Imitation Learning: Behavior Cloning, Distribution Shift, & Distribution Matching

Sham Kakade and Wen Sun

CS 6789: Foundations of Reinforcement Learning

Announcements

HW3 is out last night and is due Nov 24th 11:59pm

(Lots of bonus questions, but try them out!)

No classes on Nov 17, 19, & 24 (university semi-final week)

Recap

Offline RL

$\mathscr{D} = \{s_i, a_i, r_i, s_i'\}_{i=1}^n$, where $s_i, a_i \sim \mu, r_i = r(s_i, a_i), s_i' \sim P(\cdot \mid s_i, a_i)$

Recap

Offline RL

$\mathscr{D} = \{s_i, a_i, r_i, s_i'\}_{i=1}^n$, where $s_i, a_i \sim \mu, r_i = r(s_i, a_i), s_i' \sim P(\cdot \mid s_i, a_i)$

Fitted Q Iteration: start from $f_0 \in \mathscr{F}$

$$f_{t+1} = \arg\min_{f \in \mathscr{F}} \sum_{i=1}^n \left(f(s_i, a_i) - \underbrace{\left(r_i + \gamma \max_{a'} f_t(s_i', a') \right)}_{\mathscr{T} f_t(s_i, a_i)} \right)^2$$

Recap

Offline RL

$$\mathscr{D} = \{s_i, a_i, r_i, s_i'\}_{i=1}^{n}, \text{ where } s_i, a_i \sim \mu, r_i = r(s_i, a_i), s_i' \sim P(\cdot \mid s_i, a_i)$$

Fitted Q Iteration: start from $f_0 \in \mathscr{F}$

$$f_{t+1} = \arg\min_{f \in \mathscr{F}} \sum_{i=1}^{n} \left(f(s_i, a_i) - \underbrace{\left(r_i + \gamma \max_{a'} f_t(s_i', a') \right)}_{\mathscr{T} f_t(s_i, a_i)} \right)^2$$

Performing regression from (s_i, a_i) to $\mathscr{T} f_t(s_i, a_i)$

Recap

Offline RL

$$\mathscr{D} = \{s_i, a_i, r_i, s_i'\}_{i=1}^n, \text{ where } s_i, a_i \sim \mu, r_i = r(s_i, a_i), s_i' \sim P(\cdot \mid s_i, a_i)$$

Fitted Q Iteration: start from $f_0 \in \mathscr{F}$

$$f_{t+1} = \arg\min_{f \in \mathscr{F}} \sum_{i=1}^n \left(f(s_i, a_i) - \underbrace{\left(r_i + \gamma \max_{a'} f_t(s_i', a') \right)}_{\mathscr{T} f_t(s_i, a_i)} \right)^2$$

Performing regression from (s_i, a_i) to $\mathscr{T} f_t(s_i, a_i)$

$\sup_{\pi, s, a} \frac{d^\pi(s, a)}{\mu(s, a)} < \infty, \quad \forall f \in \mathscr{F}, \mathscr{T} f \in \mathscr{F}, \Rightarrow$ FQI learns near-optimal policy in polynomially sample complexity

Today: Imitation Learning

1. Introduction of Imitation Learning

2. Offline Imitation Learning: Behavior Cloning

3. The hybrid Setting: Statistical Benefit and Distribution Matching

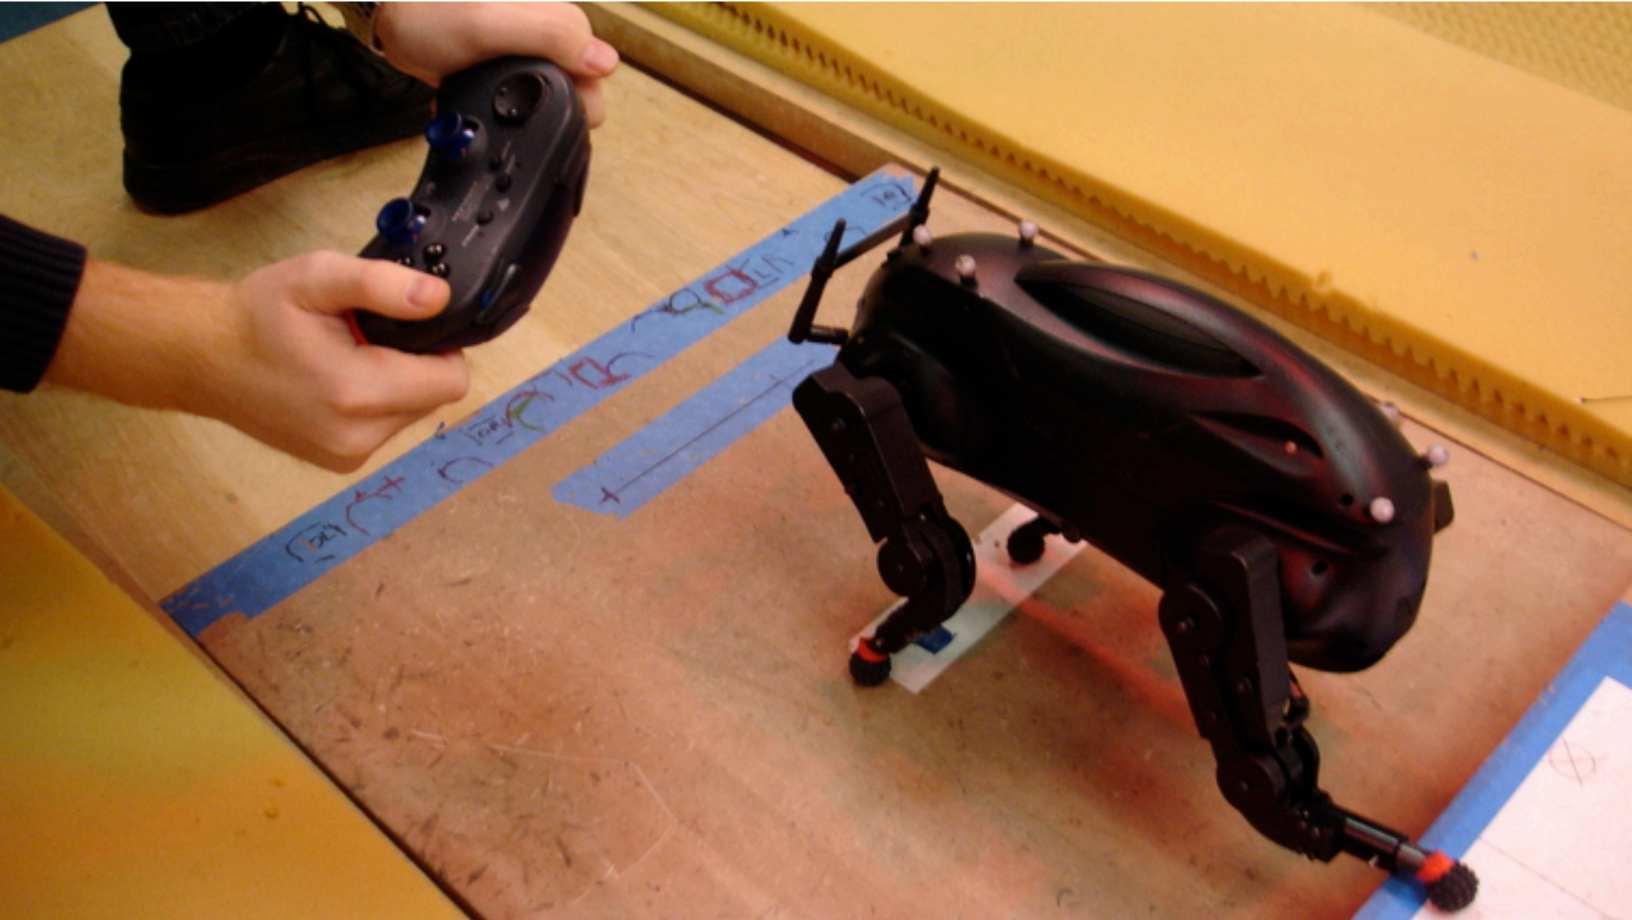

An Autonomous Land Vehicle In A Neural Network [Pomerleau, NIPS '88]

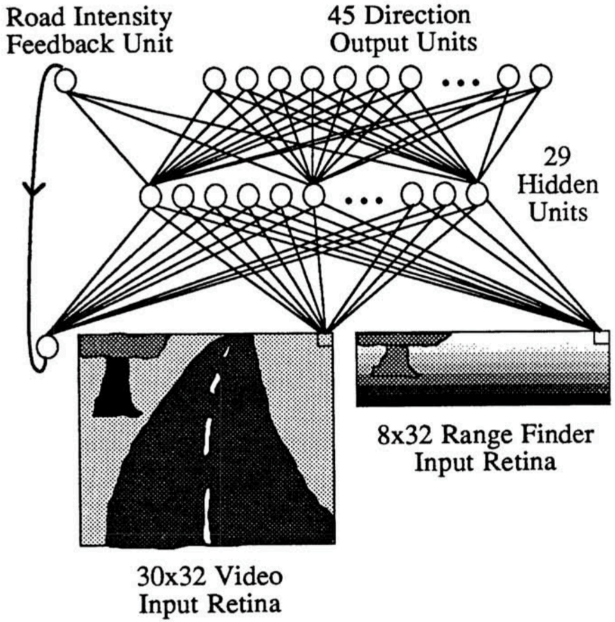

Figure 1: ALVINN Architecture

Imitation Learning

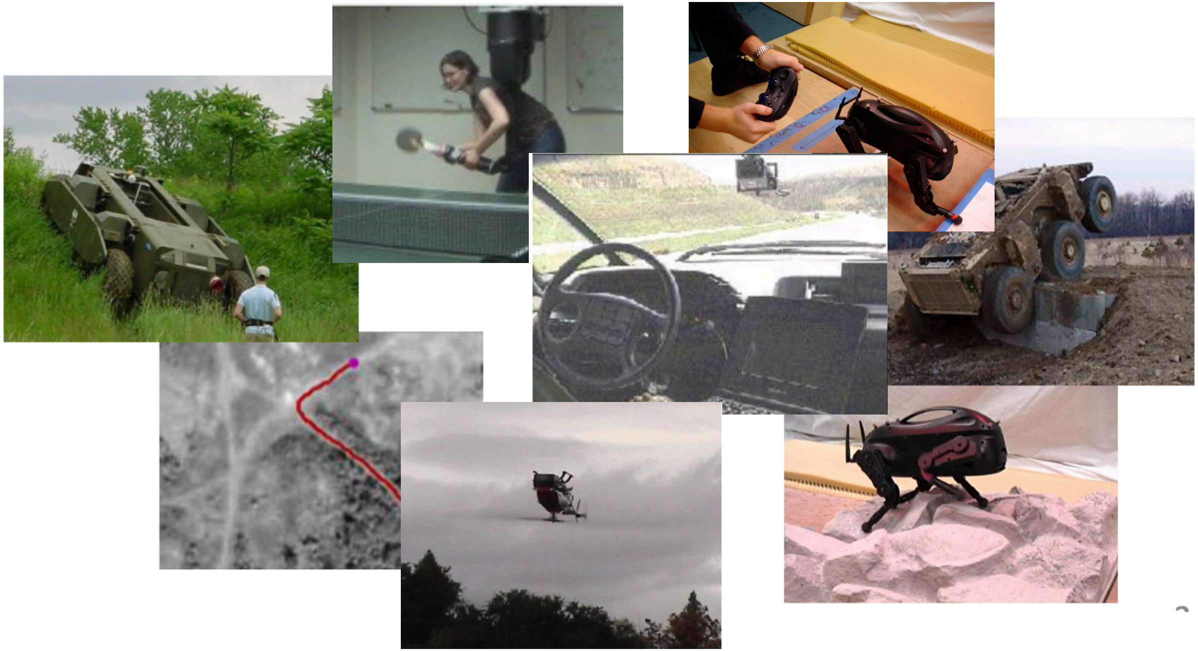

Imitation Learning

Imitation Learning

Imitation Learning

Expert Demonstrations

Imitation Learning

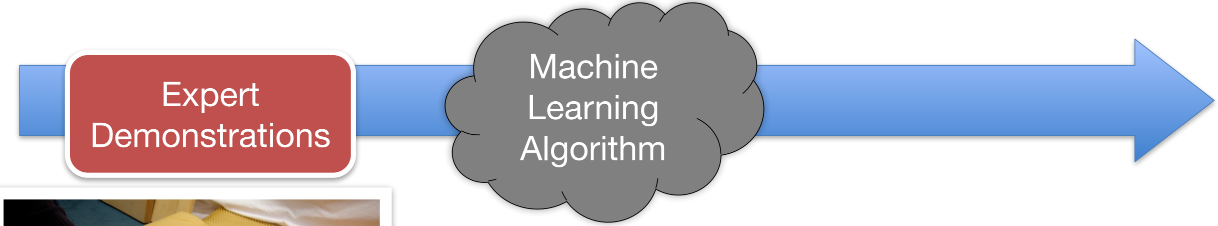

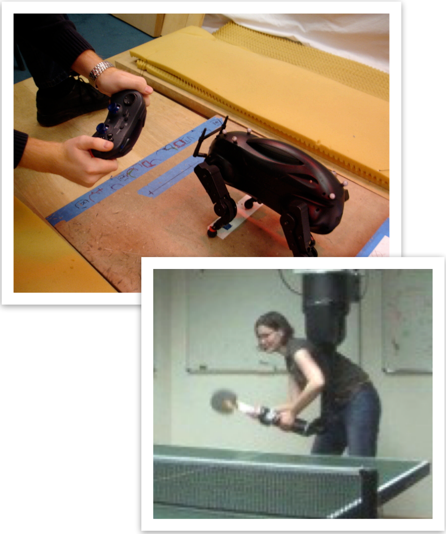

- SVM
- Gaussian Process
- Kernel Estimator
- Deep Networks
- Random Forests
- LWR
- …

Imitation Learning

Expert Demonstrations → Machine Learning Algorithm → Policy π

- SVM
- Gaussian Process
- Kernel Estimator
- Deep Networks
- Random Forests
- LWR
- …

Maps *states* to <u>actions</u>

Learning to Drive by Imitation

[Pomerleau89, Saxena05, Ross11a]

Input:

Camera Image

Policy

Output:

Steering Angle
in [-1, 1]

Supervised Learning Approach: Behavior Cloning

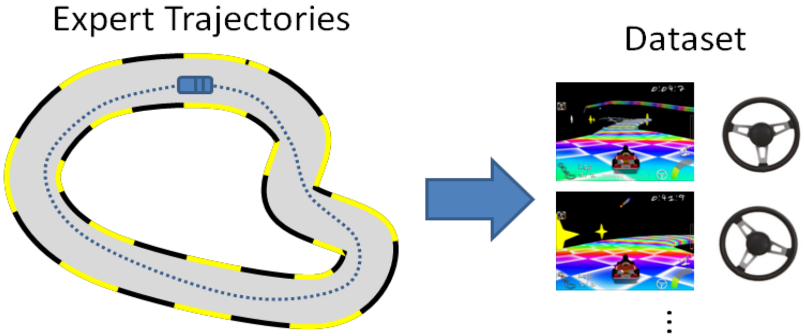

[Widrow64,Pomerleau89]

Supervised Learning Approach: Behavior Cloning

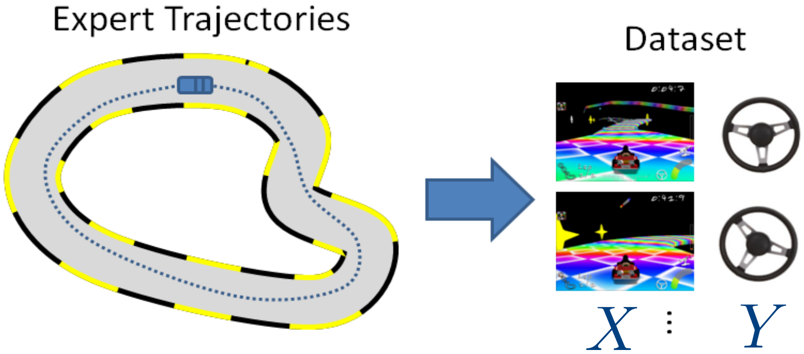

[Widrow64,Pomerleau89]

Supervised Learning Approach: Behavior Cloning

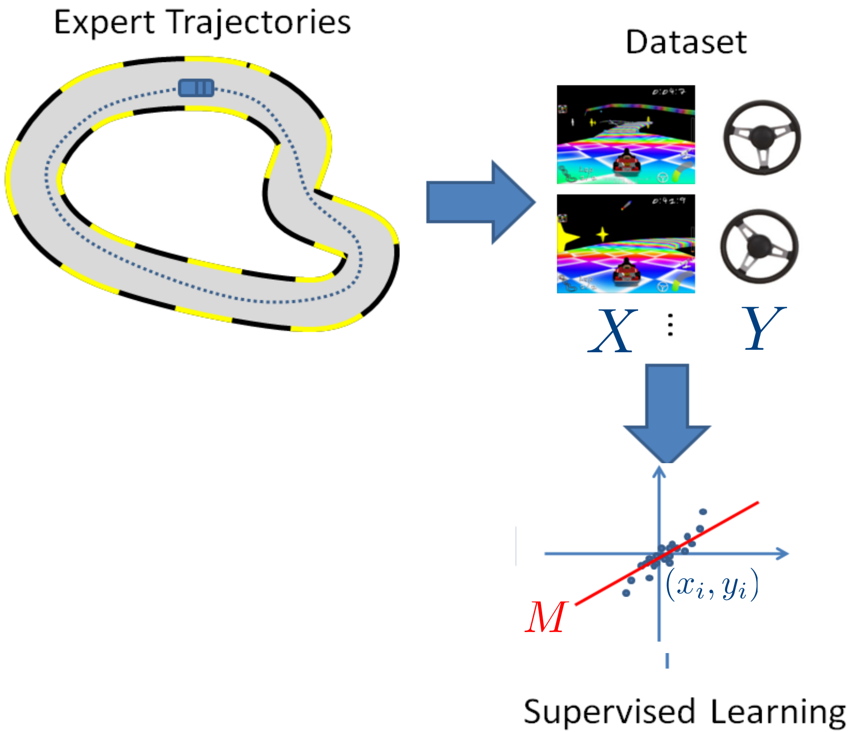

[Widrow64,Pomerleau89]

Supervised Learning Approach: Behavior Cloning

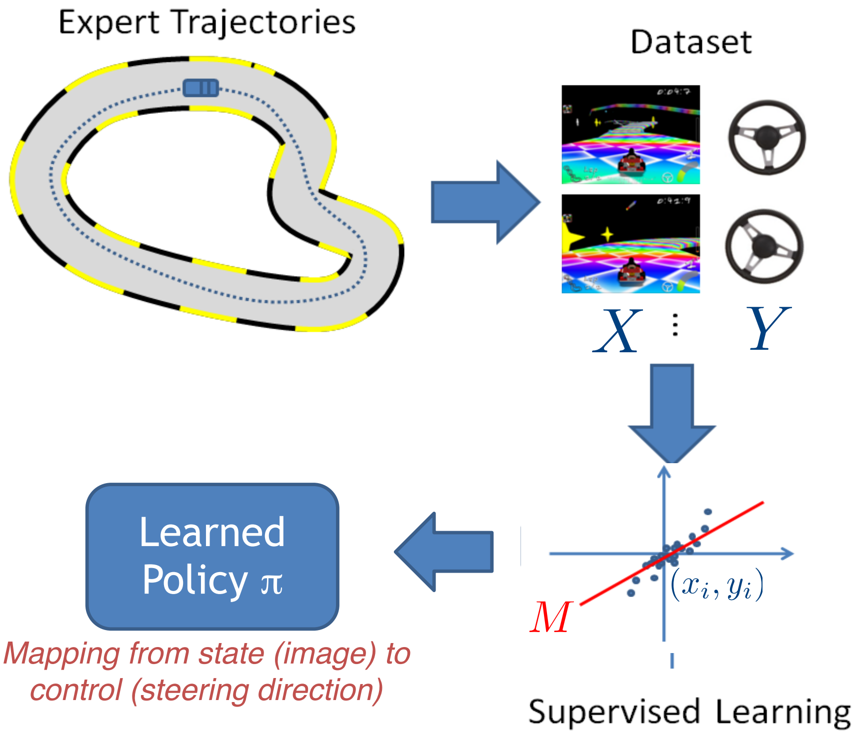


Ready!


Ready!


Ready!

Let's formalize the offline IL Setting and the Behavior Cloning algorithm

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Let's formalize the offline IL Setting and the Behavior Cloning algorithm

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is a near optimal policy $\pi^\star$

Let’s formalize the offline IL Setting and the Behavior Cloning algorithm

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is a near optimal policy $\pi^\star$

We have a dataset $\mathscr{D} = (s_i^\star, a_i^\star)_{i=1}^M \sim d^{\pi^\star}$

Let’s formalize the offline IL Setting and the Behavior Cloning algorithm

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is a near optimal policy $\pi^\star$

We have a dataset $\mathscr{D} = (s_i^\star, a_i^\star)_{i=1}^M \sim d^{\pi^\star}$

Goal: learn a policy from $\mathscr{D}$ that is as good as the expert $\pi^\star$

optimize r

Let's formalize the Behavior Cloning algorithm

BC Algorithm input: a restricted policy class $\Pi = \{\pi : S \mapsto \Delta(A)\}$

A

BC with Maximum Likelihood Estimation (MLE):

Let's formalize the Behavior Cloning algorithm

BC Algorithm input: a restricted policy class $\Pi = \{\pi : S \mapsto \Delta(A)\}$

BC with Maximum Likelihood Estimation (MLE):

$$\hat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Let's formalize the Behavior Cloning algorithm

BC Algorithm input: a restricted policy class $\Pi = \{\pi : S \mapsto \Delta(A)\}$

BC with Maximum Likelihood Estimation (MLE):

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

(We can reduce it to other supervised learning oracles such as classification, regression)

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

$\log |\Pi|$

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: Supervised learning (MLE) guarantee (see the book for reference to the classic MLE analysis):

Analysis

$$\hat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: Supervised learning (MLE) guarantee (see the book for reference to the classic MLE analysis):

Theorem [MLE Guarantee] With probability at least $1 - \delta$, we have:

$$\mathbb{E}_{s \sim d^{\pi^\star}} \left\| \hat{\pi}(\cdot \mid s) - \pi^\star(\cdot \mid s) \right\|_{tv} \leq \sqrt{\frac{\ln(|\Pi|/\delta)}{M}}$$

Analysis

$$\hat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: Supervised learning (MLE) guarantee (see the book for reference to the classic MLE analysis):

Theorem [MLE Guarantee] With probability at least $1-\delta$, we have:

$$\mathbb{E}_{s \sim d^{\pi^\star}} \left\| \hat{\pi}(\cdot \mid s) - \underset{\Delta}{\pi^\star(\cdot \mid s)} \right\|_{tv} \leq \sqrt{\frac{\ln(|\Pi|/\delta)}{M}}$$

$s_i^\star, a_i^\star \sim d_{\pi^\star} \cdot \pi^\star$

This $1/\sqrt{M}$ rate should be expected:
no training and testing mismatch at this stage!

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

$\hat{\pi} \approx \pi^\star$

under $d^{\pi^\star}$

Step 2: Transfer supervised learning error to policy's performance gap

$V^{\hat{\pi}} \Leftarrow d^{\hat{\pi}}$

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 2: Transfer supervised learning error to policy's performance gap

Theorem [BC Sample Complexity] With probability at least $1 - \delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

supervision learning error

$$\mathbb{E}_{s \sim d^{\pi^\star}} \| \widehat{\pi}(\cdot|s) - \pi^\star(\cdot|s) \|_{TV} \leq \sqrt{\frac{\ln \cdots}{M}}$$

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Assumption: Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 2: Transfer supervised learning error to policy's performance gap

Theorem [BC Sample Complexity] With probability at least $1 - \delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

Note that $1/(1-\gamma)^2$ quadratic dependency on effective horizon

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Theorem [BC Sample Complexity] With probability at least $1-\delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

M: # of expert samples

$$\widehat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Theorem [BC Sample Complexity] With probability at least $1-\delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

$$(1-\gamma)\left(V^\star - V^{\widehat{\pi}}\right) = \mathbb{E}_{s \sim d^{\pi^\star}} \mathbb{E}_{a \sim \pi^\star(\cdot|s)} A^{\widehat{\pi}}(s,a)$$

$$\hat{\pi} = \arg\max_{\pi \in \Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Theorem [BC Sample Complexity] With probability at least $1-\delta$, BC returns a policy $\hat{\pi}$:

$$V^{\pi^\star} - V^{\hat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

$$(1-\gamma)\left(V^\star - V^{\hat{\pi}}\right) = \mathbb{E}_{s\sim d^{\pi^\star}}\mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\hat{\pi}}(s,a)$$

$$\underset{\Delta}{=} \mathbb{E}_{s\sim d^{\pi^\star}}\mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\hat{\pi}}(s,a) - \mathbb{E}_{s\sim d^{\pi^\star}}\mathbb{E}_{a\sim\hat{\pi}(\cdot|s)} A^{\hat{\pi}}(s,a)$$

$$\mathbb{E}_{a\sim\hat{\pi}} A^{\hat{\pi}}(s,a) = 0$$

$$Q^{\hat{\pi}}(s,a) - V^{\hat{\pi}}(s')$$

$$\widehat{\pi} = \arg\max_{\pi\in\Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Theorem [BC Sample Complexity] With probability at least $1-\delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \le \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

$$(1-\gamma)\left(V^\star - V^{\widehat{\pi}}\right) = \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\widehat{\pi}}(s,a)$$

$$= \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\widehat{\pi}}(s,a) - \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\widehat{\pi}(\cdot|s)} A^{\widehat{\pi}}(s,a)$$

$$\le \mathbb{E}_{s\sim d^{\pi^\star}} \frac{1}{1-\gamma} \left\| \pi^\star(\cdot \mid s) - \widehat{\pi}(\cdot \mid s) \right\|_1$$

$$\left| \mathbb{E}_{x\sim P_1} f(x_0) - \mathbb{E}_{x\sim P_2} f(x) \right| \le \sup_x |f(x)| \; \| P_1 - P_2 \|_1$$

$$\widehat{\pi} = \arg\max_{\pi\in\Pi} \sum_{i=1}^{M} \ln \pi(a_i^\star \mid s_i^\star)$$

Analysis

Theorem [BC Sample Complexity] With probability at least $1-\delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \frac{2}{(1-\gamma)^2} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{MLE error}}$$

$\varepsilon \Rightarrow \frac{1}{(1-\gamma)^2} \cdot \varepsilon$

ideal: $\frac{1}{1-\gamma} \cdot \varepsilon$

$$(1-\gamma)\left(V^\star - V^{\widehat{\pi}}\right) = \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\widehat{\pi}}(s,a)$$

$A^{\pi} = Q^{\pi} - V^{\pi}$

$$= \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\pi^\star(\cdot|s)} A^{\widehat{\pi}}(s,a) - \mathbb{E}_{s\sim d^{\pi^\star}} \mathbb{E}_{a\sim\widehat{\pi}(\cdot|s)} A^{\widehat{\pi}}(s,a)$$

$$\leq \mathbb{E}_{s\sim d^{\pi^\star}} \frac{1}{1-\gamma} \left\| \pi^\star(\cdot \mid s) - \widehat{\pi}(\cdot \mid s) \right\|_1$$

$$= \frac{2}{1-\gamma} \underbrace{\mathbb{E}_{s\sim d^{\pi^\star}} \|\pi^\star(\cdot \mid s) - \widehat{\pi}(\cdot \mid s)\|_{tv}}_{\text{error-term from MLE}}$$

What could go wrong?

[Pomerleau89,Daume09]

- Predictions affect future inputs/ observations

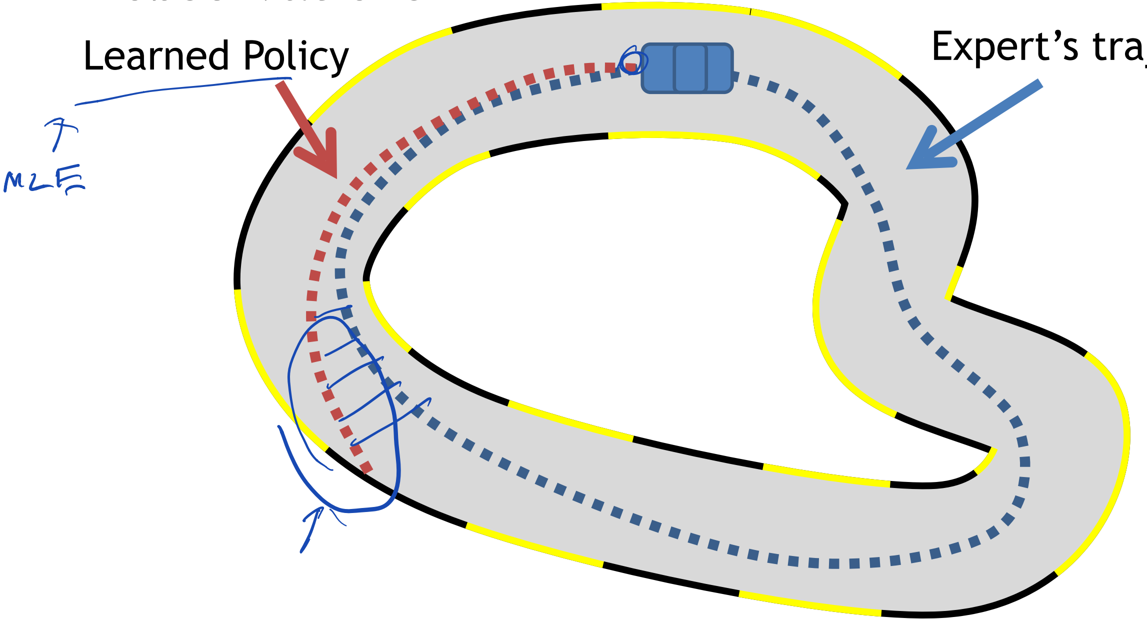

Distribution Shift: Intuitive Explanation

Let's just focus on finite horizon (H) and deterministic policies here:

$$\mathbb{E}_{s\sim d_h^{\pi^\star}} \widehat{\pi}(s) \neq \pi^\star(s) \leq \epsilon, \forall h$$

Distribution Shift: Intuitive Explanation

Let's just focus on finite horizon (H) and deterministic policies here:

$$\mathbb{E}_{s\sim d_h^{\pi^\star}} \widehat{\pi}(s) \neq \pi^\star(s) \leq \epsilon, \forall h$$

s_0

Distribution Shift: Intuitive Explanation

Let's just focus on finite horizon (H) and deterministic policies here:

$$\mathbb{E}_{s\sim d_h^{\pi^\star}} \widehat{\pi}(s) \neq \pi^\star(s) \leq \epsilon, \forall h$$

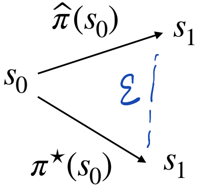

Distribution Shift: Intuitive Explanation

Let's just focus on finite horizon (H) and deterministic policies here:

$$\mathbb{E}_{s \sim d_h^{\pi^\star}} \widehat{\pi}(s) \neq \pi^\star(s) \leq \epsilon, \forall h$$

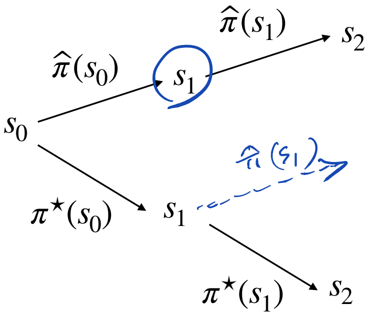

Distribution Shift: Intuitive Explanation

Let's just focus on finite horizon (H) and deterministic policies here:

$$\mathbb{E}_{s \sim d_h^{\pi^\star}} \widehat{\pi}(s) \neq \pi^\star(s) \leq \epsilon, \forall h$$

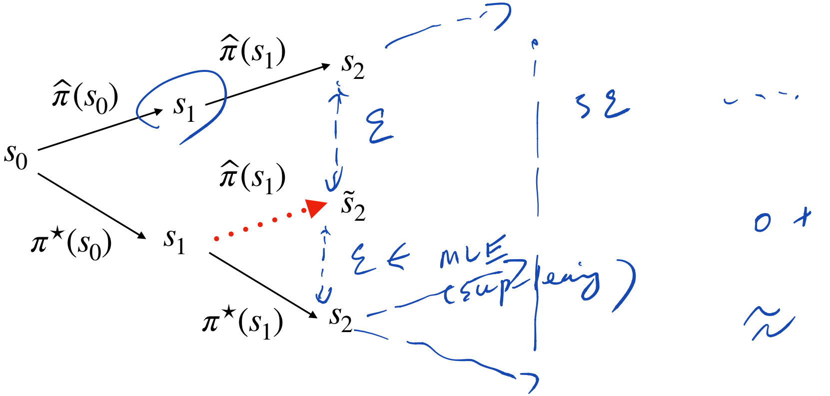

$0 + \varepsilon + 2\varepsilon + \cdots + (H-1)\varepsilon$

$\approx \frac{H^2}{2} \cdot \varepsilon$

An Autonomous Land Vehicle In A Neural Network *[Pomerleau, NIPS '88]*

An Autonomous Land Vehicle In A Neural Network *[Pomerleau, NIPS '88]*

"If the network is not presented with sufficient variability in its training exemplars to cover the conditions it is likely to encounter…[it] will perform poorly"

A potential Fix

A potential Fix

Access to P (maybe a simulator)

Let's roll out our policy in the real world, and compare our trajectories to the expert's trajectories, and then refine our learned model.

The Hybrid Imitation Learning Setting:

Recall BC: we only use offline expert data—no interaction with the environment

Hybrid setting: offline expert data + simulator (e.g., known transition P)

The Hybrid Imitation Learning Setting:

Recall BC: we only use offline expert data—no interaction with the environment

Hybrid setting: offline expert data + simulator (e.g., known transition P)

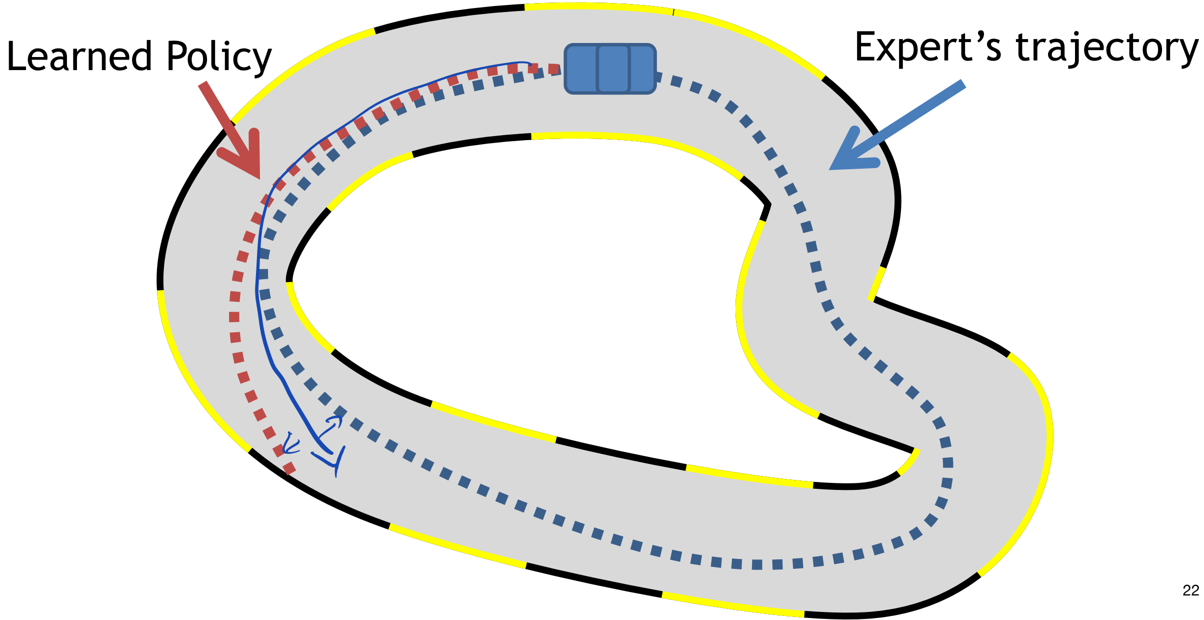

Let's formalize the Hybrid Setting

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is the optimal policy $\pi^\star$

We have a dataset $\mathscr{D} = (s_i^\star, a_i^\star)_{i=1}^M \sim d^{\pi^\star}$

Let's formalize the Hybrid Setting

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is the optimal policy $\pi^\star$

We have a dataset $\mathscr{D} = (s_i^\star, a_i^\star)_{i=1}^M \sim d^{\pi^\star}$

This time, we have a known transition P (but we cannot plan because r is unknown)

Let's formalize the Hybrid Setting

Discounted infinite horizon MDP $\mathscr{M} = \{S, A, \gamma, r, P, \rho, \pi^\star\}$

Ground truth reward $r(s, a) \in [0,1]$ is unknown; assume expert is the optimal policy $\pi^\star$

We have a dataset $\mathscr{D} = (s_i^\star, a_i^\star)_{i=1}^M \sim d^{\pi^\star}$

This time, we have a known transition P (but we cannot plan because r is unknown)

Key Q: can we do better than offline IL Behavior Cloning (statistically at least—assuming infinite computation power)?

Key Idea: distribution matching

Integral probability metric (IPM)

Key Idea: distribution matching

Integral probability metric (IPM)

Metric measures the divergence between two distributions

Key Idea: distribution matching

Integral probability metric (IPM)

Metric measures the divergence between two distributions

Given a discriminator class: $\mathscr{F} = \{f : X \mapsto \mathbb{R}\}$, and two distributions p_1 and p_2

Key Idea: distribution matching

Integral probability metric (IPM)

Metric measures the divergence between two distributions

Given a discriminator class: $\mathscr{F} = \{f : X \mapsto \mathbb{R}\}$, and two distributions p_1 and p_2

$$\text{IPM}_{\mathscr{F}}(p_1, p_2) = \max_{f \in \mathscr{F}} \left[\mathbb{E}_{x \sim p_1} f(x) - \mathbb{E}_{x \sim p_2} f(x) \right]$$

Key Idea: distribution matching

Integral probability metric (IPM)

Metric measures the divergence between two distributions

Given a discriminator class: $\mathcal{F} = \{f : X \mapsto \mathbb{R}\}$, and two distributions p_1 and p_2

$$\text{IPM}_{\mathcal{F}}(p_1, p_2) = \max_{f \in \mathcal{F}} \left[\mathbb{E}_{x \sim p_1} f(x) - \mathbb{E}_{x \sim p_2} f(x) \right]$$

$$\mathcal{F} = \{f : \|f\|_\infty \leq 1\} \Rightarrow \text{IPM}_{\mathcal{F}}(p_1, p_2) := \|p_1 - p_2\|_{tv}$$

Key Idea: distribution matching

Integral probability metric (IPM)

Metric measures the divergence between two distributions

Given a discriminator class: $\mathcal{F} = \{f : X \mapsto \mathbb{R}\}$, and two distributions p_1 and p_2

min
p_1

$$\text{IPM}_{\mathcal{F}}(p_1, p_2) = \max_{f \in \mathcal{F}} \left[\mathbb{E}_{x \sim p_1} f(x) - \mathbb{E}_{x \sim p_2} f(x) \right]$$

real

$$\mathcal{F} = \{f : \|f\|_\infty \le 1\} \Rightarrow \text{IPM}_{\mathcal{F}}(p_1, p_2) := \|p_1 - p_2\|_{tv}$$

$$\mathcal{F} = \{f : f \text{ is 1-Lipschitz}\} \Rightarrow \text{IPM}_{\mathcal{F}}(p_1, p_2) := \text{wasserstein dis}(p_1, p_2)$$

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step 2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathscr{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step 2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathscr{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Set refined discriminator class $\widetilde{\mathscr{F}} := \{f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi'\}$

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathscr{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step 2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathscr{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Set refined discriminator class $\widetilde{\mathscr{F}} := \{f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi'\}$

Q: what is the size of $\widetilde{\mathscr{F}}$?

$|\Pi|^2$

Algorithm: Distribution Matching TV distance

Consider the Discriminator class: $\mathcal{F} = \{f : \|f\|_\infty \leq 1\}$ (IPM corresponds to TV distance)

Same assumption as we had in BC:
Π is discrete, and realizable, i.e., $\pi^\star \in \Pi$

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step 2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathcal{F}} \left[\mathbb{E}_{s,a\sim d^\pi} f(s,a) - \mathbb{E}_{s,a\sim d^{\pi'}} f(s,a)\right]$$

Q: what is the size of $\widetilde{\mathcal{F}}$?

Set refined discriminator class $\widetilde{\mathcal{F}} := \{f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi'\}$

$$\forall \pi \,\&\, \pi' \in \Pi,\ \|d^\pi - d^{\pi'}\|_{tv} = \max_{f \in \widetilde{\mathcal{F}}} \mathbb{E}_{s,a\sim d^\pi} f(s,a) - \mathbb{E}_{s,a\sim d^{\pi'}} f(s,a)$$

Algorithm: Distribution Matching TV distance

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathscr{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Set refined discriminator class $\widetilde{\mathscr{F}} := \{ f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi' \}$

Algorithm: Distribution Matching TV distance

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathcal{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Set refined discriminator class $\widetilde{\mathcal{F}} := \{ f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi' \}$

Step 3: Select a policy using expert dataset $\mathcal{D} = \{ s_i^\star, a_i^\star \}_{i=1}^M \leftarrow d^{\pi^\star}$

Algorithm: Distribution Matching TV distance

Step 1: for each $\pi \in \Pi$, compute $d^\pi \in \Delta(S \times A)$ (recall P is known);
(This step is computationally inefficient)

Step2: select useful discriminators: for all pair π & π', with $\pi \neq \pi'$

$$f_{\pi,\pi'} = \arg\max_{f \in \mathcal{F}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \mathbb{E}_{s,a \sim d^{\pi'}} f(s,a) \right]$$

Set refined discriminator class $\widetilde{\mathcal{F}} := \{ f_{\pi,\pi'} : \pi \,\&\, \pi' \in \Pi, \pi \neq \pi' \}$

Step 3: Select a policy using expert dataset $\mathcal{D} = \{ s_i^\star, a_i^\star \}_{i=1}^M$

$$\hat{\pi} := \arg\min_{\pi \in \Pi} \left[\overbrace{\max_{f \in \widetilde{\mathcal{F}}} \left[\mathbb{E}_{s,a \sim d^\pi} f(s,a) - \frac{1}{M} \sum_{i=1}^{M} f(s_i^\star, a_i^\star) \right]}^{\mathrm{IPM}_{\widetilde{\mathcal{F}}}(d^\pi, \hat{d}^{\pi^\star})} \right]$$

Δ

Theorem: Distribution Matching TV distance

Theorem [Dis-match w/ TV dist] With probability at least $1 - \delta$, our algorithm finds a policy $\hat{\pi}$, s.t.,

$$V^{\pi^\star} - V^{\hat{\pi}} \leq \mathcal{O}\left(\frac{1}{1-\gamma} \underbrace{\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}}_{\text{samp-ling error}} \right)$$

Theorem: Distribution Matching TV distance

Theorem [Dis-match w/ TV dist] With probability at least $1-\delta$, our algorithm finds a policy $\widehat{\pi}$, s.t.,

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \mathcal{O}\left(\frac{1}{1-\gamma} \sqrt{\frac{\ln(|\Pi|/\delta)}{M}} \right)$$

1. Key step is to prove: $\left\| d^{\widehat{\pi}} - d^{\pi^\star} \right\|_{tv} \leq \sqrt{\frac{\ln(|\Pi|/\delta)}{M}}$

Theorem: Distribution Matching TV distance

Theorem [Dis-match w/ TV dist] With probability at least $1-\delta$, our algorithm finds a policy $\hat{\pi}$, s.t.,

$$V^{\pi^\star} - V^{\hat{\pi}} \leq \mathcal{O}\left(\frac{1}{1-\gamma}\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}\right)$$

$|\mathcal{F}| = |\Pi|^2$

1. Key step is to prove: $\| d^{\hat{\pi}} - d^{\pi^\star} \|_{tv} \leq \sqrt{\frac{\ln(|\Pi|/\delta)}{M}}$

2. For performance: $V^{\pi^\star}_{\mu} - V^{\pi}_{\mu} \leq \frac{1}{1-\gamma}\left[\mathbb{E}_{s,a\sim d^{\pi^\star}} r(s,a) - \mathbb{E}_{s,a\sim d^{\hat{\pi}}} r(s,a)\right]$

$\leq \underbrace{r_{max}}_{=1} \| d^{\pi^\star} - d^{\hat{\pi}} \|_1$

Theorem: Distribution Matching TV distance

Theorem [Dis-match w/ TV dist] With probability at least $1-\delta$, our algorithm finds a policy $\widehat{\pi}$, s.t.,

$$V^{\pi^\star} - V^{\widehat{\pi}} \leq \mathcal{O}\left(\frac{1}{1-\gamma}\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}\right)$$

1. Key step is to prove: $\left\| d^{\widehat{\pi}} - d^{\pi^\star} \right\|_{tv} \leq \sqrt{\frac{\ln(|\Pi|/\delta)}{M}}$

2. For performance: $V^{\pi^\star} - V^{\pi} \leq \frac{1}{1-\gamma}\left[\mathbb{E}_{s,a\sim d^{\pi^\star}} r(s,a) - \mathbb{E}_{s,a\sim d^{\widehat{\pi}}} r(s,a)\right]$

Theorem [Offline BC] With probability at least $1-\delta$, BC returns a policy $\widehat{\pi}$:

$$V^{\pi^\star} - V^{\widehat{\pi}} = \mathcal{O}\left(\frac{1}{(1-\gamma)^2}\sqrt{\frac{\ln(|\Pi|/\delta)}{M}}\right)$$

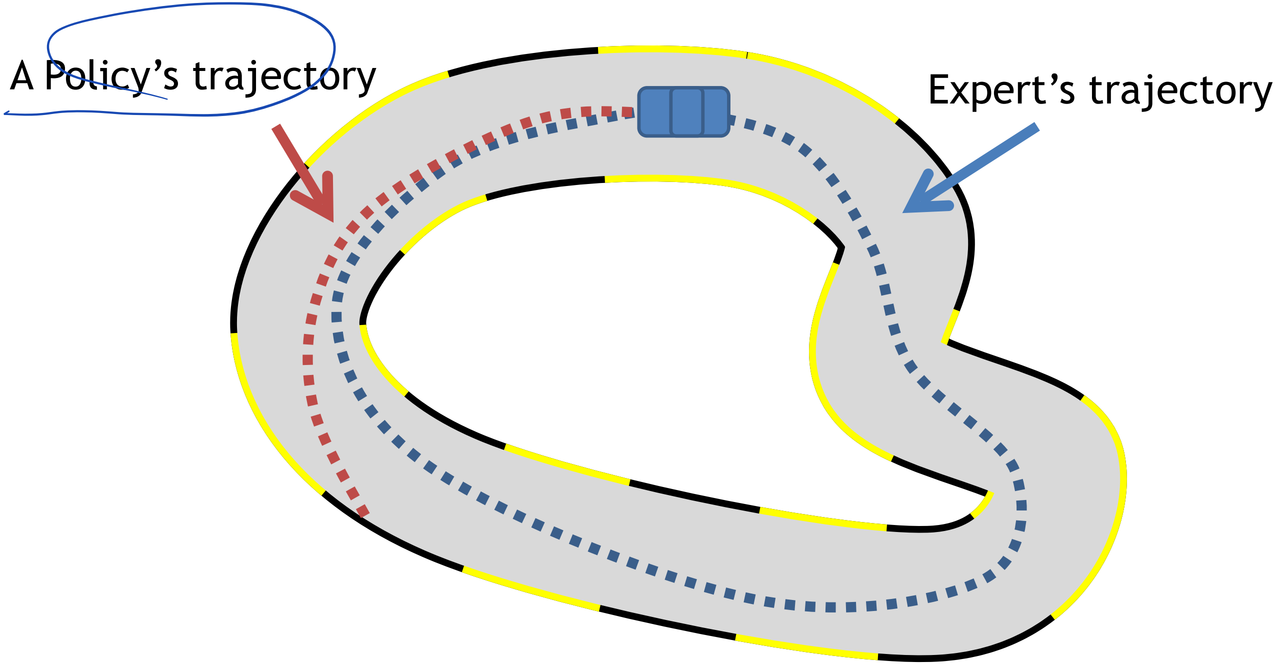
A Policy's trajectory
Expert's trajectory

d^{π} : generator that generators state-action pairs

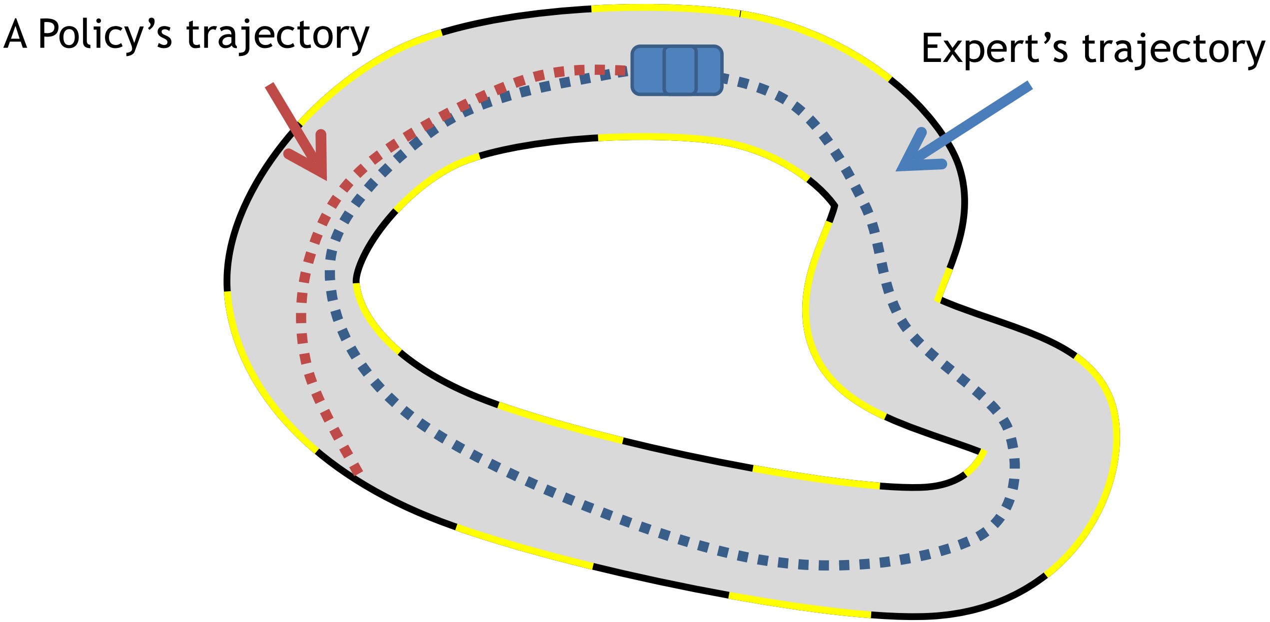

d^{π} : generator that generators state-action pairs

$\{s_i^\star, a_i^\star\}_{i=1}^{M} \sim d^{\pi^\star}$: Ground truth state-action distribution (we have samples from it)

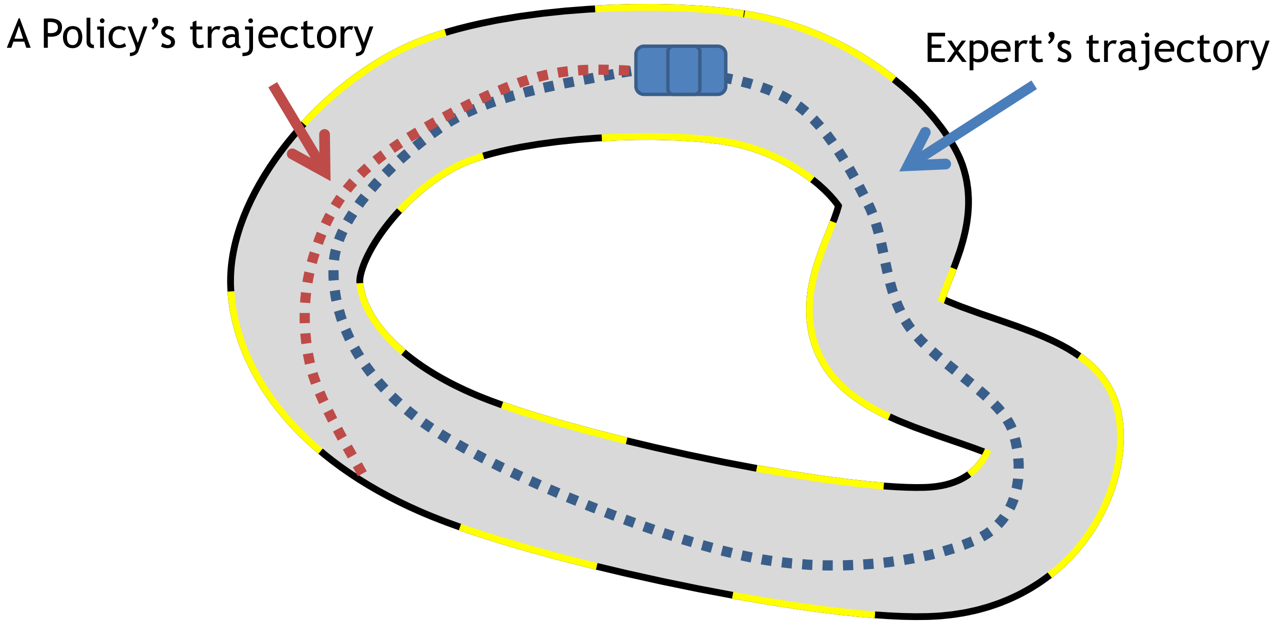

d^{π} : generator that generators state-action pairs

$d^{\pi^\star}$: Ground truth state-action distribution (we have samples from it)

$\widetilde{\mathscr{F}}$: discriminators which distinguish red and blue

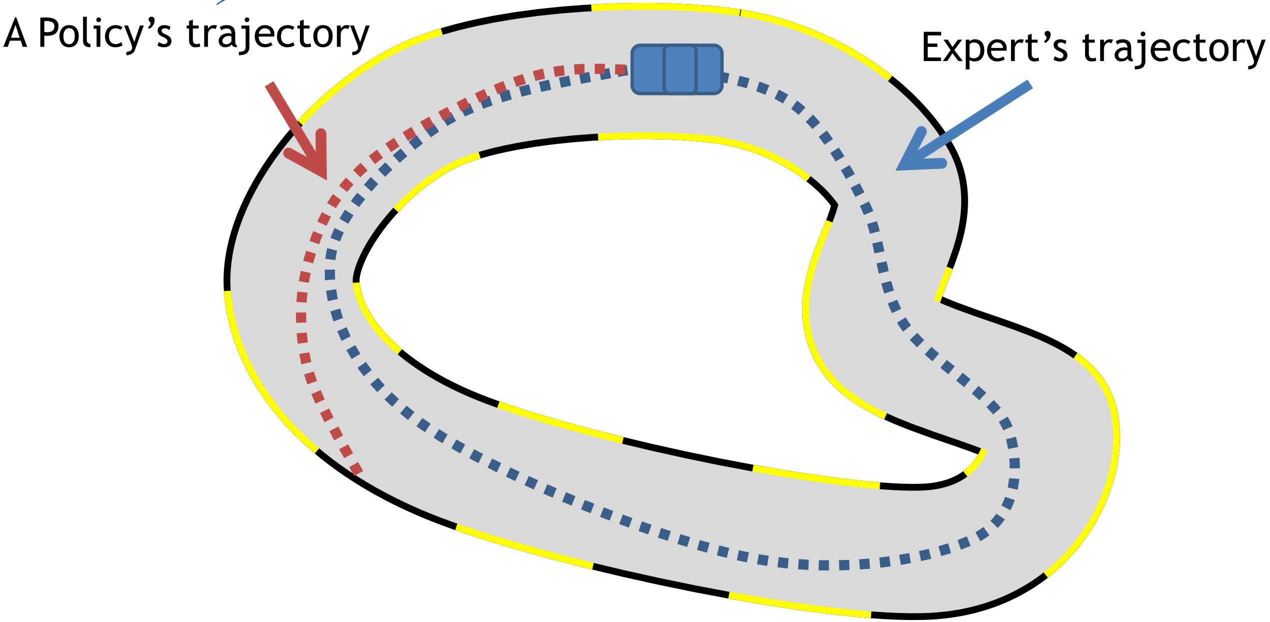

Next lecture we will talk about a computationally efficient algorithm in the hybrid setting

Conclusion:

1. Offline RL: only use offline expert data

BC is simple and easy to implement, has reasonable guarantees; but the quadratic dependency on horizon could cause real problems

Conclusion:

1. Offline RL: only use offline expert data

BC is simple and easy to implement, has reasonable guarantees; but the quadratic dependency on horizon could cause real problems

2. Hybrid RL: offline expert data + known transition (simulator)

Statistically, Distribution-matching has linear dependency on horizon, but the algorithm is computationally inefficient

$\forall \pi \in \Pi$

$\Rightarrow d^{\pi} \in \Delta(S \times A)$

Sham's Alg: (MLE)

$\forall \pi \Rightarrow d^{\pi}$

$\arg\max_{\pi \in \Pi} \hat{\sum_{i=1}} \ln d^{\pi}(s_i^*, a_i^*)$

$\{P_1 \dots P_N\}$ $P^* \sim \{x_i\}_{i=1}$

$(P^* \notin \{P_1 \dots P_N\})$

$\hat{P} \in \{P_1 \dots P_N\}$

$\|\hat{P} - P^*\|_{TV} \le 3 \min_{P \in \{P_1 \dots P_N\}} \|P - P^*\|_{TV} + \sqrt{\frac{\ln N/\delta}{M}}$

~~sche~~ sheffé estimator

Conclusion:

1. Offline RL: only use offline expert data

BC is simple and easy to implement, has reasonable guarantees; but the quadratic dependency on horizon could cause real problems

2. Hybrid RL: offline expert data + known transition (simulator)

Statistically, Distribution-matching has linear dependency on horizon, but the algorithm is computationally inefficient

Take home message:

There is a provable statistical benefit from the hybrid setting!

Ps: the distribution matching algorithm is very new (it was discovered when I was writing the book chapter…)