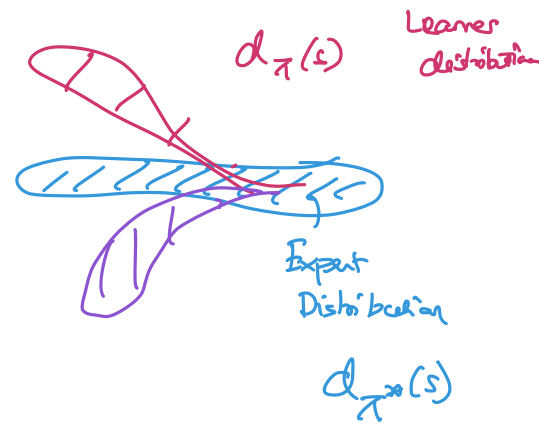


CHICKEN-OR-EGG
LEARNER DISTRIBUTION



ALGORITHM

- ① Start with some policy π
- ② Rollout policy π
- ③ See state it visits $\sim d_{\pi}(s)$
- ④ Train on new data.

QUERY EXPERT
 $a^* \leftarrow \pi^*(s)$

PROPOSED ALGORITHM

INIT: Random policy π_0


for $i = 0 \dots (N-1)$

- Collect data of state π_i visits

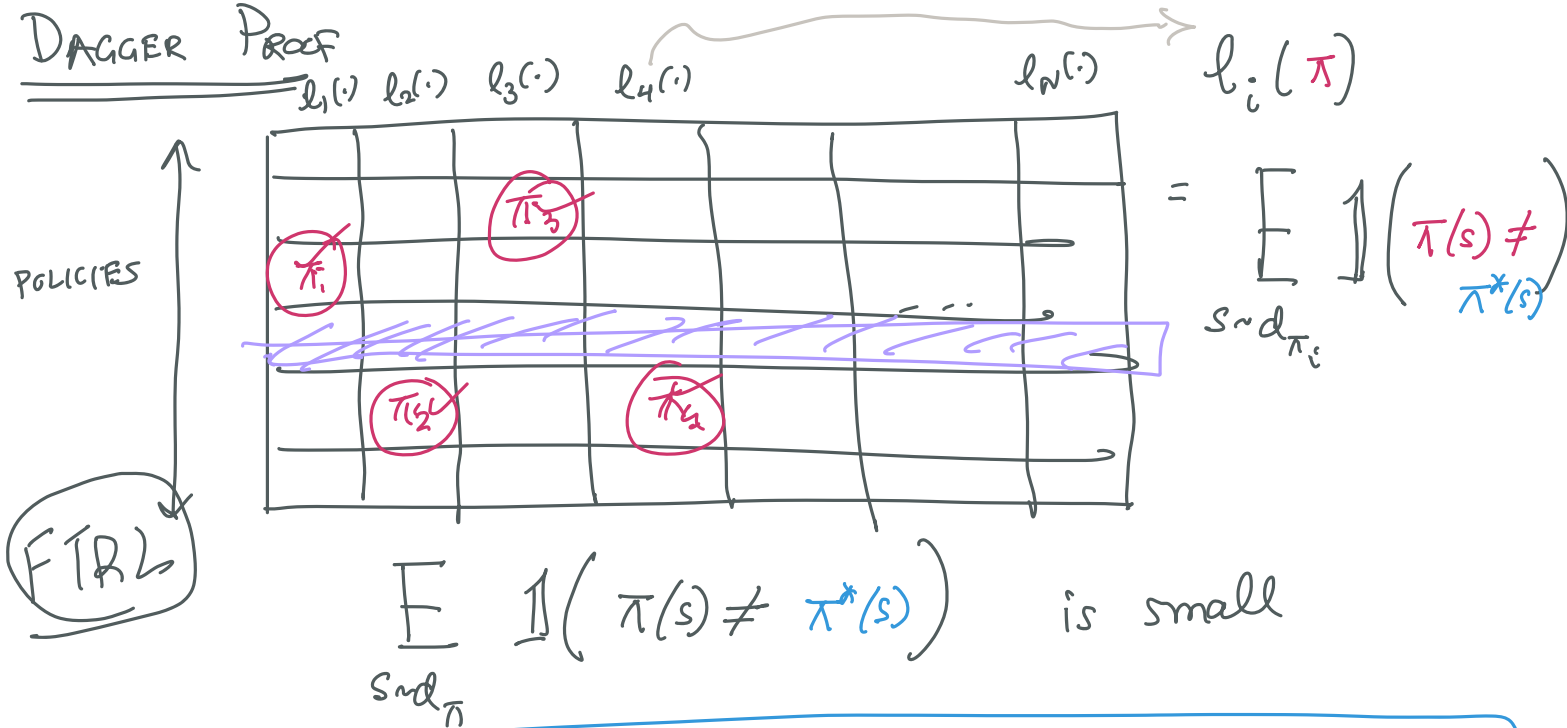
$$s \sim d_{\pi_i}(s)$$

- Query interactive expert to get $\pi^*(s) \rightarrow a^*$

- TRAIN π_{i+1} on (s, a^*)



DAGGER PROOF



Is there a round i where $\underline{l_i(\pi_i)}$ is small.

$$\begin{aligned}
 \min_{i=1, \dots, N} l_i(\pi_i) &\leq \frac{1}{N} \sum_{i=1}^N l_i(\pi_i) \\
 &\leq \left[\frac{1}{N} \sum_{i=1}^N l_i(\pi_i) - \frac{1}{N} \min_{\pi \in \Pi} \sum_{i=1}^N l_i(\pi) \right] + \frac{1}{N} \min_{\pi \in \Pi} \sum_{i=1}^N l_i(\pi) \\
 &\leq \frac{1}{N} \text{REG}(\cdot) + \frac{1}{N} \min_{\pi \in \Pi} \sum_{i=1}^N l_i(\pi) \\
 \lim_{N \rightarrow \infty} \frac{1}{N} \text{REG}(\cdot) &\rightarrow 0 \qquad \leq \epsilon
 \end{aligned}$$

