

Referring to the design of the *characteristica universalis*, a logical language capable of expressing knowledge about logic, science, and metaphysics:

And so I repeat, what I have often said, that a man who is neither a prophet nor a prince can ever undertake any thing of greater good to mankind or more fitting for divine glory.

— *Gottfried Leibnitz*

Principles of AI

- ▶ Undergraduate Computer Science and Data Science Majors
- ▶ Not required for CS but is required for Data
- ▶ Course limit is set at 22 but currently has 28 people.
- ▶ Only taught in the Spring.

Course Requirements

- ▶ To satisfy course requirements, I teach Genetic Algorithms, Neural Nets, and Natural Language Processing topics in a survey format.
- ▶ I use ASP to teach the rest of the required topics: knowledge representation, intelligent agents, reasoning, and AI programming.

ASP Text and Resources

- ▶ Book: *Knowledge Representation and Reasoning and the Design of Intelligent Agents: The Answer-Set Programming Approach*
- ▶ You can find course slides, programs from the book, tips for students, exercises, and other resources for the book at <http://pages.suddenlink.net/ykahl/>
- ▶ Homework: email me at ygkahl@gmail.com and reference this talk.

Other Resources

- ▶ I supplement with the chapters on Genetic Algorithms and Neural Nets from *The Nature of Code* by Daniel Shiffman, free online at <http://natureofcode.com/book>, but I don't put much emphasis on this.
- ▶ For Neural Nets, I highly recommend three videos from Grant Sanderson of 3Blue1Brown:
 - ▶ Introduction to Neural Nets:
<https://www.youtube.com/watch?v=aircAruvnKk&t>
 - ▶ Gradient Descent:
<https://www.youtube.com/watch?v=IHZwWFHwa-w&t>
 - ▶ Introduction to Backpropagation:
<https://www.youtube.com/watch?v=Ilg3gGewQ5U&t>

Don't Miss the Forest for the Trees

- ▶ In class, students focus heavily on new vocabulary and notation.
- ▶ My goal is to keep them focused on *ideas* represented by this vocabulary and tie them to types of agent reasoning.
- ▶ Taken all together, these types of reasoning fit into the agent loop needed for intelligent action in the world.

How Should an Agent Reason?

One suggestion is the ASP way:

1. Satisfy the rules of its program.
2. Not believe in contradictions.
3. Adhere to the “Rationality Principle” which says: “Believe nothing you are not forced to believe.”

This intuition forms the basis for the formal definition of Answer Set.

Defaults

An intelligent agent should:

- ▶ Be able to reason with general statements and
- ▶ override its generalizations in favor of exceptions.

This idea is realized via the technical notion of default and various forms of exceptions.

Action Languages

- ▶ An agent should be able to model the effects of actions.
 - We can do this with causal laws.
- ▶ It should be able to know about things that remain unchanged. (The Frame Problem)
 - Dramatic pause. Students thinking to themselves, “Oh man! That is a problem.”
 - Guess what? Here’s a solution! We use a default.
- ▶ An agent should be able to understand effects of indirect actions. (Ramification Problem)
 - We got that covered too: use state constraints.
- ▶ An agent must know that there are constraints on actions.
 - This is done with executability conditions.

Planning and Diagnostics

- ▶ Now that you've learned to represent some basic laws of the world, we just add this module, and we get planning.
- ▶ Things going too slow? Add domain-specific heuristics.
- ▶ Now a few more tweaks, and we get diagnostics.
- ▶ Notice how each of these things fits into the agent loop.

In Conclusion

- ▶ I believe tying the technical aspects of ASP to ideas of correct reasoning makes the class more interesting and provides a good framework for learning.
- ▶ Most people were able to pass the class and many did well.