Neuromodulators and the Neural Representation of Uncertainty

Peter Dayan

Gatsby Computational Neuroscience Unit
University College London

Sham Kakade   Angela Yu

Computational Neuromodulation

- general: excitability, signal/noise ratios
- specific: prediction errors, uncertainty signals

Learning and Inference

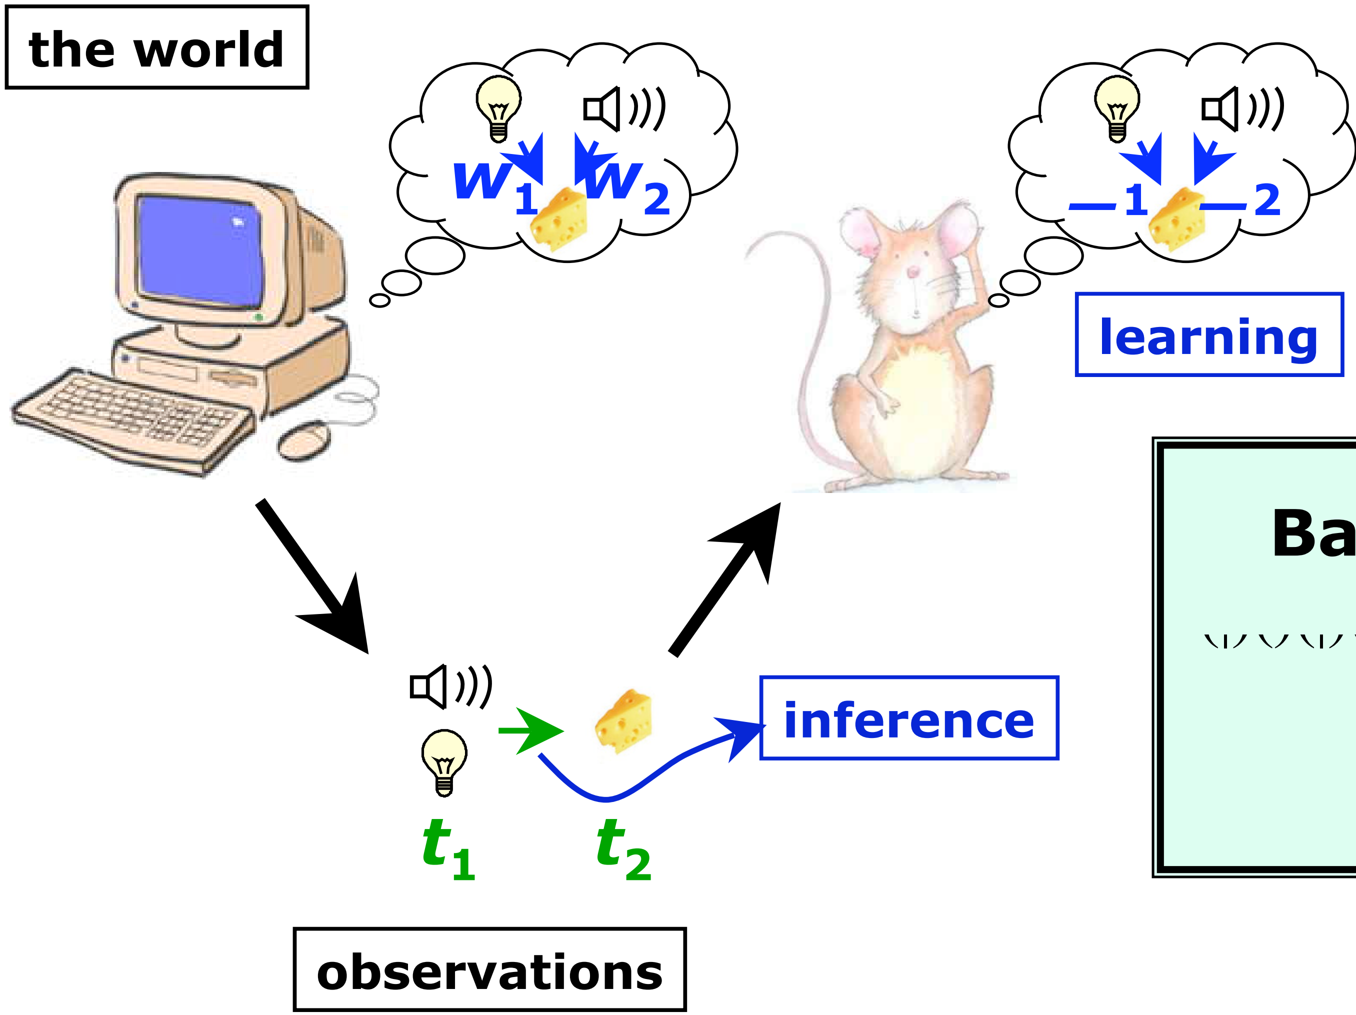

Bayes' Rule

Bayesian Uncertainty

- arises from:
  - stochasticity
  - ignorance
    - transient
    - possibility of ongoing change
- affects:
  - learning (the more the merrier)
  - inference (the less the merrier)

Experimental Data

ACh & **NE** have similar ***physiological*** effects

- ***suppress*** recurrent & feedback processing
  (*e.g.* Kimura *et al*, 1995; Kobayashi *et al*, 2000)
- ***enhance*** thalamocortical transmission
  (*e.g.* Gil *et al*, 1997)
- ***boost*** experience-dependent plasticity
  (*e.g.* Bear; Shulz; Merzench)

ACh & **NE** have distinct ***behavioral*** effects:

- ACh ***boosts*** learning to stimuli with uncertain consequences (*e.g.* Bucci, Holland, & Gallagher, 1998)

- NE ***boosts*** learning upon encountering global changes in the environment (*e.g.* Devauges & Sara, 1990)

Learning

- Learning: predict; control

Δ weight α (error) x (learning rate) x (stimulus)

- dopamine
  phasic prediction error for future reward
- serotonin
  phasic prediction error for future punishment

- acetylcholine
  expected uncertainty boosts learning
- norepinephrine
  unexpected uncertainty boosts learning

Conditioning

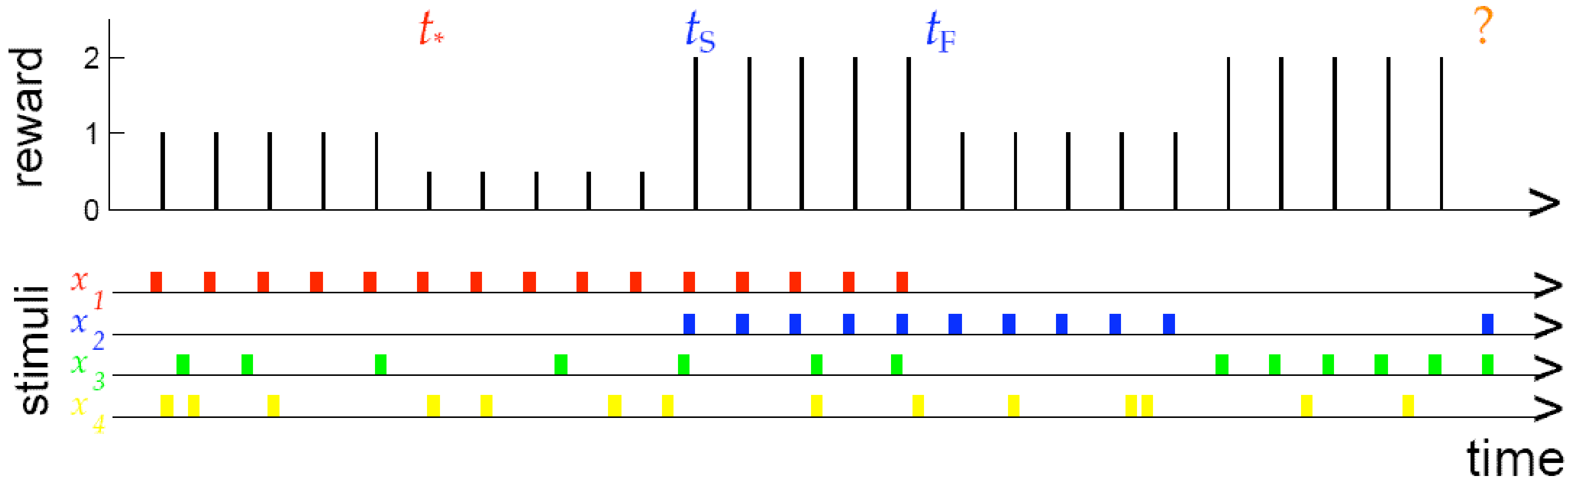

reward given $$r = \mathbf{w} \cdot \mathbf{x} + \epsilon$$

prediction weights $$\mathbf{w}' = \mathbf{w} \qquad + \eta$$

where

output noise $$\epsilon \sim N[0, \rho^2]$$

allowable drift $$\eta \sim N[0, \sigma^2 \mathbb{I}]$$

Single Stimulus

Intuition

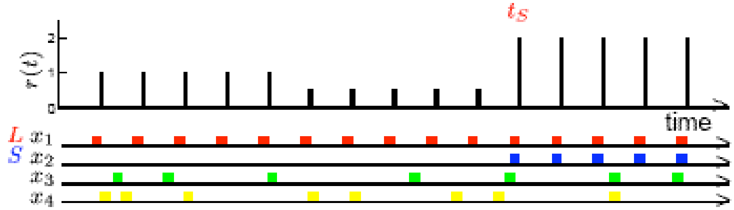

window model: w_L is the *average* of last three rewards:

mean goes like

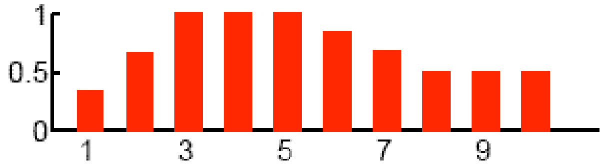

standard deviation goes like

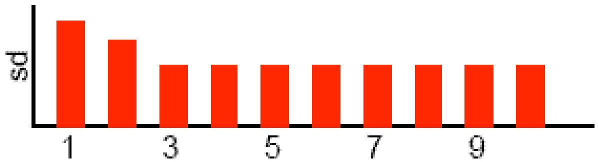

- more data $\Rightarrow$ more certain
- use of 3 rewards controls speed of adaptation
- asymptotic sd shouldn't ignore changes in r

Multiple Stimuli

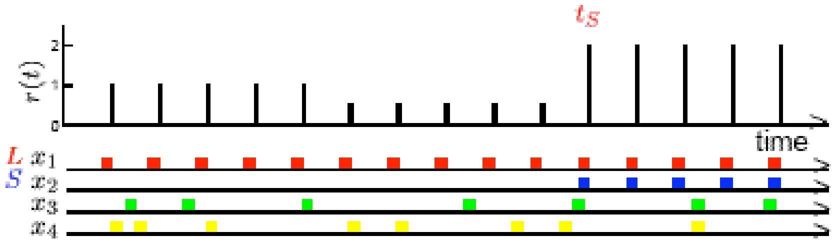

Competition for prediction error between L and S:

sound should 'win' since less well established

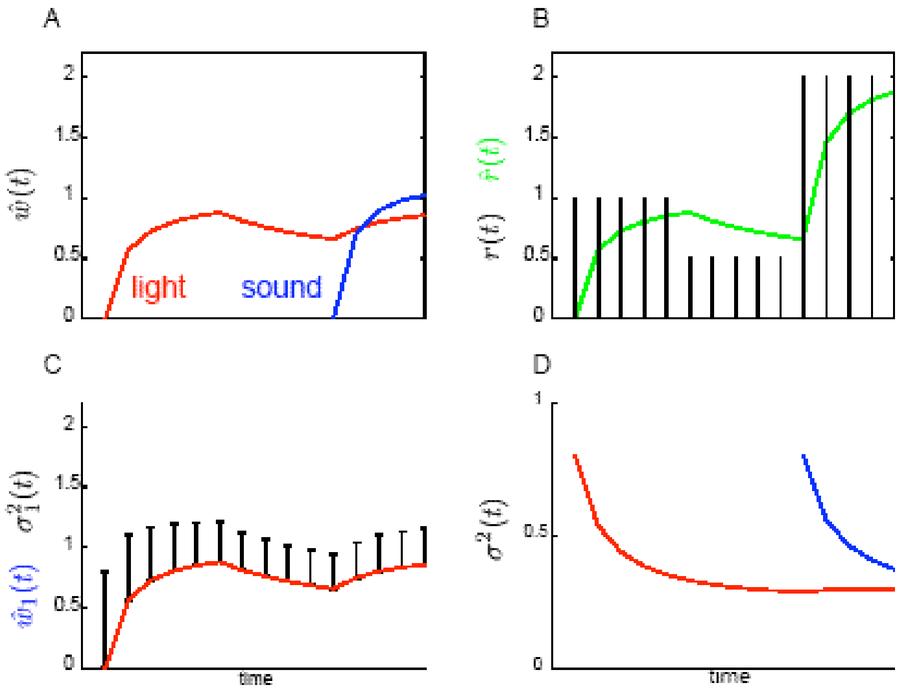

Formally: Kalman Filter

$$\begin{aligned} \text{observation:} \quad r &= \mathbf{w} \cdot \mathbf{x} + \epsilon \\ \text{state:} \quad \mathbf{w}' &= \mathbf{w} \quad\;\; + \eta \end{aligned}$$

$$\mathbf{w} \sim N[\hat{\mathbf{w}}, \Sigma]$$

$$\delta = r - \hat{\mathbf{w}} \cdot \mathbf{x}$$

$$\hat{\mathbf{w}}' = \hat{\mathbf{w}} + \delta \frac{\Sigma \cdot \mathbf{x}}{\rho^2 + \mathbf{x} \cdot \Sigma \cdot \mathbf{x}}$$

$$(\mathbf{w}' = \mathbf{w} + \delta\alpha\mathbf{x} \quad \text{delta rule})$$

- like the delta rule, bar compression and rotation
- compression is competition for learning

$$\Sigma_{ii} / \left(\rho^2 + \sum_j \Sigma_{jj}\right)$$

- ρ^2 sets learning rate according to noise (KKT)
- rotation allows backwards blocking
- update for Σ

$$\Sigma' = \Sigma + \sigma^2 \mathbb{I} - \frac{\Sigma \cdot \mathbf{x} \cdot \mathbf{x} \cdot \Sigma}{\rho^2 + \mathbf{x} \cdot \Sigma \cdot \mathbf{x}}$$

Δ weight α (error) x (learning rate) x (stimulus)

ACh and Learning

- Holland; Gallagher showed

  amygdala CEN $\longrightarrow$ nBM $\xrightarrow{\text{ACh}}$ parietal cortex

  to be critically involved in boosted learning
- hippocampal/ACC ACh involved in suppressed learning

Table 1. Outline of procedures for Experiment 1

Treatment condition (groups)	Phase 1: consistent L–T relation	Phase 2: experimental change in L–T relation	Phase 3: test of conditioning to L
Consistent (CTL-C, PPC-C)	L → T → food; L → T	L → T → food; L → T	L → food
Shift (CTL-S, PPC-S)	L → T → food; L → T	L → T → food; L	L → food

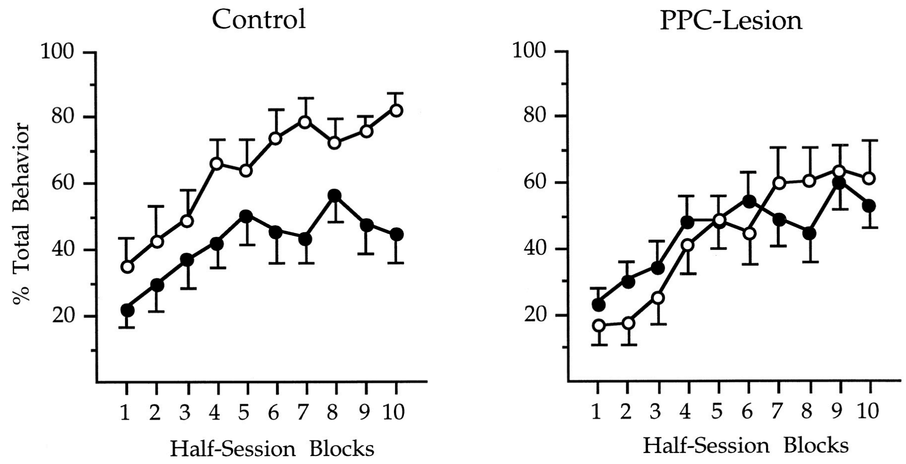

(Bucci, Holland, & Galllagher, 1998)

expected uncertainty

Learning and Inference

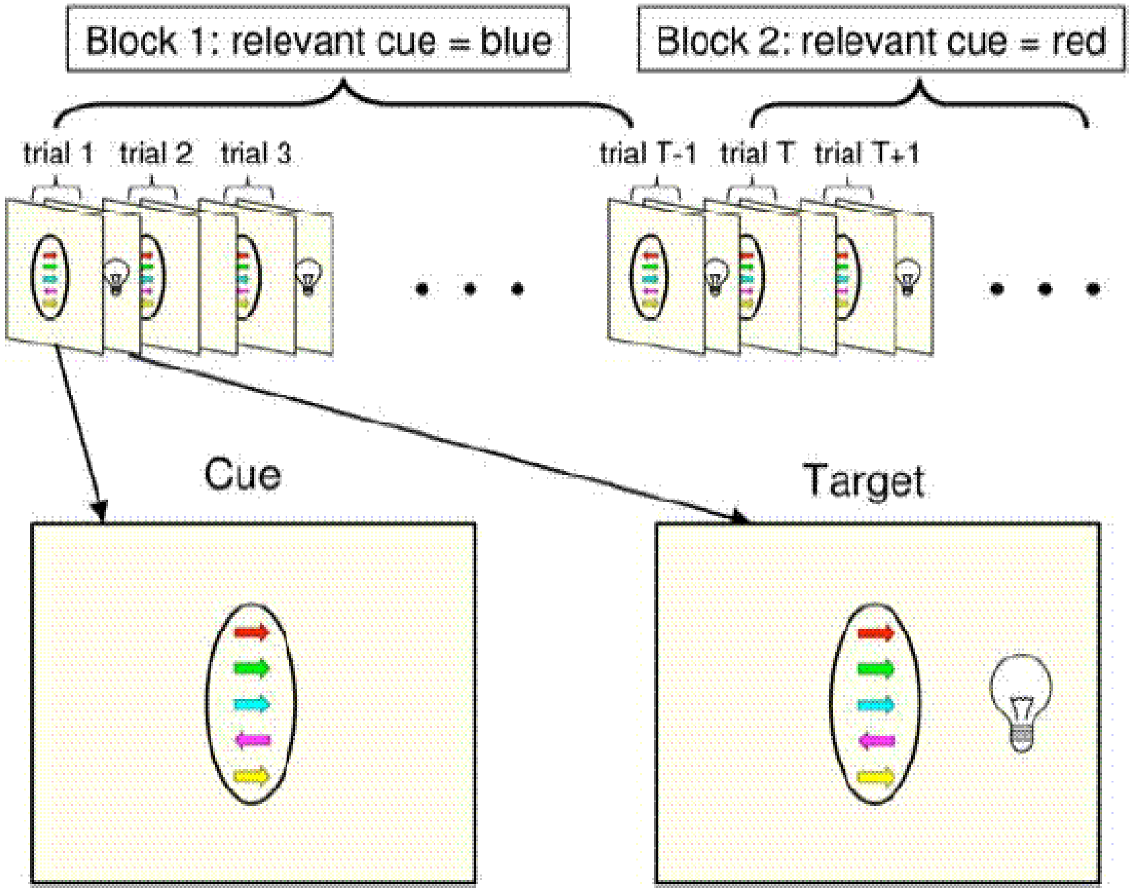

- Posner task with unsignalled cue/validity changes

  ACh

- `Reversal' task with stable validities

  NE

Formal Framework

NE | **ACh**

variability in **identity** of relevant cue

$1-\lambda_t^*$

variability in **quality** of relevant cue

$1-\gamma_t^*$

cues: vestibular, visual, ...

$$\mu_t^* = i$$

$$P^*(\mu_t^* \mid D_t) = \lambda_t^*$$

$$P^*(\mu_t = j \neq i \mid D_t) = \frac{1-\lambda_t^*}{h-1}$$

target: stimulus location, exit direction...

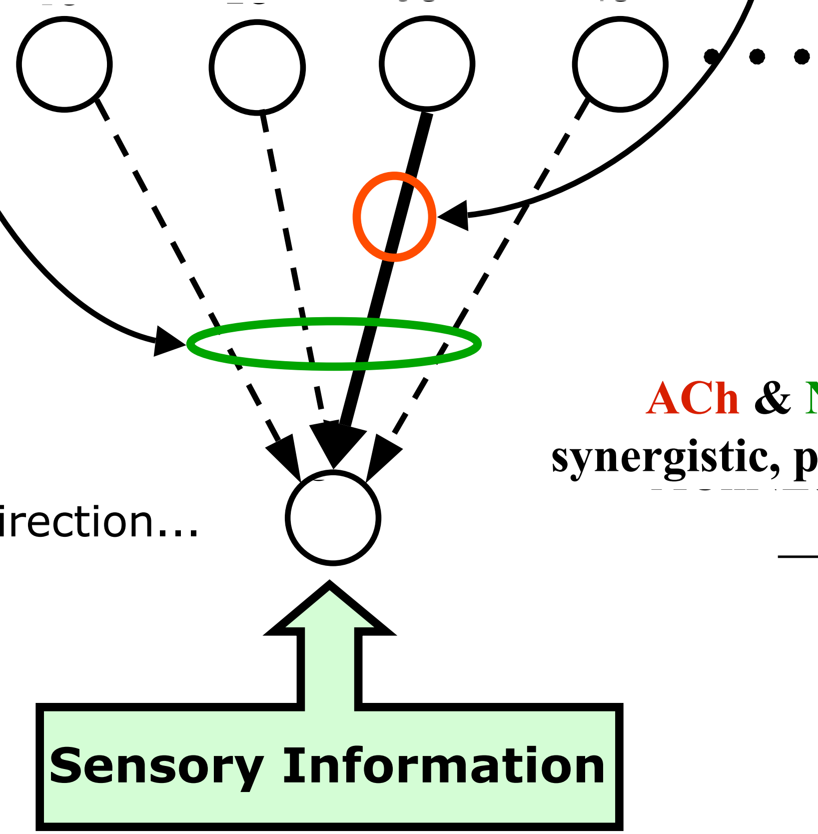

ACh & NE are part-synergistic, part-antagonistic:

avoid representing

full uncertainty

Phasic NE: A-J; Sara

detect and react to a rare target amongst common distractors

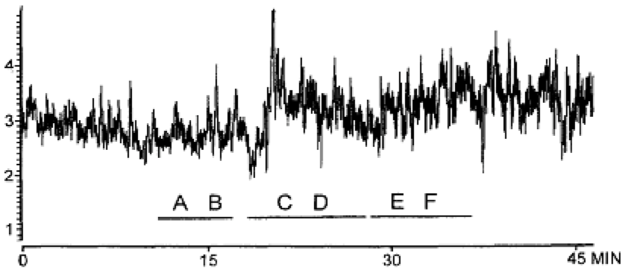

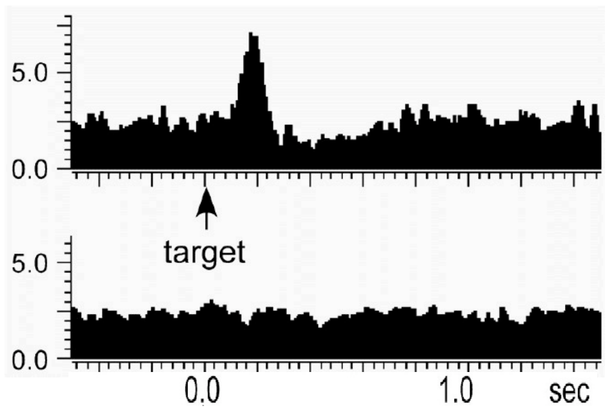

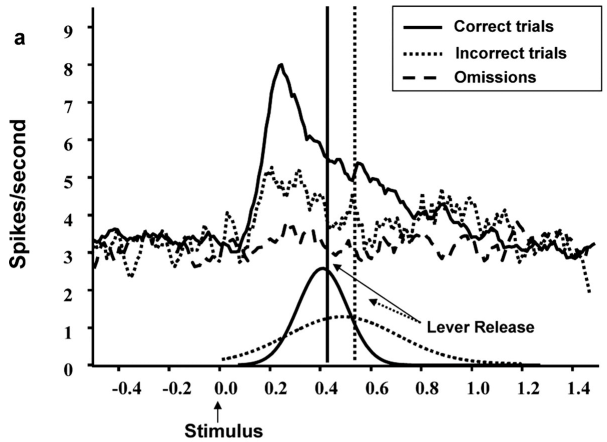

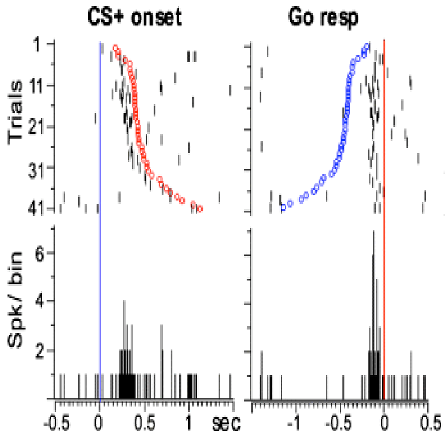

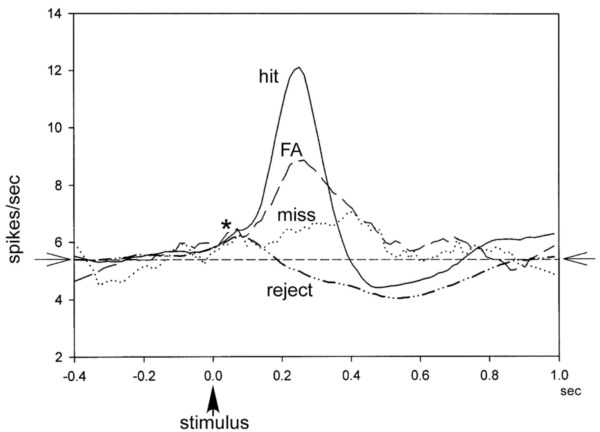

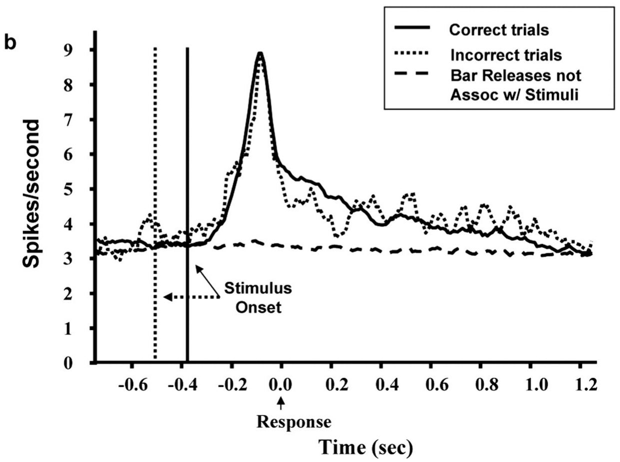

- elevated tonic activity for reversal
- activated by rare target (and reverses)
- not reward/stimulus related? more response related?

Clayton, *et al*

Vigilance Model

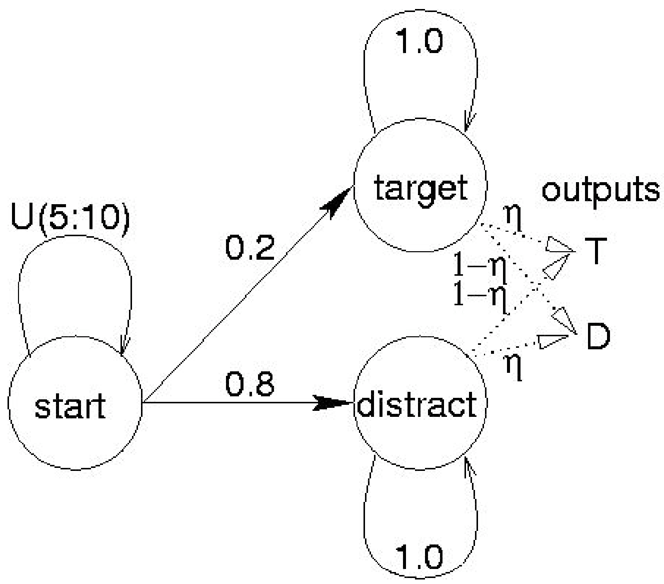

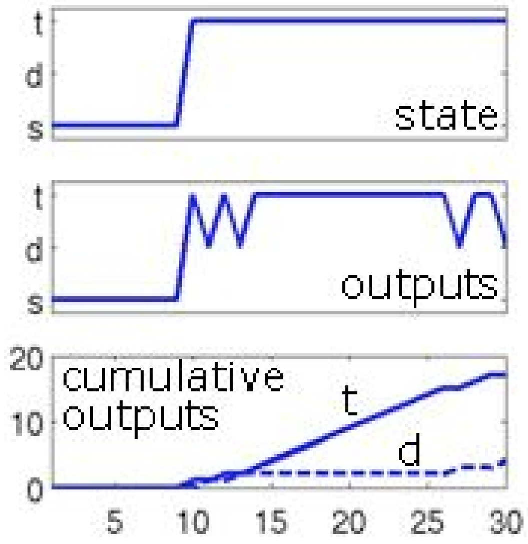

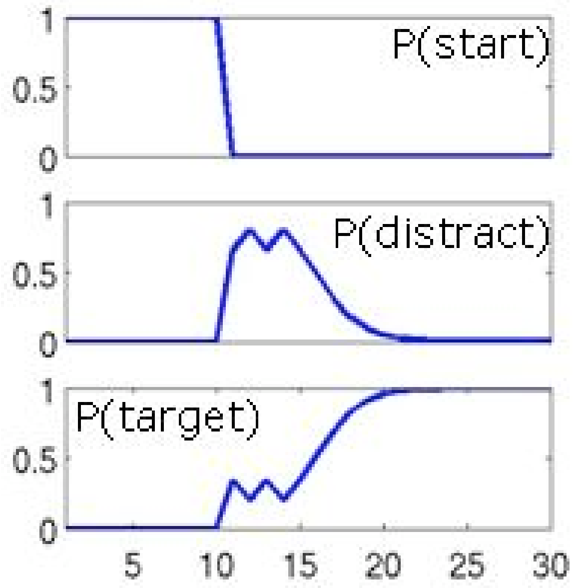

- variable time in start
- _ controls confusability

- one single run
- cumulative is clearer

- exact inference
- effect of 80% prior

Phasic NE

- NE reports uncertainty about current state
  - state in the model, not state of the model
  - divisively related to prior probability of that state

- NE measured relative to default state sequence

  start _ distractor

- temporal aspect - start _ distractor

- structural aspect target *versus* distractor

Phasic NE

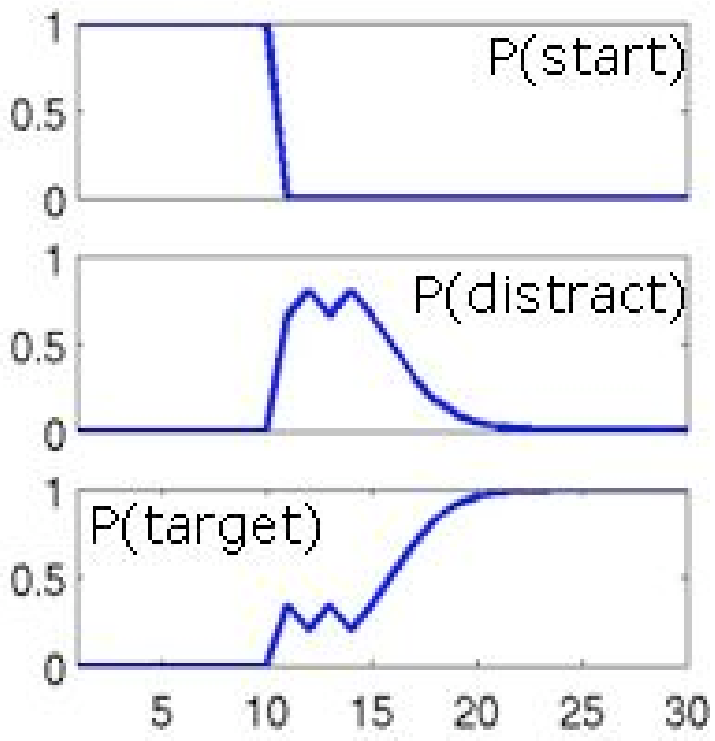

- onset response from timing uncertainty (SET)
- growth as P(target)/0.2 rises
- act when P(target)=0.95
- stop if P(target)=0.01
- arbitrarily set NE=0 after 5 timesteps

(small prob of reflexive action)

Four Types of Trial

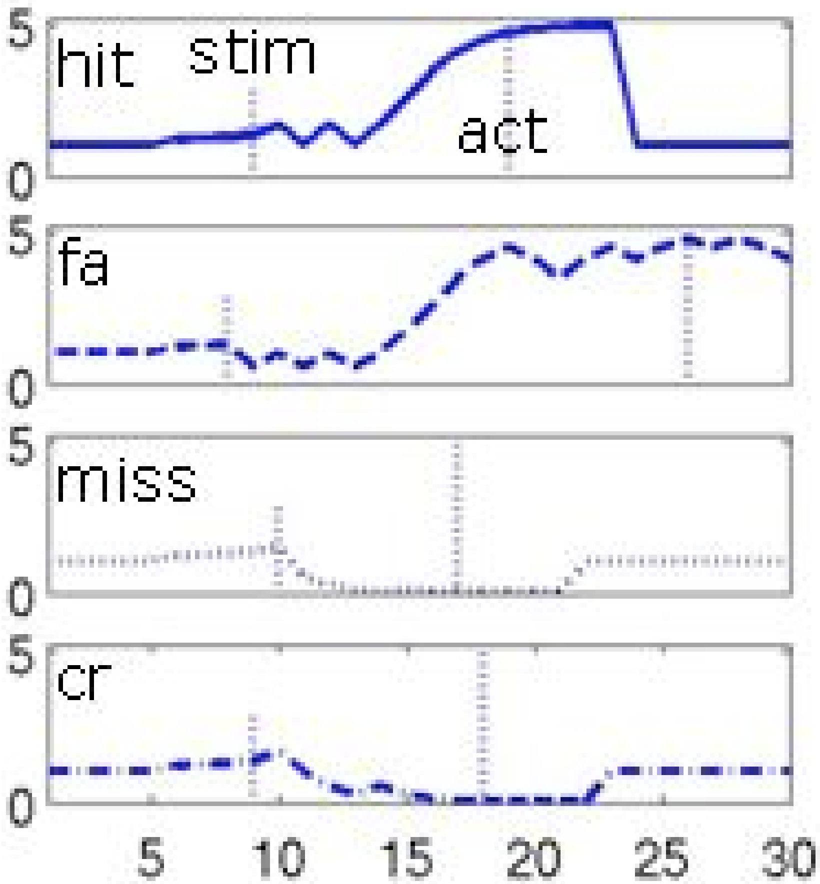

19%

1.5%

1%

77%

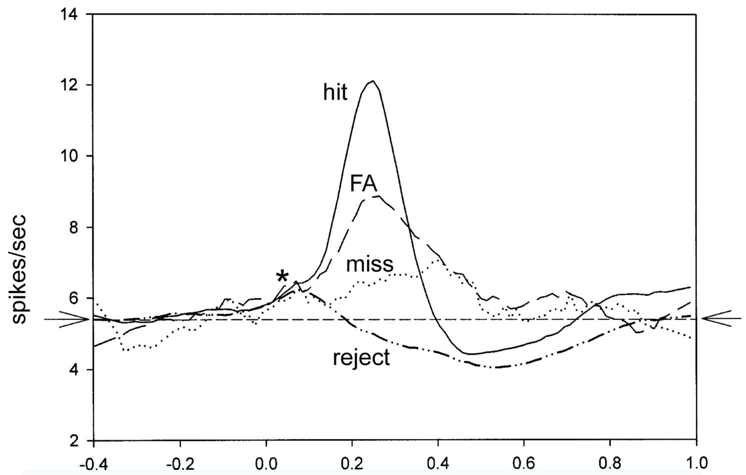

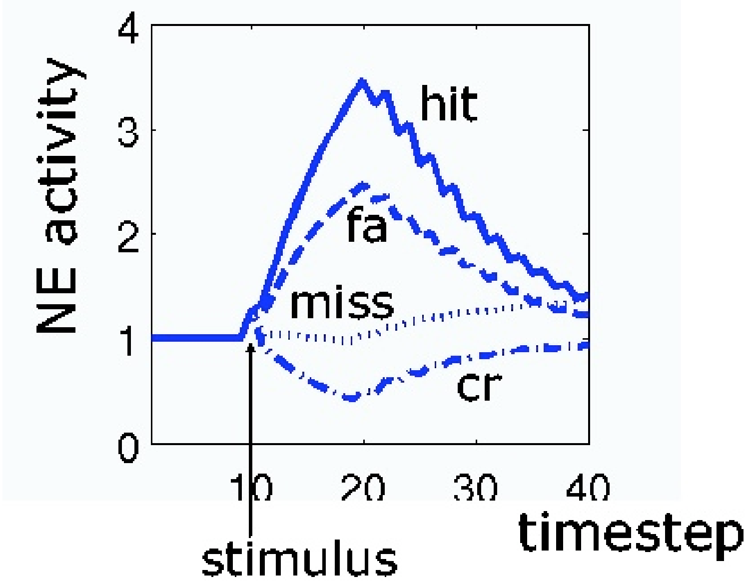

fall is rather arbitrary

Response Locking

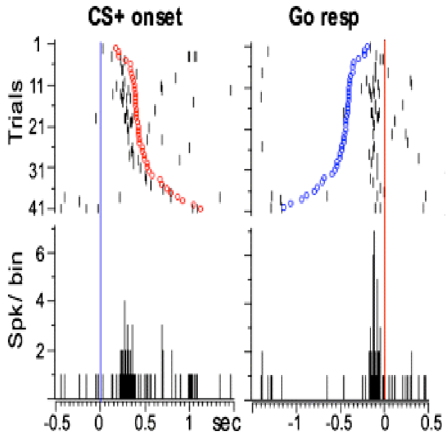

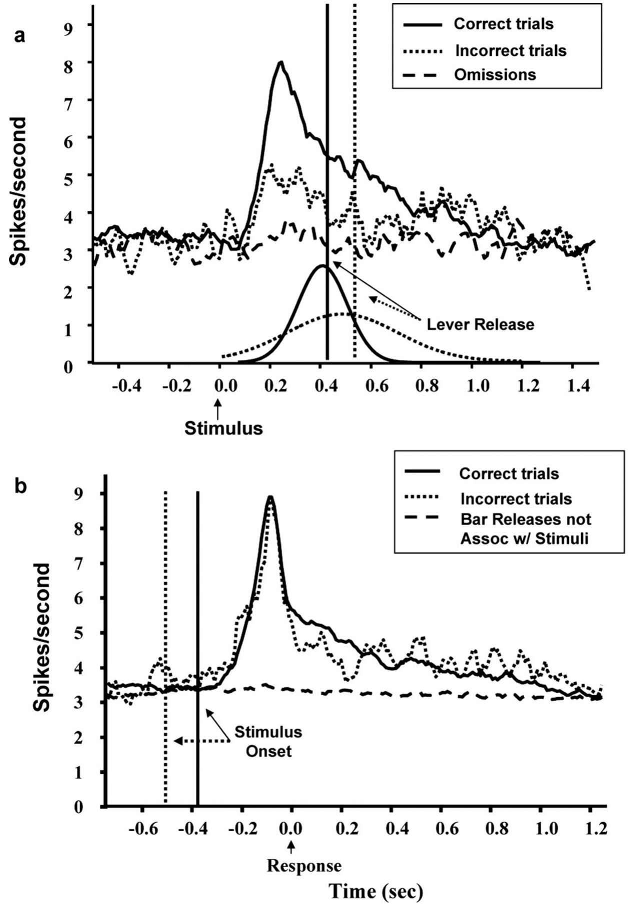

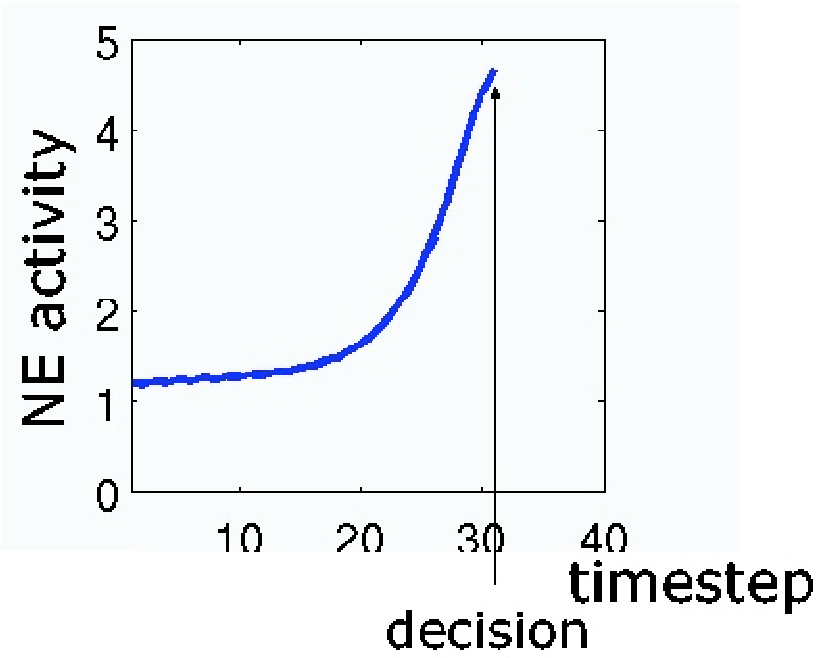

slightly flatters the model – since no further response variability

Task Difficulty

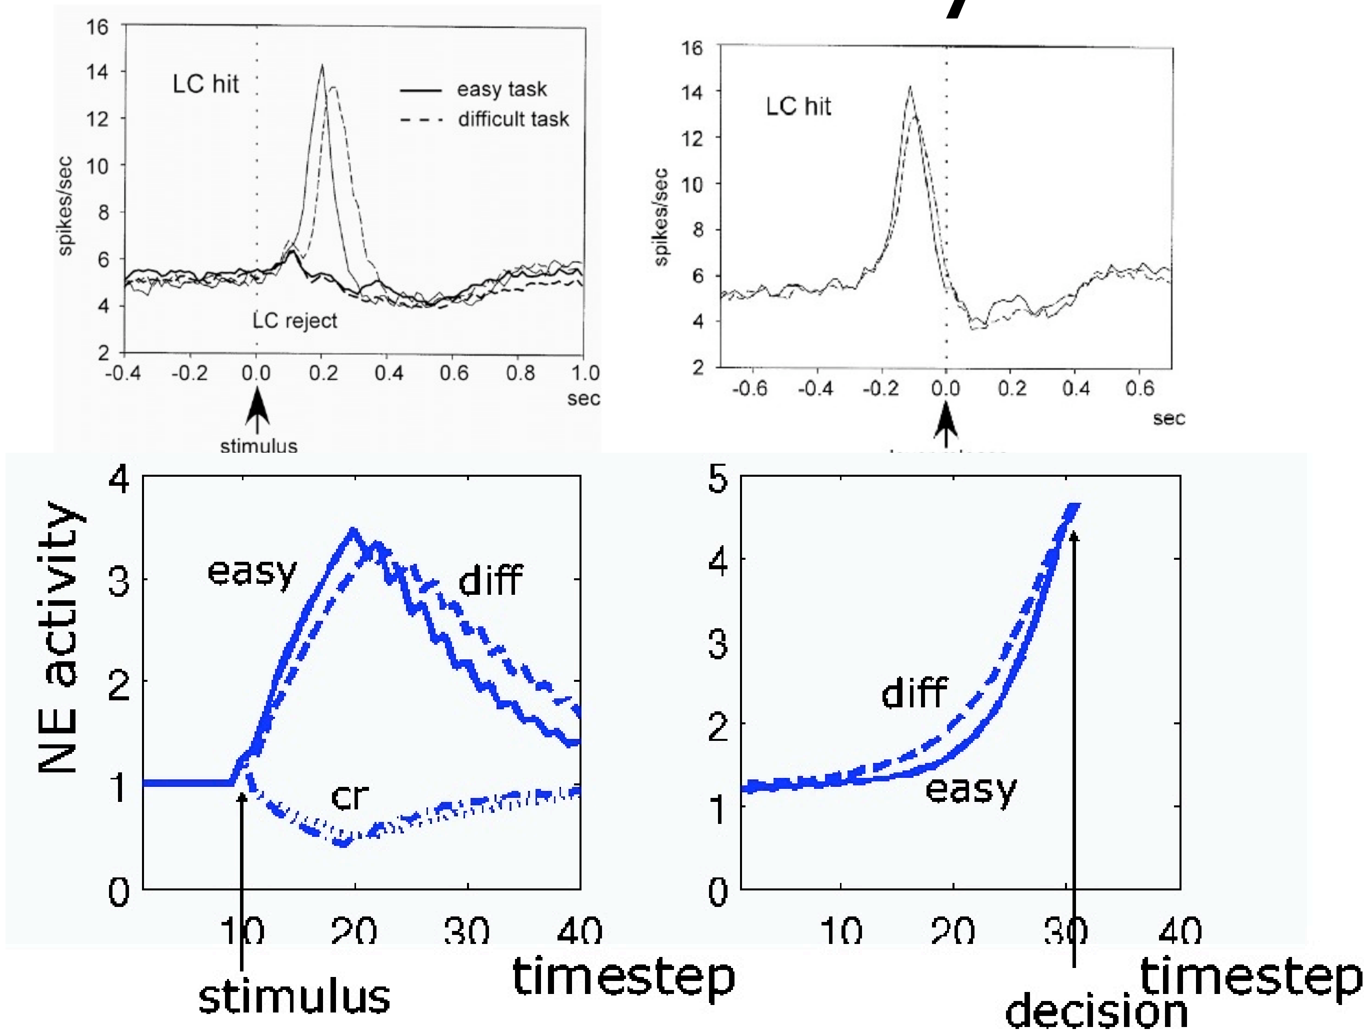

- set _=0.65 rather than 0.675
- information accumulates over a longer period
- hits more affected than cr’s
- timing not quite right

Interrupts/Reset (Shulz)

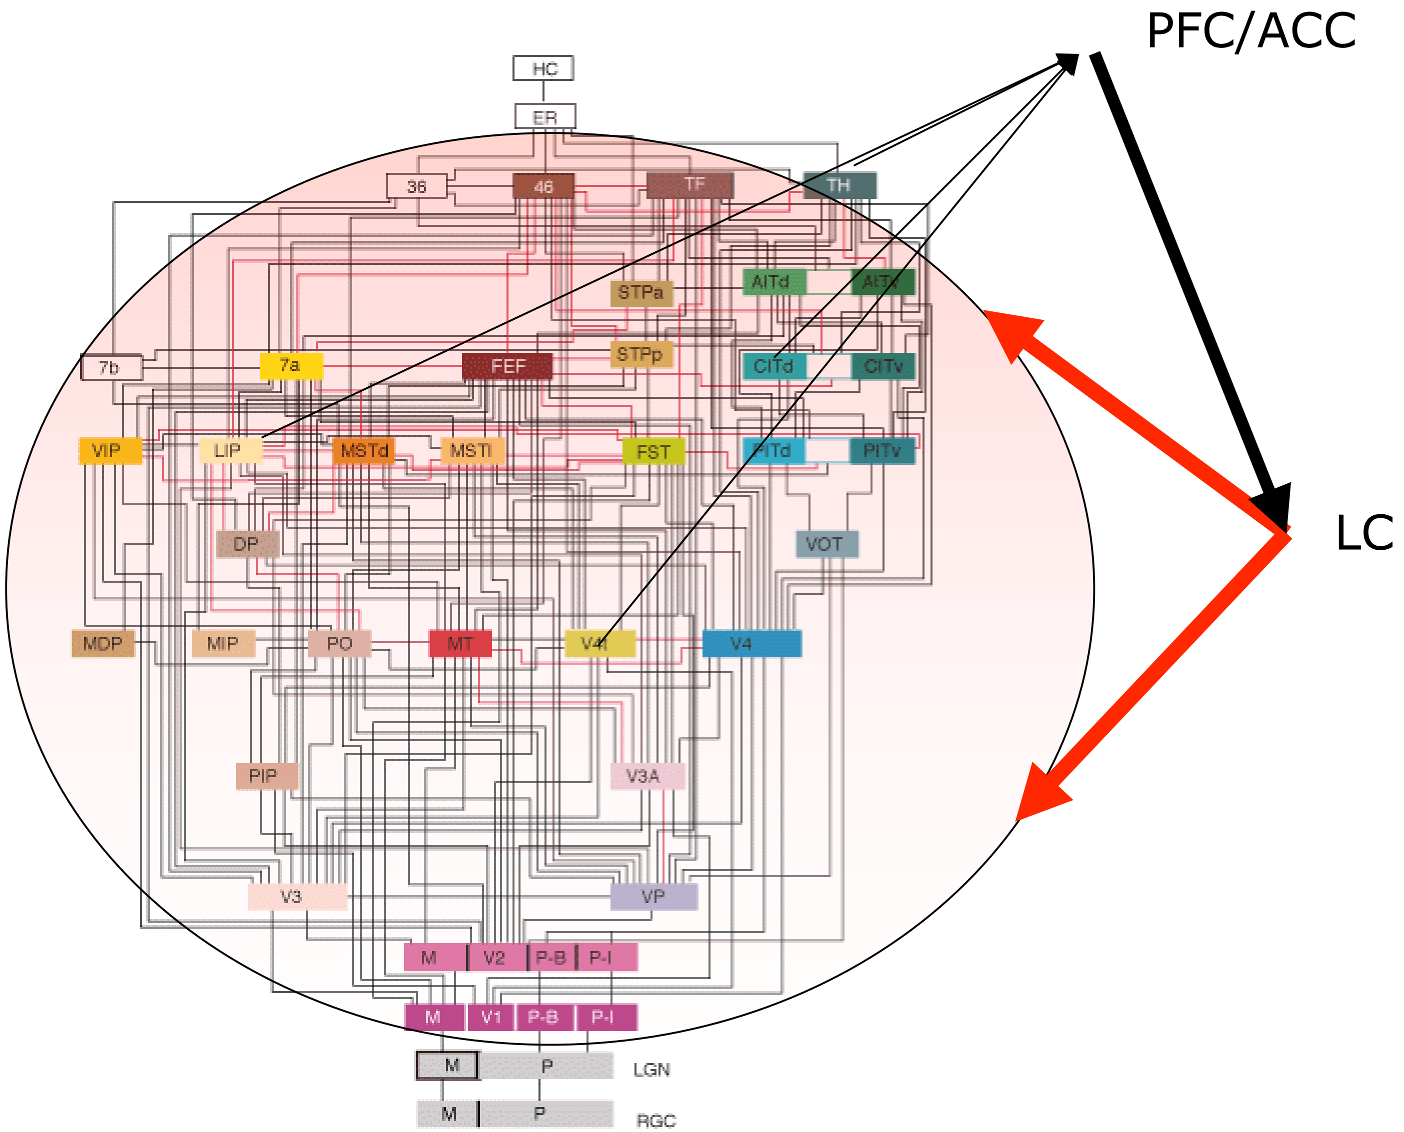

Discusssion

- phasic NE as unexpected state change *within* a model
- relative to prior probability; against default
- interrupts ongoing processing
- tie to ADHD?
- close to alerting (AJ) – but not necessarily tied to behavioral output (onset rise)
- close to behavioural switching (PR) – but not DA
- close to instability (EB)
- phasic ACh: aspects of known variability within a state?

Neuromodulation and Uncertainty

acetylcholine

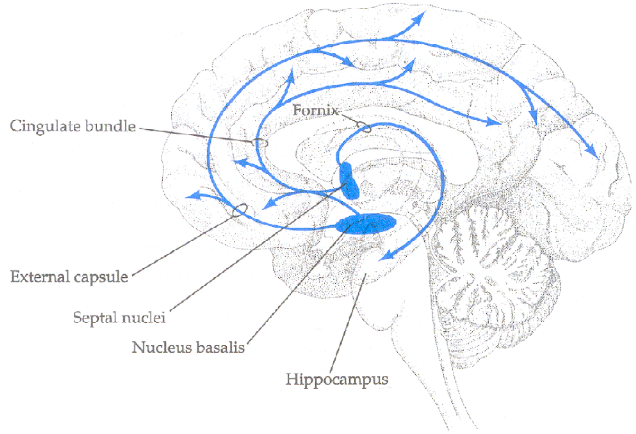

norepinephrine

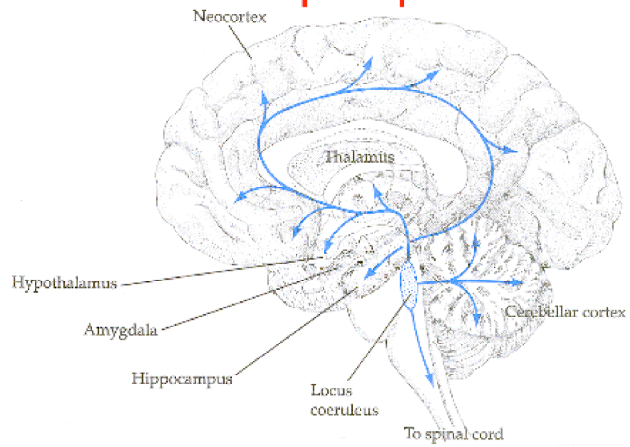

- ACh/NE as expected/unexpected uncertainty signals
- experimental psychopharmacological data replicated by simulations
- implications from complex interactions between ACh & NE
- predictions at the cellular, systems, and behavioral levels
- activity vs weight vs neuromodulatory vs population representations
- irreducible uncertainty vs ignorance

Simulation Results: Posner's Task

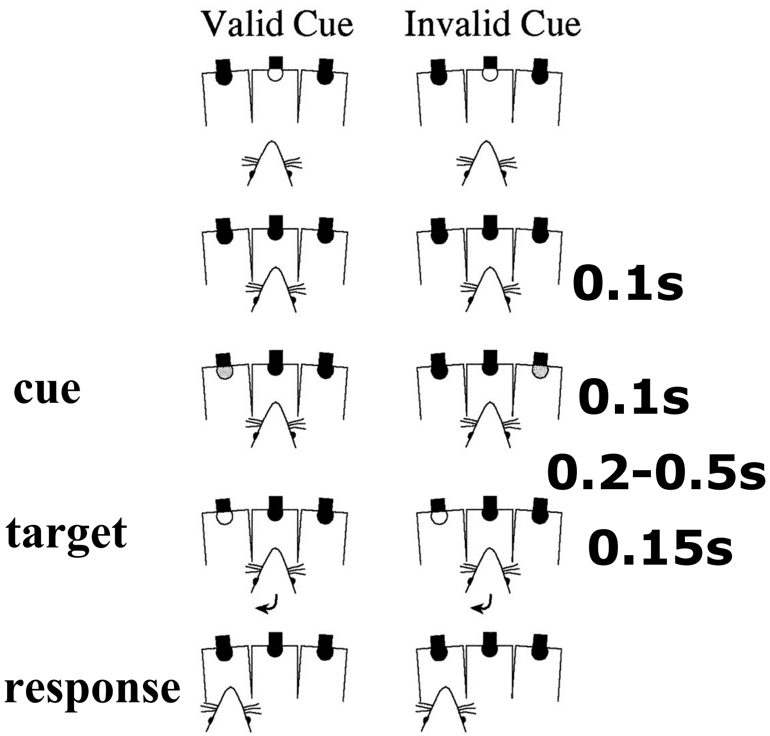

no unexpected change

nicotine

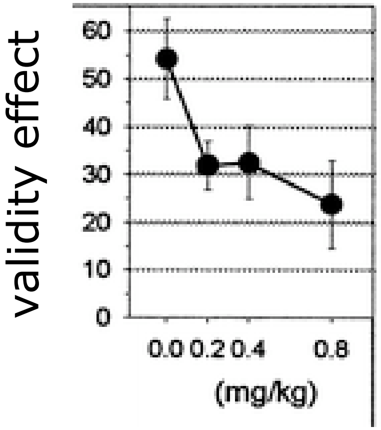

concentration

scopolamine

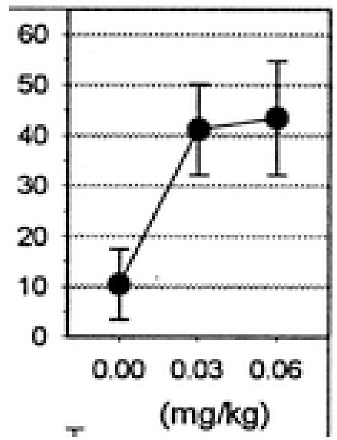

concentration

(Phillips, McAlonan, Robb, & Brown, 2000)

increase ACh

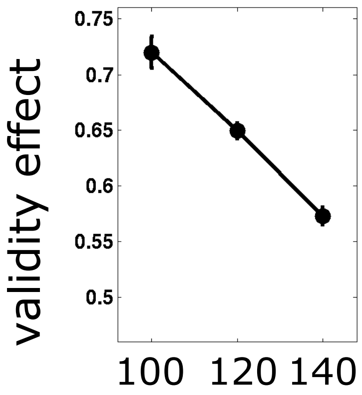

% normal level

decrease ACh

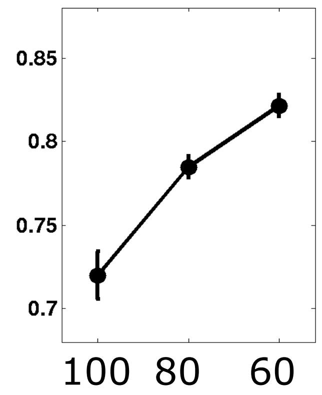

% normal level

Simulation Results: Maze Navigation

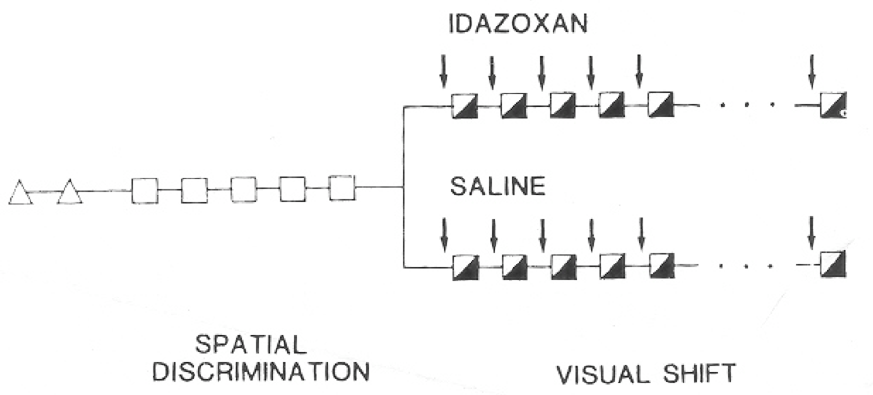

no invalidity

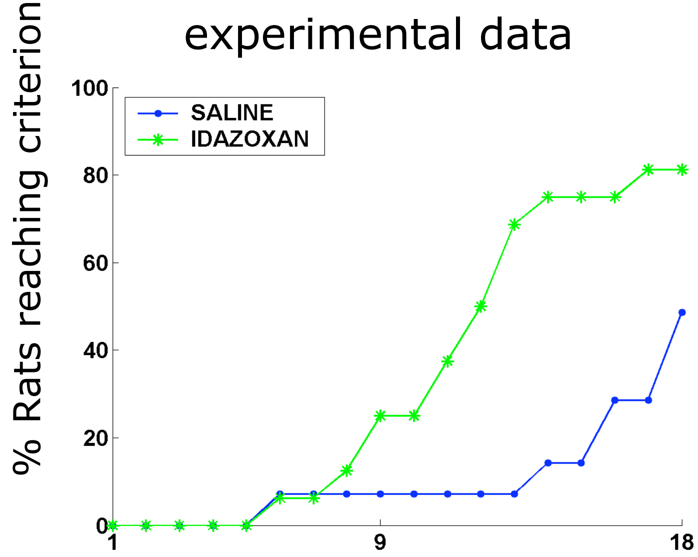

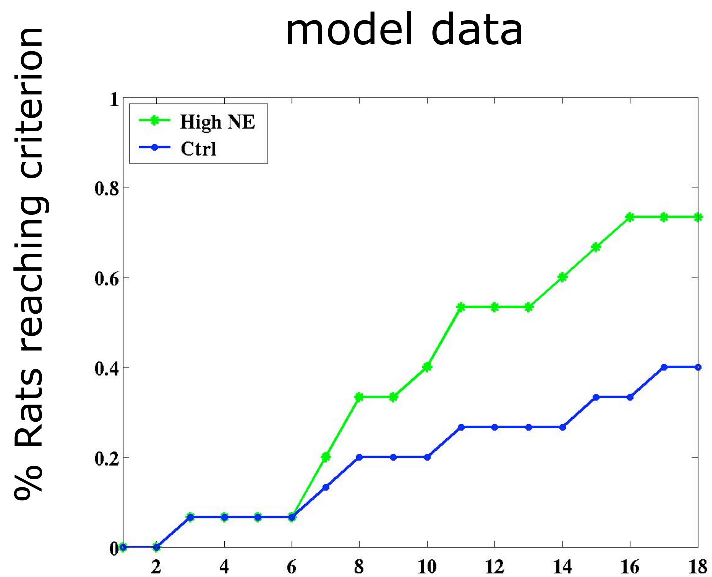

No. days after shift from spatial to visual task

(Devauges & Sara, 1990)

ACh/NE

Typical Run

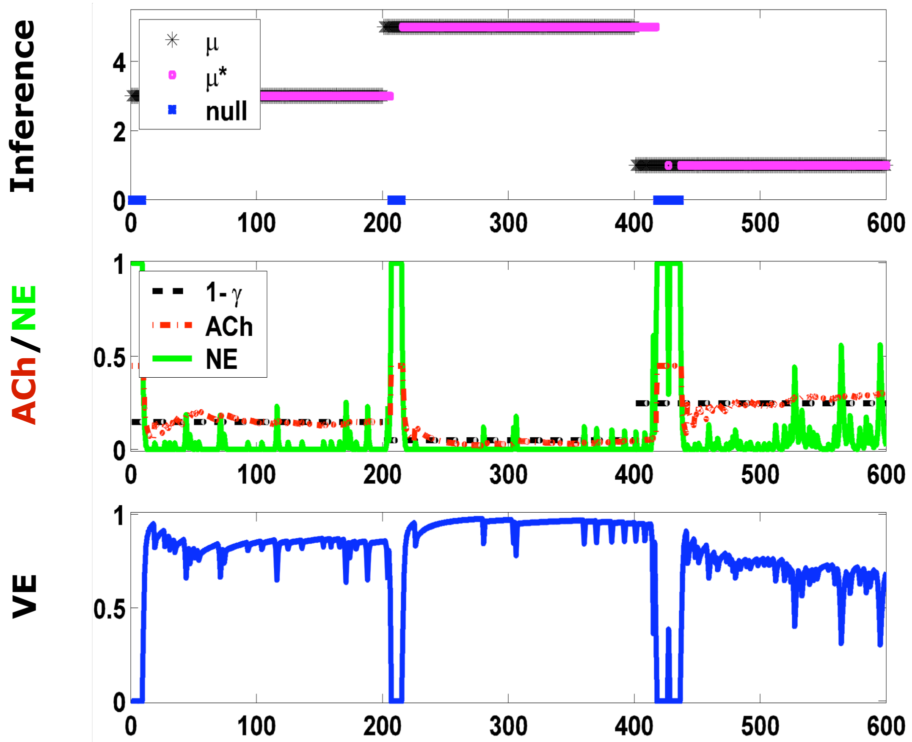

Simulated Pharmacology

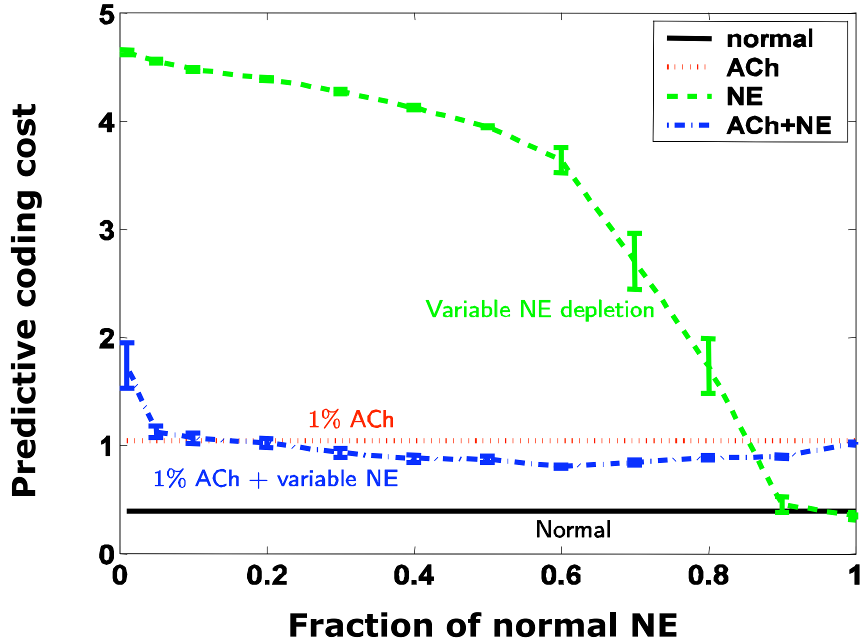

Summary

- single framework for understanding ACh, NE and some aspects of attention and learning

- ACh/NE as expected/unexpected uncertainty signals

- experimental psychopharmacological data replicated by model simulations

- implications from complex interactions between ACh & NE

- predictions at the cellular, systems, and behavioral levels

- activity vs weight vs neuromodulatory vs population representations of uncertainty

- Kalman filter; added 'shock' process; also competitive combination

- irreducible uncertainty vs ignorance

ACh in Hippocampus

Given ***unfamiliarity***,

ACh:

• ***boosts*** bottom-up, ***suppresses*** recurrent processing

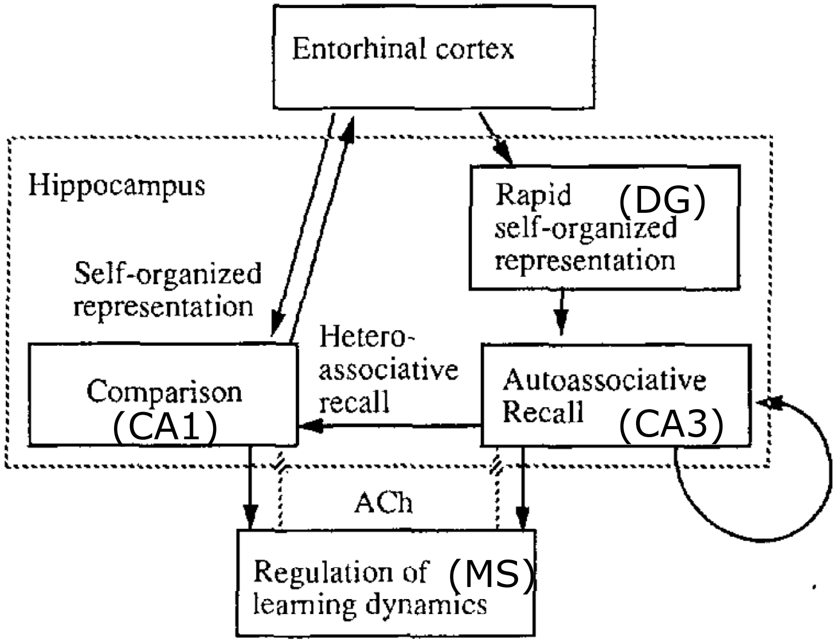

(Hasselmo, 1995)

ACh in Conditioning

Given ***uncertainty***,

ACh:

• ***boosts*** learning to stimuli of

Table 1. Outline of procedures for Experiment 1

Treatment condition (groups)	Phase 1: consistent L–T relation	Phase 2: experimental change in L–T relation	Phase 3: test of conditioning to L
Consistent (CTL-C, PPC-C)	L → T → food; L → T	L → T → food; L → T	L → food
Shift (CTL-S, PPC-S)	L → T → food; L → T	L → T → food; L	L → food

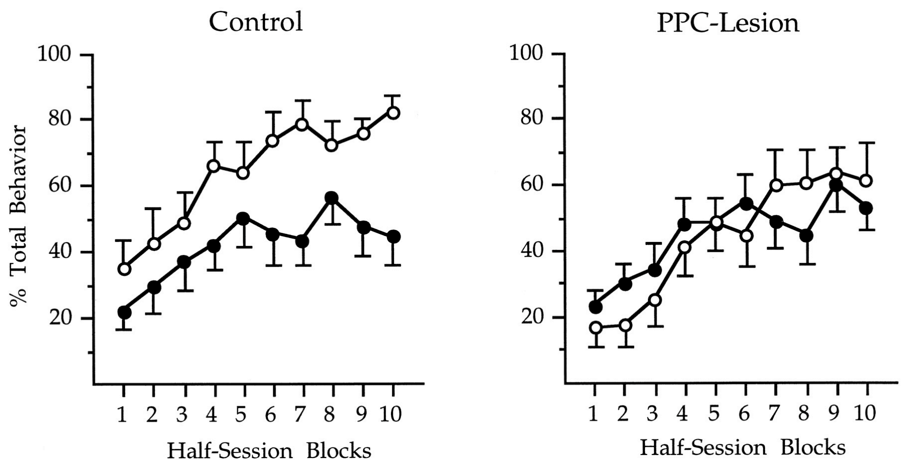

(Bucci, Holland, & Galllagher, 1998)

Cholinergic Modulation in the Cortex

Electrophysiology Data

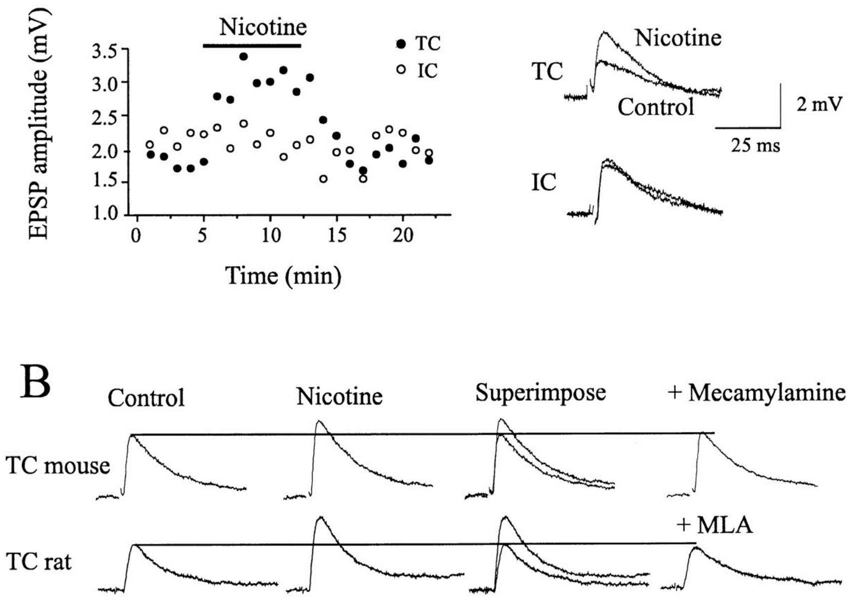

(Gil, Conners, & Amitai, 1997)

Examples of Hallucinations Induced by Anticholinergic Chemicals

Scopolamine in normal volunteers	**Integrated, realistic hallucinations with familiar objects and faces**	**Ketchum et al. (1973)**
Intravenous atropine in bradycardia	**Intense visual hallucinations on eye closure**	**Fisher (1991)**
Local application of scopolamine or atropine eyedrops	**Prolonged anticholinergic delirium in normal adults**	**Tune et al. (1992)**
Side effects of motion-sickness drugs (scopolamine)	**Adolescents hallucinating and unable to recognize relatives**	**Wilkinson (1987) Holland (1992)**

(Perry & Perry, 1995)

ACh agonists:

• ***facilitate*** TC transmission

• ***enhance*** stimulus-

ACh antagonists:

• ***induce*** hallucinations

• ***interfere*** with stimulus processing

Norepinephrine

Something similar may be true for NE
(Kasamatsu *et al*, 1981)

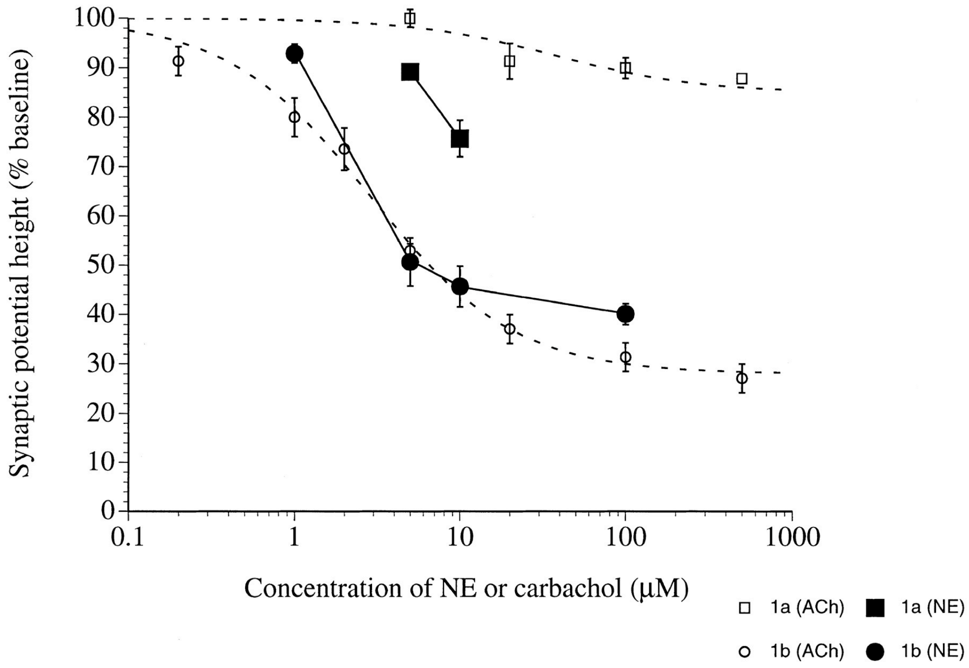

(Hasselmo *et al*, 1997)

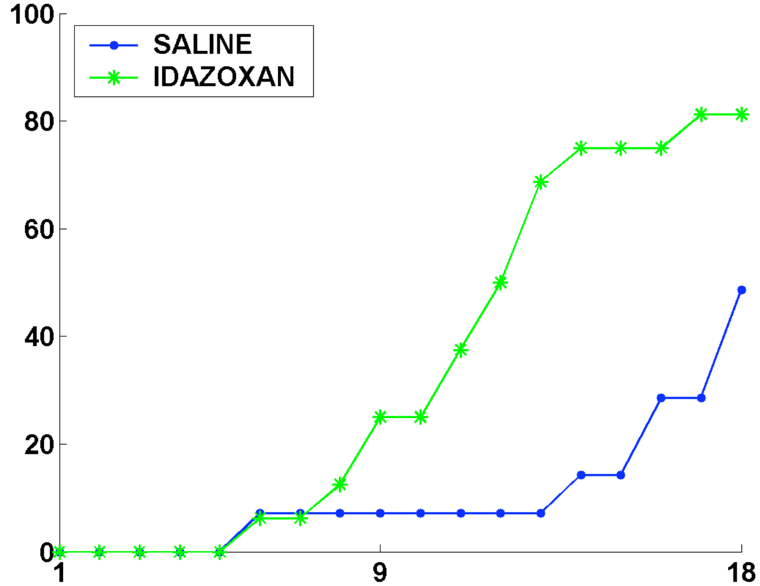

Days after task shift

(Devauges & Sara, 1990)

NE specially involved in **novelty,** confusing association with attention, vigilance,
selective attention

Approximation

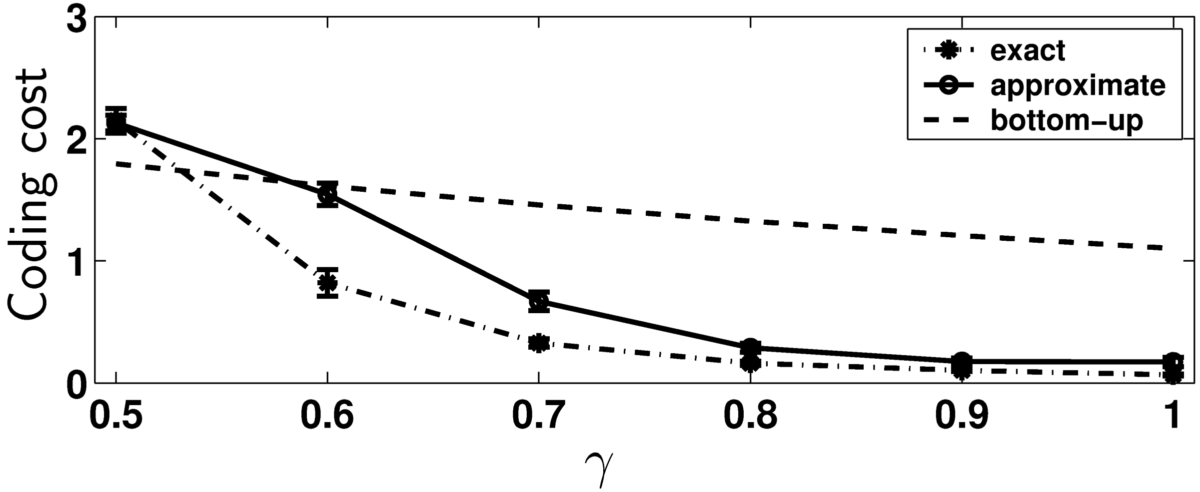

approximation is not catastrophic compared with a simpler,
algorithm