

Exam II in Nov. 23.

LCR 1-7, 10. TCS 1-18. Same format as Exam I.

1 page of notes allowed.

Review should be available tonight or tomorrow.

There will be a review session.

Artificial Intelligence (just a taste):

Many important programming techniques and programming ideas had their origin in AI programming. So new ideas often are invented and tried first in AI, and then applied to other areas.

Definition: AI is the art of getting the computer to do something that, if done by a human, would be considered intelligent.

Examples: Playing chess and other board games, proving theorems in mathematics or logic, using language, solving problems, etc.

(Some AI researchers complain that as soon as you succeed in getting a computer to do something, it's not considered AI anymore!)

A broader definition of AI is getting a computer to exhibit human-like abilities.

Notice that all the examples I listed are intellectual types of activities (things that take place primarily in the mind). Chess, math., logic, problem solving, even language. These are all characteristic of human intelligence (i.e., things humans do that most other animals don't). AI researchers were most interested in human intelligence; animal intelligence was considered less interesting. Also it allowed AI researchers to ignore all the messiness and complication of the real, physical world.

This whole approach is sometimes called GOFAI = "Good, old-fashioned AI."

It didn't work very well.

There are newer approaches to AI (last 20 years), which recognize the importance of the body and of sensorimotor coordination as a foundation for intelligence.

Like other other animals, we can move competently in the world, handle objects (play sports), recognize faces, find our way around, etc. etc.

There is a "cognitive inversion": what animals do very well and easily, is difficult for traditional AI systems, and what is easy for trad. AI is difficult for people and other animals. This suggests that brains and trad. AI systems do things very differently. So there are new approaches to AI based on neural networks, which try to do things more like the brain.

Also, embodied AI recognizes the essential role that the body plays as a foundation for higher levels of intelligence. We know implicitly a lot about the world because we have a body that is part of it. This is one reason that robotics is more important in recent AI work.

Be that as it may, in ch. 10 we have two examples of GOFAI:
Tic-tac-toe and a "smart" Paper, Scissors, Rock.

Psychologists have recognized the value of autotelic activities — activities that are their own end, or an end in themselves, i.e., play. Humans and other animals maintain a balance between:

- success (which in excess leads to boredom)

- challenge (which in excess leads to frustration)

So when we become successful, we "raise the bar," but not to much, and if the challenge is to great, we lower it.

Games:

There principal goals (and stages):

- (1) Get the program to play legally.
- (2) Get the program to play well.

Board game playing is a very stripped-down kind of problem solving. They avoid the messiness of physical reality, but still address problem solving, strategies, etc.

Key Techniques:

(1) Generate and Test: Generate possible moves and test them to see which are good.

If possible, generate only only legal moves, so you don't waste time testing illegal moves.

(2) Heuristics: Rules of thumb, or evaluation techniques, that are not guaranteed to work, but often do.

(3) Lookahead: Look multiple "moves" into the future to see which are likely to turn out best. We generally speak (for board games) in term of plies (pairs of turns, you and your opponent). So for example 5-ply lookahead.

The enemy is combinatorial explosion. The number of possibilities to be generated

and tested grows exponentially (or worse) with the number of plies. This is why the ability of chess and similar programs has tended to improve with computer speed. But ultimately combinatorial explosion defeats the fastest computers.

It doesn't seem like human chess masters do this. They can focus in on the most relevant or interesting moves, and pursue them to great depth. Less interesting ones are ignored. We really don't understand this.

Look at the Tic-tac-toe program from LCR 10.