Cognitive Science Guest Lecture

Overall Topic: AI, Robotics

Introduction to AI & Robotics

History of AI

The early years

- AI started not long after programmable computers were developed (50's)
  - Herb Simon and Alan Newell did the first AI program using punch cards
- Early focus of AI were problems with clean representations
  - checkers, chess, theorem provers
- They were successful in these endeavors largely because of the domains they chose

The glory years

- With AI's success in game domains and other expert systems, AI researchers turned to the fundamental percption/action problems
  - vision
  - speech
  - navigation
- Attempts to apply high-level reasoning to these problems achieved limited success
  - Navigation programs were slow, if they ran on real robots at all
  - Vision processing couldn't make the leap from images to concepts
  - Speech recognition failed because low level representations were poor
- It turns out that action/perception problems in the real world are some of the hardest problems that exist
  - The amount of data is incredibly large
  - The data is incredibly rich and dynamic
  - Sensors contain noise
- What happened is that each of these areas because its own field
  - Speech recognition & synthesis
  - Computer vision & image processing
  - Robotics & navigation

Modern AI

- Modern AI has found a niche in both industry and academic research
  - In one sense, AI is a set of tools that can be applied to hard problems
    - The problems tend to be ones that have abstract/clean inputs
  - In another sense, AI is starting to look again at issues in cognitive science
    - Using computational models to examine the emergence of behaviors
    - Integrating low-level subsymbolic computation (neural networks)
  - AI is trying to apply itself to real-world problems
    - Tough to get funded or build a product unless it works on a real problem

History of Robotics (non-industrial)

The early years

- Most of the early research in mobile robots was by AI researchers
  - Some of it in simulation, some on real robots
- Beginnings of integrating real sensors with AI techniques
  - Computational constraints and hard algorithms made the robot real slow

Sense-plan-act

- General method of running a robot was to do the following in a loop
  - sense: take the sensor readings and add them to a model of the world
  - plan: based on the world model, plan some actions
  - act: execute one or more actions
- Sensing and adding the sensors to a world model took time
- Planning based on the world model takes more time (can be exponentially hard)
- Acting can occur well after the sensing operation (hope the world doesn't change)

The reactionist years

- Rod Brooks at MIT said AI was a bunch of people deluding themselves
- Proposed ditching internal representations and connecting controllers directly to sensors
  - Let the world be your internal representations, and sensors be the interface
- Proposed that robots must be
  - Embodied: working on simulated robots is worse than useless
  - Situated: robots must be able to sense the complex world around them
  - Motivated: robots should be motivated to act (somehow)
- Brooks built several behavior-based robots
  - These robots were able to do simple tasks in real time with simple computing
  - Using more complex systems, the robots could do apparent goal-directed behavior

Comments on Brooks' paper

- There are a couple of things to think about
  - Can you really do complex intelligent sequences with no internal representations?
  - People have both short-term and long-term memory = internal representations
  - When you let the world be your model you depend on having good sensors
    - Especially if you want to have speech or vision as input
- Connectionist models have driven robots
  - Researchers (and students) have trained neural networks to drive a robot while avoiding obstacles, as well as other more complex behaviors
  - The most successful connectionist robots have used recurrent neural networks
    - These have internal representations of state

Behavior-based robotics

- General method: have many simple behaviors running in parallel
  - Each behavior is connected directly to the sensors
  - Each behavior outputs a set of motor control signals
  - Behaviors can turn themselves on and off based on the sensor input
  - Behaviors have priority: the highest priority behavior that is on wins
- Example robot
  - Avoid objects behavior: takes over when the sensors say there is something close
  - Move forward behavior: keeps the robot moving forward when it can
  - Explore behavior: gets the robot to head in a different direction based on sensors

The reaction to reactionism

- Behavior-based robots are great at avoiding obstacles
- Don’t hire one to drive your car
  - Can’t necessarily predict what the robot will do
  - Funny sensor readings will cause the robot to do unpredictable things
- After Brooks’ work, people started to look at more sophisticated architectures
  - Continue to decompose the problem into behaviors
  - Put reactive behaviors at the low level
  - Put delibrative behaviors at higher levels
  - Allow for more sophisticated control signal combinations
    - combine signals from different layers more intelligently
    - e.g. as you get close to an obstacle, incrementally give control to the avoid module
- Deliberative-reactive hybrids show the most promise for robotics
  - Combine the strengths of deliberative AI and reactionist approach
    - planning and expert systems
    - immediate responsiveness to the environment
  - Result is an anytime-architecture that can both plan and react

Wrap-up

The connection to Cognitive Science

- Complex behaviors can emerge from simple mechanisms
  - robotics
  - flocking behavior
- Deliberative behaviors seem to need something at a higher level
  - Multi-layer architectures have been quite successful
  - Difficulty is in learning
    - Learning at the deliberative level involves adding to knowledge
    - Learning at the reactive level can be done in a connectionist manner
    - How do you connect the two?

Where is robotics headed?

- We want robots to continue learning after they’ve acquired basic behaviors
- We want to integrate deliberative modules with reactive modules in such a way that the reactive module can pass more abstract concepts to the deliberate module.
- Robots can now start to behave appropriately in a dynamic environment