

CS470/570

Exam 1 Topics sheet

Exam 1 will focus on Ch1-4; you are responsible for all material covered in those chapters.

Although there may be some “factual” name-and-explain answers just to give away some easy points, most questions will ask you to demonstrate that you have read *and understood* the material in more than a superficial memorizing fashion. This is typically tested by asking you to *apply* techniques to sample data, discussing implications, or comparing and contrasting approaches. If you've been reading carefully, working thoughtfully through the reading/homework exercises, and engaging in the lecture, you should just need some thorough review to succeed.

As a reminder, here are some sample of the sort of things I might ask, and some skills you should have ready to demonstrate.

Ch1: Intro

- What is AI? What are the two main perspectives and how does each perspective shape what you'd accept as an AI "solution"?
- Rational agents as one perspective on designing AI systems. What is rationality anyway?
- History of AI: major milestones/phases of development. Significance of each.

Ch2: Rational Agents

- Rationality vs perfection. What's the diff?
- PEAS descriptions: what they are and how to use them. Be able to give and rationalize a solid PEAS description for a hypothetical application.
- Characterizing task environments along the six dimensions outlined. Be able to characterize given hypothetical application environments, including solid rationale.
- Types of agents, how they work, and differences between them. They vary based on the elements/functions that drive them internally...be able to describe each.

Ch3 and Ch4: The basic and advanced searching

- The basic tree/graph search algo: frontier queue, closed lists, successor insertion. How all searches are just variants on the same basic algorithm
- Basic search algorithms: be able "run" them on a sample search space, showing progression of each algo. Be able to intelligently discuss (versus just regurgitate) pros, cons, complexity, optimality, completeness
- Local search: hill-climbing and its variations, genetic approaches, simulated annealing. Motivation/limitations for each, what sorts of problems they are good for
- Being able to choose the right tool for the job is a key skill. Given a scenario, be able to analyze it to choose the right algorithm for the job...and explaining your reasoning.

Searching in uncertain worlds:

- Be able to construct search trees for non-deterministic actions/sensors. Do you understand the AND/OR tree spaces...and the “solutions” that are produced?
- Searching with partial observations: Be able to construct (partial) belief state spaces and show how belief states evolve based on availability/quality of percepts.
- Online Search: When is it useful? Pros/cons versus the others. Competitive ratios, reversibility of actions, and other concerns special to online search.
- Adaptation of hill-climbing to online search, e.g. LRTA*.