Robotics and Human-Robot Interaction

AI Class 27 (no reading)

Slides based in part on www.jhu.edu/virtlab/course-info/ei/ppt/robotics-part1.ppt and -part2.ppt and Intro to AI, Dr. Paula Matuszek, Villanova 2013

Bookkeeping

- Closing in! Almost there!
- Doodle poll for review date (tentative: 16th)
- Last schedule slips
  - Phase II: due 11:59 Dec 12
- Final survey
  - How did the project go? Who contributed what?
  - Due before final
- TTBOMK, all Phase II materials are up

Today's Class

- What's a robot (really)?
- What parts do they have?
- What are they used for?
- What kind of AI do they need?
- HRI
- Future Questions

Familiar Robots

ED-209. Robocop: 2014

Optimus Prime. Transformers: 2007-current

Sentinel. X-Men, Days of Future Past: 2014

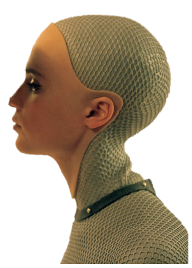

Ava. Ex Machina: 2016

Wall•E: 2008

Some Current Robots

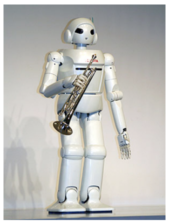





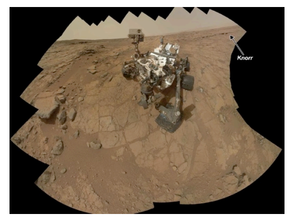

What is a Robot?

- "A robot is a reprogrammable, multifunctional manipulator designed to move … through variable programmed motions for the performance of a variety of tasks." (Robot Institute of America)
- "A robot is a one-armed, blind idiot with limited memory and which cannot speak, see, or hear."
- **In practice:** robotics intersects with any space in which computers move into the physical world.

What Are They Good At?

- What is hard for humans is easy for robots.
  - Repetitive tasks.
  - Continuous operation.
  - Complicated calculations.
  - Referring to huge databases/knowledge sources.
- What is easy for a human is (sometimes) hard for robots.
  - Reasoning.
  - Adapting to new situations.
  - Flexible to changing requirements.
  - Integrating multiple sensors.
  - Resolving conflicting data.
  - Synthesizing unrelated information.
  - Creativity.

What Should They Do?

- Boring and/or repetitive
  - welding car frames
  - part pick and place
  - manufacturing parts
- High precision / speed
  - electronics testing
  - surgery
  - precision machining
- Dangerous
  - chemical spill cleanup
  - disarming bombs
- Inaccessible
  - space exploration
  - disaster cleanup
- All of the Above
  - Continuous reef monitoring
  - Military surveillance

Categories of Robot Systems

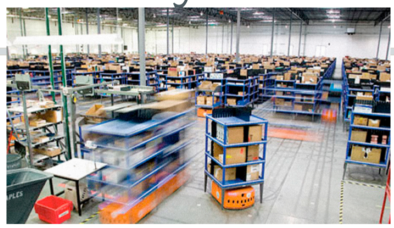

- Manipulators
  - Anchored somewhere
  - Factory assembly lines
  - International Space Station
  - Hospitals
  - Common industrial robots
- Mobile Robots
  - Move around environment
  - UGVs, UAVs, AUVs, UUVs
  - Mars rovers, delivery bots, ocean explorers
- Mobile Manipulators
  - Both move and manipulate
  - Packbot, humanoid robots

Subsystems

Robots have:

- Sensors
  - Some way of *detecting* the world
- Effectors
  - Some way of *affecting* things in the world
  - Manipulation
  - Mobility
- Control/Software

Sensors

- Perceive the world
  - **Passive** sensors capture signals from environment. (cameras)
  - **Active** sensors probe the environment (sonar)
- What are they sensing?
  - The environment (range finders, obstacle detection)
  - The robot's location (gps, wireless stations)
  - Robot's own internals: *proprioceptive* sensors
    - Stop and think about that one for a moment. Close your eyes - where's your hand? Move it - where is it now?

What Are Sensors Used For?

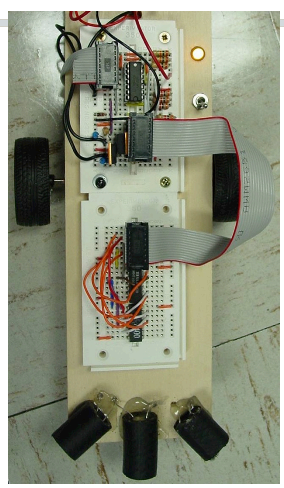

- Feedback
  - Closed-loop robots use sensors in conjunction with actuators to gain higher accuracy – servo motors.
- Decision making
  - Mobile robotics
  - Telepresence
  - Search and rescue
  - Pick and place (with vision)
- Human interaction

Some Sensors

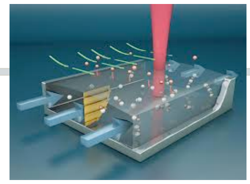

- Optical
  - Laser / radar
  - 3D
  - Color spectrum
- Pressure
- Temperature
- Chemical
- Motion & Accelerometer
- Acoustic
  - Ultrasonic
- E-field Sensing

Actuators / Effectors

- Take some kind of action in the world
- Involve movement of robot or subcomponent of robot
- Robot actions include
  - Pick and place: Move items between points
  - Continuous path control: Move along a programmable path
  - Sensory: Employ sensors for feedback (e-field sensing)

Mobility

- Legs
- Wheels
- Tracks
- Crawls
- Rolls

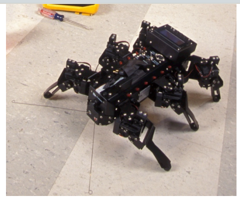

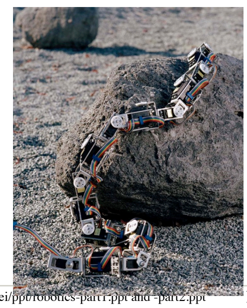

Control: The Brain

- Open loop, i.e., no feedback, deterministic
  - Instructions
  - Rules
- Closed loop, i.e., feedback
  - Learn
  - Adapt

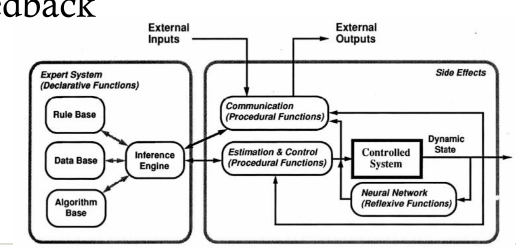

Where Is AI Needed?

- Sensing:
  - Interpreting incoming information
    - Machine vision, signal processing
    - Language understanding
- Actuation:
  - What to do with manipulators and how
    - Motion planning and path planning
- Control:
  - Managing large search spaces and complexity
    - Accelerating masses produce vibration, elastic deformations in links.
    - Torques, stresses on end actuator
    - Feedback loops
- Firmware and software:
  - Especially with more intelligent approaches!

Robotic Perception

- Sensing isn't enough: need to *act* on data sensed
  - Data are noisy
  - Environment is dynamic and partially observable
- Must be mapped into an internal *representation*
- Good representations:
  - Contain enough information for good decisions
  - Are structured for efficient updating
  - Are a natural (usable) mapping between representation and real world

Belief State

- Belief state: model of the state of the environment (including the robot)
  - X: set of variables describing the environment
  - X_t: state at time t
  - Z_t: observation received at time t
  - A_t: action taken after Zt is observed
- After A_t, compute new belief state X_{t+1}
- Probabilistic, because uncertainty in both X_t and Z_t.

Some Perception Problems

- Localization: where is the robot, where are other things in the environment
  - Landmarks
  - Range scans
- Mapping: no map given, robot must determine both environment and position.
  - SLAM: Simultaneous localization and mapping
- Probabilistic approaches typical
  - Especially machine learning!
- What about common sense? Learning?

Software Architectures

- Low-level, reactive control
  - Bottom-up
  - Sensor results directly trigger actions
- Model-based, deliberative planning
  - Top-down
  - Actions are triggered based on planning around a state model
- Which is an *intelligence* approach?
  - A? B? Neither? Both?

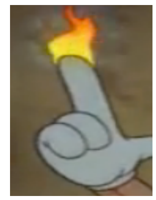

Low-Level, Reactive Control

- Augmented finite state machines
- Sensed inputs and a clock determine next state
- Build bottom up, from individual motions
- Subsumption architecture synchronizes AFSMs, combines values from separate AFSMs.
- Advantages: simple to develop, fast
- Disadvantages: Fragile for bad sensor data, don't support integration of complex data over time.
- Typically used for simple tasks, like following a wall or moving a leg.

Model-Based Deliberative Planning

- Belief State model
  - Current State, Goal State
  - Any of planning techniques
  - Typically use probabilistic methods
- Pros:
  - Can handle uncertain measurements and complex integrations
  - Can be responsive to change or problems.
- Cons:
  - Slow!
  - Developing models for, e.g., driving, is cumbersome.
- Typically used for high-level actions
  - Whether to move and in which direction.

Hybrid Architectures

- Usually, actually doing anything requires both reactive and deliberative processing.
- Typical architecture is three-layer:
  - Reactive Layer: low-level control, tight sensor-action loop, decision cycle of milliseconds
  - Deliberative layer: global solutions to complex tasks, model-based planning, decision cycle of minutes
  - Executive layer: glue. Accepts directions from deliberative layer, sequences actions for reactive layer, decision cycle of a second

Performance Metrics

- Speed and acceleration
- Resolution (in space)
- Working volume
- Accuracy
- Cost
- …plus all the evaluation functions for any AI system.

Where Are Robots *Now?*

- Healthcare and personal care
  - surgical aids, intelligent walkers, eldercare
- Personal services
  - Roomba!
  - Information kiosks, lawn mowers, golf caddies, museum guides
- Entertainment
  - sports (robotic soccer)
- Human augmentation
  - walking machines, exoskeletons, robotic hands, etc.

And More…

- Industry and Agriculture
  - assembly, welding, painting, harvesting, mining, pick-and-place, packaging, inspection, …
- Transportation
  - Autonomous helicopters, pilot assistance, materials movement
- Cars (DARPA Grand Challenge, Urban Challenge)
  - Antilock brakes, lane following, collision detection
- Exploration and Hazardous environments
  - Mars rovers, search and rescue, underwater and mine exploration, mine detection
- Military
  - Reconnaissance, sentry, S&R, combat, EOD
- Household
  - Cleaning, mopping, ironing, tending bar, entertainment, telepresence/surveillance

Tomorrow's Problems

- Mechanisms
  - Morphology: What should robots look like?
  - Novel actuators/sensors
- Estimation and Learning
  - Reinforcement Learning
  - Graphical Models
  - Learning by Demonstration
- Manipulation (grasping)
  - What does the far side of an object look like? How heavy is it? How hard should it be gripped? How can it rotate? Regrasping?

And more…

- Medical robotics
  - Autonomous surgery
  - Eldercare
- Biological Robots
  - Biomimetic robots
  - Neurobotics
- Navigation
  - Collision avoidance
  - SLAM/Exploration

Self-X Robots

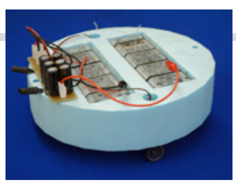


- Self-feeding
  - Literally
  - Electrically
- Self-replicating
- Self-repairing
- Self-assembly
- Self-organization
- Self-reconfiguration


Human-Robot Interaction

- Social robots
  - In care contexts
  - In home contexts
  - In industrial contexts
- Comprehension
  - Natural language
  - Grounded knowledge acquisition
  - Roomba: “Uh-oh”
- Basic idea: Human-centric environments

Why?

- Robots are getting smaller, cheaper, and more ubiquitous
- Humans need to interact with and instruct them, naturally
  - Language, gesture, demonstration, …
- Key requirements:
  - Language understanding learned from data
  - Follow instructions in a previously unseen world
  - Learn to parse natural language into robot-usable commands


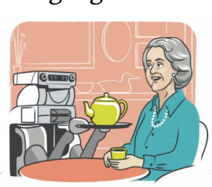


Robots in Human Spaces

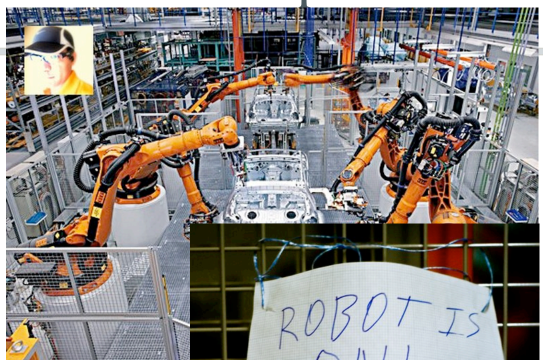

- Robots now:
  - Expensive
  - Complex
  - Special-purpose
- Environments
  - Dedicated
  - Constrained
- Use and Management
  - Controlled by trained experts
  - Slow and expensive to reconfigure/repurpose

Some current problems

HRI
World Learning
Ethical Questions

Human-Robot Interaction

- How do *humans* handle human interaction?
  - Assumptions about retention and understanding
  - Anthropomorphization
- How do robots make it easier?
  - Apologize *vs.* back off
  - Convey intent
  - Cultural context (implicit vs. explicit communication)

Use Cases: Games

Direction Following

- **G[...] Acquisition:**
  - ...sforn... ...e into
  - ...repr...
  - ...ation to
- **Learn a parser**
  - Produce robot-executable commands from NL instructions

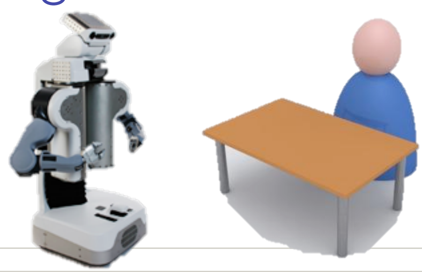

Novel Concepts

- **Grounded La[...]ition**
  - "Understan[...]ing na[...] semantically[...]sentati[...]
  - Mapping that i[...] ...erceive...
- BUT, this assumes we already know what things exist to map to!
- World modeling: learn **new concepts** from interactions

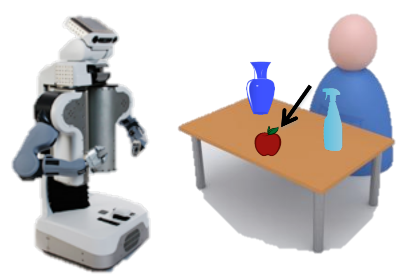

Learning is required

- Robotic systems see new physical things
- Jointly model perceptions and language to create a new, consistent world model
- Learn previously unknown attributes from descriptions
  - Yellow: new word describing new idea

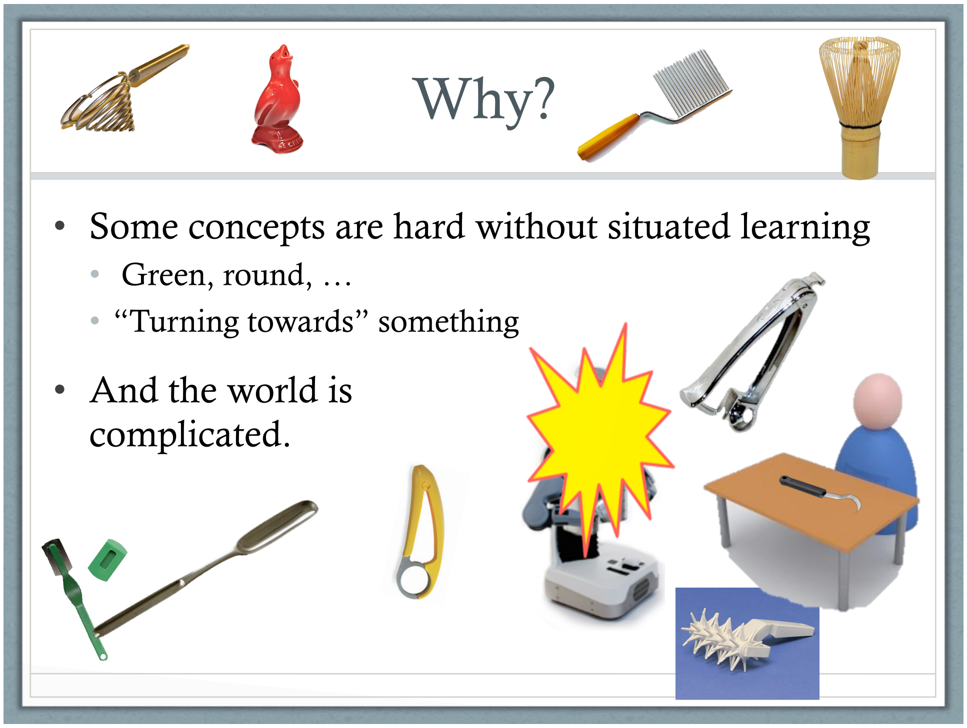
Why?
Some concepts are hard without situated learning
Green, round, …
“Turning towards” something
And the world is complicated.

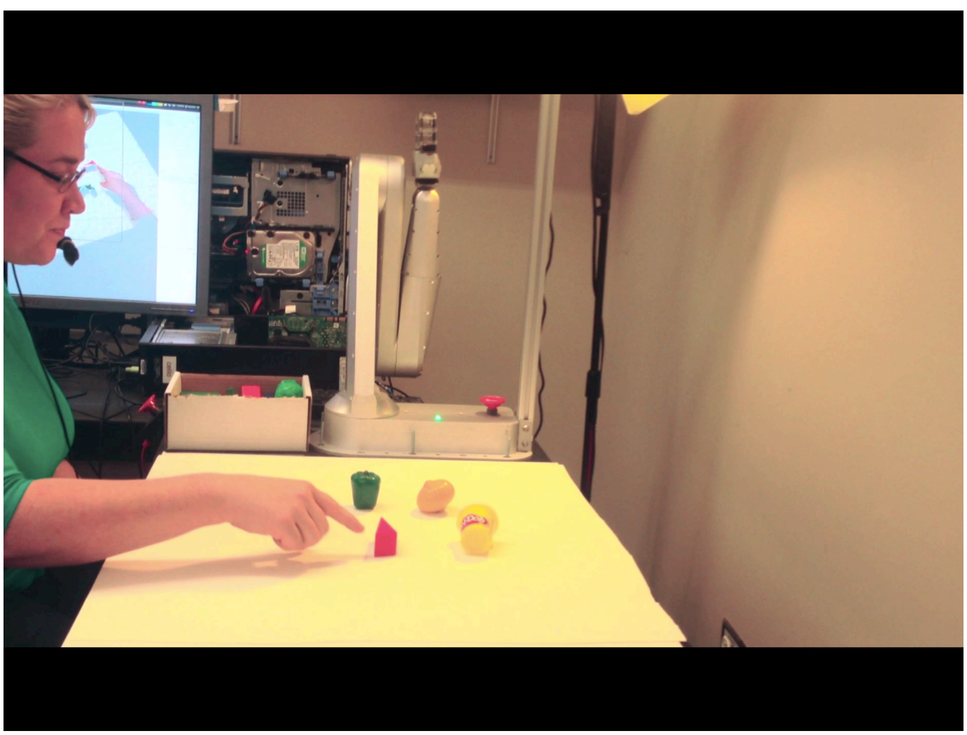

What is the Parent Saying?

Watch the video, then **describe what the parent is saying to the child**, in complete sentences.

- Pretend you are a parent teaching a child about something.
- The question is:

How does the parent describe this group of objects?

Your answer should be the sentence(s) the parent said while pointing to these things.

Submit

“This one's an orange ball.”

Showing HIT 1 of 3 Next HIT

$$\lambda x . \text{orange}(x) \wedge \text{spheroid}(x)$$

Multimodal Interactions

- Larger data set of interactions
- Capturing:
  - Speech
  - Gesture
  - RGB-D
- How do data sources combine?
- Can we model
  - …world?
  - …language?
  - …user intention?

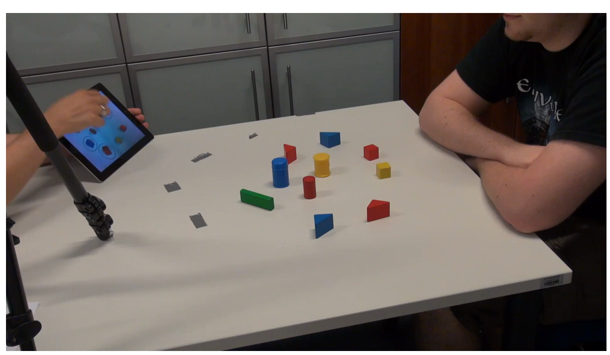

Multimodal Human Input

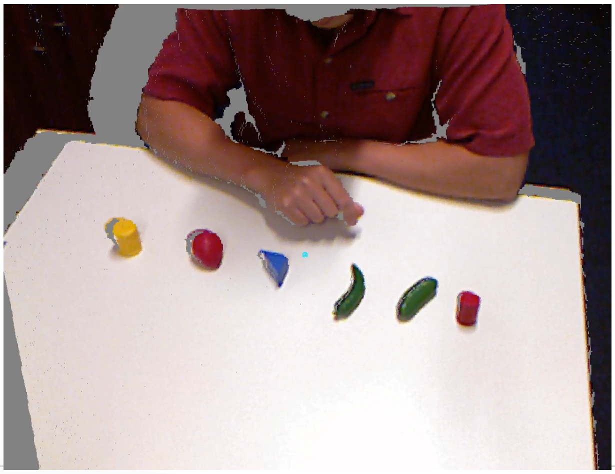

"These are green objects seeming like vegetables. This one is a … a cucumber … or a dull oval thing. And this one is a pepper. Like slightly rounded … high cone."

What Should They Do?

- Boring and/or repetitive
  - welding car frames
  - part pick and place
  - manufacturing parts
- High precision / speed
  - electronics testing
  - surgery
  - precision machining
- Dangerous
  - chemical spill cleanup
  - disarming bombs
- Inaccessible
  - space exploration
  - disaster cleanup
- All of the Above
  - Continuous reef monitoring
  - Military surveillance

What Shouldn't They Do?

- What decisions can be made without human supervision?
- May machine-intelligent systems make mistakes (like humans can)?
- May intelligent systems gamble when uncertain (as humans do)?
- Can (or should) intelligent systems exhibit personality?
- Can (or should) intelligent systems express emotion?
- How much information should the machine give the human?

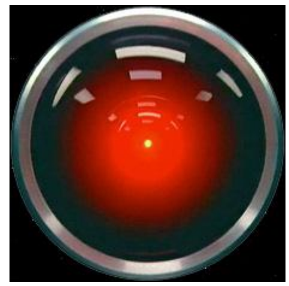

HAL - 2001 Space Odyssey

Jobs For Robots

- Eldercare
- Law enforcement
- Politics
- Space exploration
- Underwater exploration
- Monitoring
- Military surveillance
- Military monitoring
- Domestic surveillance
- Unsupervised surgery
- Unsupervised driving
- Child care

The Future

- Robots that can learn.
- Robots that interact smoothly with people.
- Robots that do ticklish things autonomously.
- Robots that make other robots.
- Robots with “strong” AI.

..?