CS 391L Machine Learning Final

This exam is open book. You may bring in your homework, class notes and text- books to help you. You will have 1 hour and 15 minutes. Write all answers in the space provided. Please make sure YOUR NAME is on anything you turn in. Square brackets [] denote the points for a question.

1. **Graphical Models** A linear graph inference problem with variables $\{x_1, \ldots, x_N\}$ has been augmented with observed measurements $\{y_1, \ldots, y_N\}$ as shown below. each of the x variables has K values and each of the y variables has M values.

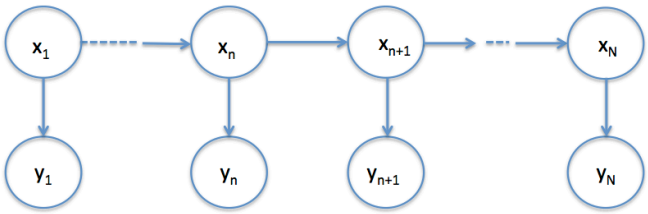

   (a) [5] Augment the graphical model to include plate graphical model conventions.

   (b) [10] Sketch the order of calculations using arrows that would be done to estimate the marginal distribution $p(\mathbf{x}_N)$.

   (c) [10] Suppose that we include an additional pair of variables $\{\mathbf{x}_{N+1}), \mathbf{y}_{N+1})\}$. Show that there is an online algorithm that uses the previous result.

2. **Reinforcement Learning**

   In a standard reinforcement learning setting that can be described by a Markov Decision Process $MDP = \{S, A, T, R\}$ where S is the state space, A is the action space, T is the transition function and R is the reward function.

   (a) [10] To reduce the cost of high-dimensional search spaces one can approximate the value function with low degree-of-freedom parameterized functions. Suppose that the value function can be written as

$$V(s) = \sum_{i=1}^{n} \exp(\theta_i - b(s))^2$$

   where the functions $b_i(s)$ are known. The bootstrapping value function approximation method attempts to choose the θ_i so as to minimize the difference of samples $\hat{V}$,

$$(\hat{V}(s_{t-1}) - (r_t + \gamma \hat{V}(s_t))^2$$

   so how this can be done by computing the appropriate derivative.

   (b) [15] You are in charge of designing an RL algorithm for driving a car. Suppose that instead of having a single large state space, you divide the problem into smaller, specialized state spaces. For example, in following a car while obeying the speed limit, you combine an RL system for maintaining speed with an RL system for following a car head within a certain distance.

      i. Sketch *BRIEFLY* how this might be done.

      ii. Given an example of a problem you might encounter with this strategy.

3. **Genetic Algorithms**
   (a) [10] In the human genome, crossover points are not arbitrary, but are constrained to a subset of all the possible points.
      i. Why do you think this is the case?
      ii. What would be a reason to *not* do this for a particular GA?
   (b) [15] In the animal kingdom speciation occurs when groups of individuals can no longer mate with each other and produce a sucessful offspring. Suppose you are charged with designing a GA that models the speciation process. How would you adapt a standard GA to incorporate a concept like speciation? How would you pick parents? How would you design a way of checking that the offspring of two individuals was unsuccessful? Sketch the details of your solution.

4. **Games** The game of soccer has penalty kicks on goal where a striker faces off with the goalie. The ball travels so fast that the goalie has to guess which side of the net the ball will come to and commit to defend that side. Similarly the striker tries to guess which side the goalie will defend and kick to the other side. Depending on how well they guess determines the success rate. The table shows some estimated data for two top Chinese players. The numbers are the probabilities of success based on their choices.

Goalie
Wang Dalei

Striker
Yu Hai

	Left	Right
Left	.8 / .2	.6 / .4
Right	.3 / .7	.9 / .1

   (a) [5] Is their a *Nash Equilibrium* for this game? Say why or why not.
   (b) [5] Write down expressions for the expected success for each player if Wang chooses Left with probability x and Yu chooses Left with probability y
   (c) [15] By analyzing the payoffs to individual players, derive the optimal probabilistic strategies for both players.