

The “Structured Matcher” Paper

Work in groups of three or four to...

- Summarize the paper in two sentences or less.
- Identify the three most important ideas in the paper.
- Identify three places in the paper that most need improvement, and suggest specific changes.

One Big Idea: Exploit Regularities

Structured Matcher describes a regularity in the world:

- a problem that is common in many domains, and
- a solution that is common to a lot of problem solvers.

The Principle of Convergent Intelligence tells us that intelligent agents must exploit regularities in the world.

- They recognize the regularity.
- They use what they know about the regularity to perform better.

The world of problem solving is full of regularities of this sort.

1. Selecting an acquisition structuring plan.
2. Determining what is wrong with a patient.

The Regularities Behind Structured Matching

What is common?

- We often know all possible answers to a question, or nearly all.
- We often know that some answers are more specific varieties of others.
- We use generality and specificity as a way to organize our knowledge about the domain.

One way to try to understand how intelligent agents solve problems is to try to identify what regularities they use—and how they use them.

Now that you know all about AI...

We say that computer science has **solved a problem** if we can build a program to solve the problem or do the task.

Work together to construct two lists for me:

- Which activities (or phenomena) that you consider *intelligent* do you think AI will be able to solve, that is, be able to construct programs that carry out the activities (or exhibit the phenomena) as well as or better than any human?
- Which activities (or phenomena) that you consider *intelligent* do you think AI will **not** be able to solve?

Since this is the last day of class, please humor me and spend all of the allotted time actually working on the exercise. :-)

A Few More Questions

Are there any problems that you think we should not use AI to study?

Has your opinion of the Turing Test changed at all?

Will a computer ever be *conscious* in the way humans are?
Will a computer ever have *emotions* in the way humans do?

Who will be legally responsible when the first intelligent program commits a crime?

The AI Enterprise

Artificial intelligence is the computational study of how a system can perceive, reason, and act in complex environments.

- agents that solve problems
- agents that reason logically
- agents that plan their actions
- agents that can handle uncertainty
- agents that use knowledge of world and self
- agents that learn
 - from observation
 - from experience
 - using existing knowledge
- agents that communicate
- agents that use physical effectors

Three Views of AI

There are at least three ways to view AI:

- as a set of programming techniques
- as a means for simulating intelligence
- as a means for studying (real) intelligence

At the simplest level, each ultimately relies on the same notion: the concepts of AI constitute a **language** for describing the phenomena of intelligent behavior.

Progressing through the list, one makes increasingly stronger commitments to what “describe” means in that last sentence.

For the Exam

What is Artificial Intelligence?

Agents that Search for Solutions

Agents and their Goals

Uninformed Search

Search and Knowledge

Heuristic Search

Search in Adversarial Environments

Knowledge in Adversarial Search

Agents that Reason Logically

Knowledge Representation

Logic as Knowledge Representation

Modus Ponens as a Reasoning Mechanism

Resolution as a Reasoning Mechanism

Agents that Reason in the Real World

Reasoning in the Face of Uncertainty

Explaining Conclusions

Agents that Learn

Machine Learning Concepts

Learning by Induction

Learning by Evolution (GAs, GP)

Knowledge-Based Learning

Agents that Plan their Actions

Planning Concepts

Goal-Stack Planning

Plan-Space Planning

Structured Matching

Wednesday, December 19, 1:00 PM - 2:50 PM