Machine Learning

Machine Learning: Jordan Boyd-Graber
University of Maryland
POLICY METHODS

Adapted from slides by David Silver, Pieter Abbeel, and John Schulman

- RL used to be niche subfield ...
- Now it's all over the place
- Part of much of ML hype
- But what is reinforcement learning?

- RL used to be niche subfield . . .
- Now it's all over the place
- Part of much of ML hype
- But what is reinforcement learning?
  - RL is a general-purpose framework for decision-making
  - RL is for an agent with the capacity to act
  - Each action influences the agent's future state
  - Success is measured by a scalar reward signal
  - Goal: select actions to maximise future reward

Approaches to RL

Value-based RL

- Estimate the optimal value function $Q(s,a)$
- This is the maximum value achievable under any policy

Policy-based RL

- Search directly for the optimal policy π^*
- This is the policy achieving maximum future reward

Model-based RL

- Build a model of the environment
- Plan (e.g. by lookahead) using model

Deep Q Learning

- Optimal Q-values should obey equation

$$Q^*(s,a) = \mathbb{E}_{s'}\left[r + \gamma Q(s',a') \,|\, s,a\right] \tag{1}$$

- Treat as regression problem
- Minimize: $\left(r + \gamma \max_a Q(s',a',\vec{w}) - Q(s,a,\vec{w})\right)^2$
- Converges to Q using table lookup representation
- But diverges using neural networks due to:
  - Correlations between samples
  - Non-stationary targets

Deep RL in Atari

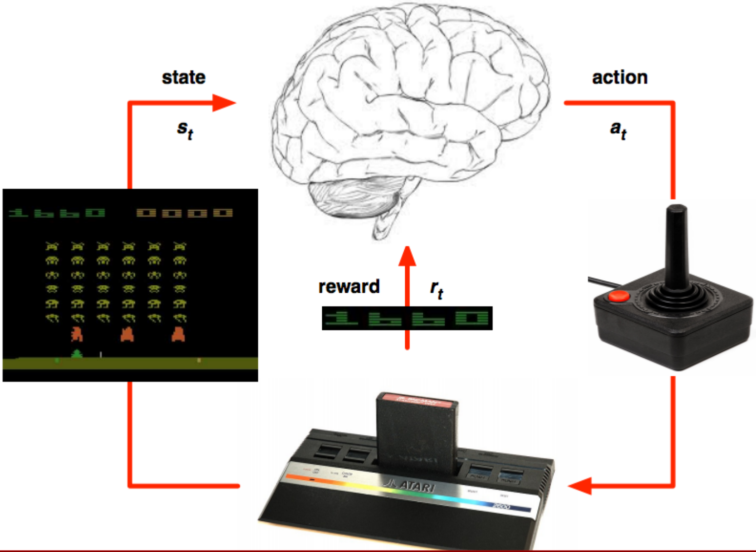

DQN in Atari

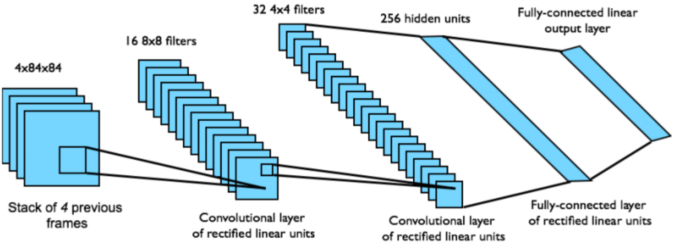

- End-to-end learning of values $Q(s,a)$ from pixels s
- Input state s is stack of raw pixels from last four frames
- Output is $Q(s,a)$ for 18 joystick/button positions
- Reward is change in score for that step

Atari Results

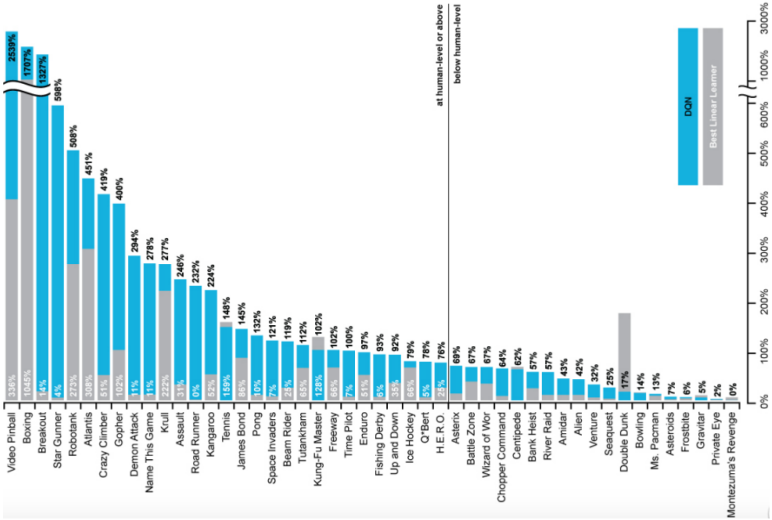

Policy-Based RL

- Advantages:
  - Better convergence properties
  - Effective in high-dimensional or continuous action spaces
  - Can learn stochastic policies
- Disadvantages:
  - Typically converge to a local rather than global optimum
  - Evaluating a policy is typically inefficient and high variance

Optimal Policies Sometimes Stochastic

Optimal Policies Sometimes Stochastic

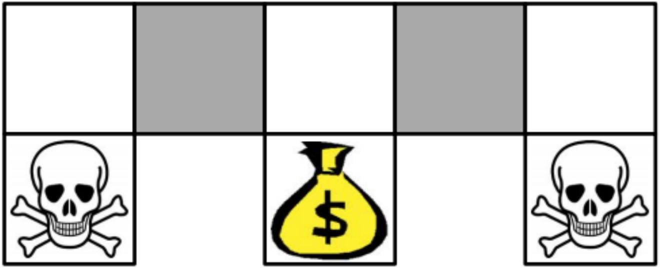

(Cannot distinguish gray states)

Optimal Policies Sometimes Stochastic

Deterministic

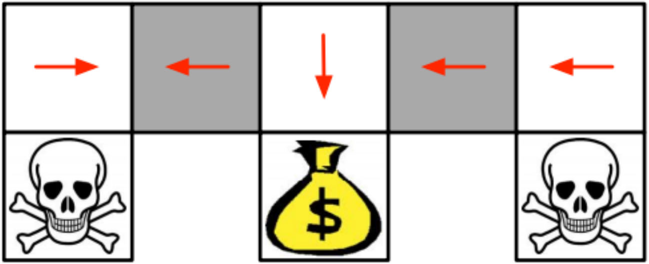

(Cannot distinguish gray states)

Optimal Policies Sometimes Stochastic

Deterministic

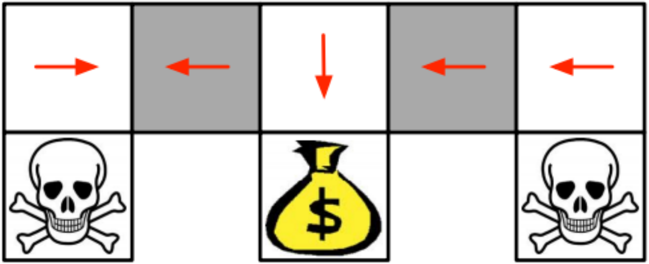

(Cannot distinguish gray states)

Value-based RL learns near deterministic policy!

Optimal Policies Sometimes Stochastic

Stochastic

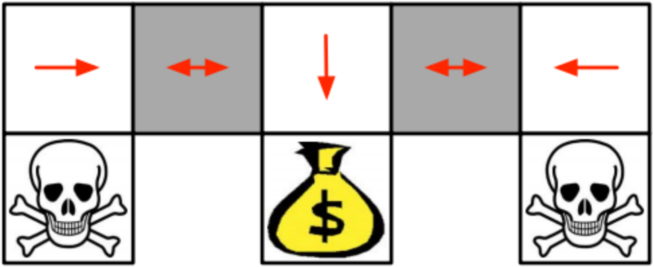

(Cannot distinguish gray states, so flip a coin!)

Likelihood Ratio Policy Gradient

Let τ be state-action $s_0, u_0, \dots, s_H, u_H$. Utility of policy π parametrized by θ is

$$U(\theta) = \mathbb{E}_{\pi_\theta, U}\left[\sum_t^H R(s_t, u_t); \pi_\theta\right] = \sum_\tau P(\tau;\theta) R(\tau). \quad (2)$$

Our goal is to find θ:

$$\max_\theta U(\theta) = \max_\theta \sum_\tau p(\tau;\theta) R(\tau) \quad (3)$$

Likelihood Ratio Policy Gradient

$$\sum_t p(\tau;\theta)R(\tau) \tag{4}$$

Taking the gradient wrt θ:

$$\tag{5}$$

Likelihood Ratio Policy Gradient

$$\sum_t p(\tau;\theta)R(\tau) \tag{4}$$

Taking the gradient wrt θ:

$$\nabla_\theta U(\theta) = \sum_\tau R(\tau)\frac{P(\tau;\theta)}{P(\tau;\theta)}\nabla_\theta P(\tau;\theta) \tag{5}$$

$$\tag{6}$$

Move differentiation inside sum (ignore $R(\tau)$ and then add in term that cancels out

Likelihood Ratio Policy Gradient

$$\sum_t p(\tau;\theta)R(\tau) \tag{4}$$

Taking the gradient wrt θ:

$$\nabla_\theta U(\theta) = \sum_\tau R(\tau)\frac{P(\tau;\theta)}{P(\tau;\theta)}\nabla_\theta P(\tau;\theta) \tag{5}$$

$$= \sum_\tau P(\tau;\theta)\frac{\nabla_\theta P(\tau;\theta)}{P(\tau;\theta)}R(\tau) \tag{6}$$

$$\tag{7}$$

Move derivative over probability

Likelihood Ratio Policy Gradient

$$\sum_t p(\tau;\theta)R(\tau) \tag{4}$$

Taking the gradient wrt θ:

$$\nabla_\theta U(\theta) = \sum_\tau R(\tau)\frac{P(\tau;\theta)}{P(\tau;\theta)}\nabla_\theta P(\tau;\theta) \tag{5}$$

$$= \sum_\tau P(\tau;\theta)\frac{\nabla_\theta P(\tau;\theta)}{P(\tau;\theta)}R(\tau) \tag{6}$$

$$= \sum_\tau P(\tau;\theta)\nabla_\theta \left[\log P(\tau;\theta)\right] R(\tau) \tag{7}$$

Assume softmax form ($\nabla_\theta \log z = \frac{1}{z}\nabla_\theta z$)

Likelihood Ratio Policy Gradient

$$\sum_{t} p(\tau;\theta)R(\tau) \tag{4}$$

Taking the gradient wrt θ:

$$= \sum_{\tau} P(\tau;\theta)\nabla_{\theta}\left[\log P(\tau;\theta)\right]R(\tau) \tag{5}$$

Approximate with empirical estimate for m sample paths from π

$$\nabla_{\theta} U(\theta) \approx \frac{1}{m}\sum_{i}^{m} \nabla_{\theta} \log P(r^{i};\theta)R(\tau^{i}) \tag{6}$$

Policy Gradient Intuition

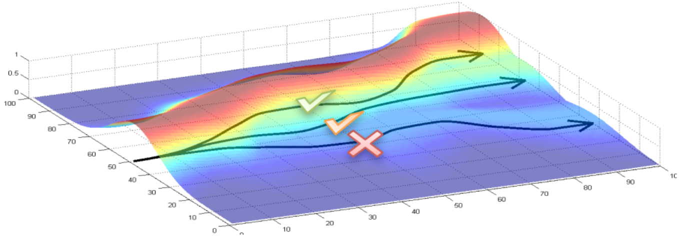

- Increase probability of paths with positive R
- Decrease probability of paths with negagive R

Extensions

- Consider baseline b (e.g., path averaging)

$$\nabla_\theta U(\theta) \approx \frac{1}{m} \sum_{1}^{m} \nabla_\theta \log P(r^i; \theta)(R(\tau^i) - b(\tau)) \tag{7}$$

- Combine with value estimation (critic)
  - Critic: Updates action-value function parameters
  - Actor: Updates policy parameters in direction suggested by critic

Recap

- Reinforcement learning is active subfield of ML
- Deep learning option for learning policy / value functions
- Representation learning helps cope with large state spaces
- Still requires careful engineering and feature engineering