

Assignment 1: Imitation Learning

Due September 11, 11:59 pm

The goal of this assignment is to gain familiarity with imitation learning, including direct behavioral cloning and the DAGGER algorithm. In lieu of a human demonstrator, demonstrations will be provided via an expert policy that we have trained for you. Your goals will be to set up behavior cloning and DAGGER, and compare their performance on a few different continuous control tasks from the OpenAI Gym benchmark suite. Turn in your report and code as described in Section 6.

The starter-code for this assignment can be found at

https://github.com/berkeleydeeprlcourse/homework_fall2023/tree/main/hw1

You have the option of running the code either on Google Colab or on your own machine. Please refer to the README for more information on setup.

Note: The Colab is only used as a source of GPU compute, so you will be editing the same code regardless of what option you choose. For this assignment, since GPU will not be necessary, we strongly recommend running the code locally to gain some familiarity with installing the necessary packages. This will be extremely beneficial for later homeworks, so you can run experiments in parallel.

If you are running locally we *strongly* recommend you use Conda to manage your Python environment and dependencies. Instructions for installing Conda and setting up an environment are included.

1 Analysis

Consider the problem of imitation learning within a discrete MDP with horizon T and an expert policy π^* . We gather expert demonstrations from π^* and fit an imitation policy π_θ to these trajectories so that

$$\mathbb{E}_{p_{\pi^*}(s)} \pi_\theta(a \neq \pi^*(s) \mid s) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{p_{\pi^*}(s_t)} \pi_\theta(a_t \neq \pi^*(s_t) \mid s_t) \leq \varepsilon,$$

i.e., the expected likelihood that the learned policy π_θ disagrees with the expert π^* within the training distribution p_{π^*} of states drawn from random expert trajectories is at most ε .

For convenience, the notation $p_\pi(s_t)$ indicates the state distribution under π at time step t while $p(s)$ indicates the state marginal of π across time steps, unless indicated otherwise.

1. Show that $\sum_{s_t} |p_{\pi_\theta}(s_t) - p_{\pi^*}(s_t)| \leq 2T\varepsilon$.

Hint 1: in lecture, we showed a similar inequality under the stronger assumption $\pi_\theta(s_t \neq \pi^*(s_t) \mid s_t) \leq \varepsilon$ for every $s_t \in \text{supp}(p_{\pi^*})$. Try converting the inequality above into an expectation over p_{π^*} .

Hint 2: use the union bound inequality: for a set of events E_i , $\Pr[\bigcup_i E_i] \leq \sum_i \Pr[E_i]$

2. Consider the expected return of the learned policy π_θ for a state-dependent reward $r(s_t)$, where we assume the reward is bounded with $|r(s_t)| \leq R_{\max}$:

$$J(\pi) = \sum_{t=1}^T \mathbb{E}_{p_\pi(s_t)} r(s_t).$$

- (a) Show that $J(\pi^*) - J(\pi_\theta) = \mathcal{O}(T\varepsilon)$ when the reward only depends on the last state, i.e., $r(s_t) = 0$ for all $t < T$.
- (b) Show that $J(\pi^*) - J(\pi_\theta) = \mathcal{O}(T^2\varepsilon)$ for an arbitrary reward.

2 Editing Code

The starter code provides an expert policy for each of the MuJoCo tasks in OpenAI Gym. Fill in the blanks in the code marked with `TODO` to implement behavioral cloning. A command for running behavioral cloning is given in the README file.

We recommend that you read the files in the following order.

- `scripts/run_hw1.py` (training loop)
- `policies/MLP_policy.py` (policy definition)
- `infrastructure/replay_buffer.py` (stores training trajectories)
- `infrastructure/utils.py` (utilities for sampling trajectories from a policy)
- `infrastructure/pytorch_utils.py` (utilities for converting between NumPy/Pytorch)

For some files, some important functionality is missing and are marked with `TODO`. Specifically, you are asked to implement parts of the following:

- `policies/MLP_policy.py`: `forward` and `update` functions
- `infrastructure/utils.py`: `sample_trajectory` function
- `scripts/run_hw1.py`: `run_training_loop` function (most of your code will be in here)

3 Behavioral Cloning

1. Run behavioral cloning (BC) and report results on two tasks: one where a behavioral cloning agent should achieve at least 30% of the performance of the expert, and one environment of your choosing where it does not. Here is how you can run the Ant task:

```
python cs285/scripts/run_hw1.py \  
  --expert_policy_file cs285/policies/experts/Ant.pkl \  
  --env_name Ant-v4 --exp_name bc_ant --n_iter 1 \  
  --expert_data cs285/expert_data/expert_data_Ant-v4.pkl \  
  --video_log_freq -1
```

When providing results, report the mean and standard deviation of your policy’s return over multiple rollouts in a table, and state which task was used. When comparing one that is working versus one that is not working, be sure to set up a fair comparison in terms of network size, amount of data, and number of training iterations. Provide these details (and any others you feel are appropriate) in the table caption.

Note: What “report the mean and standard deviation” means is that your `eval_batch_size` should be greater than `ep_len`, such that you’re collecting multiple rollouts when evaluating the performance of your trained policy. For example, if `ep_len` is 1000 and `eval_batch_size` is 5000, then you’ll be collecting approximately 5 trajectories (maybe more if any of them terminate early), and the logged `Eval_AverageReturn` and `Eval_StdReturn` represents the mean/std of your policy over these 5 rollouts. Make sure you include these parameters in the table caption as well.

Tip: To generate videos of the policy, remove the flag `-video_log_freq -1` However, this is slower, and so you probably want to keep this flag on while debugging.

2. Experiment with one set of hyperparameters that affects the performance of the behavioral cloning agent, such as the amount of training steps, the amount of expert data provided, or something that you come up with yourself. For one of the tasks used in the previous question, show a graph of how the BC agent’s performance varies with the value of this hyperparameter. In the caption for the graph, state the hyperparameter and a brief rationale for why you chose it.

4 DAGGER

1. Using the same code, you should be able to run DAGGER by modifying the runtime parameters as follows:

```
python cs285/scripts/run_hw1.py \  
  --expert_policy_file cs285/policies/experts/Ant.pkl \  
  --env_name Ant-v4 --exp_name dagger_ant --n_iter 10 \  
  --do_dagger --expert_data cs285/expert_data/expert_data_Ant-v4.pkl \  
  --video_log_freq -1
```

2. Run DAGGER and report results on the two tasks you tested previously with behavioral cloning. Report your results in the form of a learning curve, plotting the number of DAGGER iterations vs. the policy's mean return, with error bars to show the standard deviation. Include the performance of the expert policy and the behavioral cloning agent on the same plot (as horizontal lines that go across the plot). In the caption, state which task you used, and any details regarding network architecture, amount of data, etc. (as in the previous section).

5 Discussion

Please answer the following short questions so we can improve future assignments.

1. How much time did you spend on each part of this assignment?
2. Any additional feedback?

6 Turning it in

1. **Submitting the PDF.** Make a PDF report containing: Responses for parts 1.1 and 1.2, Table 1 for a table of results from part 3.1, Figure 1 for part 3.2, Figure 2 with results from question part 4.2, and responses for parts 5.1 and 5.2.

See the handout at

<http://rail.eecs.berkeley.edu/deeprlcourse/static/misc/viz.pdf>

for notes on how to generate plots.

2. **Submitting the code and experiment runs.** In order to turn in your code and experiment logs, create a folder that contains the following:
 - A folder named `run_logs` with experiments for both the behavioral cloning (part 2, not part 3) exercise and the DAGGER exercise. Note that you can include multiple runs per exercise if you'd like, but you must include at least one run (of any task/environment) per exercise. These folders can be copied directly from the `cs285/data` folder into this new folder. **Important: Disable video logging for the runs that you submit, otherwise the files size will be too large! You can do this by setting the flag `-video_log_freq -1`**
 - The `cs285` folder with all the `.py` files, with the same names and directory structure as the original homework repository. Also include the commands (with clear hyperparameters) that we need in order to run the code and produce the numbers that are in your figures/tables (e.g. `run "python run_hw1_behavior_cloning.py --ep_len 200"` to generate the numbers for Section 2 Question 2) in the form of a README file.

As an example, the unzipped version of your submission should result in the following file structure. Make sure that the `submit.zip` file is below 15MB and that they include the prefix `q1_` and `q2_`.

```
submit.zip
├── run_logs
│   ├── q1_bc_ant
│   │   └── events.out.tfevents.1567529456.e3a096ac8ff4
│   ├── q1_dagger_ant
│   │   └── events.out.tfevents.1567529456.e3a096ac8ff4
│   └── ...
├── cs285
│   ├── scripts
│   │   ├── run_hw1.py
│   │   └── run_hw1.ipynb
│   ├── policies
│   │   └── ...
│   └── ...
├── README.md
└── ...
```

3. If you are a Mac user, **do not use the default “Compress” option to create the zip**. It creates artifacts that the autograder does not like. You may use `zip -vr submit.zip submit -x "*.DS_Store"` from your terminal.
4. Turn in your assignment on Gradescope. Upload the zip file with your code and log files to **HW1 Code**, and upload the PDF of your report to **HW1**.