Actor - Critic

Hyperparams:

Initial policy params θ (ANN? Linear?)
Policy representation
Actor step size α
Critic step size β.
Value function representation, v_w.
Initial value function weights w.

For each episode:
  For each time t
    Agent observes S_t
    Agent selects action A_t using π_θ
    Env responds with S_{t+1} and R_t
    $\delta_t = R_t + \gamma v_w(S_{t+1}) - v_w(S_t)$ // TD-error
    $\forall i,\ \theta_i \leftarrow \theta_i + \alpha \gamma^t \delta_t \frac{\partial \ln(\pi_\theta(S_t, A_t))}{\partial \theta_i}$ // Actor update
    $\forall j,\ w_j \leftarrow w_j + \beta \delta_t \frac{\partial v_w(S_t)}{\partial w_j}$ // critic update.

Theory says to include
In practice it is bad. Almost nobody includes this term.

$$v^\pi(s) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k R_{t+k} \,\middle|\, S_t = s;\ \pi\right]$$

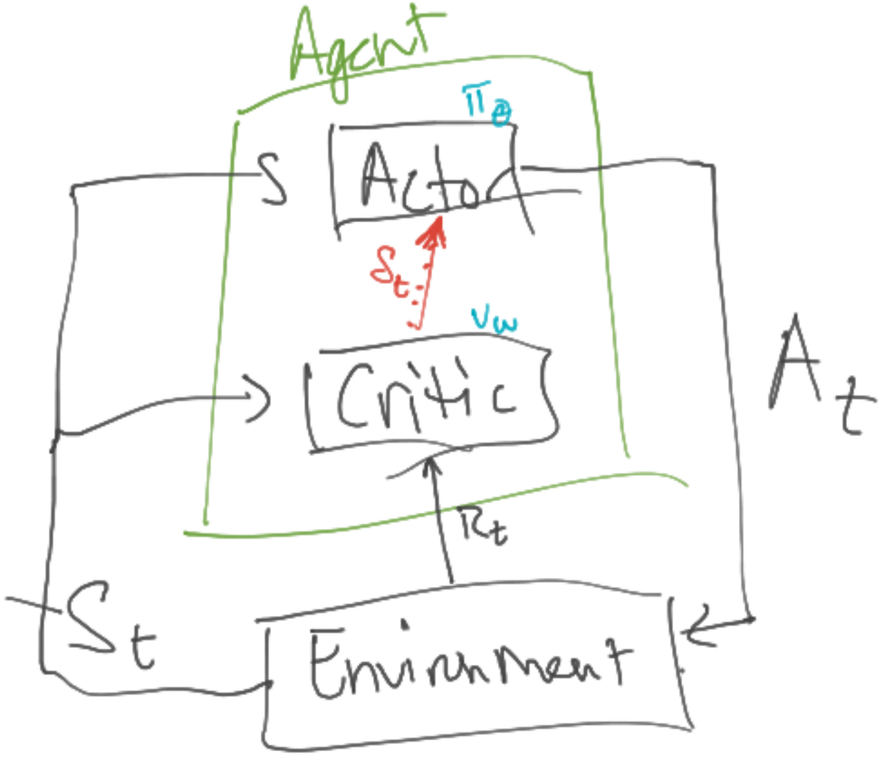

Psychology

Sutton & Barto
2nd edition.

RLDM

Operant conditioning: learning process through which the strength of a behavior is modified by reward or punishment.

(control)

(searching for a better policy)

- learning a policy.

Classical conditioning: learning procedure in which a biologically potent stimulus (e.g., food) is paired with a previously neutral stimulus (e.g., bell).

(prediction)

- learning a value function.

- Thorndike's Puzzle Boxes.

1898

Neuroscience

Dopamine

- Contraction of 3,4 - dihydroxyphenethylamine
- Chemical
- Neurotransmitters.

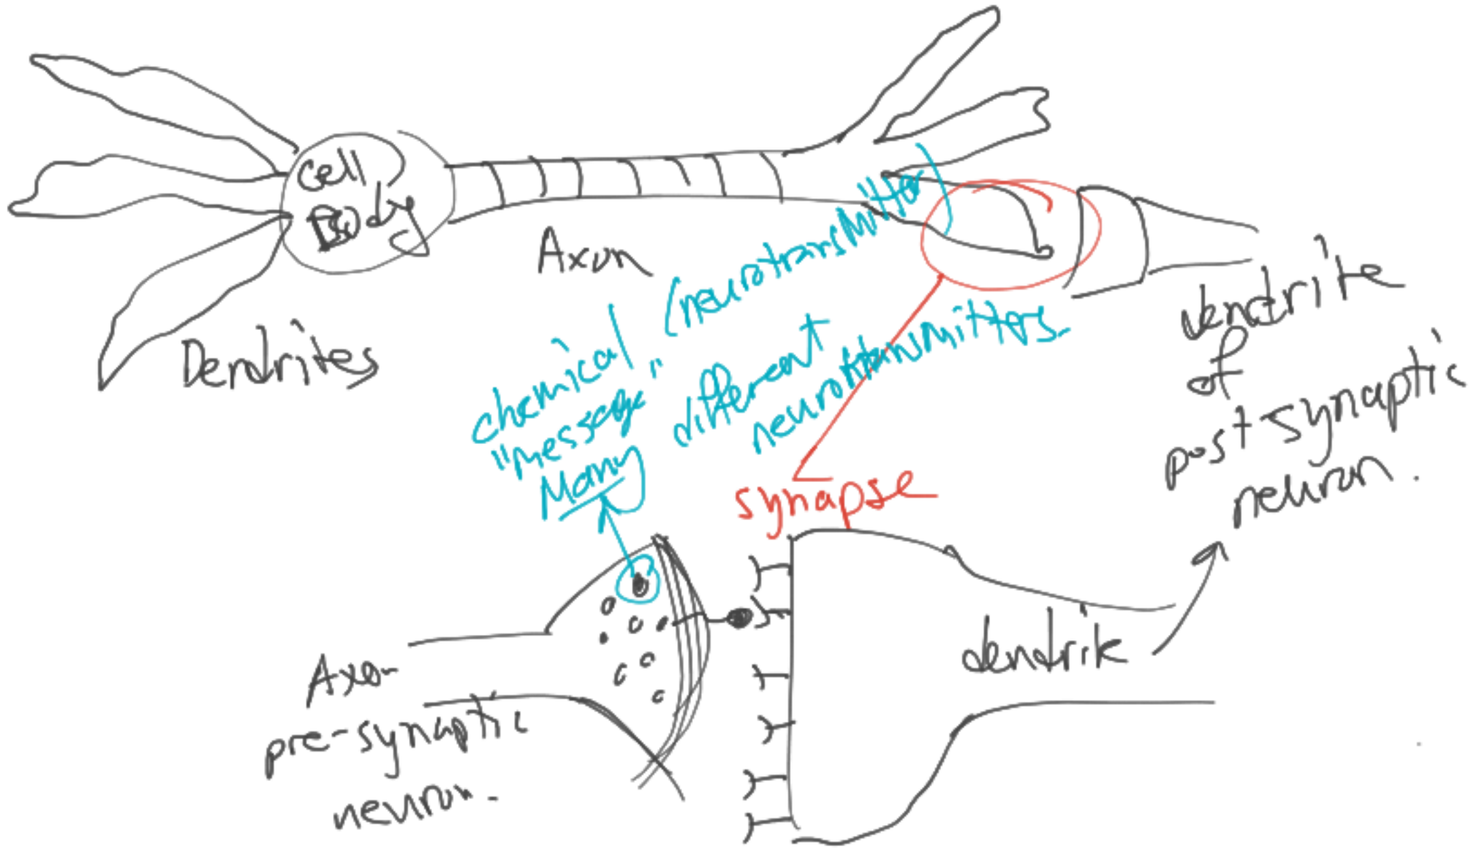

Dopamine neuron
(Dopaminergic neuron)
is a neuron that emits the neurotransmitter dopamine.

Two clusters in mammals:

SNpc and VTA

↓ ↓

Substantia nigra pars compacta — Ventral tegmental area.

SNpc —dopamine→ Striatum

↳ coordinates motor & action planning, decision making, motivation.

VTA —dopamine→ numerous areas including prefrontal cortex → planning, personality, decision making

Reward Prediction Error Hypothesis for dopamine.

TD error.

Olds & Milner 1954 : Dopamine & reward.

Brains probabably do not implement backpropagation.

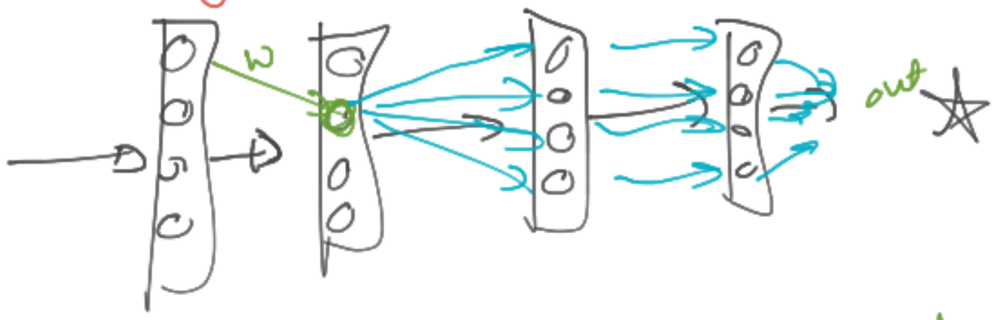

$\frac{\partial out}{\partial w}$ Δ_j terms propagated backwards through the network.

→ Information does not seem to pass backwards down the axon.

Each "neuron" can update using only its input, output, and δ_t.

Duplicate of network in reverse.

Coagent Networks

Extention of "learning automata"

→ training RL networks without backprop.

Reward devaluation Studies.

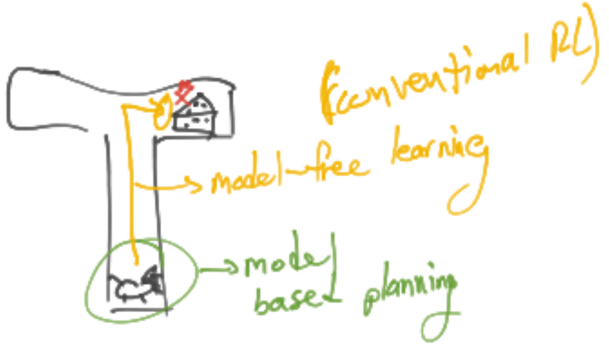

→ early learning is model-based planning

→ transitions to a model-free policy over time.

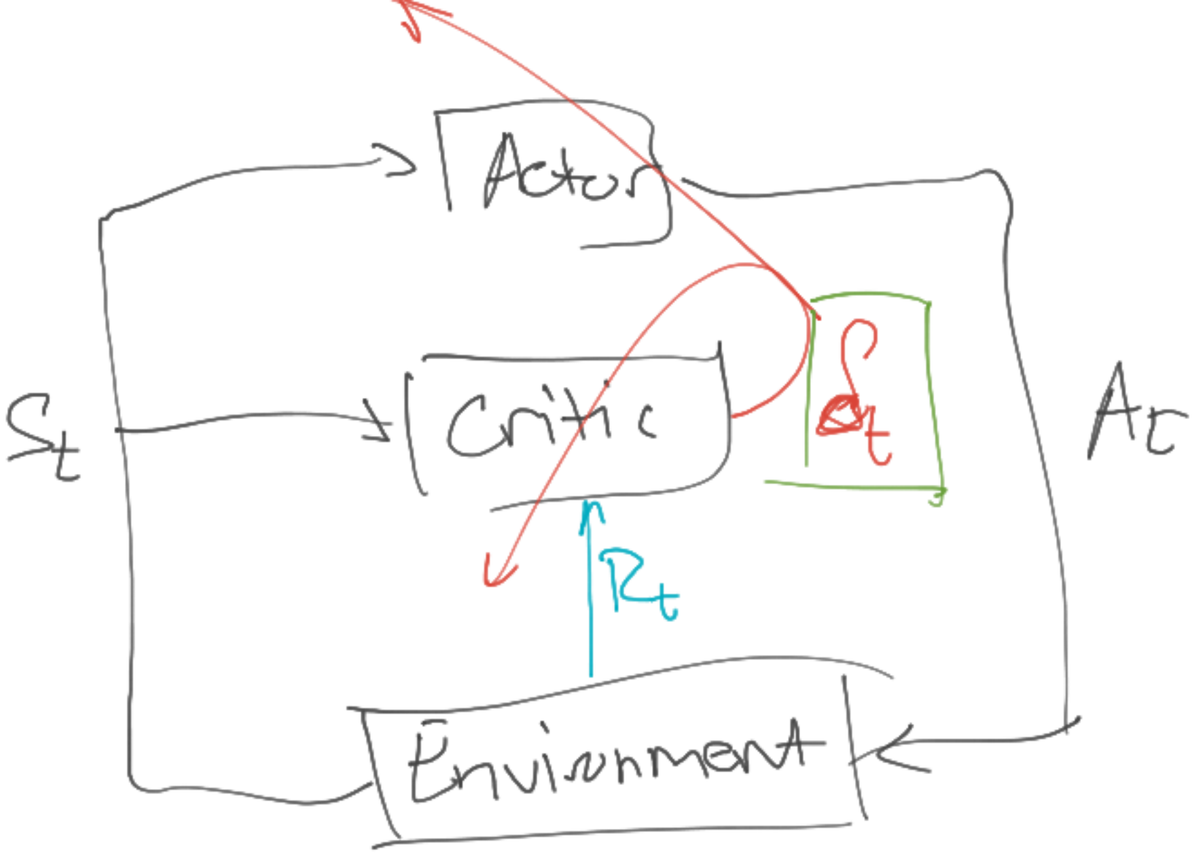

$\pi_{eat\ sugar}$

V(cake in front)

$V^{\pi_{heroin}}$(drugs) =

Addiction

octopi