Bootstrap Learning of Real-World Semantics

Benjamin Kuipers
Computer Science & Engineering
University of Michigan

Two Problems in AI: Distinct, but Overlapping

- Problem 1:
  - What must a robot know, to be useful?
  - How can we harvest human knowledge, and use it for robots?

- Problem 2:
  - What is *mind*?
  - How can a physical object have a mind?
  - What knowledge is foundational to a mind?

An early triumph for GOFAI!

- Terry Winograd's SHRDLU (1970)
  - perception
    - avoided, not solved
  - understanding
    - syntax, semantics, pragmatics, reference
  - planning
    - autonomous
  - acting in the world
    - simulated, but real
- A huge step forward!

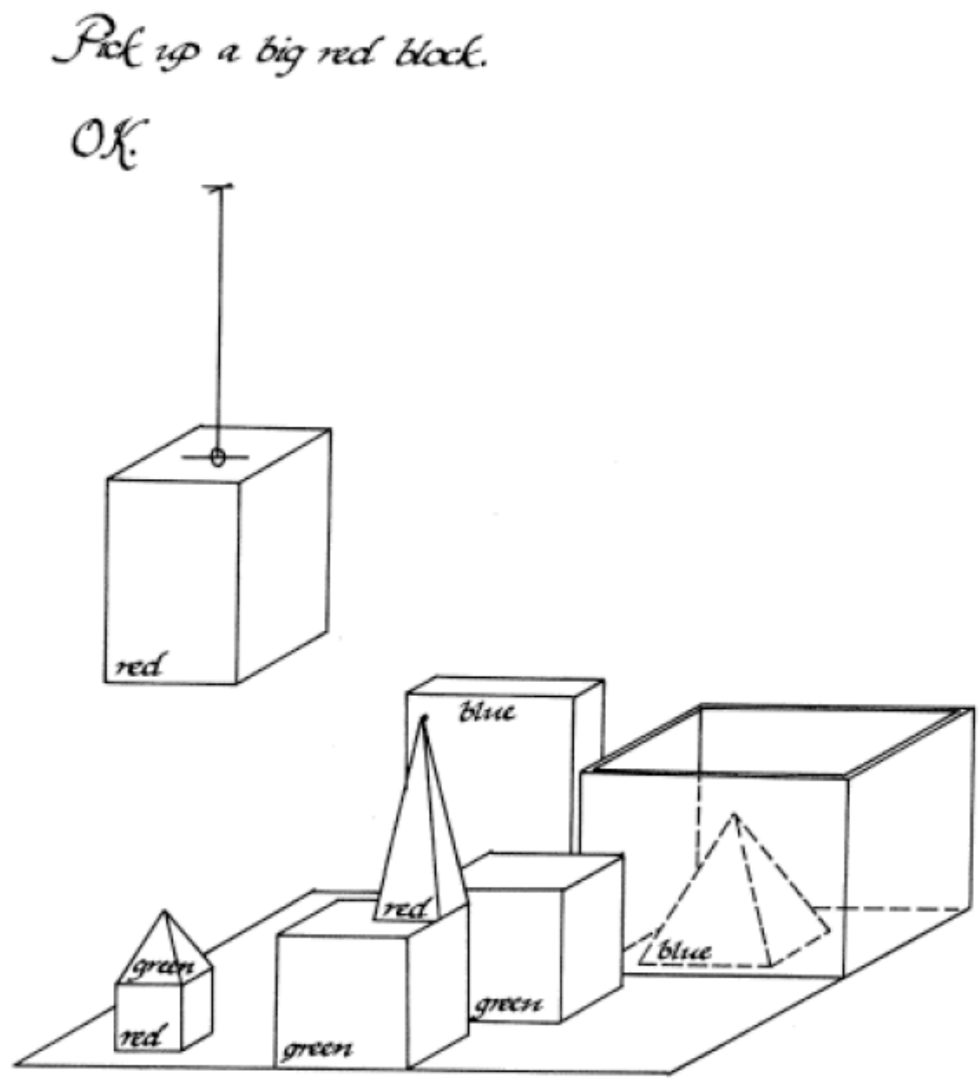

A Lack of Foundational Knowledge

- **But** . . .
  - What is a *block*?
  - or a *pyramid* or a *box*?
    - **Objects**
  - What is *Put* … *on* … ?
  - or *Pick* … *up* … ?
    - **Actions**
  - Where is *on*? or *in*?
    - **Places**
  - What are *red*? *green*? *blue*? *big*?
    - **Properties**

- These must be **learned**!

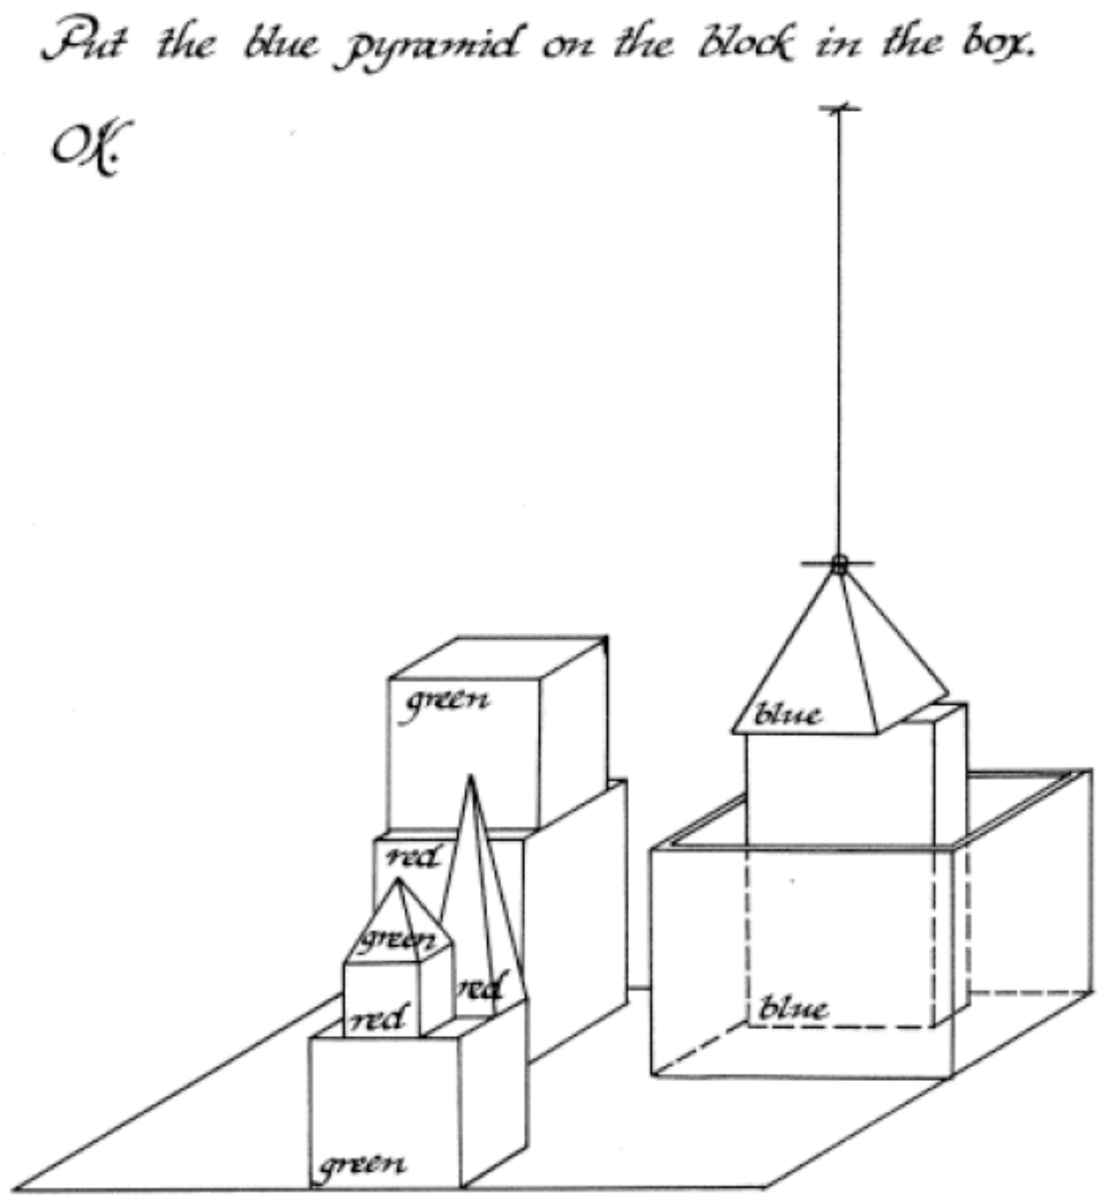

Foundational Knowledge Must Be Learned

- Symbolic representation and inference (GOFAI) are not enough.
  - We need probability theory, machine learning, statistics, control theory, dynamical systems, and much else.

- The foundations of knowledge must be *learned*, by an embodied agent, embedded in its world.
  - We are unlikely to be able to program by hand adequate representations of the complex world.

John Searle's Objection to AI

- The essence of his "Chinese Room" argument:
  - An intelligent agent's knowledge has *meaning*: reference to objects in the external world.
  - Computation, as such, is syntactic (*meaning-free*) manipulation of symbol structures.
    - Symbol structures can only refer to other symbols.
  - Thus, the mind *cannot* be explained by computation.
- This argument has genuine weight.
  - Enough to refute the possibility of AI? ***No***.
    - But it does say something about the nature of mind.
  - *Where does the meaning come from?*

The Constructivist Agent Reply

- OK! Knowledge (in robots *or in humans*) cannot *refer* to objects in the external, physical world.
  - An agent *senses* the world, and *acts* within it.
  - It has no direct access to things in the physical world!

  - It constructs *its own* internal knowledge structures to track and explain those patterns of interaction.
  - *Reference* is to constructed knowledge structures.

  - These internal structures must *correspond* to the world reasonably well.
  - If not, the agent could not survive.
- Our task is to show how this can actually work.

Learning Foundational Knowledge Like a Baby

- *The baby, assailed by eyes, ears, nose, skin, and entrails at once, feels it all as one great blooming, buzzing confusion . . .*
    - [William James, 1890]

- We've learned a lot about development since then.
  - Some knowledge is learned over evolutionary time.
  - We're going to pretend it's learned by the individual.
  - Our focus is on the growing richness of the **ontology**.
    - Ontology = What can be represented.

A Bootstrap Learning Agenda

- Learning **body space**:
  - Structure of sensory input and motor output;
  - Control laws to achieve and maintain feature values;
- Learning the **local spatial model**:
  - Structure of local space; places and paths;
  - Describe static local space; treat dynamics as noise.
- Learning about dynamics --- **objects** and **actions**:
  - Objects, and the actions that affect them;
  - Relevant object properties, affordances, and tool use;
- Learn about **goals**, **beliefs**, and **plans**:
  - A sequence of actions, in context, achieves a goal;
- Learn about **other agents**:
  - . . . with their own beliefs and goals; communication;

Background

- Our cognitive mapping research provided a key insight.
  - *High-level concepts* (places and paths) *can be abstracted from the behavior of low-level control laws, operating at the pixel level*.

- Spatial Semantic Hierarchy
  - Kuipers & Byun, JRAS, 1991
  - Kuipers, AIJ, 2000
  - Beeson, Modayil & Kuipers, IJRR, 2010

Distinctive States

- A *distinctive state* (location plus orientation) is the isolated fixed-point of a hill-climbing control law.

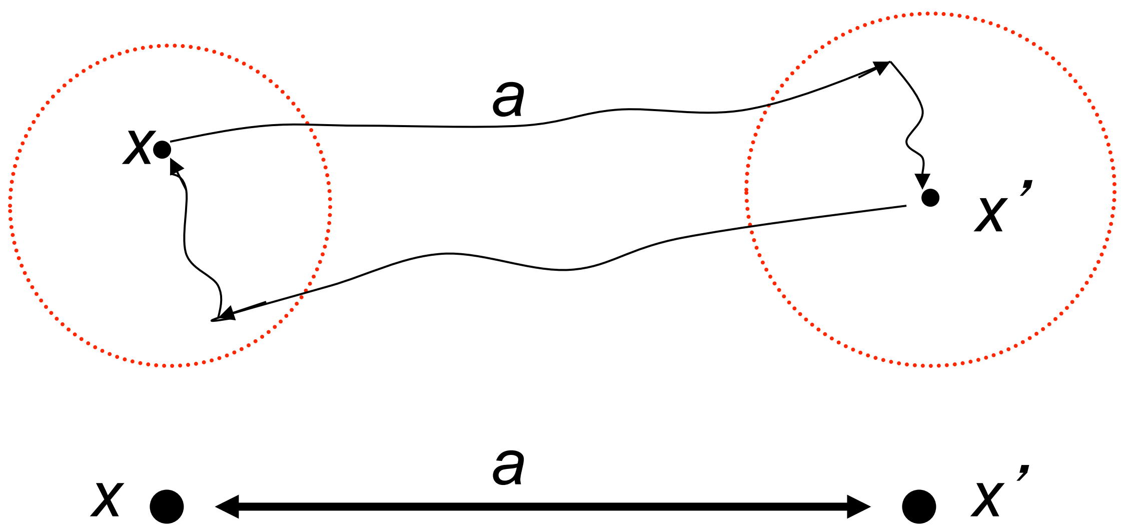

  – *High-level concepts* (places and paths) *can be abstracted from the behavior of low-level control laws, which operate at the pixel level*.

Topological Abstraction

- A control law defines an attractor
  - that represents its basin of attraction

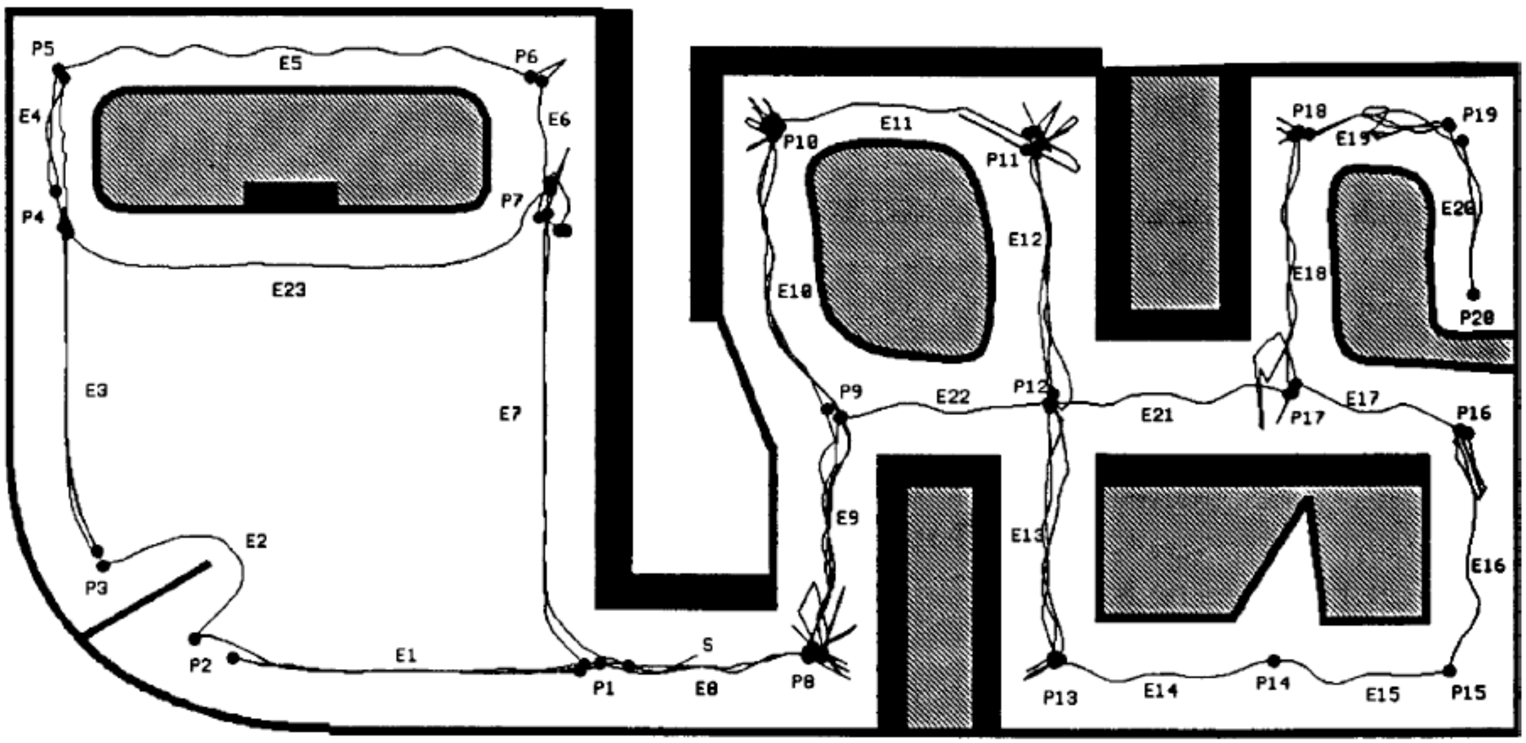

[Kuipers & Byun, JRAS, 1991]

Topological Abstraction

- A small, finite graph represents the structure of the behaviors in a continuous space.

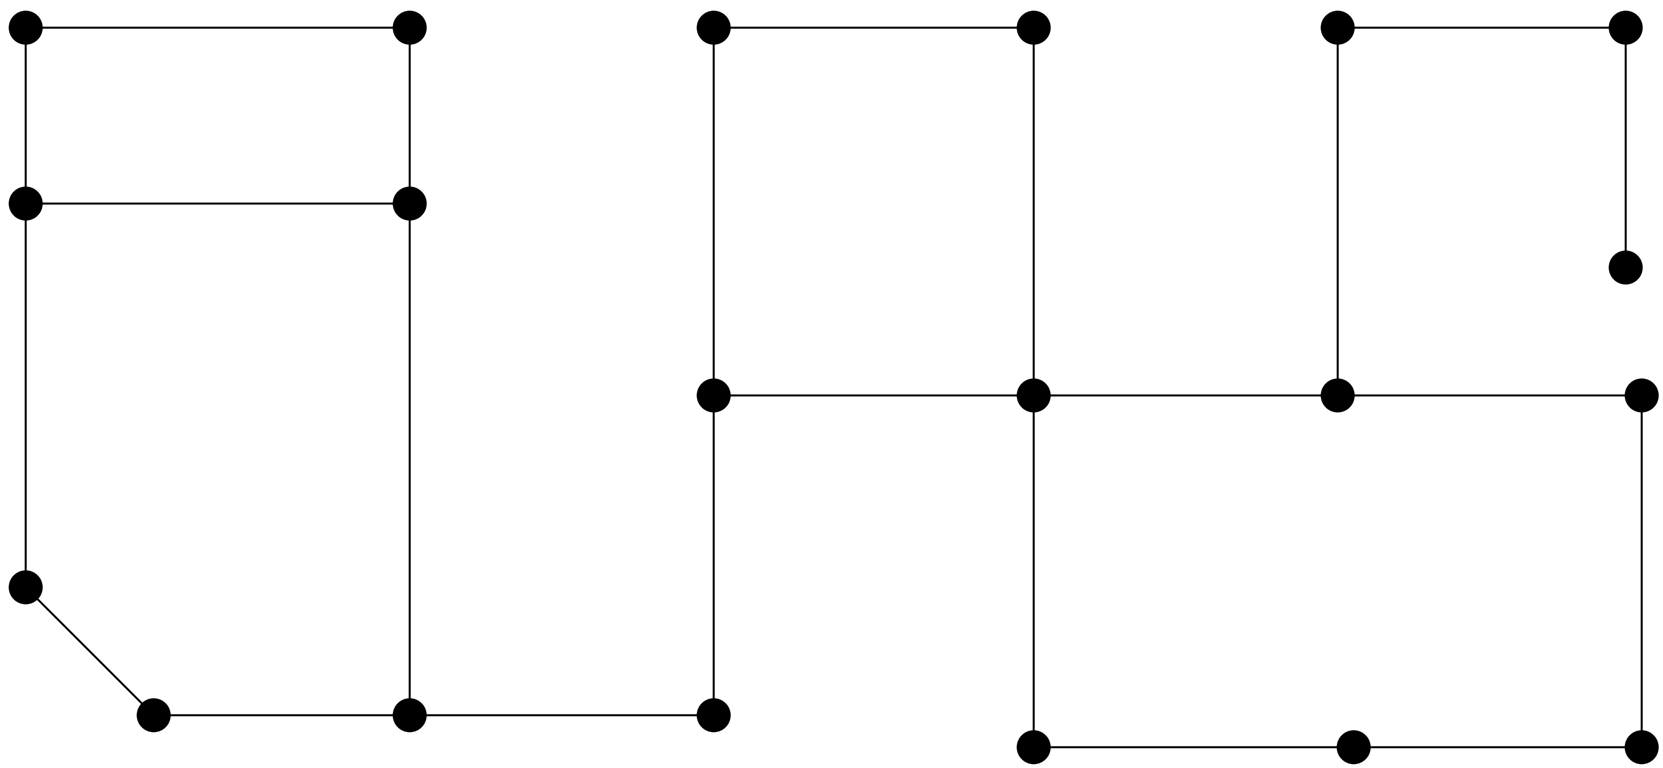

A Bridge Across the Canyon

- Years ago, we built a fragile, rickety bridge across a huge deep canyon.
  - Every part of that bridge can (and should) be improved.
  - But it got to the other side.

- How do we get from pixel-level sensors and effectors, without higher-level semantics,
  - to space, objects, actions, goals, and plans?

Lassie “sees” the world with a laser rangefinder

- 180 ranges over 180° planar field of view
- About 13” above the ground plane
- 10 - 12 scans per second

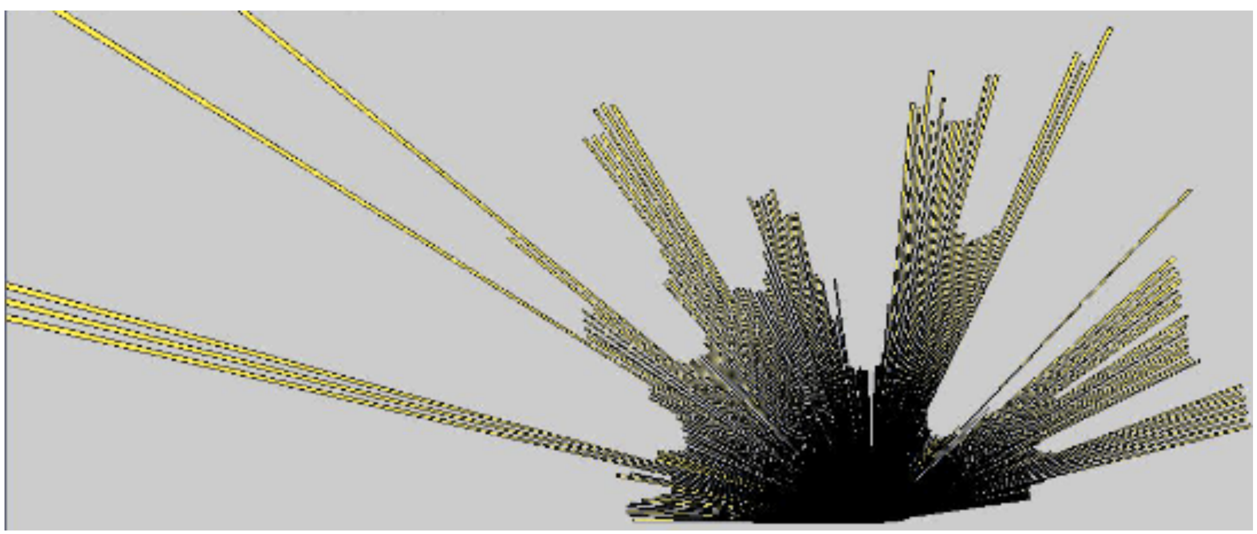

The First Problem

- Given many disorganized pixel-level sensors, how does the agent learn to make sense of them?

- Related problems we won't have time for:
  - Separating distinct sensor modalities
  - Understanding pixel-level effectors

[Pierce & Kuipers, AIJ, 1997]

Disorganized sensor: 180 “pixels”

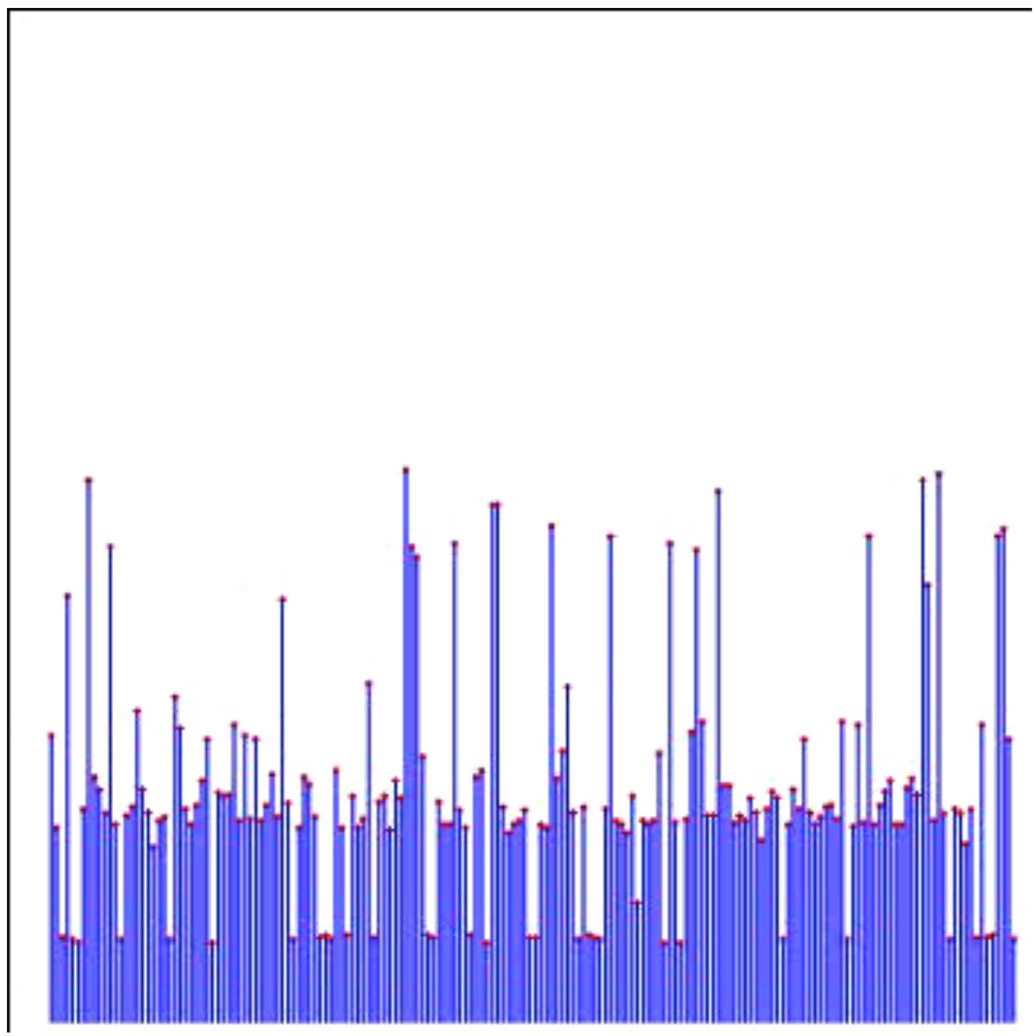

Structured Sensor Array

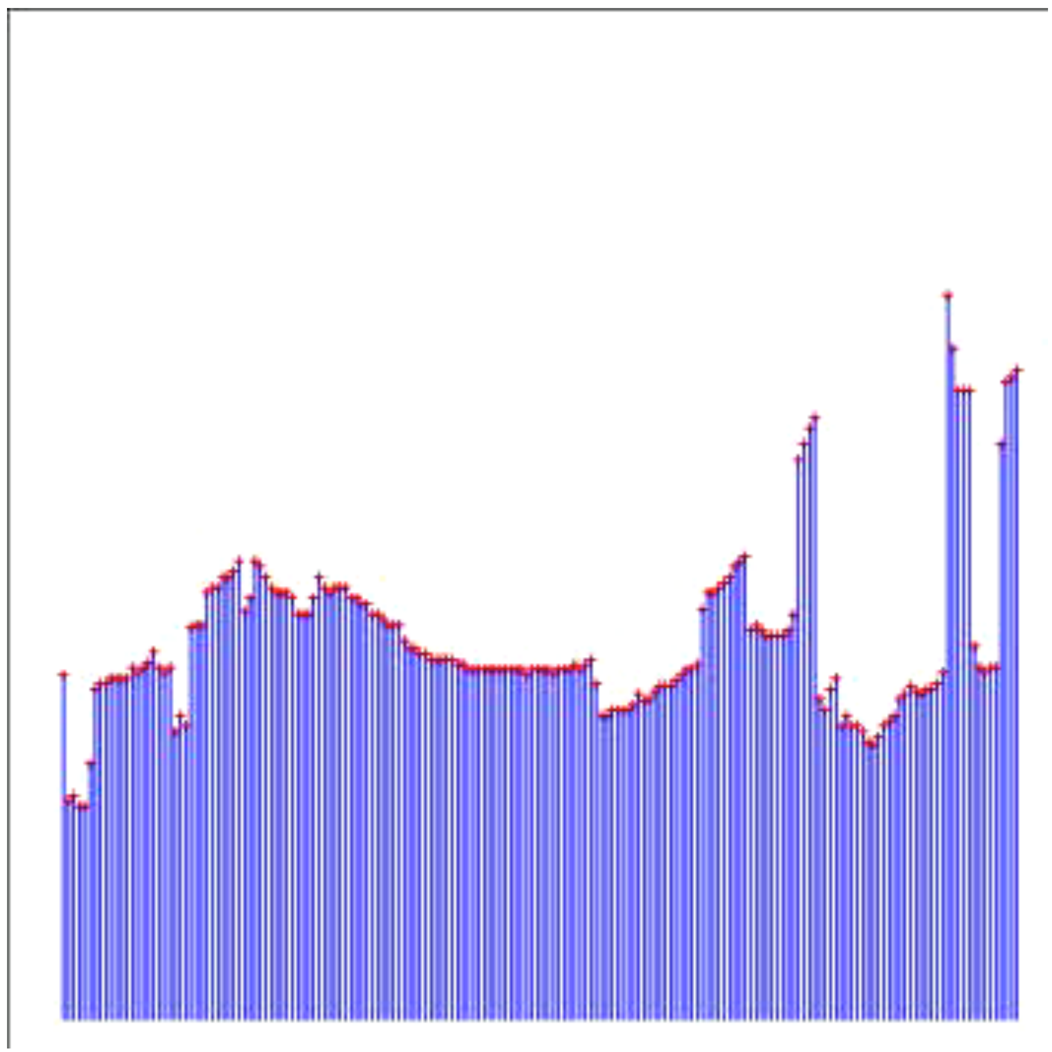

The Egocentric Range Image

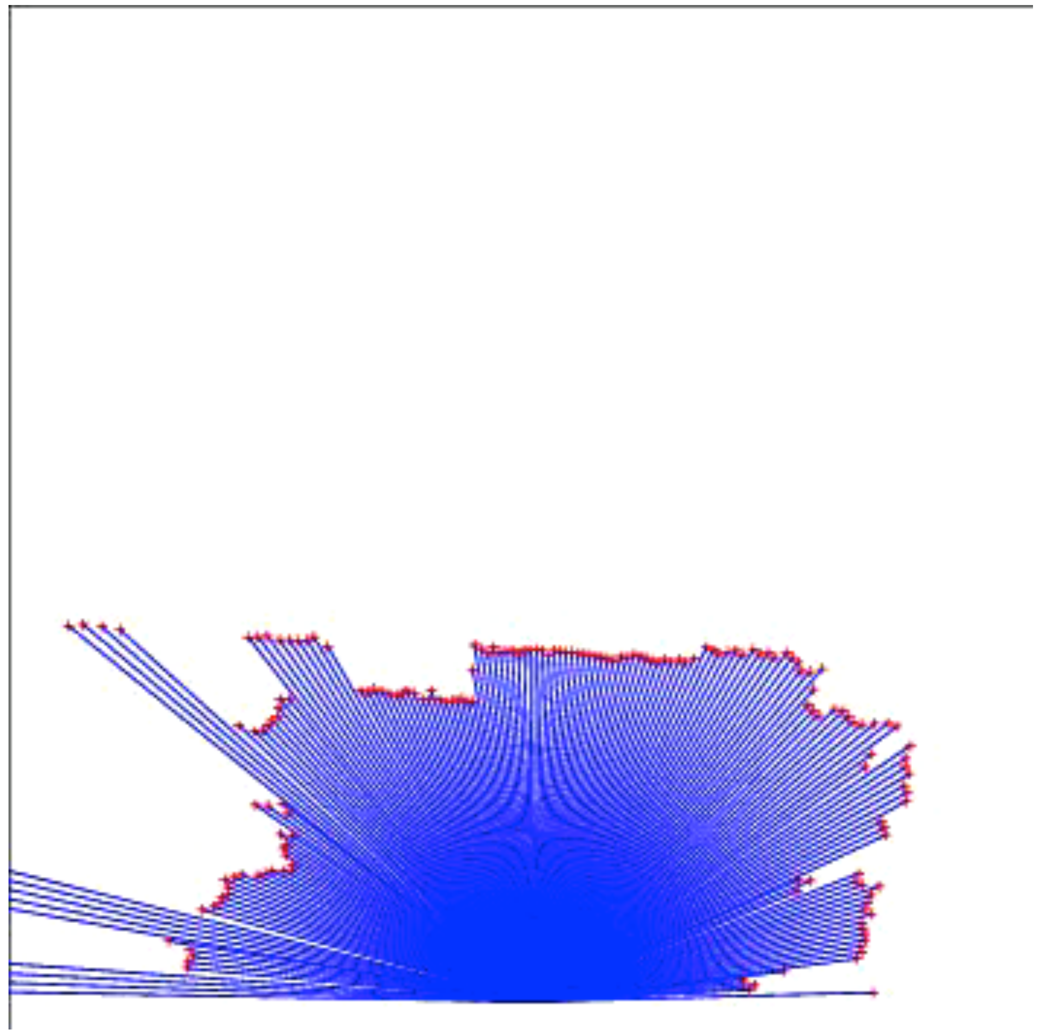

Estimate Sensor Similarities

- Start with a disorganized "bag of pixels"
- Determine pairwise sensor distances

$$d_1(s_i, s_j) = \frac{1}{N} \sum_{t=1}^{N} |s_i(t) - s_j(t)|$$

- This is l^1 distance.
  - Isomap uses l^2 (*Euclidean*) *distance*.
    - [Tenenbaum, et al, 2000]
  - Olsson, et al [2006] uses *mutual information*.
  - Modayil [2010] uses *Gaussian processes* to identify sensor embeddings.

Organization of each sensor array

- Place $s_0 \ldots s_k$ in $\Re^k$ according to $d_1(s_i,s_j)$
- Use PCA to find dominant eigenvectors
  – Project into low-dimensional space ($\Re^2$)
  – Relax sensor positions to best match $d_1(s_i,s_j)$

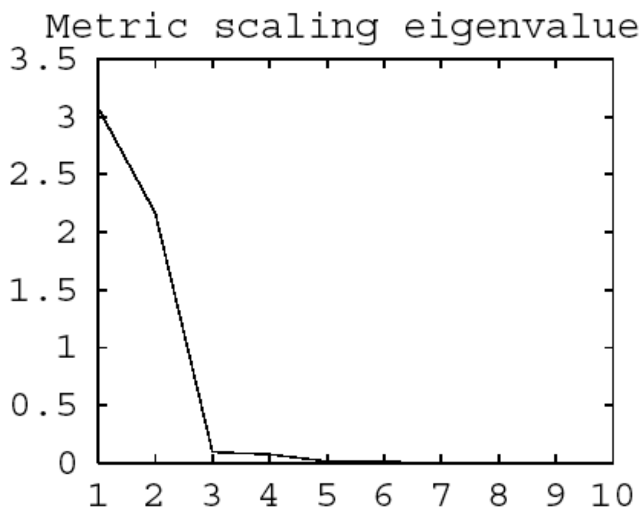

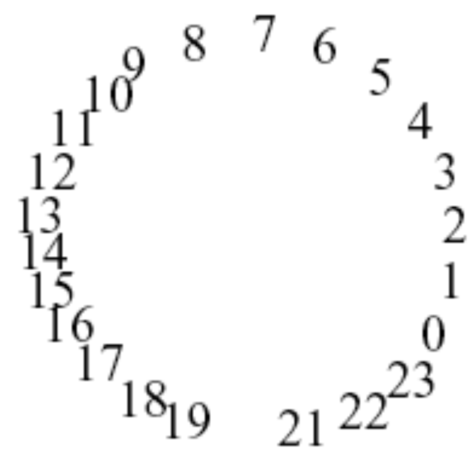

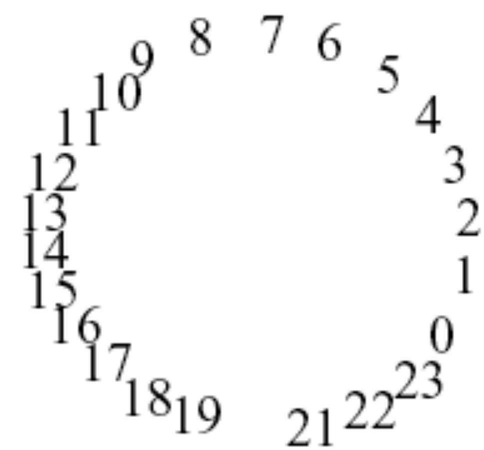

Laser Rangefinder array

- The same method works, applied to real data from the laser rangefinder array (180 rays)

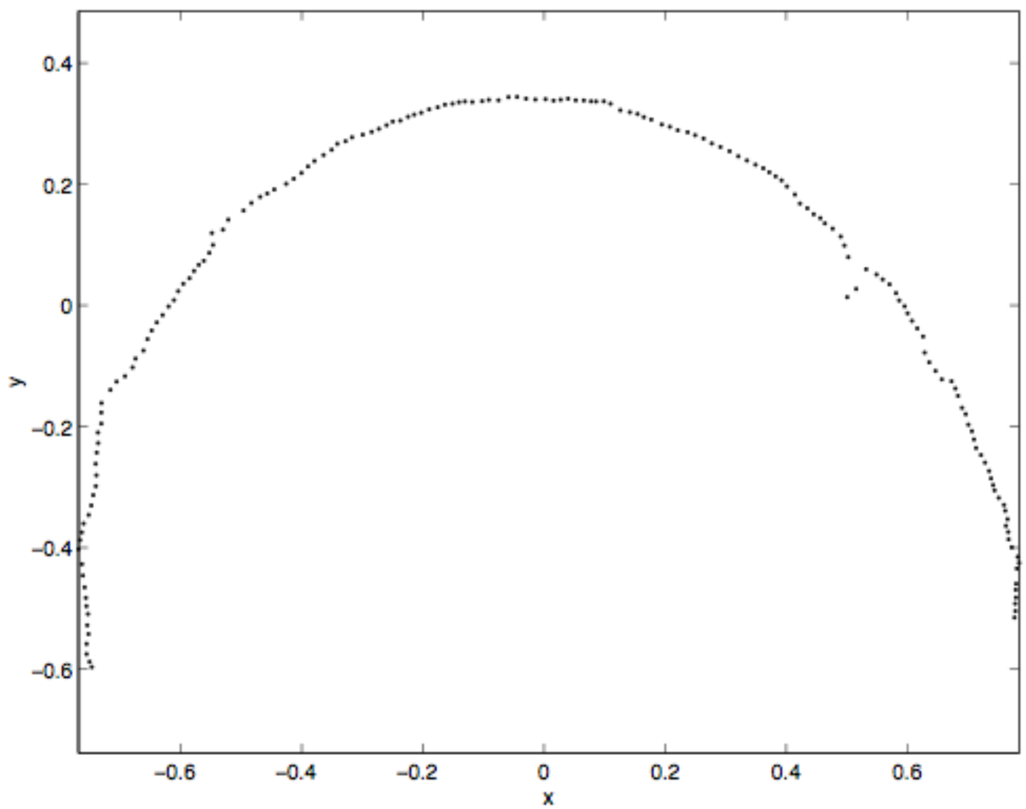

The “Roving Eye”

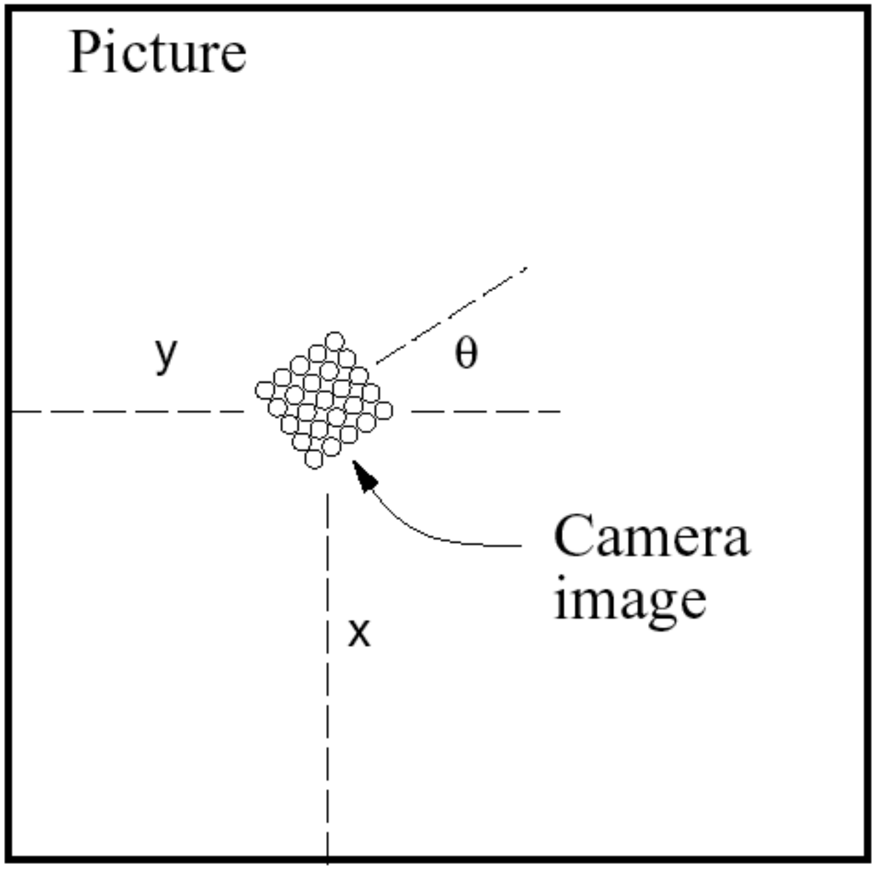

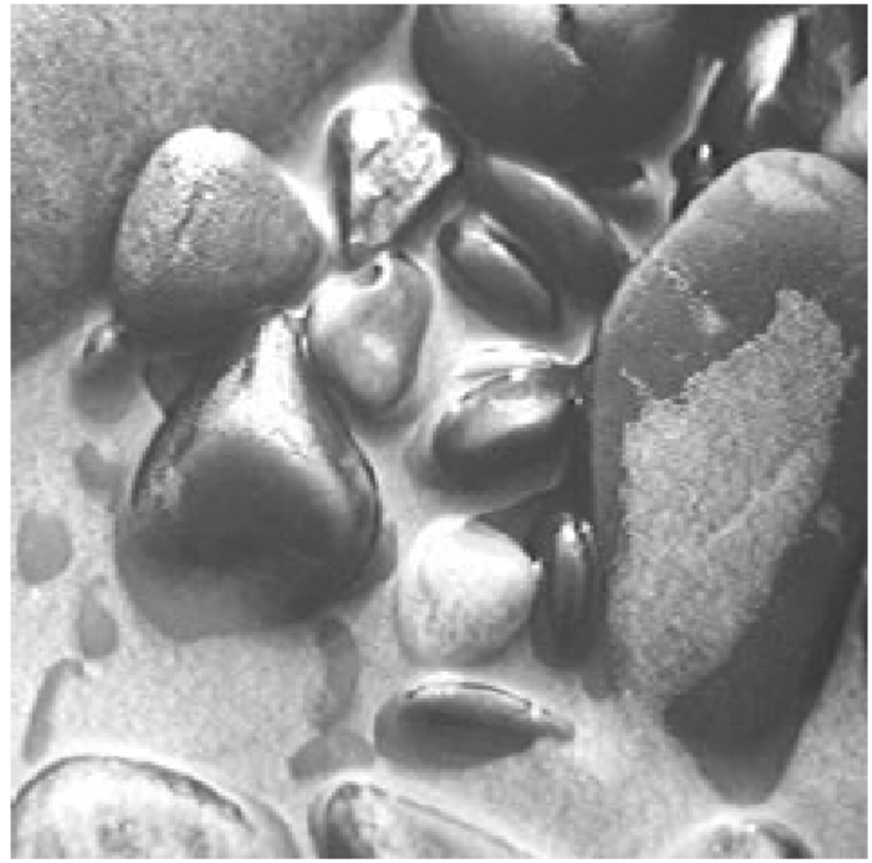

Structure of the "Roving Eye"

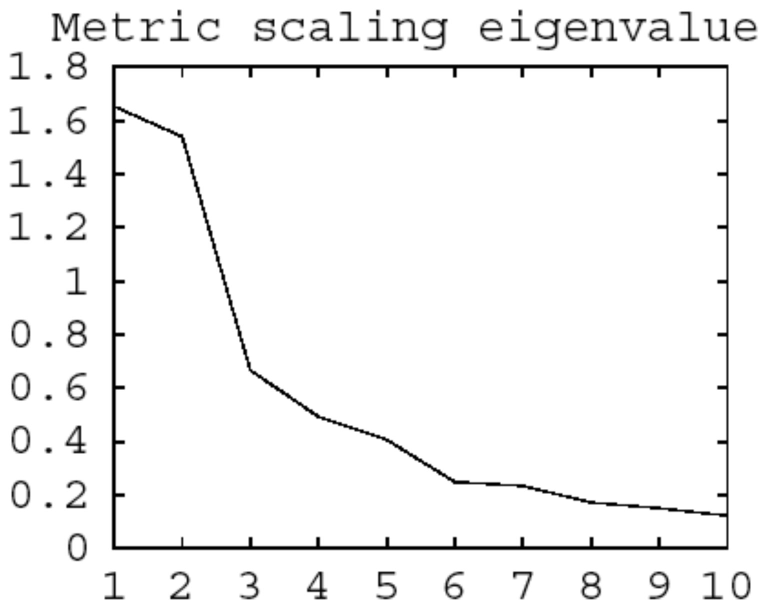

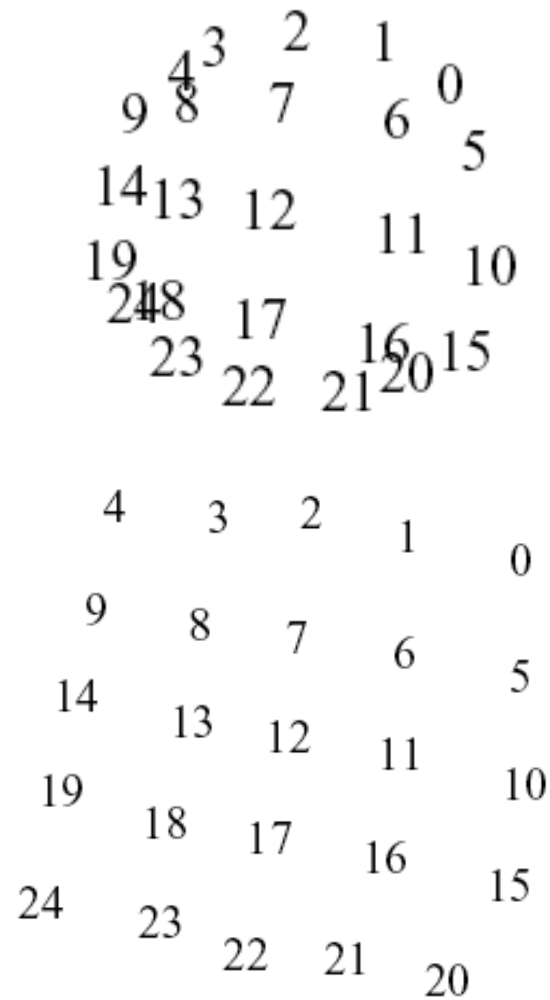

The Egocentric Range Image

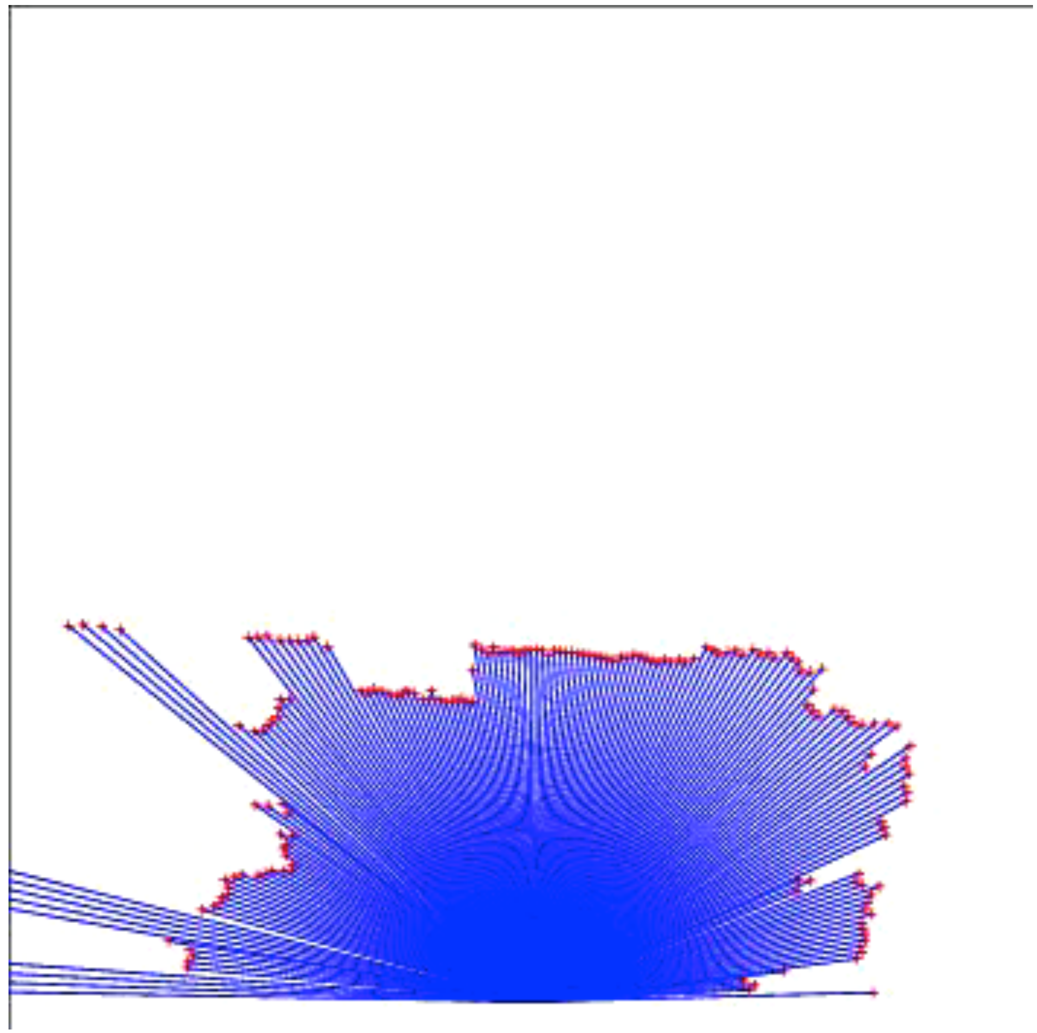

The World-Centered Range Image

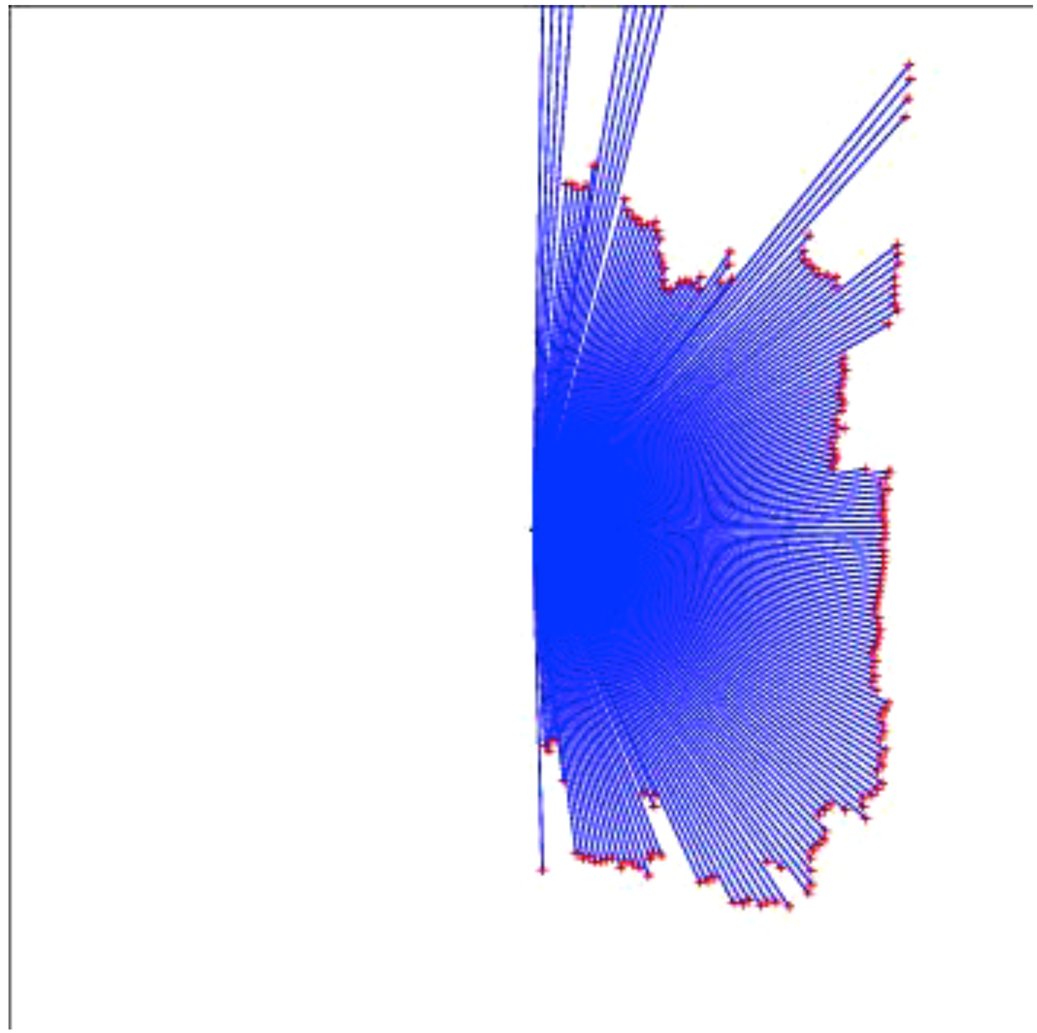

The egocentric origin now has a pose and a trajectory!

The World-Centered Range Image

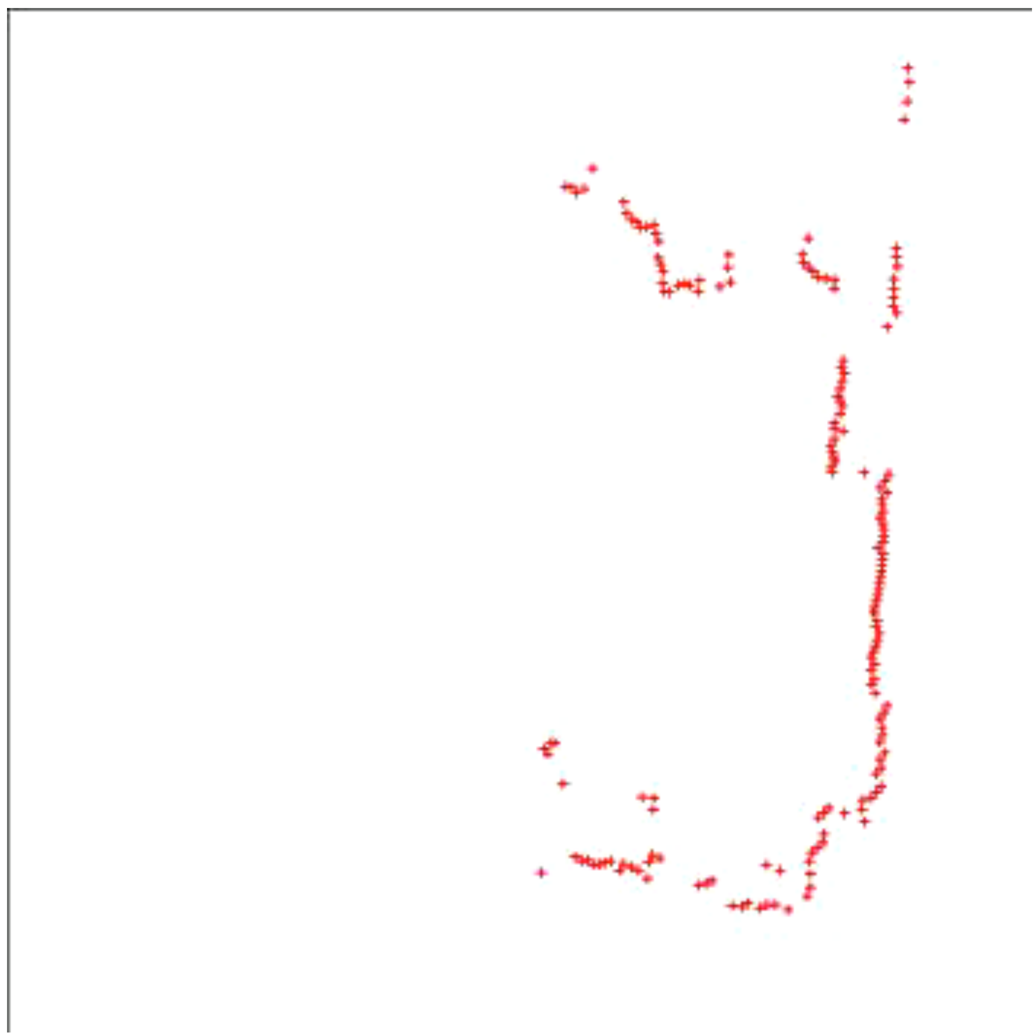

An occupancy grid includes history, making a static-world model

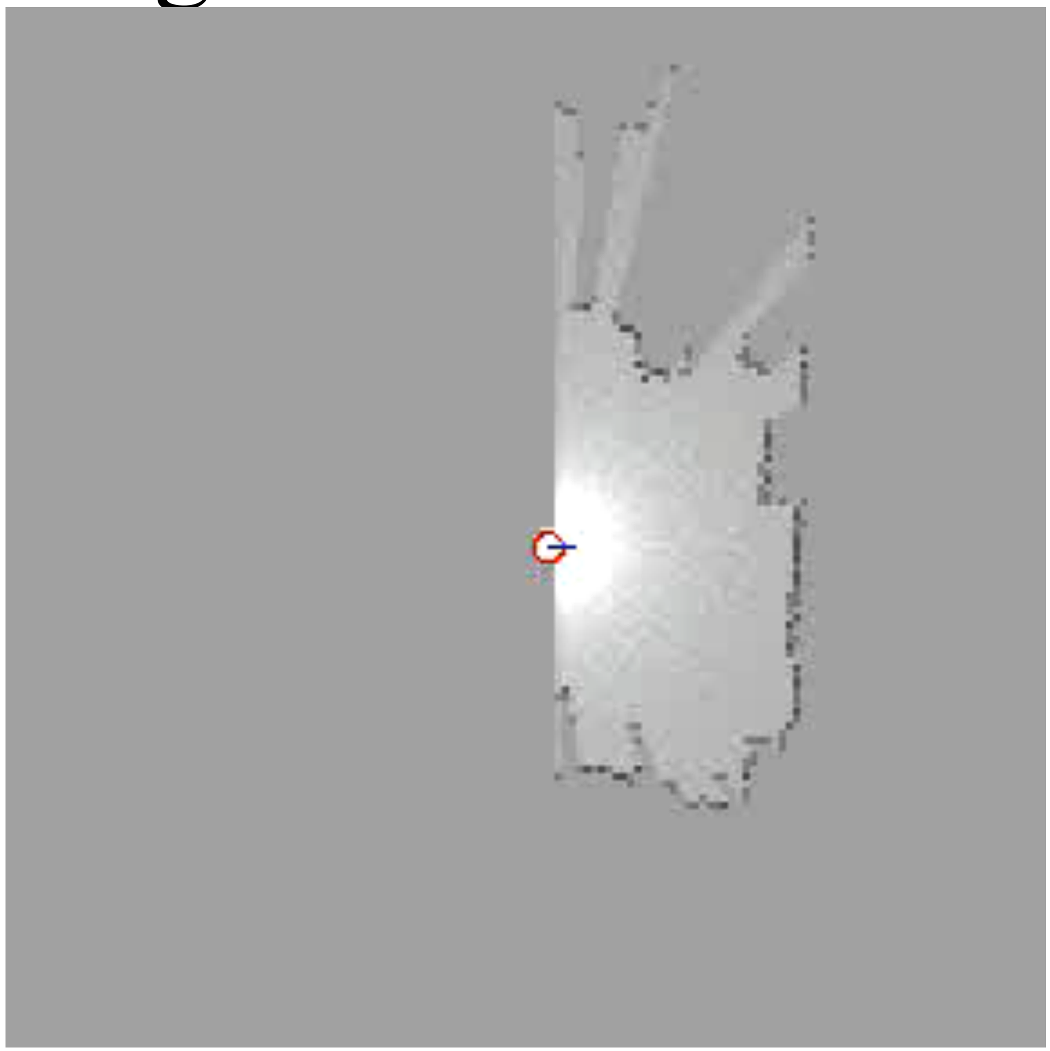

A Static Model of Space

- Given a high-dimensional sensory stream $\mathbf{z}_t$,
  - we have learned a static model $\mathbf{M}$ of the world,
  - and the trajectory $\mathbf{x}_t$ of the egocentric origin.

$$\mathbf{z}_t = G(\mathbf{M}, \mathbf{x}_t) + \epsilon$$

- $G(\mathbf{M}, \mathbf{x}_t)$ predicts and partially explains $\mathbf{z}_t$.
  - Greatly compresses the information in $\mathbf{z}_t$, in terms of smaller structures $\mathbf{M}$ and $\mathbf{x}_t$.
  - The error ε is the difference between prediction and observation.

Statistical Learning Methods Used

- Correlation (time-series and histograms)
- *k*-means and agglomerative clustering
- Multidimensional scaling
- Dimensionality reduction (PCA, Isomap)
- Sensory flow
- Image matching (ICP)
- Markov localization (max likelihood pose)
- . . .

The Second Problem: We need to learn *Objects*

- The static model **M** explains most observations.

$$\mathbf{z}_t = G(\mathbf{M}, \mathbf{x}_t) + \epsilon$$

  - Dynamic objects appear in the *discrepancies*.

- Cluster *noise pixels* in space.
  - Track the clusters over time: interpret as *objects*.
  - Merge images to make shape models

- Modayil & Kuipers [2004, 2006, 2007, 2008].

Identify Discrepancies

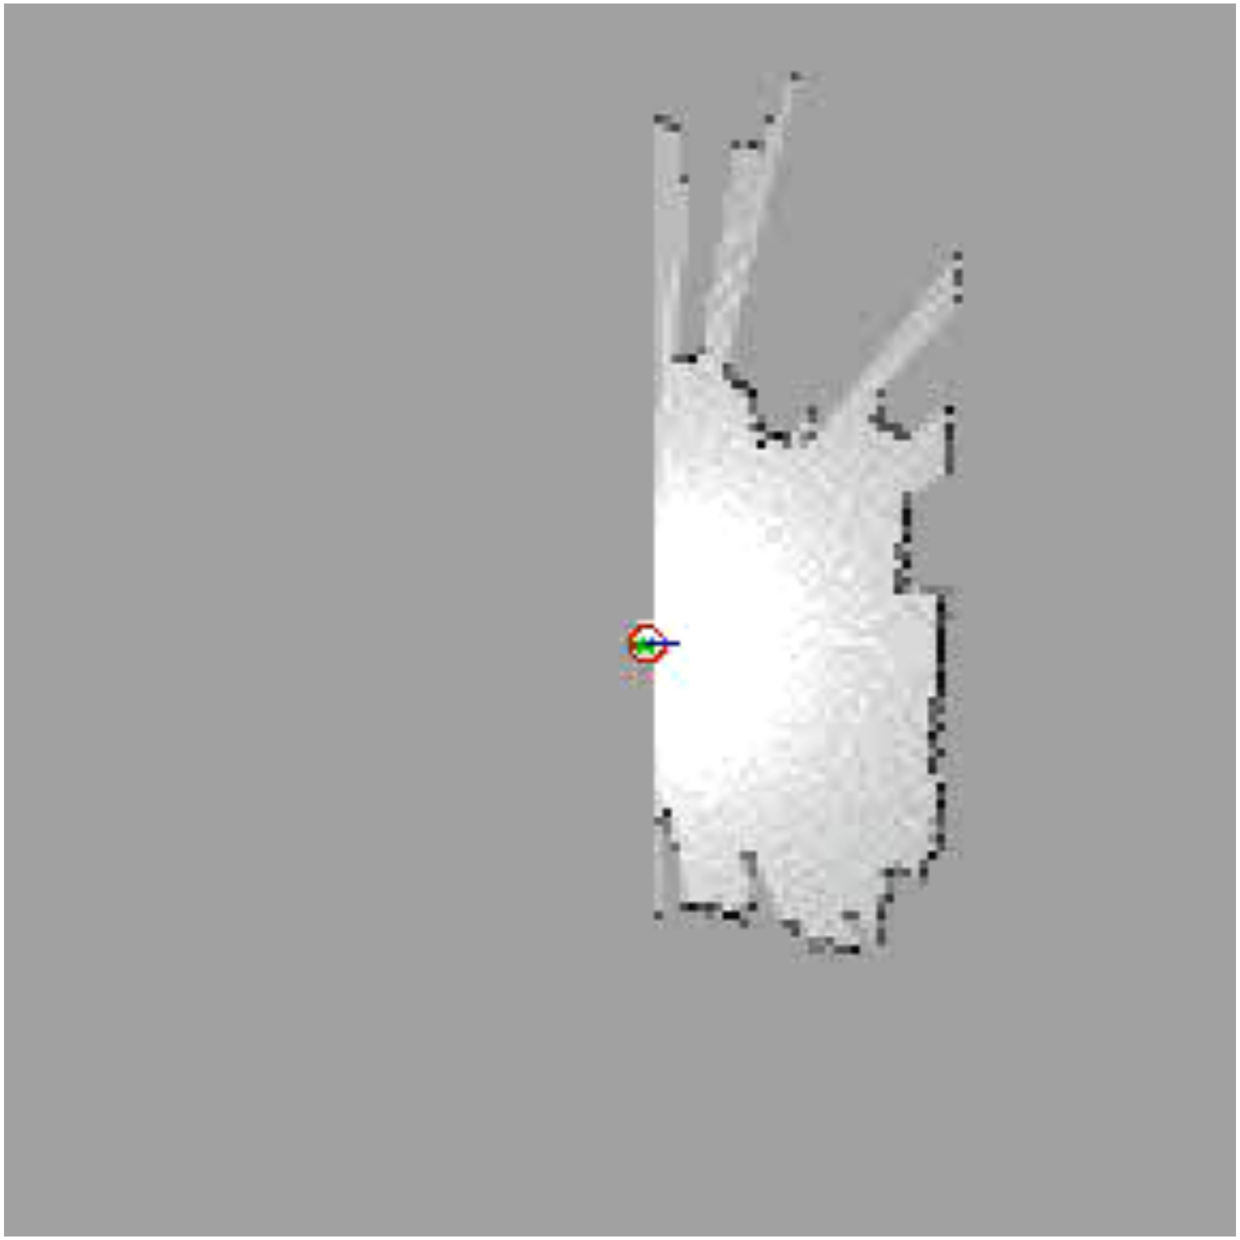

Clustering into Objects

Track Objects over Time

Describe the Scene

- Describe the scene in terms of:
  - Static world
  - Robot's own pose
  - Object in a fixed position
  - Object and trajectory
- *Individual* objects
  - Categories

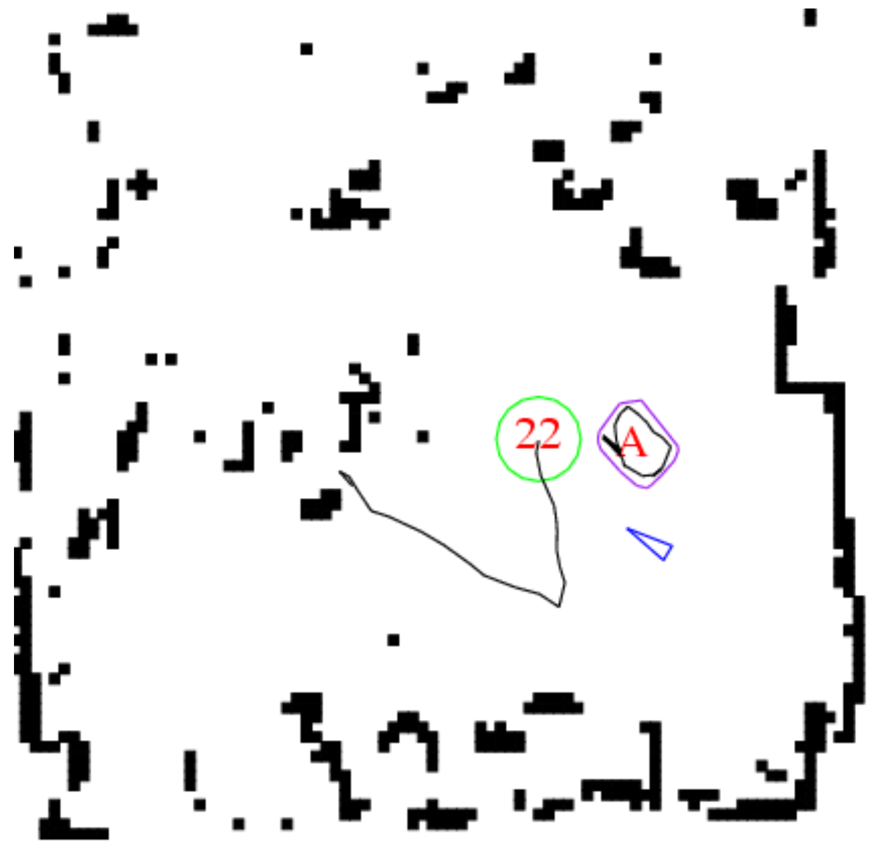

The Third Problem: Learn How *Actions* Affect Objects

- We have learned object attributes:
  - position, orientation, shape
  - in egocentric and world-centered frames.
- Do random movements (“motor babbling”)
- Record all interactions (attribute changes)
- Identify clusters in the data, describing:
  - *Effect*: qualitative change of attribute
  - *Prerequisite*: bounds on previous state
  - *Motor*: signal to perform the action

Learning Actions for Planning

- Find clusters in the data to describe:
  - Action = $\langle Prereq,\ Motor,\ Effect \rangle$
    - Essentially the STRIPS representation for actions.
- For this mobile robot, the learnable actions:
  - **Turn** to face object or desired point.
  - **Move** to desired point in nearby space.
  - **Push** (Move, to get object to move also)
- Demonstrate use of learned objects & actions
  - Given: goal location for specified object
  - Given: simple back-chaining planner

Using the Learned Actions

- Learn action properties. Do a simple plan.

The Constructivist Agent

- The agent constructs *its own* models of itself and its world:
  - Agent's own sensors and motor controls
  - Ego- and world-centered spatial frames
  - Objects, shapes, and localization
  - Actions and their properties
- These are useful *self-created* abstractions, not knowledge of the environment provided by an external authority.
  - Because they correspond well with the actual properties of the world, they guide effective action.

A Bootstrap Learning Agenda

- Learning **body space**:
  - Structure of sensory input and motor output;
  - Control laws to achieve and maintain feature values.
- Learning the **local spatial model**:
  - Structure of static local space; places and paths.
- Learning about **objects** and **actions**:
  - Objects, and actions that affect them;
  - Relevant object properties, affordances, and tool use.
- Learn about **goals**, **beliefs**, and **plans**:
  - A sequence of actions, in context, achieves a goal.
- Learn about **other agents**:
  - . . . with their own beliefs and goals; communication.

Moving from 2D to 3D

- Onward! From mobile robots with lasers to humanoid robots with eyes, arms and hands.

  – Object Semantic Hierarchy: a hierarchy of representations for learning models of 3D objects from streams of visual observations.
    - Changhai Xu & B. Kuipers, OSH

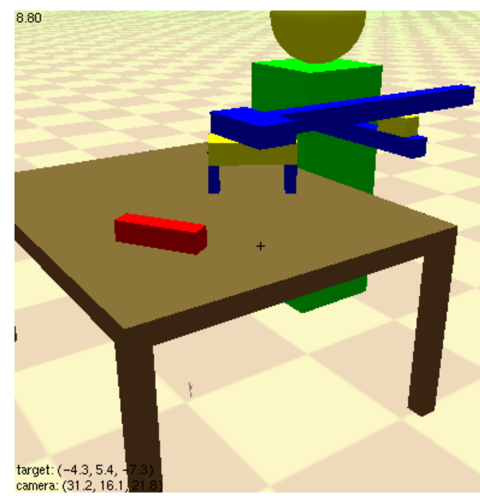

  – A hierarchy of actions: Simultaneously learning improved models of actions, and a qualitative representation with the right distinctions.
    - Jonathan Mugan & B. Kuipers, QLAP

Other Agents Have Goals and Beliefs

- To predict the behavior of an **inanimate object**,
  - If you know the forces on that object, you can
  - Infer the resulting acceleration and velocity.
    - Or from its motion, you can infer the forces.

- To predict the behavior of an **agent**,
  - If you know an agent's goals and beliefs about the world,
  - You can infer the actions it will take.
    - Or from its actions, you can infer its goals and beliefs.

- A developing agent must learn a Theory of Mind.

Living in a Society of Agents

- Learn to understand other agents
  - Theory of Mind supports improved predictions.
- Learn by imitating other agents
  - Observe the behavior of another agent.
  - Infer the state space and reward for an RL problem.
  - Practice to solve the RL problem.
- Learn rules to encourage cooperation with others.
  - Cooperation produces better outcomes for everyone.
  - Discourage or punish “free riders”.

References

- Mugan & Kuipers, Autonomous learning of high-level states and actions in continuous environments. *IEEE Trans Autonomous Mental Development*, 2012.
- Tsai, Xu, Liu & Kuipers, Real-time indoor scene understanding using Bayesian filtering with motion cues. *ICCV*, 2011.
- Xu & Kuipers, Towards the Object Semantic Hierarchy. *IEEE Int. Conf. on Development & Learning (ICDL)*, 2010.
- Kuipers, Drinking from the firehose of experience. *Artificial Intelligence in Medicine*, 2008.
- Modayil & Kuipers, The initial development of object knowledge by a learning robot. *Robotics & Autonomous Systems*, 2008.
- Kuipers, The Spatial Semantic Hierarchy, *AIJ*, 2000.
- Pierce & Kuipers, Map-learning with uninterpreted sensors and effectors, *AIJ*, 1997

A learning agent in action