Intelligent Control
and Cognitive Systems

Design and Learnability

Joanna J. Bryson
University of Bath, United Kingdom

Learning

- Learning requires:
  - A representation.
  - A means of acting on current evidence.
  - A means of incorporating feedback concerning the outcome of the action / guess from evidence.

state / accumulated evidence

prediction

learning algorithm

The "No Free Lunch" Theorems

- No learning algorithm magically dismisses combinatorial complexity, but...
  - Wolpert, D.H., (1996) The lack of *a priori* distinctions between learning algorithms, *Neural Computation*, pp. 1341-1390.
- The representation is part of the bias ∴ some types of learning may converge faster or more reliably than others in a particular problem space.
  - Wolpert, D.H., (1996) The existence of *a priori* distinctions between learning algorithms, *Neural Computation*, pp. 1391-1420.

Evolvability

- One of the things that evolves is the capacity to evolve better.
- However, any bias makes some things easier to achieve and therefore others harder.
- Full-time theoretical biologists still find the "harder" part hard to comprehend.

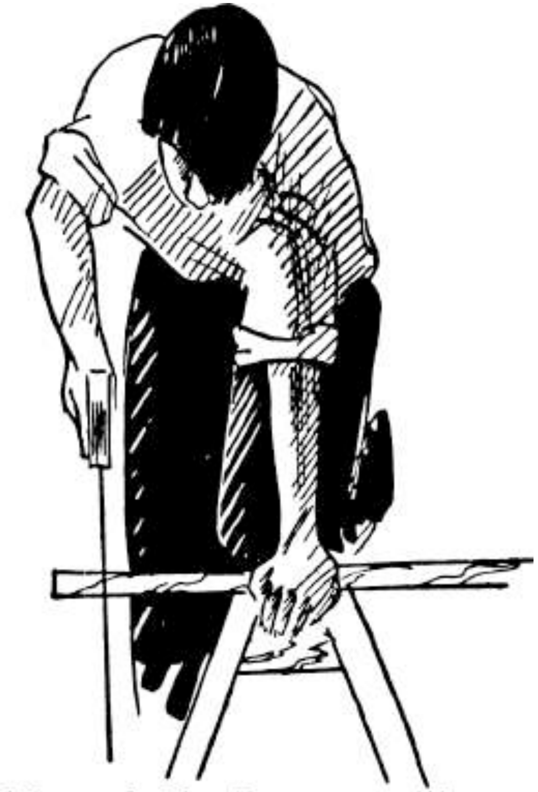

Figure 1-64.—Proper position for sawing a board to size.

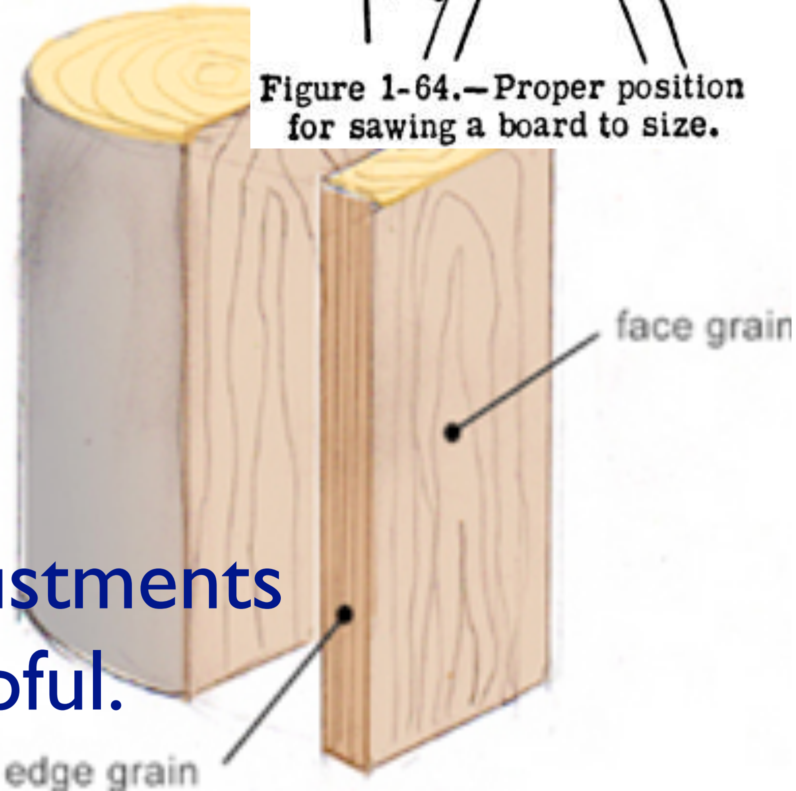

evolvability – "with the grain": adjustments phylogeny has often found helpful.

Speed of Evolution

- How likely is evolution to account for all the variety of nature?
  - Baldwin Effect history (mostly next lecture).
- If each cell is controlled by one gene: new features very unlikely.

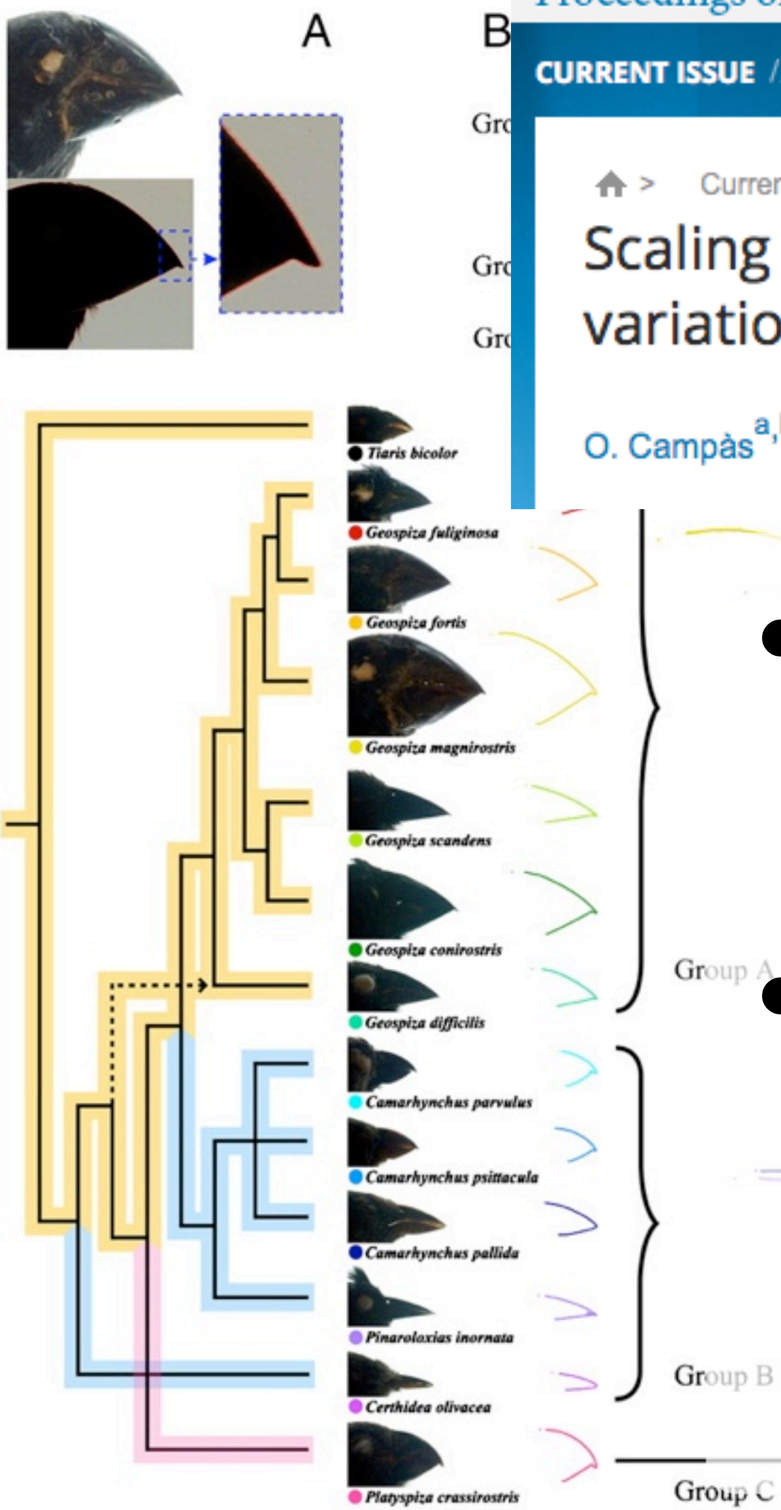

Proceedings of the National Academy of Sciences of the United States of America

CURRENT ISSUE // ARCHIVE // NEWS & MULTIMEDIA // FOR AUTHORS // ABOUT PNAS COLLECTED ARTIC

Current Issue > vol. 107 no. 8 > O. Campàs, 3356–3360

Scaling and shear transformations capture beak shape variation in Darwin's finches

O. Campàs[a,b], R. Mallarino[b], A. Herrel[b], A. Abzhanov[b,1,2], and M. P. Brenner[a,1,2]

- Two transforms capture all of the variation in Darwin's finches' beaks.
- These are correlated with the extent of expression of one gene each (along with appropriate head size etc.)

(Campás et al 2010)

subjects of future ICCS lectures

Evolvability & EvoDevo

- The standard model of evolution taught in schools (and my earlier lecture) was developed in the 1950s to be clear.
    - Variation, transmission, selection.
- Study of epigenetic effects such as evolvability, niche construction, the Baldwin Effect, maternal effects, horizontal transmission (of information, including DNA): sometimes called Evolutionary Developmental Biology.

No Free Lunch

- NN: set of inputs, set of outputs, set of weights between them, set of experiences that produce error signals, nudge weights around to try to get inputs to determine right outputs.
- GA: set of inputs, set of outputs, body connecting them described by genes, flip & switch genes to try to get inputs to determine right outputs.

subjects of past ICCS lectures

No Free Lunch

Besides the learning algorithm!

- NN: set of inputs, set of outputs, set of weights between them, set of experiences that produce error signals, nudge weights around to try to get inputs to determine right outputs.
- GA: set of inputs, set of outputs, body connecting them described by genes, flip & switch genes to try to get inputs to determine right outputs.

all stuff you have to design!

Trying to learn two types of things at once.

Two Kinds of Supervised Learning

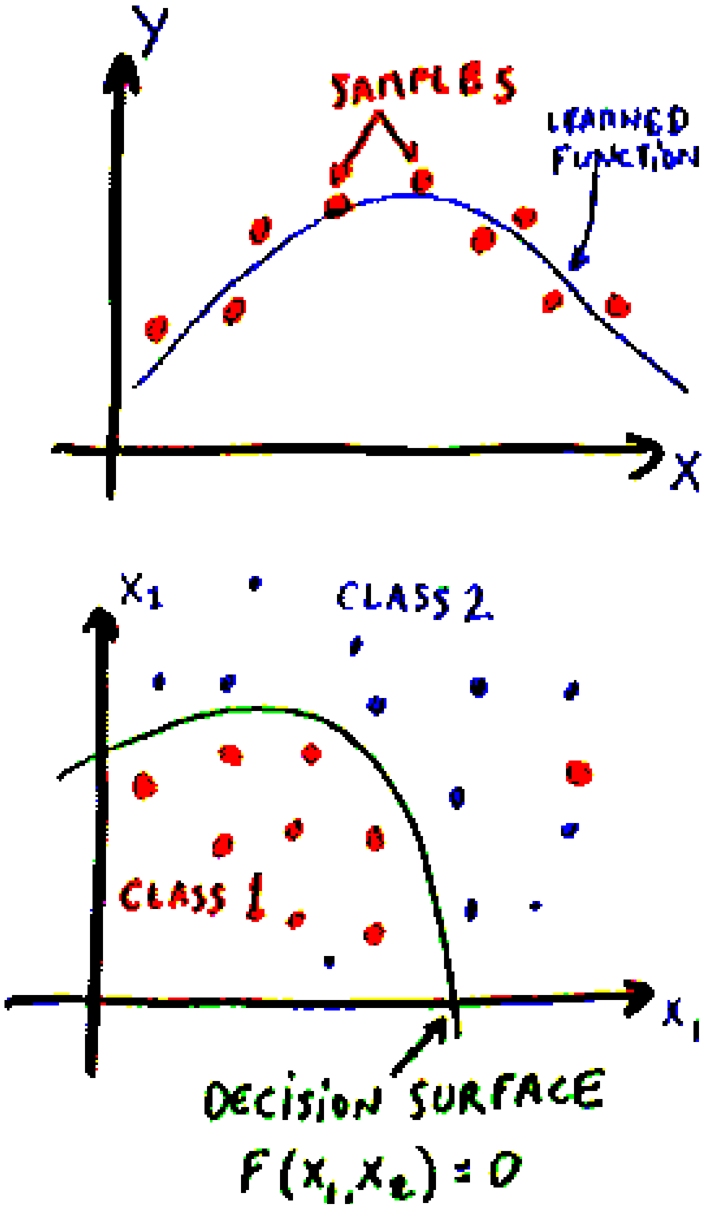

Remember I said I wouldn't cover classification?

- Regression: also known as "curve fitting" or "function approximation". Learn a continuous input-output mapping from a limited number of examples (possibly noisy).
- Classification: outputs are discrete variables (category labels). Learn a decision boundary that separates one class from the other. Generally, a "confidence" is also desired (how sure are we that the input belongs to the chosen category).

I lied

Classification

- Typical machine-learning course: Lots of data, some hypothesised causes (gaussians),
  - want to know which cause accounts for which data so you can reason about it (or something).
- Cognitive Systems: Want to know which contexts best map to which actions.

Would like to learn the context categories and actions too!

Expectation Maximisation (EM) Algorithm

- Have data points & models of causes, but
- The models have parameters, and you aren't sure they're right.
- E step: adjust E – figure out which data is probably accounted for by which cause.
- M: adjust model parameters to improve E.

iterate!

Iterative Learning and Mapping

- Building a navigational map is a good demonstration of the problem of learning in general.
- Continuously update your belief about / model of the world given your perception and knowledge of your own action.

SLAM: Simultaneous Localization and Mapping: Part I

Chang Young Kim

These slides are based on:
Probabilistic Robotics,
S. Thrun, W. Burgard,
D. Fox, MIT Press, 2005
and
Zane Goodwin's Slide from the
previous class

Many images are also taken from
Probabilistic Robotics.
http://www.probabilistic-robotics.com

SLAM

- No Map Available and No Pose Info

You don't know where you are or what you're doing, but you know what the robot CAN do & have a bunch of data.

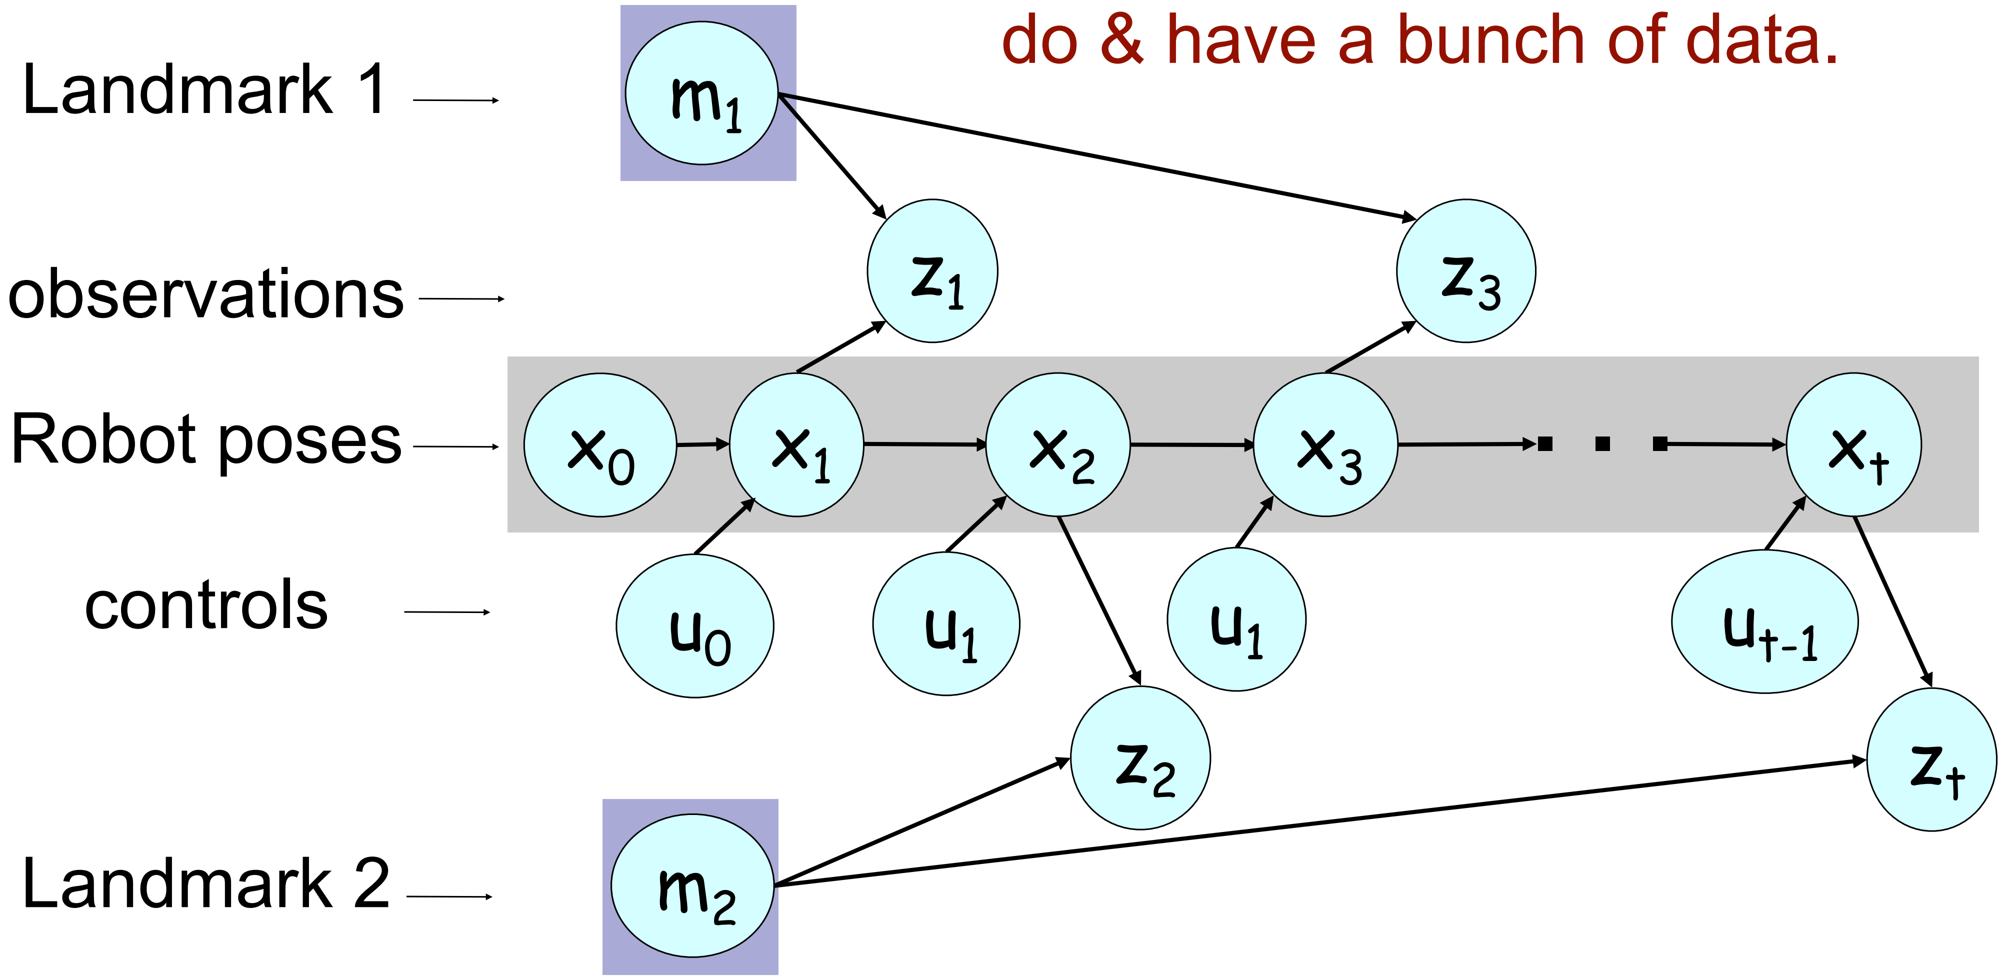

$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t})$$

JB: write vars on whiteboard

Terminology

- Robot State (**including** pose): $x_t = [x, y, \theta]$
  - Position and heading
  - $x_{1:t} = \{x_1, \ldots, x_t\}$
- Robot Controls: u_t
  - Robot motion and manipulation
  - $u_{1:t} = \{u_1, \ldots, u_t\}$
- Sensor Measurements: z_t
  - Range scans, images, etc.
  - $z_{1:t} = \{z_1, \ldots, z_t\}$
- Landmark or Map: $m_i \quad or \quad l_i$
  - Landmarks or Map

$$m = \{m_1, \ldots, m_n\} \; or \; l = \{l_1, \ldots, l_n\}$$

Terminology

- Observation model: $P(z_t \mid x_t)$ or $P(z_t \mid x_t, m)$
  - The probability of a measurement z_t given that the robot is at position x_t *and map m*.
- Motion Model: $P(x_t \mid x_{t-1}, u_t)$
  - The posterior probability that action u_t carries the robot from x_{t-1} to x_t.

Terminology

- Belief: $bel(x_t)$
  - Posterior probability
  - Conditioned on available data
  - $$Bel(x_t) = p(x_t \mid z_t, u_t)$$
- Prediction: $\overline{bel}(x_t)$
  - Estimate before measurement data
  - $$\overline{Bel}(x_t) = p(x_t \mid z_{t-1}, u_t)$$

SLAM algorithm

$$p(x_{1:t}, m \mid z_{1:t}, u_{1:t}) = bel(x_t, m)$$

- Prediction

$$\overline{bel}(x_t, m) = \int p(x_t \mid u_t, x_{t-1})\, bel(x_{t-1}, m)\, dx_{t-1}$$

- Update

$$bel(x_t, m) = \eta\; p(z_t \mid x_t, m)\, \overline{bel}(x_t, m)$$

This is just Bayes' Theorem again.
Kim used a notational abbreviation for the normalisation term.

A Third Iterative Learning Algorithm

- Particle filters build a model of what you believe about an object that's changing.
  - Currently mostly used for object tracking in vision research.
- Use a GA-like algorithm to handle the fact that the probability distributions aren't normal / gaussians.

Particle Filters

- Represent belief by random samples
- Estimation of non-Gaussian, nonlinear processes

- Sampling Importance Resampling (SIR) principle
  - Draw the new generation of particles
  - Assign an importance weight to each particle
  - Resampling

Note: sort of a GA

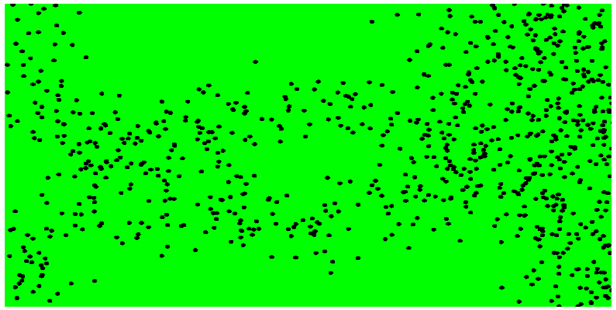

Weighted samples

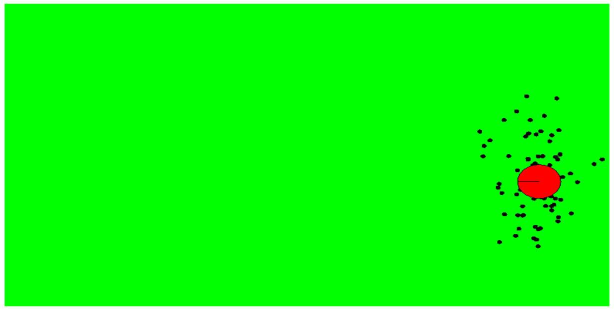

After resampling

Particle Filter Algorithm

$$Bel(x_t) = \eta \; p(z_t \mid x_t) \int p(x_t \mid x_{t-1}, u_{t-1}) \; Bel(x_{t-1}) \; dx_{t-1}$$

draw x^i_{t-1} from $Bel(x_{t-1})$

draw x^i_t from $p(x_t \mid x^i_{t-1}, u_{t-1})$

Importance factor for x^i_t:

$$
\begin{aligned}
w^i_t &= \frac{\text{target distribution}}{\text{proposal distribution}} \\
&= \frac{\eta \; p(z_t \mid x_t) \; p(x_t \mid x_{t-1}, u_{t-1}) \; Bel(x_{t-1})}{p(x_t \mid x_{t-1}, u_{t-1}) \; Bel(x_{t-1})} \\
&\propto p(z_t \mid x_t)
\end{aligned}
$$

What good is a map?

- Need to know not only location, but actions.
- Mammals learn maps and actions with hippocampus, have different maps for different task contexts.

Problem Spaces

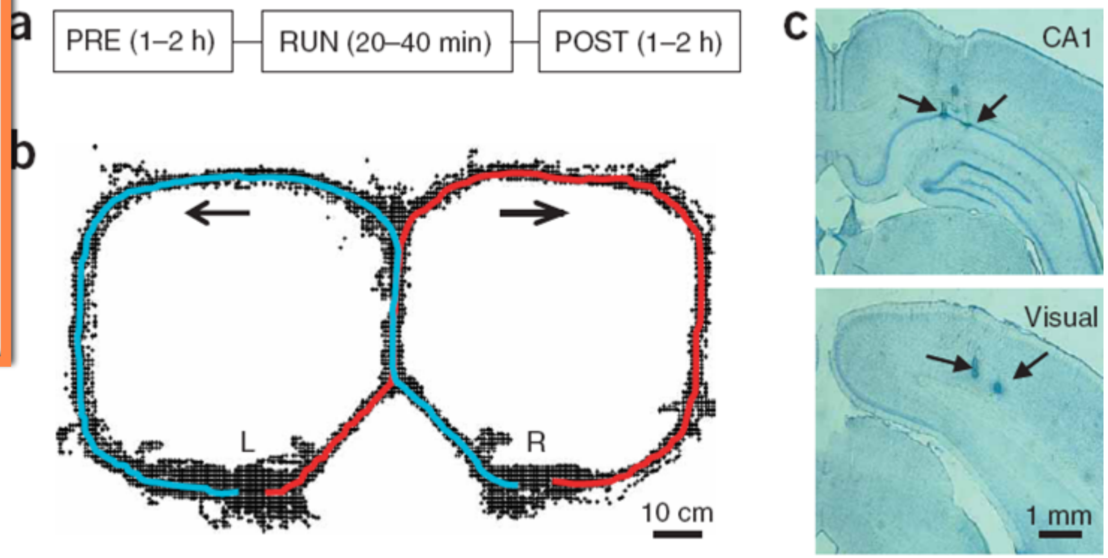

Hippocampal Learning

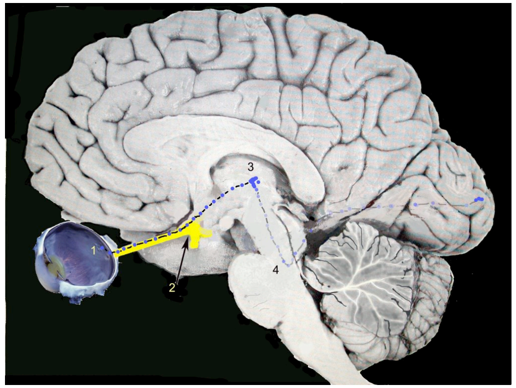

- Sparse representation.
- Useful for episodic memory – data.
- Sparse representation indexes into concept memory (in neocortex) – model.
- Data used to update model e.g. in sleep, resting (like EM).

McClelland, McNaughton, O'Reilly 1995;
Rogers & McClelland 2004; Foster & Wilson 2006.

Learning in Modular Systems

- OOD recommends modular decomposition along state / representation requirements.
- Modular approaches can facilitate learning as well as design.
- Hippocampus & Neocortex are one example of such decomposition in nature.

Simple Robot Sensor Fusion & Map Learning

(Bryson 1997, 2001)

- Nomad 2000
- Sonars, IR, Bumpers, Odometry

Joanna J. Bryson "The Behavior-Oriented Design of Modular Agent Intelligence", *Agent Technologies, Infrastructures, Tools, and Applications for e-Services*, R. Kowalszyk, J. P. Müller, H. Tianfield and R. Unland, eds., pp. 61–76, Springer, 2003.

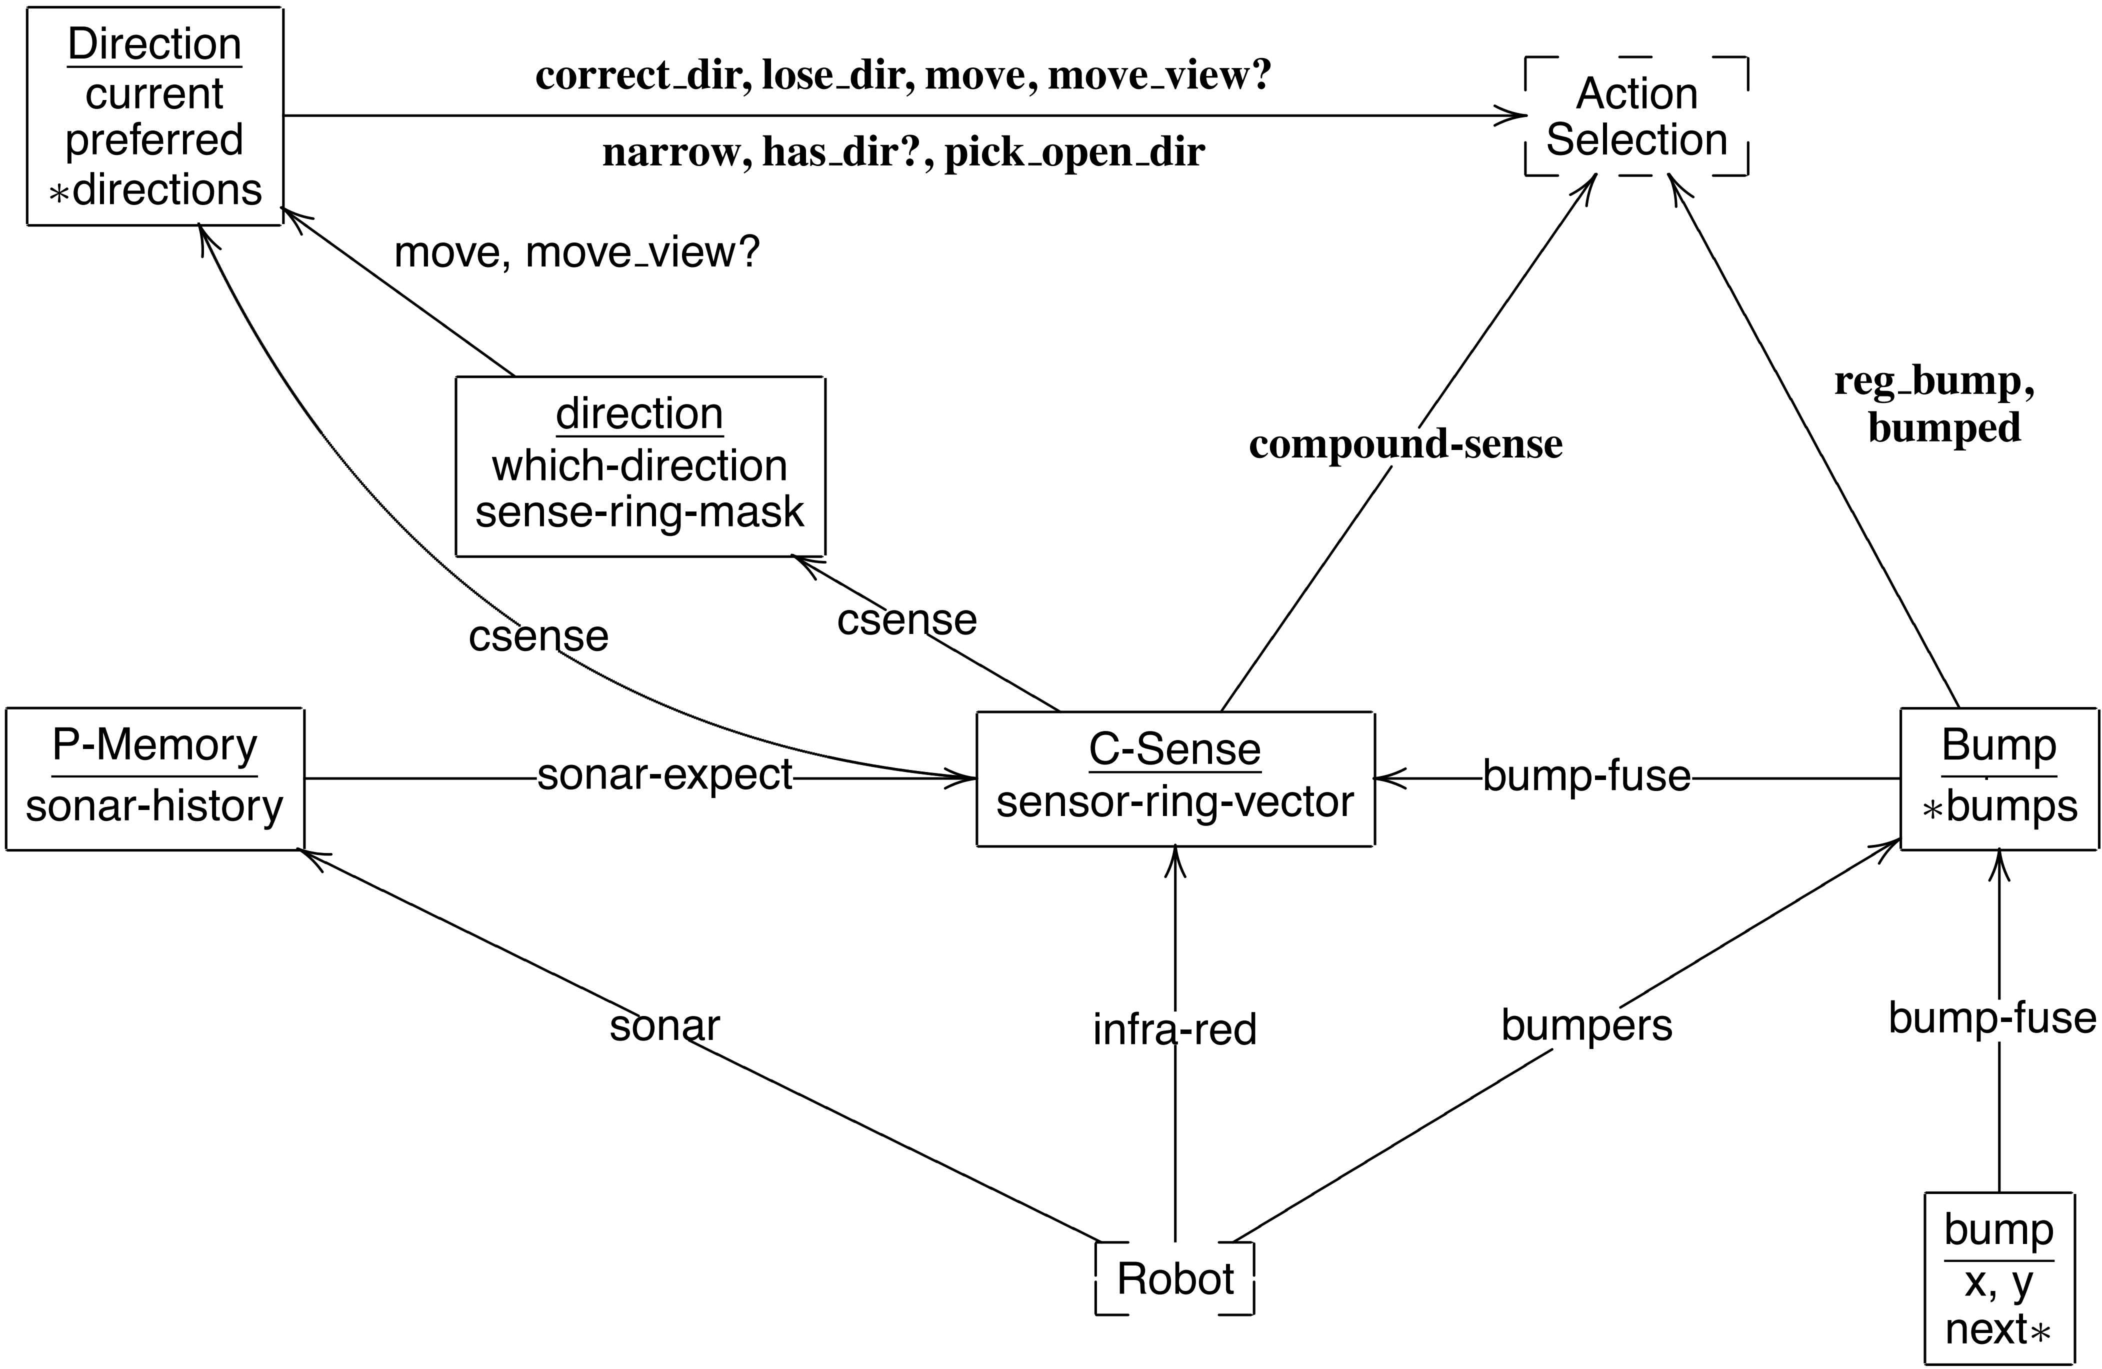
Direction
current
preferred
∗directions
correct_dir, lose_dir, move, move_view?
narrow, has_dir?, pick_open_dir
Action
Selection
move, move_view?
direction
which-direction
sense-ring-mask
compound-sense
reg_bump,
bumped
csense
csense
P-Memory
sonar-history
sonar-expect
C-Sense
sensor-ring-vector
bump-fuse
Bump
∗bumps
sonar
infra-red
bumpers
bump-fuse
Robot
bump
x, y
next∗

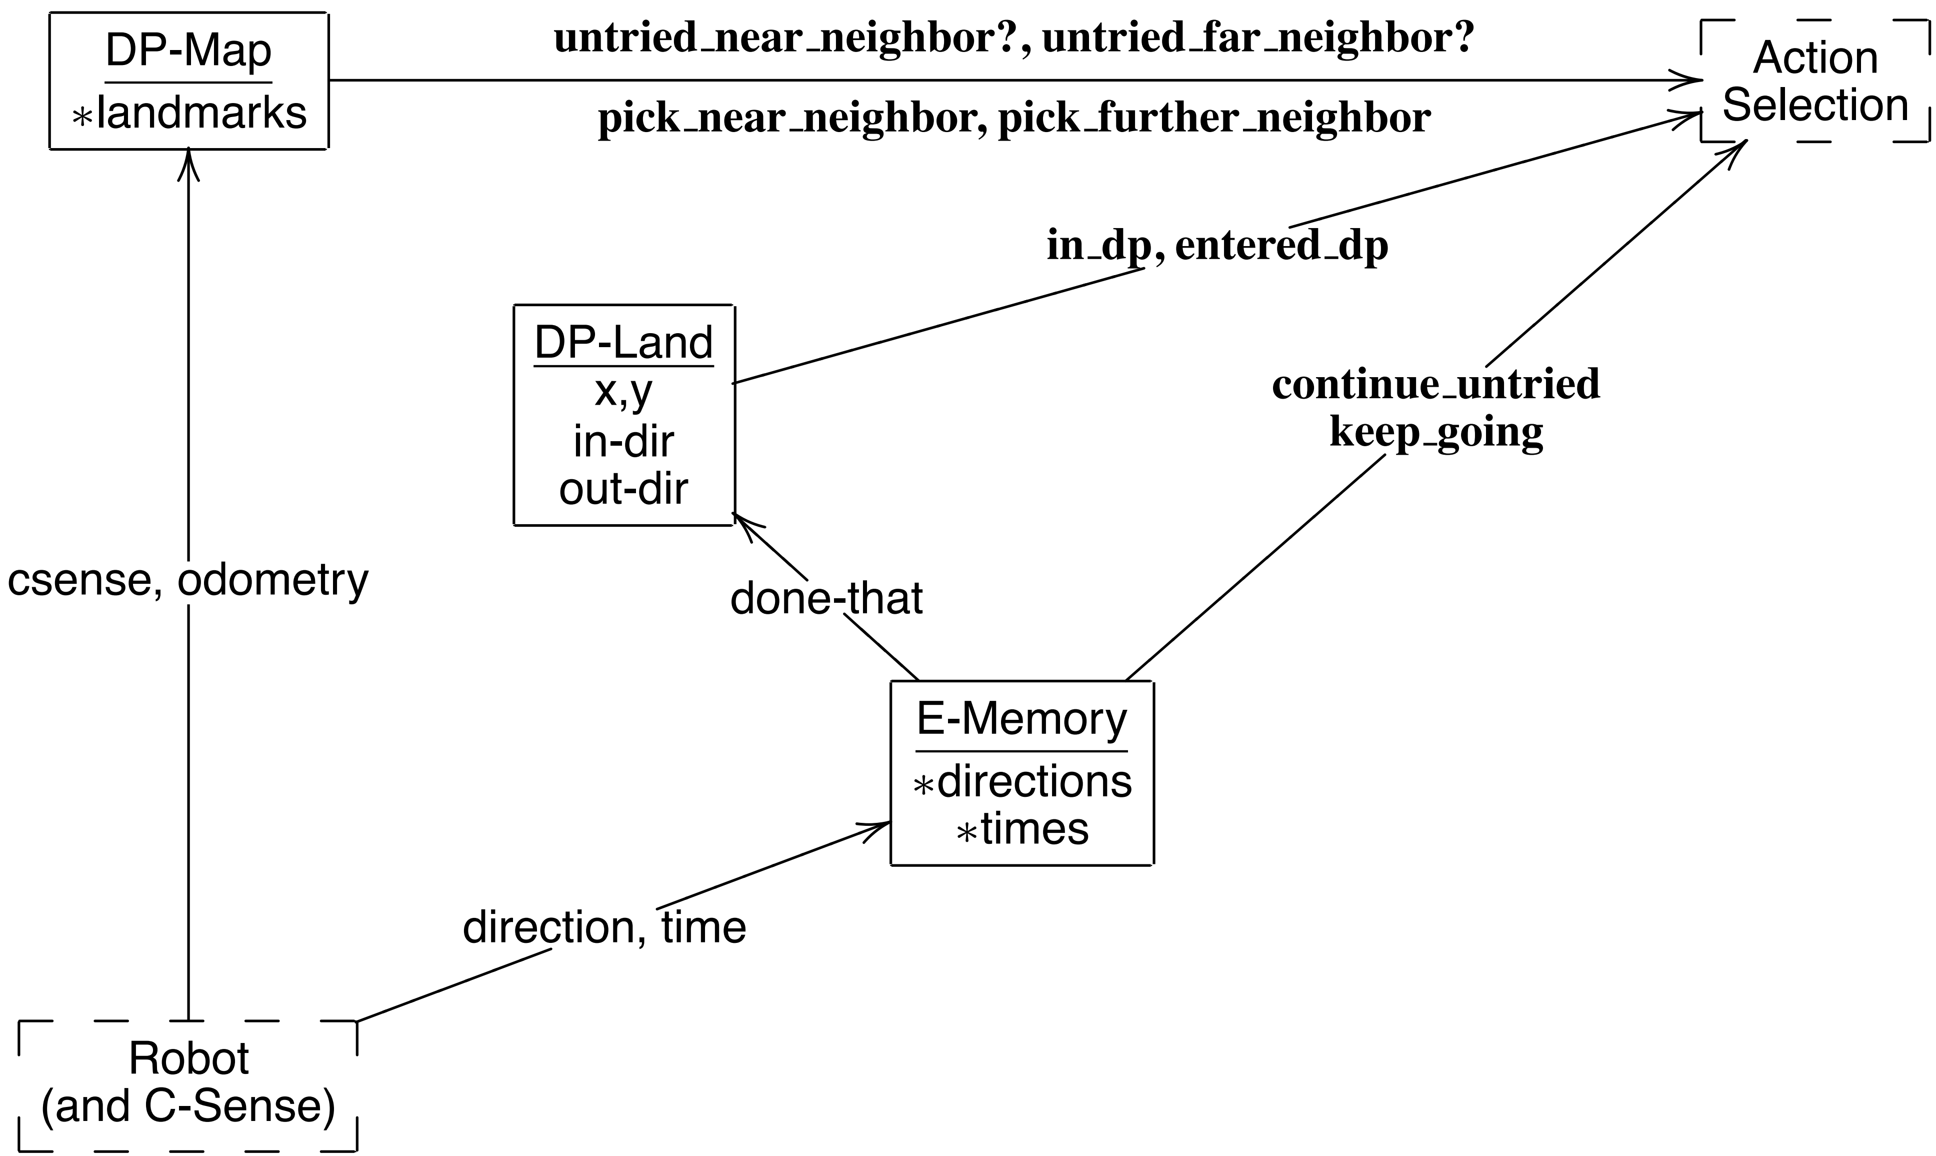
DP-Map
*landmarks
untried_near_neighbor?, untried_far_neighbor?
pick_near_neighbor, pick_further_neighbor
Action Selection
in_dp, entered_dp
DP-Land
x,y
in-dir
out-dir
continue_untried
keep_going
csense, odometry
done-that
E-Memory
*directions
*times
direction, time
Robot
(and C-Sense)

Video (from NTSC)

3 Stages of Development:

1. totally flexible obstacle avoidance,
2. limit the flexibility, to
3. allow learning maps.

3 minutes, 1998

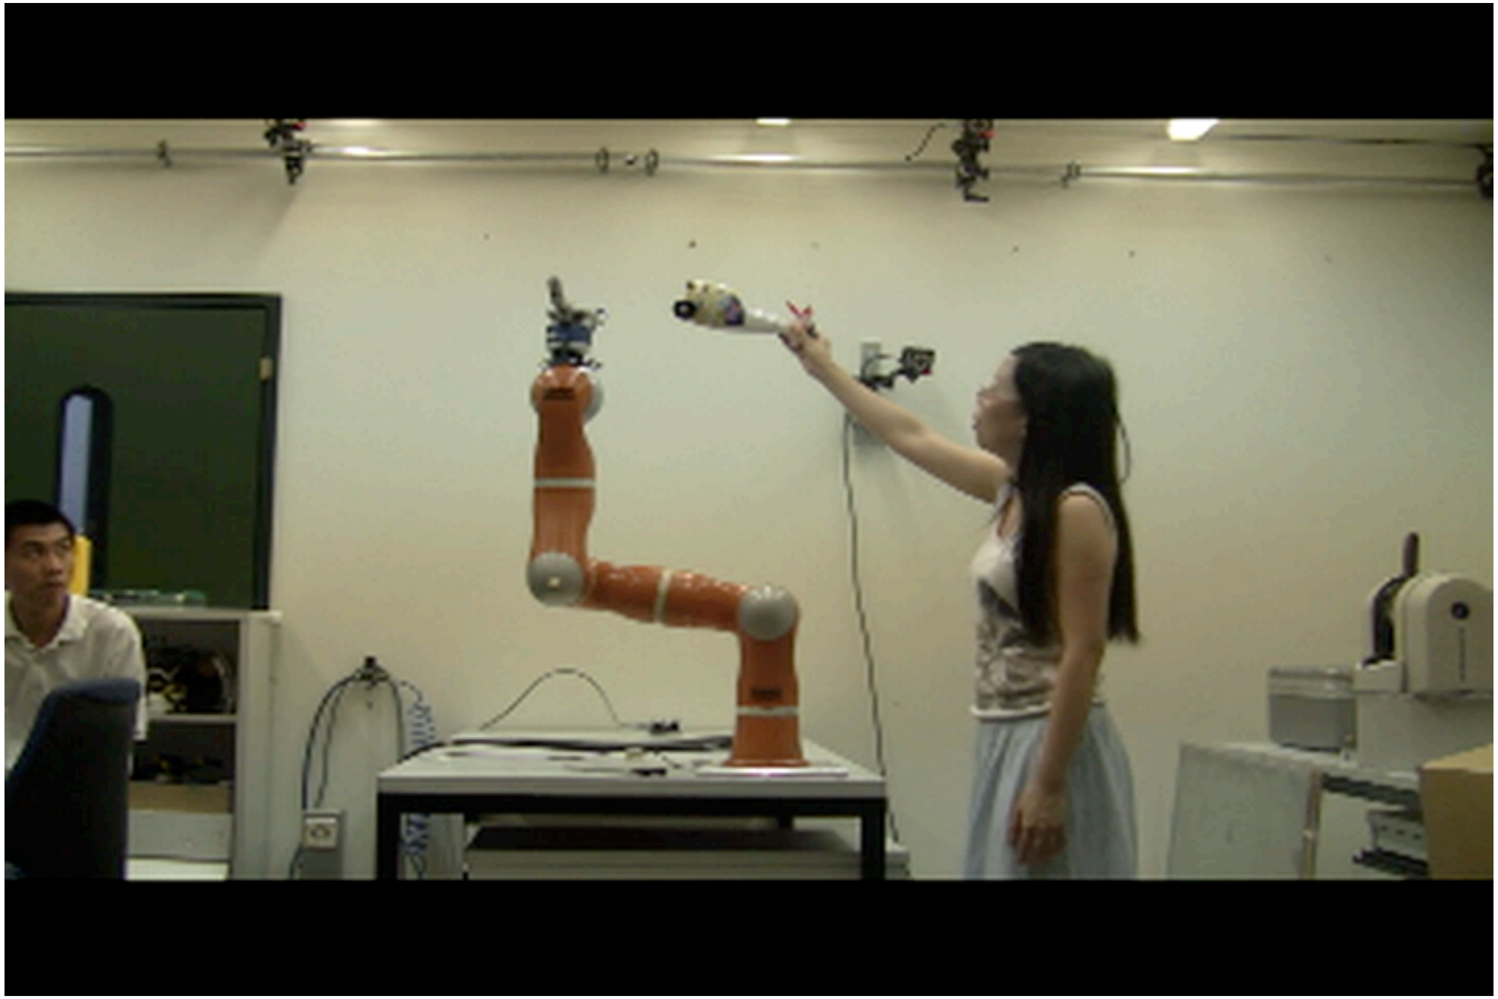

Bidan Huang, The Use of Modular Approaches For Robots to Learn Grasping and Manipulation, PhD 2015

System overview

1. Human demonstration

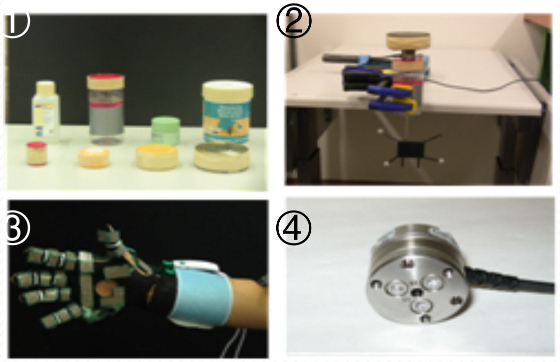

① Bottles and caps

② Motion trackers

③ Dataglove mounted with
tactile sensors

④ Force torque sensor

2. Model learning

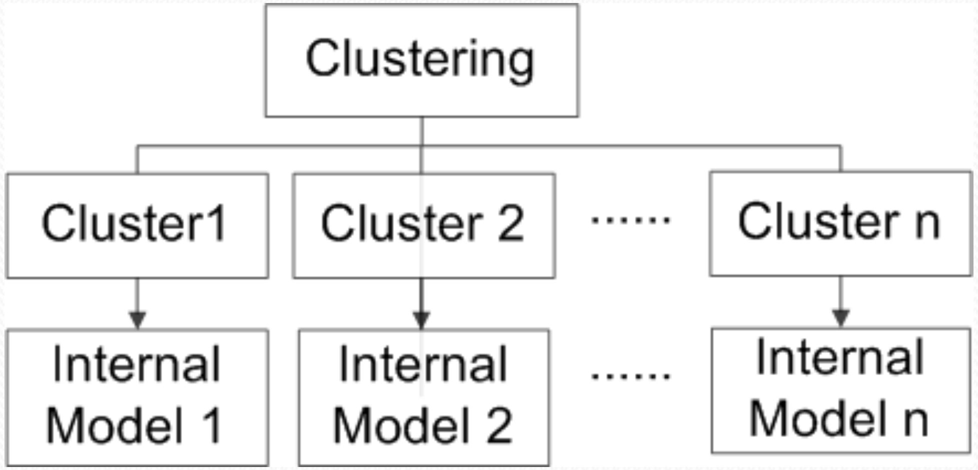

3. Robot control

Learning Multiple Module Control Model (1)

- Adaptive Control for changing task context
- Multiple controllers

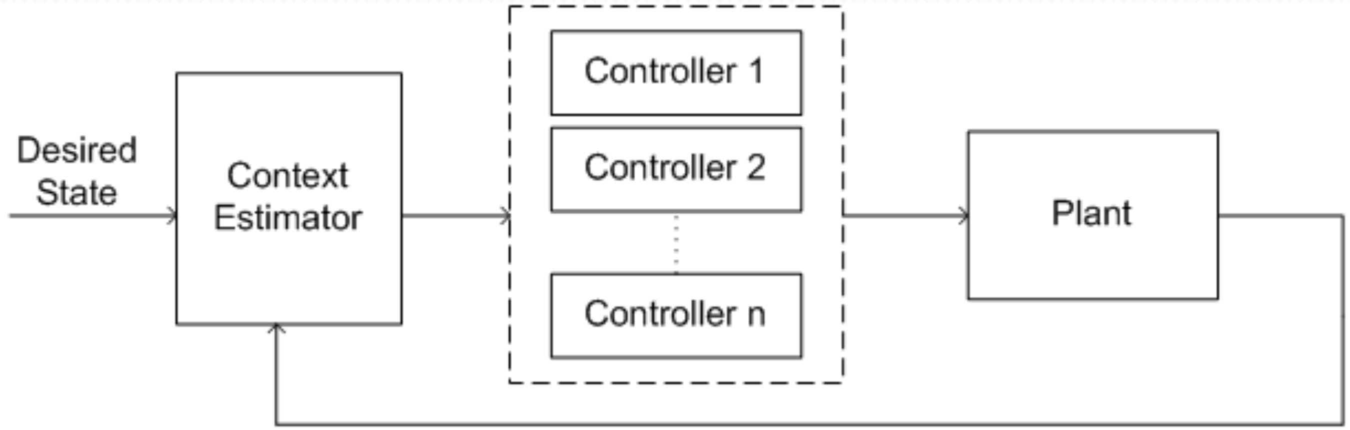

Learning Multiple Module Control Model(2)

- How many modules?
  - Hierarchical clustering
- How to represent each module?
  - Internal Model
    - Pair of Forward and Inverse Models

$$s_{t+1} = f(s_t, a_t) \qquad a_t = g(s^*_{t+1}, s_t)$$

s_t: object status at time t

a_t : action applied on object at time t

s^*_{t+1}: desired object status at time t+1

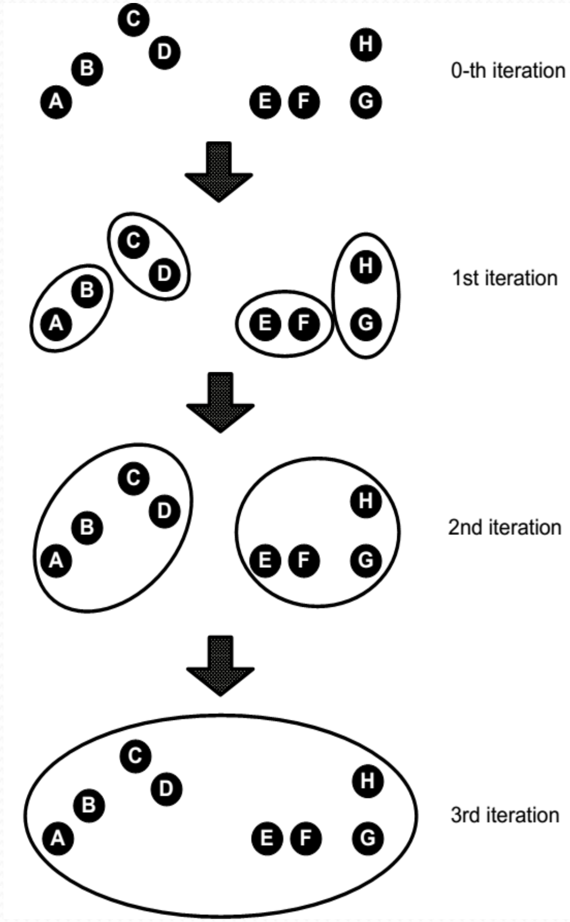

Hierarchical Clustering

Robot Control

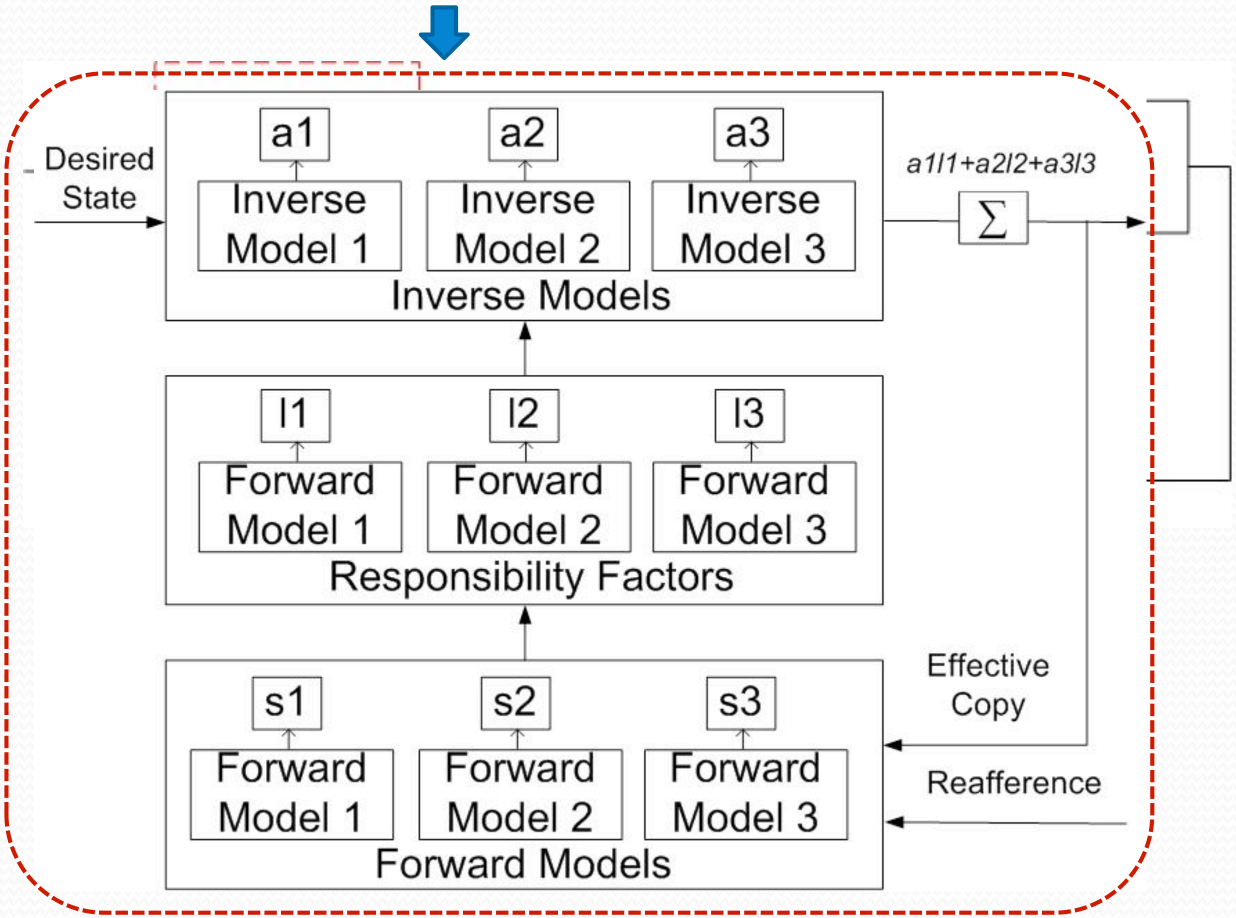

Experiment (1)

- Task:
  - opening bottle cap

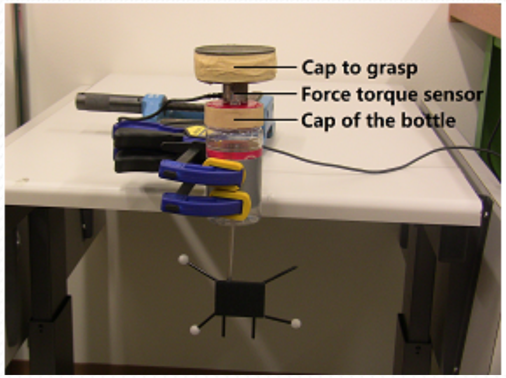

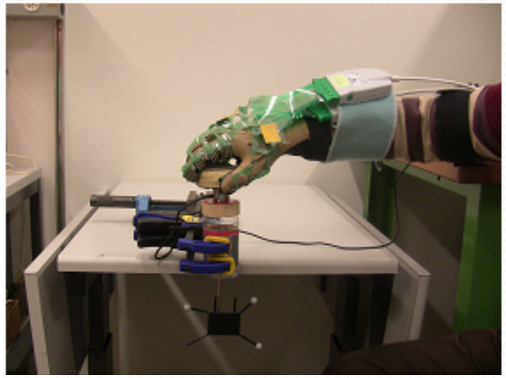

- Human demonstrate in different task contexts
  - Different bottles
  - Different cap sizes

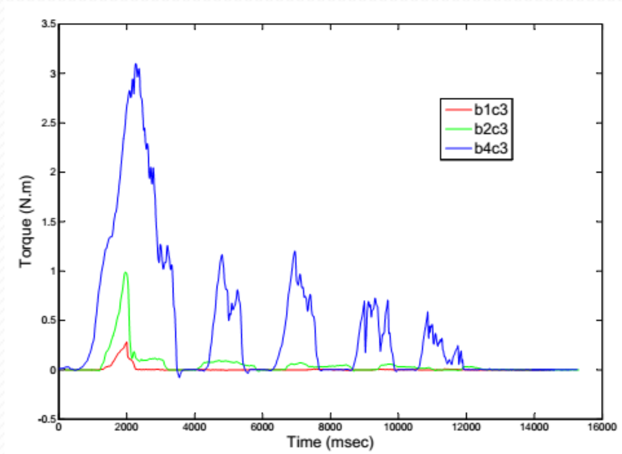

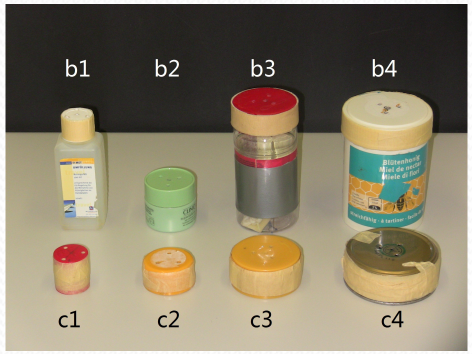

Experiment (2)

- Dynamic Time Warping
- Clustering

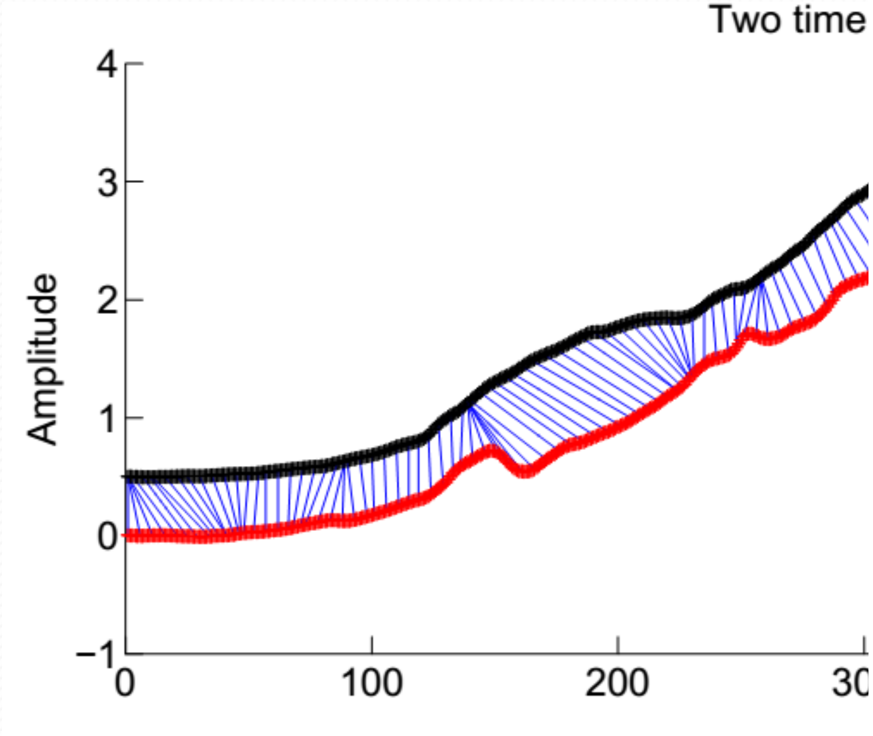

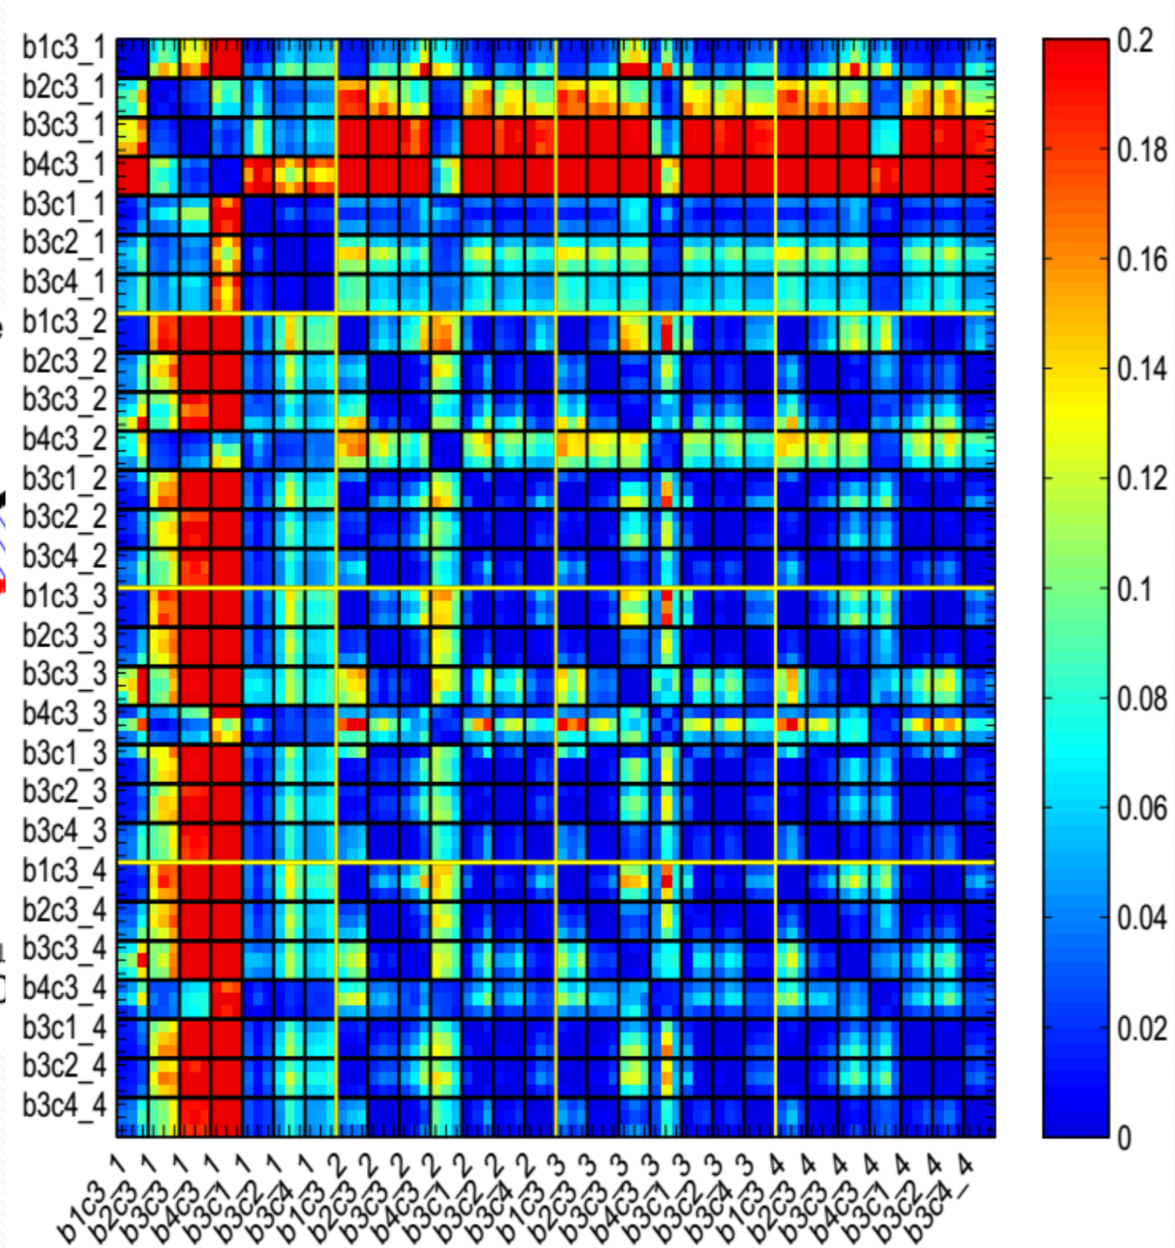

Experiment (3)

- Learning Modules
  - Clustering Control Strategies
  - Encoding by GMM
    - Forward model

$$p\{s_t, s_{t-1}, a_{t-1} \mid \Omega_F\}$$

    - Responsibility factor

$$\eta_t = \{s_t, s_{t-1}, a_{t-1}\}$$

$$\lambda_t^k = \frac{p(\eta_t|\Omega_F^k)}{\sum_{j=1}^{J} p(\eta_t|\Omega_F^j)}$$

    - Inverse model

$$p\{s_t, st+1, a_t, a_{t-1} \mid \Omega_I\}$$

$$a_t = \sum_{k=1}^{K} \lambda_t^k a_t^k = \sum_{k=1}^{K} \lambda_t^k E\left(a_t \mid s_{t+1}^*, s_t, a_{t-1}, \Omega_I^k\right)$$

		 Cap 1	 Cap 2	 Cap 3	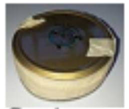 Cap 4
	Phase I			(b1c3) Cluster 3	
Bottle 1	Phase II			Cluster 3	
	Phase I			(b2c3) Cluster 2	
Bottle 2	Phase II			Cluster 3	
	Phase I	(b3c1) Cluster 2	(b3c2) Cluster 2	(b3c3) Cluster 2	(b3c4) Cluster 2
Bottle 3	Phase II	Cluster 3	Cluster 3	Cluster 3	Cluster 3
	Phase I			(b4c3) Cluster 1	
Bootle 4	Phase II			Cluster 2	

Aggressive Quadrocopter

https://www.youtube.com/watch?feature=player_embedded&v=geqip_0Vjec

https://www.grasp.upenn.edu/

Very Fast SLAM!

Simple Solutions: The Polly Algorithms

- First autonomous robot to use vision to navigate at animal-like speeds (1 m/sec).
- Ground plane assumption / 3-way navigation choice.
- Map following via odometry × low-resolution landmarks.

Specialised Perception

Hand-coded Map & Task

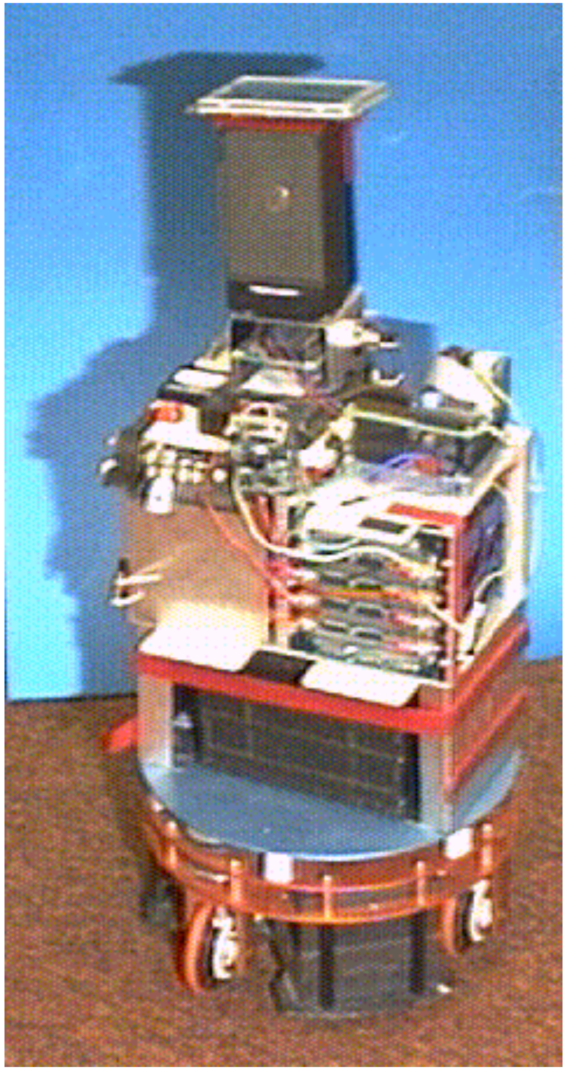

(Horswill 1993)

The Polly Algorithm

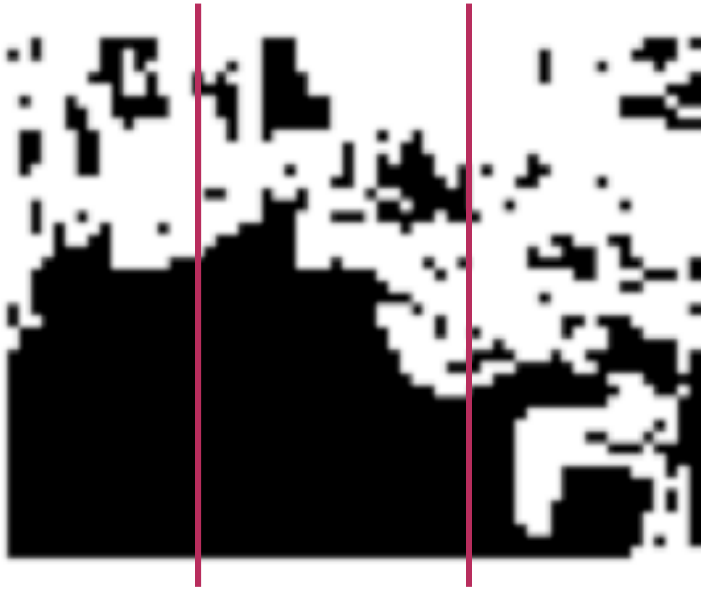

- Very low res frames give fast processing.
- Ground Frame Assumption: What is in front of you is the ground, what looks like it is too.
- Go one of three ways. Works at 60Hz.

Polly Navigation

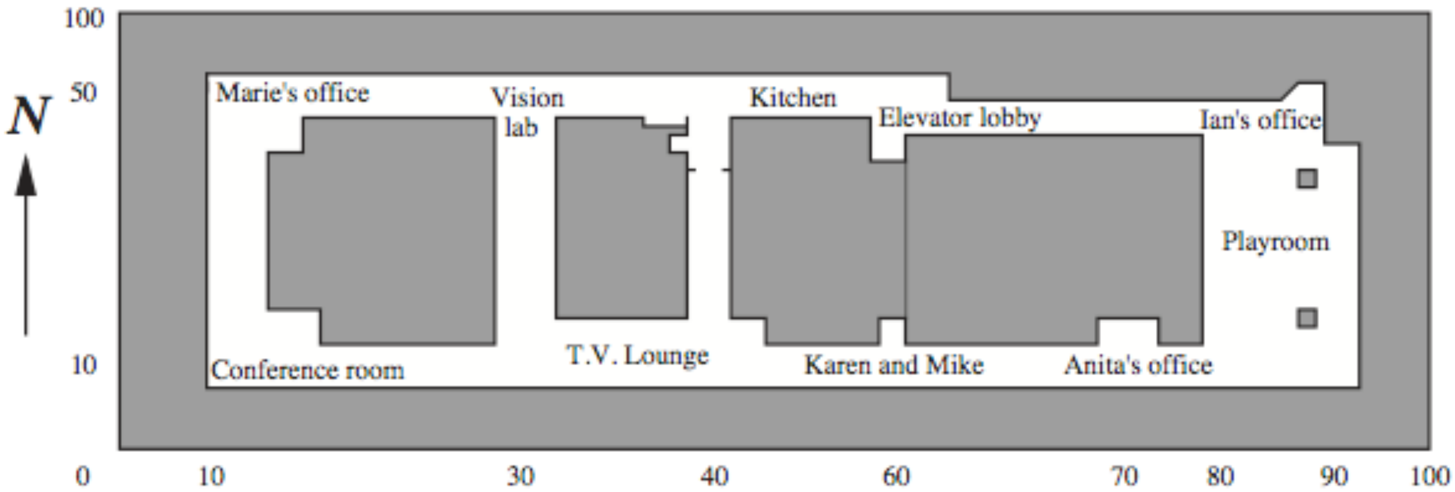

Figure 2.5: Polly's habitat, the 7th floor of the MIT AI laboratory. The diagram is not to scale, and the direction north has been defined for convenience, rather than geological accuracy.

- Hand-coded rough map of 7th floor of AI Lab
- A few landmarks recognised by 5x7 pixel images give certain location.
- Otherwise approximate using odometry.

- I have the video in NTSC, I've asked for either digital or permission for Bath to convert it.
- In the mean time, a bit less exciting: Cerebus winning a prize at IJCAI 2001.

Summary

- Learning has a significant role in cognition, and is largely what advances AI right now.
  - Due to algorithm, systems engineering and hardware advances.
- Learning never happens in a vacuum.
  - Even provably-equivalent systems will learn different things. Path dependency is part of bias / priors making it tractable.