Telluride final thoughts

A metacognitive robot that occasionally recognizes things as staplers

...("stapler")

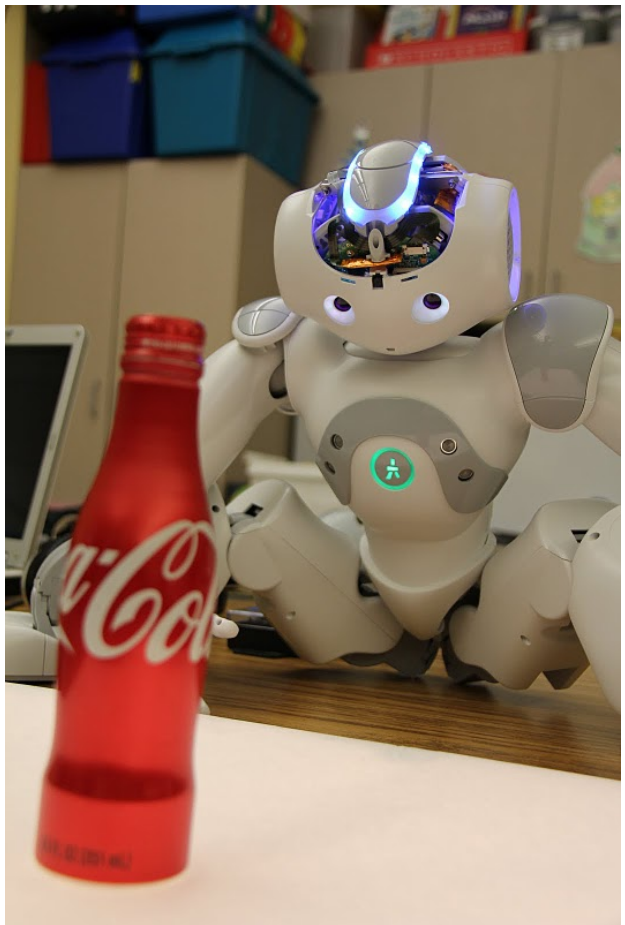

- Leabra Vision model
- Train on 1000 objs from 100 cats
- Generalization from virtual world to real world
- Cross-modal. Vision and speech recognition

Results

- Vision
  - Successfully recognized ~10 real world objects if we very carefully configured his view
  - Otherwise, mostly scissors and staplers.
  - Voting algo didn't really help - mostly revealed his massive confusion
- Audition
  - Step 1: Train on 5 Telluride speakers generalize to 1: error = 40%
  - Step 2: Train on 10 Telluride speakers generalize to 1: error = 30%

Lessons learned

- Real robots are good for "surface validity" demos - funding!
- Virtual worlds++
  - Focus on learning about the brain
  - Relatively low development effort. ODE / OpenGL
- Generalization to real world - ultimate goal:
  - Train on affine shear transforms
  - Train on shadows
  - Train on ambient lighting of surface the object is on
  - Train on textures (ouch!) Requires figure-ground.
  - Model makes really dumb mistakes - need metacognition

Credit assignment

- **Janelle Szary -** Crafting precise environments for our "generalizing" robot, Webots
- **Adam McLeod** - Image processing code / engineering geek goto guy. Auditory brainstorming.
- **Samantha Adams -** NaoSim, Webots, Nao motor control
- **Daniel Fasnacht** - Made us a really long ethernet cable
- **Dave Noelle** - Overall vision. Moral support.
- **Chris Kello** - Crisis comfort pizza. Sound eXchange expert. FFT / Lyon's Cochlear Model
- **Malcolm Slaney** - Auditory brainstorming - get rid of the pitch!
- **Randy O'Reilly -** Nao robot $$$. Opaque pointers.
- **Katie Zhuang -** Webots virtual world
- **Emmit McQuinn -** Helped grok complicated C++ code
- **Kailash Patil -** Auditory saliency reps
- **Roi Kliper -** Comparing LVis to LIBLINEAR on TI20
- **Brian Mingus -** Systems integration