Name: _____

15-681 Machine Learning — Final Exam

December 12, 1994

Avrim Blum and Tom Mitchell

Instructions:

- This final is open book, open notes.
- Write your answers directly on the exam. If you need extra space, use the back side of the previous page.
- Concise, clear answers will receive more credit than long, rambling answers.

Good luck!

1. Short answers [25 pts] _____
2. PAC model [15 pts] _____
3. Algorithm design [30 pts] _____
4. Explanation-based learning [15 pts] _____
5. Q-learning [15 pts] _____

Total: [100 pts] _____

1. **Short answers** [25 pts] (5 points each). Answer each of the following in a few sentences.

   (a) A set of 100 training examples is provided as input to each of two learning systems, L1 and L2. Afterwards, L1 is able to classify 91 of these training examples correctly, whereas L2 is able to classify only 89 of the training examples correctly. Ronnie Researcher claims, therefore, that the results of L1 should be used and L2 should be ignored. Describe two distinct flaws in Ronnie's methodology and reasoning.

   (b) Consider the following four training examples of some unknown target concept:

	Color	Size	Classification
Example1:	Blue	Large	+
Example2:	Red	Small	-
Example3:	Blue	Small	+
Example4:	Blue	Large	-

   Which attribute (Color or Size) has the highest information gain relative to the target classification? What is its information gain (you may express your answer in terms of logarithms)?

(c) Is overfitting a problem for Genetic Algorithms? If so, how would you minimize it? If not, why not?

(d) What is the VC-dimension for the hypothesis class $H = \{\text{interiors of circles}\}$, where examples are points in the plane? For instance, a typical hypothesis might be "all points within distance 4.5 of $(1, 3)$ are positive." Explain your reasoning.

(e) Consider the following HMM. Numbers in boldface represent the probability of taking an edge, and numbers not in boldface represent the probability of producing the output characters.

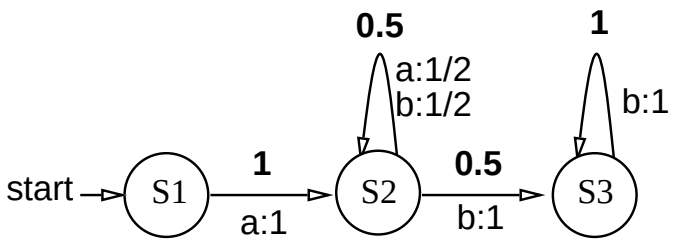

What is the probability this HMM will generate the sequence `aabb` as its first four characters?

What is the most likely sequence of states for that output sequence.

2. **PAC model.** [15 pts]

(a) In class we proved a statement of the following form:

> For any finite hypothesis space and any distribution D on examples, if we see enough data, then with high probability no hypothesis with high true error (with respect to D) will have zero observed error.

Make this statement precise by quantifying "enough," "with high probability,' and "high true error".

(b) Describe at least one specific reason you don't usually need in practice as many examples as specified by the above bound.

(c) In one sentence, give one reason why "Occam's razor" is a good idea.

(d) Consider a distribution and target concept where the probability a random example is positive is 0.5. Suppose your hypothesis space H is just $\{true, false\}$ (all depth-0 decision trees). Is it true that "If I have seen a lot of data and some hypothesis has zero observed error, then with high probability that hypothesis also has low true error?" Why or why not? How is this statement different from the statement in (a)?

3. **Algorithm design.** [30pt]

(a) How would you design a hill-climbing "gradient descent-like" algorithm for decision tree learning? (Your algorithm won't be a true "gradient-descent" since you won't have real-valued parameters, but it should find a local optimum)

(b) Give a mistake-bound learning algorithm for the following situation. Examples have n boolean features, and the target concept requires that for any positive example, certain pairs of features must be equal, and certain pairs must not be equal. For instance, a typical target concept might be: "positive if $x_1 = x_4$ and $x_2 = x_7$ and $x_7 \neq x_9$".

What is your worst-case mistake bound?

4. **Explanation based learning** [15 pts].

(a) Use Prolog-EBG to learn the target concept "days on which Sue goes sailing", denoted by Sailing-Day(?d). Assume the following domain theory:

```
Sailing-Day(?d) :- Raining?(?d,no), School?(?d,no)
Sailing-Day(?d) :- Raining?(?d,no), Windy(?d)
School?(?d,no) :- Weekend(?d)
Weekend(?d) :- Sunday(?d)
Weekend(?d) :- Saturday(?d)
Raining?(?d,no) :- Sunny?(?d,yes)
```

Show the result of running the EBG algorithm on the following positive training example: in particular, show the explanation and the resulting learned rule. The output hypothesis may only use those attributes used to describe the training example.

```
Sailing-Day(May8), During-Finals-Week(May8),  Sue's-Birthday(May8),
Saturday(May8),  Sunny?(May8,yes),  Day-After(May8, May9)
```

(b) Describe a situation in which you would prefer KBANN to EBNN. Describe a situation in which you would prefer EBNN to KBANN. For each situation, explain why (in one sentence).

5. Q-**learning** [15pt]

Consider playing Tic-Tac-Toe against an opponent who plays randomly. In particular, assume the opponent chooses with uniform probability any open space, unless there is a forced move (in which case it makes the obvious correct move).

(a) Formulate the task of learning an optimal Tic-Tac-Toe strategy as a Q-learning task. Specify the states, transitions, and reward (you don't need to list all the states :-)

(b) Describe in one (short) paragraph how explanation-based learning could be used to speed up convergence of Q learning for tic-tac-toe.

(c) Can you imagine, more generally, a game-playing situation where it is *not* best to play the move with the highest *actual* Q value?