

Reinforcement Learning

ML Techniques / Paradigms

- ★ • Supervised learning (labeled data)
- Unsupervised " (clustering)
- Reinforcement (delayed labeling)
- Hybrid

Applications

- Playing games
- Recommender systems
- Traffic control
- Finance (Stock trading)

Example


Robot moves from point A to point B

Sensors/Capabilities

- Camera
- Distance sensor
(sonar, radar, lidar)
- Wheels / motors / drivers
- GPS
- Compass
- Payload
- Map

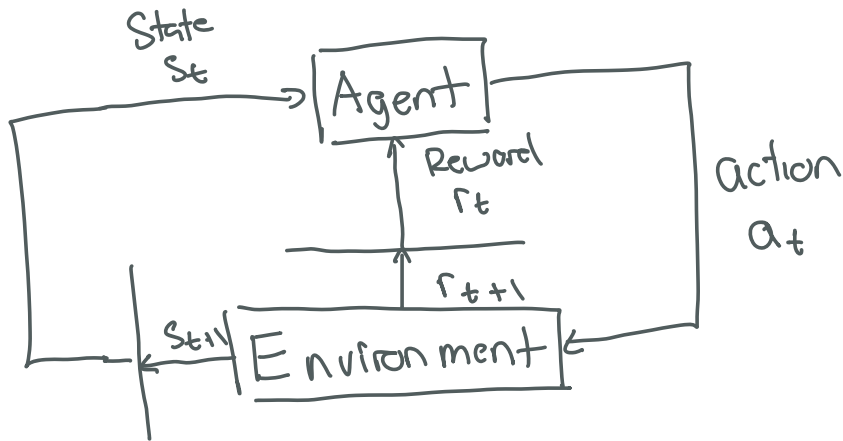
Concerns / Issues

- Weight & size
- Battery life
- * Safety
- * Obstacles
- * Terrain



How do I take
this camera frame
(and GPS coords)
to decide how
to act?

RL




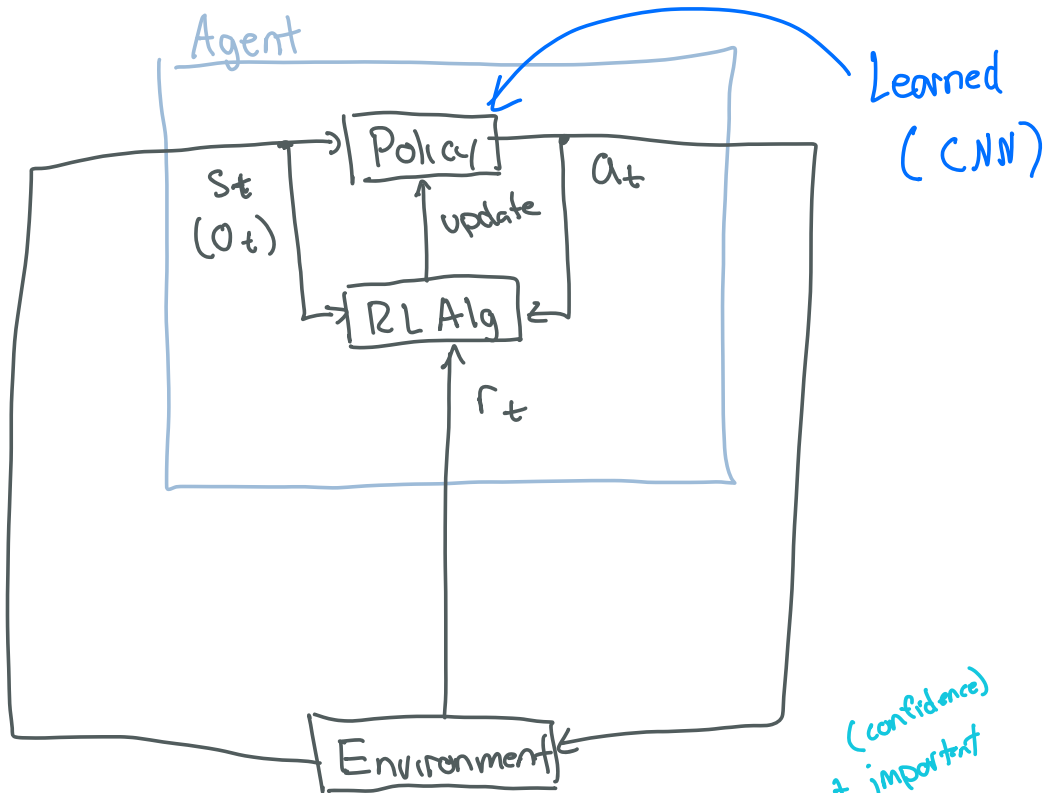
Robot

- Agent → Robot
- Env → Campus
- State → Observed parts of env
- Action → Motion
- Reward → At goal

Stock Trading

- Agent → Trader
- Env → World's Economy
- State → Stock prices
- Action → Buy, sell, hold, cancel?
- Reward → \$





- State $s \in S$
- Action $a \in A$
- Reward $r \in R$
- Discount factor $0 \leq \gamma \leq 1$
- Total reward

$$G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$$

0.85 ¹⁰⁰ ↓

