

CSE 473: Artificial Intelligence

The Future of Artificial Intelligence





Outline

Pros and Cons of high-level AI
Asimov's Three Laws of Robotics
Will they be like us?
Tools vs Agents
Technological Challenges
Social Challenges



Pros and Cons of Achieving High Levels of AI

Pros...

Powerful tools, solutions to tough problems, better standards of living(?)

Cons...

Tools might be used against people; technology may create worse problems than it solves; standards of living might get worse; we might feel we lose some aspect of our humanity.



Isaac Asimov's Three Laws of Robotics (1940)

First Law: A robot may not injure a human or through inaction, allow a human to come to harm.

Second Law: A robot must obey the orders given it by human beings, unless such orders would conflict with the first law.

Third Law: A robot must protect its own existence, as long as such protection does not conflict with the first or second law.



Are the 3 Laws the Answer?

Carrying out the laws requires very sophisticated judgment.

Extending the Laws(?!)

Zeroth law: A robot may not injure humanity or through inaction allow humanity to come to harm.

(due to Asimov, Olivaw, and Calvin).

David Langford's tongue-in-cheek extensions, acknowledging military funding for robotics:

4. A robot will not harm authorized Government personnel but will terminate intruders with extreme prejudice.

5. A robot will obey the orders of authorized personnel except where such orders conflict with the Third Law.

6. A robot will guard its own existence with lethal antipersonnel weaponry, because a robot is bloody expensive.

Will They Be Like Us?

Like us, AI systems...

- ...will talk to us in our languages.
- ...will help us with our problems.
- ...will have anthropomorphic interfaces.

Unlike us, AI systems...

- ...will compute and communicate extremely quickly.
- ...will have bounds for learning and retention of knowledge that will soon surpass ours.
- ...might not be well modeled by the psychological models that work for people.



Tools vs Agents

Agent: Takes responsibility, takes initiative, interacts with others on behalf of a client.

Tool: Responds directly to its user. Does not take responsibility. Does not take initiative. Does not normally interact with others on behalf of a client.

Technological Challenges

Giving computers “common sense” is still an unrealized goal.

Human language is diverse (there are many languages, dialects, and idiolects) and often ambiguous. Computers don't yet understand it well.

General AI systems that can redesign themselves.



Achieving the Transition to the Next Epoch of Civilization

Ray Kurzweil, in *The Singularity Is Near*, says that it will happen in the next century; maybe in the next generation.

AI will soon reach a “singularity” in which its effectiveness suddenly increases enormously, due to a virtuous circle of improvements.

Aspects of the Singularity

Convergence: It results from multiple exponential trends – computer fab. Technology, brain science, genome science, and A.I.

Immortality: A person could just “upload” her/his personality to the system (cloud? Matrix?)

Superhuman intelligence will be benevolent.

AI will be the vehicle to spread human civilization throughout the universe.



Singularity: Questions

What aspects do you consider believable?

Does it help to make a goal out of it?

How would economics relate to it?

Won't politics or religion get in the way?

How do we make sure it will be good?

Where are the best opportunities to contribute?

Social Challenges

Users need to understand the limits of their tools and agents. (Expert systems tend to be brittle)

AI applications need to be created that help bring harmony to the world rather than which intensify battles.

AI applications are needed which enhance the economy rather than reduce economic competition.

AI extends the reach of automation and threatens to eliminate, if not change many white-collar jobs.

AI raises the bar for information literacy and computer literacy.