Do bees have Bayesian brains?

Barbara Webb
School of Informatics
University of Edinburgh

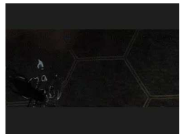

What are the important problems of my field?

What field?

- Roughly, knowing enough about insects to build robots of comparable competence

Some important problems I won’t discuss:

- Do some low level details really matter, e.g. molecular basis of learning and memory?
- What biophysical principles are needed to get comparable small, powerful efficient, robust actuation?

‘Bees’ = clever insects in general…

E.g. Giurfa et al (2001) Bees can learn concept of ‘same’ and ‘different’ in delayed match to sample task

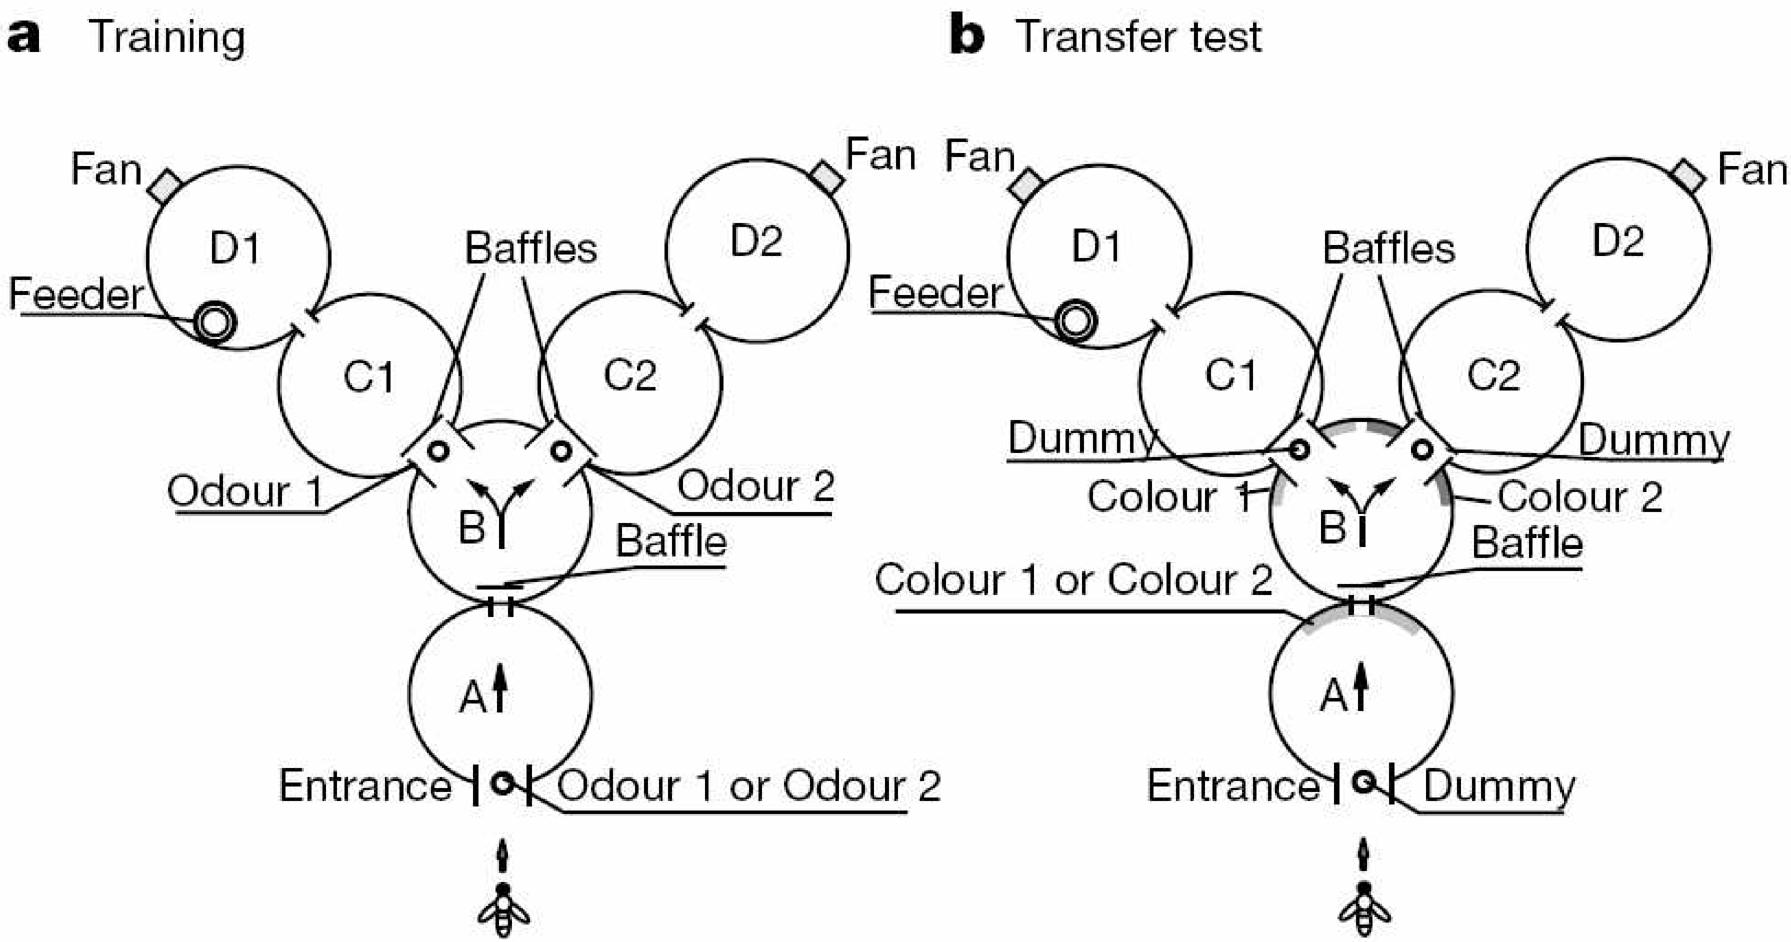

‘Bayesian brains?’

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

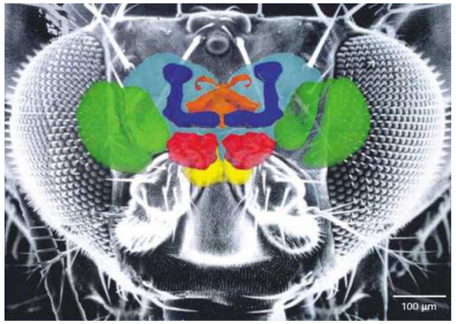

Heisenberg (2003)

Descriptive? Animals operating under uncertainty behave in ways near to the Bayesian optimum

Do animals operating under uncertainty behave in ways near to the Bayesian optimum?

Example: Bayesian foraging *(Valone, 2006)*

•Animals indicate estimated quality of a food patch by when they stop searching within it and move on to another.

•Could just use current information – e.g. falling rate of food item encounters, or fixed 'giving up' time since last food encounter

•Actual behaviour better fitted by Bayesian models where current information is combined with a prior representing the distribution of patch quality; for high variance environments this is the optimal strategy.

•Training animals in different environments (manipulating the prior) leads to changes in quitting rates consistent with Bayesian predictions *(e.g. shown for bumblebees in Biernaske et al. 2009)*

•However, to date, the evidence supports use of prior information but (arguably) not explicit numerical confirmation of Bayesian estimation.

‘Bayesian brains?’

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

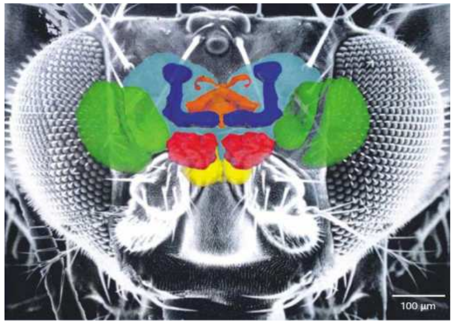

Heisenberg (2003)

Descriptive? Animals operating under uncertainty behave in ways near to the Bayesian optimum

Mechanistic? Above behaviour is explained by assuming nervous systems use predictive internal models –

- that are probabilistic or represent uncertainty
- that are updated by methods that at least approximate Bayesian probability

Two routes to a Bayesian bee

Starting point of modelling ‘simple’ insect behaviour:

- Take mechanistic view to identify key problems
- Look at state of the art ‘engineered’ solutions

→ Find answer is (very often) Bayesian

High level theory of the ‘Bayesian brain’ (e.g. Friston):

- Proposal is general, so should apply to insects
- Key is to map to neurophysiological mechanisms, but for humans/primates/vertebrates this is poorly constrained

→ Perhaps insects would be a good first target?

Example 1: multimodal cue integration in crickets

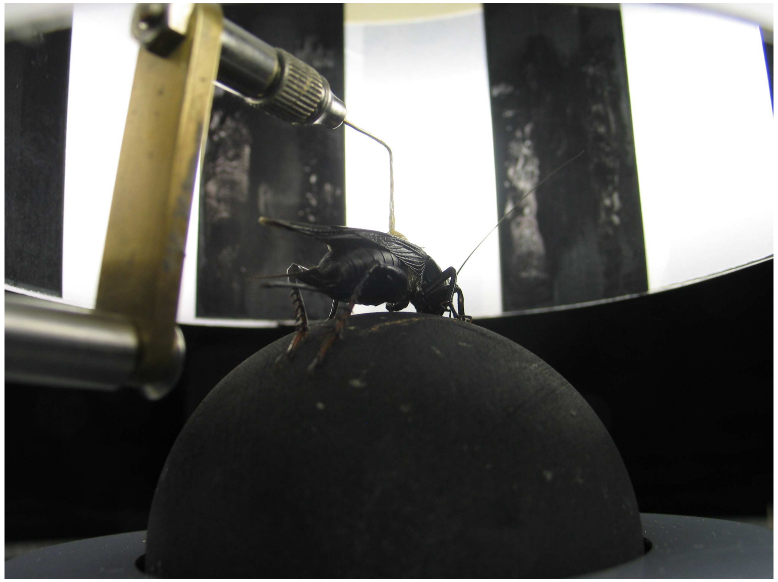

- How does a cricket's turning response reflect combined auditory and visual stimuli?
- Closed loop behavioural data initially suggests it just 'adds' the two turning tendencies

- Data from Mark Payne

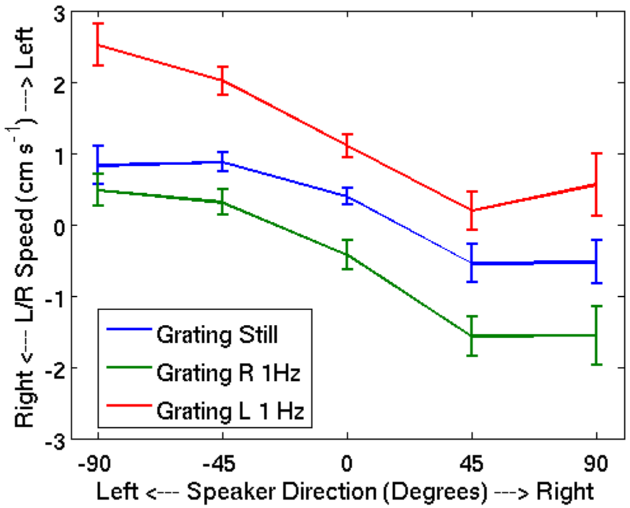

Not additive

- Auditory and optomotor models tuned to dynamics of cricket behaviour
- Simple addition is problematic in closed loop: robot repeatedly 'corrects' for the visual flow induced by turns towards sound.
- Arena tests on crickets show they do not exhibit this problem

Adding optomotor control to robot

Mark Payne

- Not simple inhibition – cricket reacts to change of visual feedback during auditory-evoked turns

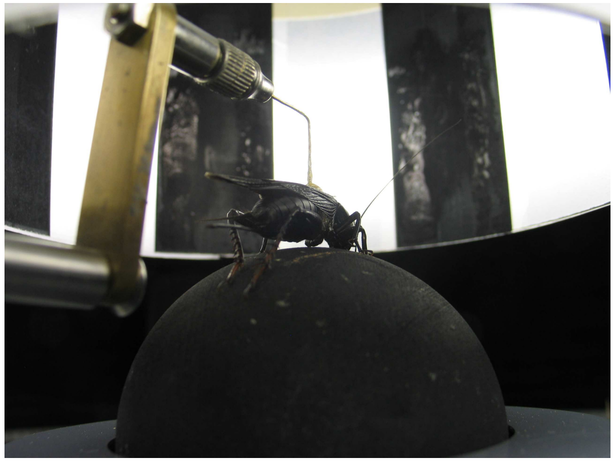

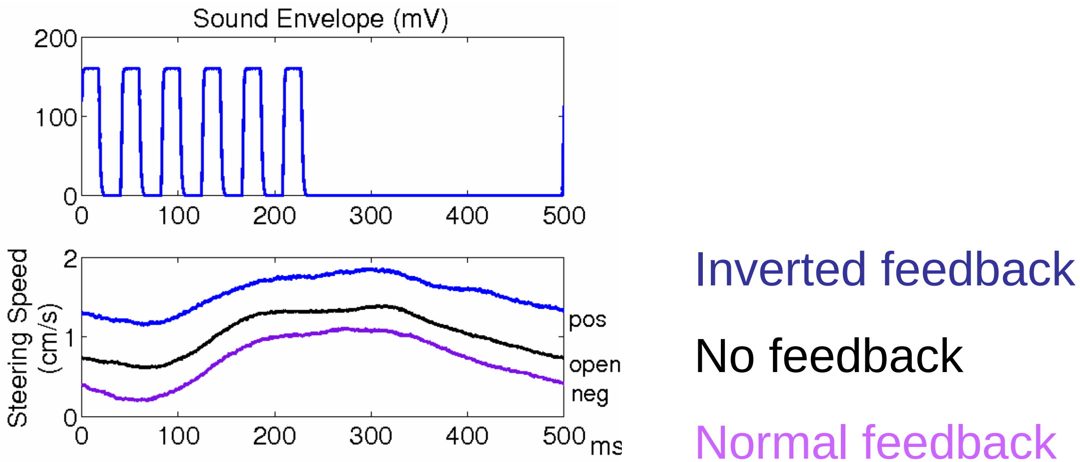

- Possibly involves efference copy and forward model?

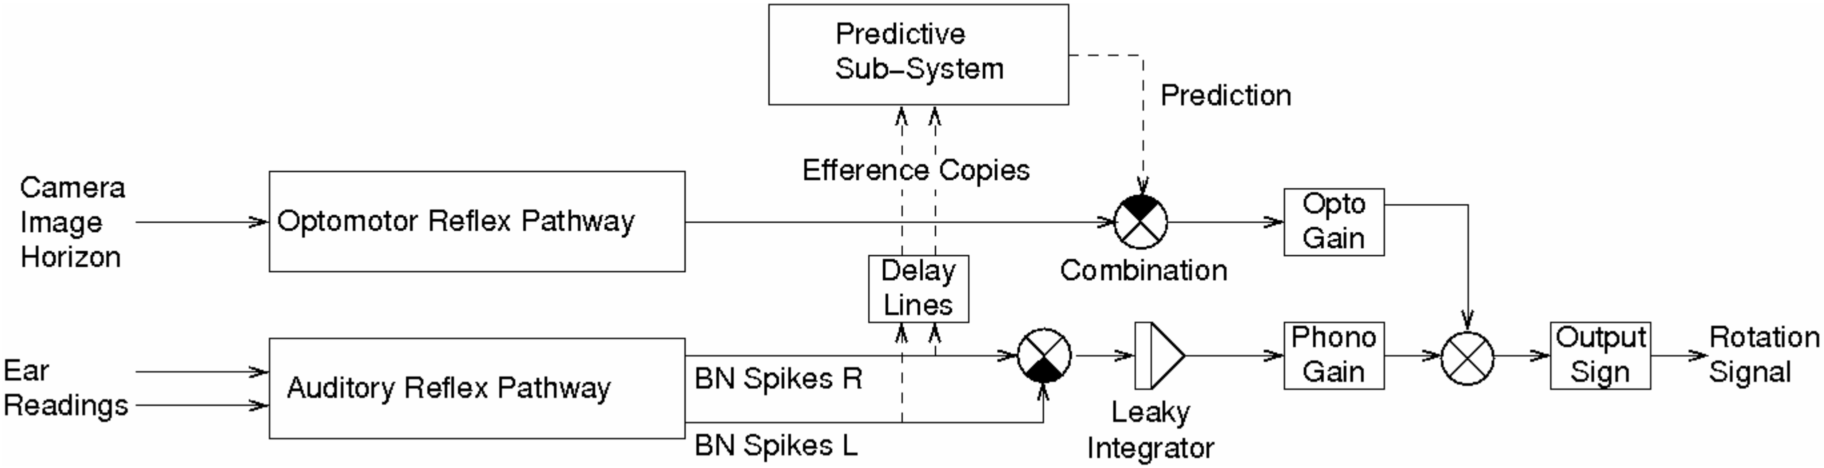

Von Holst & Mittlestadt (1950)

"efference copy"

"efference leaves an 'image' of itself somewere in the CNS, to which the reafference of this movement compares as the negative of a photograph to its print…movement will continue until the reafference exactly nullifies the efference copy" Von Holst (1954)

Explicitly claims to be discussing at functional not neural level

Sperry (1950)

"corollary discharge"

"[movement] may have a corollary discharge into the visual centres…an anticipatory adjustment in the visual centers specific for each movement with regard to its direction and speed"

Describes as providing 'neural basis' for 'effort of will'

Neither address problem of translation between motor commands and sensory signals: the function of the 'forward model'

Forward model: predicts the future state of a system given the current state and the control signals. E.g. for sensorimotor control:

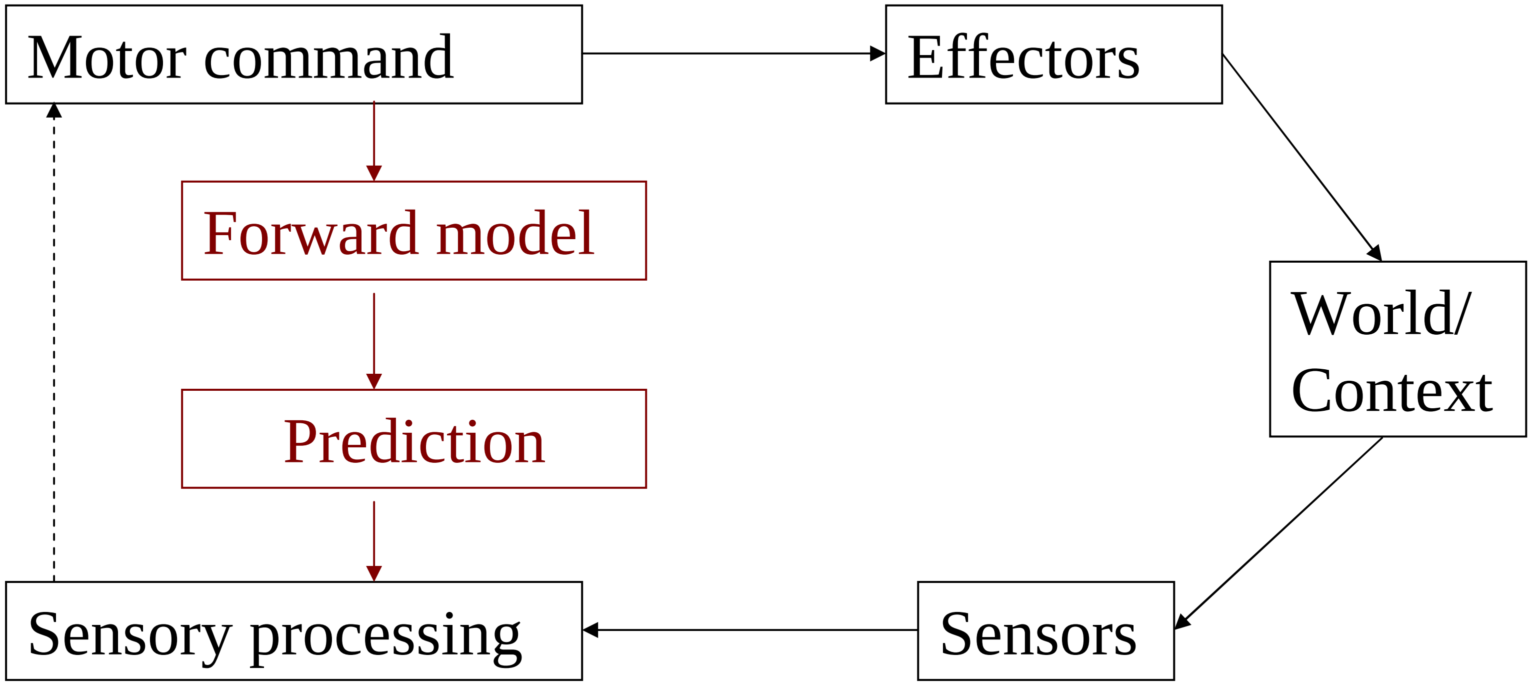

Forward model: in principle represents the external processes, in practice may be highly simplified, or use look-up table, or be learned.

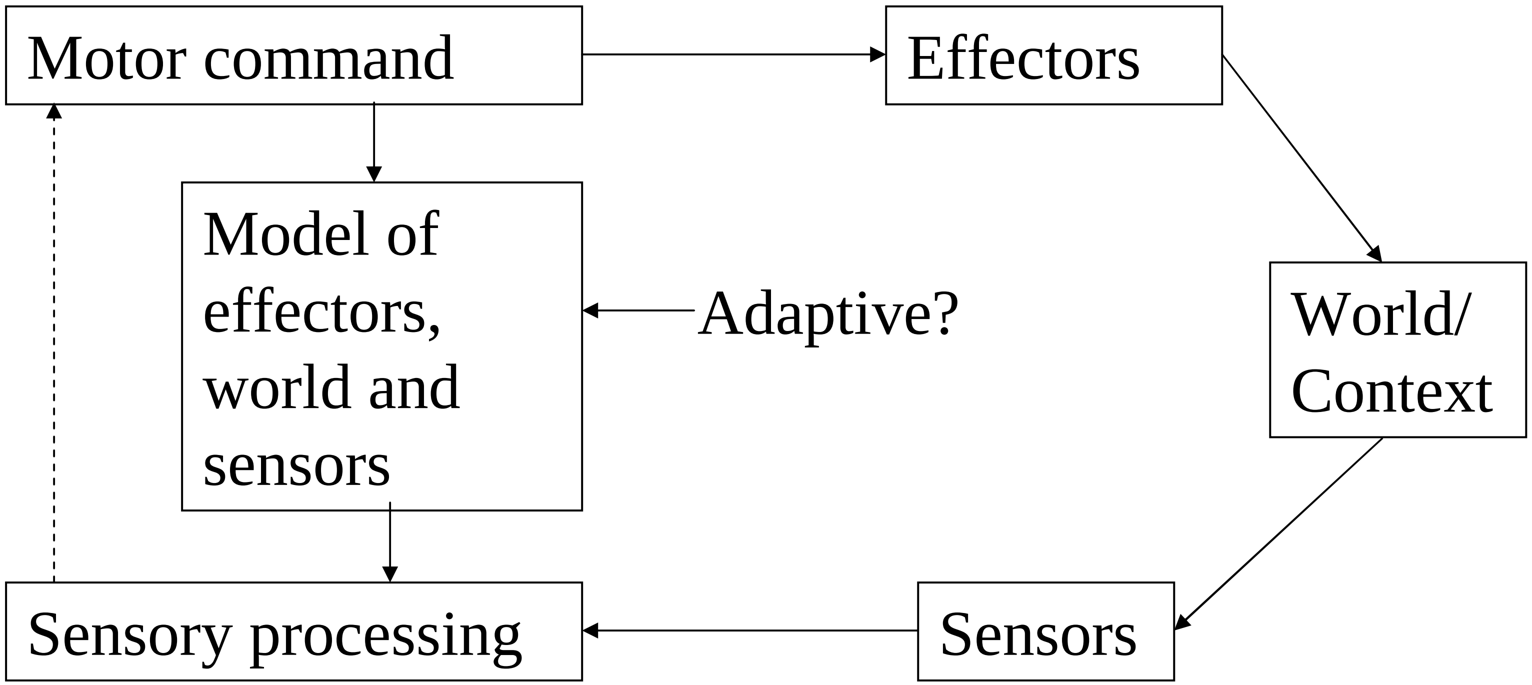

Are insects doing forward modelling?

- Is there modulation of sensory processing, linked to the current behaviour?
- Is this modulation the result of internal connections from motor to sensory areas?
- Does the observed modulation require non-trivial predictive processes?

- Does the modulation resemble a Bayes filter?

Bayes Filter Summary

z = observation
u = action
x = state

$$\boxed{Bel(x_t)} = P(x_t \mid u_1, z_1 \ldots, u_t, z_t)$$

Bayes
$$= \eta\ P(z_t \mid x_t, u_1, z_1, \ldots, u_t)\ P(x_t \mid u_1, z_1, \ldots, u_t)$$

Markov
$$= \eta\ P(z_t \mid x_t)\ P(x_t \mid u_1, z_1, \ldots, u_t)$$

Total prob.
$$= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_1, z_1, \ldots, u_t, x_{t-1})\ P(x_{t-1} \mid u_1, z_1, \ldots, u_t)\ dx_{t-1}$$

Markov
$$= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1})\ P(x_{t-1} \mid u_1, z_1, \ldots, u_t)\ dx_{t-1}$$

Markov
$$= \eta P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1})\ P(x_{t-1} \mid u_1, z_1, \ldots, z_{t-1})\ dx_{t-1}$$

$$\boxed{= \eta\ P(z_t \mid x_t) \int P(x_t \mid u_t, x_{t-1})\ Bel(x_{t-1})\ dx_{t-1}}$$

Bayes Filter is usually described as a two step process

- Prediction of state x_t given x_{t-1} and control action u_t

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

Or for discrete state values:

$$= \sum_x p(x_t \mid u_t, x_{t-1})\, bel(x_{t-1})$$

- Correction given observation (sensor input) z_t

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

Bayes filter

- Often applied to human cognition e.g. Grush (2004)
- *"The idea is that in addition to simply engaging with the body and environment, the brain constructs neural circuits that act as models of the body and environment. During overt sensorimotor engagement, these models are driven by efference copies in parallel with the body and environment, in order to provide expectations of the sensory feedback, and to enhance and process sensory information."*
- E.g., what we see depends on what we expect to see.
- Is the same true for crickets?

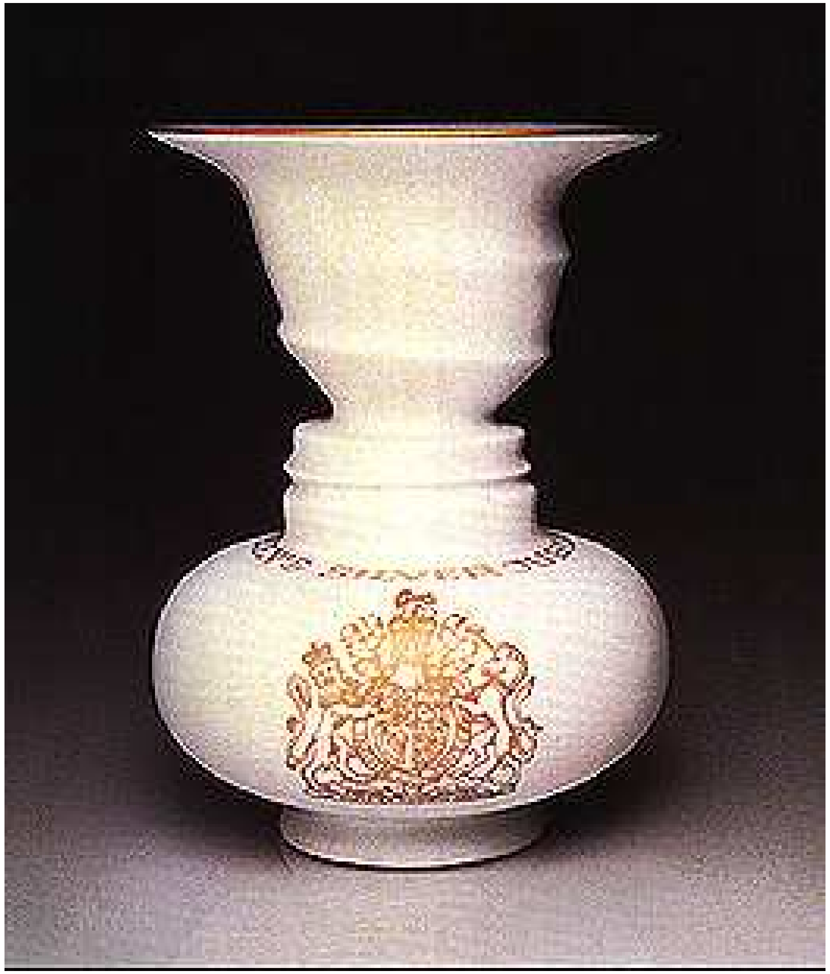

Example 2: ant navigation

Navigation in the desert ant - Wehner 1996

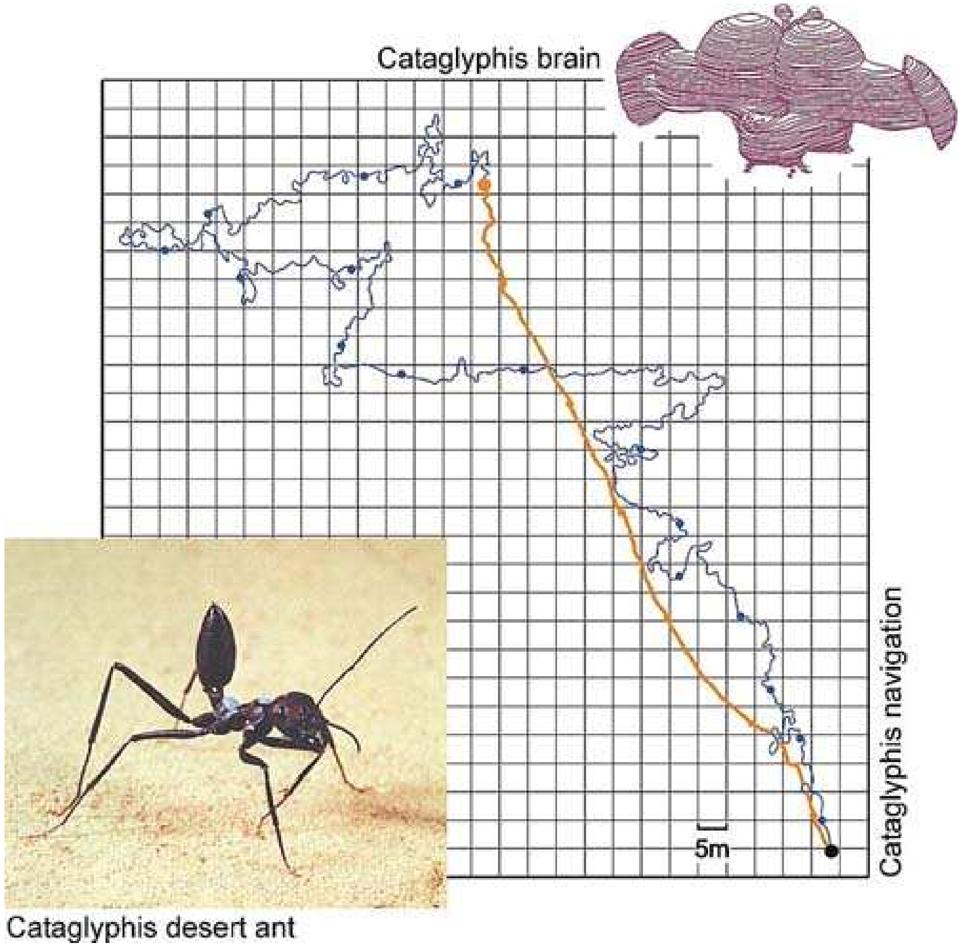

Polarised light compass

Desert ants (and many other animals) have visual receptors tuned to the polarisation plane of light.

Skylight has a natural polarisation pattern

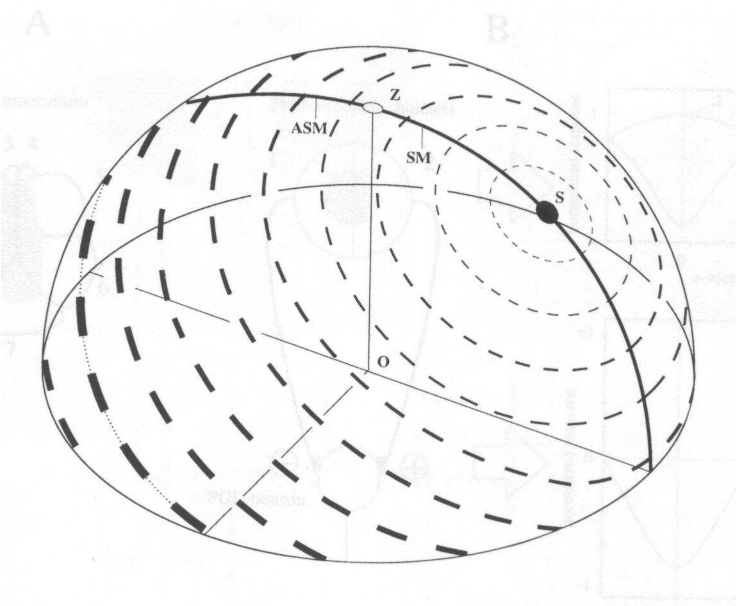

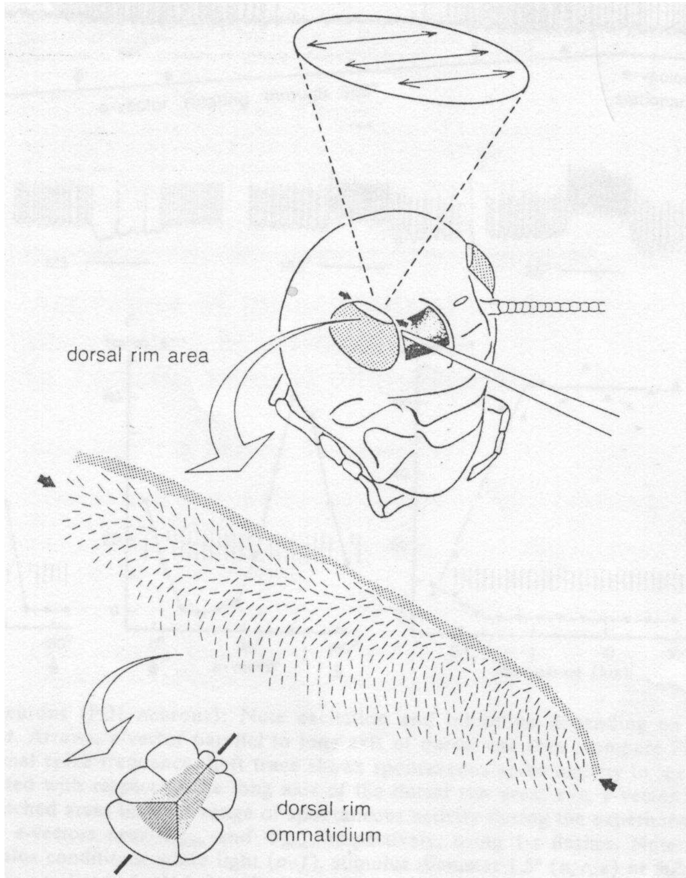

Measuring velocity

Ants counting steps: manipulating leg length results in different estimates of distance (Wittlinger et al 2006)

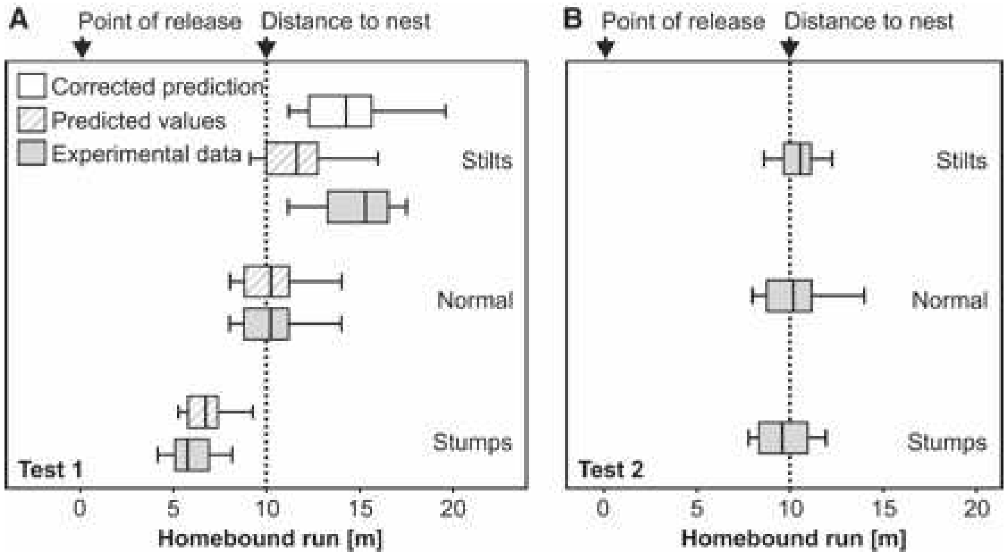

Ants also use visual homing

- In fact, ants show highly reproducible route memories, and can recognise where they are along the route from their visual surroundings (data from Michael Mangan)

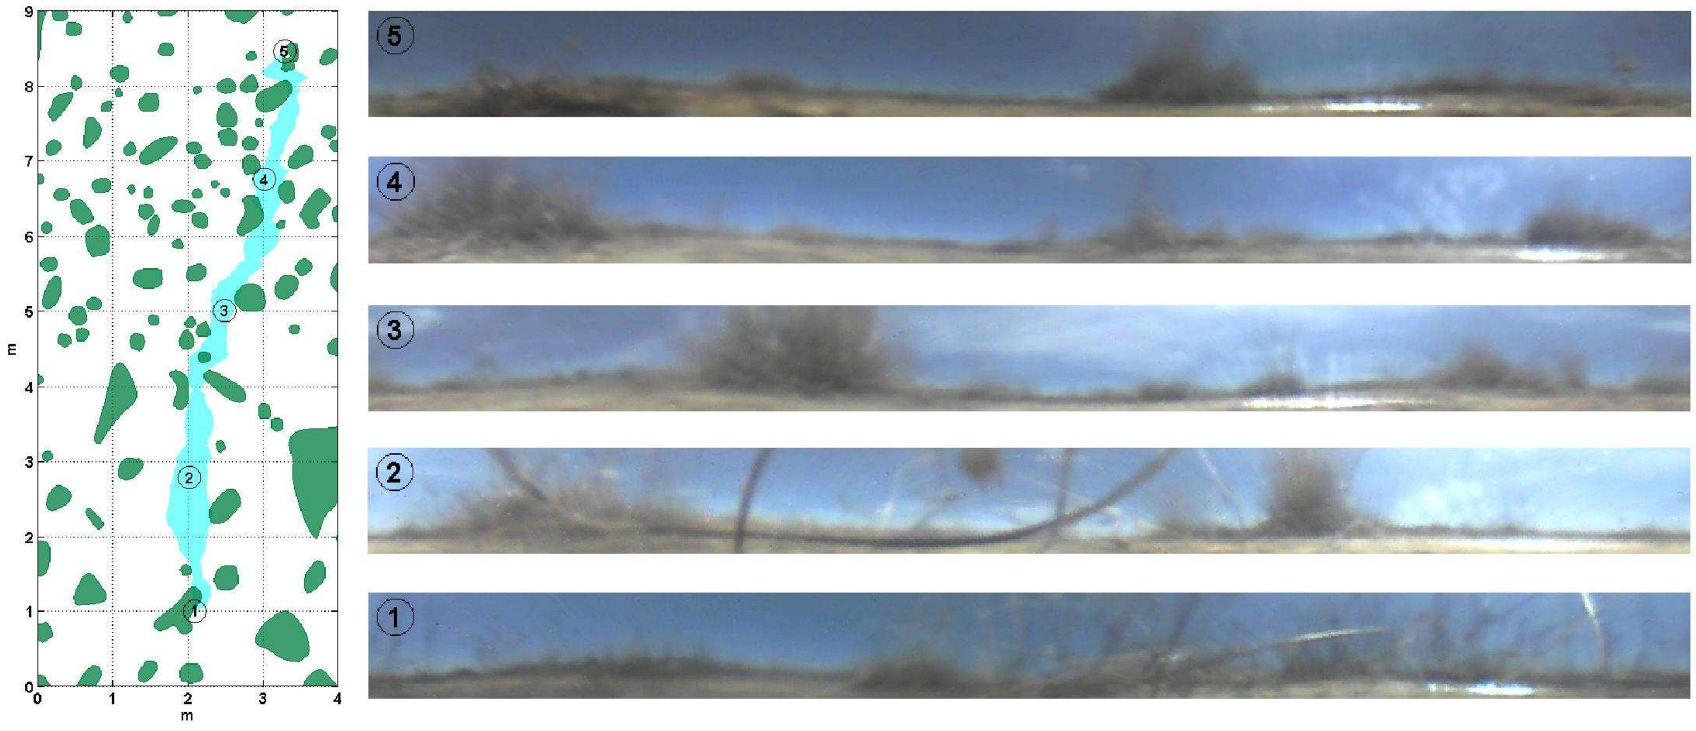

Each ant retraces its own unique route when displaced along it

Ant A: From home to feeder | Ant B: From feeder to halfway | Ant B: From home to halfway

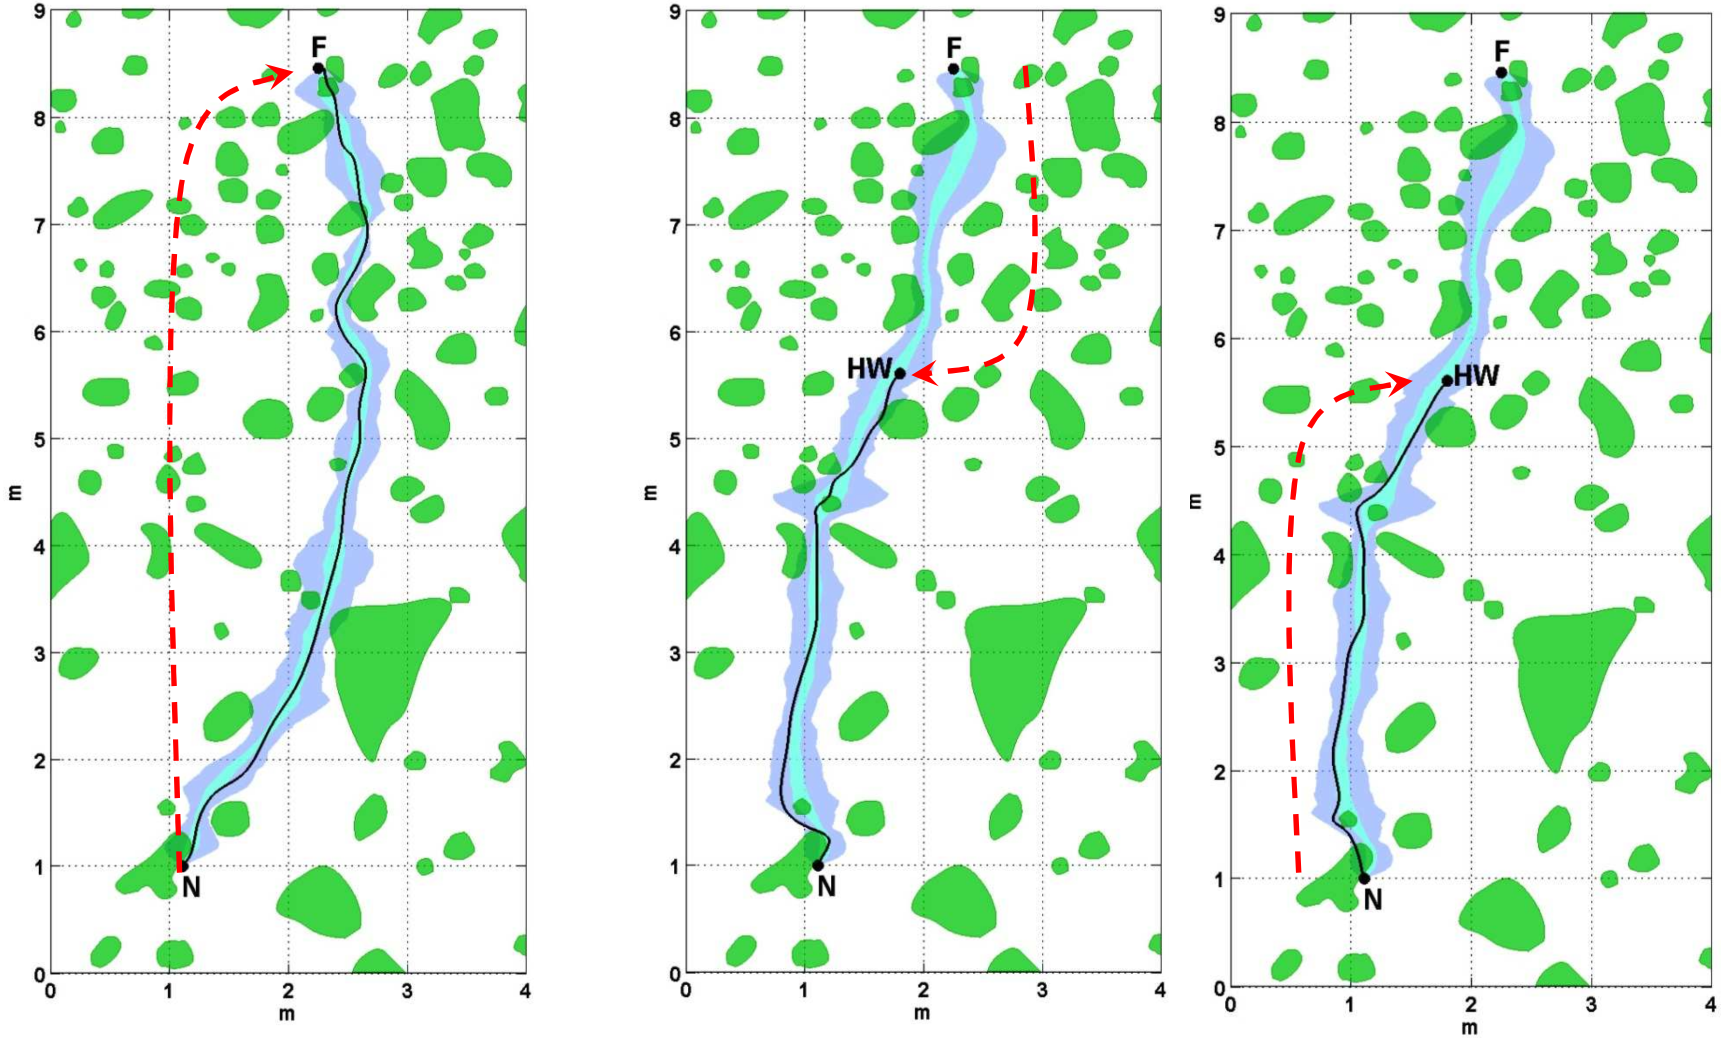

- How does ant recover from displacement?
  - Corresponds to classic ‘kidnapped robot’ problem to which the best solution is bayesian localisation
  - Can ants use priors to disambiguate landmarks?

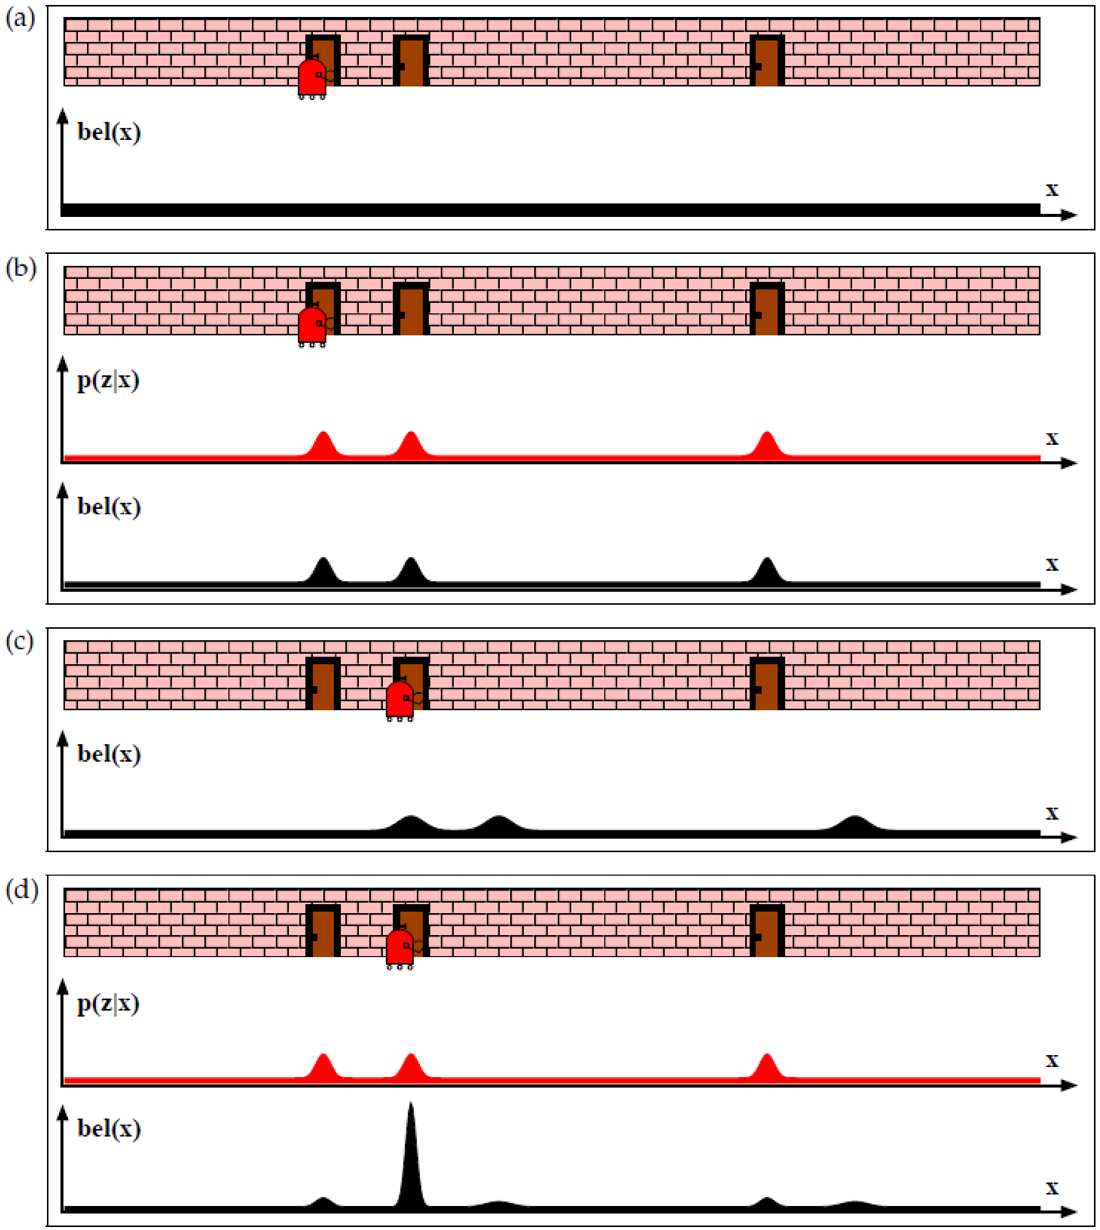

1. Robot starts with equal probability for every possible location x

2. Measurement z indicates robot is near a door: get three peaks in position estimate

3. Robot moves to the right: updates position estimate but becomes less certain

4. New measurement indicates robot is near door: this makes one possible position more likely than the others

- How does ant recover from displacement?
  - Corresponds to classic ‘kidnapped robot’ problem to which the best solution is bayesian localisation
  - Can ants use priors to disambiguate landmarks?
- Do ants optimally combine their path integration and landmark based estimates?
  - Can this be affected by manipulating the reliability of the different cues?

Example 3: Pavlovian conditioning in flies

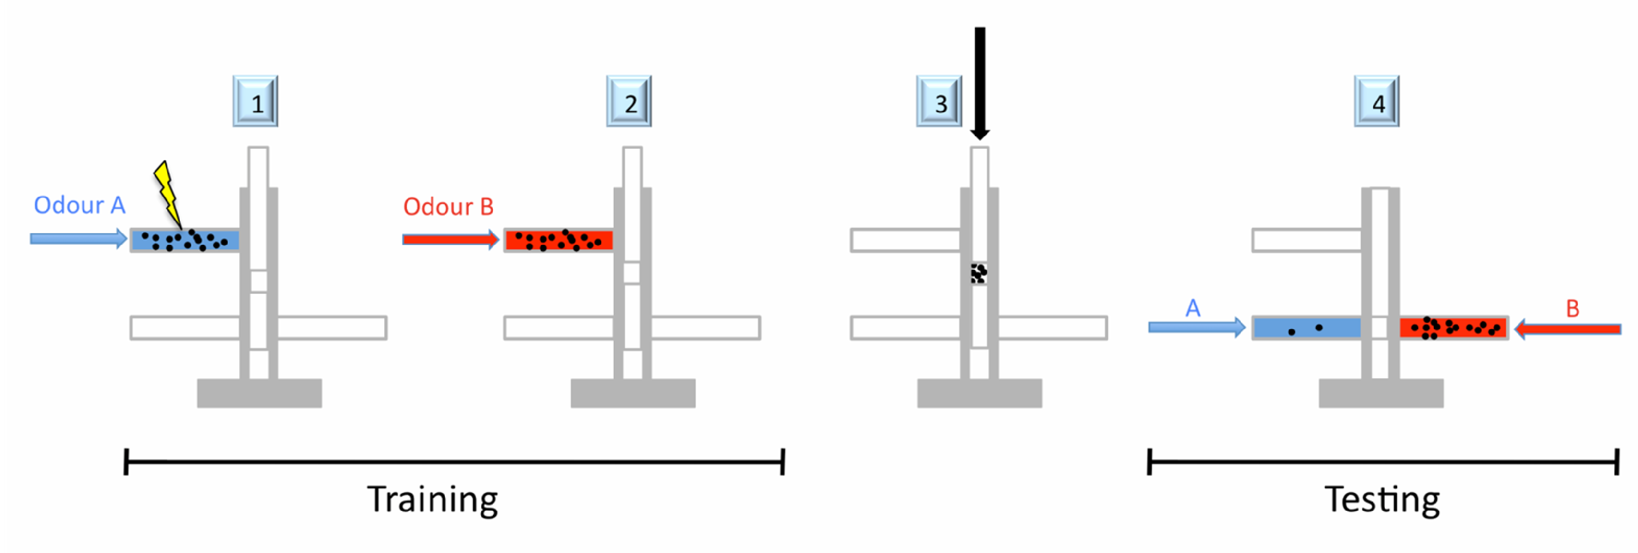

Memory trace for olfactory associations is located in the insect mushroom bodies

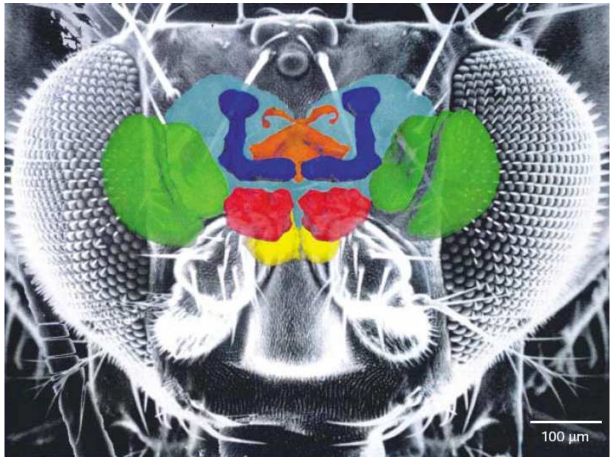

Heisenberg (2003)

Learning with multiple CS – configural stimuli

(Joanna Young, with Douglas Armstrong)

- Can flies learn AB+ CD- as easily as A+ B+? What about overlapping configurations AB+ BC- ?

XOR

- Can they learn parts vs. whole: AB+ A- B- or A+ B+ AB- ?
- What about contextual discrimination: AB+ CD+ AC- BD- ?
- Can they learn more than one thing: A+ B+ C- or A+ AB+ ?

Blocking

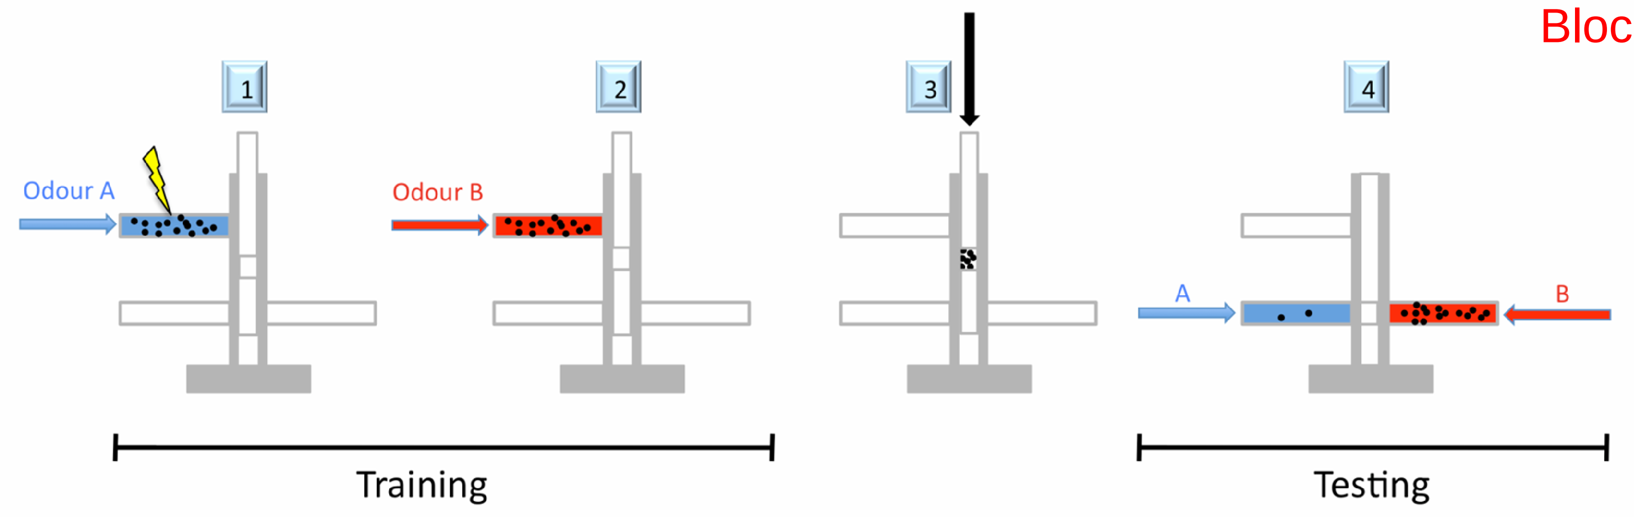

Bayesian accounts of associative learning

- Animal is learning *P(US(t)|CS(t),D(t-1,…,1))*
- Courville et al (2006) suggest that configural learning in particular is better described learning a generative model for

  P(US(t),CS(t)|D(t-1,…,1))

$$P(R \mid A-C, D) = \sum_m P(R \mid A-C, m, D) P(m, D)$$

$$P(m \mid D) \propto P(D \mid m) P(m)$$

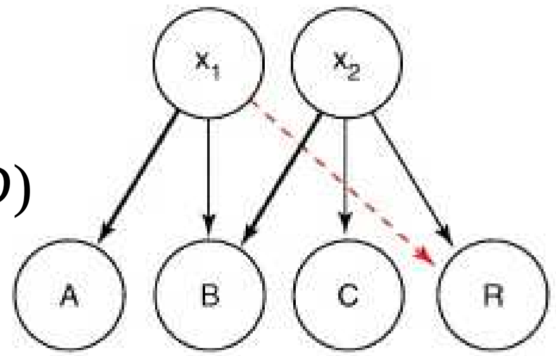

TRENDS in Cognitive Sciences

- Daw and Courville (2007) propose particle filter version in which individuals maintain one hypothesis, and either smoothly update or jump to new one depending on current inputs

Bayes in the brain

Several ideas: e.g. Ma et al 2006 – Poisson-like noise allow populations of neurons to represent distributions and support Bayesian inference by simple linear combination

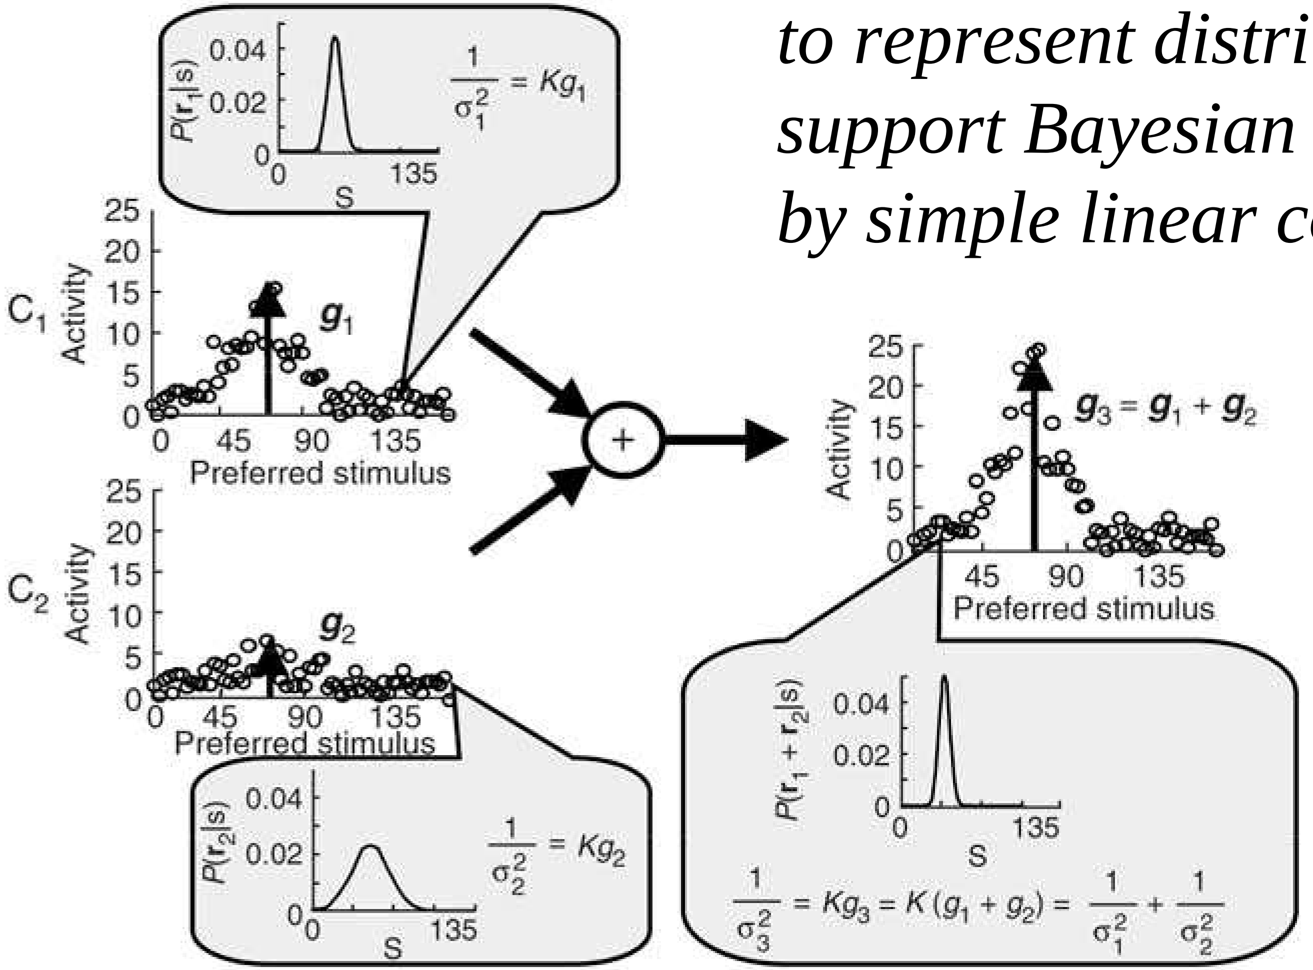

Bayes in the brain

"If one were to try and summarize all brain areas which have so far been mentioned as incorporating some aspect of predictive processing, these would include: unimodal sensory cortices, lateral and medial parietal and temporal areas, orbitofrontal, medial frontal and dorsolateral prefrontal cortex, premotor cortex, insula, cerebellum, basal ganglia, amygdala and thalamus….

…In other words, the whole brain." (Bubic et al 2010)

Bayes in the bee brain?

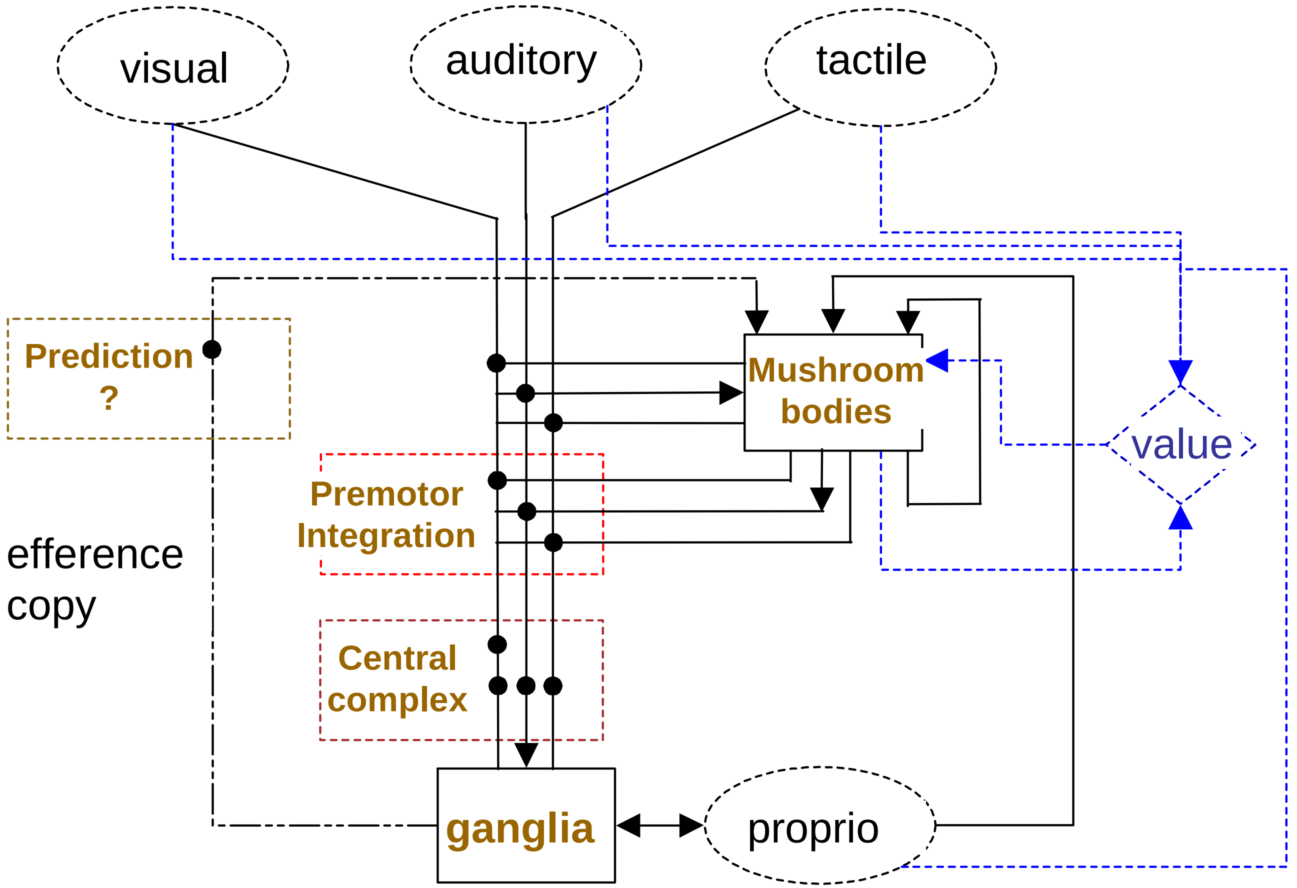

Do bees have Bayesian brains?

- Obviously, there is influence of innate or learned ‘priors’ on processing of current inputs
- In some cases this may involve non-trivial prediction, and/or ‘optimal’ merging of prediction and observation (still to prove?)
- In many cases there are alternative, plausible, specialised, simple solutions (can we distinguish cases when these don’t suffice?)
- Maybe bee (and human) brains are just collections of hacks after all…