Types of AI & ML

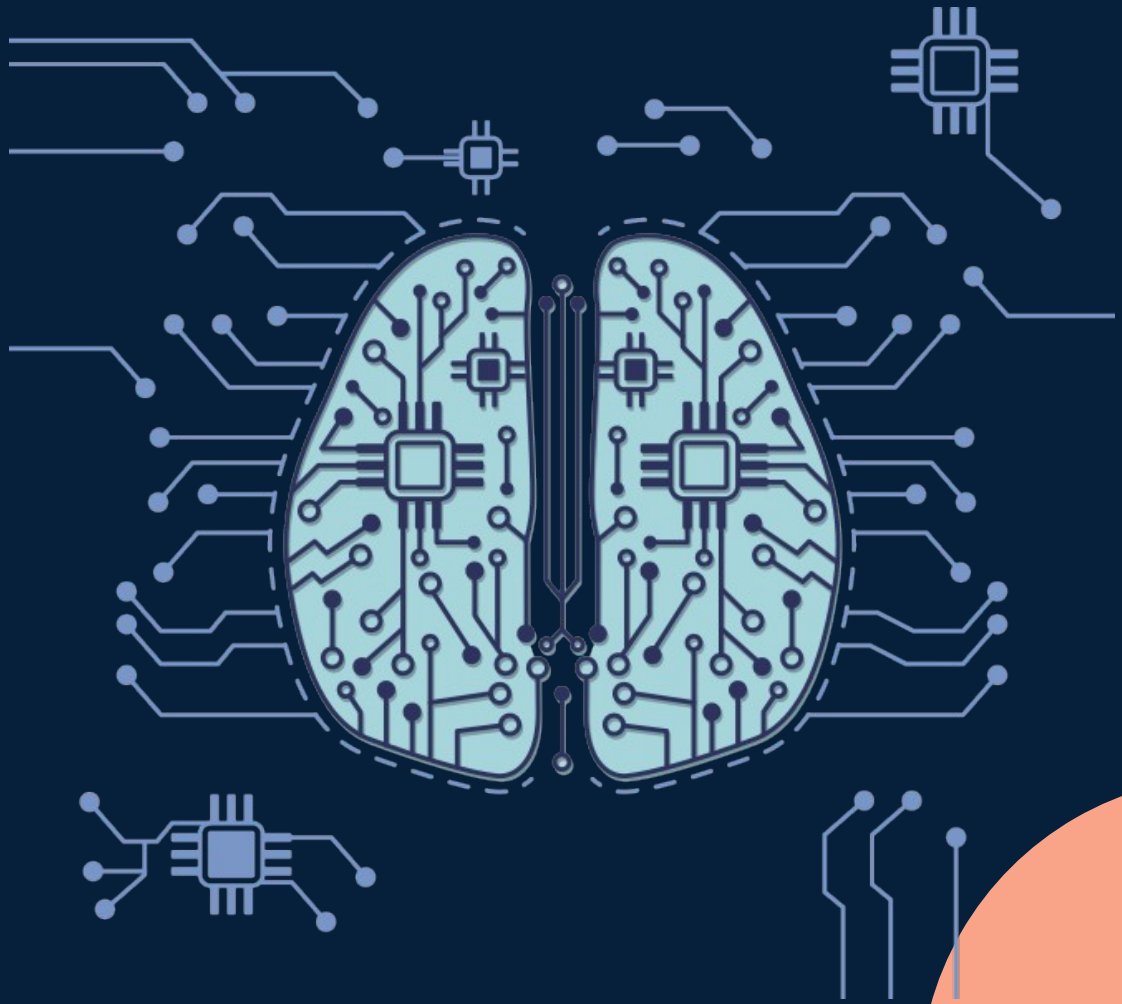

PLAY · EXPERIENCE · LEARN

A POSTERIORI

Types of AI (Summary)

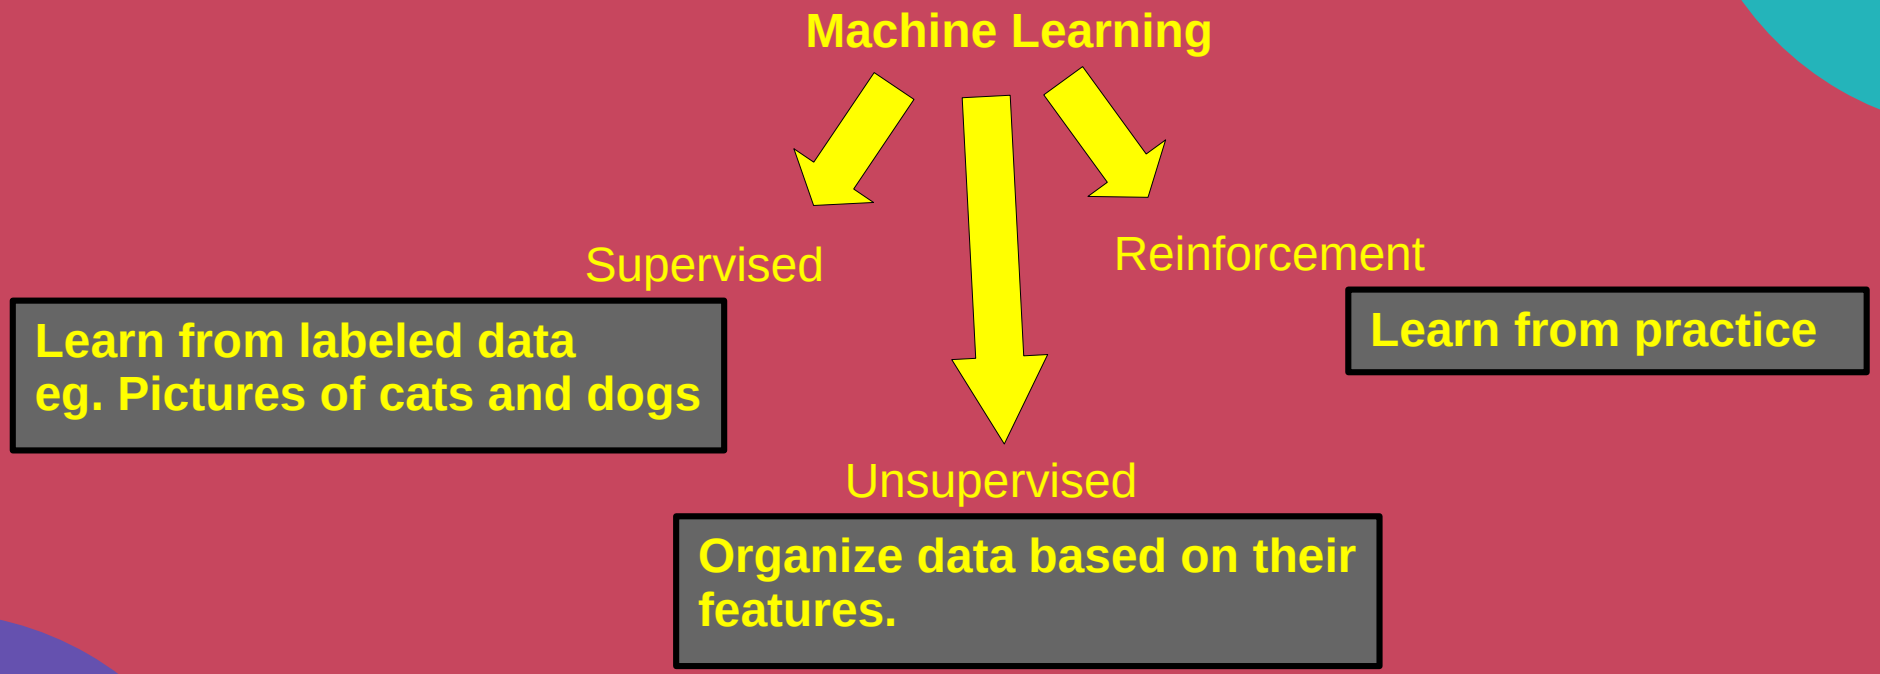

Supervised Learning (Review)

Training

"Dog" "Cat"

Machine Learning Model

Performing

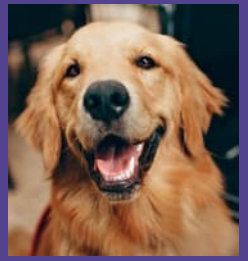

Machine Learning Model

"It's a **Dog**"

Unsupervised Learning (Review)

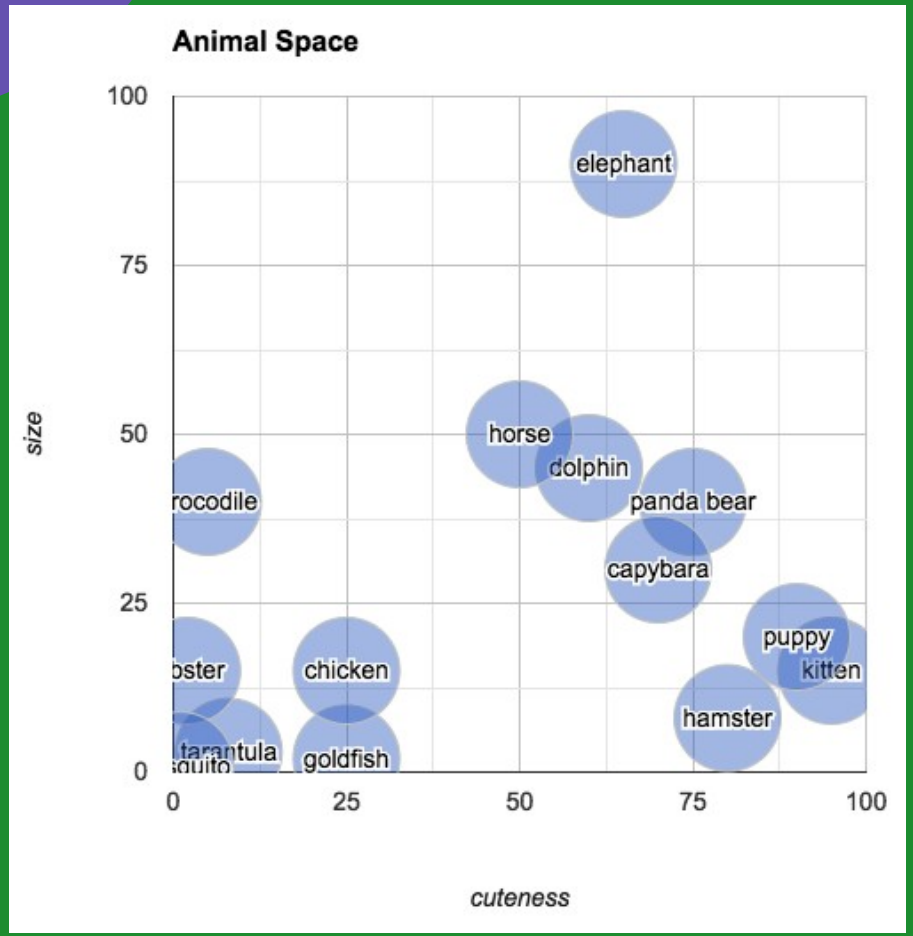

- Organize data by their features
- Doesn’t understand what the data mean

Reinforcement Learning

Answer T / F

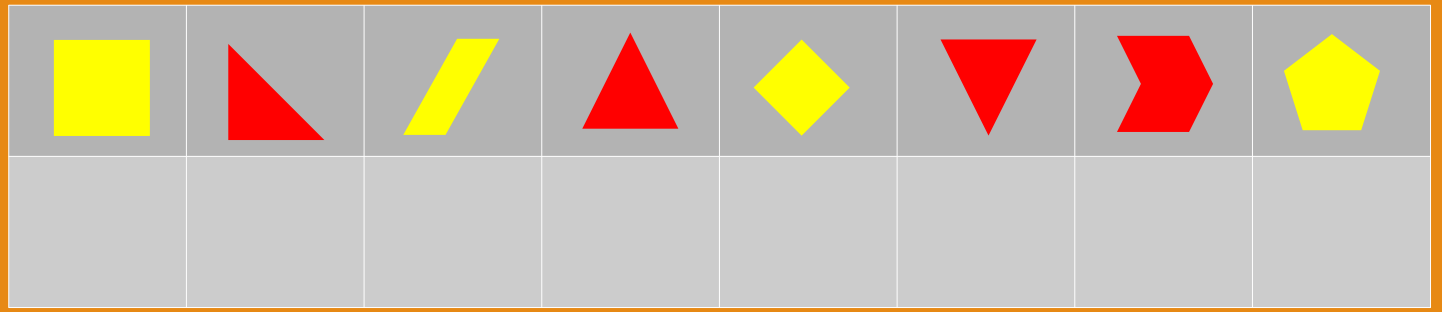

Guess the answer and I’ll let you know if it’s right or wrong.

Reinforcement Learning

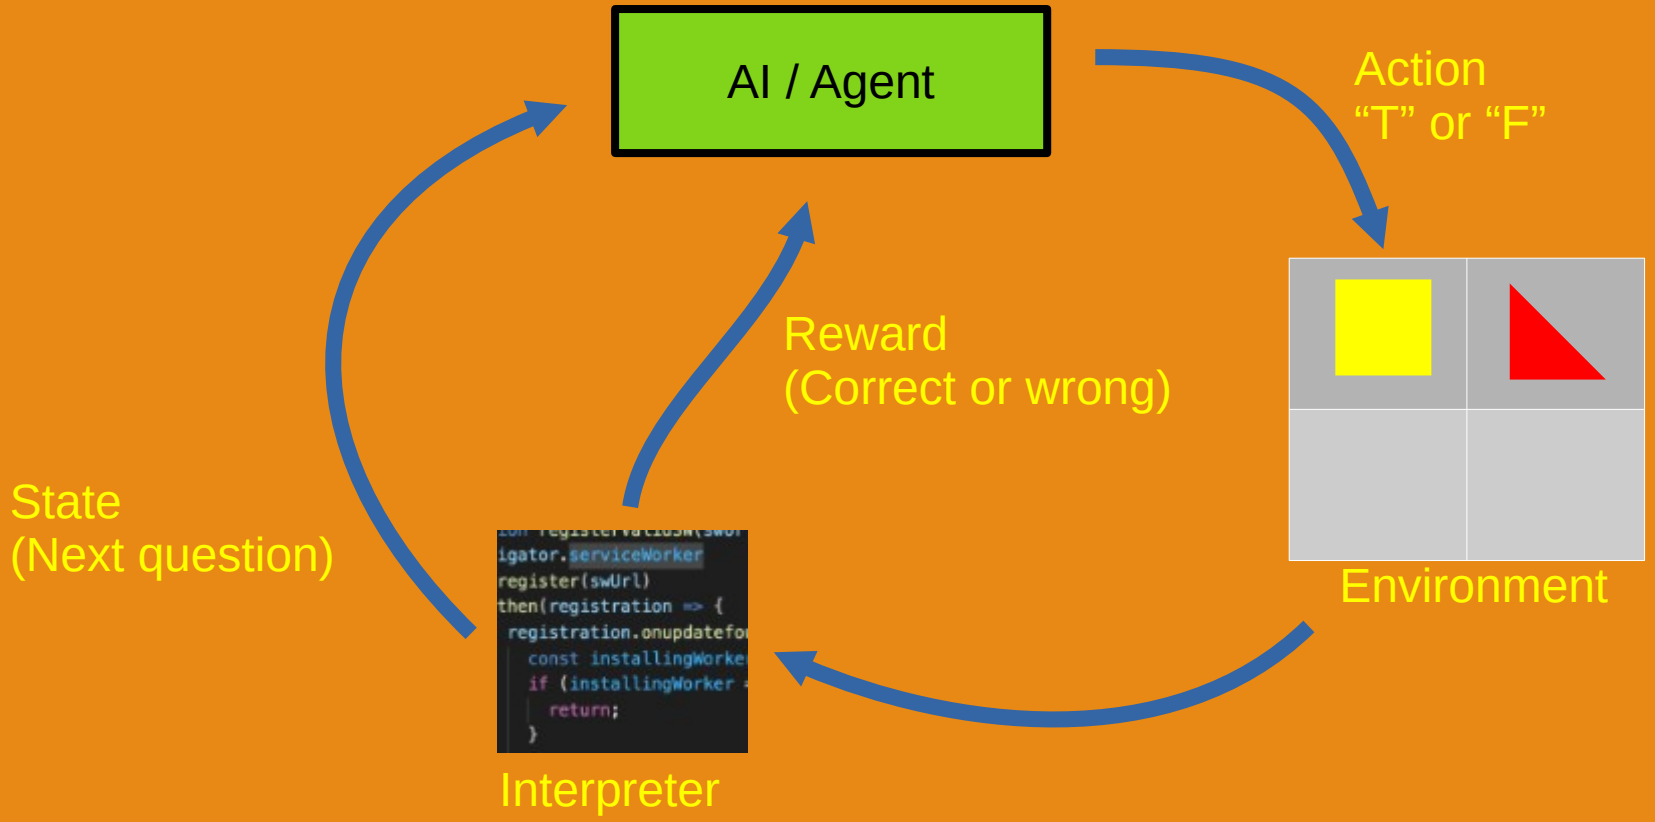

Reinforcement Learning

- AI / Agent
  - Machine Learning Algorithm
  - Many different types...
  - Needs to be tuned
  - Choice of algorithm and tuning affects learing speed

Algorithm	Description
Monte Carlo	Every visit to Monte Carlo
Q-learning	State–action–reward–state
SARSA	State–action–reward–state–action
Q-learning - Lambda	State–action–reward–state with eligibility traces
SARSA - Lambda	State–action–reward–state–action with eligibility traces
DQN	Deep Q Network
DDPG	Deep Deterministic Policy Gradient
A3C	Asynchronous Advantage Actor-Critic Algorithm
NAF	Q-Learning with Normalized Advantage Functions
TRPO	Trust Region Policy Optimization
PPO	Proximal Policy Optimization
TD3	Twin Delayed Deep Deterministic Policy Gradient
SAC	Soft Actor-Critic

Reinforcement Learning

- Interpreter
  - Provides rewards for correct behaviour
  - Usually need to be custom written for each application
  - Writing a good reward function is essential; AI will always try to maximise rewards
    - eg. If we reward a self-driving car based only on travelling time, then the AI will try to reach as quickly as possible, even if it means speeding, running red lights, or driving over pedestrians

Reinforcement Learning

- Environment
  - Reinforcement learning needs a lot of practice; this isn’t always practical or safe
  - Synthetic data may be used to provide safe training opportunity

Amazon DeepRacer:

Virtual environment for training

Testing with real robot

Reinforcement Learning

- Demo
  - Cart Pole problem
  - Compare the purely random solution and the ML solution

Applications of Reinforcement Learning

- Self-driving cars
- Robotics
- Image processing
- Generate art

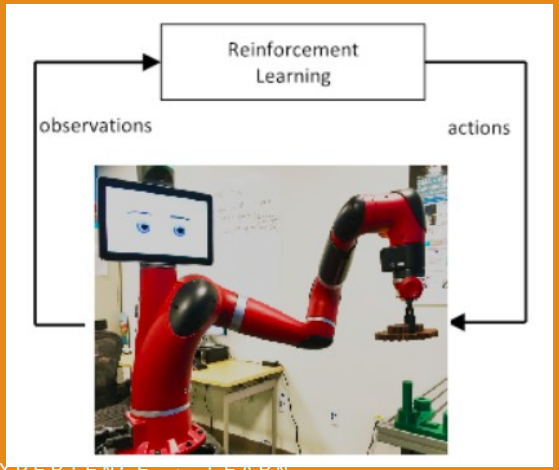

Benefits and Challenges

- Pro
  - Good for solving complex problems where humans may not know the solution
  - Works great when it can practice without limits (eg. Playing video games)
- Con
  - Simple problems can be better solved with a human written solution
  - Needs to practice a lot to learn (eg. Need to run over a lot of pedestrains before learning to avoid them)

Ethics in Reinforcement Learning

- How we balance rewards will affect the AI's behaviour
- eg. In a self-driving car, which should we reward more? Protecting the driver or pedestrians?
  - Which is the "right" thing to do?
  - Which will your customers prefer?

A POSTERIORI

Play · Experience · Learn

- Created by A Posteriori LLP
- Visit http://aposteriori.com.sg/ for more tips and tutorials
-
- Some vector art created by freepik - www.freepik.com

PLAY · EXPERIENCE · LEARN

A POSTERIORI