ELEC 5530/6530
Robot Control
Lecture 10: Reinforcement Learning: Introduction

Instructor: Dr. Bosen Lian
Electrical and Computer Engineering
Auburn University
Spring 2026

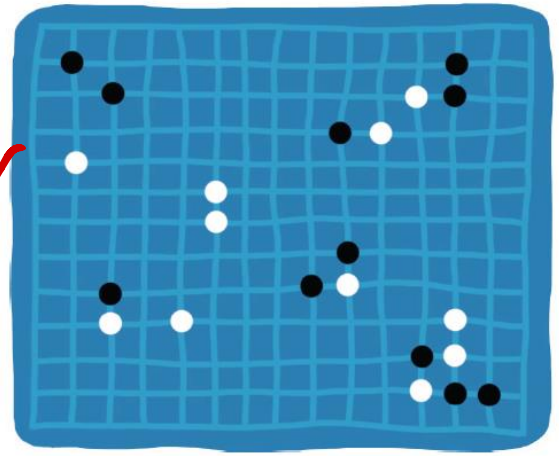

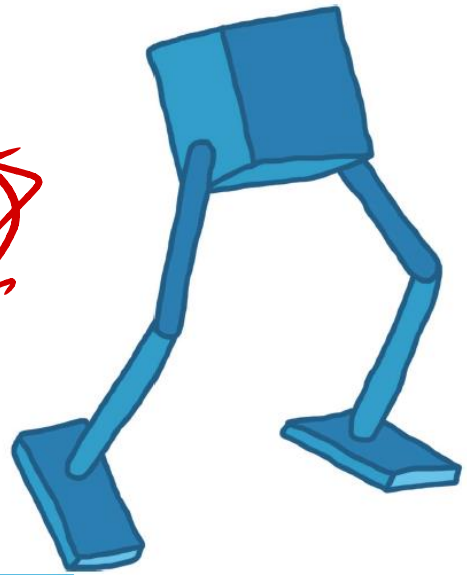

- What is Reinforcement Learning (RL)?

- Reinforcement learning focuses on learning **how to act** in different situations to **maximize numerical rewards**. Instead of being told what to do, the learner must discover which actions yield the highest returns through **trial and error**.
- **RL** has already mastered games by finding optimal moves in complex, uncertain environments. Engineers face identical challenges with real-world systems, leading to a key question: can RL also solve complex control tasks like **robotic walking** or **autonomous driving**?

Control Objective

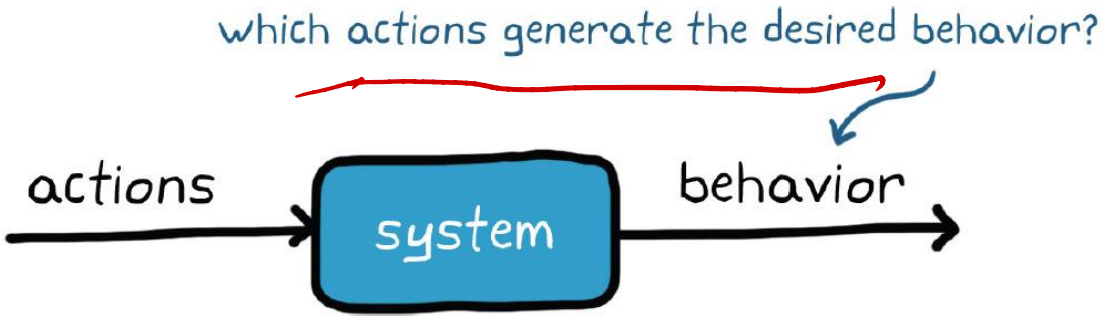

In feedback control, the controller uses state observations to improve performance and correct for random disturbances. Engineers utilize feedback signals and system models to design controllers that meet specific requirements.

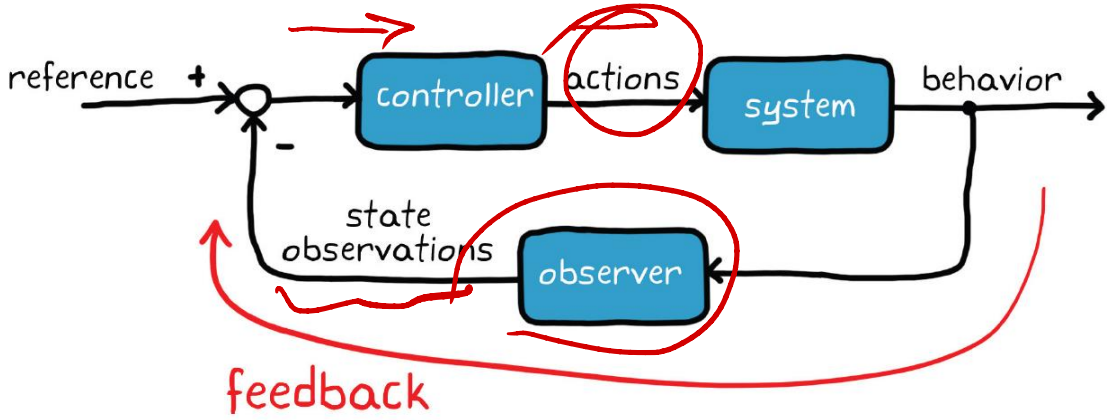

Control Problem

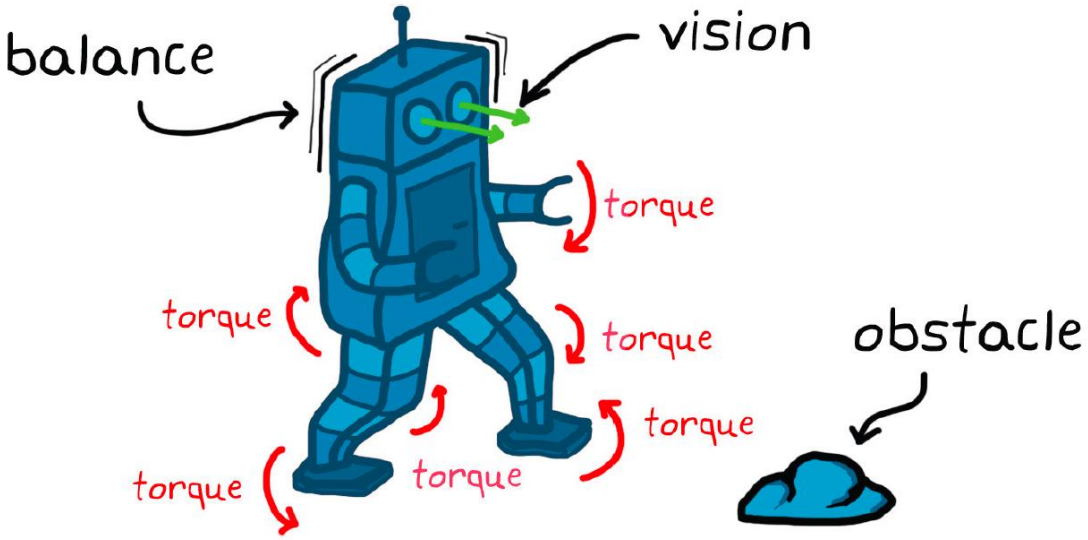

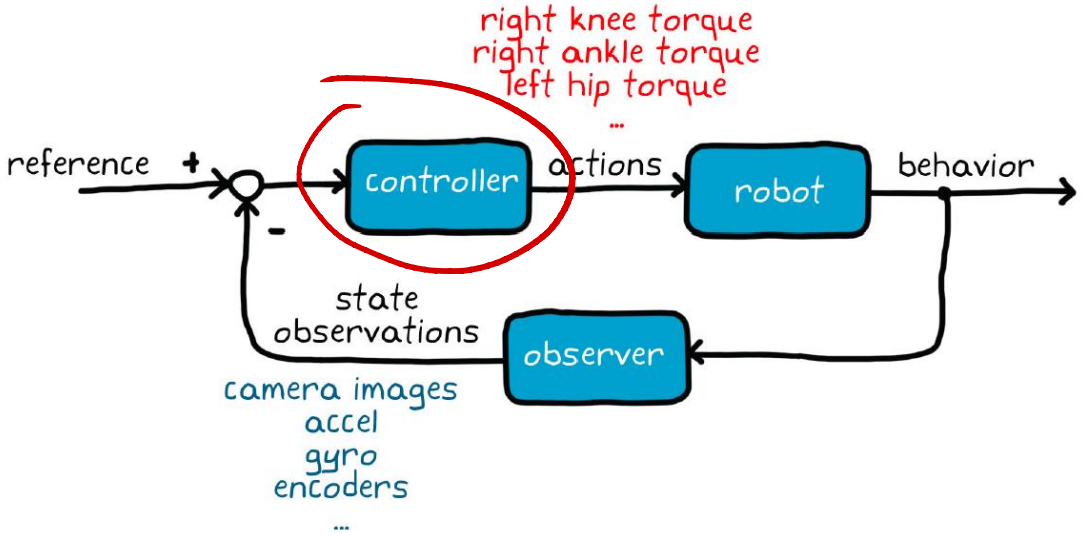

Control Solution

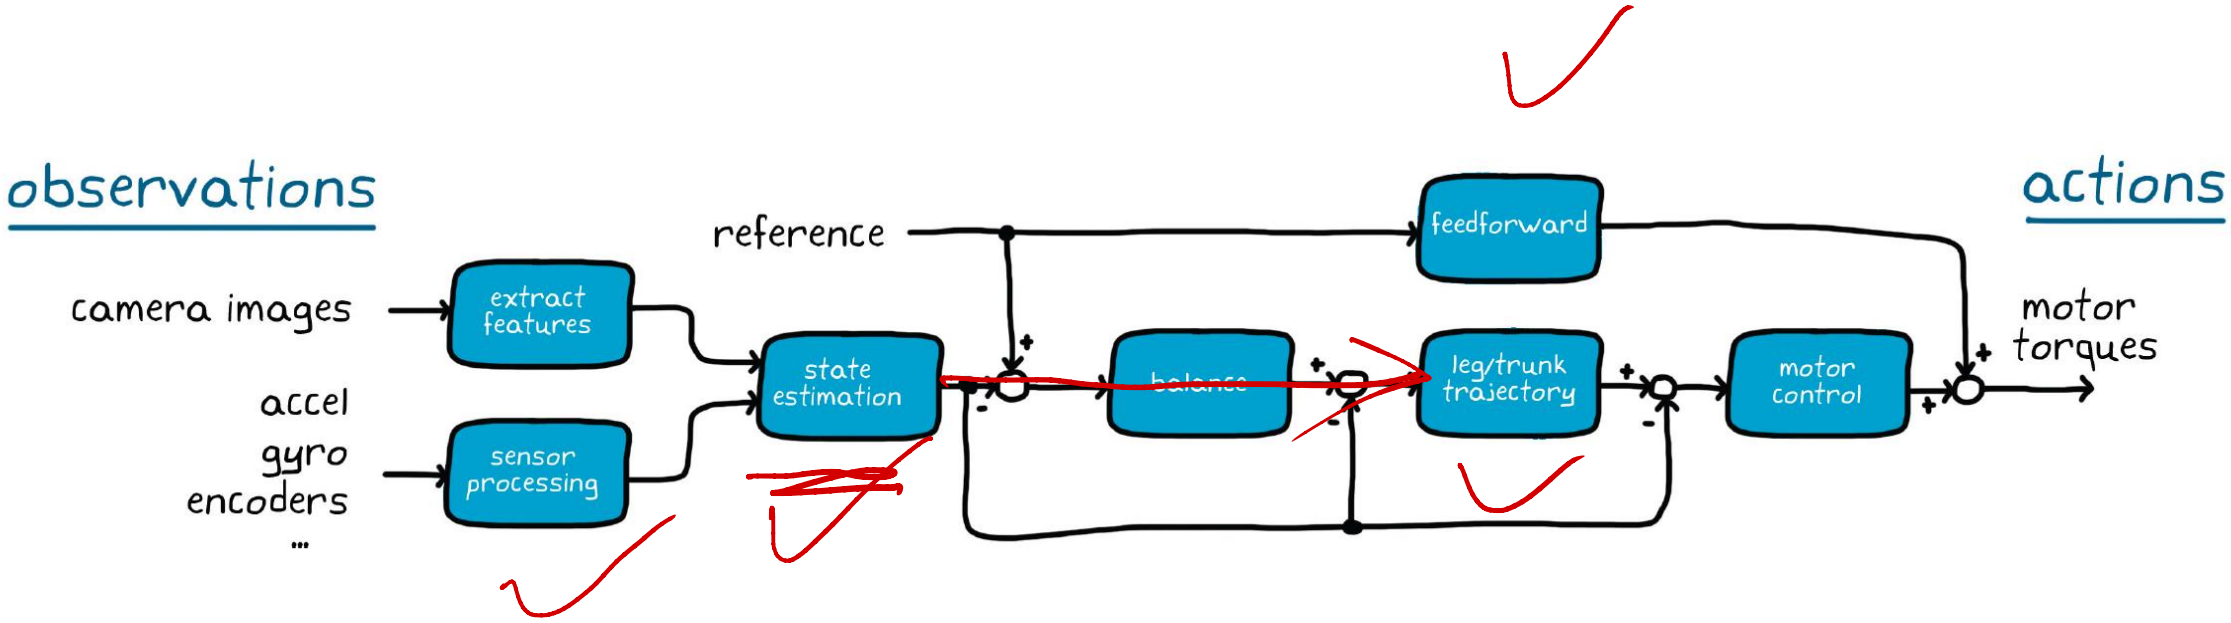

Why RL?

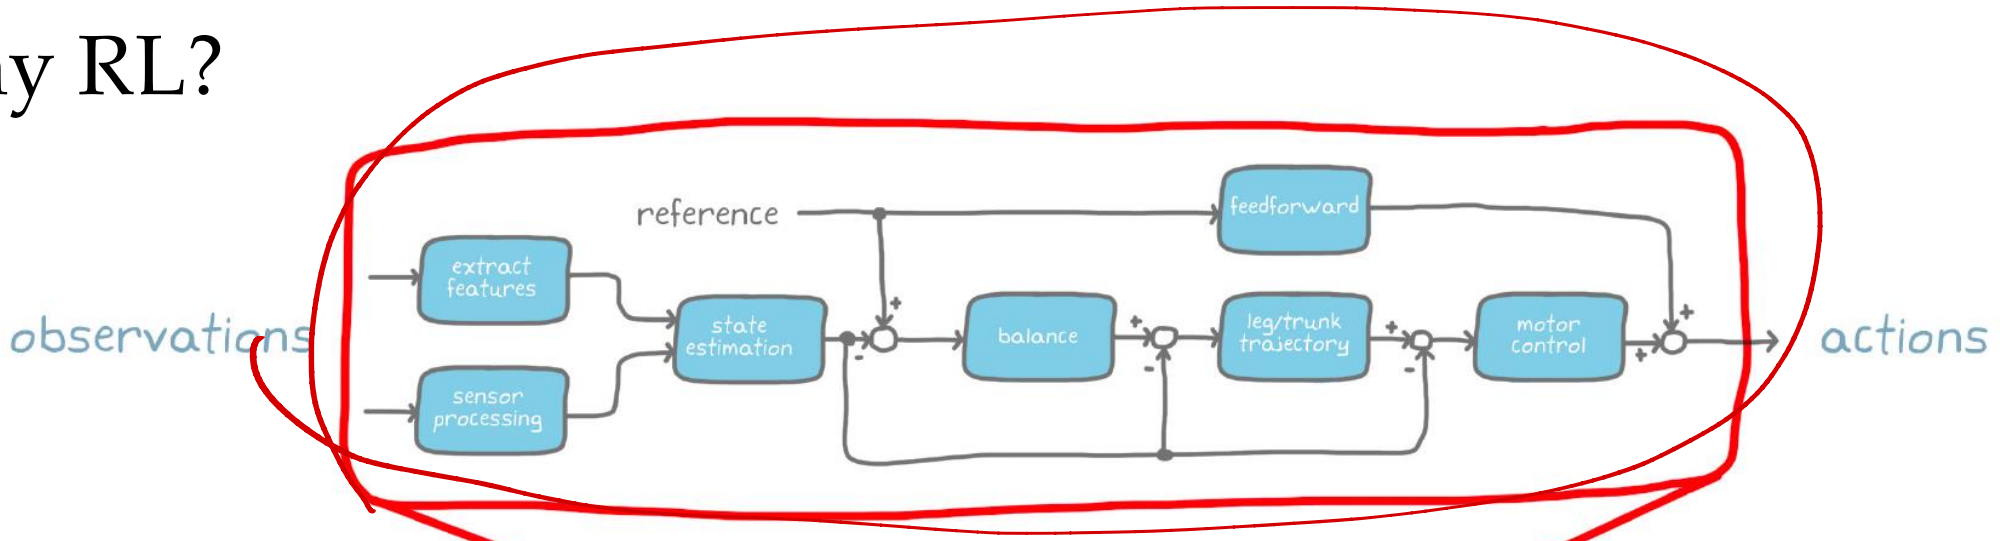

observations → function → actions

what does this function look like?

how do you design it?

Subset of Machine Learning

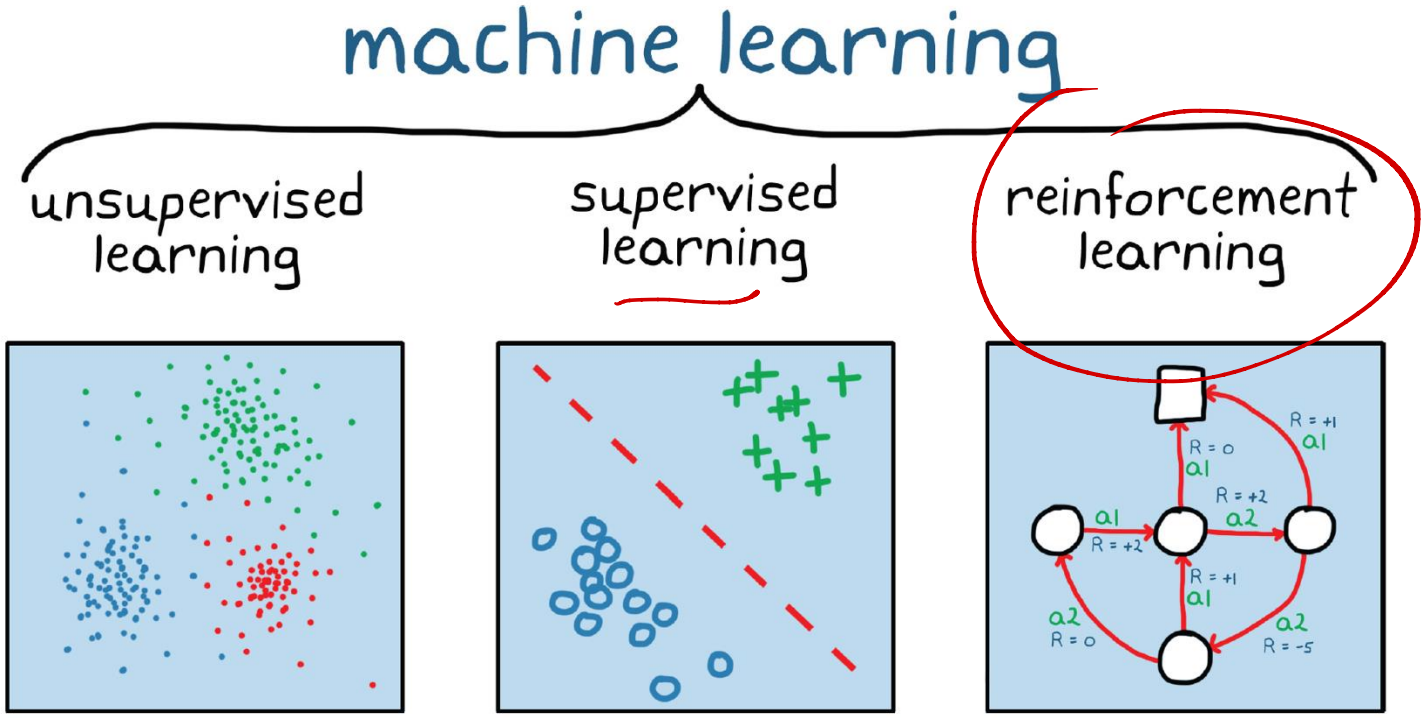

Unsupervised Learning

Identify patterns or hidden structures in unlabeled or uncategorized datasets.

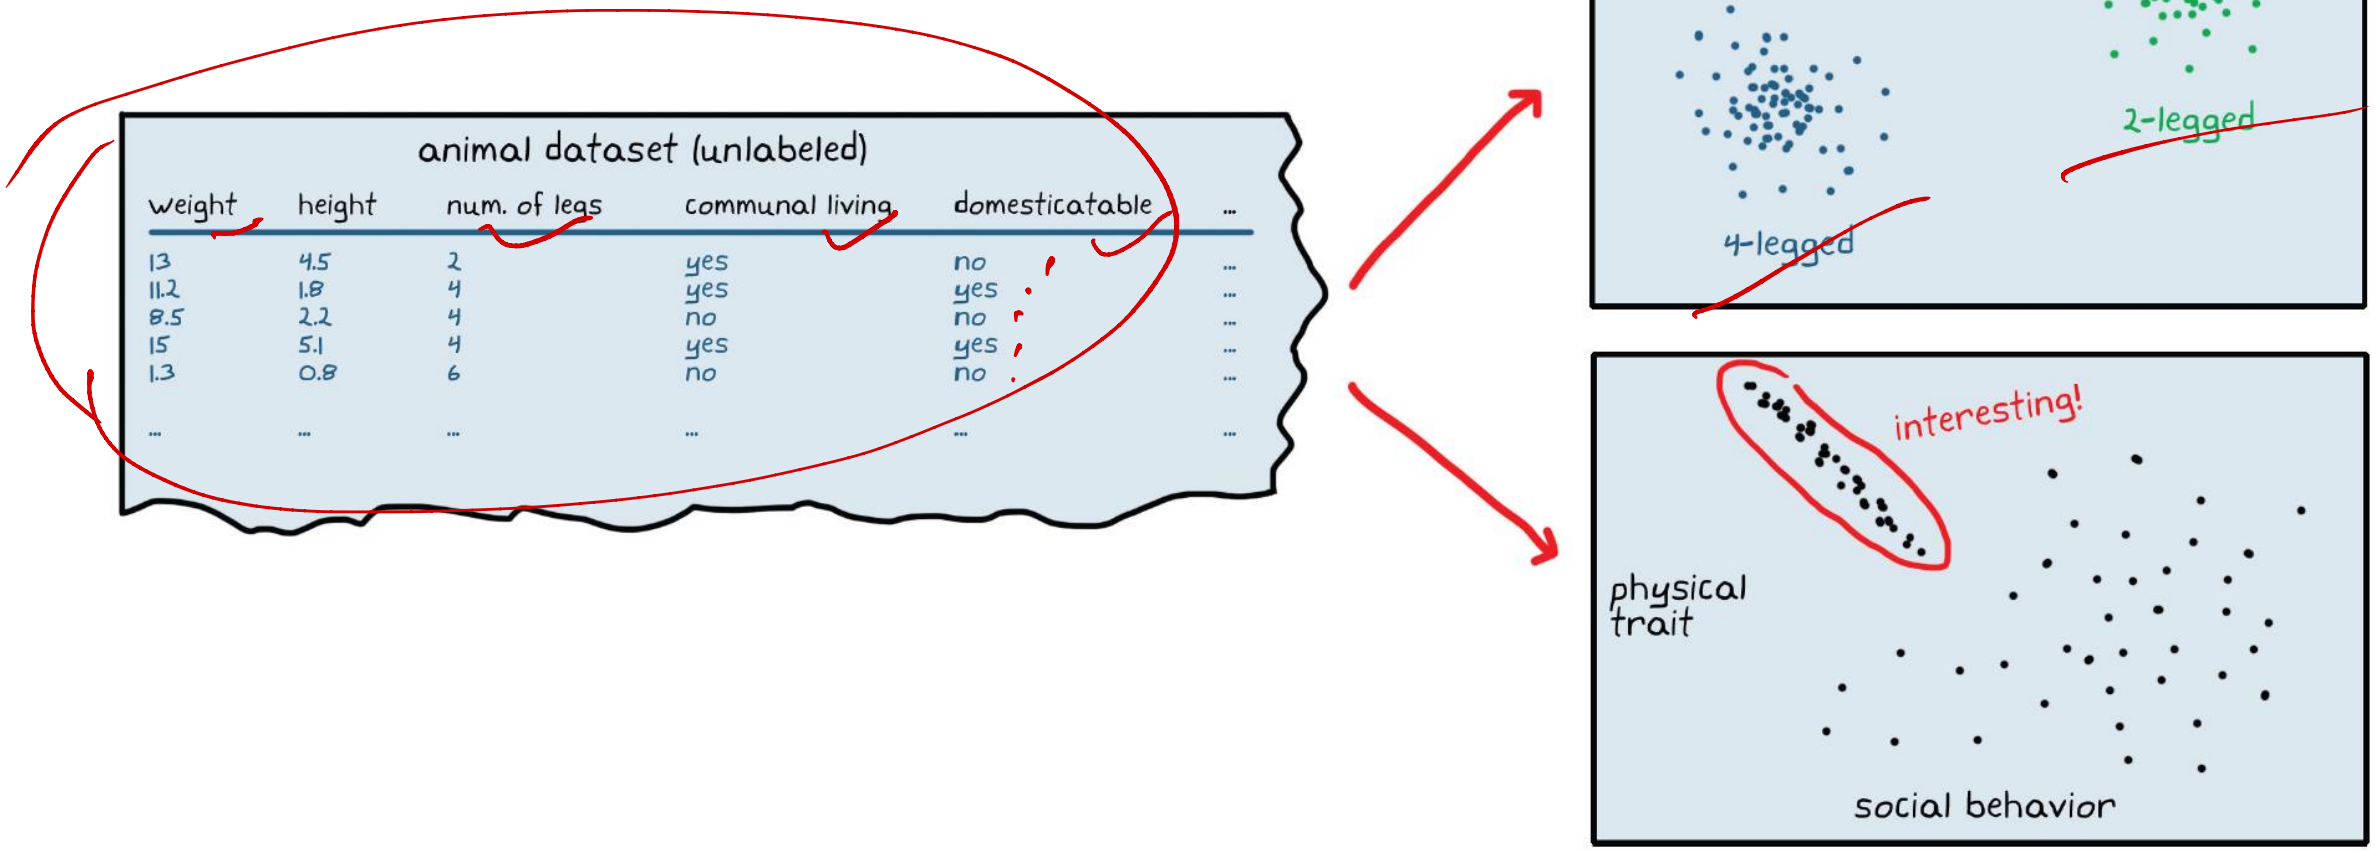

Unsupervised Learning

Identify patterns or hidden structures in labeled or categorized datasets to output the lables.

animal dataset (labeled)

species	weight	height	num. of legs	communal living	domesticatable	...
rat	1.3	1.1	4	yes	yes	...
robin	1.2	0.8	4	no	no	...
elephant	48.5	12.2	4	yes	no	...
rabbit	25	2.1	4	yes	yes	...
spider	0.1	0.2	8	no	no	...
...	...	...	...	...	...	...

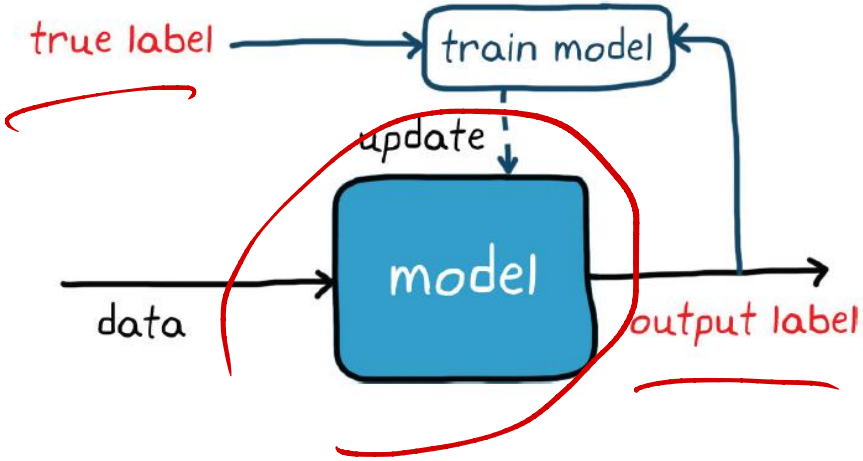

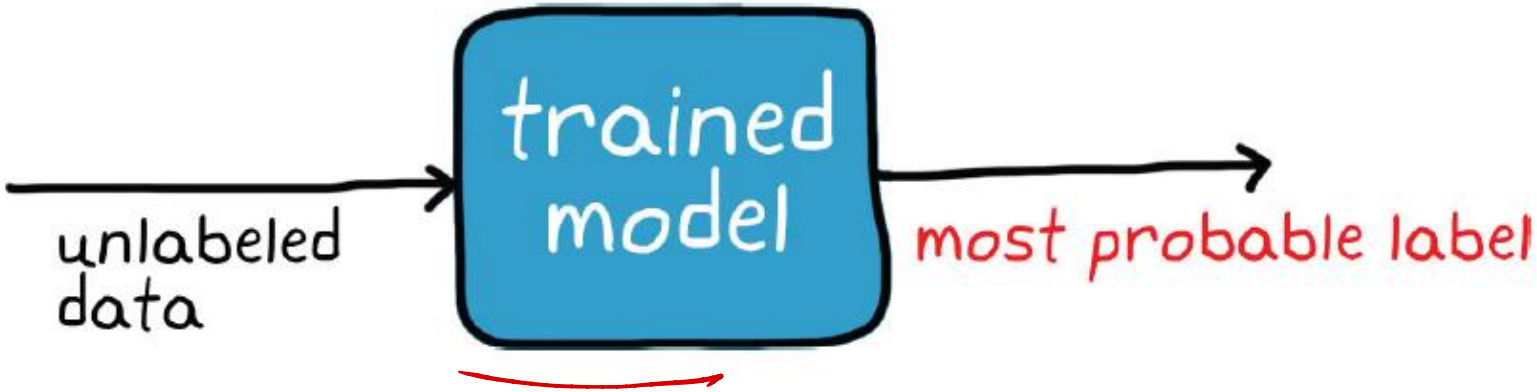

RL

Use dynamic environmental data. Instead of labeling data to determine the optimal sequence of actions to achieve the best results.

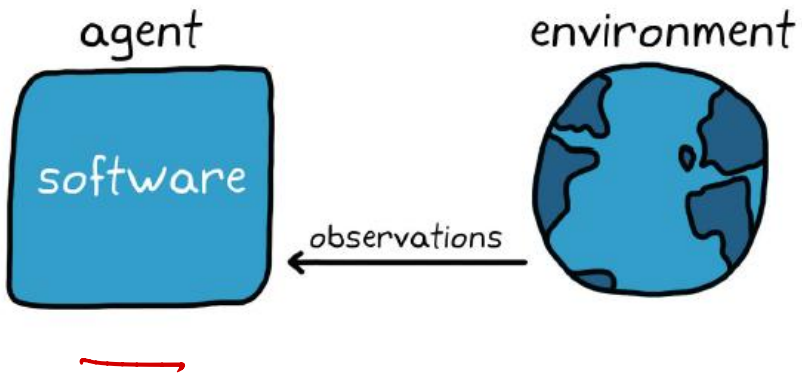

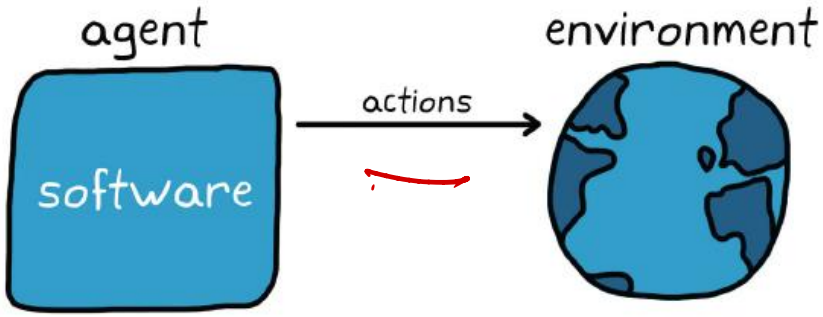

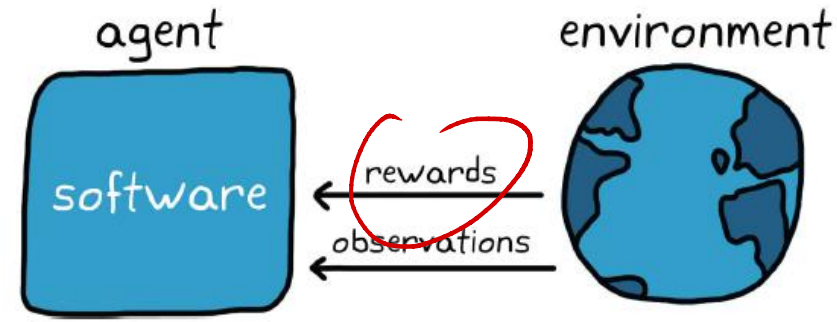

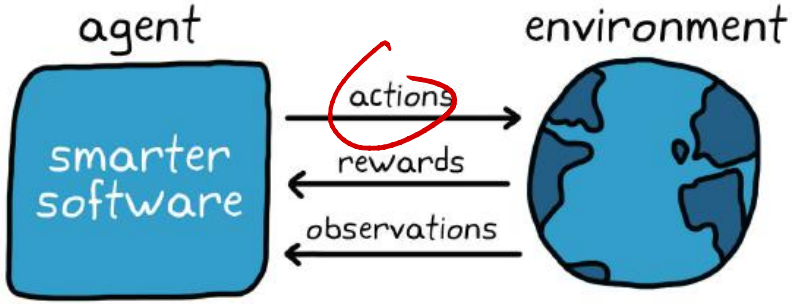

How RL works?

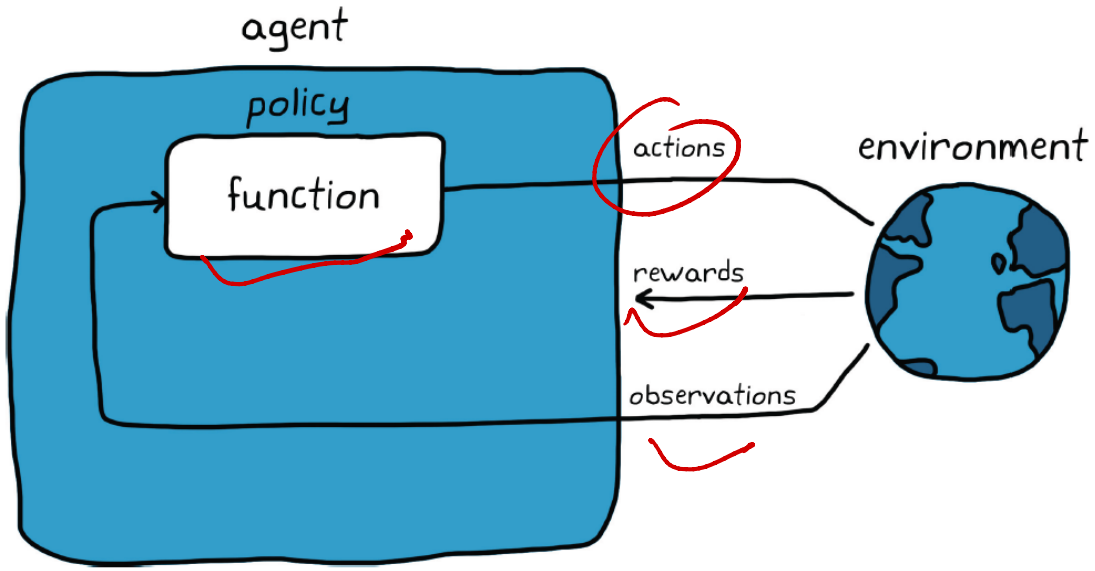
agent
policy
function
actions
environment
rewards
observations

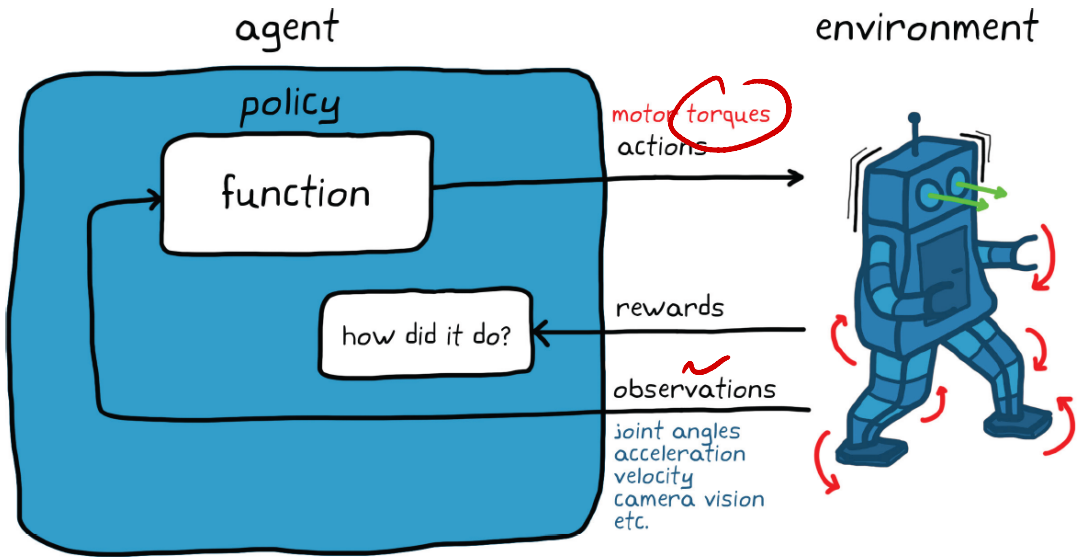
agent
environment
policy
function
how did it do?
motor torques
actions
rewards
observations
joint angles
acceleration
velocity
camera vision
etc.

Learn Optimal Control/Action Policy

Dynamic Environment
Need Update Frequently

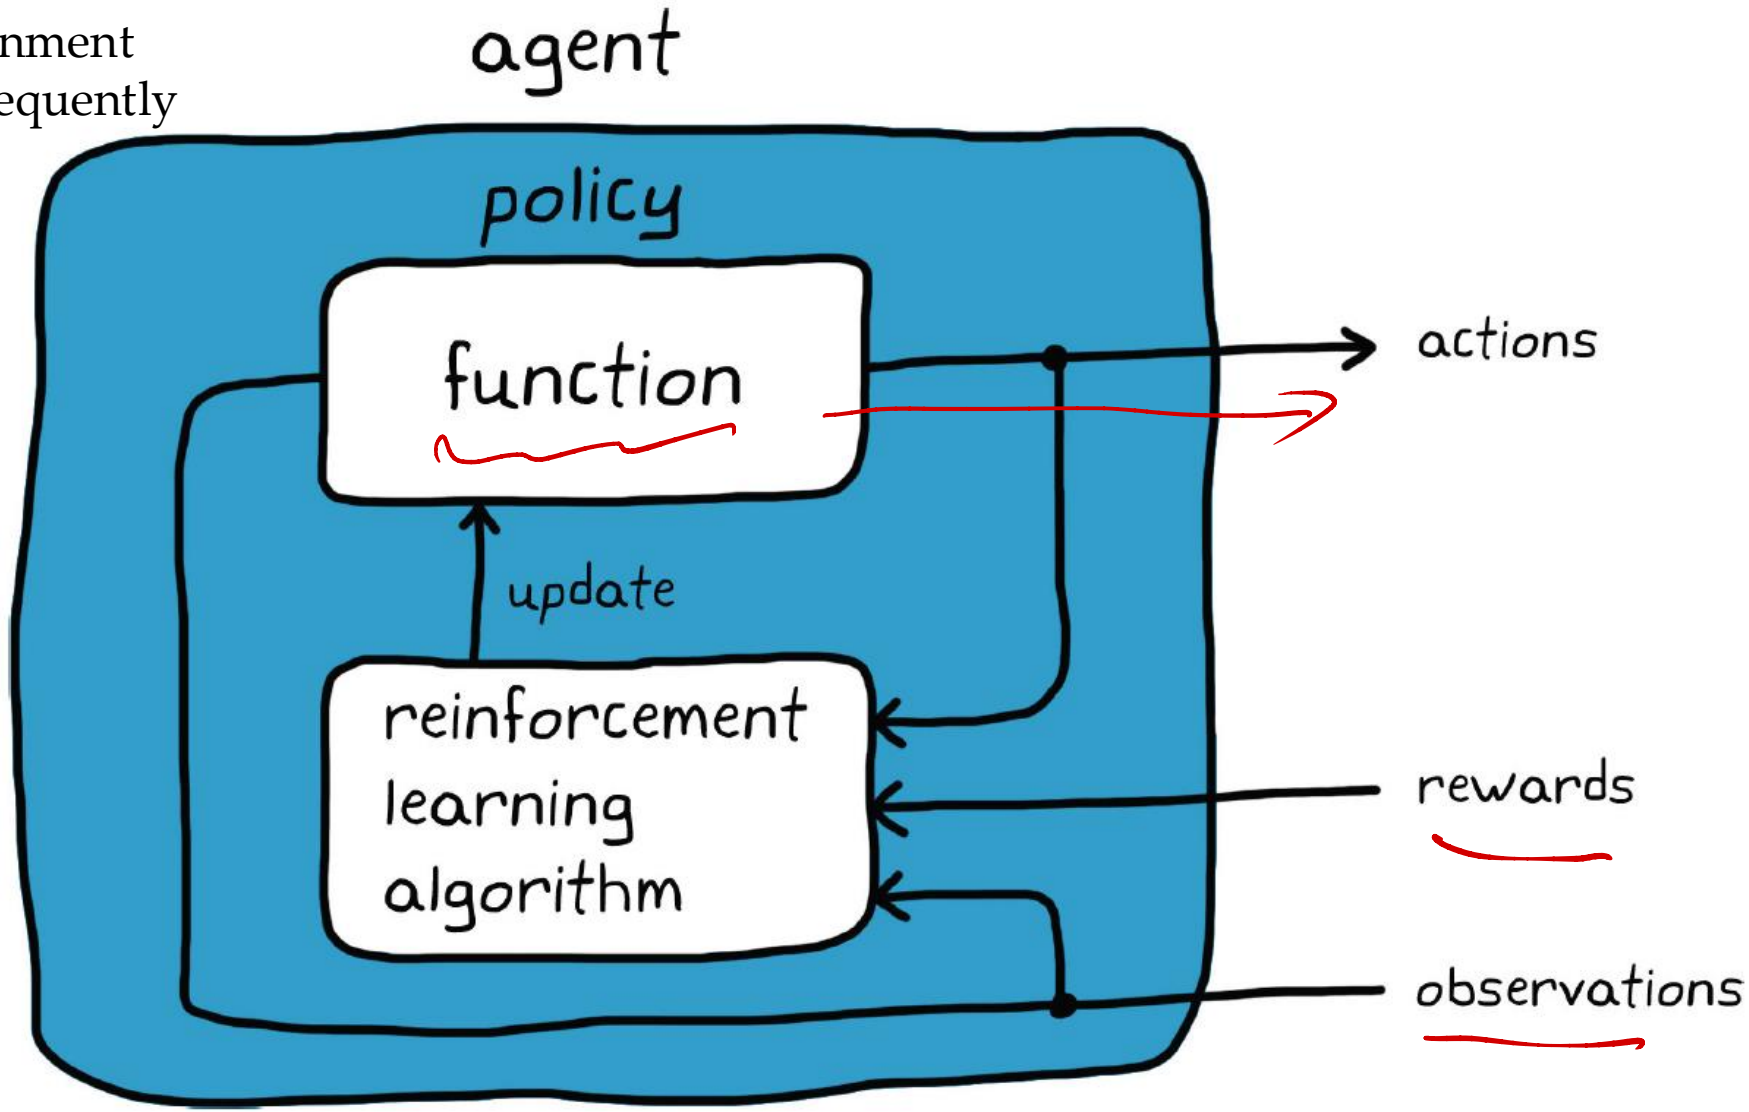

What is learning in RL?

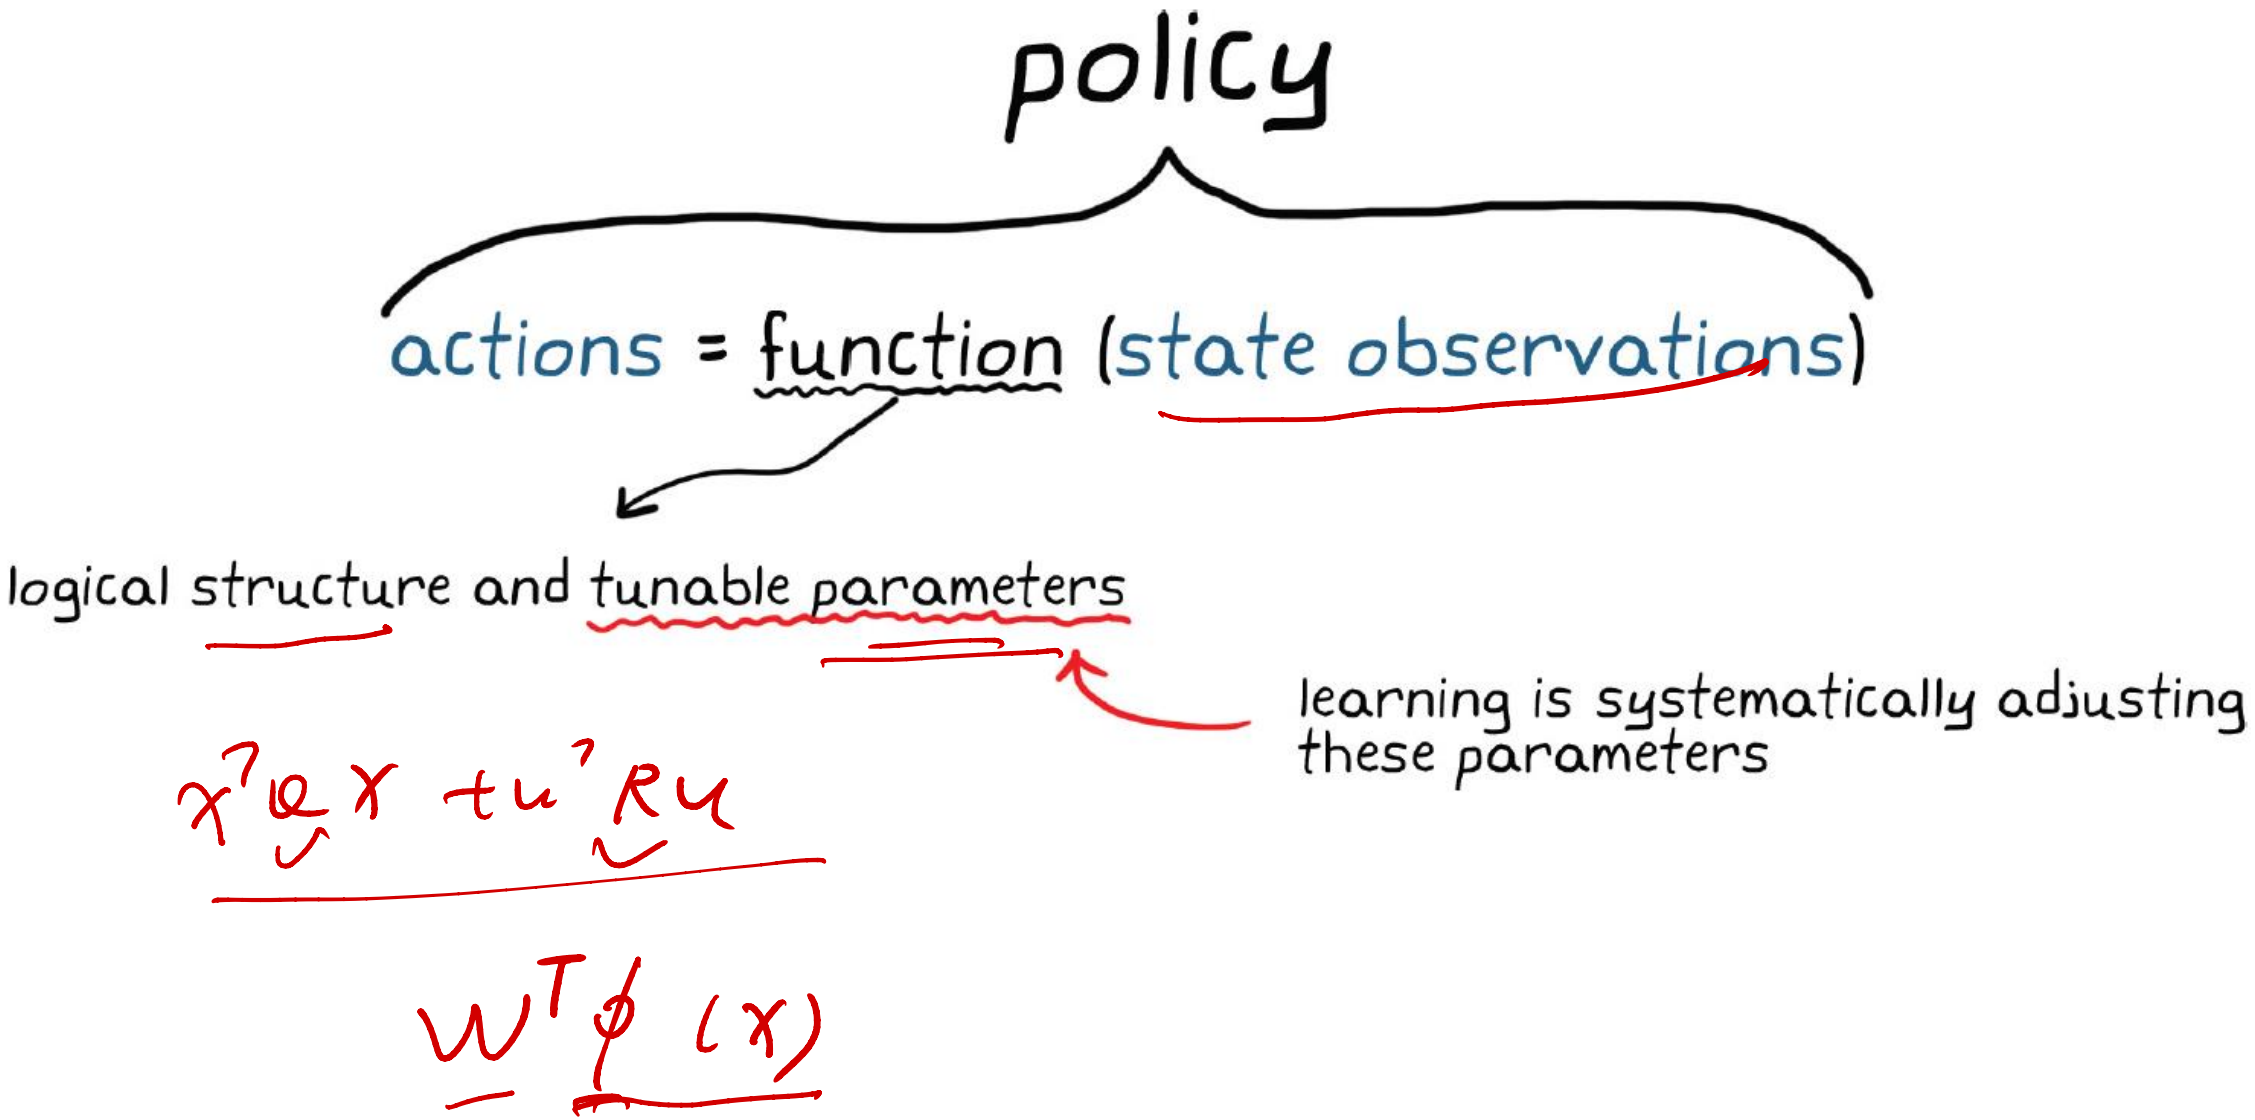

How is RL similar to feedback control?

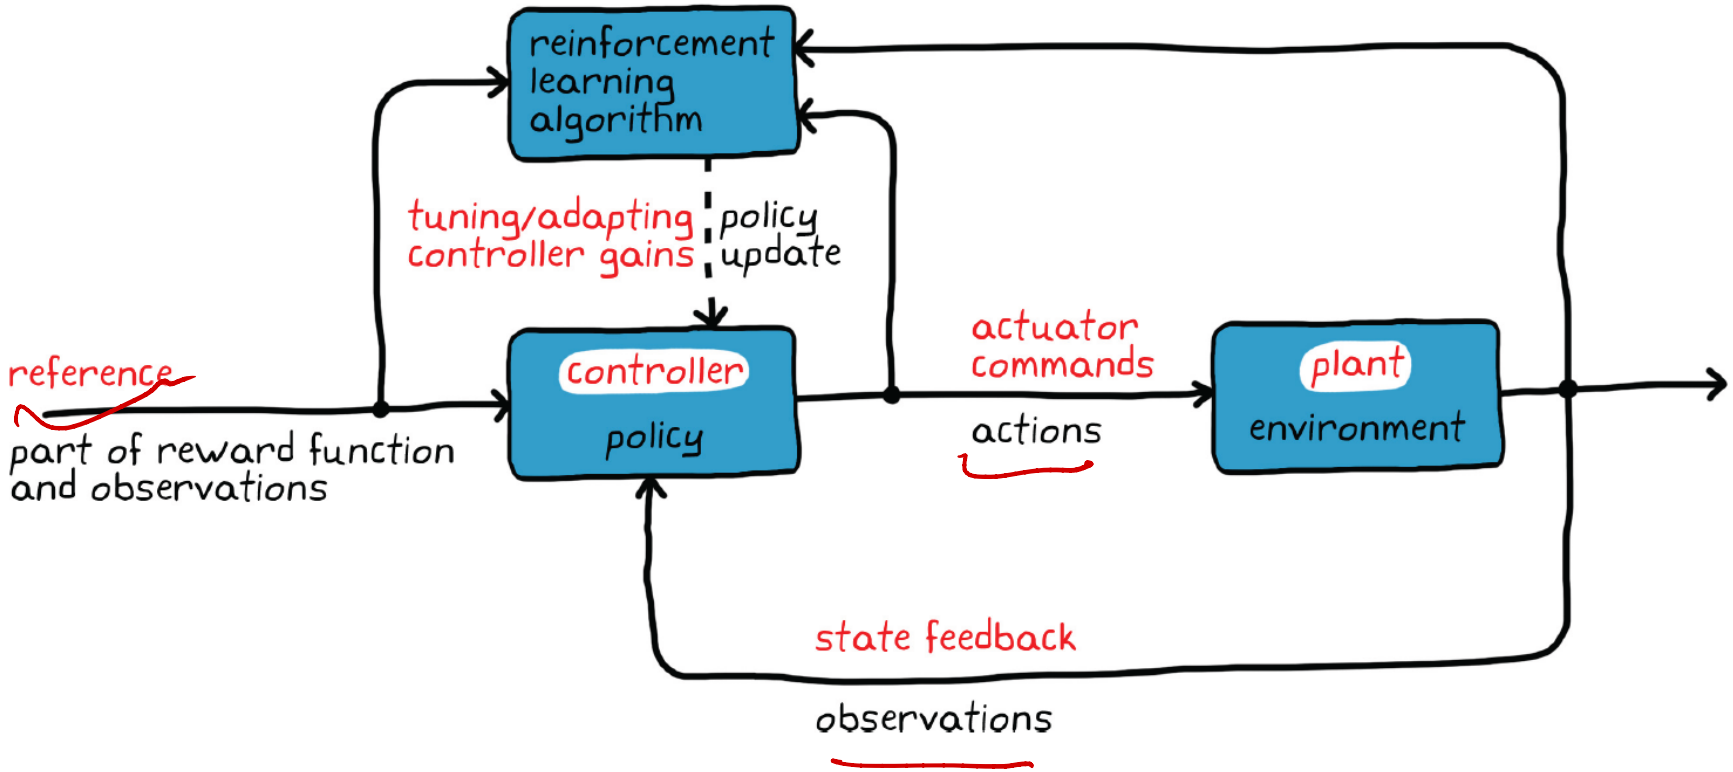

Process of RL

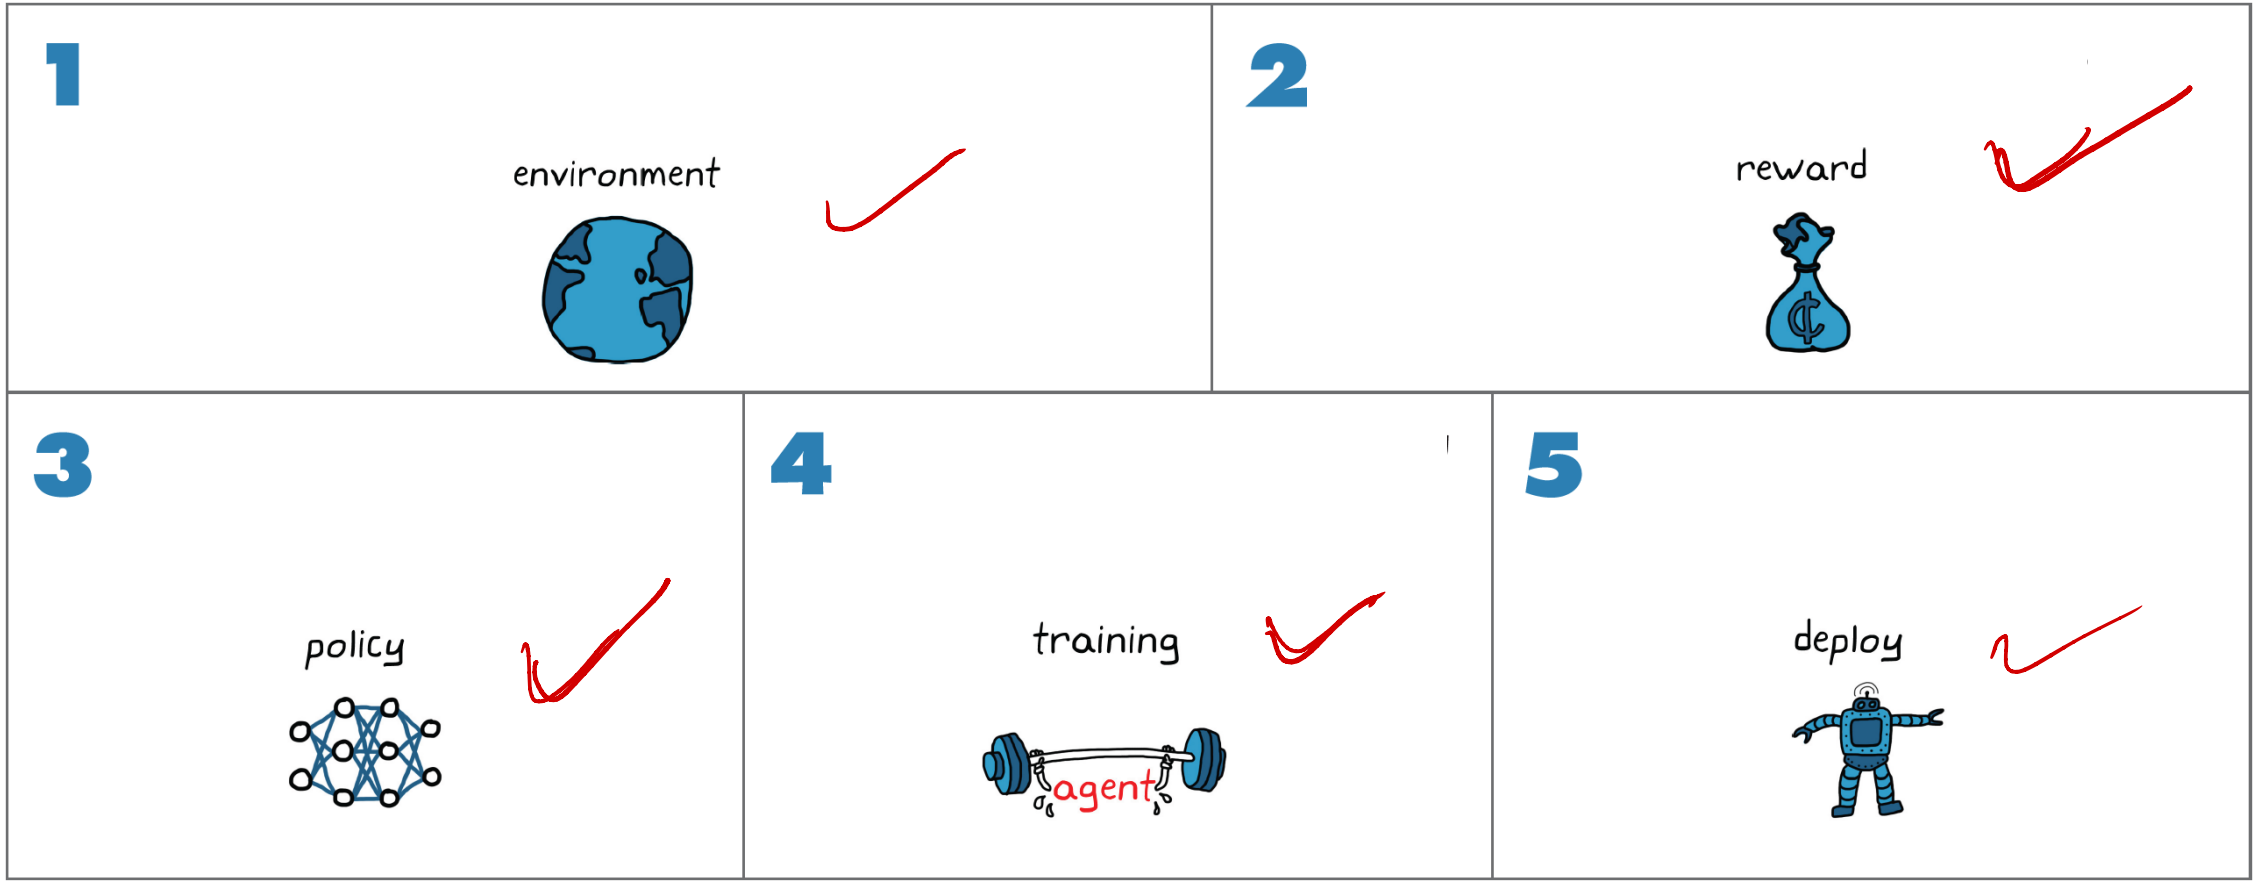

Environment

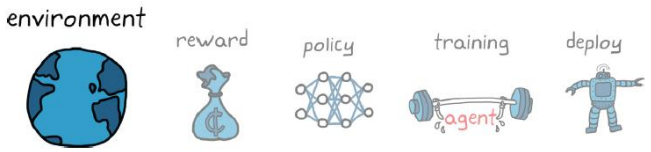

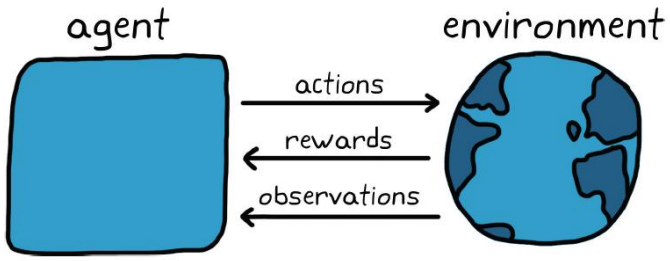

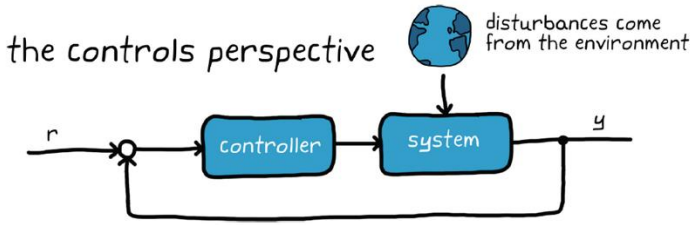

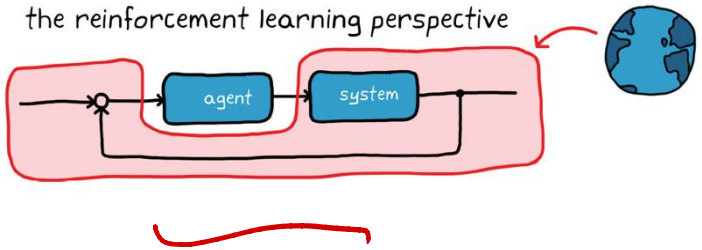

Environment: Model-Free RL

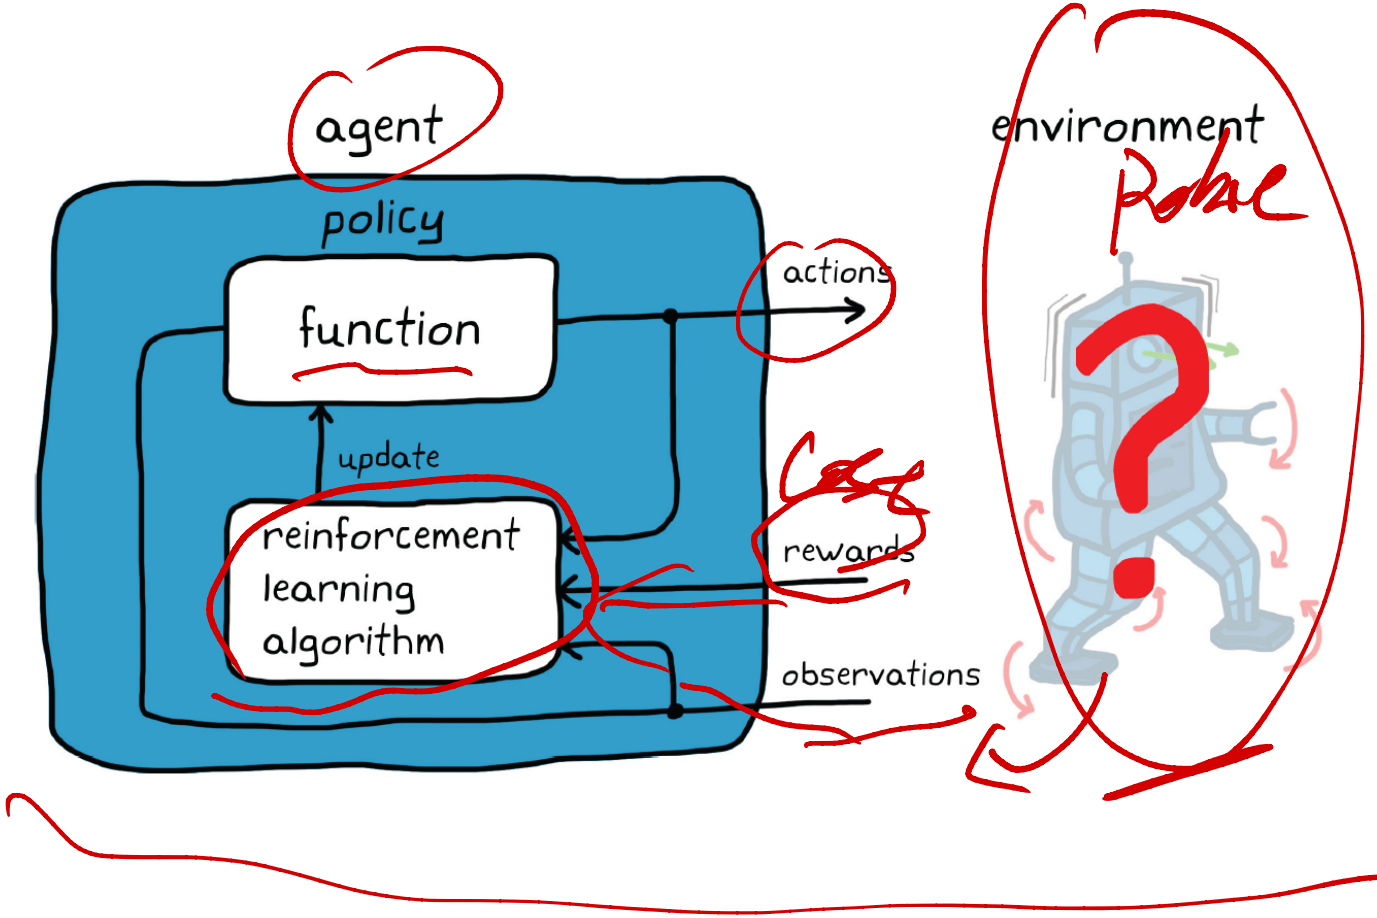

Environment: Model-Based RL

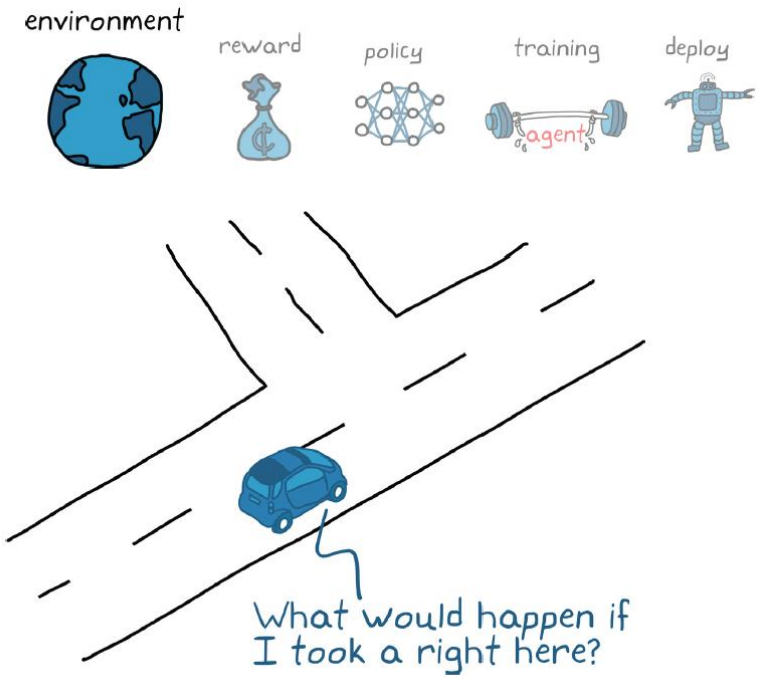

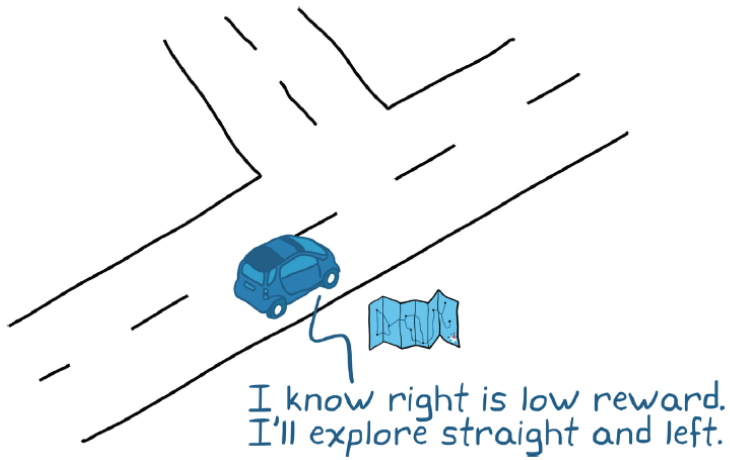

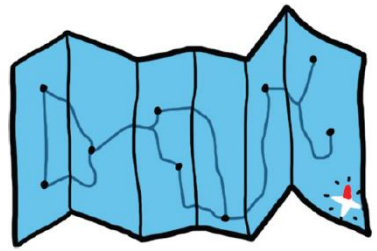

this road map should help!

Model-Free RL vs Model-based RL