UCI University of California, Irvine

CS 277: Control and Reinforcement Learning

Winter 2022

Lecture 7: Exploration

Roy Fox

Department of Computer Science

Bren School of Information and Computer Sciences

University of California, Irvine

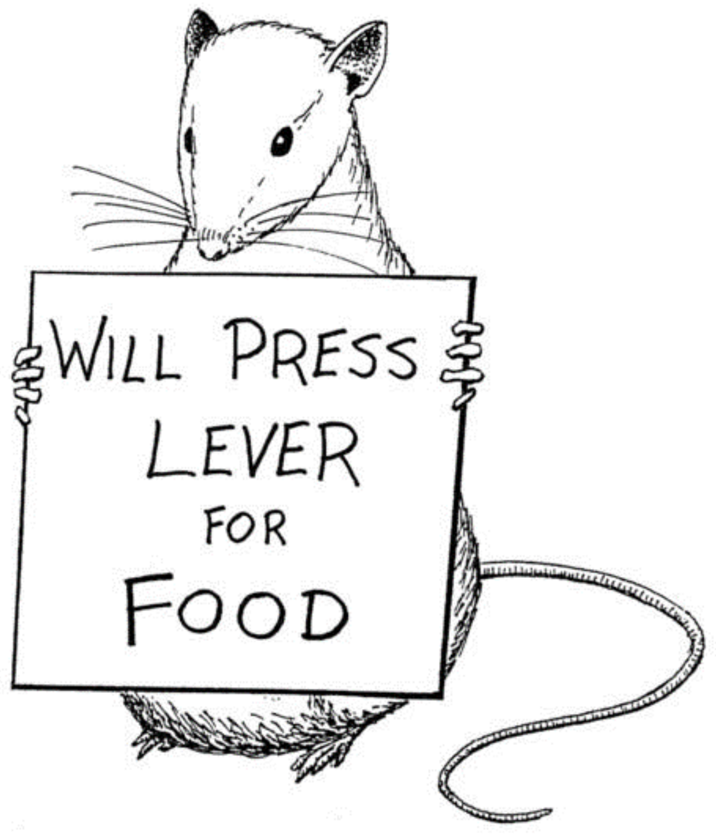

Logistics

assignments

- Assignment 2 due next Tuesday

quizzes

- Quiz 4 to be published soon

Today's lecture

Multi-Armed Bandits

Exploration in Deep RL

Multi-Armed Bandits (MABs)

- Basic setting: single instance x, multiple actions $a_1, \ldots, a_k$
  - Each time we take action a_i we see a noisy reward $r_t \sim p_i$
- Can we maximize the expected reward $\max_i \mathbb{E}_{r \sim p_i}[r]$?
  - We can use the mean as an estimate $\mu_i = \mathbb{E}_{r \sim p_i}[r] \approx \frac{1}{n(i)} \sum_{t \in \mathcal{T}_i} r_t$
- Challenge: is the best mean so far the best action?
  - Or is there another that's better than it appeared so far?

One-armed bandit

Multi-armed bandit

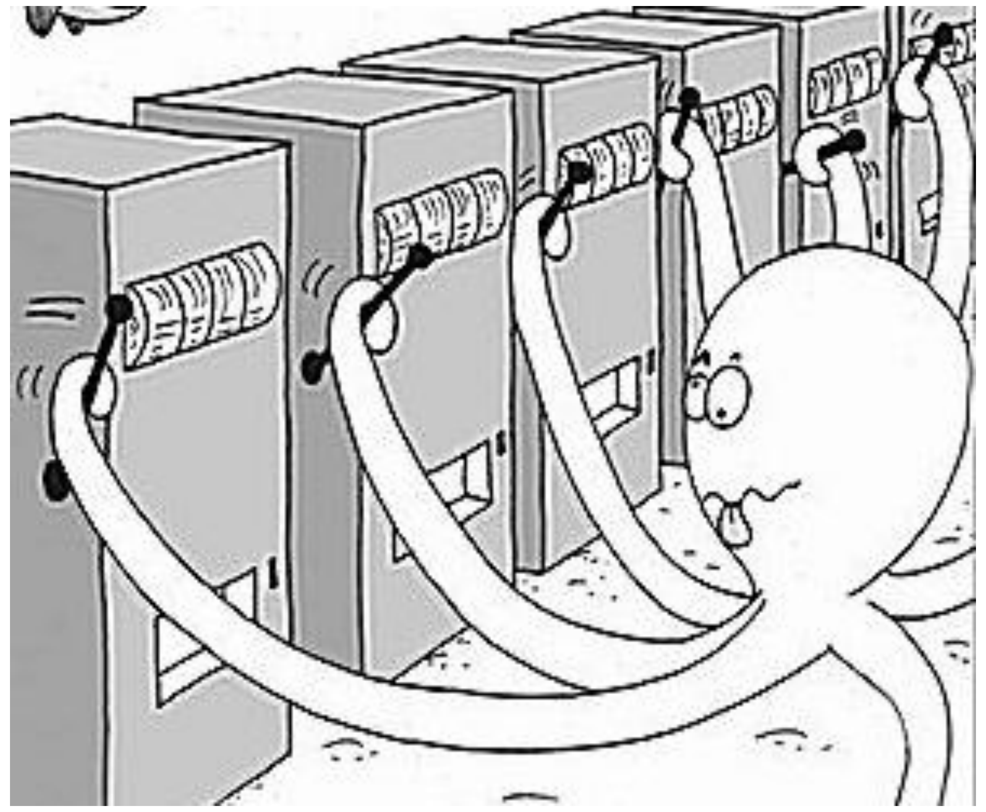

Exploration vs. exploitation

- Exploitation = choose actions that seems good (so far)
- Exploration = see if we're missing out on even better ones
- Naïve solution: learn r by trying every action enough times
  - Suppose we can't wait that long: we care about rewards while we learn
- Regret = how much worse our return is than an optimal action

$$\rho(T) = T\mu_{a^*} - \sum_{t=0}^{T-1} r_t$$

  - Can we get the regret to grow sub-linearly with T? $\implies$ average goes to 0: $\frac{\rho(T)}{T} \to 0$

Let's play!

- http://iosband.github.io/2015/07/28/Beat-the-bandit.html

Simple exploration: ϵ-greedy

- With probability ϵ:
  - Select action uniformly at random
- Otherwise (w.p. $1 - \epsilon$):
  - Select best (on average) action so far
- Problem 1: all non-greedy actions selected with same probability
- Problem 2: must have $\epsilon \to 0$, or we keep accumulating regret
  - But at what rate should ϵ vanish?

Boltzmann exploration

- Keep an average of past rewards $\hat{\mu}_i = \frac{1}{n(i)} \sum_{t \in \mathcal{T}_i} r_t$
- Boltzmann (softmax) exploration: $\pi(a_i) = \operatorname{softmax}_\beta \hat{\mu}_i = \frac{\exp(\beta \hat{\mu}_i)}{\sum_j \exp(\beta \hat{\mu}_j)}$
- Obviously bad actions $\hat{\mu}_i \ll \max_j \hat{\mu}_j$ are unlikely to be used (but can!)
  - Problem: still must have $\beta \to \infty$, or we keep accumulating regret

Optimism under uncertainty

- Tradeoff: explore less used actions, but don't be late to start exploiting what's known
  - Principle: optimism under uncertainty = explore to the extent you're uncertain, otherwise exploit
- By the central limit theorem, the mean reward $\hat{\mu}_i$ of arm i quickly $\to \mathcal{N}\left(\mu_i, O\left(\frac{1}{n(i)}\right)\right)$
- Be optimistic by slowly-growing number of standard deviations:

$$a = \arg\max_i \hat{\mu}_i + \sqrt{\frac{2\ln T}{n(i)}}$$

  - Upper confidence bound (UCB): likely $\mu_i \leq \hat{\mu}_i + c\sigma_i$; unknown variance $\Longrightarrow$ let c grow
  - But not too fast, or we fail to exploit what we do know
- Regret: $\rho(T) = O(\log T)$, provably optimal

Thompson sampling

- Consider a model of the reward distribution $p_{\theta_i}(r \,|\, a_i)$
- Suppose we start with some prior $q(\theta)$
  - Taking action a_t, see reward $r_t \implies$ update posterior $q(\theta \,|\, \{(a_{\leq t}, r_{\leq t})\})$
- Thompson sampling:
  - Sample $\theta \sim q$ from the posterior
  - Take the optimal action $a^* = \max_i \mathbb{E}_{r \sim p_{\theta i}}[r]$
  - Update the belief (different methods for doing this)
  - Repeat

Other online learning settings

- What is the reward for action a_i?
  - MAB: random variable with distribution $p_i(r)$
  - Adversarial bandits: adversary selects r_i for every action
    - The adversary knows our algorithm! And past action selection! But not future actions
      - Learner must be stochastic (= unpredictable) in choosing actions
    - Amazingly, there are learners with regret guarantees
- Contextual bandits: we also get instance x, make decision $\pi(a|x)$
  - Can we generalize to unseen instances?

Today's lecture

Multi-Armed Bandits

Exploration in Deep RL

Learning with sparse rewards

- Montezuma's Revenge
  - Key = 100 points
  - Door = 500 points
  - Skull = 0 points
    - Is it good? Bad? Affects something off-screen? Opens up an easter egg?
  - Humans have a head start with transfer from known objects
- Exploration before learning:
  - Random walk until you get some points — could take a while!

RL exploration is more complicated...

- Need to consider states and dynamics
- Need coordinated behavior to get *anywhere*
  - E.g., cross a bridge to get the game started...
  - Random exploration will kill us with high probability
    - Structured exploration: noise over time has joint distribution, temporal structure
- How to define regret?
  - With respect to constant action? We can outperform it
  - With respect to optimal policy? May be too hard to learn $\implies$ linear regret
  - Most approaches are heuristic, no regret guarantees; often train-time rewards don't matter

Count-based exploration

- Generalizing UCB exploration $a = \arg\max_i \hat{\mu}_i + \sqrt{\frac{2 \ln T}{n(i)}}$ from MAB to RL
- Count visitations to each state $n(s)$ (or state-action $n(s, a)$)
- Optimism under uncertainty: add exploration bonus to scarcely-visited states

$$\tilde{r} = r + r_e(n(s))$$

  - r_e should be monotonic decreasing in $n(s)$
  - Need to tune its weight

Density model for count-based exploration

- How to represent "counts" in large state spaces?
  - We may never see the same state twice
  - If a state is very similar to ones we've seen often, is it new?
- Train a density model $p_\phi(s)$ over past experience
- Unlike generative models, we care about getting the density correctly
  - But we don't care about the quality of samples

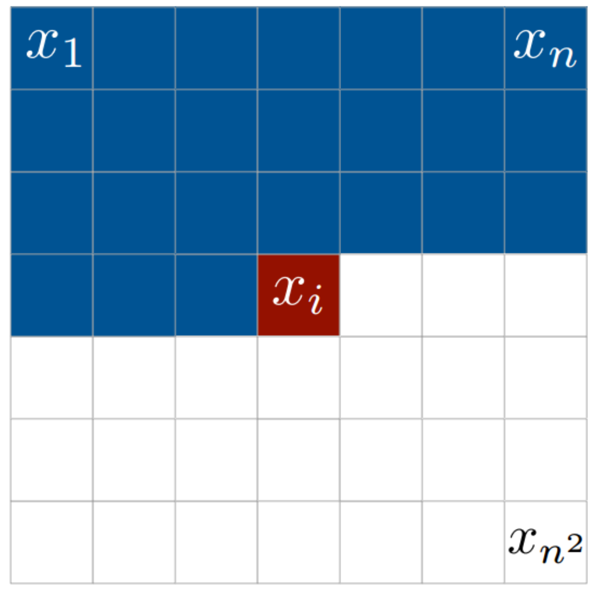

- Density models for images:
  - CTS, PixelRNN, PixelCNN, etc.

Pseudo-counts

- How to infer pseudo-counts from a density model?

$$p_\phi(s) = \frac{n(s)}{N}$$

- After another visit:

$$p_\phi(s) = \frac{n(s)+1}{N+1}$$

- To recover the pseudo-count:
  - $p_{\phi'} \leftarrow$ mock-update the density model with another visit of s
  - Compute

$$\hat{N} = \frac{1 - p_{\phi'}(s)}{p_{\phi'}(s) - p_\phi(s)} p_\phi(s) \qquad \hat{n}(s) = \hat{N} p_\phi(s)$$

Exploration bonus

- What's a good exploration bonus?
- In bandits: Upper Confidence Bound (UCB)
  - $r_e(n(s)) = \sqrt{\dfrac{2 \ln N}{n(s)}}$
- In RL, often:
  - $r_e(n(s)) = \sqrt{\dfrac{1}{n(s)}}$

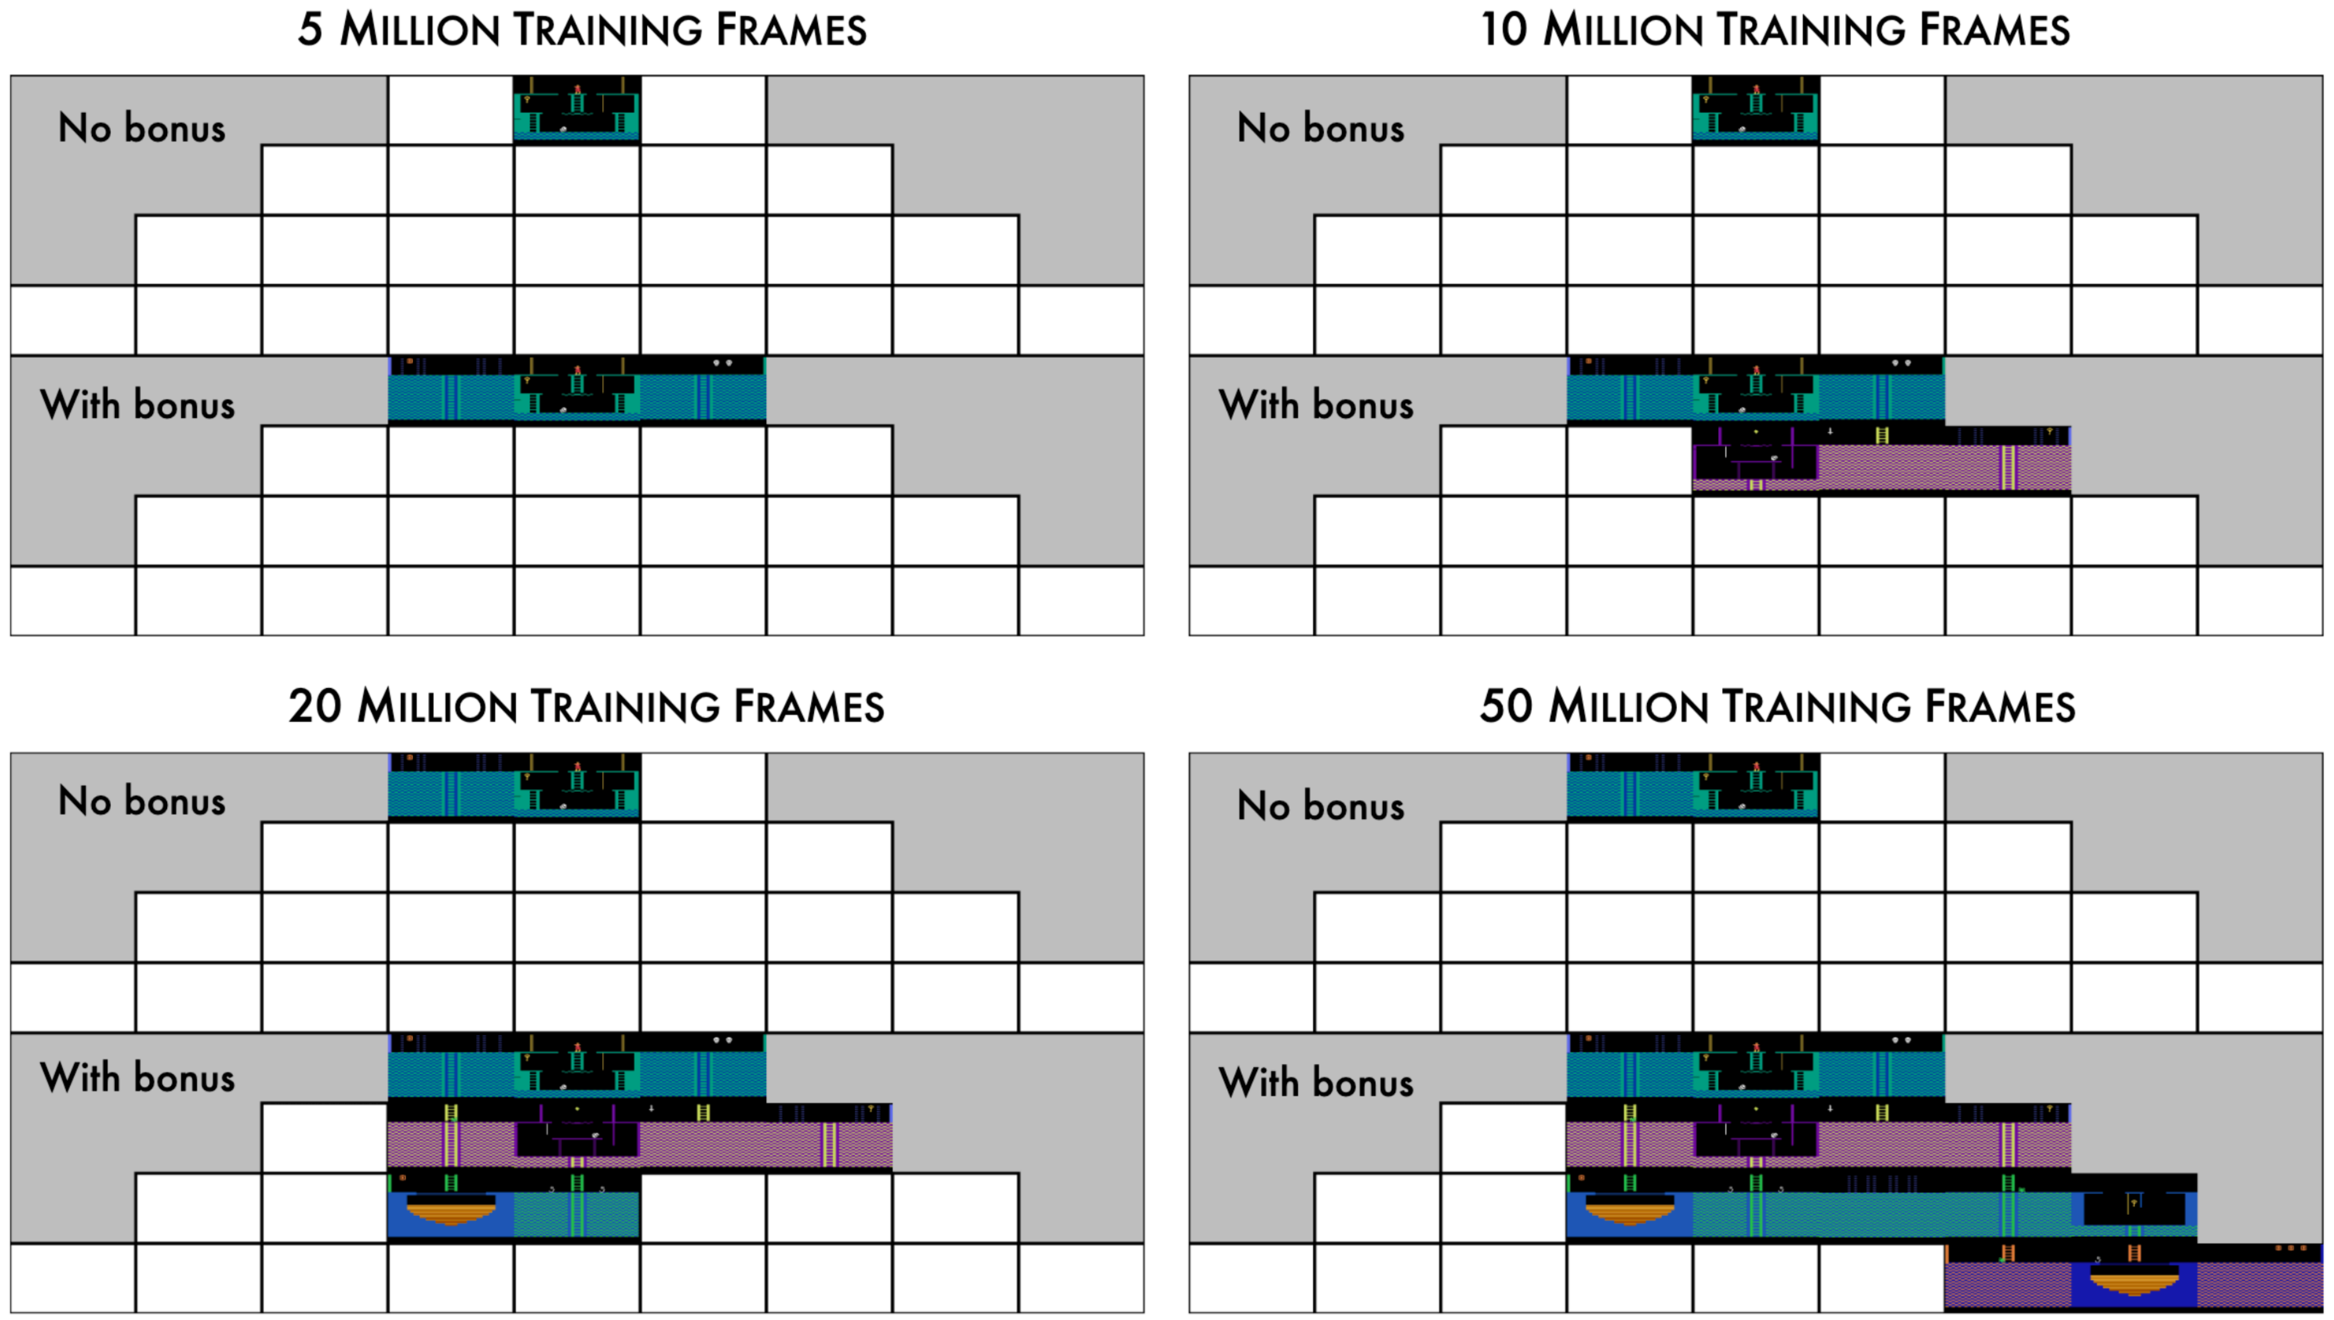

[Bellemare et al., 2016]

Thompson sampling for RL

- Keep a distribution over models $p_\theta(\phi)$
- What's our "model"? Idea 1: MDP; Idea 2: Q-function
- Thompson sampling over Q-functions:
  - Sample $Q \sim p_\theta$
  - Roll out an episode with the greedy policy $\pi(s) = \arg\max_a Q(s, a)$
  - Update p_θ to be more likely for Q' that gives low empirical Bellman error
  - Repeat

Recap

- Online learning = getting good rewards while learning
  - In contrast: learn however, but deploy good policy
- Online learning requires trading off exploration–exploitation
  - Don't overfit to too little data
  - Don't be late to use what you've learned
- Optimism under uncertainty: exploration bonus for novelty
- Thompson sampling: coordinated exploration actions
- Same principles hold in RL