1) Not responsible for generalization.

lecture 11 → RL!

Let $X_1, \ldots, X_n$ be n i.i.d. random variables, let $\mu = \mathbb{E}[X_1]$, and $\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i$, and $X_i \in [a, b]$

then

$$\Pr(|\bar{X}_n - \mu| \geq t) \leq 2e^{-\frac{2nt^2}{(b-a)^2}}$$

t: any positive constant

$\bar{X}_n$

$$\ell(w_k) = \frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$

$(y_i - \hat{y}_i)^2 = X_i$

μ

$$\ell(w_k) \; \frac{1}{N}\sum_{i=1}^{N} (y_i - \hat{y}_i)^2$$

N → entire population.

2) HW dropped, all get 100%.
   → Still graded if submit.
   → Still 100% in actual gradebook.

Perhaps public grading test.

Extra credit for decent attempt?

What is reinforcement learning (RL)?

"Reinforcement learning is an area of machine learning, inspired by behaviorist psychology, concerned with how an agent can learn from interactions with an environment."

– Wikipedia / Sutton & Barto / Phil

1998

Agent - Environment Diagram

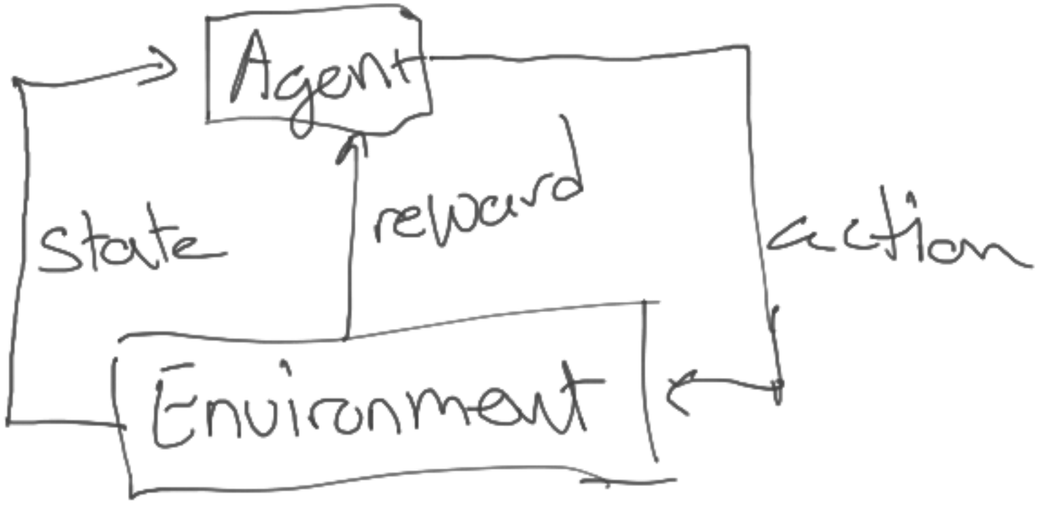

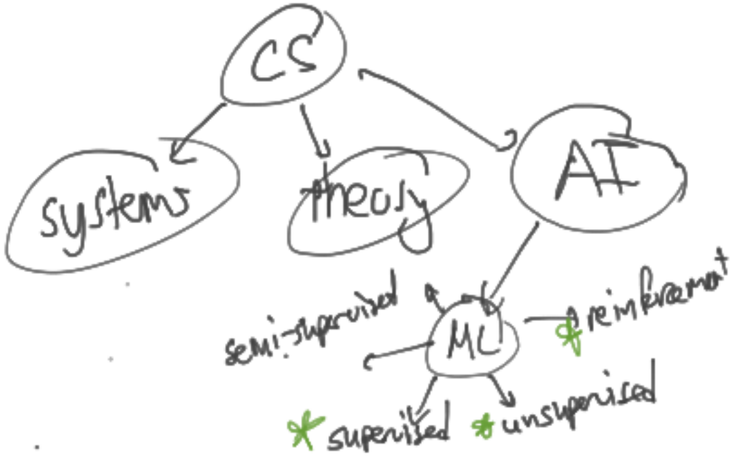

Agent: Child, dog, robot, program, ITS, diabetes treatment.

Environment: World, lab, software environment

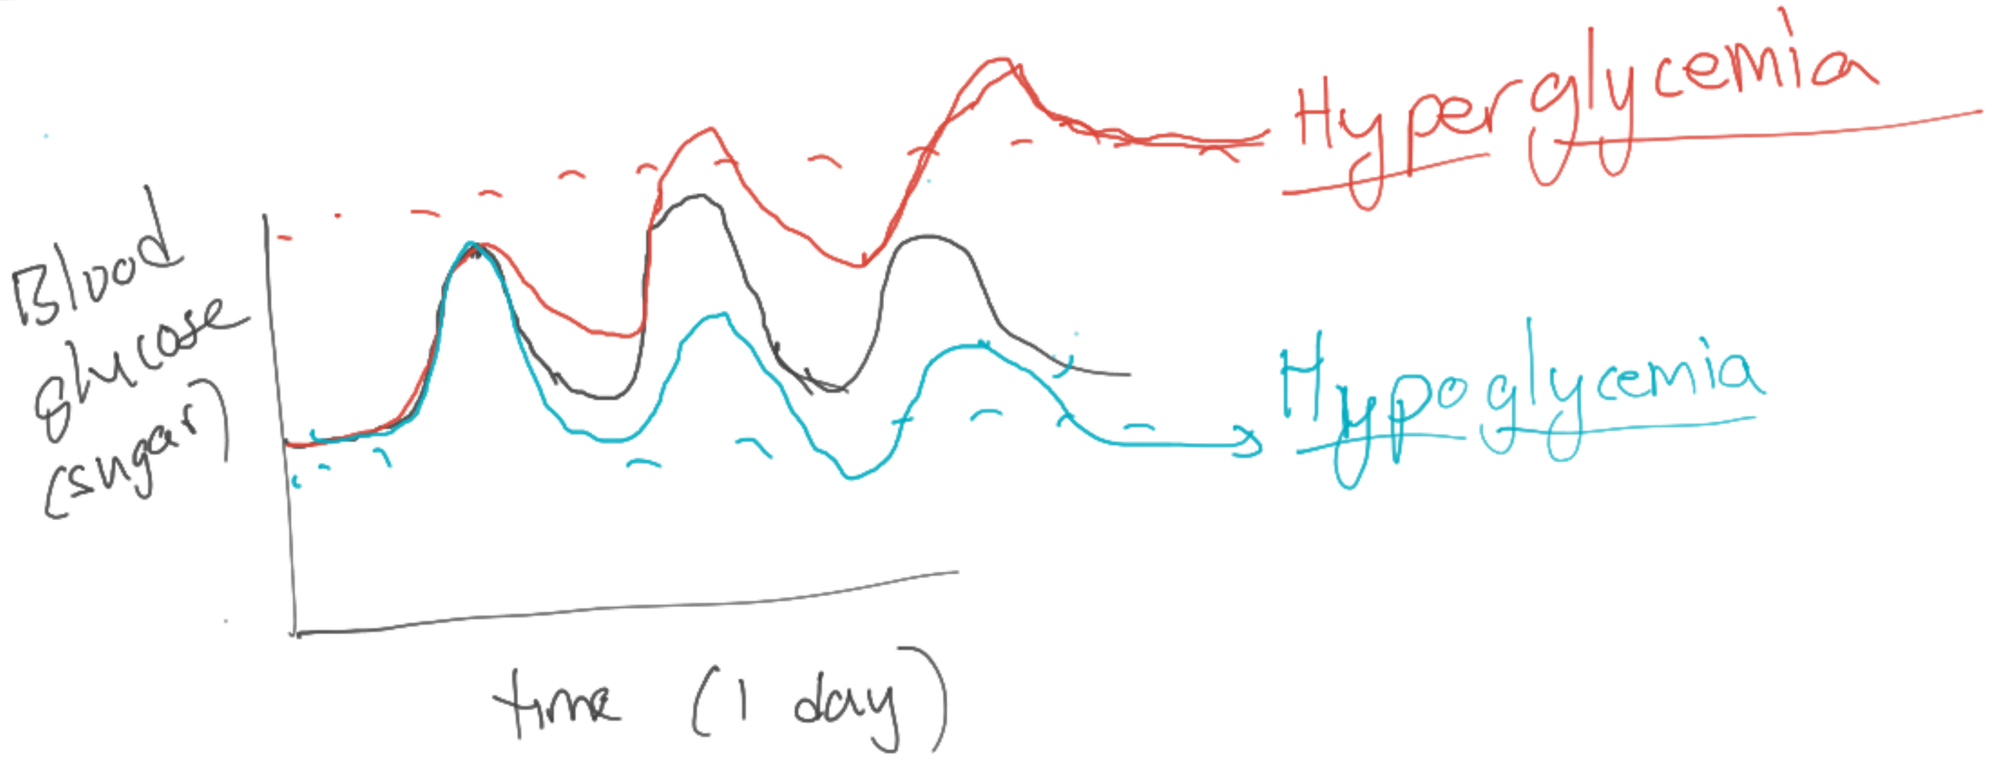

Neuroscience: How do animals learn?
A specific agent or set of agents.

- The study of some examples of learning and intelligence.

(ML)
RL : How can we make an agent that learns?

- The study of learning & intelligence. (in general)
(animal or not),

Dopamine ≈ Temporal difference error.

Two most related fields

- Operations research
- Control (adaptive / classical)

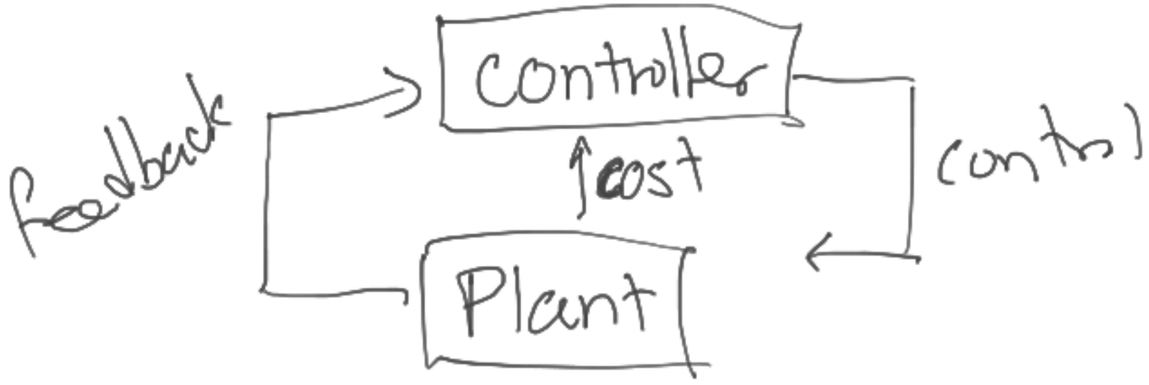

- Main difference is that these fields typically assume the environment (Plant) can and should be directly approximated.

When should you use RL?

→ As a last resort.

-

Key properties:

Next time!

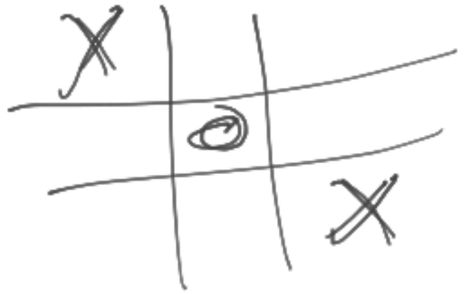

needs a colors.