




Practical Machine Learning

Reinforcement Learning




Supervised Learning



- Labeled Data Needed
- The model can not be better than the labels

Pick Up the Cup



Learning Strategies

Supervised
Learning

Unsupervised
Learning

Reinforcement
learning

Learning Strategies

Supervised
Learning


Unsupervised
Learning

Reinforcement
learning

Reinforcement Learning

*Learning from the environment what works
by getting positive feedback/reward*


Reinforcement Learning




Reinforcement Learning




Reinforcement Learning




Reinforcement Learning



Reinforcement Learning



Reinforcement Learning




$$\text{Reward } R_t = \sum_{t=0}^{\infty} \gamma^t r_t$$

γ = Discount factor


Credit Assignment Problem

Which of the action contributed to getting the reward?



Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1), 8-30.

Q – Learning




$$\text{Reward } R_t = \sum_{t=0}^{\infty} \gamma^t r_t$$

γ = Discount factor



$$\text{Q-function} \\ Q(s_t, a_t) = \mathbb{E}(R_t | s_t, a_t)$$

Deep Q Networks (DQN)



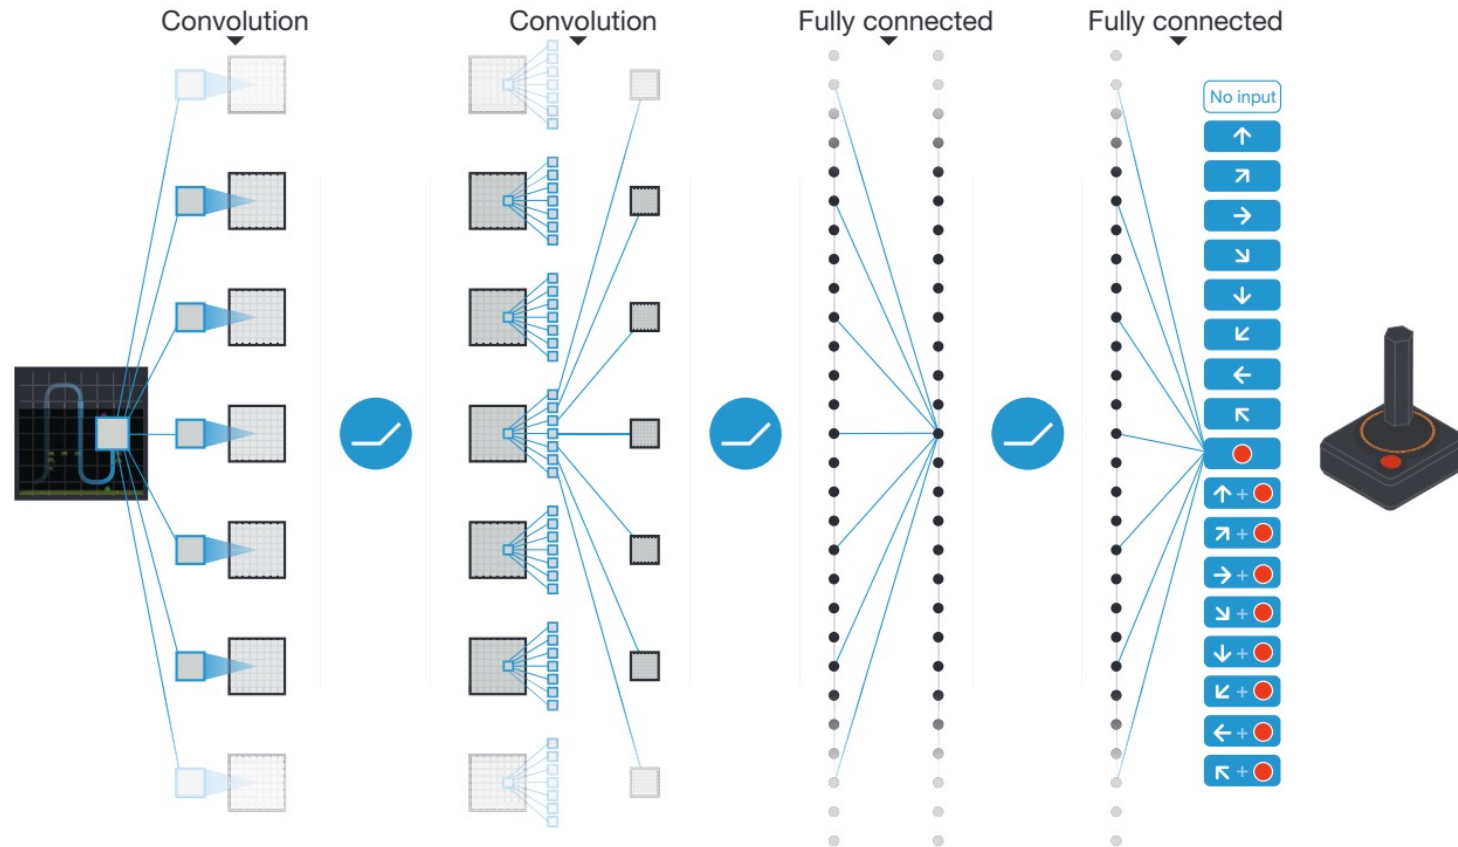
The goal is to approximate the Q-function

Deep Q Networks




Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. "Human-level control through deep reinforcement learning." *nature* 518, no. 7540 (2015): 529-533.

Deep Reinforcement Learning



Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves et al. "Human-level control through deep reinforcement learning." *nature* 518, no. 7540 (2015): 529-533.

Exercise



https://github.com/tensorflow/agents/blob/master/docs/tutorials/1_dqn_tutorial.ipynb

Conclusion

Reinforcement Learning

- What is Reinforcement Learning?
- Agent in environments
- Policy for next action
- Learning from reward

- Using Deep Neuronal Networks in RL
 - Deep Q Learning
 - Next steps: Policy Gradient
 - ... and many more approaches

License

This file is licensed under the Creative Commons Attribution-Share Alike 4.0 (CC BY-SA) license:

<https://creativecommons.org/licenses/by-sa/4.0>

Attribution: Sven Mayer

