

1 Decision Trees I (10 points)

- (a) Draw a minimal-sized decision tree for the three-input XOR function.
- (b) Prove that a decision tree for the n -input XOR function must be a full tree in general.

2 Decision Trees II (10 points)

Do problem 18.5 from the textbook (18.10 from the second edition).

3 Neural Nets I (10 points)

Construct a neural network by hand (provide the architecture, weights and activation function) that computes the XOR function of two inputs.

4 Neural Nets II (10 points)

Demonstrate that a neural network (with a sufficient number of nodes) can compute any function from n binary inputs to one binary output. (Hint: If you haven't taken a circuits class yet, you might want to Google something called *conjunctive normal form*.)

5 Regression I (10 points)

Consider a regression problem with training data t , feature matrix Φ and solution $w = (\Phi^t \Phi)^{-1} \Phi^t t$, such that $\Phi w \approx t$. Suppose we replace each feature ϕ_i (column i) of Φ with a new feature ϕ' such that $\phi'(x) = a\phi(x)$, for some constant a . Show mathematically that if w' is the solution to the regression problem with feature matrix Φ' , then $\Phi w = \Phi' w'$; i.e., scaling all of the features by a constant doesn't change your learned predictor.

6 Regression II (10 points)

A vertical line has infinite slope, something that can cause problems for linear regression.

- (a) Explain the mathematical problem that occurs if you try to fit a line to training data that exactly (or even approximately) form a vertical line.
- (b) If something like this happens, is it because you have too many features or too few features?