## Advanced conditioning: generative models

Nathaniel Daw
New York University
IPAM Summer School, July 19

## Overview

- Background and philosophy
- Multiple cues and structural assumptions
  - – associativity models
  - – unblocking
- Structure learning
  - – configural conditioning
  - – evidence/complexity tradeoffs

## Levels of analysis

Marr's (1982) famous hierarchy:

Computation
interpretation: why?

Algorithm

Implementation
simulation: how?

## Levels of analysis

Marr's (1982) famous hierarchy:

Computation
interpretation: why?

(+ priors!)

Algorithm

pretend exact, for now.
much more to say about this.

Implementation
simulation: how?

## Why study conditioning?

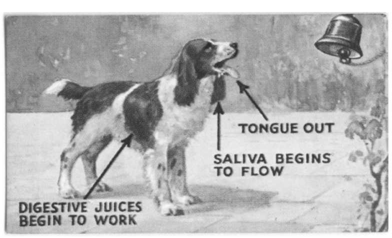


- well developed tasks, data & (nonstatistical) theory
- pure view on issues of inference, evidence, uncertainty, learning
- neural & statistical substrates now within grasp

## Cue combination

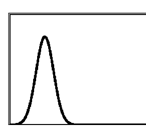

$P(r \mid A) = N(w_A, \sigma)$

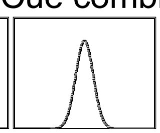

$P(r \mid B) = N(w_B, \sigma)$

what is $P(r \mid AB)$?

# Cue combination

Kalman filter: $P(r \mid AB) = N(w_A + w_B, \sigma^2)$

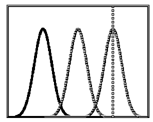

# Cue combination

Kalman filter: $P(r \mid AB) = N(w_A + w_B, \sigma^2)$

product of experts: $P(r \mid AB) \propto N(w_A, \pi_A^{-1})\, N(w_B, \pi_B^{-1})$
$= N(\pi_A\sigma^2 w_A + \pi_B\sigma^2 w_B, \sigma^2)$
$\sigma^2 = (\pi_A + \pi_B)^{-1}$

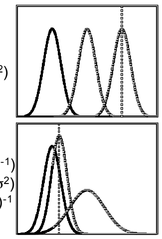

# Cue combination

Kalman filter: $P(r \mid AB) = N(w_A + w_B, \sigma^2)$

product of experts: $P(r \mid AB) \propto N(w_A, \pi_A^{-1})\, N(w_B, \pi_B^{-1})$
$= N(\pi_A\sigma^2 w_A + \pi_B\sigma^2 w_B, \sigma^2)$
$\sigma^2 = (\pi_A + \pi_B)^{-1}$

additive mixture: $P(r \mid AB) \propto \pi_A N(w_A, \sigma^2) + \pi_B N(w_B, \sigma^2)$

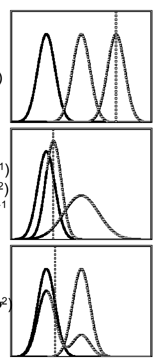

# Cue combination

E step:
$q_A \propto \pi_A \exp((r - w_A)^2 / \sigma)$

M step (partial):
$w_A \leftarrow w_A + \eta q_A (r - w_A)$
$\pi_A \leftarrow \pi_A + \eta(q_A - \pi_A)$

Expected $r$ given $AB$:
$\pi_A w_A + \pi_B w_B$

additive mixture: $P(r \mid AB) \propto \pi_A N(w_A, \sigma^2) + \pi_B N(w_B, \sigma^2)$

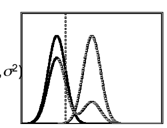

# Cue combination

E step:
$q_A \propto \pi_A \exp((r - w_A)^2 / \sigma)$

M step (partial):
$w_A \leftarrow w_A + \eta q_A (r - w_A)$
$\pi_A \leftarrow \pi_A + \eta(q_A - \pi_A)$

responsibility gates learning

stimuli seek to predict whole outcome (compare RW: $w_A \leftarrow w_A + \eta(r - w_A - w_B)$)

Expected $r$ given $AB$:
$\pi_A w_A + \pi_B w_B$

prediction also gated

additive mixture: $P(r \mid AB) \propto \pi_A N(w_A, \sigma^2) + \pi_B N(w_B, \sigma^2)$

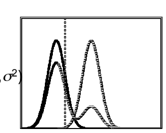

# Notes

- Resemble CS-processing theories from conditioning (Mackintosh, Pearce & Hall)
- Responsibility a form of selective attention in learning and prediction – but arises from optimal inference, not limited capacity
- Lots of work involving surprise, dimensional attention, etc.

(see Dayan & Long 1997; Kruschke 2001; Kakade & Dayan 2002; Courville et al 2006)

## One example

Blocking:

| phase 1 | phase 2 | test |
|---|---|---|
| A→r | AB→r | B ? (attenuated) |

RW: $w_B \leftarrow w_B + \eta(r - w_A - w_B)$
explains blocking by lack of prediction error

AM: $w_B \leftarrow w_B + \eta q_B(r - w_B)$
explains blocking by lack of attention

## One example

Blocking:

| phase 1 | phase 2 | test |
|---|---|---|
| A→r | AB→r | B ? (attenuated) |

Downward unblocking:

| phase 1 | phase 2 | test |
|---|---|---|
| A→rr | AB→r | B ? (restored) |

RW: $w_B \leftarrow w_B + \eta(r - w_A - w_B)$
predicts inhibitory conditioning to *B*

AM: $w_B \leftarrow w_B + \eta q_B(r - w_B)$
predicts excitatory conditioning to *B*

## Downward unblocking

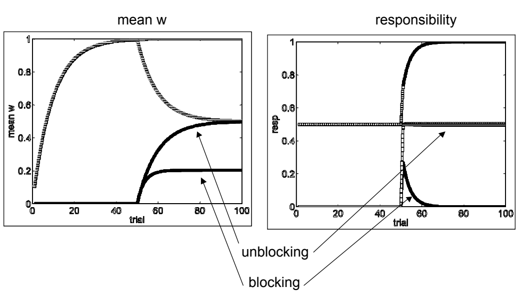


## Interim summary

- Differing structural/generative assumptions underlie difference between psychological models of conditioning
  - none wholly satisfactory though
- "Attention" related to responsibility / uncertainty
- Next: learned generative structure (work with Aaron Courville, Dave Touretzky)

# Cue combination, 2

Idea ("configural" conditioning, Pearce 1994):

net prediction should be function of stimulus configuration as a unit; smooth generalization between similar configurations

contrast to R/W ("elemental"):

net prediction to conjunction is function (sum) of elements' predictions*

(* except not really since this is untenable)

# Cue combination, 2

Idea ("configural" conditioning, Pearce 1994):

net prediction should be function of stimulus configuration as a unit; smooth generalization between similar configurations

contrast to R/W ("elemental"):

net prediction to conjunction is function (sum) of elements' predictions*

(* except this was immediately relaxed as untenable)

# Rescorla/Wagner

- Obvious problem with nonlinear discriminations

(eg XOR: A, B rewarded separately but not together)

Results

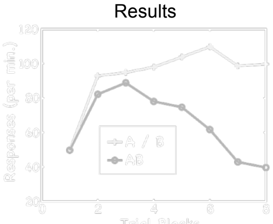

R/W model

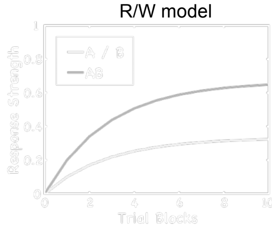

# Configurations

- Solved by adding "configural units"
- ie keep linear prediction rule
- Which units to add and when?
  (No serious answer)

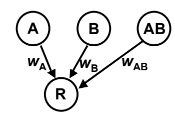


- Net prediction?

Configural theory (Pearce 94): Graded generalization

observe *AB*, predict: 1 $w_{AB}$ + 0.5 $w_A$ + 0.5 $w_B$

"Added elements" RW (WR 72): Binary

observe *AB*, predict: 1 $w_{AB}$ + 1 $w_A$ + 1 $w_B$

# Summation

Train:
- *AB*→reward
- *CD*→reward

Test:
- *AB*; *CD* (trained)
- *AC*; *BD* (transfer)
- *A*; *B*; *C*; *D* (elements)

→ summation, but sublinear (why do things this way?)

AE Rescorla/Wagner model

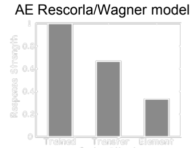

Results (Rescorla 03)

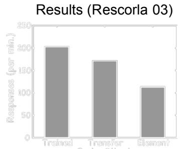

Pearce model

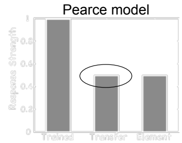

# Generalization

train:
- *A*→reward;
- *BC*→reward;
- *ABC*→nothing

generalization follows overlap?

Rescorla/Wagner model

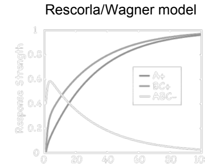

Results (Redhead & Pearce 94)

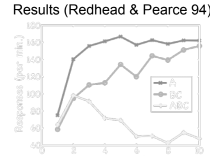

Pearce model

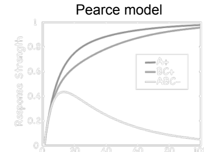

# Where are we

- No normative reason to prefer either account
  - Plus ambivalent empirical guidance
- Key issue: similarity vs. discrimination
  - How do animals respond to novel stimulus patterns; how do they discriminate between overlapping patterns?
- We recognize this as a tradeoff between generalization and data-fitting

# A new model

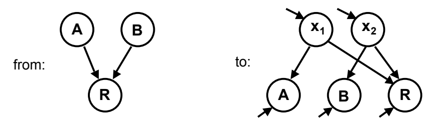


- Sigmoid belief network $P(stim_i \mid x) = (1+\exp(w_i \cdot x))^{-1}$
  - Model full joint distribution $P(A, B, R)$
    - Generative view; animals clearly learn this
  - configural units → latent causes
    - Actual structure of experiments
  - Learn weights + model structure
    - Which ‘configural units’ to include when
    - cf human causal learning (Griffiths et al)

# Uncertainty in model structure

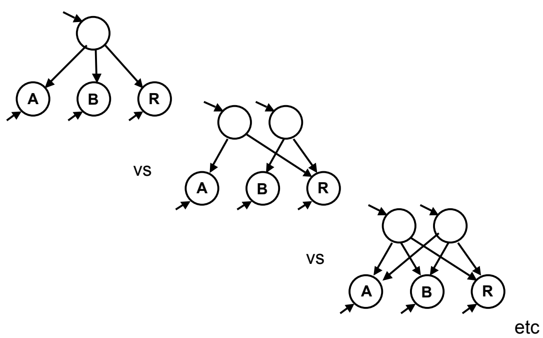


# Learning & prediction

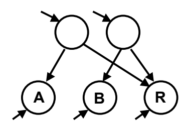


- Learning: $P(\mathbf{w}, M \mid data)$
  - Bayes’ rule $P(\mathbf{w}, M \mid data) \propto P(data \mid \mathbf{w}, M) P(\mathbf{w}, M)$
- Prediction: $P(r \mid stim, data)$
  - Marginalize uncertain weights, model, hidden activation

$$P(R \mid stim, data) = \sum_M \int d\mathbf{w}\, P(R \mid stim, M, \mathbf{w}) P(\mathbf{w} \mid M, data) P(M \mid data)$$

# Simplicity vs fidelity

- Prior preference for simple model structure, small weights
  - this arises from both prior and likelihood terms
- Progressive shift toward more accurate accounts as data accumulate
  - we’re used to this with weights
- Signature of Bayesian account: Tradeoff between model complexity and data fidelity

# Generalization

- Generalization between similar stimulus patterns → inference over latents

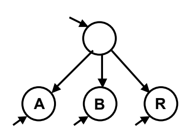

## Generalization

- Generalization between similar stimulus patterns → inference over latents

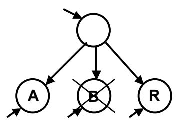


## Generalization

- Generalization between similar stimulus patterns → inference over latents

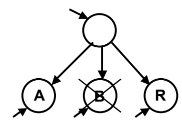


- Model learns whether to represent data as configuration or sum of subconfigurations

## Generalization

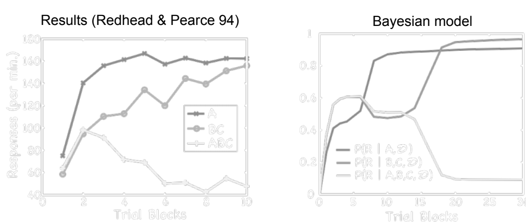


## Generalization

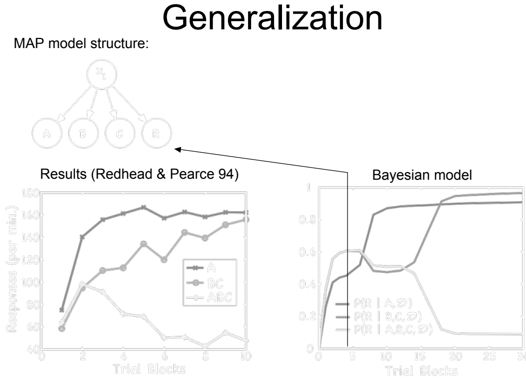


## Generalization

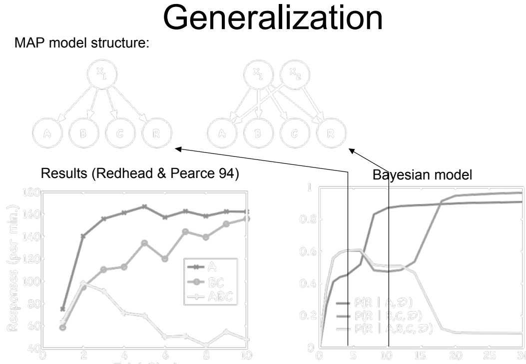


## Generalization

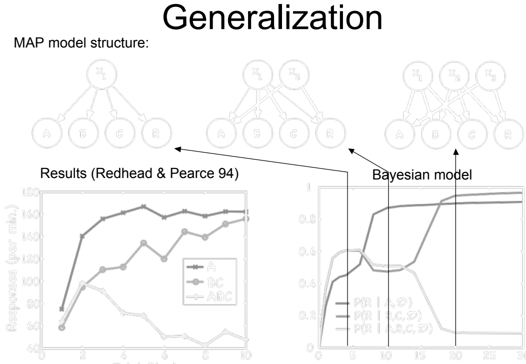

# Summation

- Train:
  - A,B→reward
  - C,D→reward
- Test:
  - A,B; C,D (trained)
  - A,C; B,D (transfer)
  - A; B; C; D (elements)

MAP model structure:

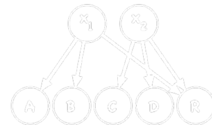

Results (Rescorla 03)

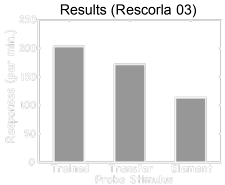

Bayesian model

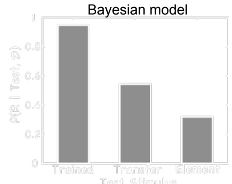

# Fit/complexity tradeoffs

- Different amounts of the same experience can produce opposite results
- Train: (Yin et al. 94)
  - A→reward
  - AB→nothing ⟵ few or many trials
- for few AB trials: “Second-order conditioning”
  - Test: B predicts reward
- for many AB trials: “Conditioned inhibition”
  - Test: B predicts absence of reward
  - e.g. cancels reward prediction when combined with a novel excitor C

# SOC vs CI

- Train:
  - A→reward, C→reward
  - AB→nothing
- Test:
  - B (second order)
  - BC (inhibition)

Bayesian model

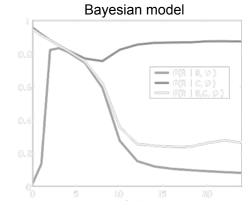

# SOC vs CI

- Train:
  - A→reward, C→reward
  - AB→nothing
- Test:
  - B (second order)
  - BC (inhibition)

MAP model structure:



Bayesian model

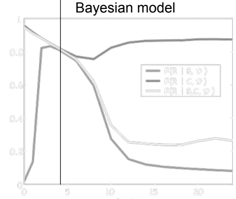

# SOC vs CI

- Train:
  - A→reward, C→reward
  - AB→nothing
- Test:
  - B (second order)
  - BC (inhibition)

MAP model structure:



Bayesian model

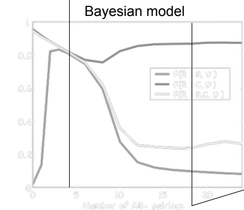

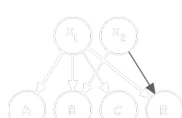

# Other issues

- Explore the priors – experimentally manipulable
- Relaxing assumptions
  - Independence between trials (change in model space; Kruschke)
  - Event timing within trial (Miller, Courville)
- Tractable/realistic approximation and implementation
- Neural substrates

# Implementation

There is no obvious tractable solution; no summary statistics etc.

- Seems unrealistic for subjects to maintain a full distribution over structural possibilities

Anonymous postdoctoral advisor:

- "Where in the brain is the reversible jump Markov chain Monte Carlo area?"

Anonymous NIPS reviewer:

- "The work is very solid. The general impression, however, is that the authors used a grenade to kill a fly, and in addition had to position the fly carefully."

# Implementation

- Perhaps sampling is a reasonable idea after all

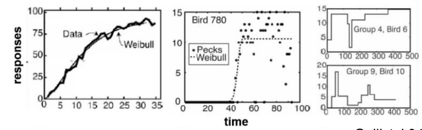


Gallistel 04

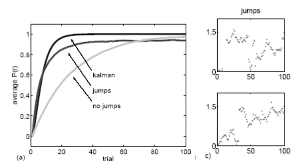


Daw & Courville 07

# Conclusions

- Generalization and discrimination recognized as the tradeoff between simplicity and data fidelity in Bayesian inference
  - Normative reinterpretation of configural conditioning
- Inference with latent variables is a natural account of classical conditioning situations
  - Interestingly different from associative learning
  - Well fit to standard experimental situation (in general, depends on situation and cover story)
  - see also Griffiths/Tenenbaum/Gopnik, Blaisdell, Cheng, etc
- Understand conditioning in Bayesian terms (priors); uncover evidence about important inferential issues
- from computational to algorithmic