Lecture 11. Bayesian inference

How the brain infers the world?
How we infers the happenings in brain?

- Bayes rule
- Bayesian inference
- Bayesian decoding of hippocampus
- Perception as Bayesian inference
- Illusion as Optimal Bayes

Bayes approach

- Alan Turing cracking the Engima code.
- Bernard Koopman narrowing down the search for German U-boats – Bayesian update.
- John Craven finding the lost Hydrogen bomb in Palomares B-52 crash incident.

Susan McGrayne: The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines, and Emerged Triumphant from Two Centuries of Controversy

Perception as an inference process based on evidence and priors

Olshausen & Field 1996

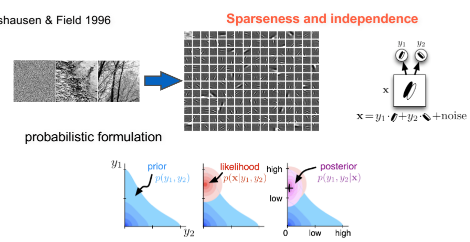

Bayes' Rule

$$P(H/E) = \frac{P(E \mid H) P_{prior}(H)}{\sum_{H'} P(E \mid H') P_{prior}(H')}$$

H = a particular hypothesis
E = evidence
H' = all the hypotheses considered.

Bayes rule:

Since $P(\omega_j, x) = P(\omega_j \mid x) P(x) = P(x \mid \omega_j) P(\omega_j)$

posterior $P(\omega_j \mid x) = \dfrac{P(x \mid \omega_j) P(\omega_j)}{P(x)}$

- where in case of two categories

$$P(x) = \sum_{j=1}^{j=2} P(x \mid \omega_j) P(\omega_j)$$

- Posterior (hypothesis given the data) = (
- Likelihood (data given the hypothesis) x Prior (hypothesis)) / Prior(Observation)

Bayes' Rule

$$P(\omega / x) \propto P(x / \omega) P(\omega)$$

models of the Environment

How data is generated by the environment

Prior belief – relative probability of what to believe in.

If one event (catching Salmon) happens more often than the other (catching Chicken) …

If $P(\omega_1)$=2/3, and $P(\omega_2)$=1/3,
without making an observation, how would I guess?

Now, if the observation indicates that x=12.2 inches, how would I decide? First find the likelihood p(x|w).

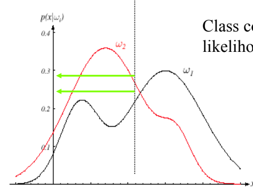

Class conditional or likelihood

$P(x=12.2\,|\,\omega_1)=0.22 \qquad P(x=12.2\,|\,\omega_2)=0.26$

x= measurement, observation, $\omega_1 = class1, \omega_2 = class2$

Bayes Decision (MAP) Rule

- Choose category *i* that has the maximum posteriori probability (MAP):

$$P(\omega_i\,|\,\mathrm{x})$$

- This produces the minimum *probability of error:*

$$p_e = 1 - P(chosen\ category\,|\,\mathrm{x})$$

Probability of making ALL the wrong choices.

Thomas Bayes 1702-1761

$$P(x)=\sum_{j=1}^{j=2} P(x\,|\,\omega_j)P(\omega_j)$$
$$=\frac{2}{3}\times 0.22+\frac{1}{3}\times 0.26 = 0.233$$

$$P(\omega_1\,|\,x=12.2)=\frac{P(x=12.2\,|\,\omega_1)P(\omega_1)}{P(x=12.2)}=\frac{0.1466}{0.233}=0.629$$
$$P(\omega_2\,|\,x=12.2)=\frac{P(x=12.2\,|\,\omega_2)P(\omega_2)}{P(x=12.2)}=\frac{0.0866}{0.233}=0.371$$

Which one would you choose?
What is the probability of error?

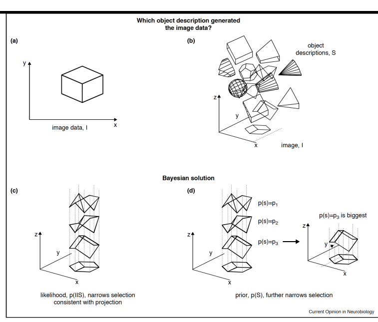

Illustration of Bayesian object perception **(a)** What 3D object caused the image of a cube? The likelihood (p(I|S)=p(image data|object descriptions)) constrains the possible set of objects to those consistent with the image data, but even this is an infinite set. **(b)** The prior knowledge probability (p(S)=p(object descriptions)) constrains the consistent set of 3D objects to those that are more probable in the world. **(c,d)** The probability over all instances is determined by the product of the likelihood and prior knowledge: p(S,I)=p(I|S)p(S)=p(object descriptions, image data). See the section on Basic Bayes for mathematical definitions. (Adapted from [64], © 1993 IEEE.)

Which object description generated the image data?

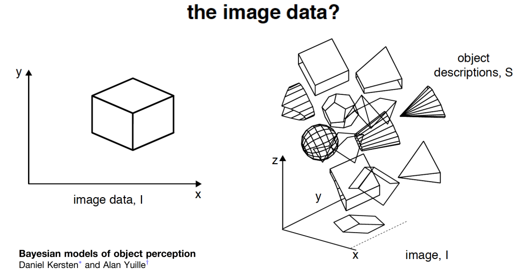

Bayesian models of object perception
Daniel Kersten[*] and Alan Yuille[†]
Current Opinion in Neurobiology (2003)

Which object description generated the image data?

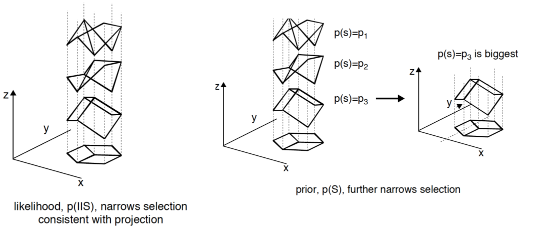

Current Opinion in Neurobiology

Motion Perception

Assuming small motion, the change in intensity is equal to the movement in the image gradient.

$$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v = -\frac{\partial I}{\partial t} \qquad \text{using} \quad u = \frac{dx}{dt}, v = \frac{dy}{dt},$$

Optical flow constraint

$$D(r) = \sum_{x,y} \left(\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \frac{\partial I}{\partial t}\right)^2$$

Velocity Space

- The optical-flow constraint: the flow velocity *(u,v)* has to lie along a straight line perpendicular to the brightness gradient.
- $$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v = -\frac{\partial I}{\partial t}$$
  $$(I_x, I_y) \cdot (u,v) = -I_t$$
- We can determine the component in the direction of the brightness gradient, but not the flow components in the direction at right angles.

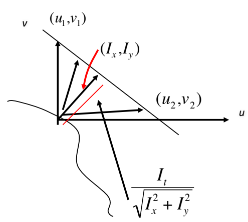

Resolving ambiguity in local motion estimates

Show aperture movie

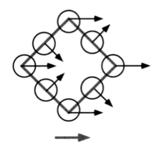

- How to integrate local cues?
  - intersection of constraints (IOC)
  - vector averaging (VA)
  - "blob" tracking

Show Yair Weiss Rhombus movie

What are the priors in motion perception?

Given the same evidence, do we prefer an interpretation rest on slow motion or fast motion? Do things tend to move fast or move slowly?

Show Ellipse.

Priors favoring slow speed:

$$P(v) \propto \exp(-\|v\|^2/2\sigma_p^2).$$

The posterior probability of a velocity was computed by combining the likelihood and prior using Bayes' rule. Because we assumed that the noise is independent over spatial location, the total likelihood function is just a product of likelihoods:

$$P(v|I) \propto P(v) \prod_i P(I(x_i,y_i,t)\,|v), \qquad (6)$$

How to encode uncertainty about the constraints?

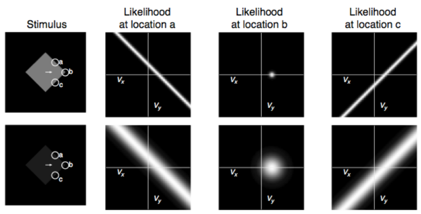

Posterior on velocity

$$P(v|I) \propto \exp\left(-\|v\|^2/2\sigma_p{}^2 - \frac{1}{2\sigma^2} \int_{x,y} \sum_i w_i(x,y)\, (I_x(x,y)v_x + I_y(x,y)v_y + I_t)^2 \,dx\,dy \right).$$

Here we assumed the entire image moves according to a single translational velocity, and so summed over all spatial positions. In this case, $\sum_i w_i(x,y)$ is a constant, so the posterior probability is given by:

$$P(v|I) \propto \exp\left(-\|v\|^2/2\sigma_p{}^2 - \frac{1}{2\sigma^2} \int_{x,y} (I(x,y)\, v_x + I_y(x,y)v_y + I_t)^2 \,dx\,dy \right)$$

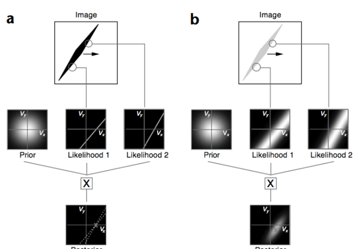

Yair Weiss[1], Eero P. Simoncelli[2] and Edward H. Adelson[3] **Motion illusions as optimal percepts**

nature *neuroscience* • volume 5 no 6 • june 2002

Hippocampus

Limbic system structure important in forming new memories and connecting emotion and senses (smell, sound) to memories. Primitive brain structure located on top of the brain stem, buried under the cortex, involved in emotion and motivation, feelings of pleasures (eating and sex).

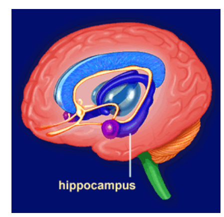

Hippocampus

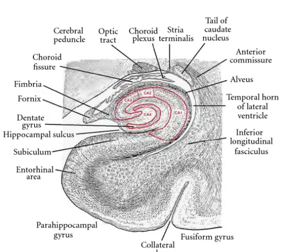

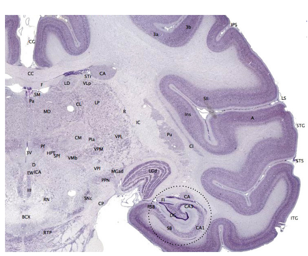

Hippocampus

- Connectivity
  - receives input from much of cortex as well as subcortical structures
  - projects back to these same structures
  - recurrent network in HC, what is it good for?

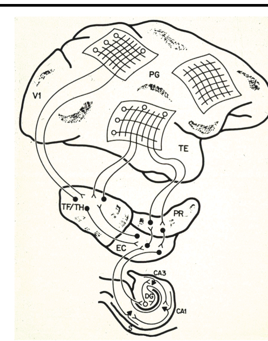

Place cells in the CA1,CA3 and dentate gyrus in rodent hippocampus

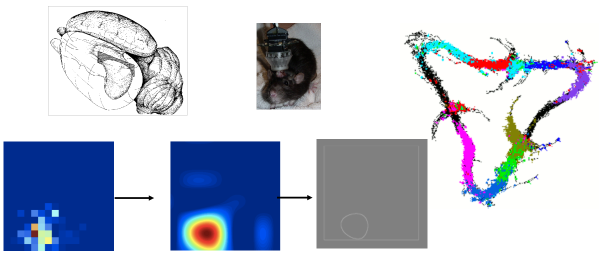

How to decode the location of the rat based on the ensemble activities?

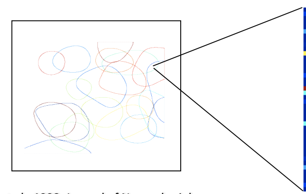

Zhang et al., 1998, Journal of Neurophysiology

Neural decoding:

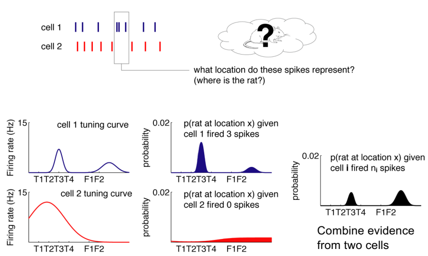

Naïve Bayes

- Bayes rule

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)}$$

$$X = (X_1, ..., X_N)$$

- Naïve Bayes assumes that Xi are conditionally independent, given Y

$$P(X|Y) = P(X_1, ..., X_N|Y) = \prod_{i=1}^{N} P(X_i|Y)$$

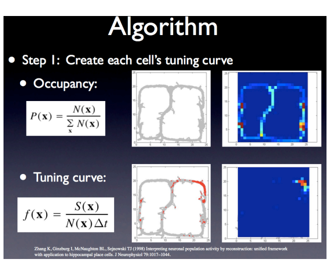

N(x) number of discrete time steps the rat spend at each location.

Δt = the duration (ms) of each discrete time step.

f(x) = average spike counts per unit time at each location.

S(x) spike counts at location x.

n= number of spikes within a time window.

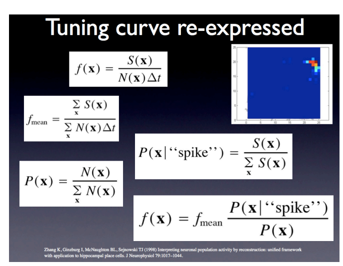

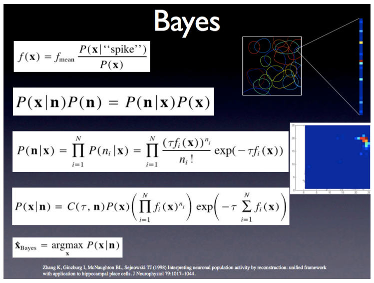

f_{mean} = mean firing rate of the neurons across all locations.

n= number of spikes within a time window.

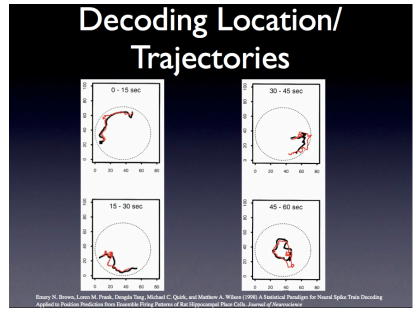

Is there replay during awake?

- During sleep and inattentive pauses in the environment, animals replay sequences of neural activity representing past experiences. Thought to reflect the process of consolidation.

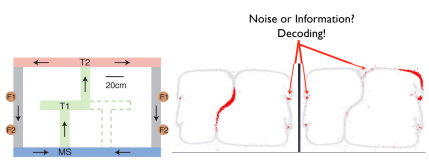

Is there replay during awake?

What is about to be experienced in the forward order (Diba and Buzsaki, 2007)

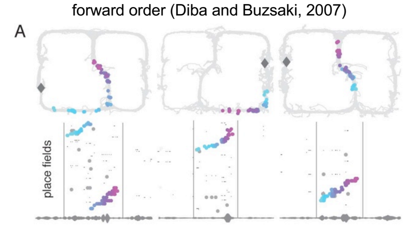

forward

Is there replay during awake?

What was just experienced in the reverse order (Foster and Wilson, 2006)

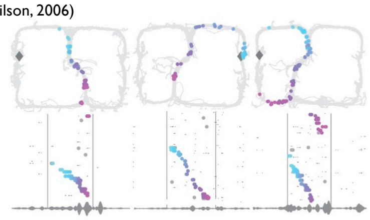

backward

Decoding cognition (sweeps)

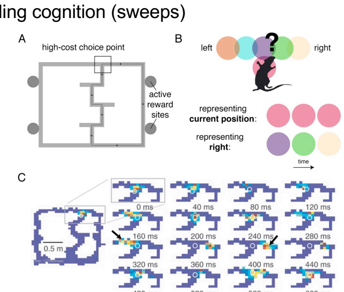

David Redish