BioE 2696/ECE 2695: Control Theory in Neuroscience

Reinforcement Learning: Physiology

April 6, 2009

Instructor: Rob Turner
Assoc. Professor of Neurobiology
University of Pittsburgh

Overview

- What is reinforcement learning ("RL")?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
  - Anatomy, Physiology, Pharmacology
- Dopamine as a RL signal
- Flaws in the dopamine-RL story
- Basal ganglia as actor/critic
- Basal ganglia as a tutor

What is reinforcement learning?

- Operant conditioning
  modification of "voluntary" behavior
  Thorndike's Law of Effect
  - behavior is modified by consequences
  - satisfying consequences → increase a behavior
  - annoying consequences → decrease a behavior

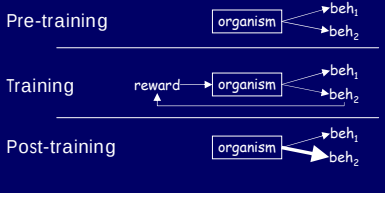

What is reinforcement learning?

- Operant conditioning = "the biasing of future actions by past outcomes"
  - requires temporal contiguity & contingency
  - conditioning is context- or stimulus-specific

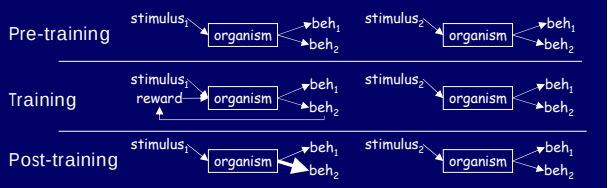

What is reinforcement learning?

- Operant conditioning = "the biasing of future actions by past outcomes"
- Requires prediction of outcomes across time.

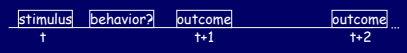

For correct prediction of outcome:

$$\text{prediction}(t) = \text{expected outcome}(t+1) + \text{prediction}(t+1)$$

Incorrect prediction yields a "temporal difference" (TD) prediction error:

$$\delta(t+1) = \text{outcome}(t+1) - \text{prediction}(t) + \text{prediction}(t+1)$$

TD can be used to improve predictions (η = rate of learning):

$$\text{prediction}(t)_{new} = \text{prediction}(t)_{old} + \eta\cdot\delta(t+1)$$

- What is RL?

Basal ganglia & dopamine roles in RL

- Impaired RL w/ BG pathology
  - Parkinson's disease (Knowlton et al, Science '96)
- Impaired win-stay learning w/ striatal lesions
  - double dissociation w/ hippocampal-based win-shift learning (Packard et al, Beh Neurosci '92)
- Schultz recording from midbrain dopamine (DA) neurons
  - unpredicted reward → response
  - reward-predicting stim. → response
  - predicted reward → no response
  - omit predicted rew → suppression

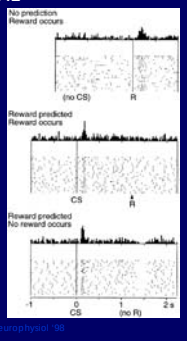

Schultz, J Neurophysiol '98

BG Anatomy I.

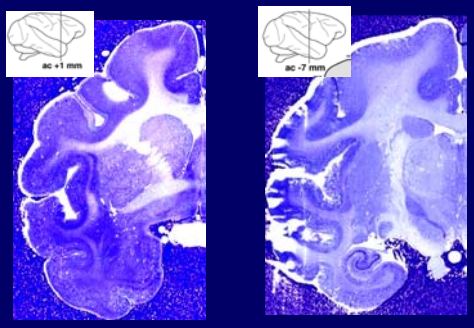

BG Anatomy II.

Direct Pathway
- GABA → internal pallidum (GPi)
- D_1 dopamine receptor

Indirect Pathways
- GABA→ external pallidum (GPe) and then multiple routes to GP
- D_2 receptor

Dopaminergic Pathway
- substantia nigra compacta (SNc)
- → $D_{1/3}$ direct pathway
- → D_2 indirect pathway

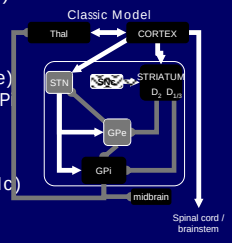

BG Anatomy III. Segregated Functional Circuits

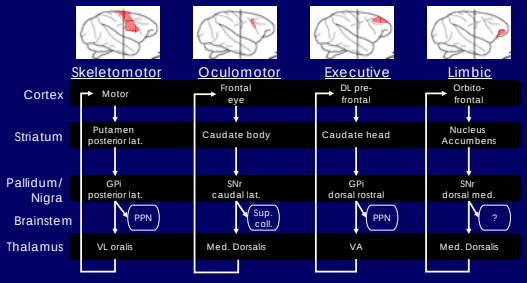

Alexander, DeLong & Strick '86 Ann Rev Neurosci.

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
    - pigmented (neuromelanin)
    - degenerate in Parkinson's disease (PD)

Human midbrain (unstained)

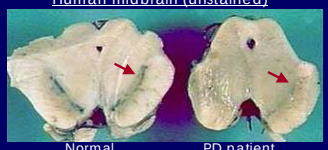

Normal PD patient

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
    - nigrostriatal

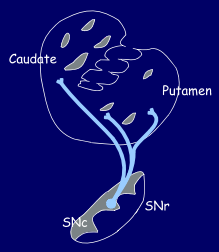

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
    - nigrostriatal
    - many extra-striatal collaterals

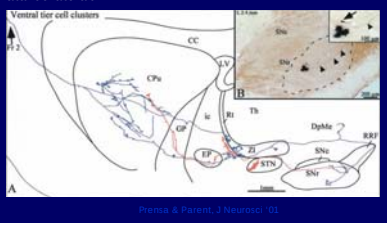

Prensa & Parent, J Neurosci '01

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
    - nigrostriatal
    - many extra-striatal collaterals
      - » organized to "broadcast"

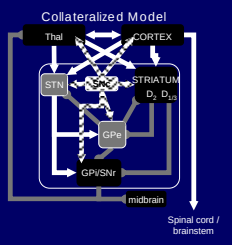

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
  - excitatory efferents
    - midbrain (pedunculopontine)
    - superior colliculus (SC)
    - subthalamic nucleus (STN)

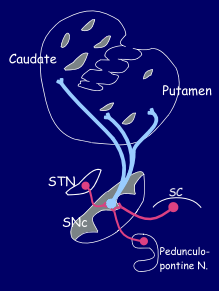

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
  - inhibitory efferents
    - SNr
    - striatal "patches"

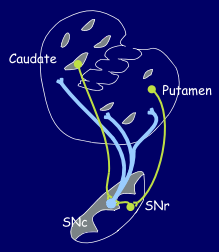

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
  - cortico-striatal efferents
    - "limbic" allocortex projects to striatal patches

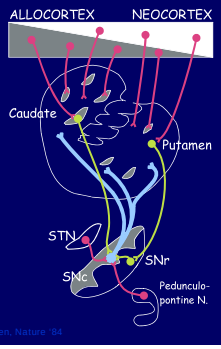

Gerfen, Nature '84

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
  - cells
  - projection patterns
  - efferents
  - circuit organization
    - part closed-loop
    - part limbic→associative→ motor "spiral"

Spiral organization

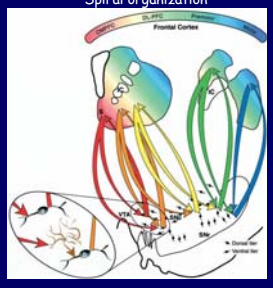

McFarland & Haber, J Neurosci '04

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
- Physiology
  - long multiphasic action potentials
  - low resting firing rate (1-8 Hz)
  - also respond to:
    - novel sensory stimuli
    - significant events (e.g., trial onset)
  - resting rate may modulate w/risk & motivation
  - no response to movement

SNc neuronal recording

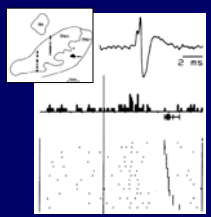

Schultz, J Neurophysiol '86

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
- Physiology
- Synaptic organization
  - DA terminals in striatum synapse on dendritic spines

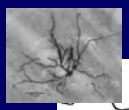

Striatal "medium spiny" neuron

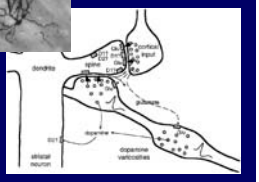

- What is RL?
- Basal ganglia & dopamine roles in RL

The midbrain dopamine system

- Anatomy
- Physiology
- Synaptic organization
  - 2 receptor families ($D_{1/5}$ & $D_{2\text{-}4}$)
  - Proj. neurons in striatum express D_1 (direct PW) or D_2 (indirect)

Striatal medium spiny neurons

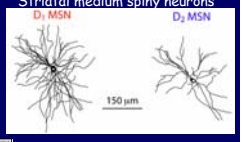

DA promotes LTP in D_1 neurons
DA promote LTD in D_2 neurons

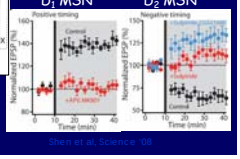

3-factor learning
1) pre-synaptic spikes
2) post-synaptic depolarization
3) + dopamine
≠Hebbian learning

Shen et al, Science '08

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system

Dopamine as a RL signal

- result confirmed multiple times
- responses stereotyped & independent of reward or stimulus type or features
- responses scale w/ reward prob. & size (but not for omissions)
- responses are similar in most DA neurons (surprising)
- no response to aversive stimuli
- "blocking effect" – new cues do not associate w/ rewards that are fully predicted

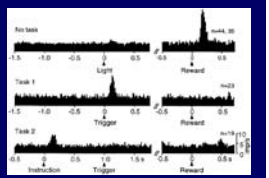

Responses scale w/rew size

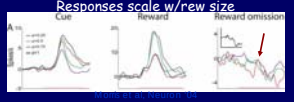

Morris et al, Nature '04

Summary so far...

- Reinforcement learning can be mediated by a "temporal difference" (TD) learning rule.
- The BG are implicated in RL.
- Midbrain DA neurons have activity that closely approximates the TD prediction.

But...

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal

Flaws in the dopamine-RL story

DA response specificity & latency (Redgrave et al.)

- responses to unexpected sensory events w/ reward value (can unexpected events mediate reinforcement learning?)
- responses don't discriminate between types of rewards (cannot mediate reward-specific predictions)
- latencies are too short for complex perceptual analysis

ALTERNATIVE: DA neurons encode a general time-stamp for occurrences of significant unpredicted events.

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal

Flaws in the dopamine-RL story

Effects of DA lesions (Berridge, Salamone & Palmiter)

- DA-deficient mice can learn an appetitive T-maze

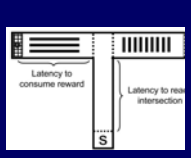

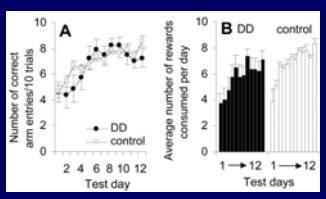

Robinson et al & Palmiter, Beh Neurosci '05

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal

Flaws in the dopamine-RL story

Effects of DA lesions (Berridge, Salamone & Palmiter)

- DA-deficient mice can learn an appetitive T-maze
- DA-deficient mice show signs that they "like" rewards (i.e., DA ≠ a hedonic signal)
- DA is necessary for motivating movement to obtain rewards (DA "wanting")

Robinson et al & Palmiter, Beh Neurosci '05
Salamone et al, Psychopharm '07 (review)
Berridge, Psychopharm '07 (review)

ALTERNATIVE: DA signals "incentive salience" used to control effort-related functions (i.e., response vigor)

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal
- Flaws in the dopamine-RL story
- Alternative models of DA function

The BG as Actor-Critic

- Parallel segregated loops (actor & critic)
- Limbic circuit signals value
- Motor circuit signals kinematics & context
- Selection is impaired in BG pathologies (PD)

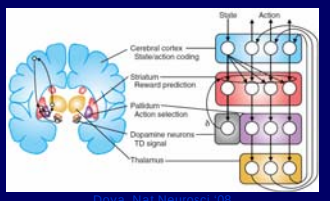

Doya, Nat Neurosci '08

But, disconnecting the BG doesn't impair selection!

GPi Inactivations: Summary Results

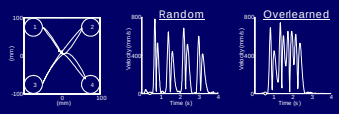

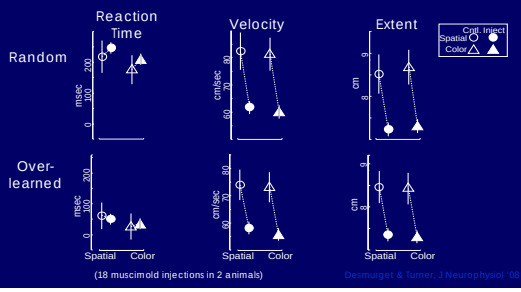

(18 muscimold injections in 2 animals)

Desmurget & Turner, J Neurophysiol '08

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal
- Flaws in the dopamine-RL story
- Alternative models of DA function

BG as a Tutor

- Learning-related neuronal activity
- BG disconnection impairs learning, but not performance (pallidotomy)
- BG plasticity is fast (DA-driven 3-factor rule)
- Cortical plasticity is slow (Hebbian 2-factor rule)

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal
- Flaws in the dopamine-RL story
- Alternative models of DA function

BG as a Tutor

BIRDSONG

- Song beh. similar to sequence skill in mammals
- Bird homologues exist for mammalian BG
- In adult birds, BG disconnection
  - leaves learned songs intact
  - interferes with experience-dependent song plasticity

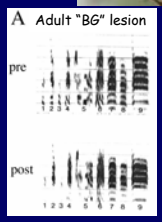

Brainard & Doupe, NY Acad Sci '04

- What is RL?
- Basal ganglia & dopamine roles in RL
- The midbrain dopamine system
- Dopamine as a RL signal
- Flaws in the dopamine-RL story
- Alternative models of DA function

BG as a Tutor

BIRDSONG

- In young birds, BG disconnection blocks song learning completely

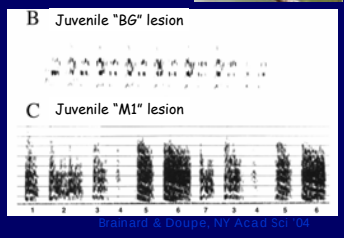

Brainard & Doupe, NY Acad Sci '04

BG Functions in Motor Control

MY WORKING HYPOTHESES

The BG may:

- modulate "motor motivation"
- mediate fast reward-dependent learning
- facilitate slow learning in thalamo-cortical circuits