10-601 Introduction to Machine Learning

Machine Learning Department
School of Computer Science
Carnegie Mellon University

Q-Learning

Matt Gormley
Lecture 24
Nov. 19, 2018

Reminders

- Homework 7: HMMs
  - Out: Wed, Nov 7
  - Due: Mon, Nov 19 at 11:59pm
- Homework 8: Reinforcement Learning
  - Out: Mon, Nov 19
  - Due: Fri, Dec 7 at 11:59pm

- Peer Tutoring
  - Sign up by Mon, Nov 19
- Schedule Changes
  - Recitation on Mon, Nov 26
  - Lecture on Fri, Nov 30
  - Recitation on Wed, Dec 5

VALUE ITERATION

Definitions for Value Iteration

Whiteboard

- State trajectory
- Value function
- Bellman equations
- Optimal policy
- Optimal value function
- Computing the optimal policy
- Ex: Path Planning

Example: Path Planning

Example: Robot Localization

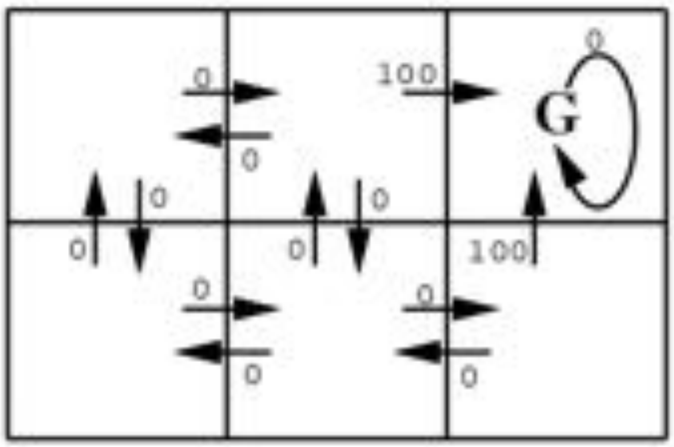

$r(s, a)$ (immediate reward) values

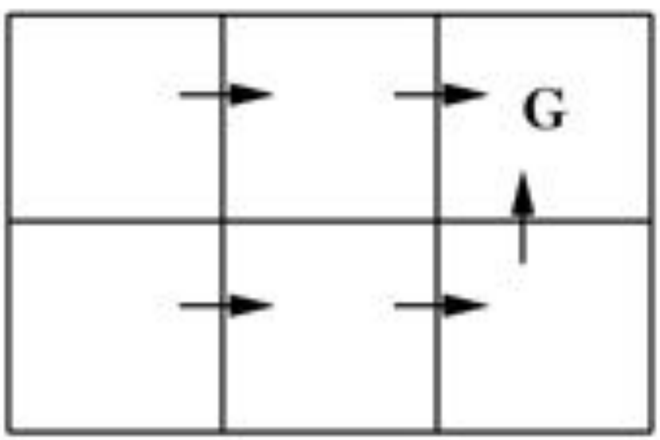

One optimal policy

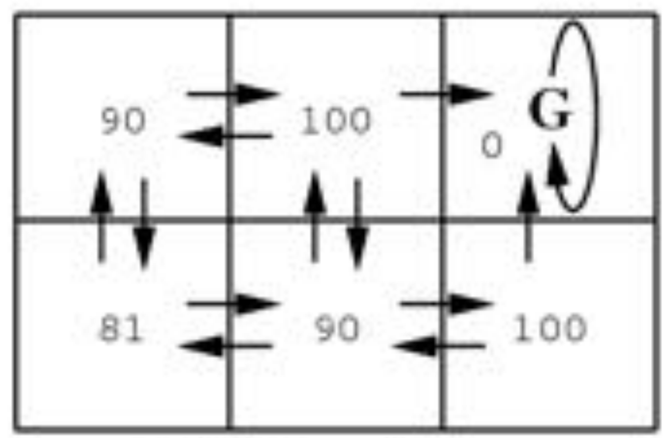

$V^*(s)$ values

Figure from Tom Mitchell

Value Iteration

Whiteboard

- Value Iteration Algorithm
- Synchronous vs. Asychronous Updates

Value Iteration

Algorithm 1 Value Iteration

procedure VALUEITERATION($R(s, a)$ reward function, $p(\cdot|s, a)$ transition probabilities)
  Initialize value function $V(s) = 0$ or randomly
  while not converged **do**
    for $s \in \mathcal{S}$ **do**
      for $a \in \mathcal{A}$ **do**
        $Q(s, a) = R(s, a) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s, a)V(s')$
      $V(s) = \max_a Q(s, a)$
  Let $\pi(s) = \text{argmax}_a Q(s, a), \ \forall s$
  return π

Value Iteration Convergence

Theorem 1 (Bertsekas (1989))
V converges to V^*, if each state is visited infinitely often

Theorem 2 (Williams & Baird (1993))
if $max_s|V^{t+1}(s) - V^t(s)| < \epsilon$
then $max_s|V^{t+1}(s) - V^*(s)| < \frac{2\epsilon\gamma}{1-\gamma}, \forall s$

Theorem 3 (Bertsekas (1987))
greedy policy will be optimal in a finite number of steps (even if not converged to optimal value function!)

Holds for both asynchronous and sychronous updates

Provides reasonable stopping criterion for value iteration

Often greedy policy converges well before the value function

POLICY ITERATION

Policy Iteration

Whiteboard

- Policy Iteration Algorithm
- Solving the Bellman Equations for Fixed Policy
- Convergence Properties
- Value Iteration vs. Policy Iteration

Policy Iteration

Algorithm 1 Policy Iteration

procedure POLICYITERATION($R(s,a)$ reward function, $p(\cdot|s,a)$ transition probabilities)

  Initialize policy π randomly

  while not converged **do**

    Solve Bellman equations for fixed policy π

$$V^\pi(s) = R(s,\pi(s)) + \gamma \sum_{s'\in\mathcal{S}} p(s'|s,\pi(s))V^\pi(s'),\ \forall s$$

    Improve policy π using new value function

$$\pi(s) = \operatorname*{argmax}_a R(s,a) + \gamma \sum_{s'\in\mathcal{S}} p(s'|s,a)V^\pi(s')$$

  return π

Policy Iteration

Algorithm 1 Policy Iteration

procedure POLICYITERATION($R(s, a$ … n, $p(\cdot|s, a)$ transition probabilities)

Initialize policy π randomly

while not converged **do**

Solve Bellman equations for fixed policy π

$$V^\pi(s) = R(s, \pi(s)) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s, \pi(s)) V^\pi(s'), \ \forall s$$

Improve policy π using new value function

$$\pi(s) = \operatorname*{argmax}_a R(s, a) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s, a) V^\pi(s')$$

return π

Compute value function for fixed policy is easy

System of |S| equations and |S| variables

Greedy policy w.r.t. current value function

Greedy policy might **remain the same** for a particular state if there is no better action

Policy Iteration Convergence

In-Class Exercise:

How many policies are there for a finite sized state and action space?

In-Class Exercise:

Suppose policy iteration is shown to improve the policy at every iteration. Can you bound the number of iterations it will take to converge?

Value Iteration vs. Policy Iteration

- Value iteration requires $O(|A|\ |S|^2)$ computation per iteration

- Policy iteration requires $O(|A|\ |S|^2 + |S|^3)$ computation per iteration
- In practice, policy iteration converges in fewer iterations

Algorithm 1 Value Iteration

procedure ValueIteration($R(s,a)$ reward function, $p(\cdot|s,a)$ transition probabilities)
    Initialize value function $V(s) = 0$ or randomly
    while not converged **do**
        for $s \in \mathcal{S}$ **do**
            for $a \in \mathcal{A}$ **do**
                $Q(s,a) = R(s,a) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s,a) V(s')$
            $V(s) = \max_a Q(s,a)$
    Let $\pi(s) = \operatorname{argmax}_a Q(s,a),\ \forall s$
    return π

Algorithm 1 Policy Iteration

procedure PolicyIteration($R(s,a)$ reward function, $p(\cdot|s,a)$ transition probabilities)
    Initialize policy π randomly
    while not converged **do**
        Solve Bellman equations for fixed policy π

$$V^\pi(s) = R(s,\pi(s)) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s,\pi(s)) V^\pi(s'),\ \forall s$$

        Improve policy π using new value function

$$\pi(s) = \operatorname*{argmax}_a R(s,a) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s,a) V^\pi(s')$$

    return π

Learning Objectives

Reinforcement Learning: Value and Policy Iteration

You should be able to…

1. Compare the reinforcement learning paradigm to other learning paradigms
2. Cast a real-world problem as a Markov Decision Process
3. Depict the exploration vs. exploitation tradeoff via MDP examples
4. Explain how to solve a system of equations using fixed point iteration
5. Define the Bellman Equations
6. Show how to compute the optimal policy in terms of the optimal value function
7. Explain the relationship between a value function mapping states to expected rewards and a value function mapping state-action pairs to expected rewards
8. Implement value iteration
9. Implement policy iteration
10. Contrast the computational complexity and empirical convergence of value iteration vs. policy iteration
11. Identify the conditions under which the value iteration algorithm will converge to the true value function
12. Describe properties of the policy iteration algorithm

Q-LEARNING

Q-Learning

Whiteboard

- Motivation: What if we have zero knowledge of the environment?
- Q-Function: Expected Discounted Reward

Example: Robot Localization

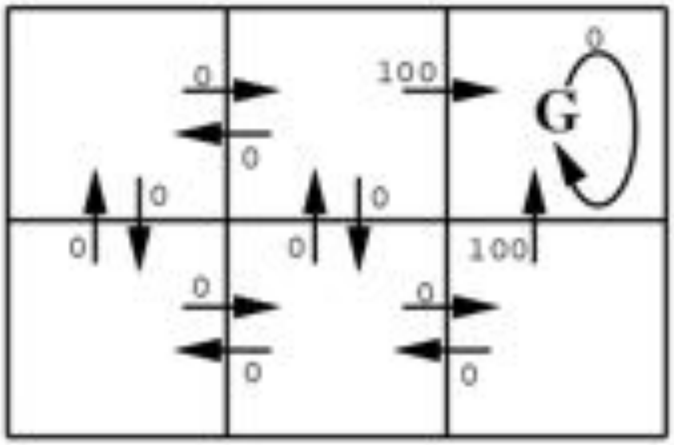

$r(s, a)$ (immediate reward) values

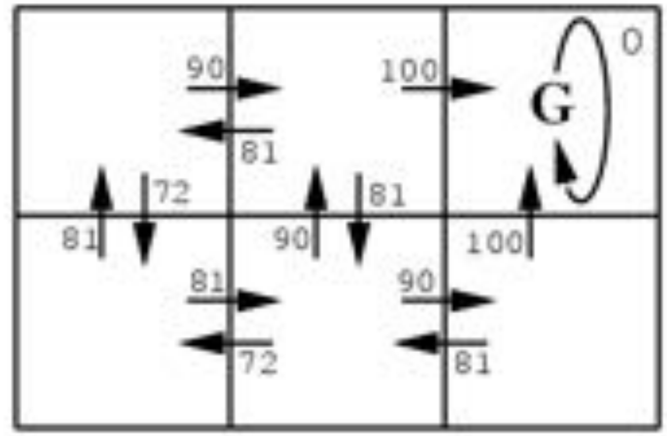

$Q(s, a)$ values

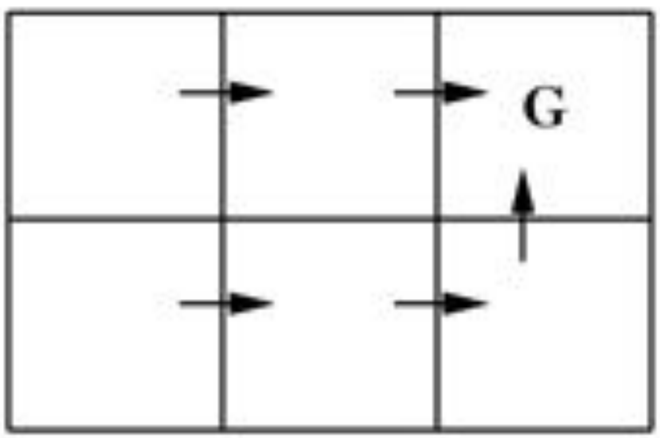

One optimal policy

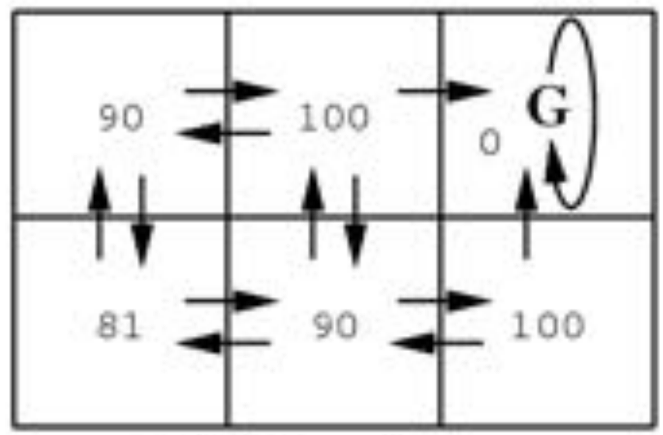

$V^*(s)$ values

Figure from Tom Mitchell

Q-Learning

Whiteboard

- Q-Learning Algorithm
  - Case 1: Deterministic Environment
  - Case 2: Nondeterministic Environment
- Convergence Properties
- Exploration Insensitivity
- Ex: Re-ordering Experiences
- ϵ-greedy Strategy

DEEP RL EXAMPLES

TD Gammon → Alpha Go

Learning to beat the masters at board games

THEN	NOW
"…the world's top computer program for backgammon, TD-GAMMON (Tesauro, 1992, 1995), learned its strategy by playing over one million practice games against itself…" (Mitchell, 1997)	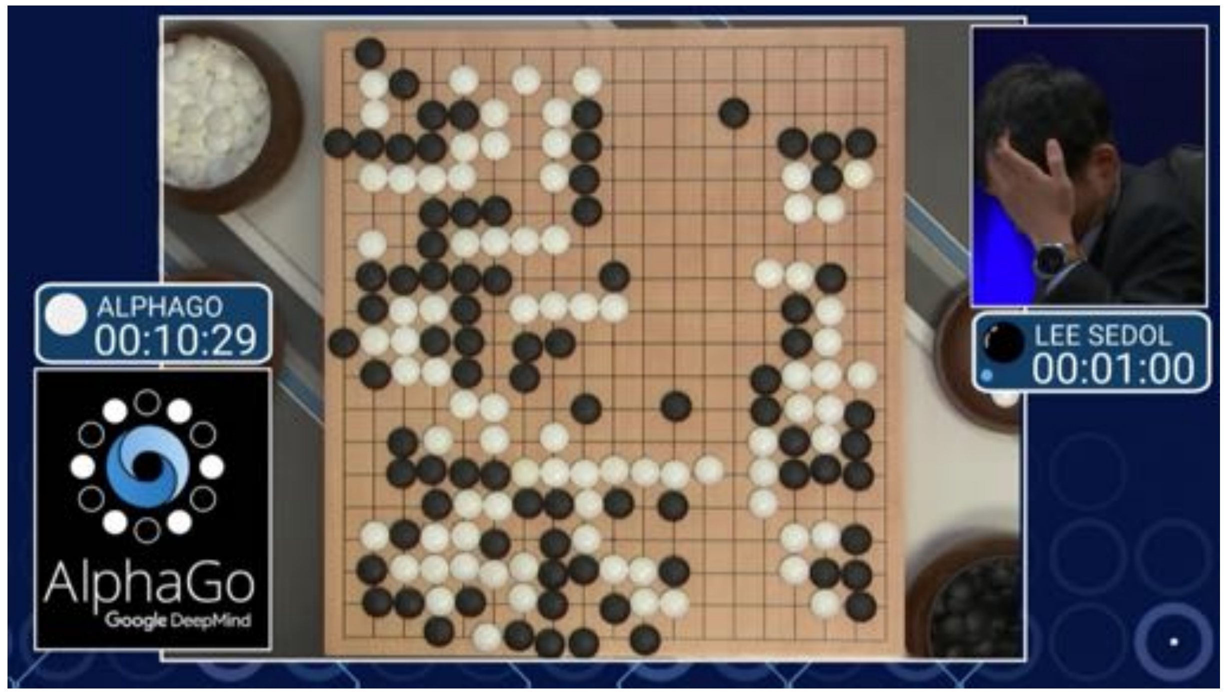

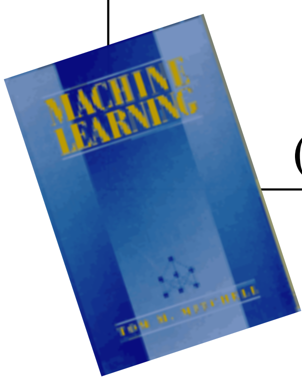

Playing Atari with Deep RL

- Setup: RL system observes the pixels on the screen
- It receives rewards as the game score
- Actions decide how to move the joystick / buttons

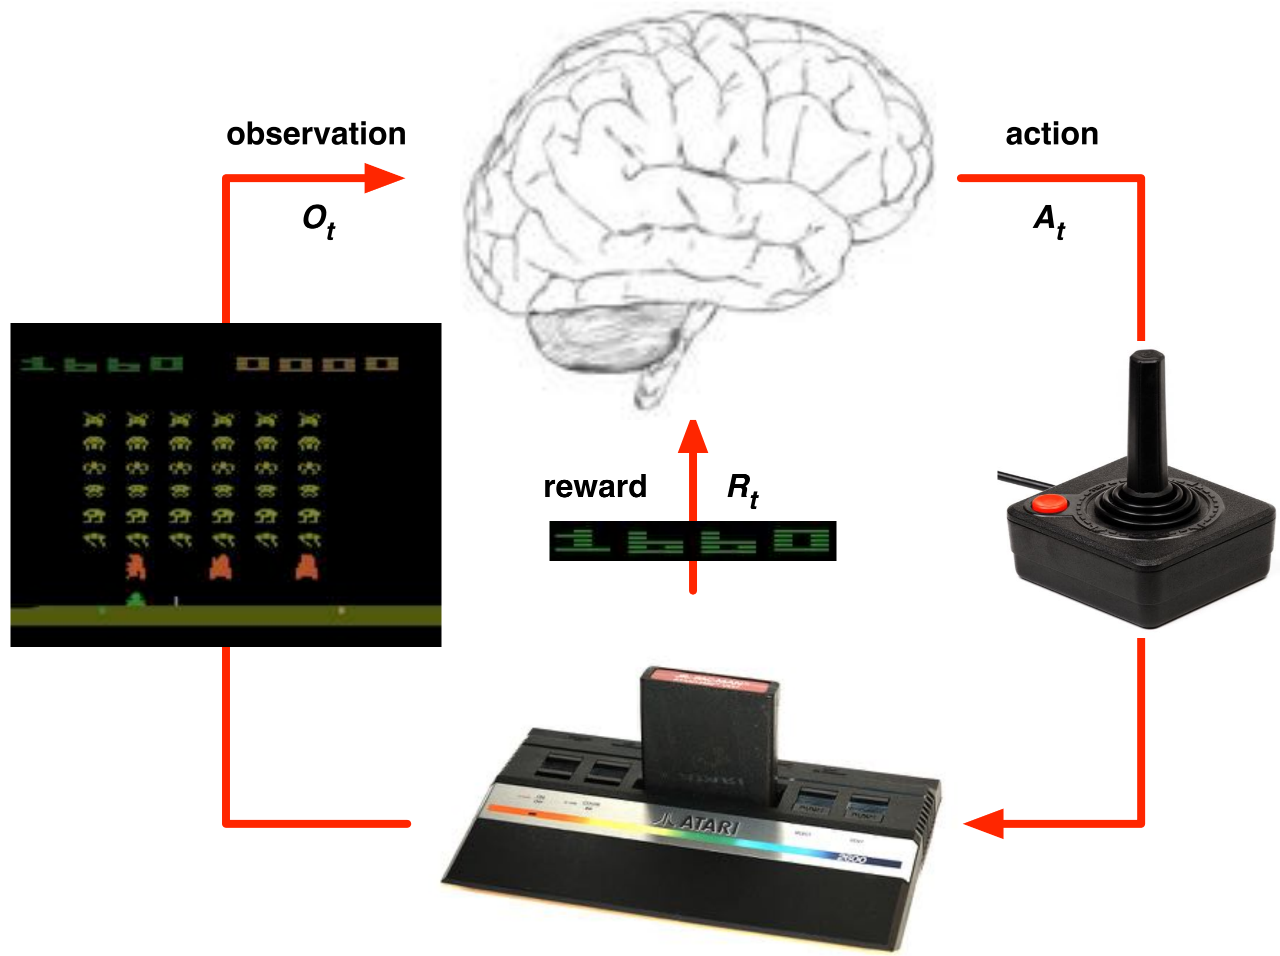

Figures from David Silver (Intro RL lecture)

Playing Atari with Deep RL

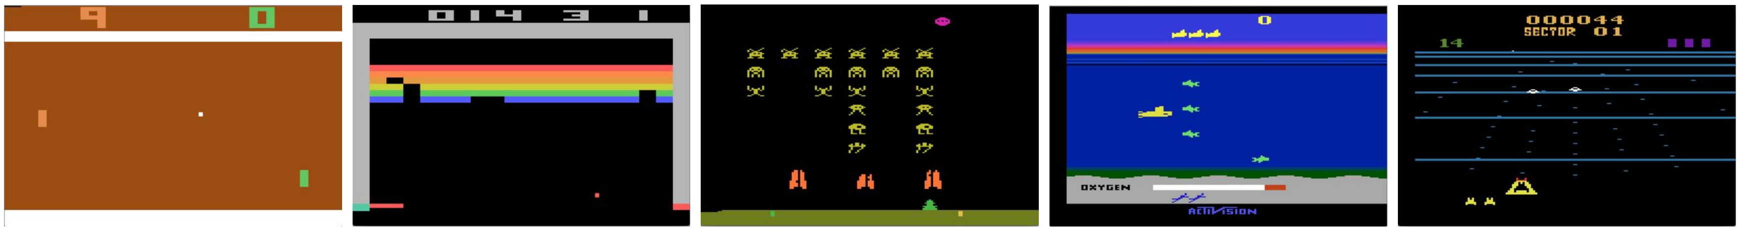

Figure 1: Screen shots from five Atari 2600 Games: (*Left-to-right*) Pong, Breakout, Space Invaders, Seaquest, Beam Rider

Videos:

- Atari Breakout: https://www.youtube.com/watch?v=V1eYniJ0Rnk
- Space Invaders: https://www.youtube.com/watch?v=ePv0Fs9cGgU

Figures from Mnih et al. (2013)

Playing Atari with Deep RL

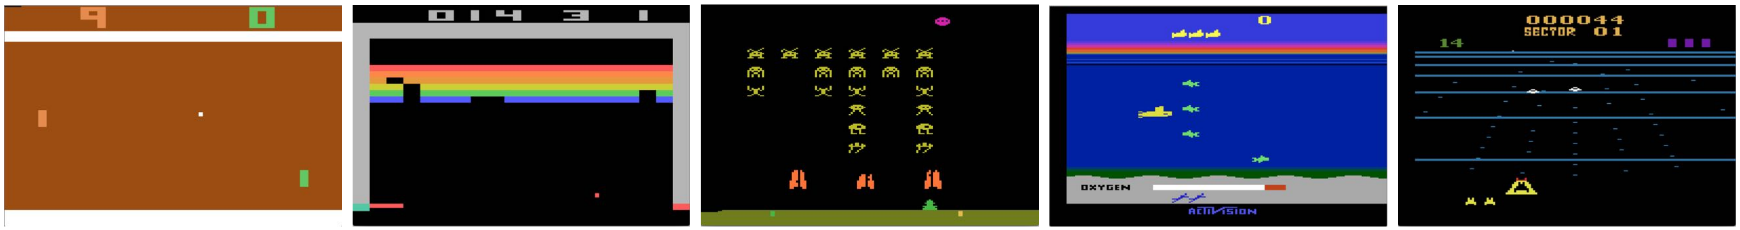

Figure 1: Screen shots from five Atari 2600 Games: (*Left-to-right*) Pong, Breakout, Space Invaders, Seaquest, Beam Rider

	B. Rider	Breakout	Enduro	Pong	Q*bert	Seaquest	S. Invaders
Random	354	1.2	0	−20.4	157	110	179
Sarsa [3]	996	5.2	129	−19	614	665	271
Contingency [4]	1743	6	159	−17	960	723	268
DQN	**4092**	**168**	**470**	**20**	**1952**	**1705**	**581**
Human	7456	31	368	−3	18900	28010	3690
HNeat Best [8]	3616	52	106	19	1800	920	**1720**
HNeat Pixel [8]	1332	4	91	−16	1325	800	1145
DQN Best	**5184**	**225**	**661**	**21**	**4500**	**1740**	1075

Table 1: The upper table compares average total reward for various learning methods by running an ϵ-greedy policy with $\epsilon = 0.05$ for a fixed number of steps. The lower table reports results of the single best performing episode for HNeat and DQN. HNeat produces deterministic policies that always get the same score while DQN used an ϵ-greedy policy with $\epsilon = 0.05$.

Figures from Mnih et al. (2013)

Q-Learning

Whiteboard

- Approximating the Q function with a neural network
- Deep Q-Learning
- Experience Replay

Alpha Go

Game of Go (圍棋)

- 19x19 **board**
- Players alternately play black/white **stones**
- **Goal** is to fully encircle the largest region on the board
- **Simple** rules, but **extremely complex** game play

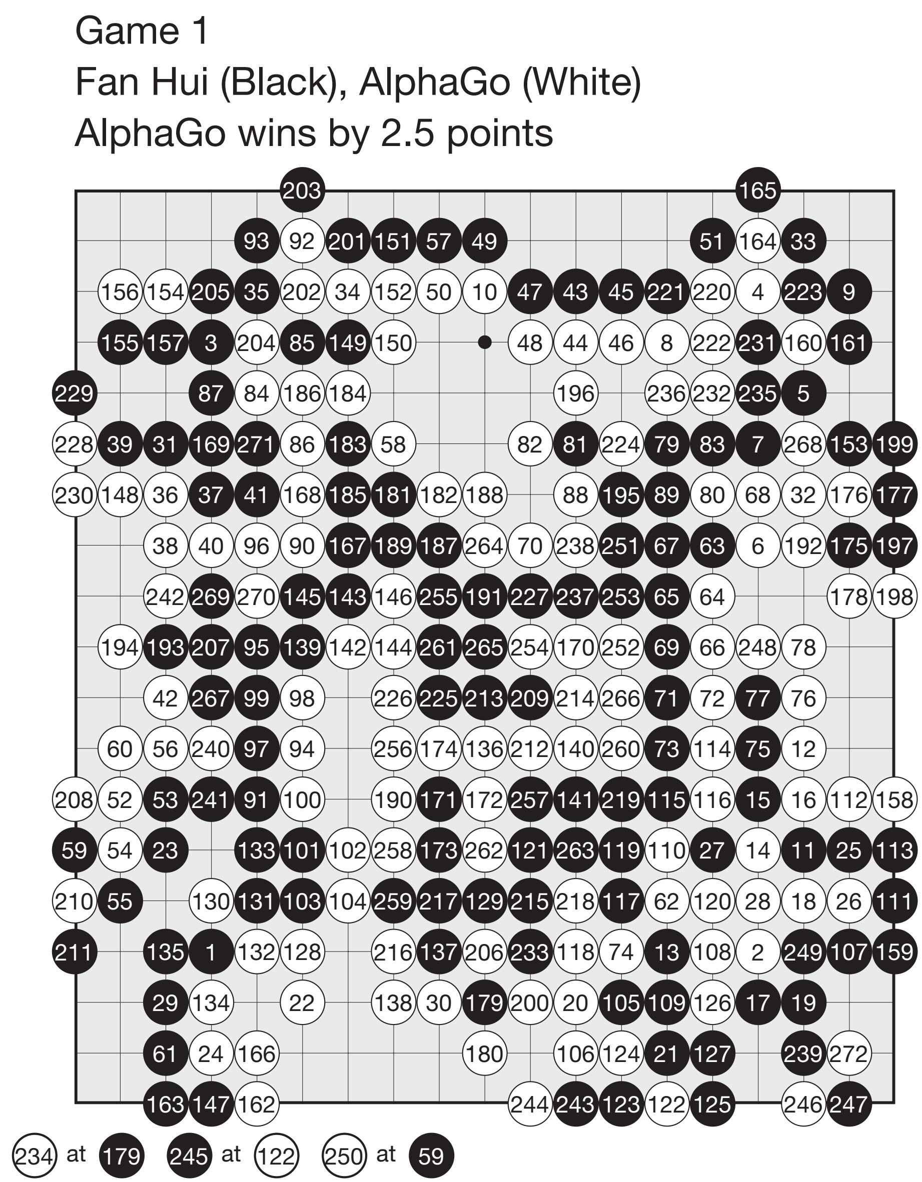

Figure from Silver et al. (2016)

Alpha Go

- State space is too large to represent explicitly since # of sequences of moves is $O(b^d)$
  - Go: $b=250$ and $d=150$
  - Chess: $b=35$ and $d=80$
- Key idea:
  - Define a neural network to approximate the value function
  - Train by policy gradient

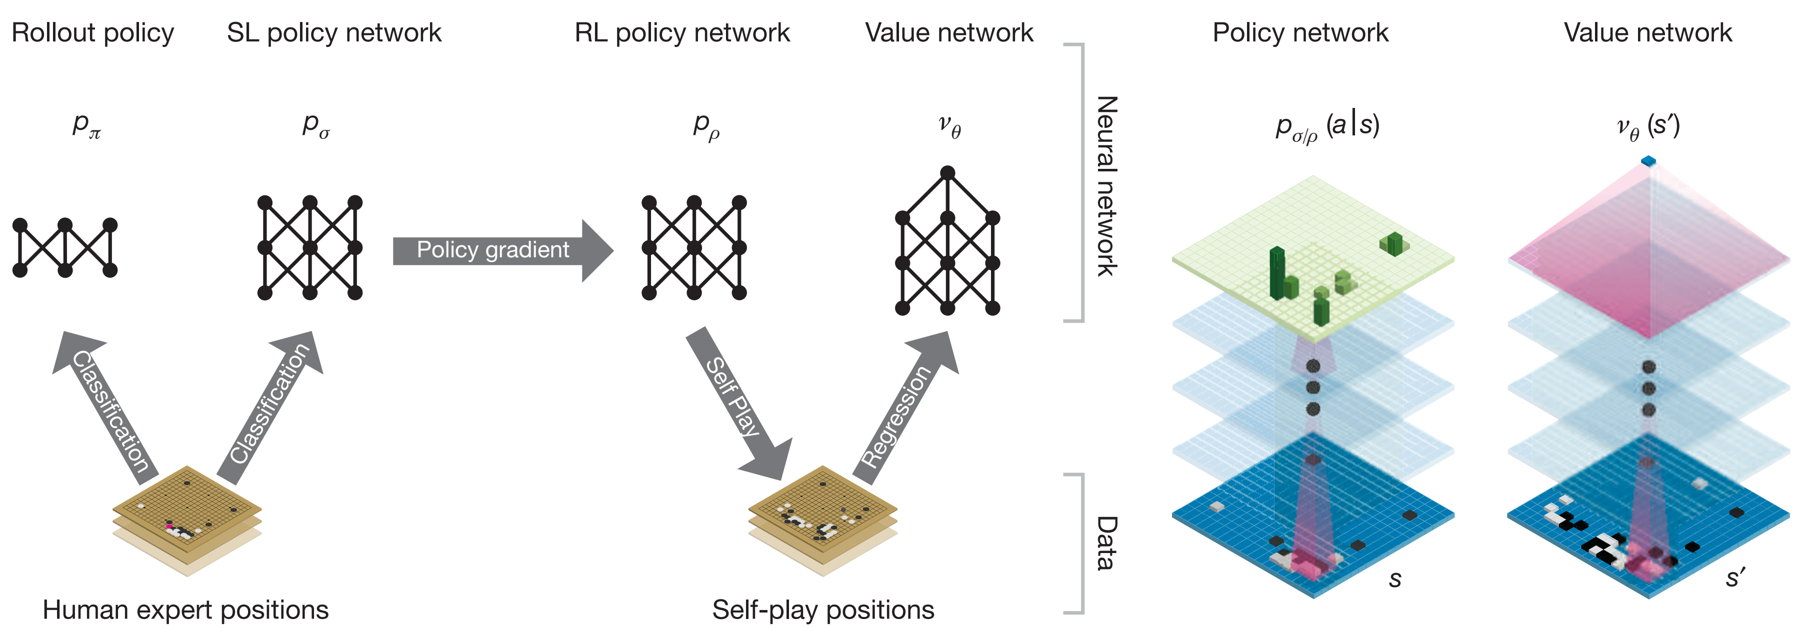

Figure from Silver et al. (2016)

Alpha Go

- Results of a tournament
- From Silver et al. (2016): “a 230 point gap corresponds to a 79% probability of winning”

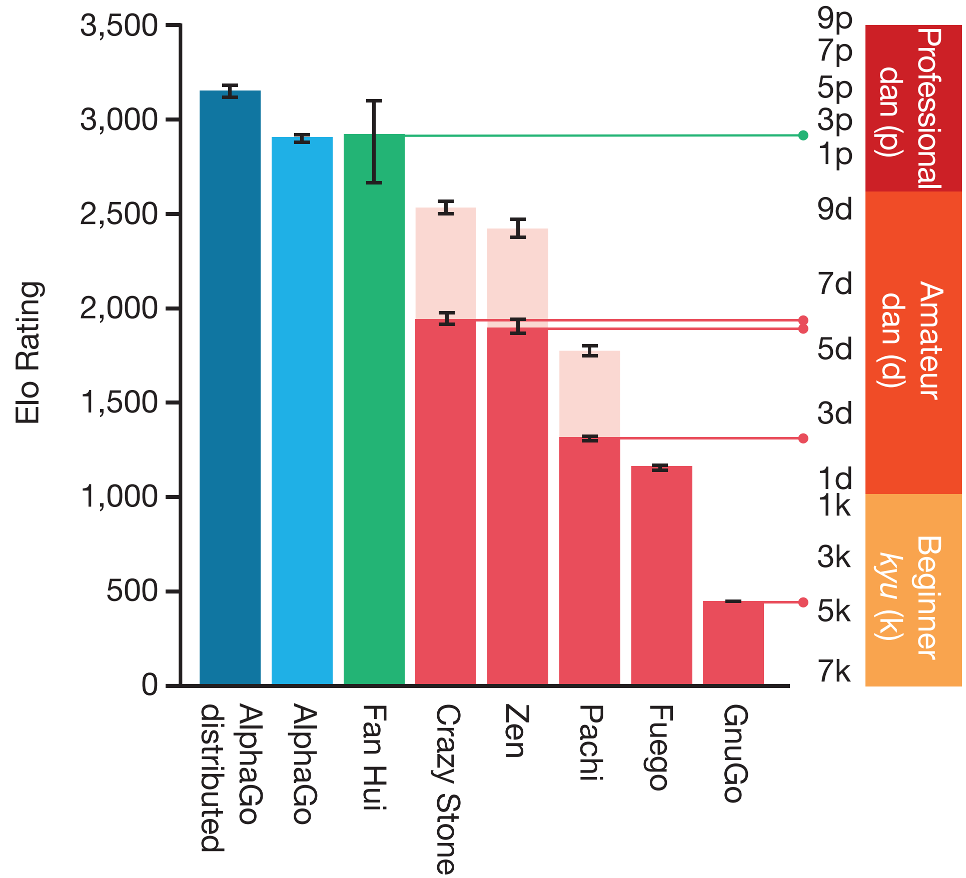

Figure from Silver et al. (2016)

Learning Objectives

Reinforcement Learning: Q-Learning

You should be able to…

1. Apply Q-Learning to a real-world environment
2. Implement Q-learning
3. Identify the conditions under which the Q-learning algorithm will converge to the true value function
4. Adapt Q-learning to Deep Q-learning by employing a neural network approximation to the Q function
5. Describe the connection between Deep Q-Learning and regression