What to Expect on the Final Exam

The exam will be held in 124 Ryder Hall on Friday, April 18, 8-10 a.m. The exam is open book: You may use your class notes, homeworks, handouts, and the textbook.

The exam will cover only material presented since the midterm, except to the extent that earlier material is relevant to this more recent material. (In particular, some of the earlier material on propositional and first-order logic is relevant to the extent that it is necessary to understand the corresponding inference methods, and, of course, depth-first search is relevant to understanding how backward chaining is implemented.)

Since the final exam is not comprehensive over the entire semester, the relative weightings of the midterm and final exams will be adjusted from what was announced in the syllabus to whatever gives you the best overall result. This means that whichever of the two scores is higher will count 40% and the lower one will count 25% when determining your overall course grade.

Relevant Material

Propositional Logic

- CNF, clause form
- inference
  - forward chaining
  - backward chaining
  - resolution

Readings: 7.3-7.5, prop-logic-equiv.handout

First-order Logic

- equivalences involving quantifiers, skolemization, unification
- inference
  - forward chaining
  - backward chaining
  - resolution

Readings: 8.1-8.3, pred-calc-equiv.handout, skolem.handout, unify.handout, 9.1-9.5 (only through p. 300), pred-calc-infer.handout, chaining.handout, bchain.handout

Representing and reasoning with uncertainty

- Probability, joint distribution, conditional probability, Bayes' rule
- Bayes nets
  - CPTs, what the presence or absence of arrows means

  - exact inference through enumeration
  - approximate inference: rejection sampling vs. likelihood weighting

Readings: 13.1-13.6, 14.1-14.2, 14.4-14.5 (only through p. 516)

Supervised learning

- naive Bayes classifiers
- decision trees
- neural networks: perceptrons, backprop nets

Readings: 18.1-18.3, 20.5, *Learning Decision Trees* & *Artificial Neural Networks* presentation slides (You will not be expected to know the details of the information gain calculation for constructing decision trees from data, nor will you be expected to know the details of the perceptron or backpropagation learning algorithms.)