University of Denver
Intro to AI; Spring, 2014
Final Review

Below are topics that should be understood for the exam. This means that you should be able to define or provide an explanation for a term, provide an example of any algorithm, and know the complexities of approaches discussed in class. The lecture notes and the book should also be reviewed in detail.

Chapter 10

PDDL
Problem Formulation
Search Approaches

Chapter 13

Basic probability
Prior Probabilities
Bayes Rule
Independence
Chain Rule
Building Joint Distributions

Chapter 14

Bayesian Networks
Semantics (independence relations)
Markov Blanket
Inference:

- Direct Sampling
- Rejection Sampling
- Likelihood Weighting
- MCMC / Gibbs Sampling

Chapter 17

Reinforcement learning
Bellman equation for value iteration
Q-learning vs. value iteration
TD(λ) - How to train given an eligibility trace for dynamic programming, Monte-Carlo, and how to mix the two with TD(λ)

Chapter 18

Decision Trees
Entropy calculation for splitting
Linear Regression (Perceptron)

- Learning rule

Logistic Regression
Neural Networks
Simple view of kernels (extra features used to make the state space linearly separable)

Chapter 22

N-gram models (characters/words)
Language prediction
IR - recall vs precision
Data extraction

Chapter 25

Robot categories
Active vs passive sensing

- Localization / mapping (SLAM)

Effectors
Workspace vs configuration space

Selected Example Questions:

Encode a PDDL action which loads an item onto an airplane.

Write and explain the equations for computing the probability that you roll three six-sided dice and do not roll a 1.

Given conditional probabilities, build the joint distribution table.

Show the Bayesian network constructed by different orderings of variables.

Given a Bayesian network, determine whether two variables are independent or dependent given provided evidence.

Describe how rejection sampling works.

What is the Bellman equation? How can you use it to determine if the values in a MDP have converged?

Given an eligibility trace, show the training that will be performed for each state using Monte-Carlo learning.

Calculate the entropy of a splitting on a given variable next in a decision tree.

Given a perceptron, show how the weights would be updated given a training example.

Why can XOR not be learned by a simple perceptron? How could you modify the inputs to allow XOR to be learned?

What is the difference between logistic and linear regression?

Given sample text, compute the bigrams in the text.

Given two bigram models (one for two different languages) and a word, how would you compute the probability that the word is in each language?

How can an algorithms maximize its precision? Its recall?

Given the workspace for a robotic arm, draw the configuration space.