Selection and reinforcement architectures in the vertebrate basal ganglia

euCognition Meeting: Munich, June 2007

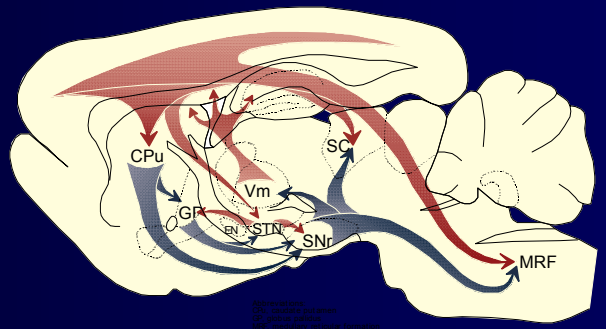

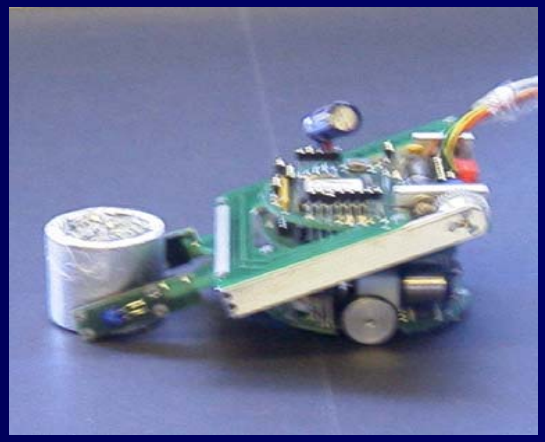

Peter Redgrave

Neuroscience Research Unit,
Dept Psychology,
University of Sheffield, UK

Overview

- A fundamental computational problem – selection
- A biological solution – the basal ganglia
- Reinforcement function(s) in the basal ganglia – biasing selections
- Insights from connectivity
- Agency and discovery of novel actions

A general architecture for a multifunctional system

…including the brain

- Largely independent parallel processing functional units
- Each with:
  - specific functional objectives
  - specialised sensory input
  - specialised behavioural output

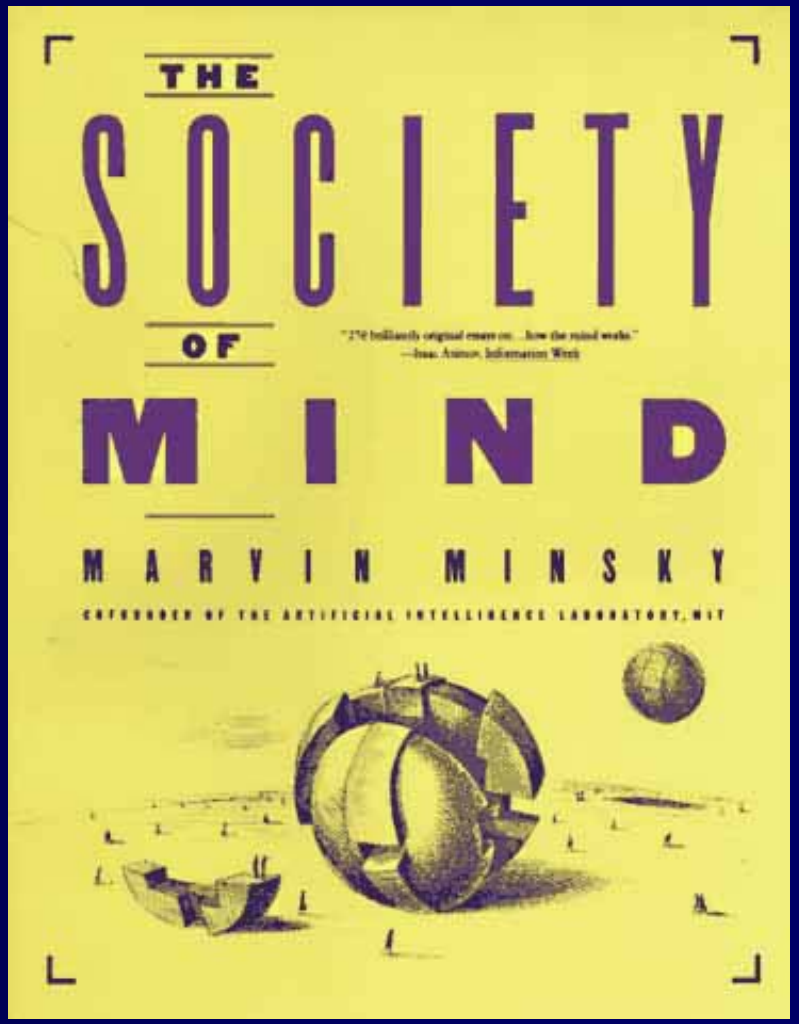

The Selection Problem

- Spatially distributed
- Processing in parallel
- All act through final common motor path

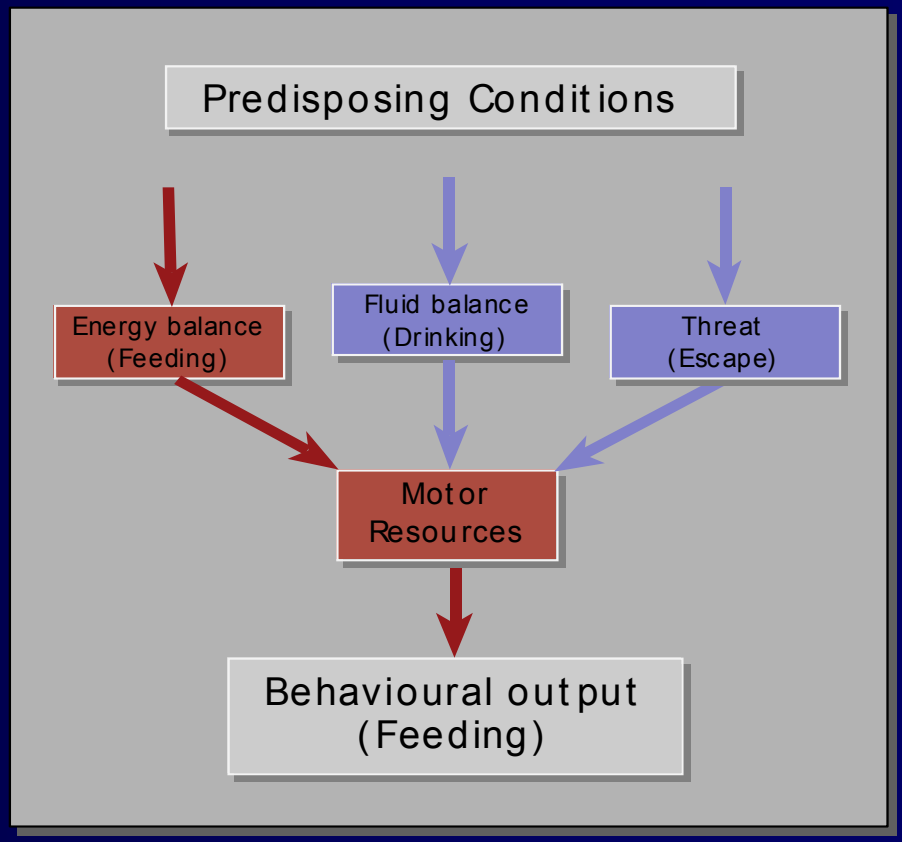

At any point in time which system should be permitted to direct motor output (behaviour)?

Control Engineering Solutions

- Recurrent reciprocal inhibition
  - Selection an emergent property
  - Positive feedback
  - Winner-take-all

- Centralised selection
  - Localised switching
  - Dissociates selection from perception and motor control

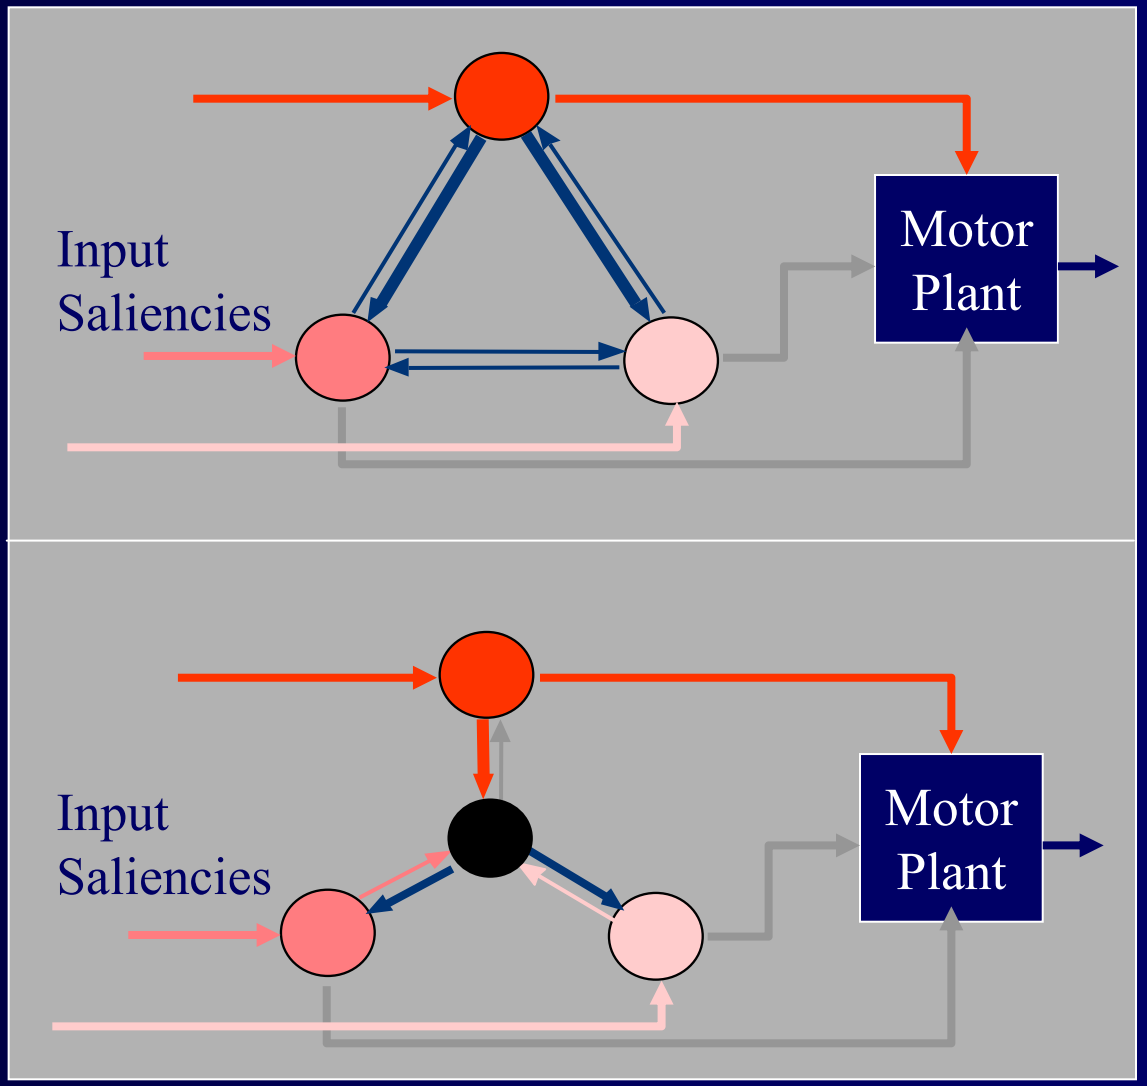

The basal ganglia as a central selector

- External command systems
  - Cortical
  - Limbic
  - Brainstem
- Command inputs
  - Sensory
  - Cognitive
  - Affective
- Command outputs
  - Converge on brainstem and spinal motor generators
- Links with basal ganglia
  - Phasic excitatory inputs
  - Tonic inhibitory outpus

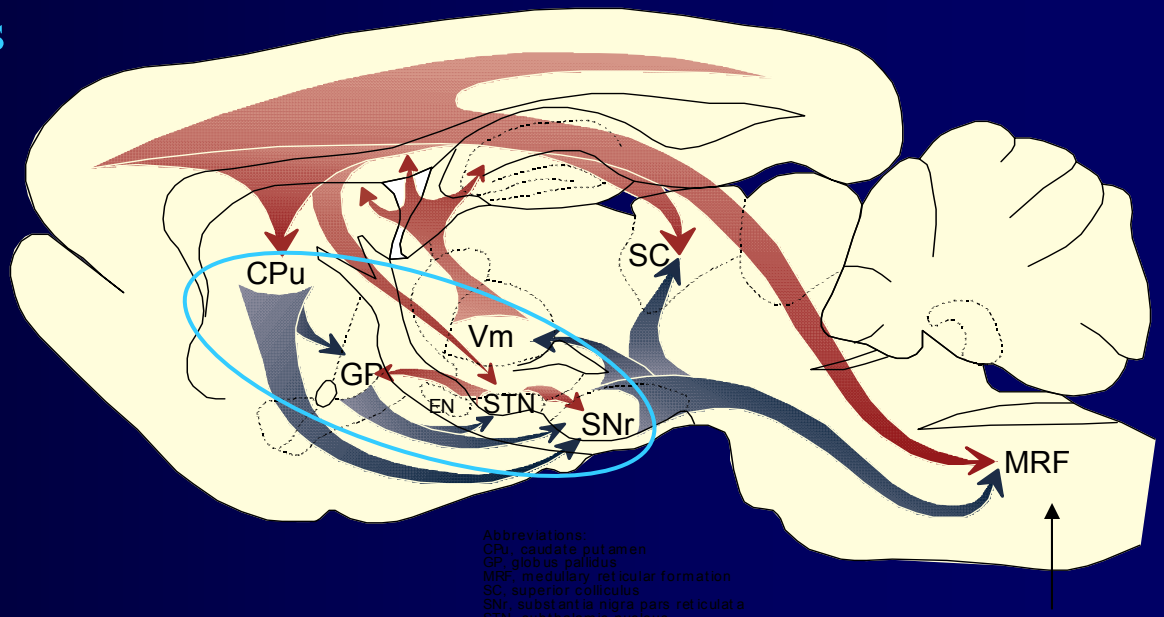

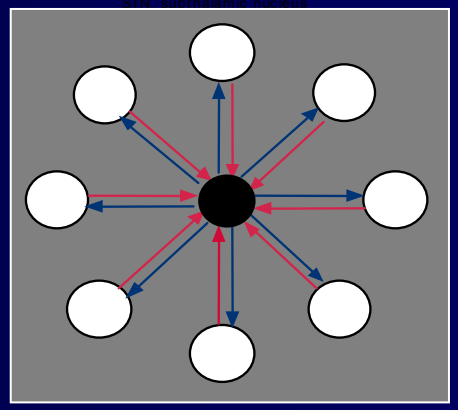

Central selection architecture

Redgrave P, Prescott T, Gurney KN. 1999. The basal ganglia: A vertebrate solution to the selection problem? Neuroscience 89:1009-1023.

Evolutionary conservatism

"The basal ganglia in modern mammals, birds and reptiles (i.e. modern amniotes) are very similar in connections and neurotransmitters, suggesting that the evolution of the basal ganglia in amniotes has been very conservative."

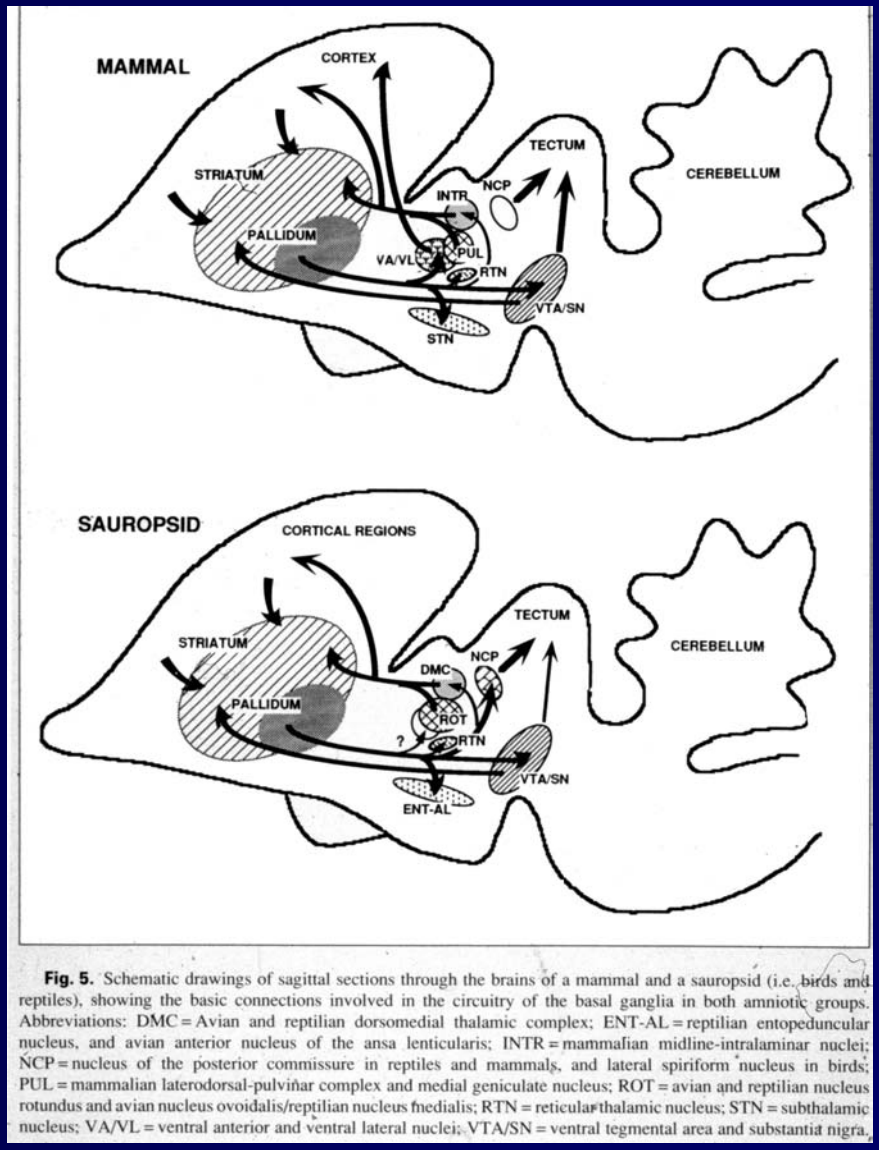

Fig. 5. Schematic drawings of sagittal sections through the brains of a mammal and a sauropsid (i.e. birds and reptiles), showing the basic connections involved in the circuitry of the basal ganglia in both amniotic groups. Abbreviations: DMC = Avian and reptilian dorsomedial thalamic complex; ENT-AL = reptilian entopeduncular nucleus, and avian anterior nucleus of the ansa lenticularis; INTR = mammalian midline-intralaminar nuclei; NCP = nucleus of the posterior commissure in reptiles and mammals, and lateral spiriform nucleus in birds; PUL = mammalian laterodorsal-pulvinar complex and medial geniculate nucleus; ROT = avian and reptilian nucleus rotundus and avian nucleus ovoidalis/reptilian nucleus medialis; RTN = reticular thalamic nucleus; STN = subthalamic nucleus; VA/VL = ventral anterior and ventral lateral nuclei; VTA/SN = ventral tegmental area and substantia nigra.

Medina, L and Reiner, A.

Neurotransmitter organization and connectivity of the basal ganglia in vertebrates: Implications for the evolution of basal ganglia. Brain Behaviour and Evolution (1995) **46**, 235-258

Looped architecture: a fundamental component

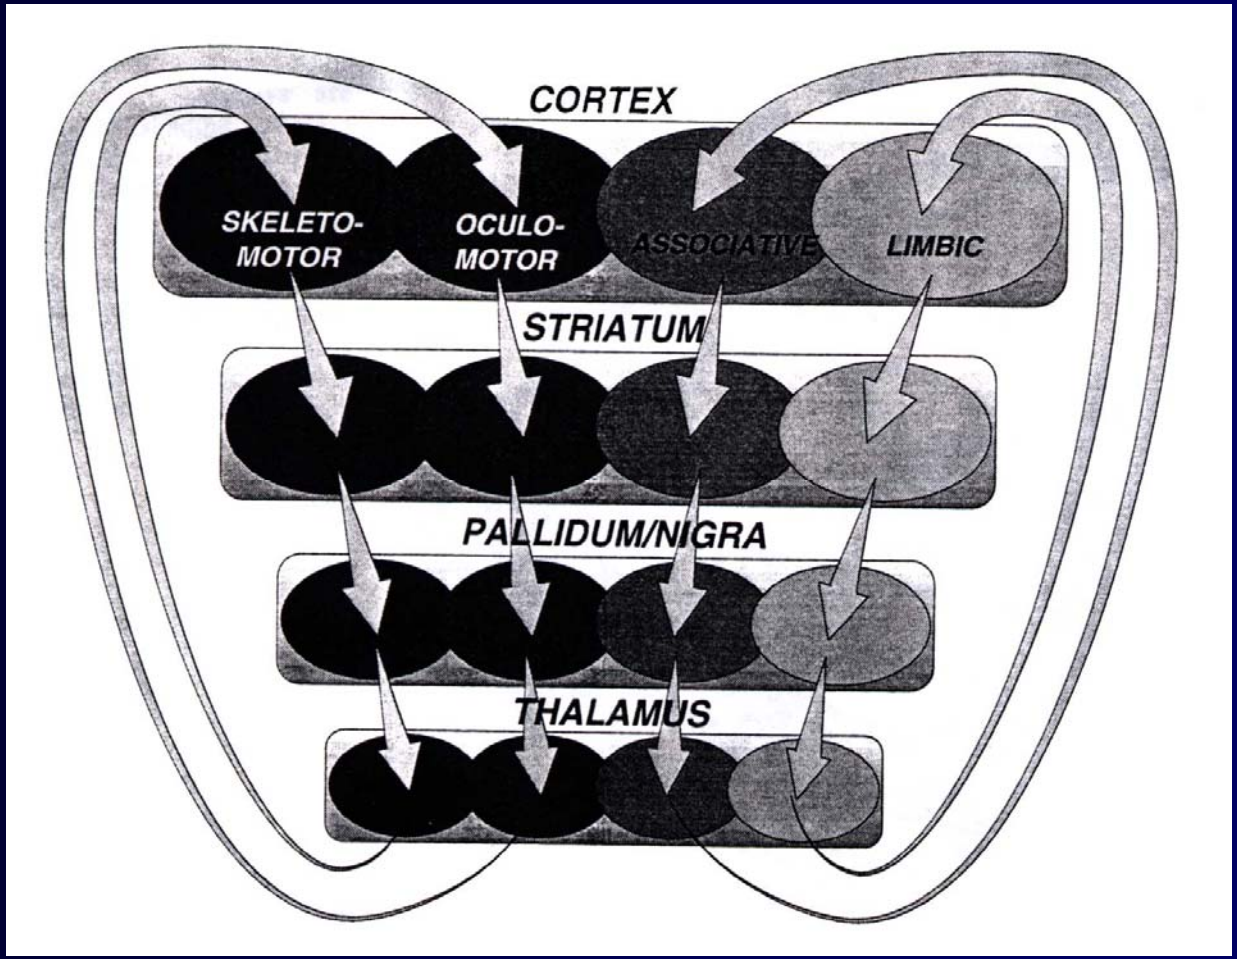

Alexander, G. E., M. R. DeLong, et al. (1986). "Parallel organization of functionally segregated circuits linking basal ganglia and cortex." Ann. Rev. Neurosci. **9**: 357-381.

Repeating microcircuitry across functional territories

- External inputs
  - Cerebral cortex
  - Limbic system
  - Brainstem via thalamus
- Input functions
  - Cognitive
  - Affective
  - Sensorimotor

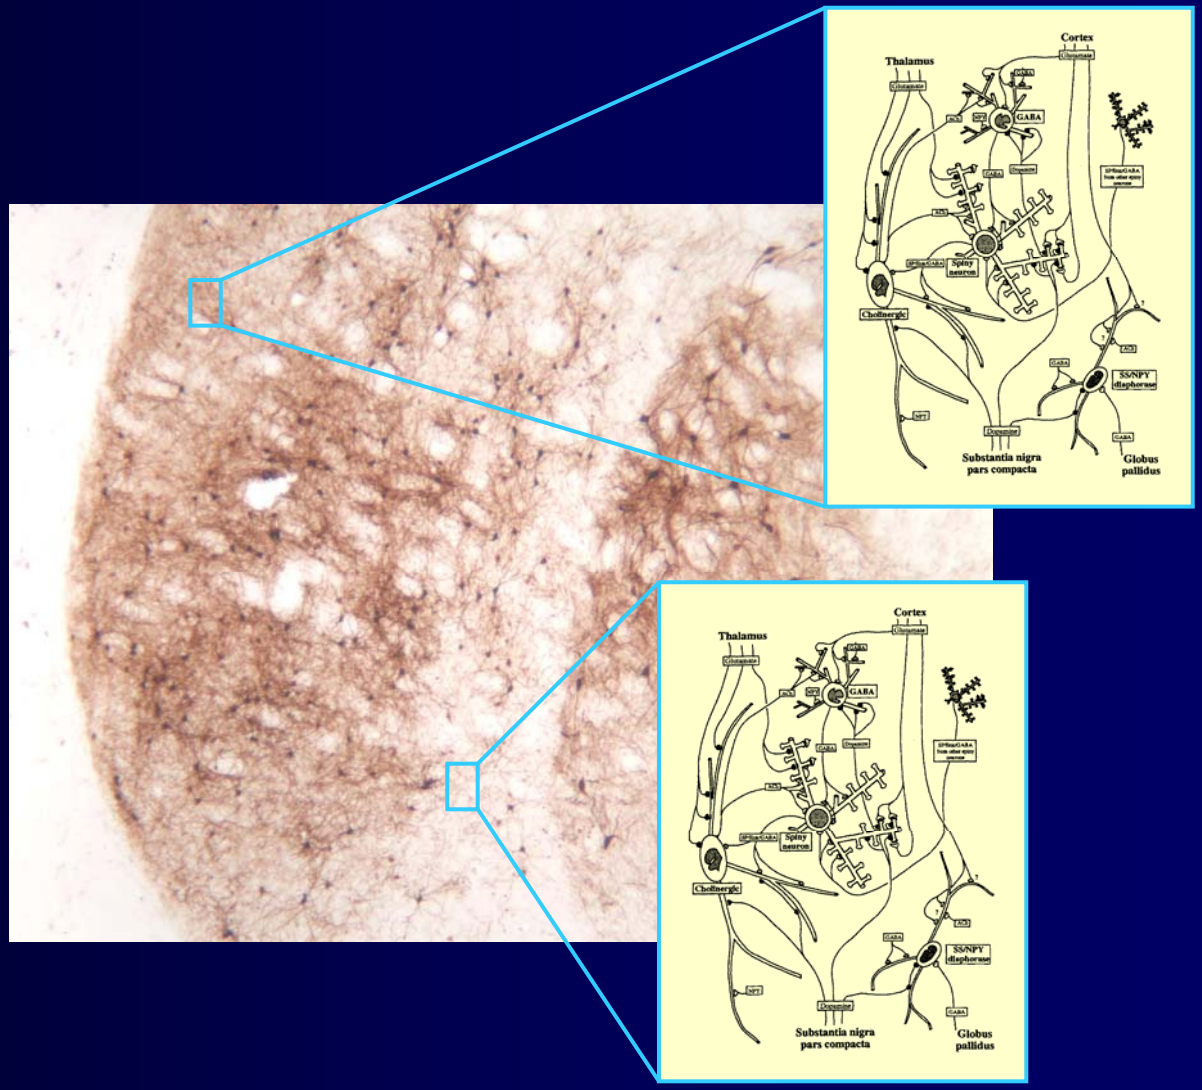

Bolam JP, Bennett BD. 1995. Microcircuitry of the neostriatum. In: Ariano MA, Surmeier DJ, editors. Molecular and cellular mechanims of neostriatal function. Austin, TX.: R.G. Landes Co. p 1-19.

Disinhibition is the critical output function

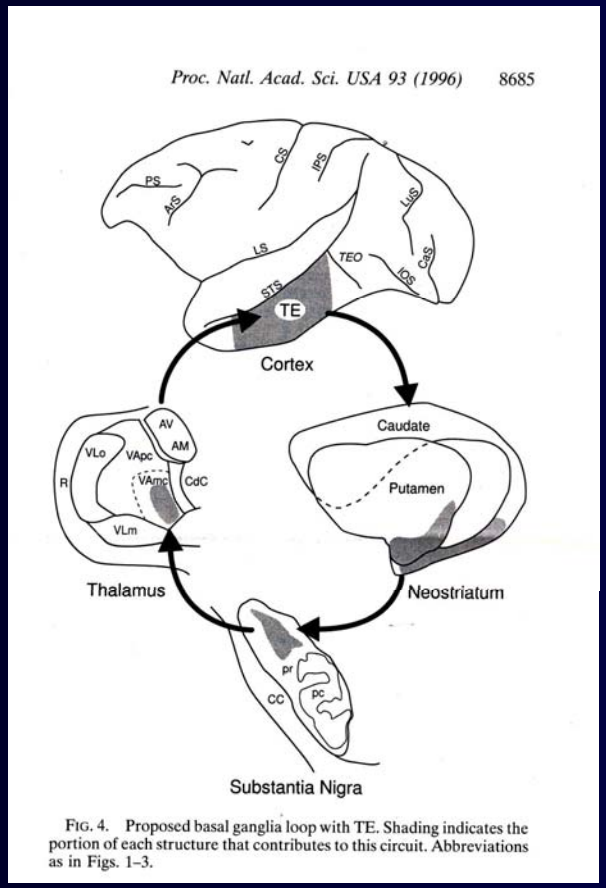

FIG. 4. Proposed basal ganglia loop with TE. Shading indicates the portion of each structure that contributes to this circuit. Abbreviations as in Figs. 1–3.

Middleton, F. A. and P. L. Strick (1996). "The temporal lobe is a target of output from the basal ganglia." Proc Natl Acad Sci USA **93**(16): 8683-8687.

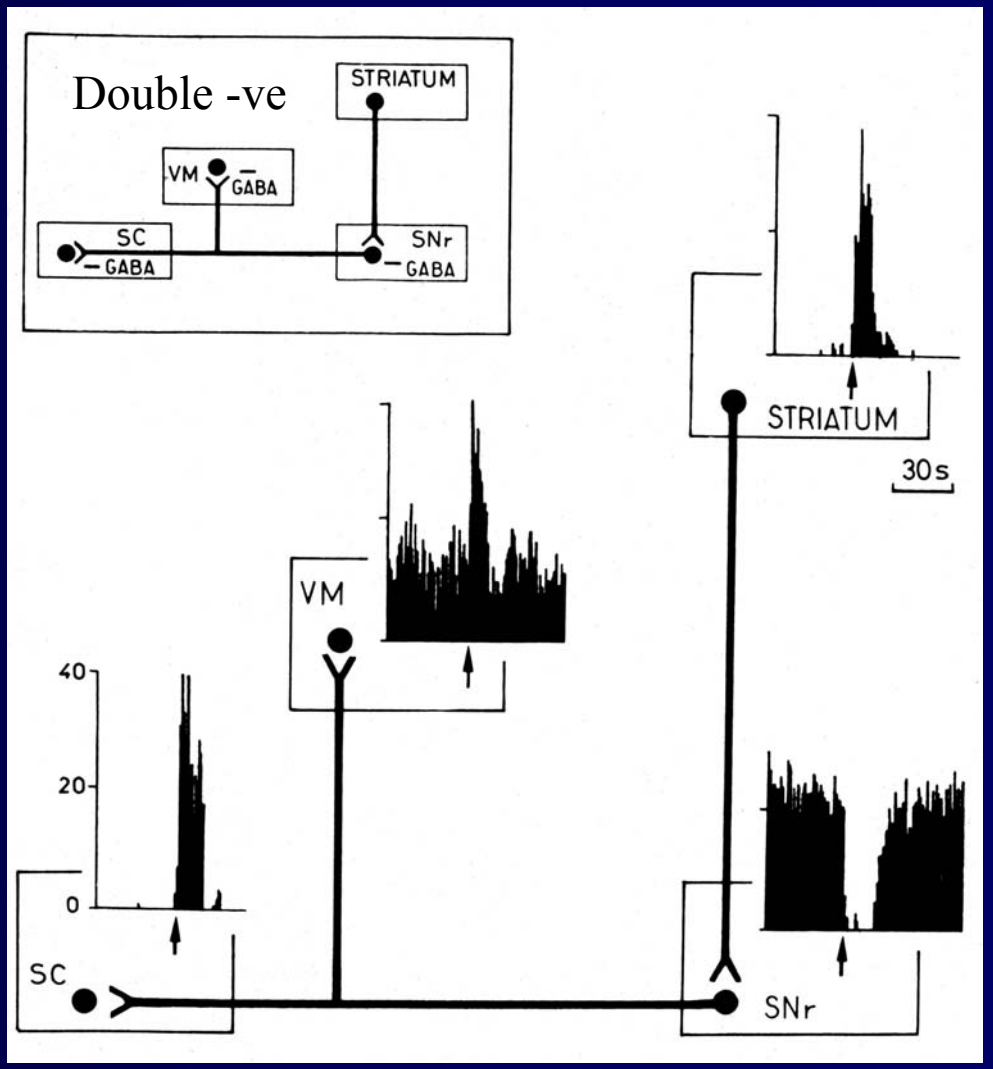

Chevalier, G. and J. M. Deniau (1990). "Disinhibition as a basic process in the expression of striatal functions." Trends Neurosci. **13**: 277-281.

Selective disinhibition is a mechanism for selection

Potential resolution →

The Selection Problem

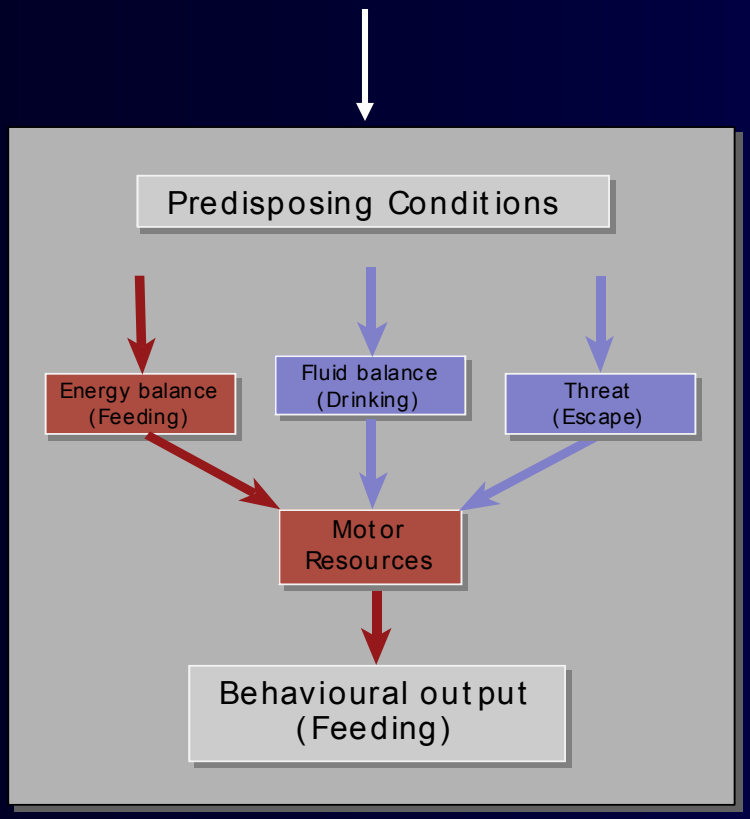

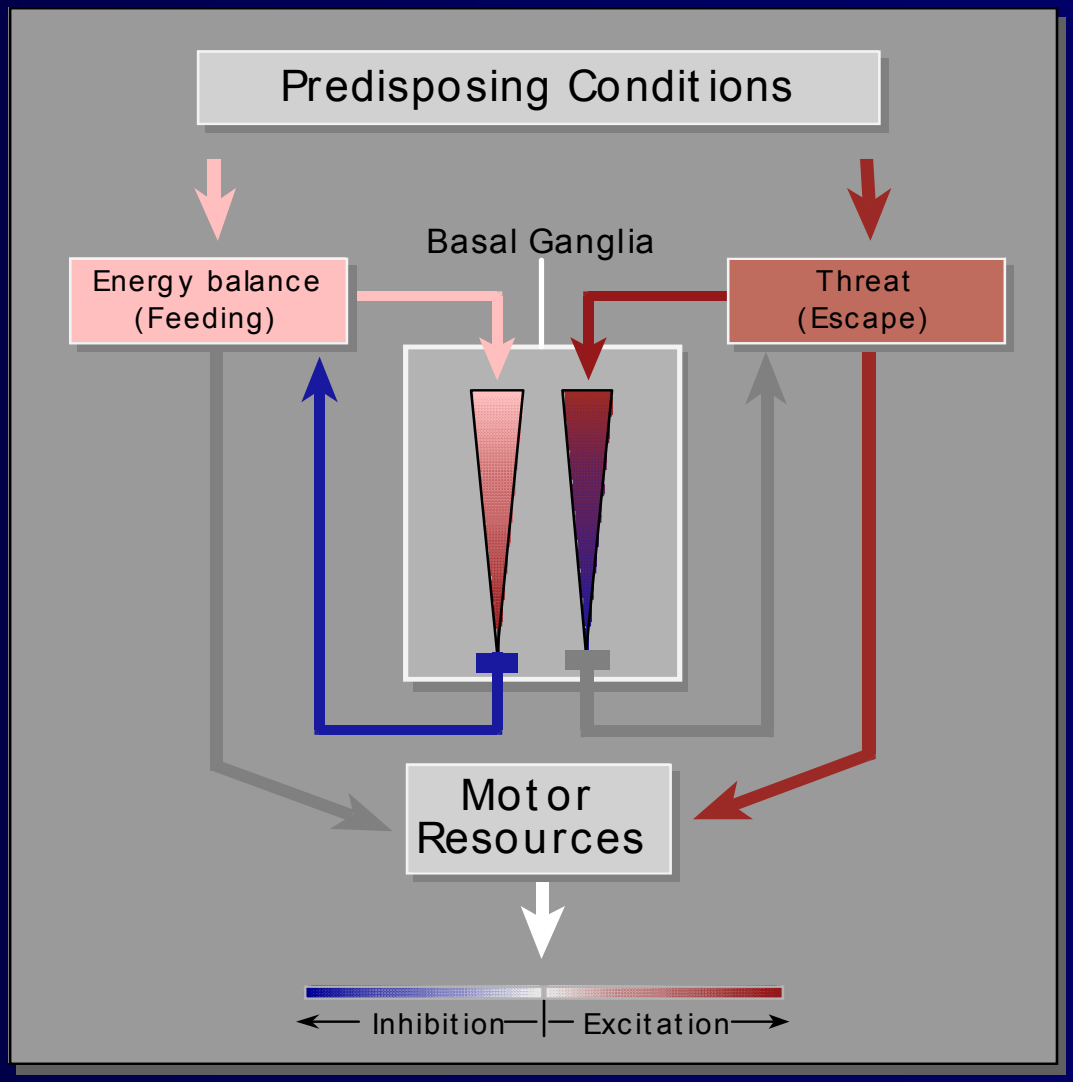

Redgrave P, Prescott T, Gurney KN. 1999. The basal ganglia: A vertebrate solution to the selection problem? Neuroscience 89:1009-1023.

Serial Selection in the Basal Ganglia

1) Up-down states of medium spiny neurones

2) Local inhibition in striatum

3) Diffuse/focused projection onto output nuclei

4) Recurrent inhibition in output nuclei

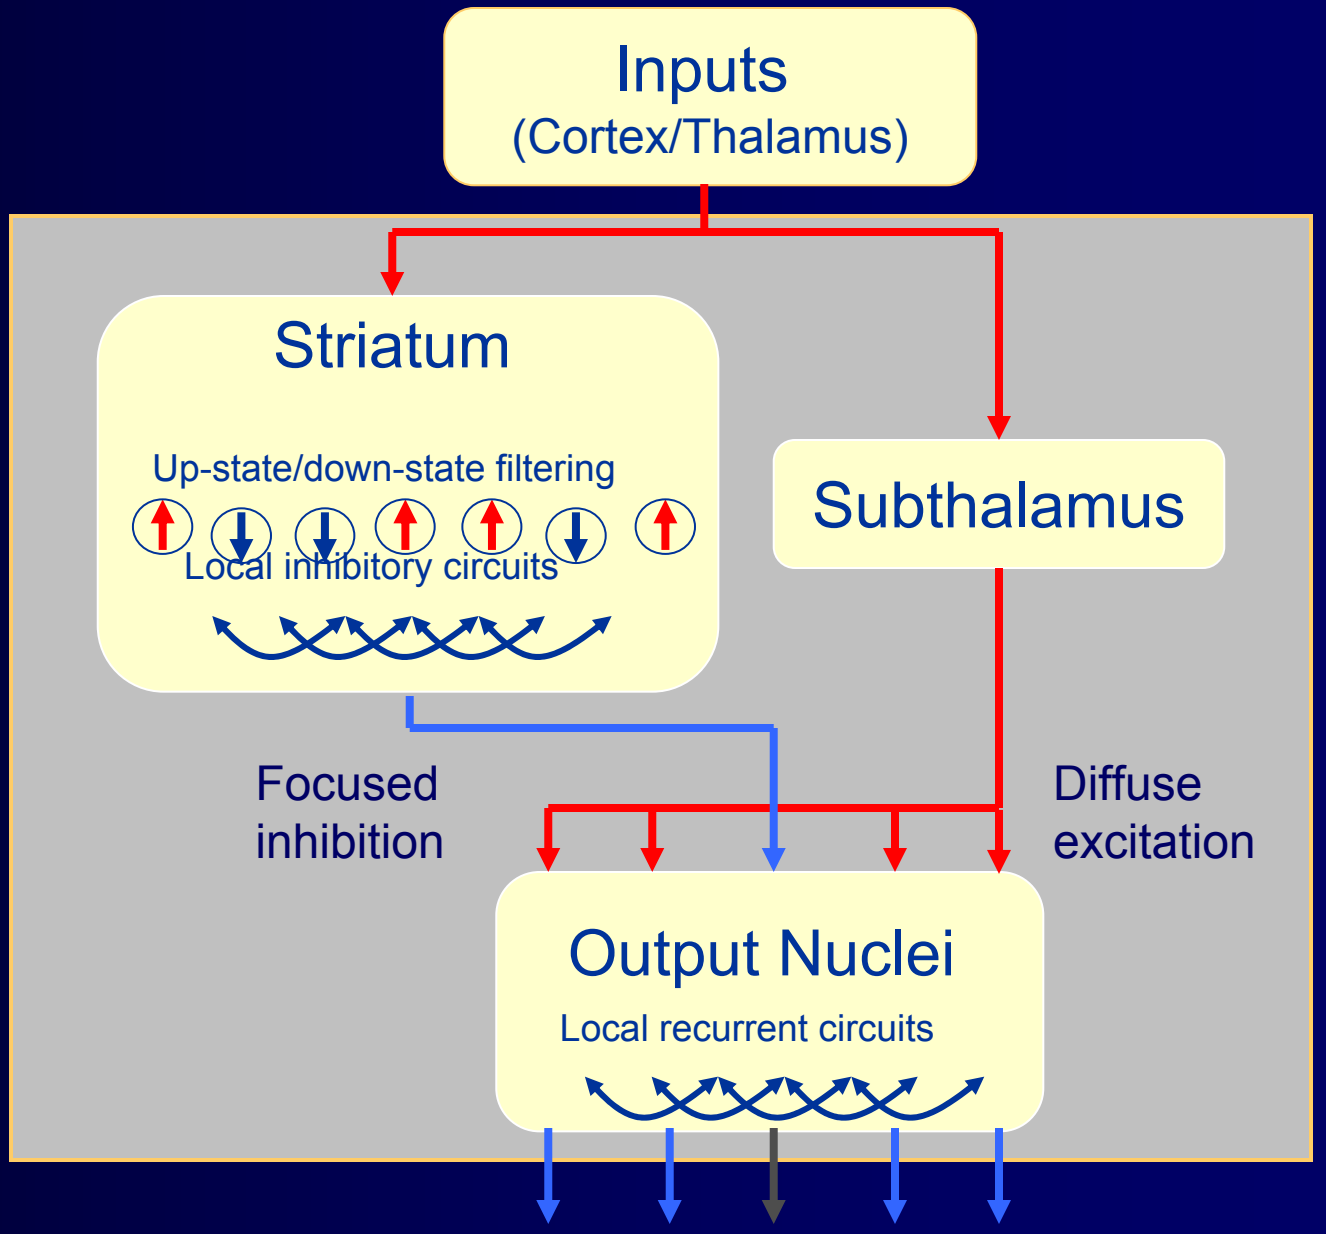

Qualitative model: PR

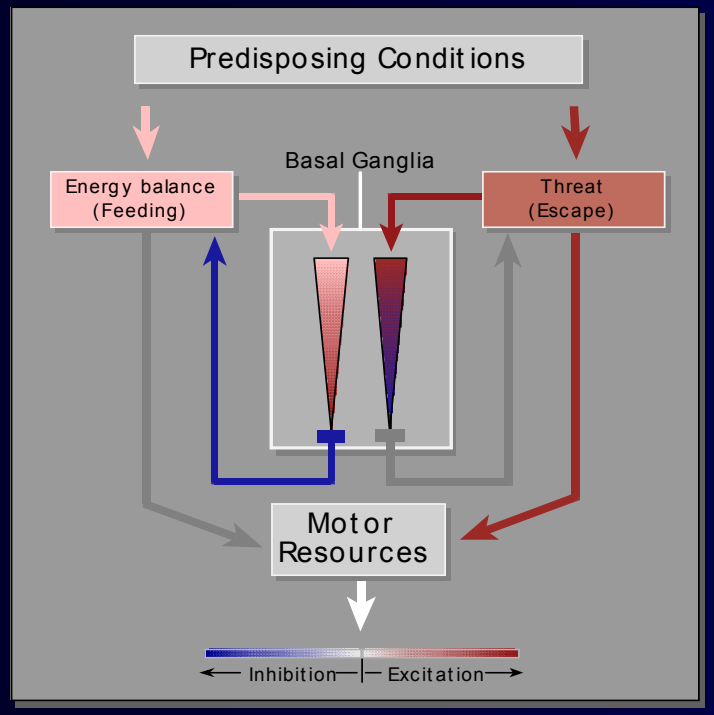

Gurney, K., T. J. Prescott, et al. (2001). "A computational model of action selection in the basal ganglia. I. A new functional anatomy." Biol Cybern **84**: 401-410.

Quantitative analysis: Kevin Gurney

Analytic equilibrium solution

Model neurons - leaky integrators with piecewise linear output

striatum - control pathway

$$H[c_i - \epsilon/w_s(1-\lambda)] \equiv H_i^{\uparrow}(\lambda)$$

$$x_i^{e-} = m^-[w_s(1-\lambda_e)c_i - \epsilon]H_i^{\uparrow}(\lambda_e)$$

striatum - selection pathway

$$x_i^{g-} = m^-[w_s(1+\lambda_g)c_i - \epsilon]H_i^{\uparrow}(-\lambda_g)$$

STN

$$\begin{aligned} x_i^{+} &= m^+(w_t c_i + \epsilon' - w_g y_i^e)H_i^{+\uparrow} \\ H_i^{+\uparrow} &= H(w_t c_i + \epsilon' - w_g y_i^e) \end{aligned}$$

GPe

$$\begin{aligned} \tilde{a}_i^e &= w^-(\delta X^+ - x_i^{e-}) + \epsilon_e \\ y_i^e &= m^e \tilde{a}_i^e H(a_i^e) \end{aligned}$$

GPi/SNr

$$\begin{aligned} \tilde{a}_i^g &= w^-(\delta X^+ - x_i^{g-}) - w_e y_i^e + \epsilon_g \\ y_i^g &= m^g \tilde{a}_i^g H(a_i^g) \end{aligned}$$

Solving for STN excitation

$$X^+ = \frac{n}{1 + \delta w_g w^- n \phi_{m,s}} \left\{ w_t \phi_{*,s} \langle c \rangle_*^s + \phi^s \epsilon' + \phi_{q,s} w_g w^- [(1-\lambda_e) w_s \langle c \rangle_{q,s} - \epsilon] - w_g \phi^s \epsilon_e \right\}$$

....

Action selection in a spiking network model

Humphries, M. D., Stewart, R. D. & Gurney, K. N. A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *J Neurosci* **26**, 12921-42 (2006).

Mean firing rate across all SNr neurons in each channel

Mean firing rate of all cortical inputs to each channel

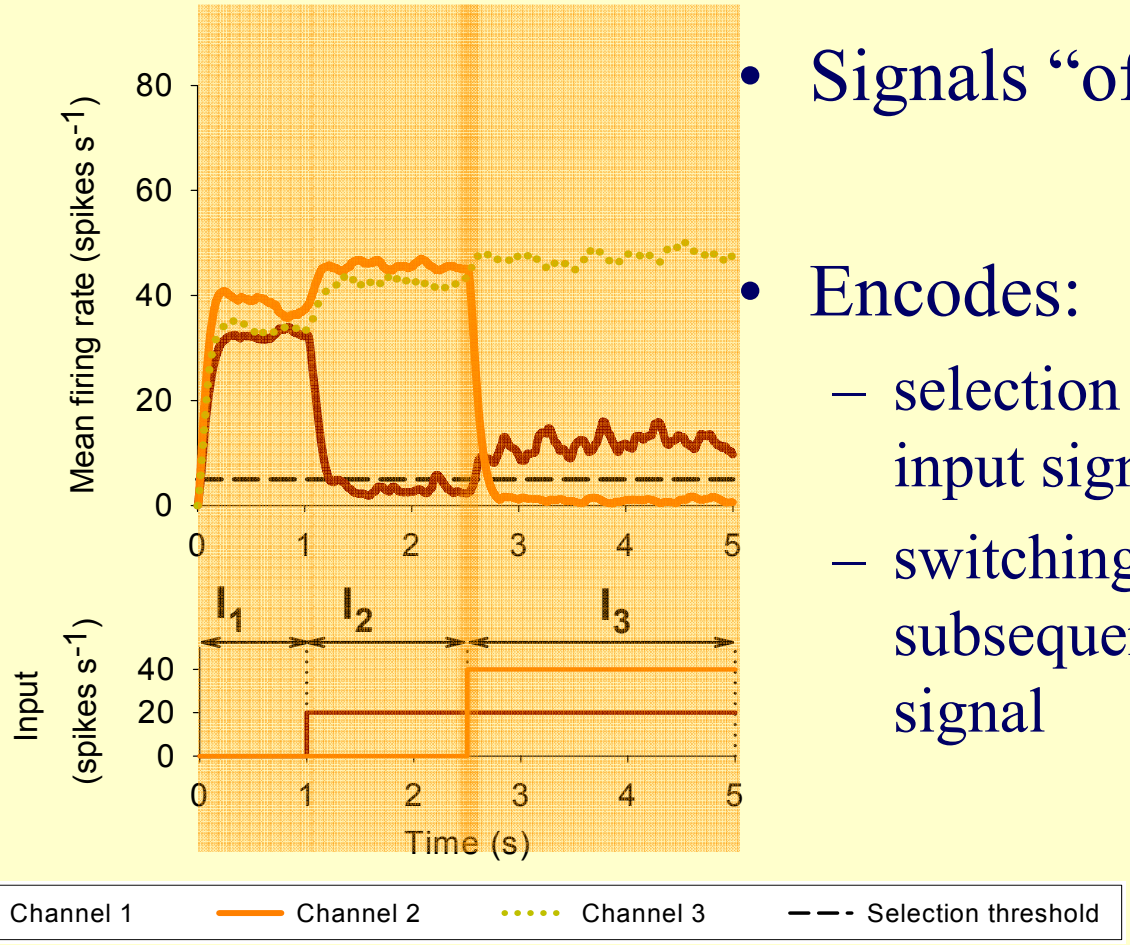

- Signals "off" state
- Encodes:
  - selection of largest input signal
  - switching to larger subsequent input signal

…but will it work in the real world ?

- The model was embodied to:
  - Generate realistic (environmentally driven) sequences of input
  - Force interpretation of outputs in terms of actions

- Aim: To test if model can generate action sequences in a behaving robot
  - Research sought to model behavioural switching in a foraging rat

Prescott TJ, Gonzalez FMM, Gurney K, Humphries MD, Redgrave P. 2006. A robot model of the basal ganglia: Behavior and intrinsic processing. Neural Networks 19(1):31-61.

Action Selection: Rat foraging

- Motivations
  - Hungry : 24hrs food deprived
  - Frightened: placed in open arena
- Behaviour
  - Initially keeps close to walls and corners
  - Collects food
  - Returns to corner to eat

Robot Action Selection

- Motivations
  - Hunger
  - Fear
- 5 behavioural sub-systems
  - Wall seek
  - Wall follow
  - Can seek
  - Can pick-up
  - Can deposit
- 8 Infra-red sensors detect
  - Walls
  - Corners
  - Cans
- Gripper sensors detect
  - Presence/absence of can

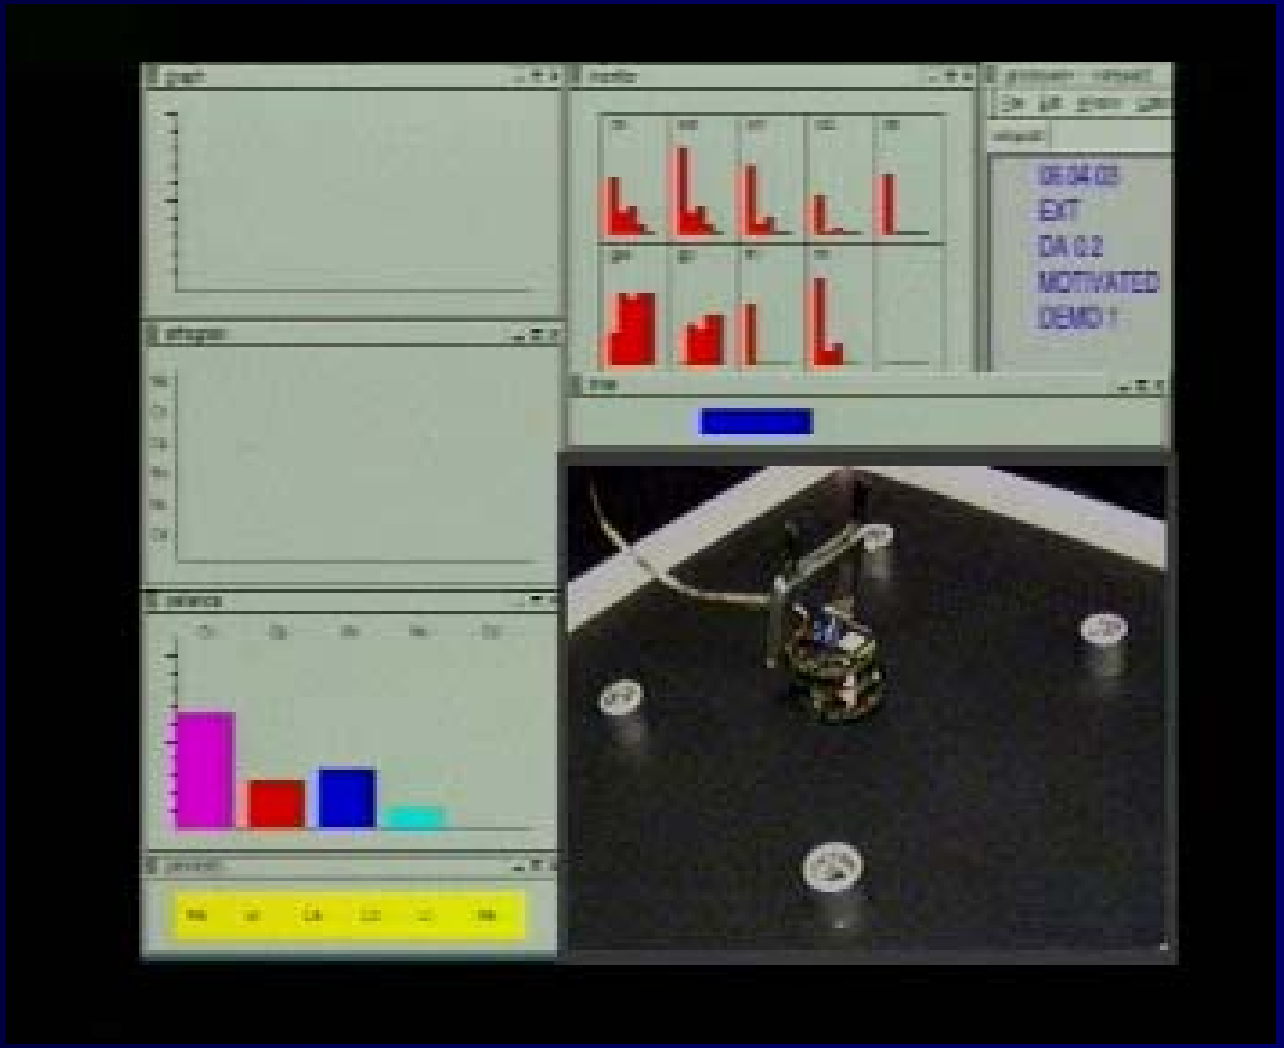

Prescott TJ, Gonzalez FMM, Gurney K, Humphries MD, Redgrave P. 2006. A robot model of the basal ganglia: Behavior and intrinsic processing. Neural Networks 19(1):31-61.

Conclusions so far….

- Selection hypothesis of basal ganglia architecture confirmed in analysis, simulation and control of robot action selection
- Consistent with early development and evolutionary conservation
- Represents a generic task performed in all functionally segregated territories of the basal ganglia
  - Selection of overall behavioural goal (limbic)
  - Selection of actions to achieve selected goal (associative)
  - Selection of movements to achieve selected actions (sensorimotor)

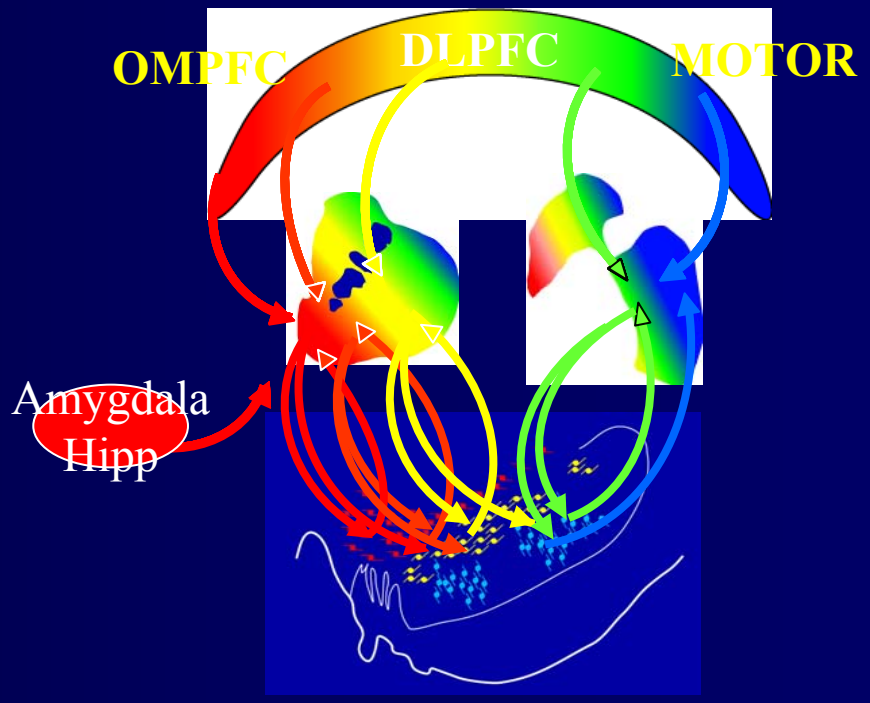

Haber et.al J.Neurosci. 2000

Part 2

Adaptive function(s) in the basal ganglia

Reinforcement biases action selection

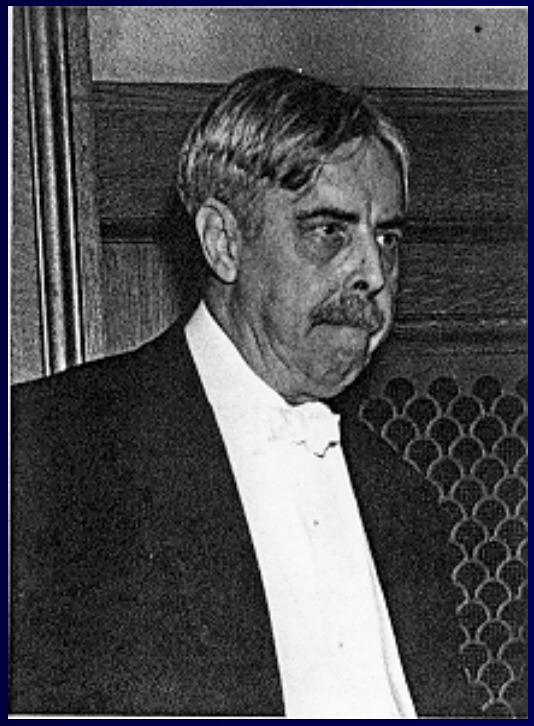

E.L. Thorndike (1898)

Law of Effect

"If the response in the presence of a stimulus is followed by a satisfying event, the association between the stimulus and the response is strengthened. If the response is followed by an annoying event, the association is weakened".

Mechanisms of bias

For action selection to adapt with experience, experience must bias future selections

Possible mechanisms

- Increase strength of reinforced 'bid' at source
- Increase sensitivity to 'bid' in BG input nuclei

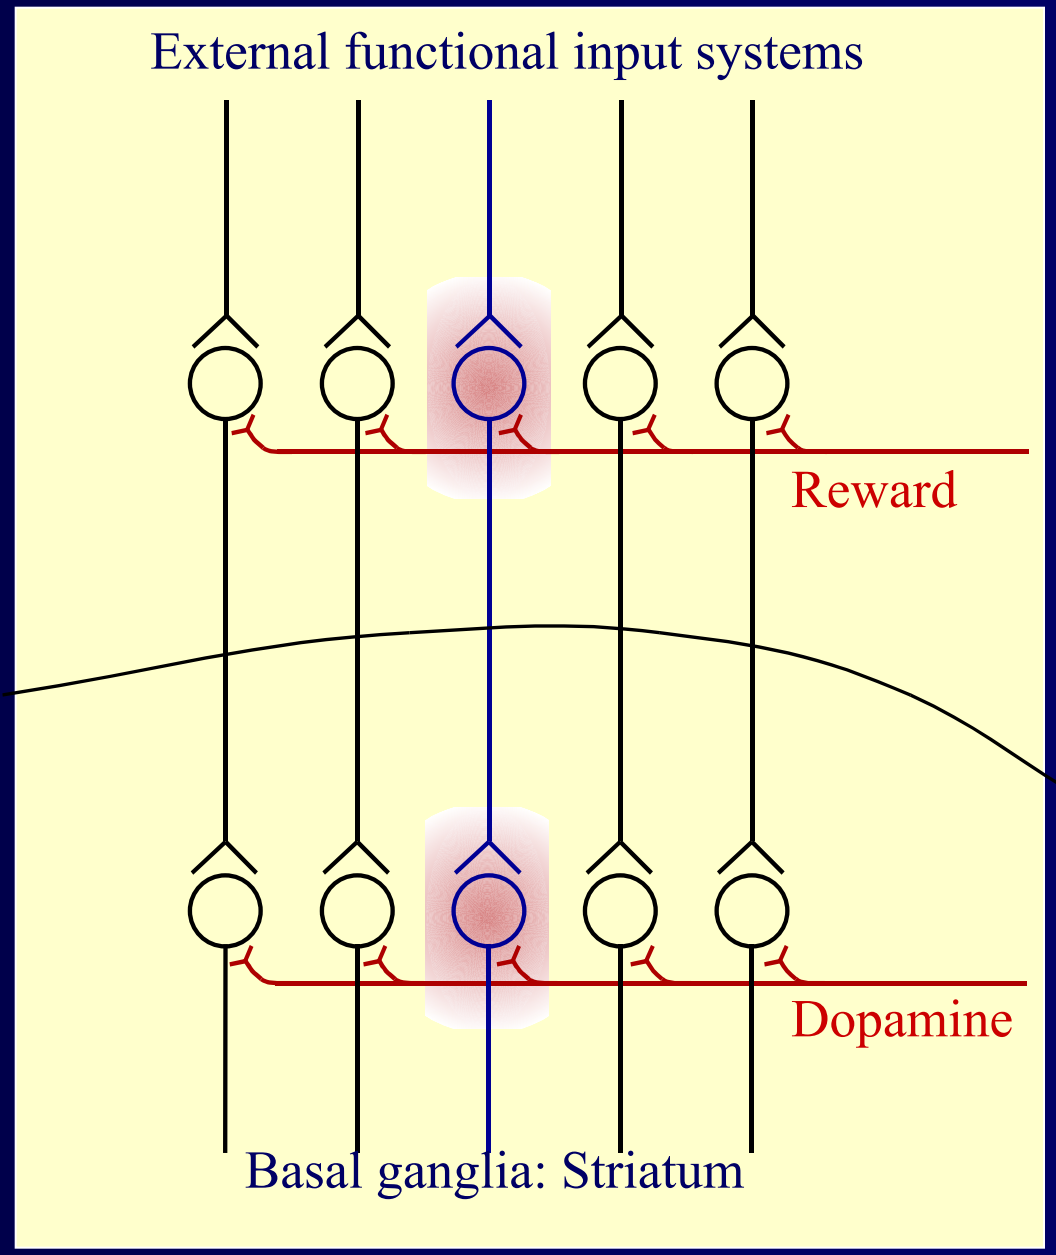

A critical role for the ascending dopamine systems

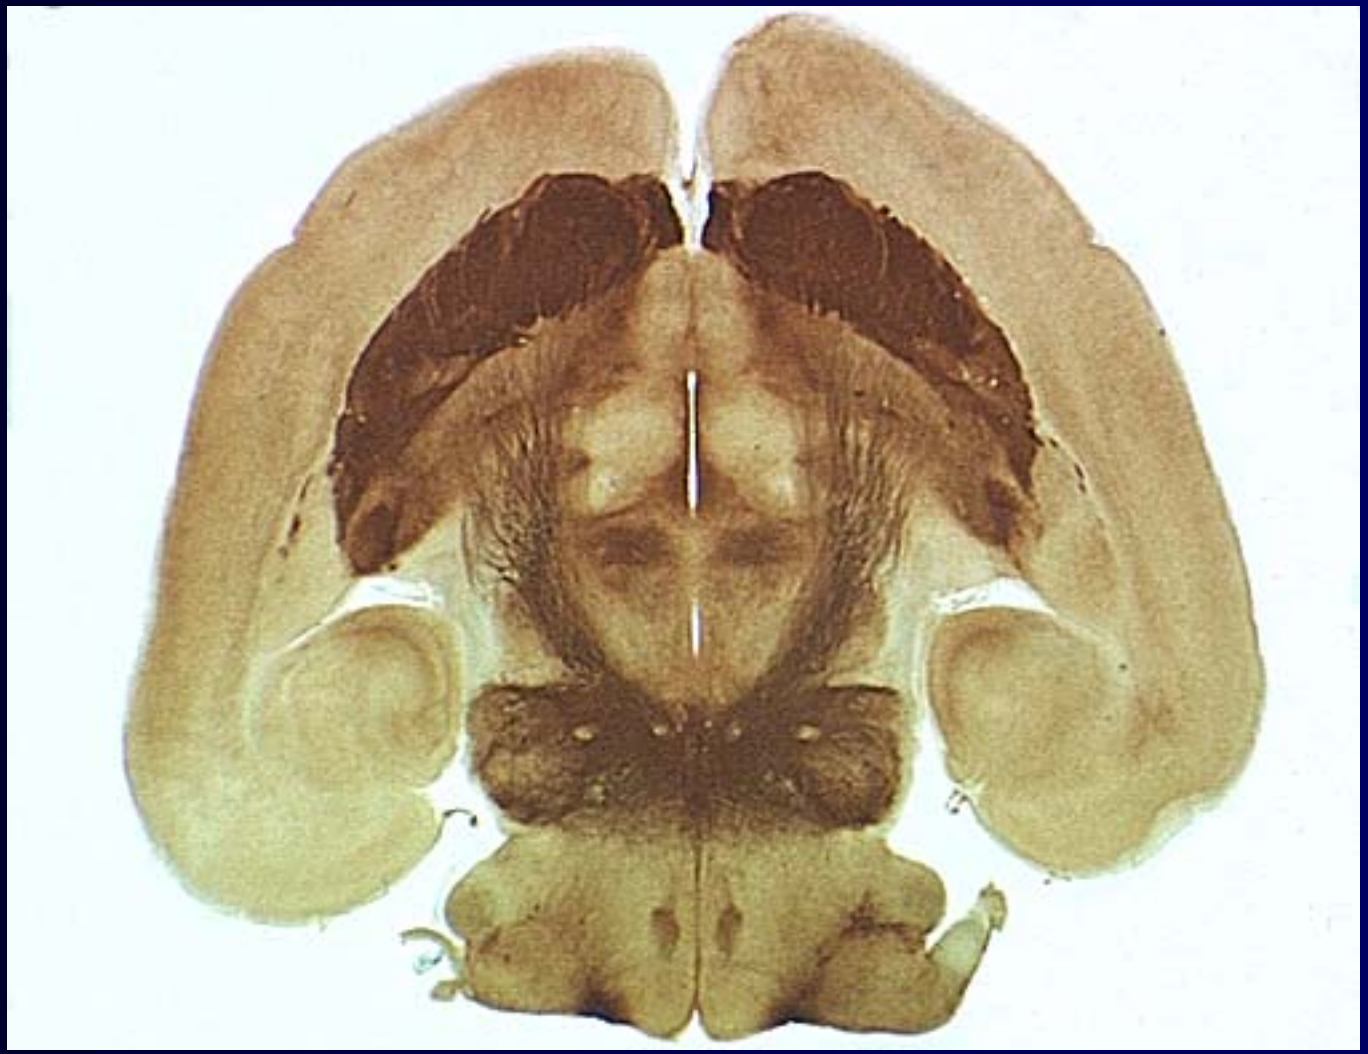

Picture by Wes Chang (Gallo centre San Francisco)

The phasic dopamine signal

- Short latency (70-100ms)
- Short duration (~ 100ms) burst of impulses
- Elicited by biologically salient stimuli

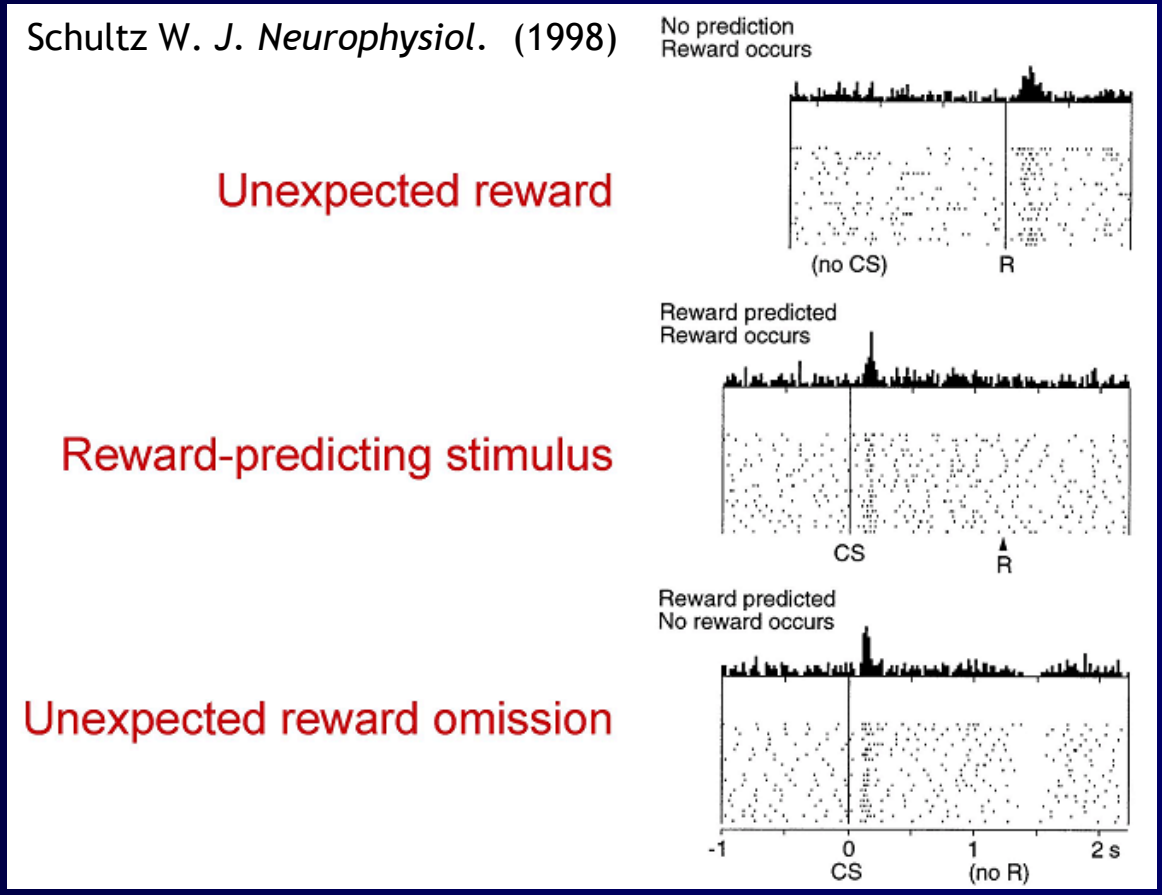

Reward prediction errors

- Phasic DA signals similar to reward prediction error term in the temporal difference reinforcement learning algorithm (Barto, Montague, Dayan)
- Reward prediction errors = unexpected sensory events that are ‘better’ or ‘worse’ than predicted
- Phasic DA responses provide training signals for both Pavlovian and instrumental associative learning
- Increase probability of selecting responses to maximise future reward

Midbrain superior colliculus

Caudal

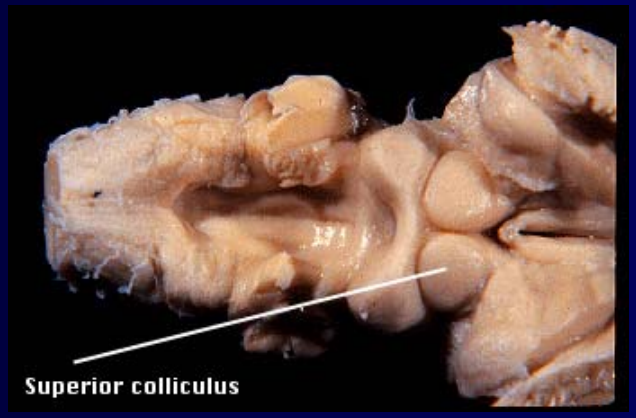

Rostral

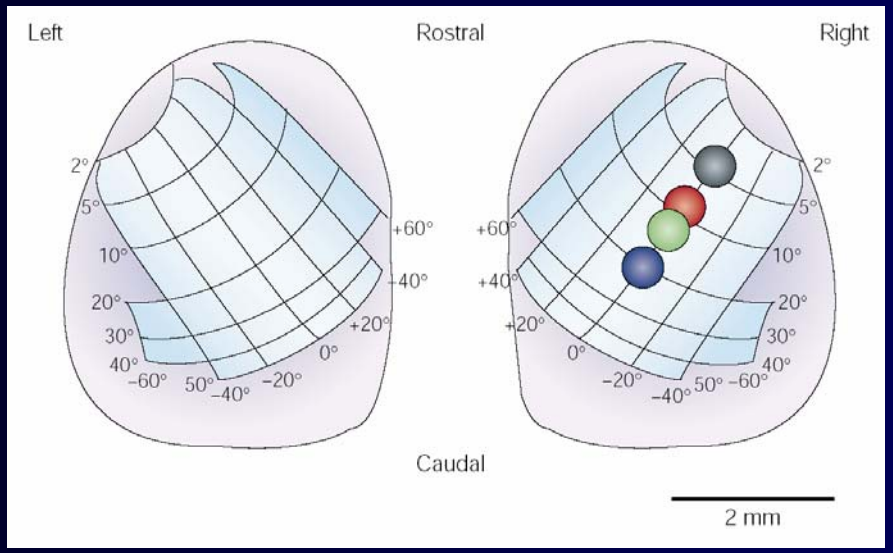

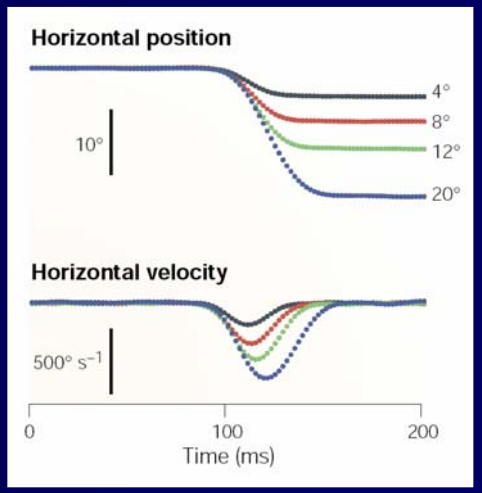

Evoked eye movements bring events onto the fovea

Sparks DL. 2002. The brainstem control of saccadic eye movements. Nature Reviews Neuroscience 3:952-964.

The latency constraint

Unexpected visual stimuli elicit sensory and motor responses in superior colliculus:
- sensory response (~40 ms)
- motor response (<150 ms) → Gaze-shift

Phasic DA responses occur before signals eliciting foveating gaze-shift
- 70-100ms after stimulus onset

Event judged 'better' or 'worse' than expected before it's brought to the fovea to be identified

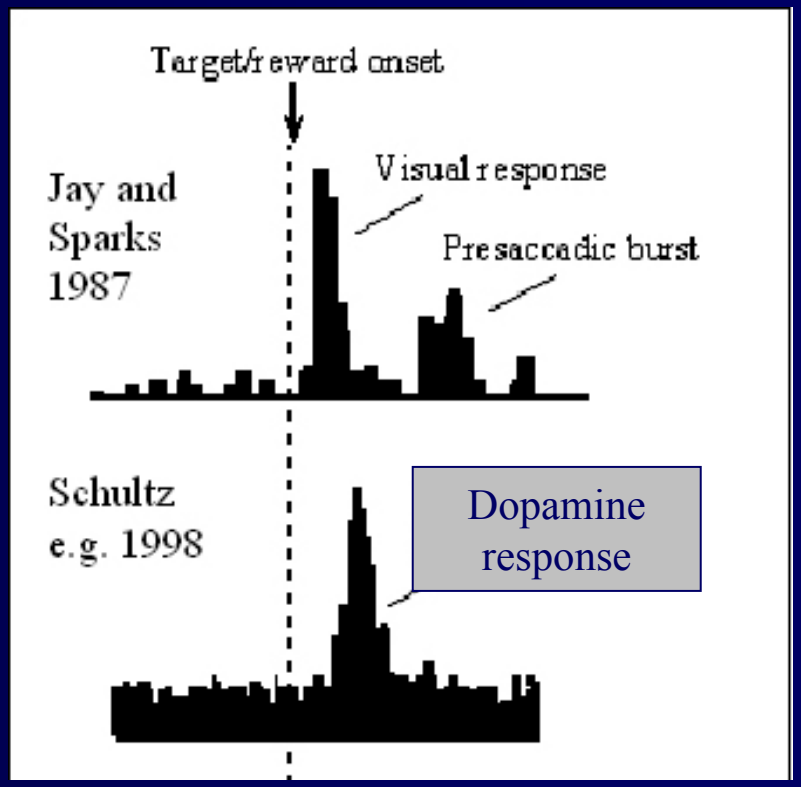

- <u>Conclusion:</u>
  Visual input to DA neurones result of pre-attentive, pre-saccadic stimulus processing

Redgrave P, Prescott TJ and Gurney K (1999). TINS 22(4): 146-151

Specific Research Question

- Since DA signals depend on perceptual capacities of pre-gaze-shift visual processing…..
- What are afferent visual projections to DA neurones ?
- Two possibilities
  - Retino-geniculo-striate cortical projection system
  - Retino-tectal projection system

Geniculostriate-cortical projection

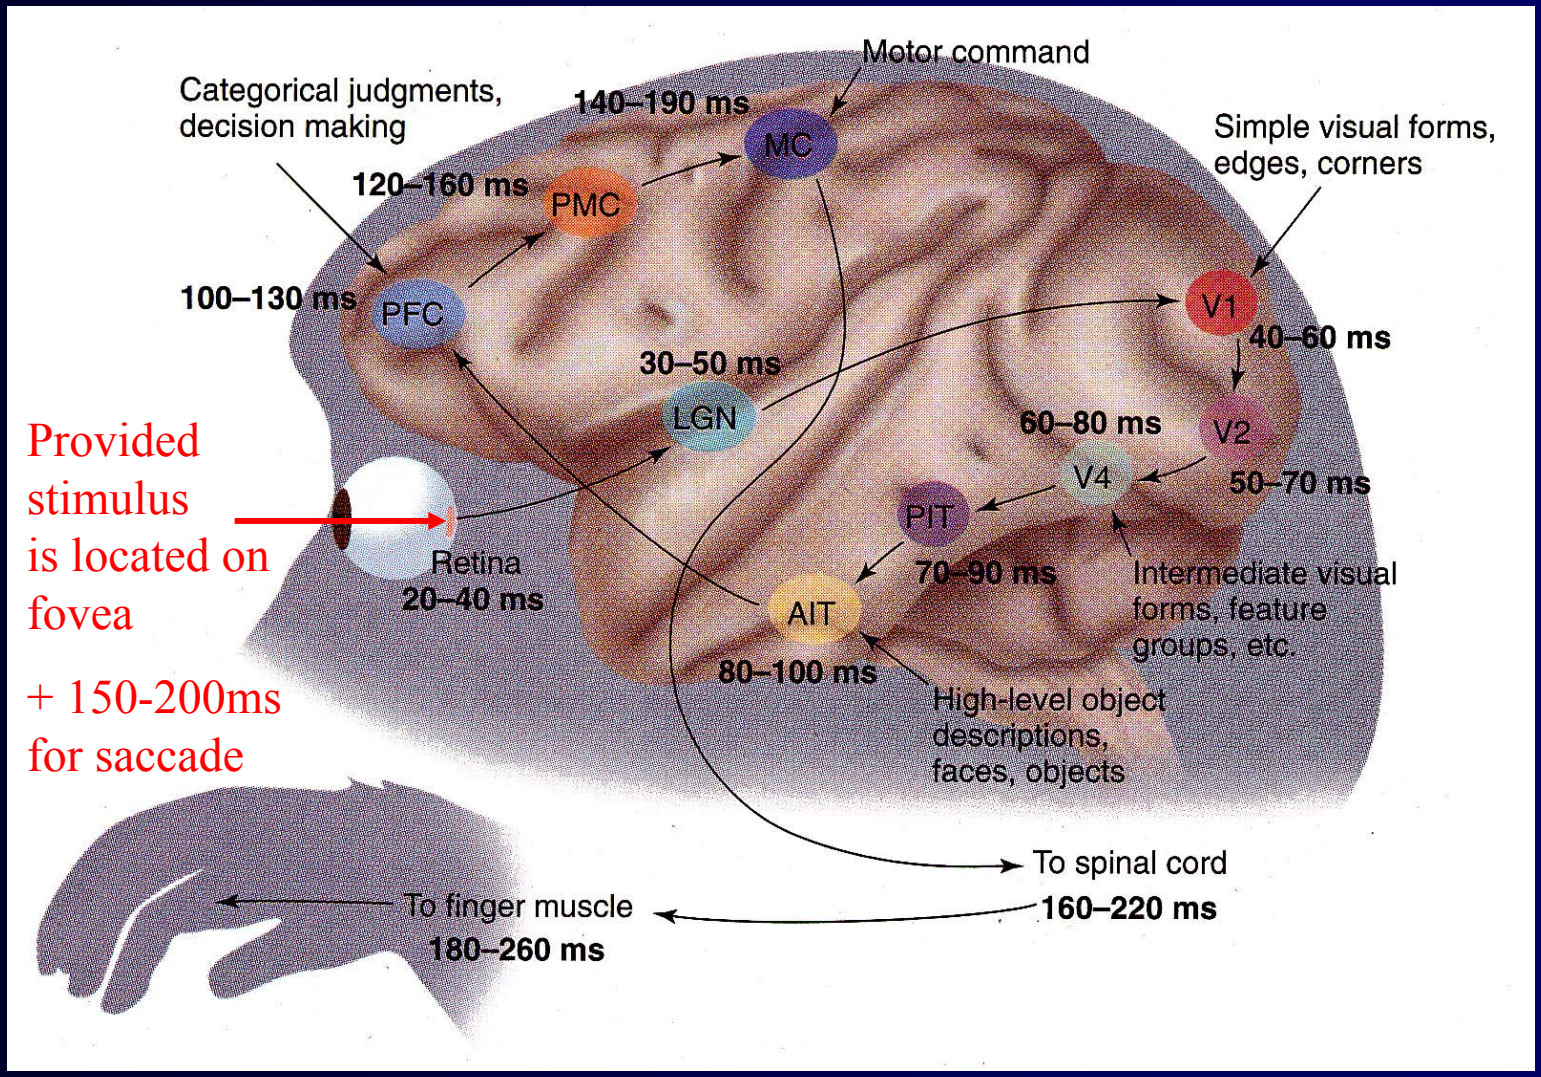

Thorpe SJ and Fabre-Thorpe (2001) Science, 291, 260-163

Subcortical retino-collicular projection

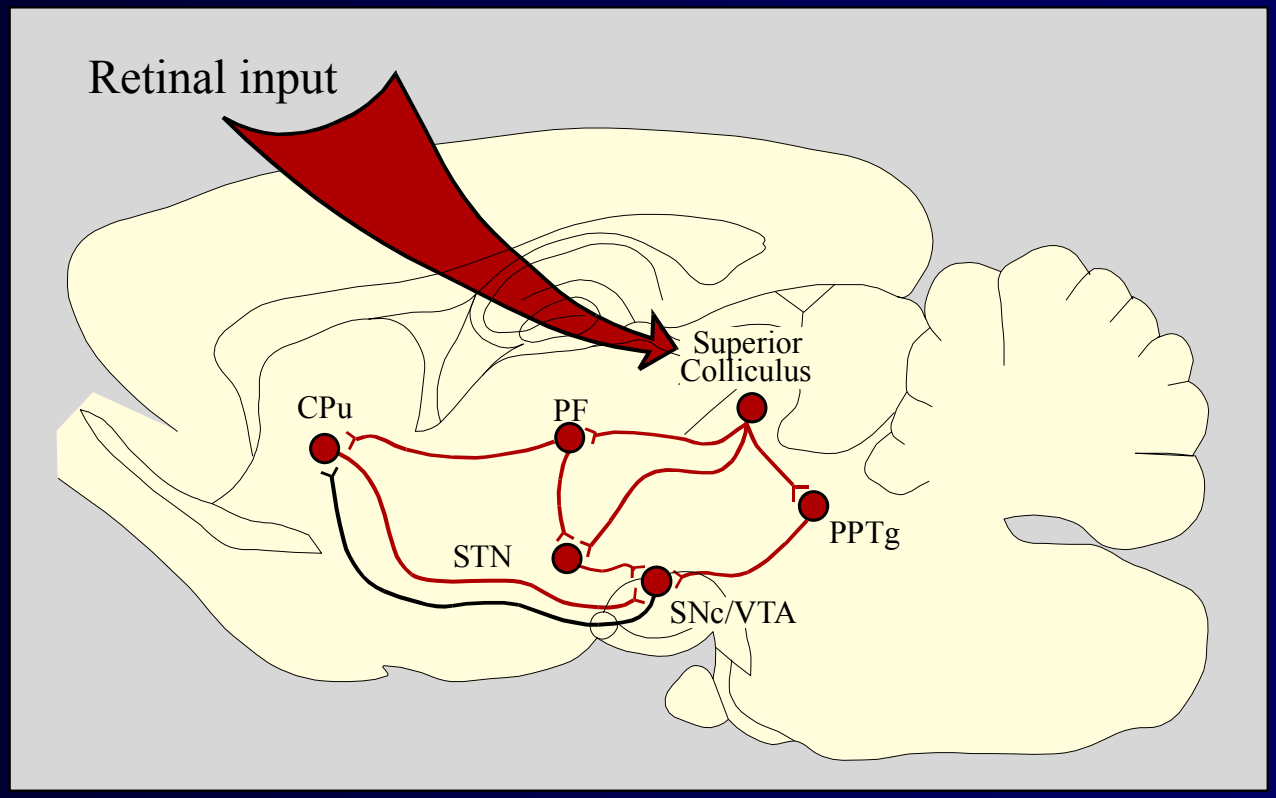

Latencies from stimulus onset in superior colliculus ~ 40-50ms

Conclusion : retino-collicular route the more likely route

Colliculus as the source of visual input: I

Anatomical Evidence

- The Tectonigral projection
- Direct pathway discovered from superior colliculus to substantia nigra pars compacta
- ...in rat (Comoli, et al. 2003 Nature Neurosci **6**: 974-980)
- …cat (McHaffie, et al 2006 Neuroscience
- …and monkey (Redgrave, Haber et al – work in progress

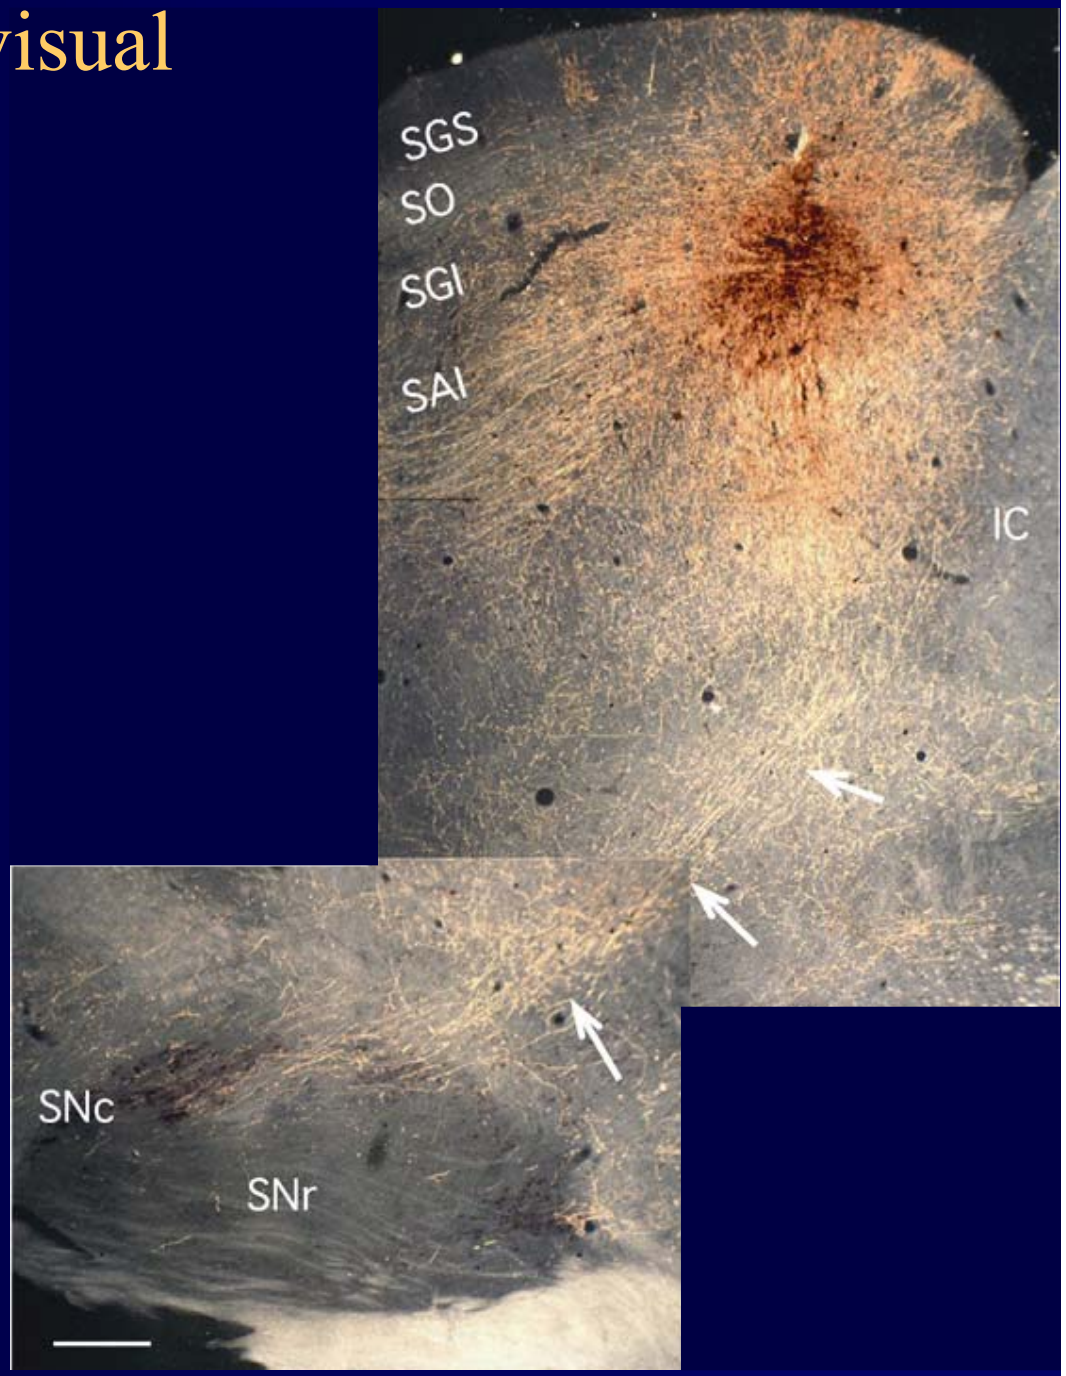

Superior Colliculus: a critical visual relay ?

Visual Evoked Potentials (VEPs)

What is source of flash-evoked potentials recorded from substantia nigra pars compacta ?

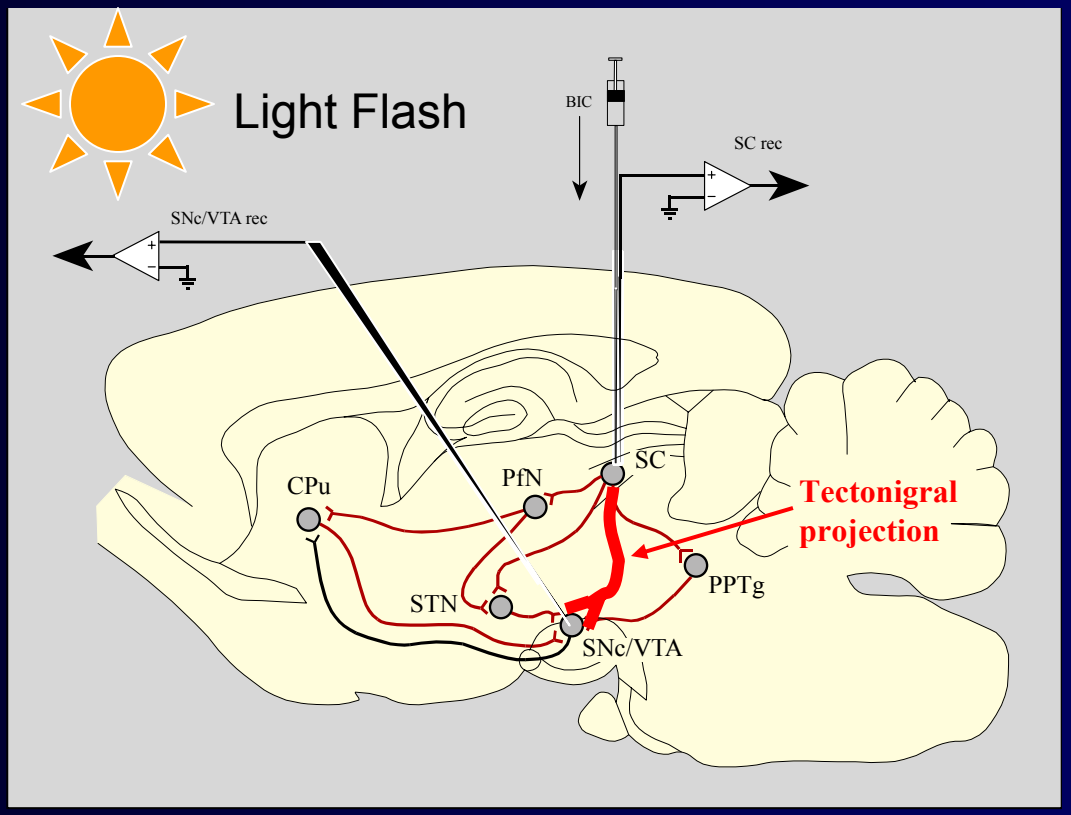

Comoli, et al. (2003). *Nature Neurosci* **6**: 974-980.

Visual relay: VEP latencies

- Whole-field light flash
  - Onset and peak latencies of VEP in SC significantly short than in SNc (n=12)
- Implication
  - Visual information in SC could be source of afferent input to SNc

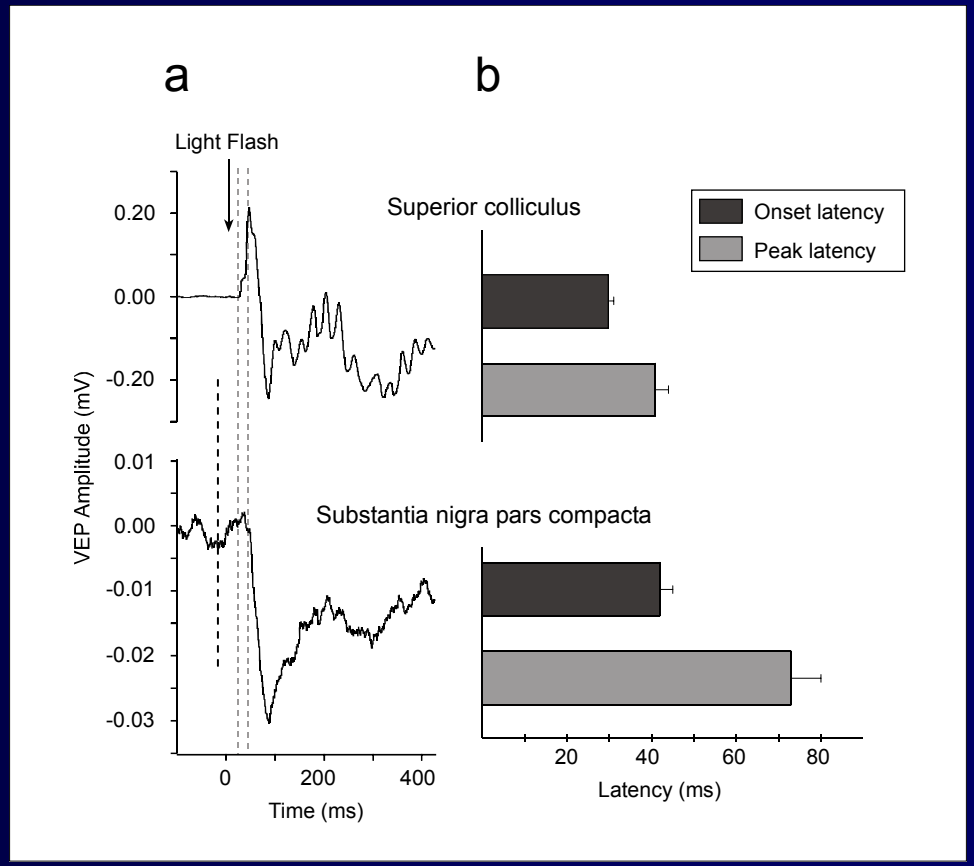

Comoli, et al. (2003). *Nature Neurosci* **6**: 974-980.

Visual system lesions

- Visual cortex aspiration
  - VEP in SC suppressed – Sprague effect
  - VEP in SNc suppressed
- Collicular bicuculline
  - In complete absence of visual cortex - VEPs reinstated in SC
  - Reinstated/potentiated VEP in SNc
- Superficial SC aspiration
  - VEP in SNc abolished
  - Not reinstated by bicuculline
- Conclusion
  - Colliculus critical relay

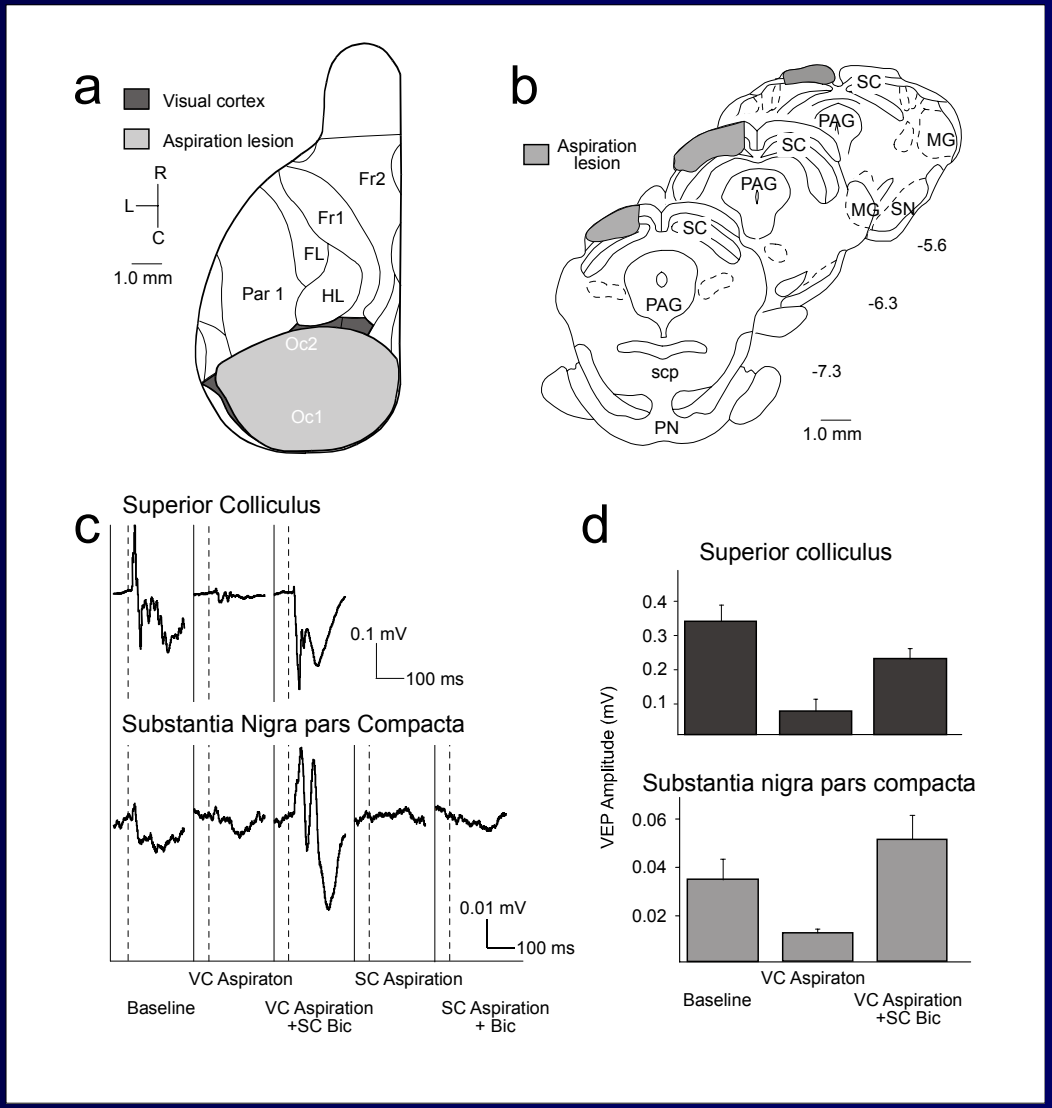

Colliculus as the source of visual input: II

Electrophysiological Evidence

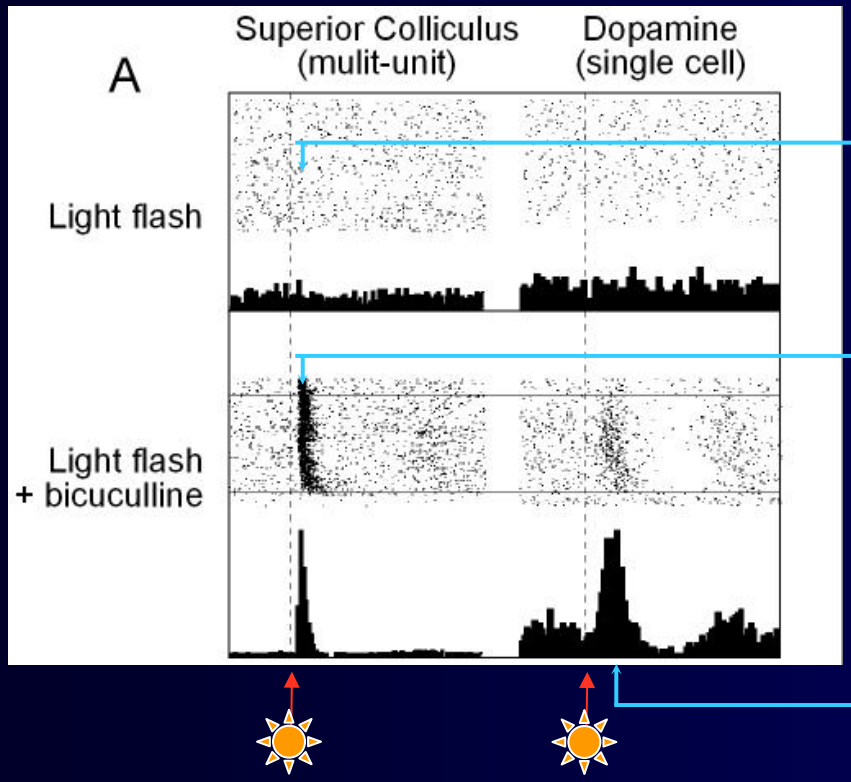

- Pre-drug baseline
  - No flash-evoked response in deep SC or DA cells
- After BIC into deep SC
  - local neurones responsive to light
- When SC cells 'see' so do DA cells
  - Excitatory responses: 17/30 (56.6%)

Dommett E, Coizet V, Blaha CD, Martindale J, Lefebvre V, Walton N, Mayhew JE, Overton PG, Redgrave P. 2005. How visual stimuli activate dopaminergic neurons at short latency. Science 307(5714):1476-1479.

Visual activation of DA: Conclusion

- Cortical visual systems neither necessary nor sufficient for phasic activation DA neurones

- Pre-tectal and accessory optic systems – ocular reflexes or responses to photoperiod

- Conclusion:
  Superior colliculus is primary if not exclusive source of short latency visual input to DA neurones

Collicular activation of DA: Implications

- Visual processing in colliculus – exquisitely sensitive to luminance onset/offset or movement within its retinotopic map

- Colliculus largely blind to static contrast, colour or geometric configuration

…but DA neurones sensitive to high spatial frequencies and colour

DA neurones discriminate magnitude/probability of reward-predicting stimuli differing in colour and high spatial frequency geometric configuration (e.g.Tobler et al 2005 *Science*, 307, 1642-5)

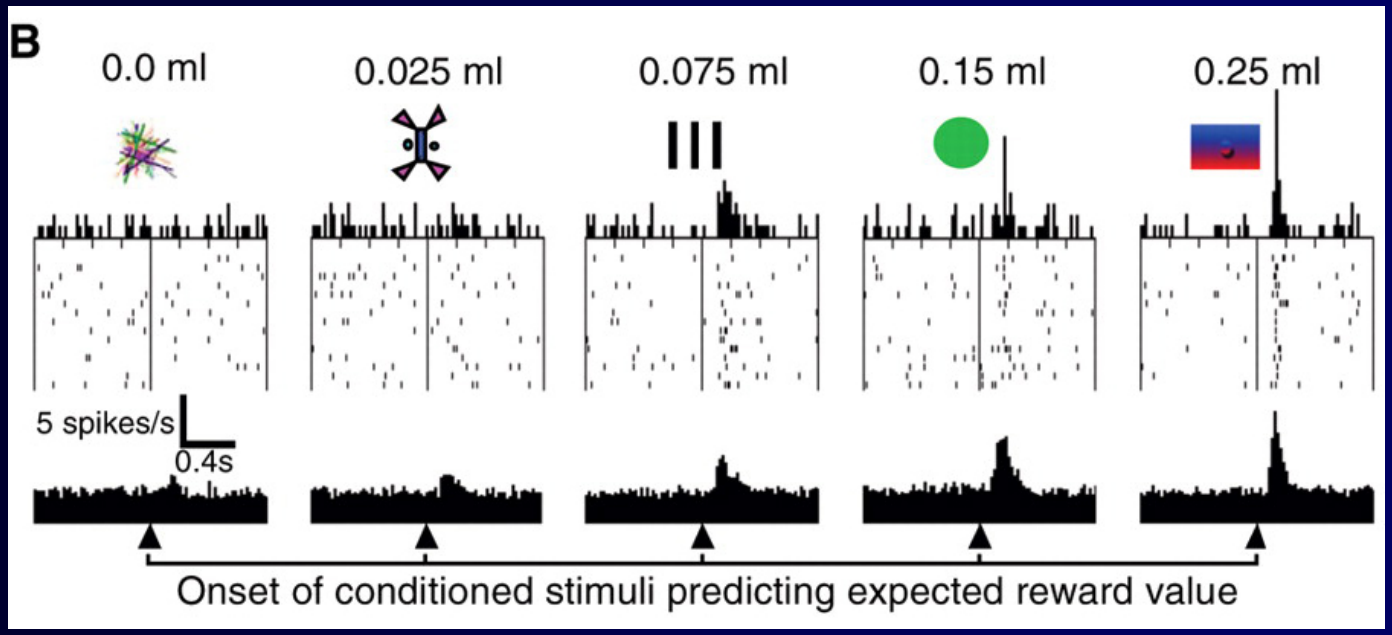

…so how do they do it ?

Read the methods sections !

- "Training consisted of 100-200 trials/stimulus/day, 5days/week, for ~ 5weeks."
  = 2500-5000 trials/stimulus
  = 12,000-25,000 stimulus/reward pairings

- "To aid discrimination, each stimulus was presented at a unique location on the computer monitor."

Training differentially sensitises different regions of the spatial map in the colliculus

…requires stimuli to be presented at the same location

DA responses to unpredicted non-reward

Responses to phasic novel events reported

…informally

"Effective stimuli include: 1) novel, unexpected stimuli eliciting orienting reactions….."

Ljungberg et al. J. Neurophysiol. 1992

"We also noticed that DA neurons typically responded to a visual or auditory stimulus when it was presented unexpectedly, but stopped responding if the stimulus was repeated; a subtle sound outside the monkey's view was particularly effective."

Takikawa et al. J. Neurophysiol. 2004

…and formally

Horvitz et al. Brain Res. 1997

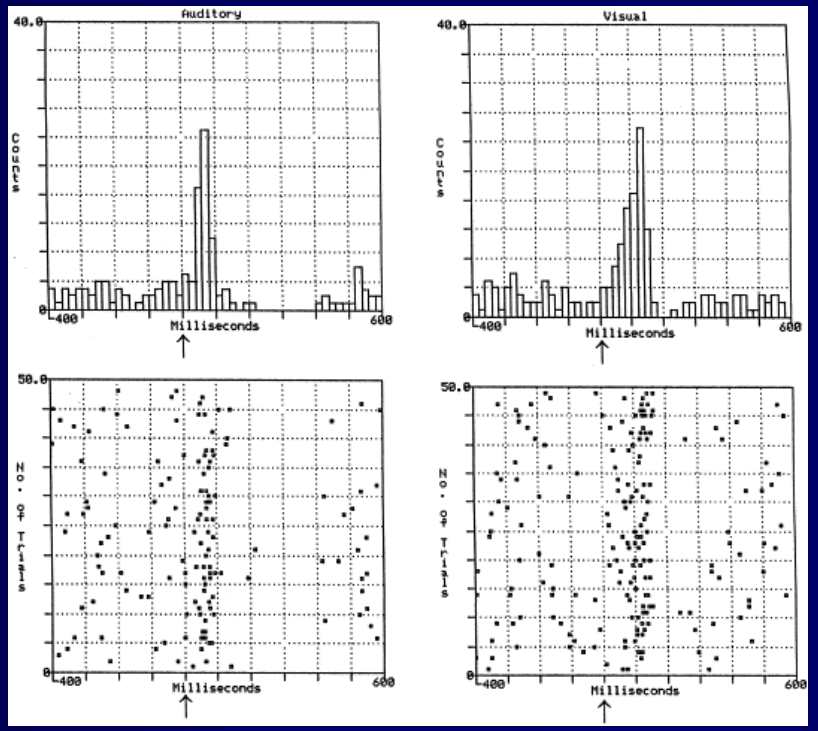

System evolved to work in natural environment

- DA signals report unpredicted novel-neutral and reward related events
- Unexpected events in ‘real world’ temporally and spatially unpredicted
- Provided with a degraded signal

Dopamine conclusions so far…..

- Colliculus registers location of luminance changes
- Afferent signals to DA communicate occurrence of biologically salient events (novel-neutral and reward related)
- …not their identity
- When stimuli are both temporally *and* spatially unpredictable…..

  ….stimulus identity will remain to be determined at the time of DA signalling

What does phasic dopamine reinforce ?

- Because afferent sensory processing limited…
- ….unlikely to reinforce selection of actions to maximise future reward
- they certainly look like teaching signals….
- ….but for what kind of learning ?

Essential characteristic of the phasic dopamine signal:
It's timing

Electrophysiology

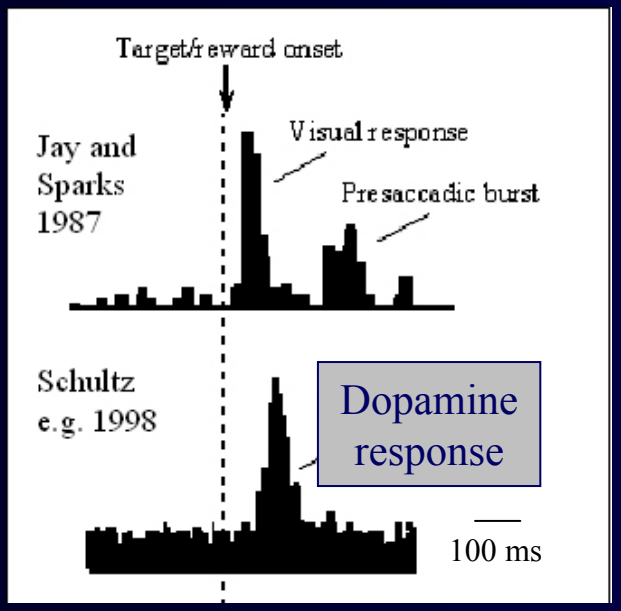

Extracellular dopamine

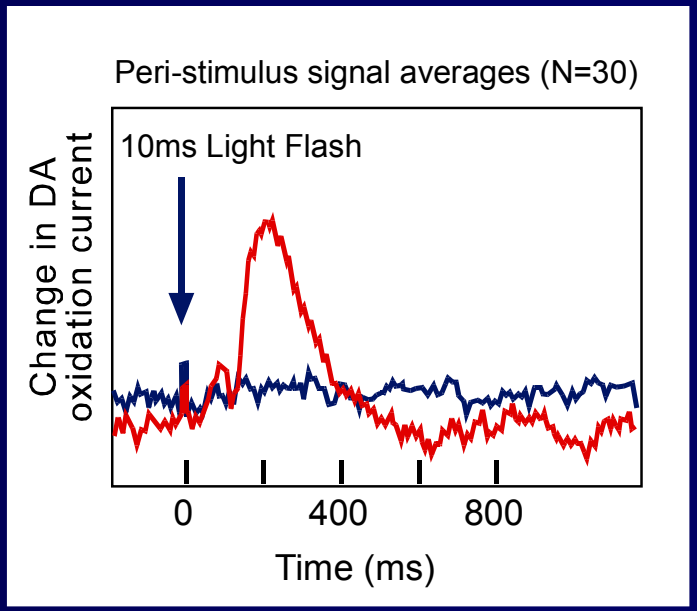

100ms latency 100ms duration response constant across:

- species
- experimental paradigms
- sensory modalities
- perceptual complexity of eliciting events

Insight

- If phasic DA responses operate like a time-stamp
- What are the signals in DA target regions at the time of the DA stamp ?
- …. these are the signals the timed dopamine input will be interacting with

1st Signal: Context

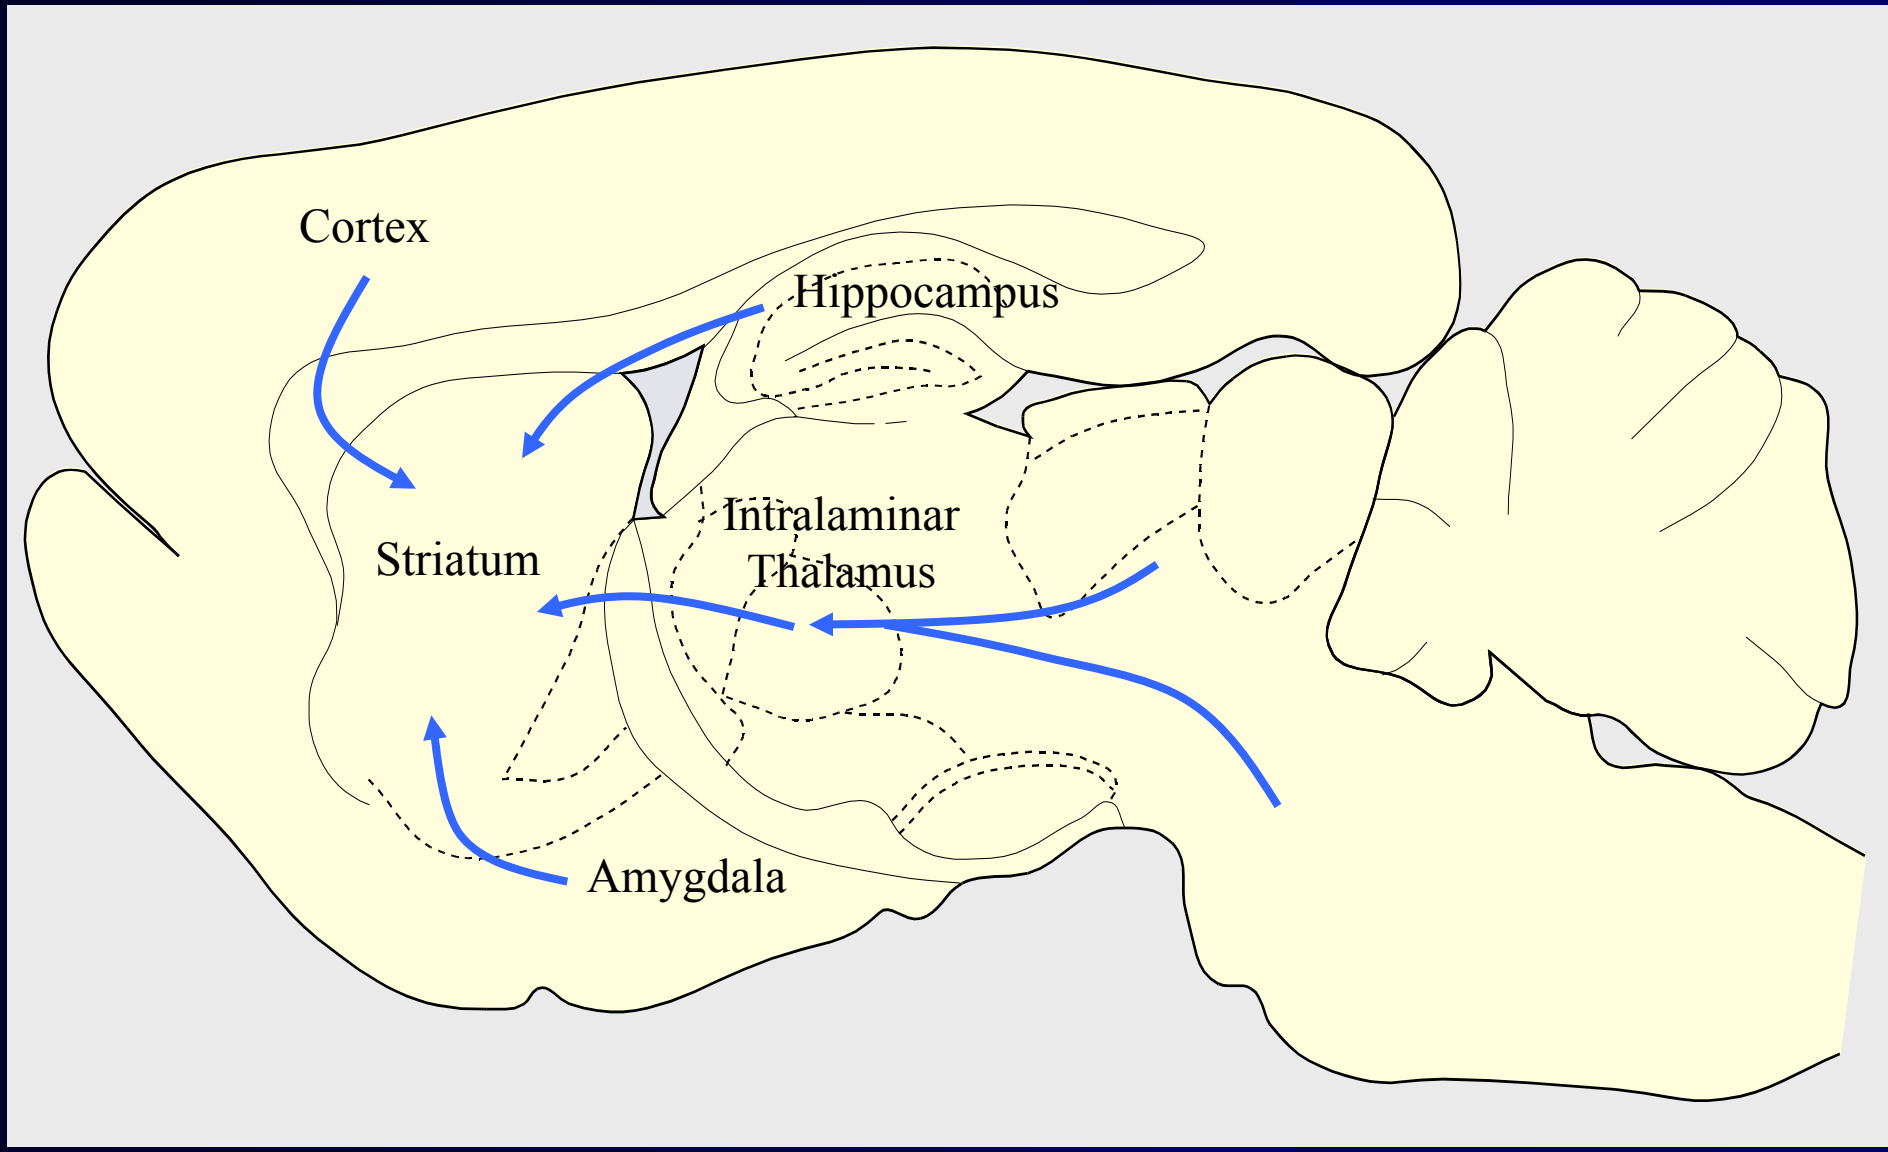

Redgrave and Gurney, *Nature Reviews Neuroscience* Dec. 2006, The short latency dopamine signal: a role in discovering novel actions

2nd Signal: Efference copies of motor commands

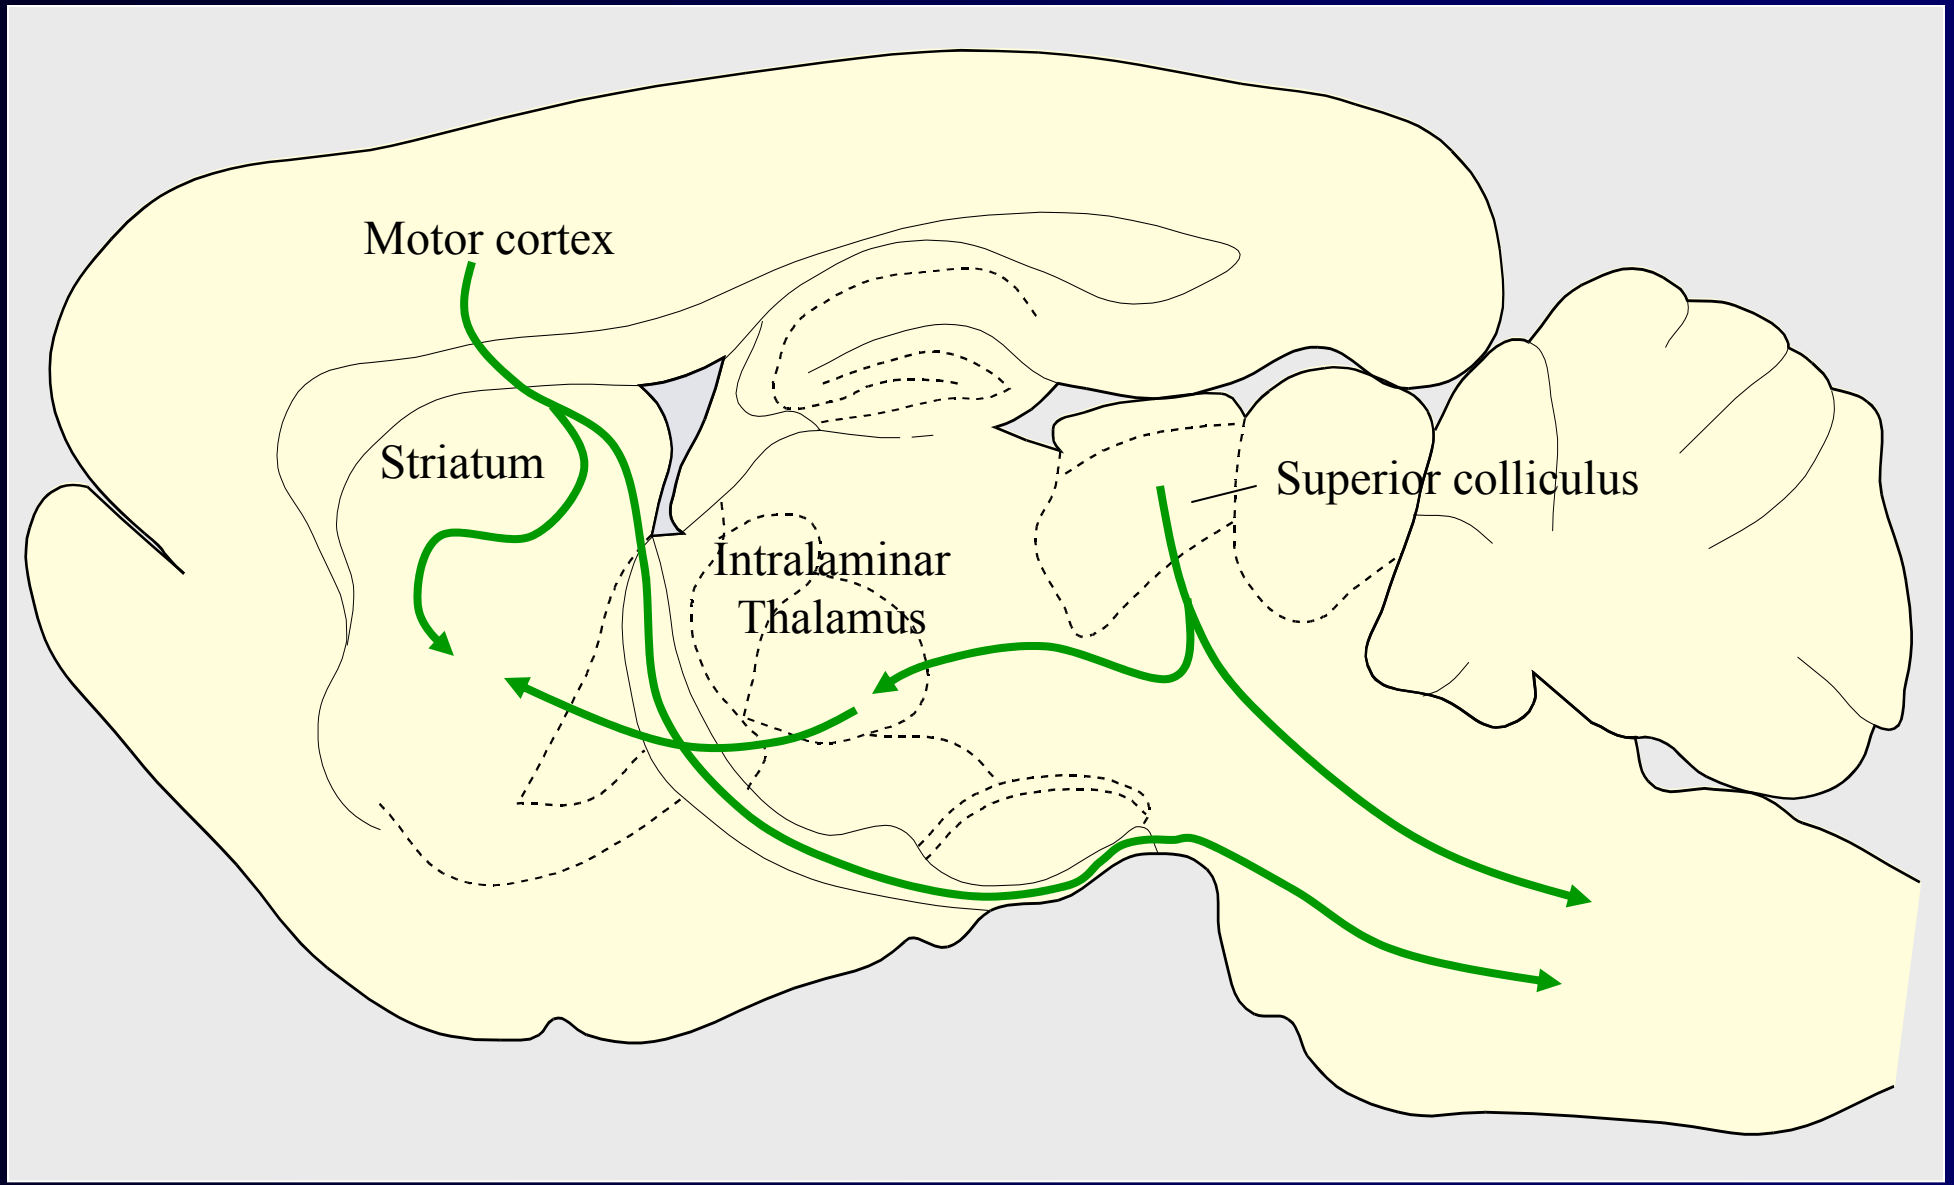

3rd Signal: Concurrent sensory signals

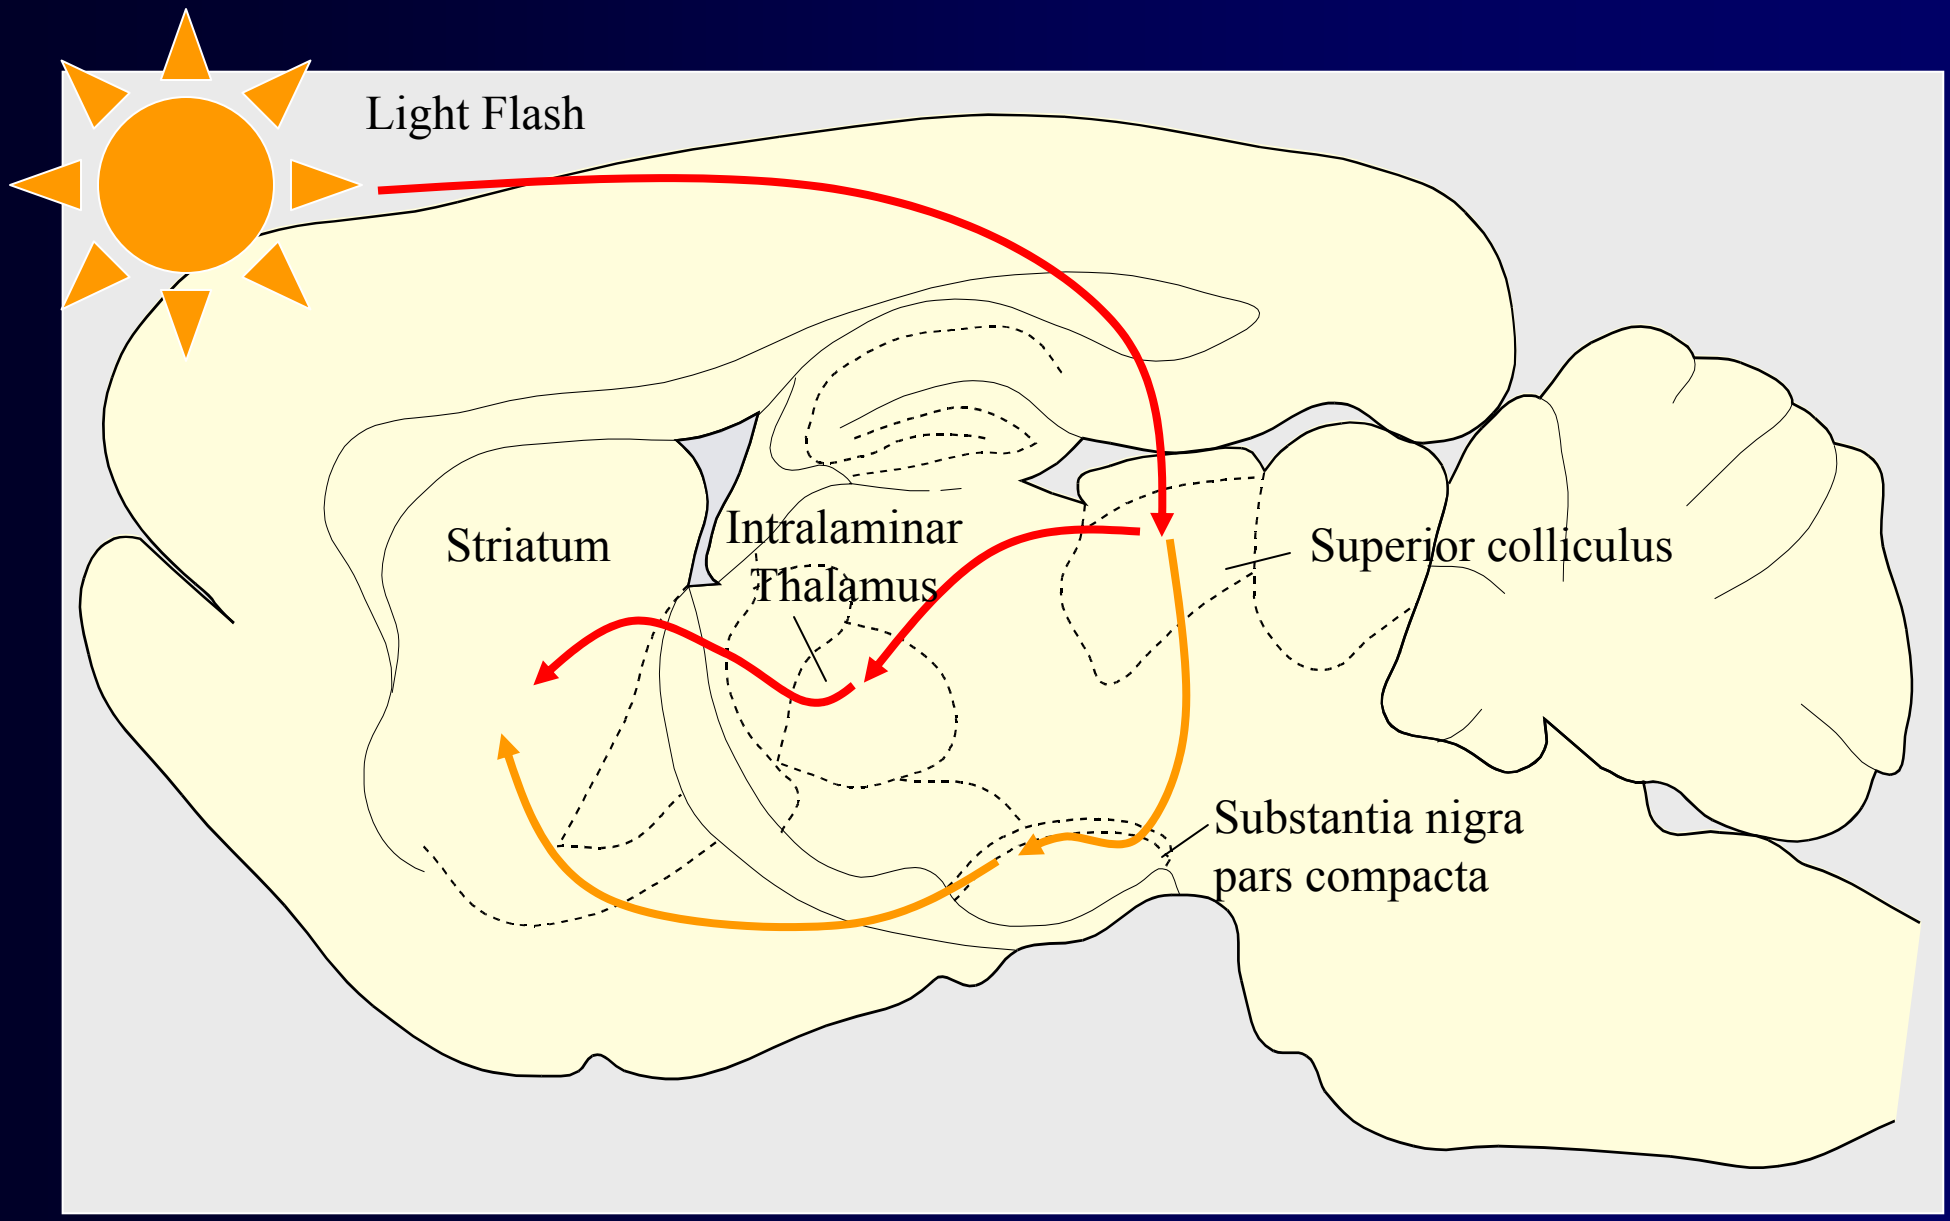

McHaffie et al *TINS* , Aug. 2005, Sub-cortical loops through the basal ganglia

4 Classes of converging Signals

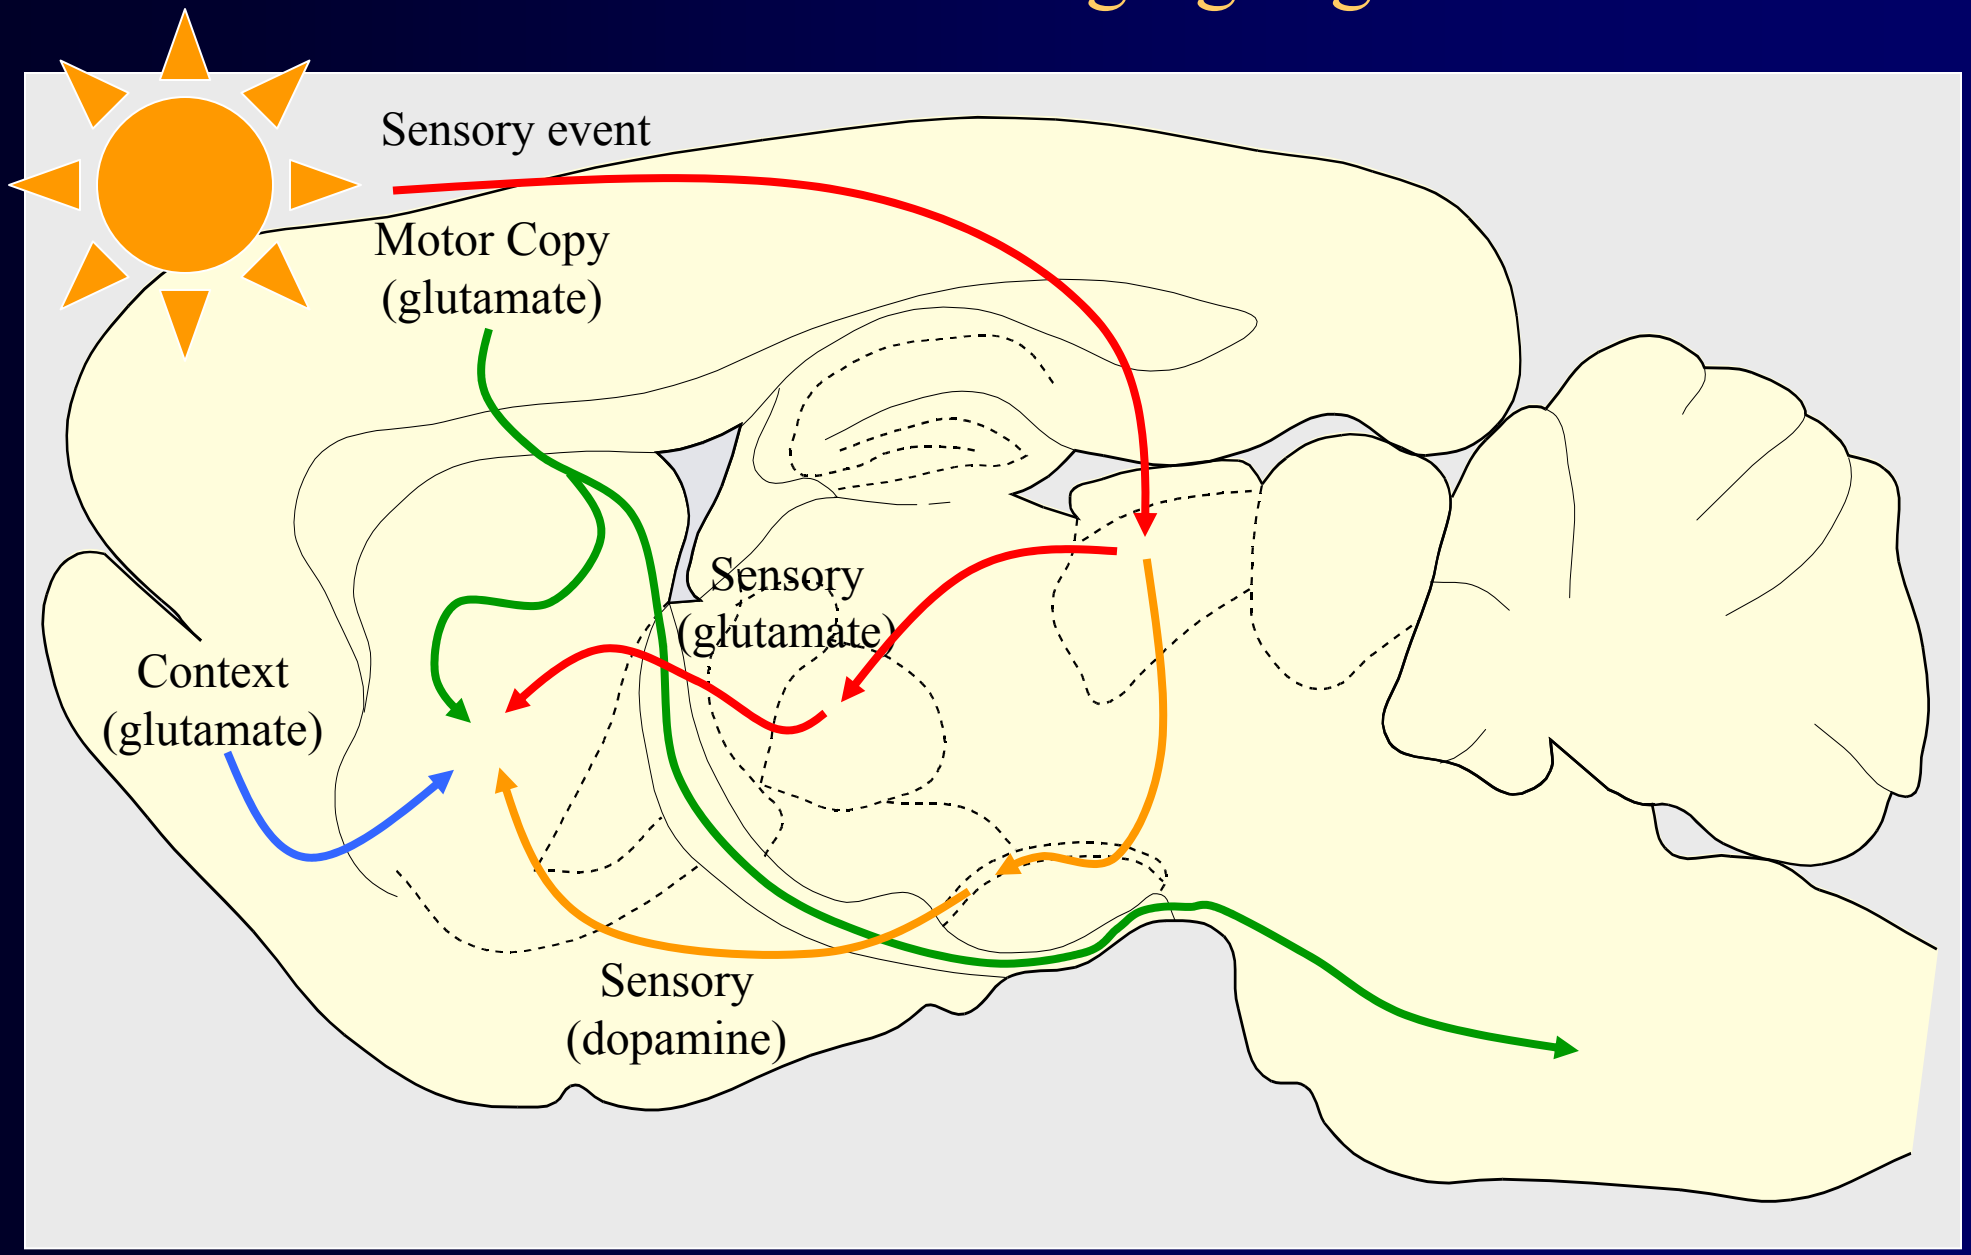

Timed convergence of signals → Agency determination

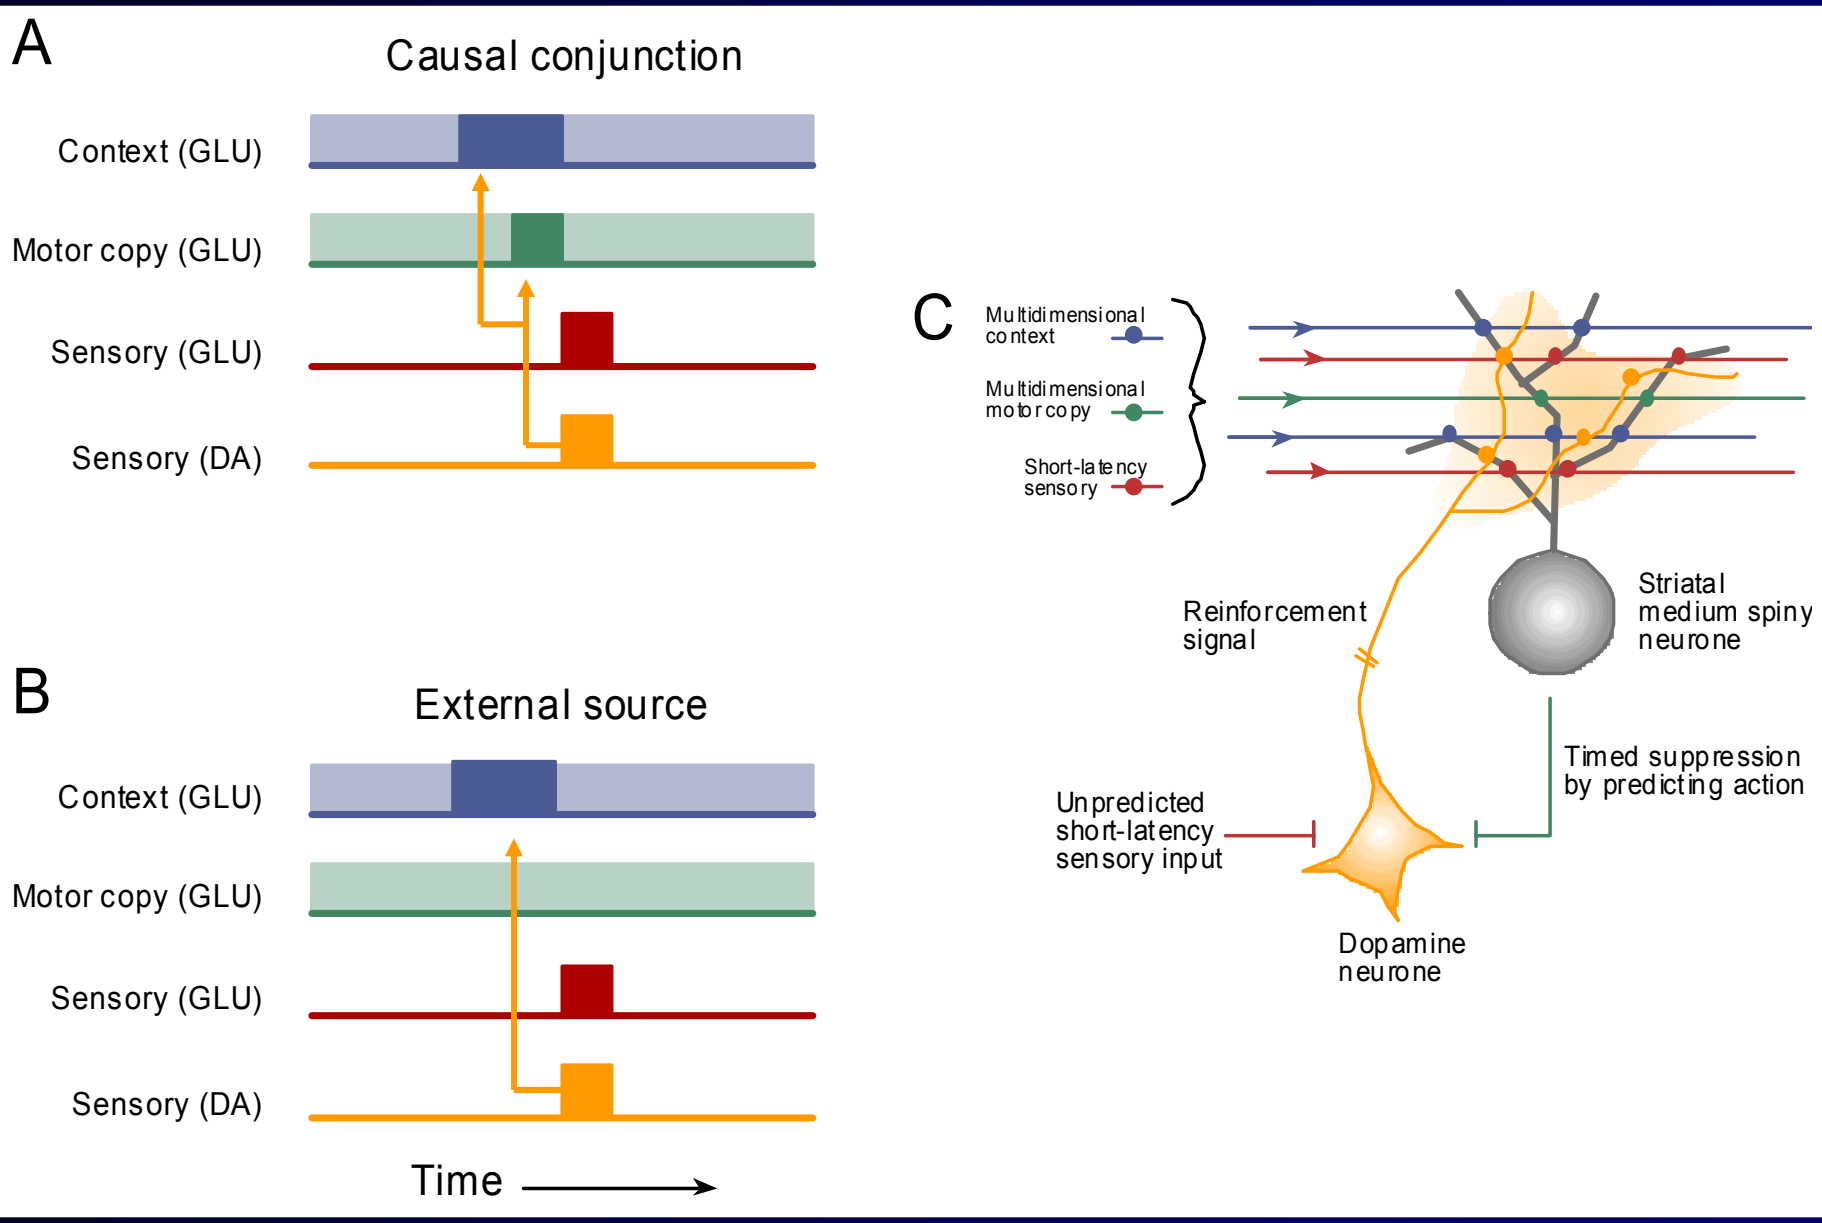

Reason for pre-gaze-shift DA signaling becomes apparent

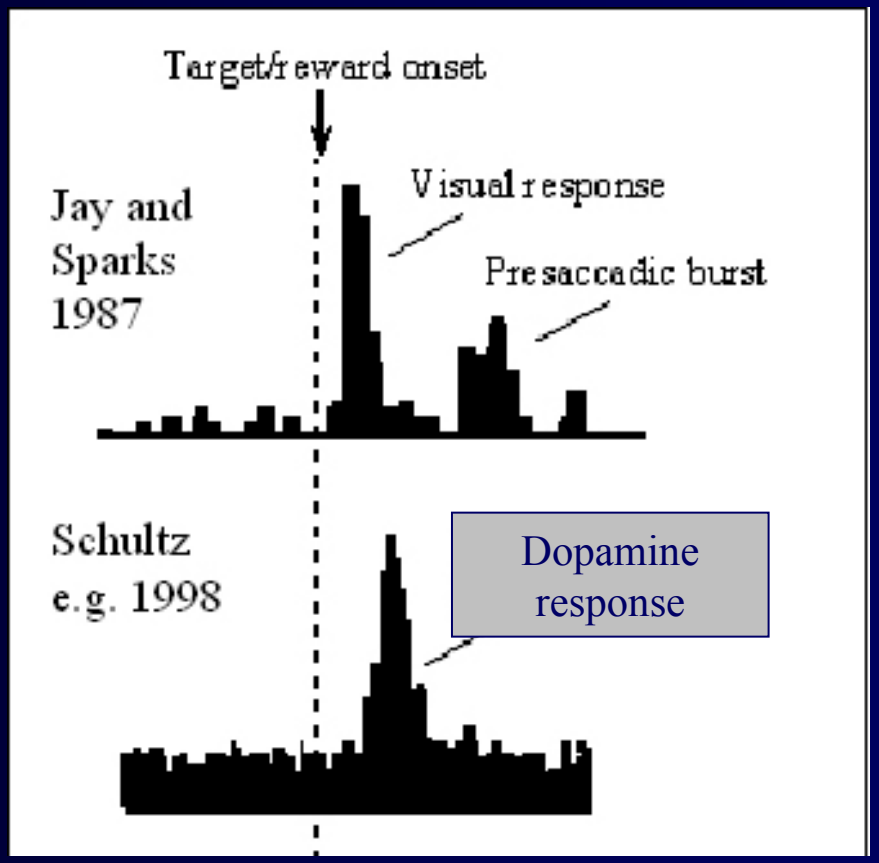

Redgrave P, Prescott TJ and Gurney K (1999). TINS 22(4): 146-151

Gaze-shift contamination → Credit-assignment problem

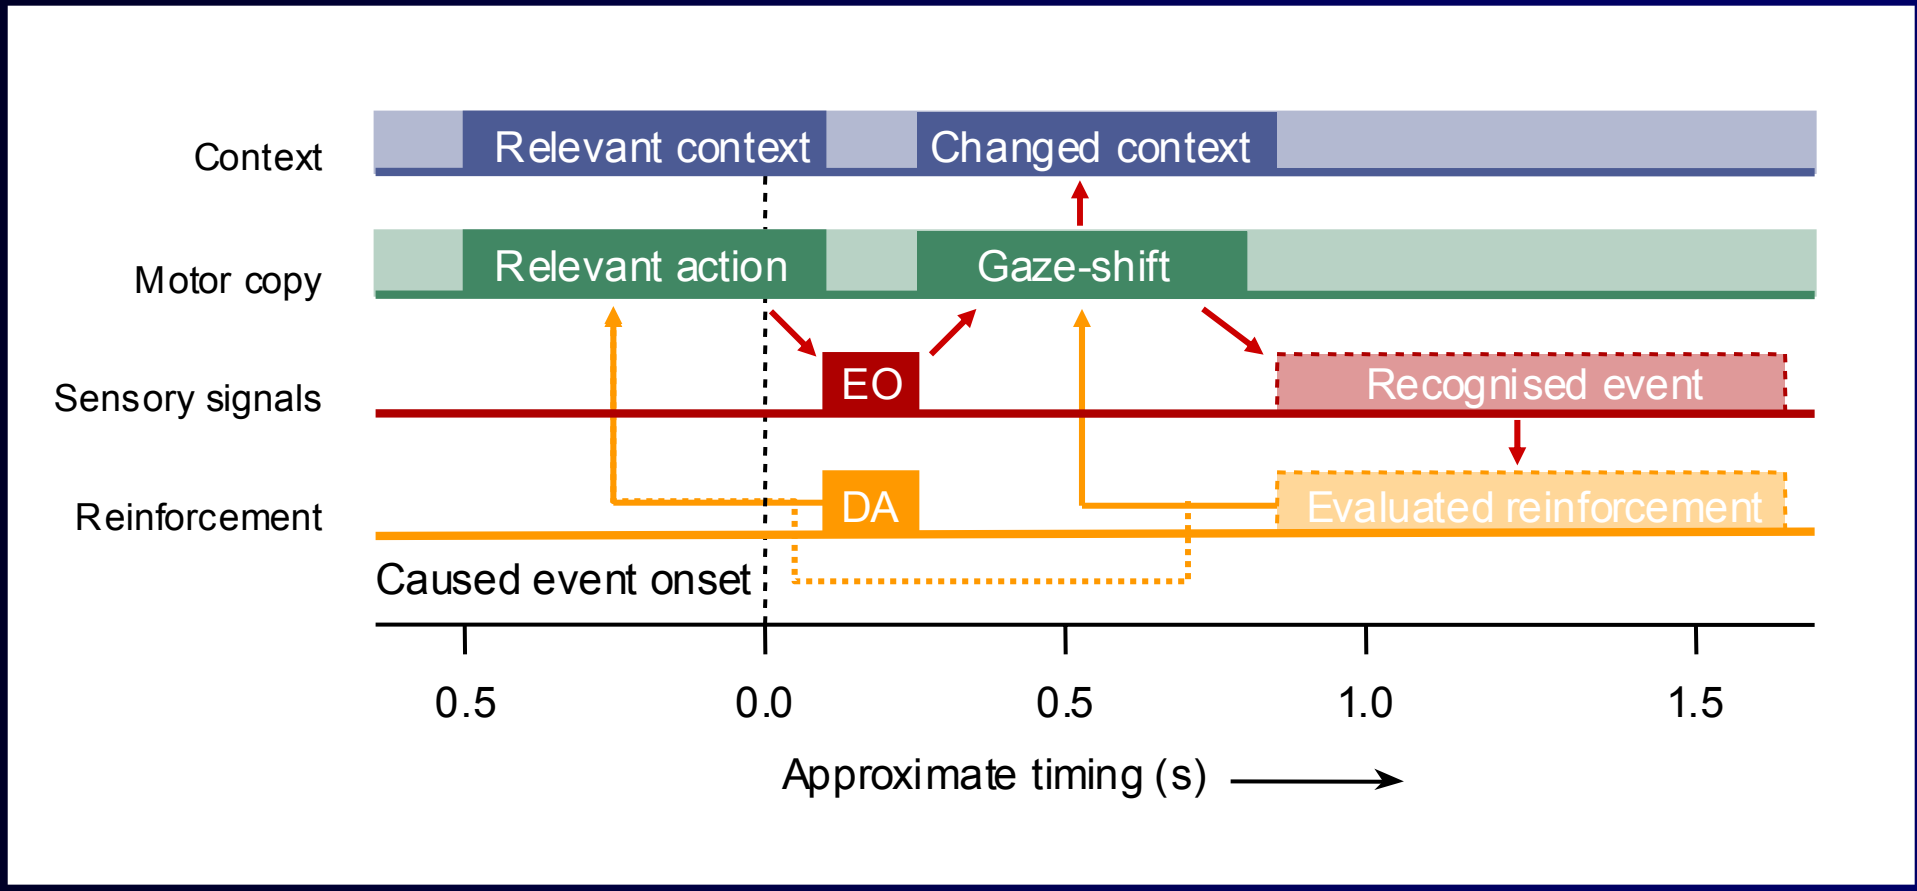

Developing new actions…not that easy

Before

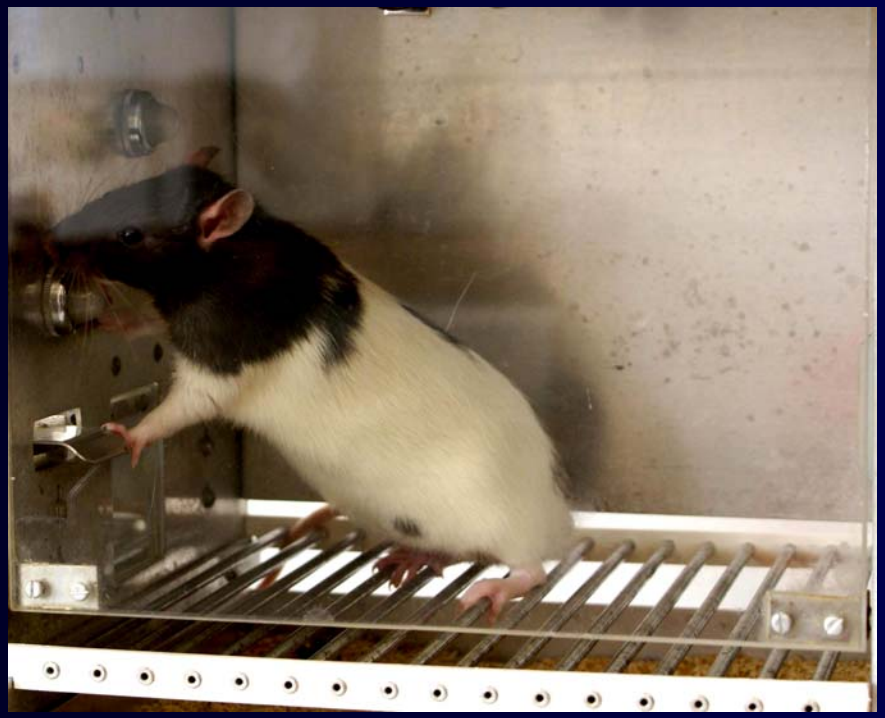

After

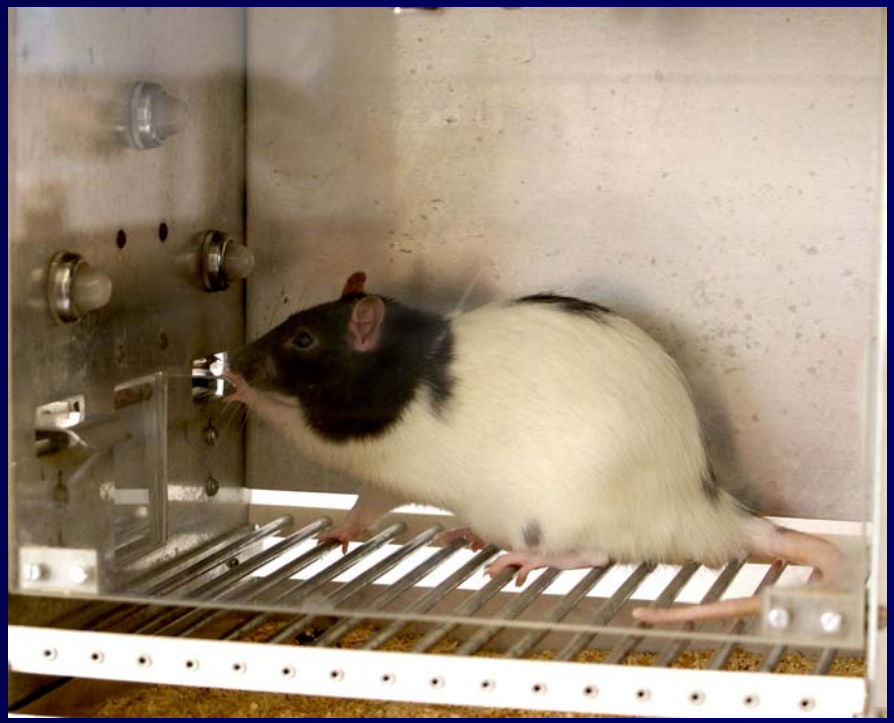

Context and motor copy – multidimensional

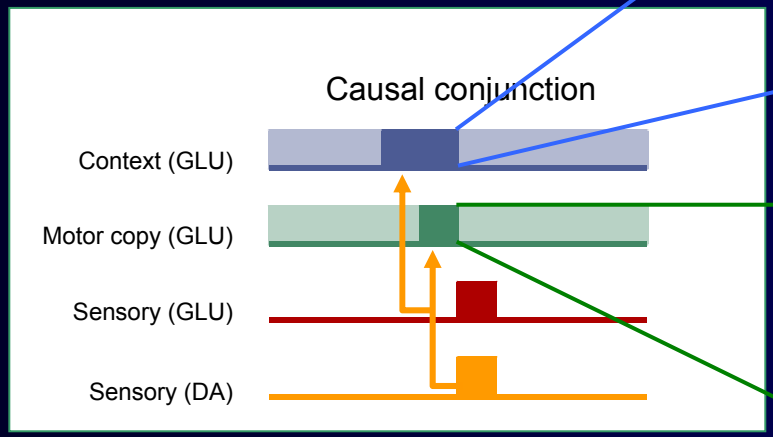

Physical location
Distal stimuli
Proximal stimuli
Timing
etc

Context

Goal
Actions
Movements
Timing
etc

Motor copy

How are critical aspects of context and movement responsible for caused events discovered ?

Development of new actions

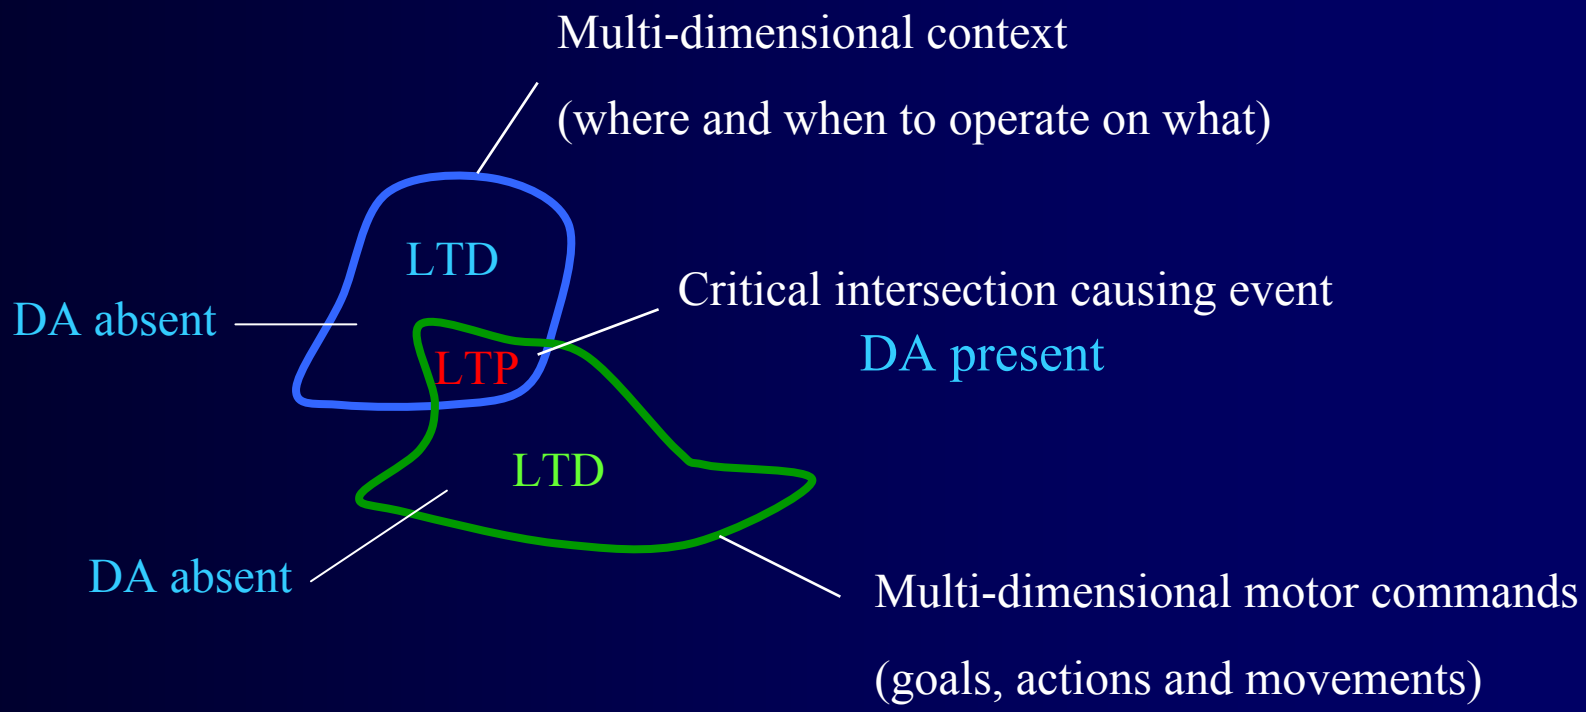

- Repetitive sampling of preceding movements in preceding contexts – with variation
- + LTP/LTD determined by presence/absence of phasic DA

→ system to converge on critical causative components

Final conclusions

- Basal ganglia connectivity provides an architecture permitting agency to be determined
- Variable repetition + DA-related plasticity enables discovery of critical components of context and movement → novel action
- Through play and exploration agent builds library of action-outcome routines = options (Barto)
- Routines later selected and assembled into sequences on the fly = novel adaptive and purposive behaviour

Collaborators

- Anatomy
  - Eliane Comoli (Sheffield, now Brazil…SP-RP)
  - Véronique Coizet (Sheffield)
  - Paul Overton (Sheffield)
- Electrophysiology
  - Véronique Coizet (Sheffield)
  - Ellie Dommett (Sheffield)
  - Paul Overton (Sheffield)
- Electrochemistry
  - Charles Blaha (Sydney – now Memphis)
  - John McHaffie (Wake Forest)
  - Barry Stein (Wake Forest)
  - Huai Jiang (Wake Forest)
  - Paul May (U. Mississippi)
  - Suzanne Haber (Rochester)
  - John Reynolds (Otago, NZ)

wellcome trust

EPSRC


100dB loss
at 4 kHz

….but what happens when it's nasty ?

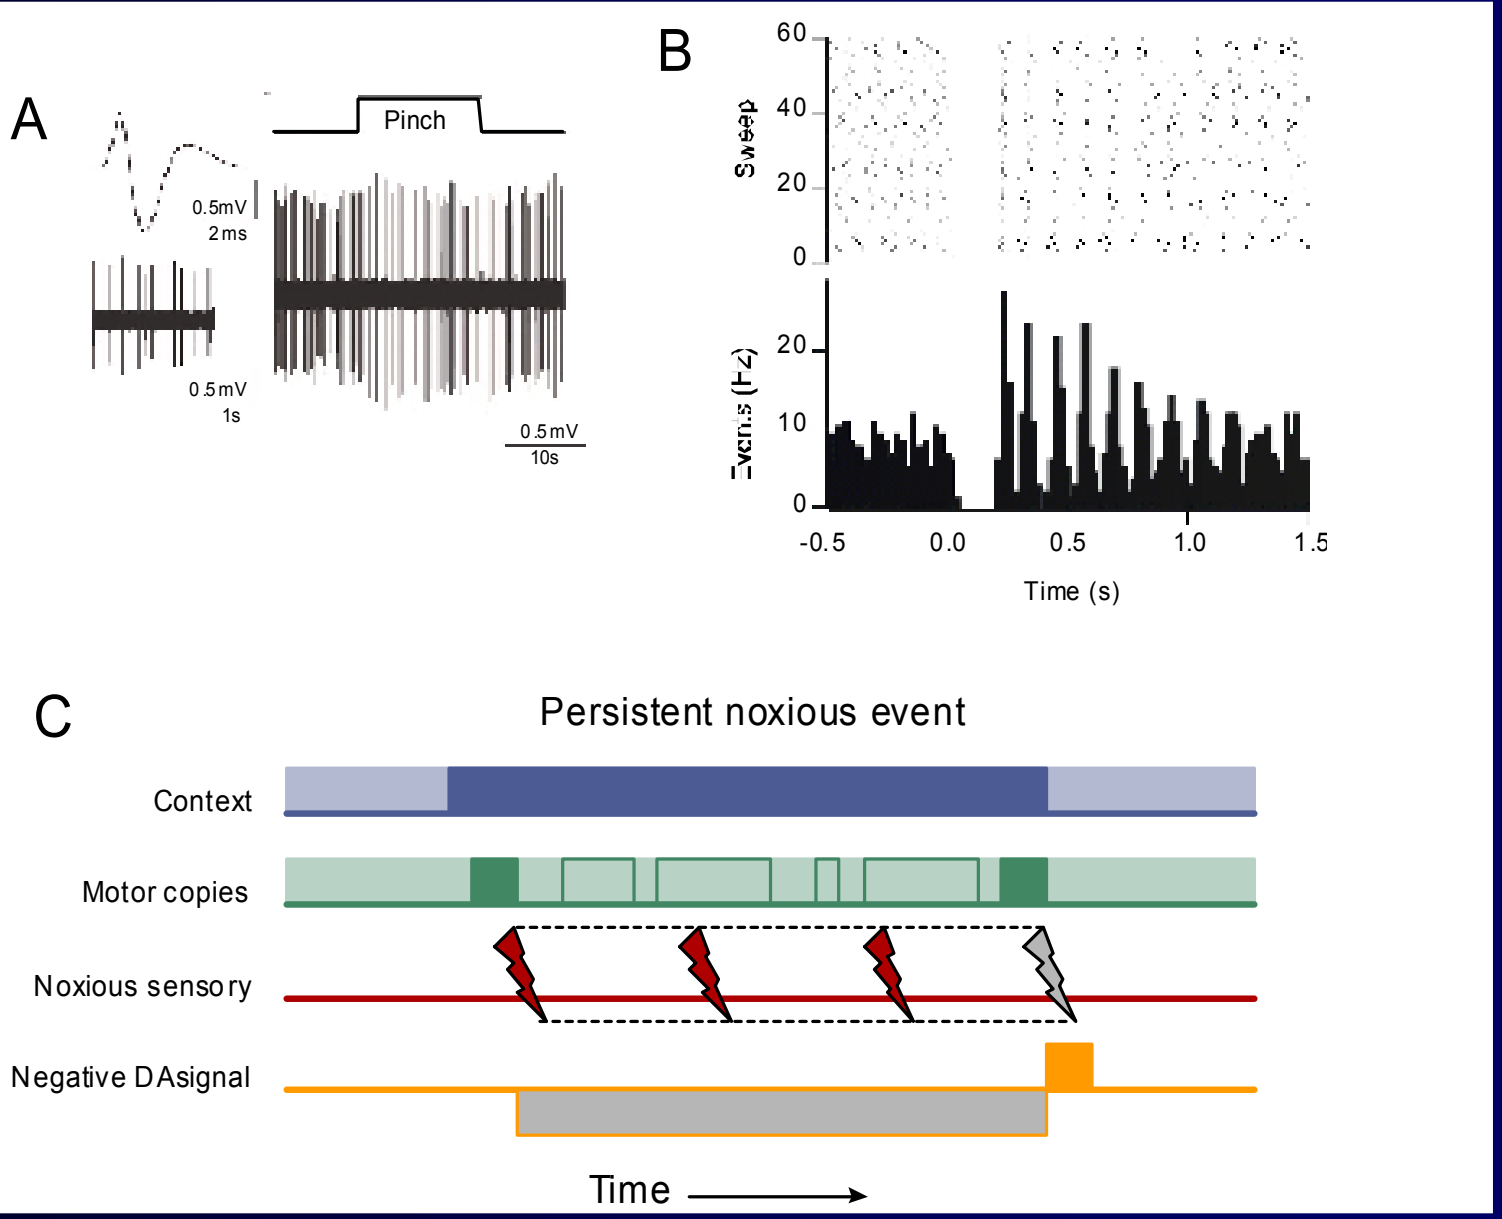

Reward prediction error → collicular deep layers

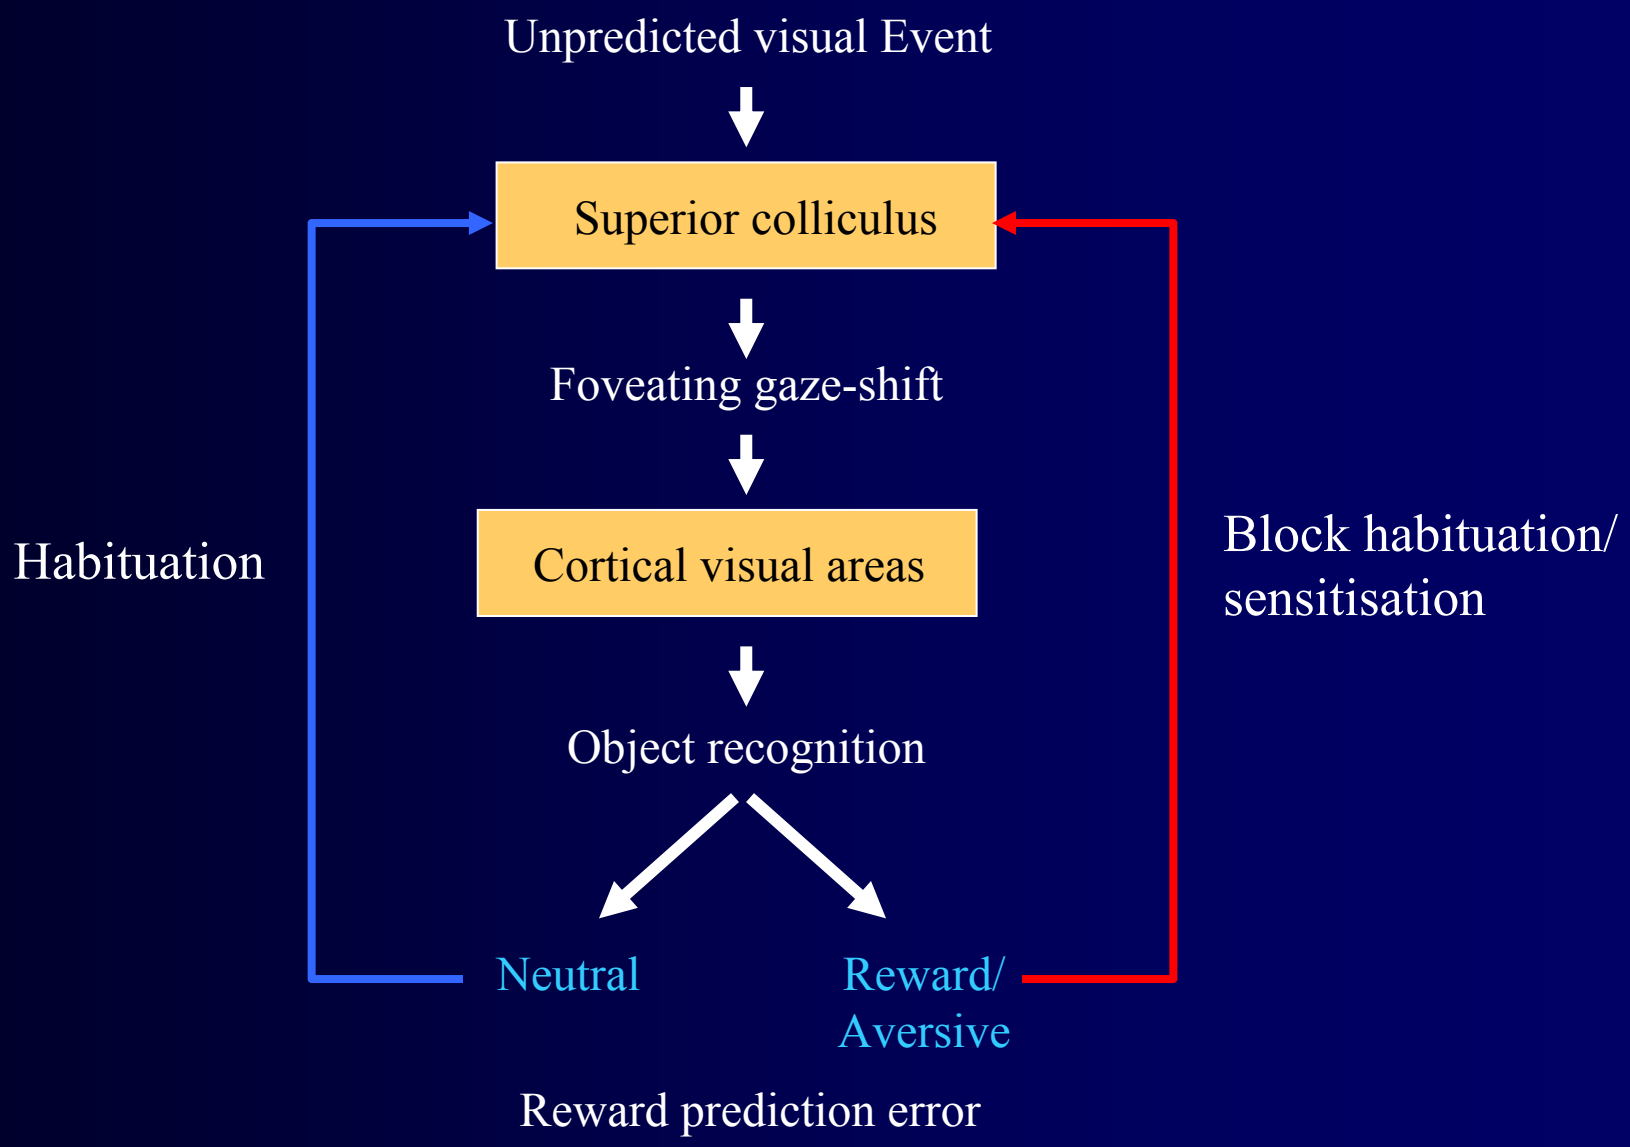

The superior colliculus responds to….

- All salient novel visual events
  - If no reinforcement consequences will habituate
- All visual events associated with reward
  - Habituation blocked/sensitised

Habituation

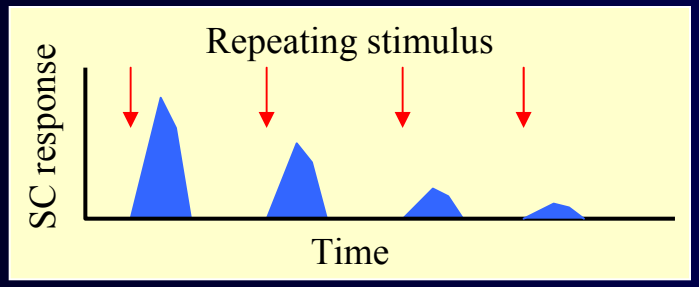

Neutral stimuli

Oyster CW, 1975 *J Neurophysiol* 38(2):301-312.

Sensitisation

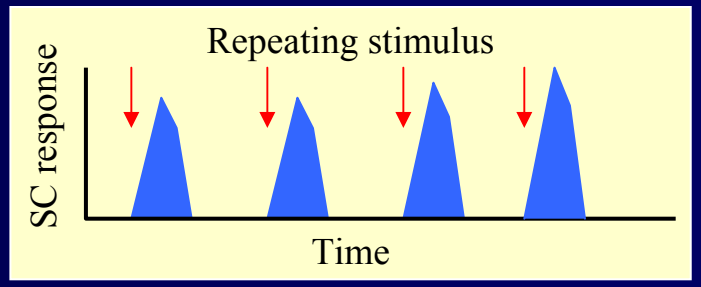

Reward, reward-predicting stimuli

Ikeda T, 2003 *Neuron* 39:693-700.

… DA neural response can partition prediction errors

Fiorillo CD, Tobler PN, Schultz W. 2003. Science 299:1898-1902.

Signal indicating probability of reward

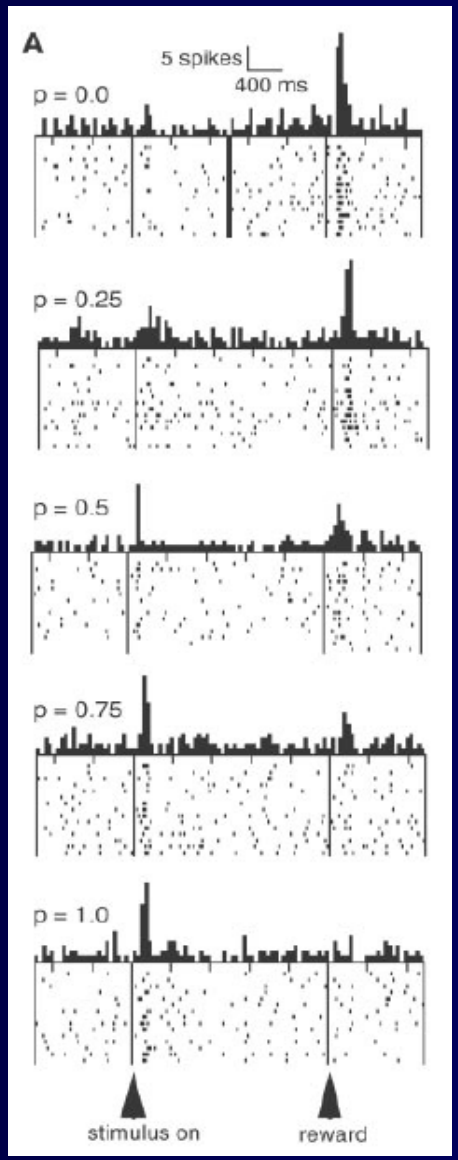

Signal indicating reward

Reward prediction error absent when based on visual search

Failure to discriminate rewarded and unrewarded door

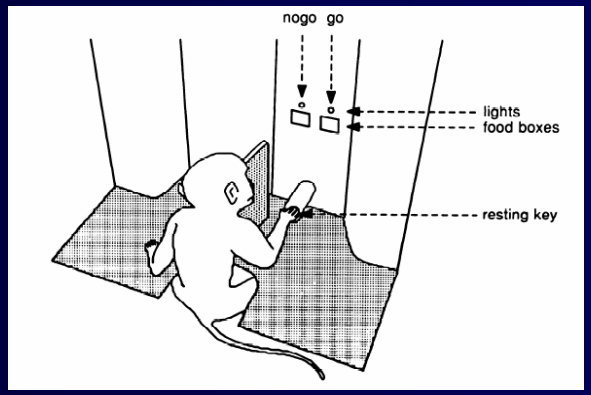

Signals only the 1st 50% reward prediction error

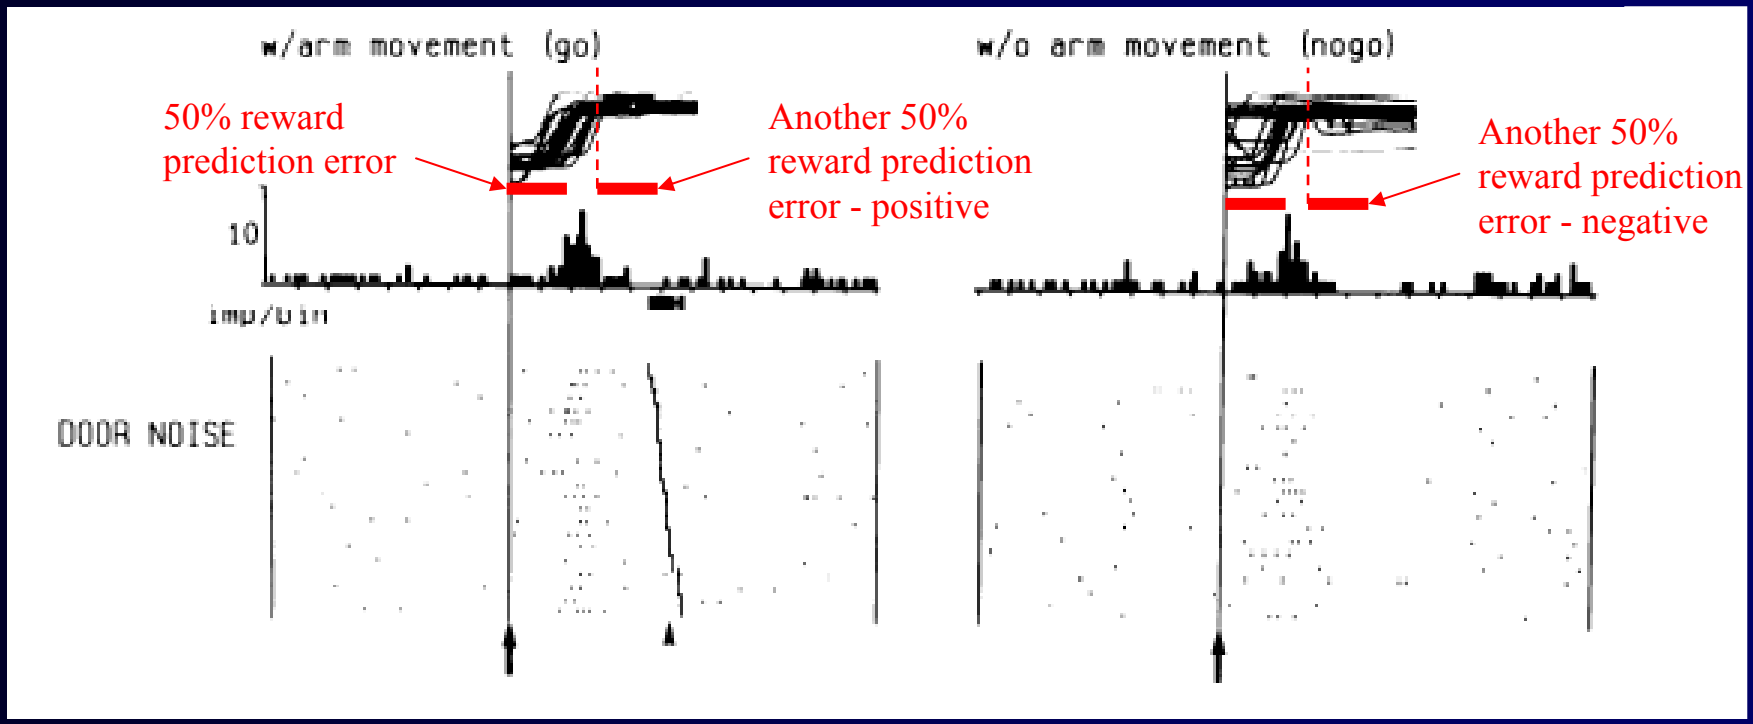

Schultz, W. & Romo, R.. *J. Neurophysiol.* **63**, 607-624 (1990).

….but what about Genela's experiment ?

Morris G, Arkadir D, Nevet A, Vaadia E, Bergman H. 2004. Neuron 43(1):133-143.

..can't DA signal reward probability when stimuli are spatially unpredictable ?

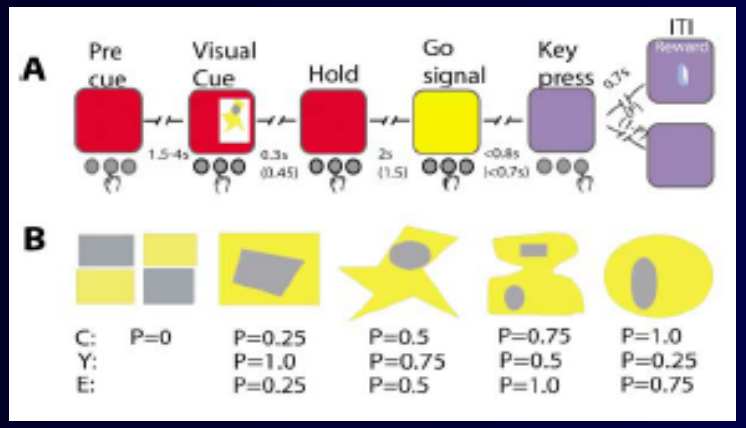

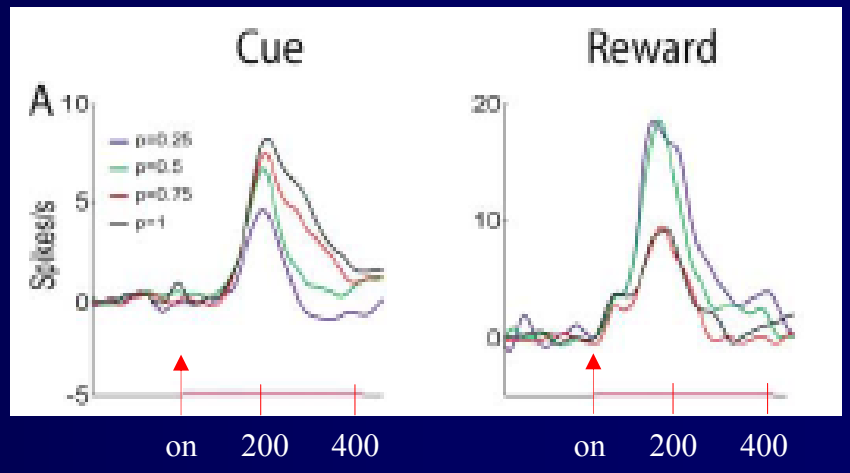

1. Large low spatial frequency stimuli – not matched for luminance
2. Only two possible locations where the stimuli might appear
3. 60-92 training days: 18,000-46,000 training trials with training stimuli….further 9000-20,500 training trials experimental stimuli
4. No fixation point and no measurement of eye movements
5. Results significant only when data for analysis extended from 200-400ms