


The End (and look ahead)

R&N Chapter 26
CS 121
August 12, 2009
Chris Archibald

Back to the beginning

- What is AI?
 - Discipline that systemizes and automates intellectual tasks to create machines that:

Act like humans	Act rationally
Think like humans	Think rationally



More formal and mathematical

Achievements of AI

- Logic reasoning systems
- Search and game playing (chess, checkers)
- Knowledge-based systems
- Bayesian networks (diagnosis)
- Machine learning and data mining
- Planning and military logistics
- Autonomous robots

Learning



THE POWER OF T⁴ TECHNOLOGY

CyberKnife[®] LightScribe Tumor (T⁴) radiosurgery with Ultimate Conformality

INTEGRATION OF TWO REVOLUTIONARY TECHNOLOGIES

Proprietary Image-Guidance System
Tracks and sets the tumor location to enable accurate cone projections for tumor treatment.

Multi-Jointed Robotic Arm
Enables access to previously unreachable tumors and reduces damage to surrounding vital structures.

Integration of these unique technologies allows physicians to treat complex-shaped tumors with clinical precision necessary that has been demonstrated to be comparable, if not superior, to frame-based radiosurgical systems.

Simple Outpatient Treatment Process

Planning: CT scanning and advanced treatment planning are utilized.

Positioning: The patient lies on a table with only a face mask or body mold used for immobilization.

Verification: The image guidance system verifies tumor location and compares it to the planned target area.

Targeting: When tumor movement is detected, the robotic arm is repositioned within a fraction of a second.

Repeat: The verification process is repeated prior to delivery of each radiation beam.


Treatment: A series of highly collimated radiation beams deliver precise radiosurgery to the tumor.

Completion: Following CyberKnife[®] treatment, the patient goes home. There is zero recovery time.

**FULL-BODY
100% Frameless
T⁴ Radiosurgery**

CyberKnife[®] T⁴ Radiosurgery
A new standard in RTT conformality

Dealing with the real world (uncertainty)



Working as a team

Constraint satisfaction and robotics

Views on AI

- Weak AI: Can machines act intelligently?
- Strong AI: Can machines actually think?
- Let's look at a few topics and arguments related to each viewpoint.

Weak AI

- Some argue that AI is impossible
 - This depends on how you define AI
 - If AI is “quest for best agent program on a given architecture”, then clearly it is possible
- Maybe “Can machines think?” isn't the right question
 - Can airplanes fly? Can boats swim?
- Turing test focuses on intelligent behavior

Weak AI

- Possible objections to the possibility of intelligent machines
 - Argument from disability
 - A machine can never do X
 - “Be kind, resourceful, beautiful, friendly, have initiative, have a sense of humor, tell right from wrong, make mistakes, fall in love, enjoy strawberries and cream, make someone fall in love with it, learn from experience, use words properly, be the subject of its own thought, have as much diversity of behavior as man, do something really new.”
 - Machines have done many things that were previously thought to require human intelligence
 - Chess, checkers, drive, fly, diagnose, discover things, even judgement

Weak AI

- Possible objections to the possibility of intelligent machines
 - Mathematical objection
 - Certain mathematical questions are unanswerable by formal systems
 - Godel's incompleteness theorem
 - People, it is argued, are not limited in this way
 - Bob cannot consistently assert that this sentence is true
 - Does this reflect poorly on Bob?
 - Often we can rigorously prove that a formal system cannot prove X , but then we intuitively argue that a human *can* do X , without any rigorous evidence or proof

Strong AI

- Some philosophers argue that a machine which passed the Turing test would still only be *simulating* thinking
- But, do we have evidence that other humans are thinking?
 - Turing feels that these mysteries need not be solved before useful things can be done
- General question: Are artificially created things real?
 - Urea, sweeteners, insemination - Yes
 - Flowers, wine, paintings – no
- Behavior / Pedigree

Chinese Room

- John Searle (1980)
 - Thought experiment attempting to answer these questions
 - English human in a room
 - Room contains rule book (in english), paper, Chinese symbols
 - Paper is passed into the room, with Chinese writing.
 - Human uses the rules to create output Chinese writing for this input.
 - From outside, the room understands Chinese
 - But, there is nothing in the room that understands Chinese
 - Not the person, not the paper and book.
 - So, a running program does not constitute understanding

Ethics of AI

- Problems AI poses
 - People might lose their jobs to automation
 - People might have too much (or too little) leisure time
 - People might lose their sense of being unique
 - People might lose some of their privacy rights
 - The use of AI systems might result in a loss of accountability
 - The success of AI might mean the end of the human race

Ethics of AI

- Loss of accountability
 - What if a AI diagnosis is wrong? Who is at fault?
 - What if a trading agent incurs debt? Who pays it?
 - What if an autonomous car wrecks, who is liable?
- The law has yet to catch up with many of these eventualities
- Society has yet to catch up as well
 - We forgive human errors
 - What about AI errors?

Conclusion

- My own feelings are that AI gives us many useful tools
- What we use these tools for depends on us
- I hope you can all take the AI techniques you have learned this quarter to be able to:
 - Solve problems you face
 - Understand new technologies better
 - Inspire you to study these topics more!

Conclusion

Thanks!