



Reinforcement Learning

What is Reinforcement?


When you do something over and over and learn from it.






What is Learning?

When information goes into your brain





How do we Combine the Two? (Reinforcement Learning)

We put information into our brain and make sure it stays there

We put information into your brain over and over again until you get it



Thank You

Questions?



Bibliography

<https://www.geeksforgeeks.org/what-is-reinforcement-learning/>

<https://deepmind.com/learning-resources/reinforcement-learning-series-2021>

https://en.wikipedia.org/wiki/Reinforcement_learning



The Interconnectedness of RL and Intelligence

- The Oxford Dictionary defines intelligence as the ability to acquire and apply knowledge and skills
- Reinforcement learning acquires knowledge as machines find the best possible behavior, which is done by learning from mistakes
- Has ties to psychology and methods of learning

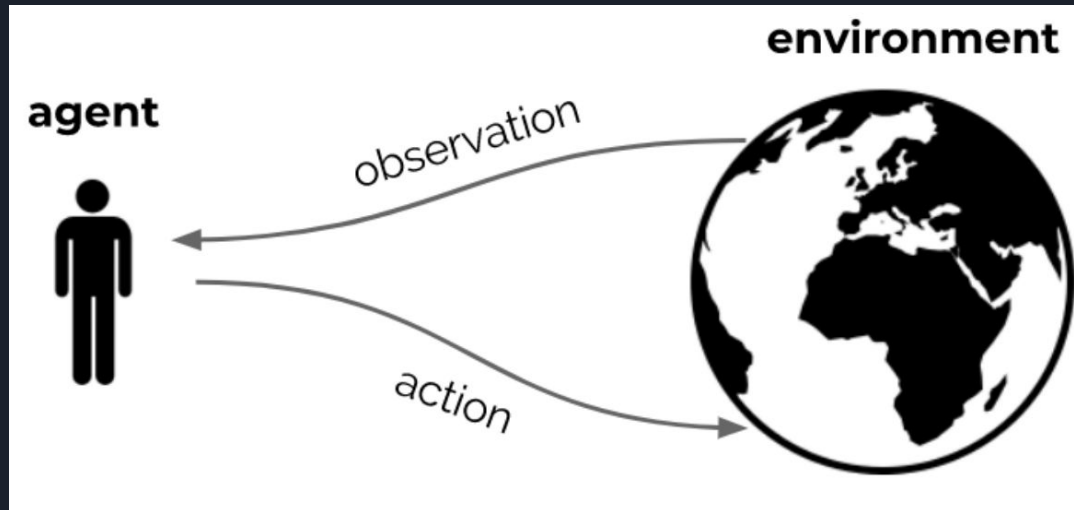
Psychology Ties

- Operant conditioning is a type of associative learning process used by many biological organisms, where the strength of a behavior is modified by reinforcement or punishment. Reinforcement learning uses the same principles of reinforcement in parent conditioning
- Positive Reinforcement is when an event follows a particular behavior, strengthening it and making sure it occurs again. I.e Giving a dog a treat after it sits
 - Maximizes performance
 - Too much can diminish results
- Negative reinforcement Strengthens behavior by removing stimuli following a behavior i.e the seat-belt sound disappearing after you put a seat belt on
 - Increases behavior
 - Helps encourage minimum behavior



What is RL in the field of CS?

“how intelligent **agents** ought to take **actions** in an **environment** in order to maximize the notion of cumulative **reward**” - Wikipedia





What is RL in the field of CS?

Based on the **reward hypothesis**:

Any goal can be formalized as the outcome of maximizing a cumulative reward

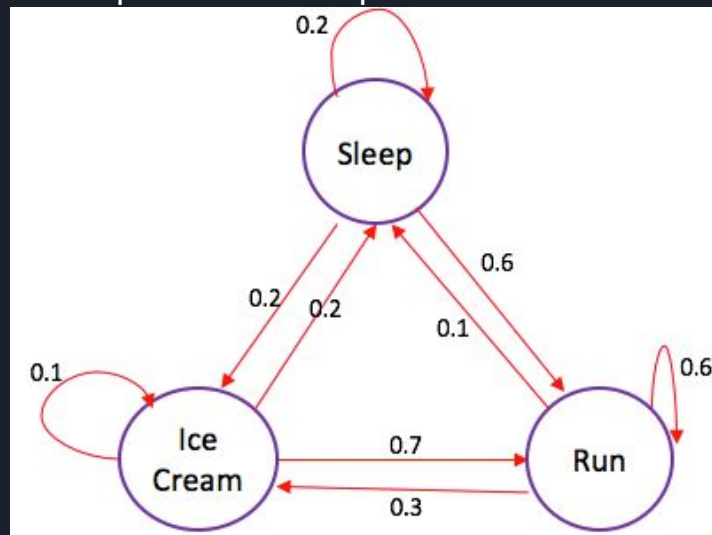
Example problems and their rewards:

- Flying a helicopter → air time, distance
- Managing an investment portfolio → gains, lack of risk
- Playing a video/board game → winning, maximizing score

What is RL in the field of CS?


Solving the problem requires finding a good balance between exploration and exploitation

- Basic model: Markov Decision Process
 - a set of environment and agent states, S ;
 - a set of actions, A , of the agent;
 - $P_a(s, s') = Pr(s_{t+1} = s' | s_t = s, a_t = a)$ is the probability of transition (at time t) from state s to state s' under action a .
 - $R_a(s, s')$ is the immediate reward after transition from s to s' with action a .



Recent Advances in RL

- DeepMind developed models to play Go, Shogi, Chess and StarCraft at the expert level (AlphaGo documentary)
- DeepMind also developed a system to fold proteins, winning them the Critical Assessment of protein Structure Prediction competition
- EquBot and Horizons launched AI platforms based in reinforcement learning to trade Exchange Traded Funds
- System to schedule traffic lights for minimal waiting time
- End to end system for controlling robot movements from raw video
- Translating texts between languages (less successful)





Project ideas + directions

- Self-driving cars
 - AWS Deepracer
- Games
 - Tetris
 - Snake
 - etc.





Libraries & Technology

- Pytorch:
 - <https://pytorch.org/>
- Tensorflow: By Google, one of the most robust frameworks
 - <https://www.tensorflow.org/>
- Keras: Simple and allows for quick prototyping
 - <https://keras.io/>



Next Steps To Learn

- Data may not be available to train complex models.
- All the possible steps the model can take to achieve the goal.
- How to build more complex training models.
- Start simple and get more complicated



Citations

[What Is Reinforcement Learning? \(simplilearn.com\)](#)

<https://medium.com/ibm-data-ai/reinforcement-learning-the-business-use-case-part-2-c175740999>

<https://celebrate150.theiet.org/media/2938/reinforcement-learning-based-multi-agent-system-for-network-traffic-signal-control.pdf>

<https://arxiv.org/pdf/1504.00702.pdf>

http://users.umiacs.umd.edu/~jbg/docs/2014_emnlp_simtrans.pdf

<https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>