Applications and Ethical Issues

CS 486/686:Introduction to Artificial Intelligence

Example Applications

- Security Games

  - Limited security resources: Selective checking

  - Adversary monitors defenses, exploits patterns

- How do you assign limited resources to defend targets?

Example Applications

ARMOR
(LAX)

PROTECT

Example Applications

- Beyond Security

  - Scheduling ticket checking/collection on the LAX metro

  - Patrolling for poachers/illegal timber harvesters

  - ...

Wildfire Resource Management

Wildfire Resource Management

Advertising

- **Targeted ad campaigns**

  - Numerous AI techniques

    - naive Bayes

    - SVM

    - decision trees

    - NLP (including bag of words model, stemming)

    - graphical models (i.e. Bayes Nets)

Health Domain

- Machine learning for detection of fetal hypoxia during labour and delivery

  - Support vector machines

- Fall detection amongst the elderly

  - Decision trees

- Rehabilitation after strokes

  - MDPs/POMDPs

- Robotic exercise coaches

Other Domains

- Recommendation Systems

- Tutoring systems

- Biodiversity and conservation research

- Transportation scheduling

- Disaster management

- ...

Ethical Issues

- In this course we have looked at whether we ***can*** develop AI

  - We also need to consider whether we ***should***

- All new technologies have unintended consequences

  - Is AI different than other fields?

  - "*We first shape the tools and thereafter our tools shape us*", Marshall McLuhan (1964)

Service Robots

- PaPeRo

  - child-minding robot

  - video-game playing, conduct conversations, face recognition, determine if a child goes out of range,...

  - tests show that children like them (bond well with them)

  - But, should children be left in the care of a robot for longer periods of time?

    - Importance of human contact? Negligence?

Service Robots

- Relative increase in the number of elderly compared to younger caregivers

- Pros: allows the elderly to maintain independence

- Cons: reduction in human contact leading to isolation

Military Uses of AI

- DARPA and BAE (for example) are big funders of AI research

- Unmanned vehicles

- Robotic weapons

- Wars between machines?

- Human decision-making taken out of the loop?

Privacy Issues

- AI has the potential to mass-produce surveillance

  - The UK has a large network of surveillance cameras

  - Other countries monitor Web traffic and telephone calls (see recent news)

  - AI technology helps in the data processing/ information gathering

Accountability Issues

- Medical systems

  - what if a doctor using a medical diagnostic system ends up making an incorrect diagnosis?

- Autonomous vehicles

  - who is at fault if an autonomous vehicle is in an accident (or is speeding)

- Monetary transactions

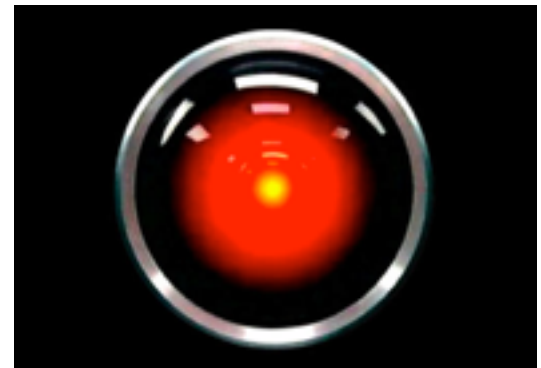 - who is responsible if a trading agent misbehaves?

AI and the End of the Human Race!

- Any technology in the wrong hands can cause harm

  - but what if the "wrong hands" is the technology itself?

  - A lot of science fiction literature on this topic

The Matrix
Terminator
etc...

AI and the End of the Human Race!

- Any technology in the wrong hands can cause harm

  - but what if the "wrong hands" is the technology itself?

  - A lot of science fiction literature on this topic

- Where can things go horribly wrong?

  - Utility function specification

  - Ultraintelligent machines and the "intelligence explosion"

The Singularity

- Will this really be an issue?

    - Technological progress might be limited

    - Limits on computability and computational complexity

    - Note: some people embrace the idea of the singularity (eg. transhumanism and Ray Kurzweil)

The Singularity

- If the Singularity is coming then maybe we should design our current AI systems so that they will treat us well

  - eg. Asimov's Three Laws of Robotics

  - What is the right design? Should it evolve?

The Singularity

- If the Singularity is coming then maybe we should design our current AI systems so that they will treat us well

  - eg. Asimov's Three Laws of Robotics

  - What is the right design? Should it evolve?

- If robots/AI systems are so intelligent then what rights should they have?