for stay or start) that is meant to keep the agent in the start state, and g (for go or goal) meant to take agent to goal. The world is completely observable to the agent (it knows what state it is in and what reward it is getting). At goal, agent exits the game; there are no more moves until its next trial. Say there is a future-reward discount factor of γ . Reward of the start state is r and the reward of the goal is R , which in this world are -0.1 and $+1.0$ respectively. Both rewards and discounts start when the agent is born, so if the agent wakes up in S and takes action g successfully, its total reward is $0.1 + \gamma \times 1.0$. The world is uncertain. In fact the transition function $T(FromState, Action, ResultState)$ giving the probabilities of transitions given actions is: $T(S, s, S) = .9$, $T(S, s, G) = .1$, $T(S, g, S) = .2$, $T(S, g, G) = .8$. There are no $T(G, ., .)$ since state G is terminal. Call $T(S, g, G) = p$, the probability of successfully going, and then $T(S, g, S) = q = 1p$, the probability of an unsuccessful go action.

1. What is the agent's optimal policy?

6 Bayes Net Inference

Consider the Bayes net shown below.

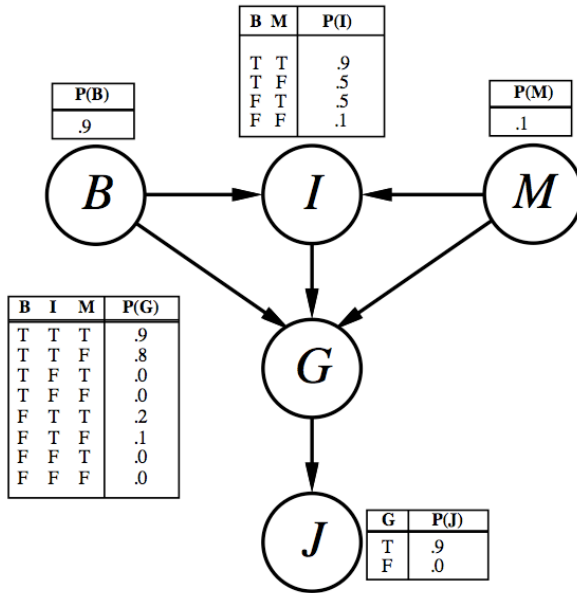


Figure 2: A simple Bayes net with boolean variables $B = \text{BrokeElectionLaw}$, $I = \text{Indicted}$, $M = \text{PoliticallyMotivatedProsecutor}$, $G = \text{FoundGuilty}$, and $J = \text{Jailed}$.

1. Which, if any, of the following are asserted by the network *structure* (ignoring the CPTs for now)?

- (i) $P(B, I, M) = P(B)P(I)P(M)$
- (ii) $P(J|G) = P(J|G, I)$
- (iii) $P(M|G, B, I) = P(M|G, B, I, J)$

2. Calculate the value of $P(b, i, \neg m, g, j)$

3. Calculate the probability that someone goes to jail given that they broke the law, have been indicted and face a politically motivated prosecutor.

7 Perceptron

- The following table shows a data set and the number of times each point is misclassified during a run of the perceptron algorithm, starting with zero weights. What is the equation of the separating line found by the algorithm, as a function of x_1 , x_2 , and x_3 ?

x_1	x_2	x_3	y	times misclassified
2	3	1	+1	12
2	4	0	+1	0
3	1	1	-1	3
1	1	0	-1	6
1	2	1	-1	11

Assume that the learning rate is 1 and the initial weights are all zero.

- Which of the four data sets below are correctly classifiable with a perceptron? The inputs have two continuous features. The two classes are labelled by + and -.

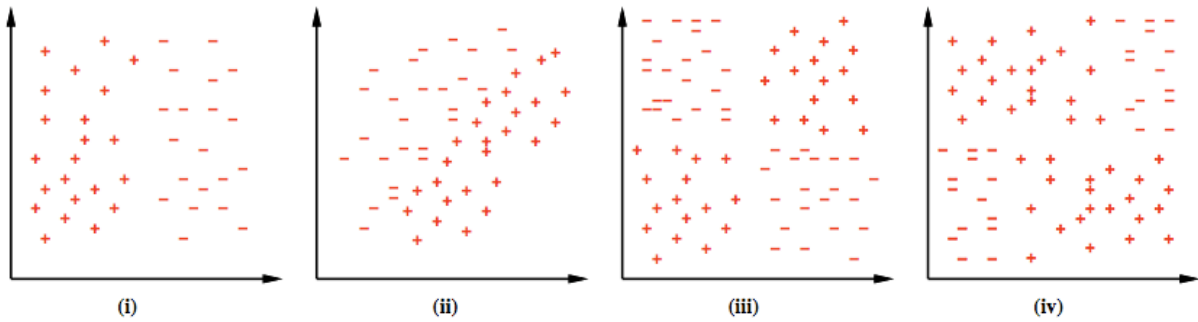


Figure 3: Four datasets for classification.

8 Overfitting and Complexity

1. How does a neural net overfit? How can you reduce overfitting in a neural net?
2. How does a Bayesian classifier overfit? How can you reduce overfitting?
3. Which of the following approaches do perceptrons use to control complexity while minimizing error? Which do feed-forward neural networks (with no weight decay or early stopping) use?
 - A: Use a fixed-complexity hypothesis class
 - B: Include a complexity penalty in the measure of error
 - C: Nothing

9 Naive Bayes

Consider a Naive Bayes problem with three boolean features, x_1, x_2, x_3 . Imagine that we have seen a total of 12 training examples, 6 positive (with $y = 1$) and 6 negative (with $y = 0$). Here is a table with some of the counts:

	$y = 1$	$y = 0$
$x_1 = 1$	5	5
$x_2 = 1$	0	0
$x_3 = 1$	2	4

1. Supply the following estimated probabilities:

$$P(x_1 = 1|y = 0) =$$

$$P(x_2 = 1|y = 1) =$$

$$P(x_3 = 0|y = 0) =$$

2. Is there anything problematic with any of the above estimates? Be specific.

3. Which feature plays the largest role in deciding the class of a new instance? Why?

10 Bayesian learning

Suppose that you want to build a program that detects whether an incoming e-mail message is spam or not. You decide to attack this using machine learning. So, you collect a large number of training messages and label them as spam or not-spam. You further decide that you will use the presence of individual words in the body of the message as features. That is, you collect every distinct word found in the training set and assign to each one an index, from 1 to N . Then, given a message, you construct a feature vector with N entries and write in each entry a 1 if the word appears in the message, and 0 if it does not.

1. What does a naive Bayes model for this problem look like?
2. Assuming a naive Bayes model, how would you compute its parameters? Give precise formulas.
3. What assumption did you make in your calculations?
4. What would be a less naive model?