

OFFLINE IMITATION LEARNING BY CONTROLLING THE EFFECTIVE PLANNING HORIZON

Anonymous authors

Paper under double-blind review

ABSTRACT

In offline imitation learning (IL), we generally assume only a handful of expert trajectories and a supplementary offline dataset from suboptimal behaviors to learn the expert policy. While it is now common to minimize the divergence between state-action visitation distributions so that the agent also considers the future consequences of an action, a sampling error in an offline dataset may lead to erroneous estimates of state-action visitations in the offline case. In this paper, we investigate the effect of controlling the effective planning horizon (i.e., reducing the discount factor) as opposed to imposing an explicit regularizer, as previously studied. Unfortunately, it turns out that the existing algorithms suffer from magnified approximation errors when the effective planning horizon is shortened, which results in a significant degradation in performance. We analyze the main cause of the problem and provide the right remedies to correct the algorithm. We show that the corrected algorithm improves on popular imitation learning benchmarks by controlling the effective planning horizon rather than an explicit regularization.

1 INTRODUCTION

Imitation learning (IL) is one of the sequential decision making problem settings that aims to solve a task from expert demonstrations instead of explicit notions of utility (Pomerleau, 1991; Ng et al., 2000). In a standard setting of IL, the agent is only given with a few number of expert trajectories, and it performs the imitation of the expert by interacting with the environment. We denote this setting as online imitation learning since it is possible to query the consequence of action during the learning with the online interactions. On the other hand, it has been recently questioned about the practicality of online interactions, and a setting called offline imitation learning that aims to learn from supplementary suboptimal dataset instead of interactions was proposed (Kim et al., 2021; 2022; Ma et al., 2022).

As a basic approach for IL, one can come up with behavior cloning (BC) (Pomerleau, 1991), which treats the problem as a supervised learning to map state to action based on expert demonstrations. However, BC does not make a use of given information other than expert demonstrations, and suffers from *compounding error* by covariate shift as it gets off track (Ross et al., 2011). To address this issue, previous studies (Ho & Ermon, 2016; Kostrikov et al., 2018; 2019) have proposed to imitate an expert policy by matching *state-action visitation distributions*, which can be estimated by interacting with the environment. State-action visitation matching allows the agent to learn how to return to the preferred states, and this approach has been successful in a wide range of domains of online IL.

Interestingly, in recent offline IL studies, it is often found that BC is quite competitive (Li et al., 2022; Kim et al., 2021), and sometimes it even performs better than visitation distribution matching algorithms depending on the experiment settings. This is because there is an inherent error in estimating visitation distributions from a finite dataset in offline cases, and BC, which completely ignores the dynamics information, does not suffer from this problem. Therefore, unlike the online cases where we can sample from true environment dynamics, BC or visitation distribution matching cannot always be the superior choice over the other in offline cases.

Motivated by this, we consider controlling the effective planning horizon, i.e. the discount factor γ , for the offline imitation learning problems. The discount factor implies how important the future is compared to the present. For visitation distribution matching algorithms, it can be interpreted as the amount of state-action distribution difference between the learned and expert policy we allow in the

future compared to now. While it has been overlooked in the online IL studies since we recover more accurate policy as we less discount the future, it will be advantageous to use an optimal discount factor in offline IL that makes the best trade-offs: using small γ to shorten the effective planning horizon helps to become robust to the errors in the inferred dynamics (Janner et al., 2019) (which also allows avoiding explicit regularization for robustness as in previous studies), whereas using large γ makes an agent less prone to the compounding errors by considering longer consequences of training.

In this paper, we start by formally analyzing the performance trade-off of a discount factor in offline IL. It turns out that, however, recently proposed offline IL algorithms show pathological behavior when the discount factor is lowered and naively controlling the discount factor does not lead to a performance gain. It is because the error from the approximations they share depends on the discount factor, and it becomes revealed and maximized as we lower γ . To this end, we propose a simple technique called Inverse Geometric Initial state sampling (IGI) to address the problem. We show that IGI enables the algorithm to properly learn in the low discount factor settings, and by tuning the discount factor, the offline IL agent with IGI outperforms previous state-of-the-art algorithms with explicit regularizations.

2 BACKGROUND

2.1 MARKOV DECISION PROCESS (MDP)

We consider environments modeled as a Markov Decision Process (MDP), which is defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, p_0, R, \gamma)$. Here, \mathcal{S} is the state space, and \mathcal{A} is the action space. P and p_0 represent the dynamics and initial state distribution, respectively. R is the reward function which is assumed to be bounded in $[0, R_{\max}]$, and $\gamma \in [0, 1)$ is the discount factor. A policy $\pi(\cdot|s)$ determines the probability of agent’s action in a state s . The goal of an agent is to maximize the sum of reward discounted by γ . The state-action visitation distribution $d^\pi(s, a)$, which is induced by a policy π , is defined as

$$d^\pi(s, a) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t Pr(s_t = s, a_t = a | s_0 \sim p_0(\cdot), s_t \sim P(\cdot|s_{t-1}, a_{t-1}), a_t \sim \pi(\cdot|s_t)).$$

We consider an offline imitation learning problem where a small number of expert demonstrations (\mathcal{D}^E) and a suboptimal dataset (\mathcal{D}^O) consisting of transitions (s, a, s') are given. We denote total dataset $\mathcal{D}^D = \mathcal{D}^E \cup \mathcal{D}^O$ as a union of two datasets. Especially, we assume that \mathcal{D}^E is sampled by following the expert policy π^E in underlying MDP \mathcal{M} . We denote empirical distribution of \mathcal{D}^E and \mathcal{D}^D as $E(s, a)$ and $D(s, a)$ respectively. On the other hand, we denote the state-action distribution of π^E as $d^E(s, a)$.

2.2 IMITATION LEARNING VIA STATE-ACTION VISITATION MATCHING

Imitation learning (IL) aims to train an agent that mimics the expert based on the expert demonstrations. In addition to demonstrations, supplementary information on environment dynamics is given, by directly interacting with environment (online IL) or by a dataset of suboptimal behaviors (offline IL). One simplest and most popular approach to imitate an expert is behavior cloning (BC), which treats IL as a supervised learning problem, ignoring any of supplementary information (Pomerleau, 1991). The objective of BC can be represented as follows:

$$\min_{\pi} -\mathbb{E}_{(s,a) \sim E} [\log \pi(a|s)]. \quad (1)$$

However, it is well known that the error of obtained policy induced by covariate shift compounds over time, and leads to an eventual failure unless \mathcal{D}^E is large enough (Ross et al., 2011).

To address the weakness of BC, a state-action visitation distribution matching objective is now widely adopted, which minimizes the divergence between d^π and d^E (Kostrikov et al., 2019; Ke et al., 2020). The distribution matching objective can be presented as follows:

$$\max_{\pi} -D_{\text{KL}}(d^\pi \| d^E) = \mathbb{E}_{(s,a) \sim d^\pi} \left[\log \frac{d^E(s, a)}{d^\pi(s, a)} \right]. \quad (2)$$

Note that the objective can be alternatively interpreted as a reinforcement learning (RL) problem:

$$\max_{\pi} -D_{\text{KL}}(d^{\pi} \| d^E) = (1 - \gamma) \cdot \mathbb{E}_{\substack{s_0 \sim p_0, \\ a_t \sim \pi, \\ s_{t+1} \sim P}} \left[\sum_{t=0}^{\infty} \gamma^t \log \frac{d^E(s_t, a_t)}{d^{\pi}(s_t, a_t)} \right]. \quad (3)$$

We can observe that the objective (3) can be seen as a RL problem with a reward $r = \log \frac{d^E}{d^{\pi}}$. When online interactions are allowed, we can sample from d^{π} by executing policy π , and Ho & Ermon (2016) proposed to estimate the newly defined reward by learning the discriminator using samples from d^E and d^{π} :

$$\max_{c: S \times A \rightarrow (0,1)} \mathbb{E}_{(s,a) \sim d^E} [\log c(s, a)] + \mathbb{E}_{(s,a) \sim d^{\pi}} [\log (1 - c(s, a))], \quad (4)$$

where the optimal discriminator $c^*(s, a)$ can be used to recover the reward as $\log c^*(s, a) - \log (1 - c^*(s, a)) = \log \frac{d^E(s,a)}{d^{\pi}(s,a)}$. On the contrary, in a fully offline setting we consider, the interaction with the environment to receive samples from d^{π} is not allowed. A common choice in this case is to train a discriminator that can distinguish between the expert demonstrations and a supplementary suboptimal dataset (Kim et al., 2021; 2022; Ma et al., 2022), and is explained in detail in Section 3.1.

3 OFFLINE IL VIA STATE-ACTION VISITATION MATCHING

In this section, we first derive a practical objective that can be used to optimize the distribution matching objective (2) in the offline IL case. After the derivation of an objective that is used throughout the paper, we provide an analysis of the effect of the discount factor on the error bound of offline IL. In particular, we show that there is a trade-off when controlling the discount factor, unlike the online case.

3.1 DERIVATION OF AN OFFLINE IL OBJECTIVE

We start from the widely used visitation distribution matching objective (2) as suggested in (Ho & Ermon, 2016; Torabi et al., 2018). Rather than optimizing for the policy π , we follow the recent approaches that optimize directly for the visitation distribution d^{π} (Lee et al., 2021; Kim et al., 2021). To optimize for d^{π} , we need to ensure that the solution we get is a valid visitation distribution. By making Bellman flow constraints and normalization constraints explicit, we have:

$$\max_{d^{\pi}} -D_{\text{KL}}(d^{\pi} \| d^E) \quad (5)$$

$$\text{s.t } \mathcal{B}_* d^{\pi}(s, a) = (1 - \gamma)p_0(s) + \gamma \mathcal{P}_* d^{\pi}(s, a) \quad \forall s, \quad (6)$$

$$\sum_{s,a} d^{\pi}(s, a) = 1, \quad d^{\pi}(s, a) \geq 0 \quad \forall s, a, \quad (7)$$

where $\mathcal{B}_* d^{\pi}(s, a) = \sum_a d^{\pi}(s, a)$ is a marginalization and $\mathcal{P}_* d^{\pi}(s, a) = \sum_{\bar{s}, \bar{a}} P(s|\bar{s}, \bar{a}) d^{\pi}(\bar{s}, \bar{a})$ is an expectation over the previous state-actions. It can be easily seen that the following equalities hold by interchanging the order of summations:

$$\sum_s \nu(s) \mathcal{P}_* d^{\pi}(s, a) = \sum_{s,a} d^{\pi}(s, a) \mathcal{P} \nu(s), \quad \text{where } \mathcal{P} \nu(s) = \sum_{s'} P(s'|s, a) \nu(s')$$

$$\sum_s \nu(s) \mathcal{B}_* d^{\pi}(s, a) = \sum_{s,a} d^{\pi}(s, a) \mathcal{B} \nu(s), \quad \text{where } \mathcal{B} \nu(s) = \nu(s).$$

The Lagrangian of the constrained problem (5-7) is

$$\begin{aligned} \max_{d^{\pi} \geq 0} \min_{\nu, \lambda} & -\mathbb{E}_{s,a \sim d^{\pi}} \left[\log \frac{d^{\pi}(s, a)}{d^E(s, a)} \right] + \lambda \left[\sum_{s,a} d^{\pi}(s, a) - 1 \right] \\ & + \sum_s \nu(s) ((1 - \gamma)p_0(s) + \gamma \mathcal{P}_* d^{\pi}(s, a) - \mathcal{B}_* d^{\pi}(s, a)) \end{aligned} \quad (8)$$

where λ and $\nu(s)$ are Lagrange multipliers. Based on the relationships shown above, we can rewrite Equation (8), where the summation is computed w.r.t. $d^\pi(s, a)$:

$$\max_{d^\pi \geq 0} \min_{\nu, \lambda} (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] + \sum_{s, a} d^\pi(s, a) \left[\gamma \mathcal{P}\nu(s) - \mathcal{B}\nu(s) + \lambda - \log \frac{d^\pi(s, a)}{d^E(s, a)} \right] - \lambda. \quad (9)$$

In the second term of Equation (9), to derive a practical algorithm, we want to make an expectation to be performed on the total dataset distribution $D(s, a)$. For this, we use the importance sampling with importance weight $\zeta(s, a) = \frac{d^\pi(s, a)}{D(s, a)}$. Then, the objective becomes

$$\max_{\zeta \geq 0} \min_{\nu, \lambda} \mathcal{L}(\zeta, \nu, \lambda) = (1 - \gamma) \mathbb{E}_{p_0} [\nu(s)] + \mathbb{E}_D [\zeta(s, a) (A_\nu(s, a) + \lambda - \log \zeta(s, a))] - \lambda. \quad (10)$$

where $A_\nu(s, a) = \gamma \mathcal{P}\nu(s) - \mathcal{B}\nu(s) + \log \frac{d^E(s, a)}{D(s, a)}$, which can be interpreted as an advantage when the action a is executed in the state s if we define the reward as a log ratio of distributions $r(s, a) = \log \frac{d^E(s, a)}{D(s, a)}$. Such a log ratio can be obtained by training a discriminator $c(s, a)$ through the following objective:

$$\max_{c: S \times A \rightarrow [0, 1]} \mathbb{E}_{(s, a) \sim d^E} [\log c(s, a)] + \mathbb{E}_{(s, a) \sim D} [\log(1 - c(s, a))], \quad (11)$$

such that the optimal discriminator $c^*(s, a)$ satisfies $\log \frac{d^E(s, a)}{D(s, a)} = \log \left(\frac{c^*(s, a)}{1 - c^*(s, a)} \right)$. The advantage $A_\nu(s, a)$ can then be estimated based on the discriminator c being trained.

We can further simplify the optimization problem (10) by noting that the strong duality holds, which enables us to change the order of optimization without affecting the optimal value. Closed-form solution of maximization of Equation (10) can be obtained by solving first-order optimality condition $\frac{\partial \mathcal{L}(\zeta^*, \nu, \lambda)}{\partial \zeta} = 0$, which gives $\zeta^* = \exp(A_\nu(s, a) + \lambda - 1)$. Similarly, $\lambda^* = -\log[\mathbb{E}_{(s, a) \sim D} \exp(A_\nu(s, a) - 1)]$ can be obtained. By substituting in ζ^* and λ^* to (10), we derive the final objective:

$$\min_{\nu} (1 - \gamma) \mathbb{E}_{p_0} [\nu(s)] + \log \mathbb{E}_{(s, a) \sim D} [\exp(A_\nu(s, a))] =: \mathcal{L}(\zeta^*, \nu, \lambda^*), \quad (12)$$

and optimize for ν . Based on the optimized ν^* , we can obtain $\zeta^* = \text{softmax}(A_{\nu^*}(s, a) - 1)$. Using this, policy can be extracted through the following objective:

$$\min_{\pi} \mathbb{E}_{(s, a) \sim d^{\pi^*}} [\log \pi(a|s)] = \mathbb{E}_{(s, a) \sim D} \left[\frac{d^{\pi^*}(s, a)}{D(s, a)} \log \pi(a|s) \right] = \mathbb{E}_{(s, a) \sim D} [\zeta^* \log \pi(a|s)]. \quad (13)$$

Note that the derived objective (12) can be seen as a special case of DemoDICE (Kim et al., 2021) and SMODICE (Ma et al., 2022), and it can be recovered by setting $\alpha = 0$ in DemoDICE (no explicit regularization toward suboptimal dataset) or by setting f -divergence to be KL divergence in SMODICE. While the algorithm in this paper is based on KL-divergence minimization of visitation distributions, note that it can be trivially extended to any f -divergence minimization case.

3.2 TRADE-OFF BETWEEN TWO DISTINCT EFFECTS IN OFFLINE IL BY DISCOUNT FACTOR

Now we show that there is indeed a trade-off when controlling the discount factor γ in the offline setting. Intuitively, an offline agent can only receive supplementary information about environment dynamics through limited finite demonstrations, and the error of estimated dynamics is inevitable. Therefore, we can expect that learning with a long planning horizon increases the risk of compounding estimation error. On the other hand, the trained agent is typically evaluated in a non-discounted environment by measuring the average reward it gets, and excessively lowering the discount factor for training will make a large train-test discrepancy. The theorem below is a formal analysis backing up the argument, showing an error bound of the imitated policy with respect to a discount factor.

Theorem 1. *Let P an underlying transition dynamics and \hat{P} an estimated transition dynamics. γ is a discount factor used for evaluating the policy and $\hat{\gamma}$ is a discount factor used for training the policy where $\hat{\gamma} \leq \gamma$. Let $d_{\hat{P}, \hat{\gamma}}^\pi$ a state-action visitation distribution induced by π under the dynamics*

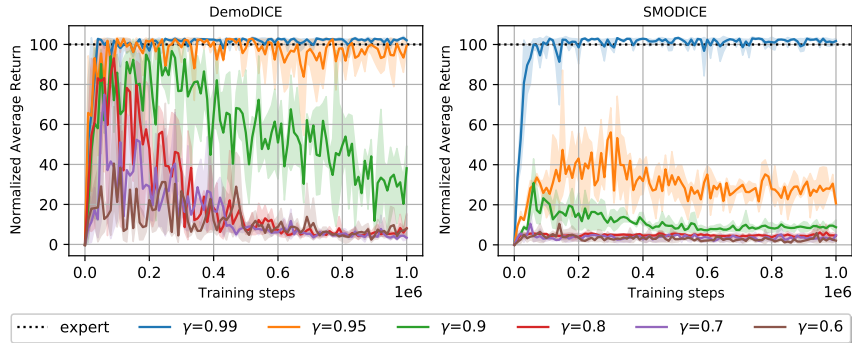


Figure 1: Learning curves of DemoDICE (Kim et al., 2021) and SMODICE (Ma et al., 2022) over various choices of discount factor γ . We use \mathcal{D}^O consisting of 400 expert trajectories and 800 random-policy trajectories, and \mathcal{D}^E consisting of 1 expert trajectory. Evaluations are averaged over 3 seeds in a Hopper-v2 environment, and they are normalized so that 0 corresponds to the average score of the random-policy dataset, and 100 corresponds to the average score of the expert policy dataset.

P using γ , and $d_{P,\gamma}^E$ a state-action visitation distribution induced by π^E under the dynamics P using γ . Assume that the reward is bounded in $[0, R_{\max}]$. Then, the error bound for imitated policy π is

$$\left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^E} [r(s, a)] \right| \leq \frac{2R_{\max}}{1 - \hat{\gamma}} \left(\frac{\gamma - \hat{\gamma}}{1 - \gamma} + \frac{\hat{\gamma}\epsilon_P}{1 - \hat{\gamma}} + \frac{\epsilon_\pi}{1 - \hat{\gamma}} \right), \quad (14)$$

where $\epsilon_P = \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [D_{TV}(\hat{P}||P)] + \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^E} [D_{TV}(\hat{P}||P)]$, and $\epsilon_\pi = \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [D_{TV}(\pi||\pi^E)]$.

The proof is deferred to Appendix A. The first term reflects the train-test discrepancy caused by optimizing policy using $\hat{\gamma}$ instead of the discount factor γ that is used for evaluation, i.e. in the LHS of inequality (14), and is minimized as we choose $\hat{\gamma}$ close to γ . The second term represents the effect of the error on estimating dynamics ϵ_P . It can be observed that this term becomes smaller if the effective planning horizon is shortened by using a lower $\hat{\gamma}$. From the definition of ϵ_P , it is the accumulation of step-wise dynamics estimation error over trajectories of both π and π^E , and there will always be an irreducible amount of ϵ_P in the offline case. The last term stands for the policy difference between π and π^E over the estimated dynamics and will be minimized if imitation learning is done properly to maximize the objective (2).

In short, the error of the imitated policy will be mainly dependent on the first two terms of (14), which have opposite dependencies to the discount factor for training $\hat{\gamma}$. This result supports our argument about the trade-off between two distinct effects by discount factor during the offline training and shows the possibility that an offline IL algorithm may benefit from choosing $\hat{\gamma}$ that is smaller than γ that we use to evaluate.

4 CONTROLLING THE DISCOUNT FACTOR IN OFFLINE IMITATION LEARNING

As suggested by Theorem 1, we expect that there would be a benefit from controlling the discount factor γ that is used for training. The algorithm derivation is based on a visitation distribution matching (5-12), as well as those of previous studies, which appears independent of and applicable to any choice of the discount factor in principle. However, it turns out that naively lowering γ in previously suggested visitation distribution matching offline IL algorithms (Kim et al., 2021; Ma et al., 2022) only results in a monotonically decreasing performance, as shown in Figure 1. The performance reduction of using lower discount factors is significant, and it seems pointless to choose any γ below 0.99 in contrast to our analysis. In this section, we analyze the cause of this pathological behavior and propose a simple scheme that can alleviate it.




Figure 2: Toy infinite horizon MDP example with 3 states and 2 actions. All transitions are deterministic and shown with the arrows. s_0 is initial state, and g is absorbing state. We indicate the probability of taking an action based on the corresponding policy on the arrows in the figure. (a) represents expert policy and (b) is for learned policy with parameter θ .

4.1 DISTRIBUTION MISMATCH IN TRAINING A DISCRIMINATOR

The main problem lies in the fact that while the visitation distribution matching objective (5) requires the estimation of $d^E(s, a)$, which is a discounted state-action distribution induced by the expert policy π^E , a discriminator $c(s, a)$ is learned to discriminate the empirical distributions of the datasets, i.e. between $D(s, a)$ and $E(s, a)$. Since the empirical distributions are undiscounted, for a state-action pair sample (s_t, a_t) at timestep t , $d^E(s, a)$ would weight γ^t on this sample compared to $E(s, a)$ that does not weight on this sample, making the two distributions significantly different as we have smaller γ .

To observe the consequence of having a discriminator trained on empirical distributions, we assume that the reward is represented with an optimal discriminator such that $r(s, a) = \log \frac{E(s, a)}{D(s, a)}$. **As a result of applying this to objective (10), it becomes:**

$$\begin{aligned} \max_{d^\pi \geq 0} \min_{\nu, \lambda} \mathcal{L}(d^\pi, \nu, \lambda) &= (1 - \gamma) \mathbb{E}_{s \sim p_0} [\nu(s)] \\ &+ \sum_{s, a} d^\pi(s, a) \left[\gamma \mathcal{P} \nu(s) - \mathcal{B} \nu(s) + \log \frac{E(s, a)}{D(s, a)} - \log \frac{d^\pi(s, a)}{D(s, a)} + \lambda \right] - \lambda. \end{aligned} \quad (15)$$

By taking the derivation steps backward, we can confirm that using the discriminator trained on empirical distributions is equivalent to solving the visitation distribution matching objective (5-7) except for the objective:

$$\max_{d^\pi} -D_{KL}(d^\pi(s, a) \| E(s, a)). \quad (16)$$

That is, by training a discriminator trained on empirical distributions, we actually have matched $d^\pi(s, a)$ and $E(s, a)$ to get a policy π . Note that there is a one-to-one correspondence between the state-action visitation distributions and policies (Ho & Ermon, 2016). It implies that unless $E(s, a) = d^E(s, a)$, even if $E(s, a)$ is a valid visitation distribution that satisfies Bellman flow constraints, the policy inferred by matching $E(s, a)$ will be different from d^E , and thus, $\pi^* \neq \pi_E$. In most cases, $E(s, a)$ would not be a valid visitation distribution, and a policy π that gives $D_{KL}(d^\pi(s, a) \| E(s, a)) = 0$ would not exist in general. While this discrepancy between E and d^E is negligible when γ is large and the length of trajectories in \mathcal{D}^E is long enough, large error on the imitated policy can be incurred in other cases. We give a simple toy example below.

Illustrative example Assume an infinite horizon MDP with 2 states and an absorbing state, and with 2 actions as shown in Figure 2. The expert policy π_E has 0.6 probability of doing a_1 in s_0 , 0.4 of doing a_2 in s_0 , and 1 of doing a_1 in the other states. Assuming that trajectories are sampled up to timestep 2, the expert demonstration \mathcal{D}^E will be consisting of two kinds of trajectories $\tau_1 = (s_0, a_2, g, a_1, g, a_1)$ and $\tau_2 = (s_0, a_1, s_1, a_1, g, a_1)$ with the ratio of 6 : 4. $E(s, a)$ can be calculated by $\frac{\text{The number of (s, a) pairs}}{\text{Total number of state-action pairs}}$ accordingly. The learned policy is parametrized as $\pi(a_1|s_0) = \theta$, $\pi(a_2|s_0) = 1 - \theta$, and $\pi(a_1|\cdot) = 1$ in other states. By computing d^π based on the above, we can express $D_{KL}(d^\pi \| E)$ in terms of θ and γ . Then, we can minimize D_{KL} to get θ^* , and see if we can recover $\theta^* = 0.6$. In this toy example, it turns out that $\theta \neq 0.6$ unless $\gamma = 0.5$ (detailed




Figure 3: Learning curves of our method in `HalfCheetah-v2`. Here, suboptimal dataset is consisted of 100 expert trajectories and 800 random-policy trajectories. We applied moving average with 3 seeds. The plot on the left is the result of sampling $\tilde{p}_0(s|t)$ from expert dataset, and the right is the result of sampling from total dataset.

derivations can be found in Appendix B). It can be seen in this example that we cannot match two distributions $d^\pi(s, a) = d^\pi(s)\pi(a|s)$, $E(s, a) = E(s)\pi^E(a|s)$ in general, and by minimizing the distribution between state visitation distributions, optimized π will be different to π_E .

4.2 INVERSE GEOMETRIC INITIAL STATE SAMPLING

As shown in the previous subsection, we cannot recover $\pi = \pi^E$ unless $d^E = E$, where E is the distribution that the discriminator is trained on. One straightforward solution to this problem is to train the discriminator to distinguish d^E and E , eliminating the root cause of the problem. We can simulate the sampling from d^E by weighting each sample in the dataset with respect to their timesteps when they were sampled. For example, we can first sample the timestep t from a geometric distribution $\text{Geom}(1 - \gamma)$ and then sampling (s_t, a_t) pair that had been sampled at t . This procedure will approximate the sampling from d^E sufficiently well, given that the trajectories stored in \mathcal{D}^E is long enough.

However, training the discriminator to distinguish d^E gives rise to another problem; if we use lower γ , the samples from d^E will mostly consist of early timestep samples of \mathcal{D}^E , since the probability assigned decreases exponentially over timesteps. This results in a significant under-usage of dataset \mathcal{D}^E by throwing away samples of later timesteps, and significantly deteriorates the performance. We need a way to ensure that we get $\pi = \pi^E$ while not hurting the effective number of data.

To this end, we propose to devise a different initial state distribution other than the actual p_0 to satisfy both conditions. Note that even if we change the initial distribution, the optimality of policies are not affected as mentioned in Kostrikov et al. (2019). d^π and d^E that π and π_E induces will be different, but one-to-one correspondence does still hold that matching visitation distributions ensure $\pi = \pi_E$ with arbitrary initial state distribution.

In particular, if we use an initial distribution that makes $d^E = E$, i.e. flattens the visitation distribution to look like an undiscounted distribution, recovering $\pi = \pi_E$ will be guaranteed even if we use the discriminator that learns from the empirical distributions. Note that sampling from d^E corresponds to sample timestep t from $\text{Geom}(1 - \gamma)$ and choose (s_t, a_t) that had been sampled at t . Hence, if we use a modified initial distribution \tilde{p}_0 that samples from all the timesteps in \mathcal{D}^E with weights inversely proportional to $\text{Geom}(1 - \gamma)$, the resultant \tilde{d}^E will have uniform weights regarding timesteps and therefore $\tilde{d}^E = E$. We call this modified initial state sampling method as Inverse Geometric Initial state sampling (IGI). In practice, there are a few more things to consider determining the weightings for each possible initial state samples; e.g. due to datasets being truncated at certain timesteps and trajectories that are terminated early. Nevertheless, it is possible to set up a system of linear equations to obtain the required weights for each sample to ensure $\tilde{d}^E = E$. Details on how to set up a linear system is shown in Appendix C.

Unfortunately, it turns out that the exact IGI that ensures $\tilde{d}^E = E$ can be too restrictive. Following the recent practices of using a single expert policy for \mathcal{D}^E in evaluating imitation performances, only the state-action pairs in the single trajectory can be sampled as an initial state, which cripples the

diversity. Furthermore, we found that sampling the initial states only from \mathcal{D}^E also limited the opportunity of our agent to learn how to act outside \mathcal{D}^E , incurring the covariate shift and compounding error problems that BC has suffered. On the other hand, using IGI to sample from the total dataset \mathcal{D}^D did not exhibit such a problem, while alleviating the discrepancy between \tilde{d}^E and E . In Figure 3, we compare the performance of two options: IGI(\mathcal{D}^E) and IGI(\mathcal{D}^D). It is clearly visible that the imitation performance of IGI(\mathcal{D}^E) is limited, whereas IGI(\mathcal{D}^D) performs robustly over a wide set of γ s. As a consequence, we used IGI with samples from \mathcal{D}^D in the following experiments.

5 RELATED WORK

Controlling the effective planning horizon In reinforcement learning (RL), the discount factor effectively controls the amount of forward planning the agent considers making a decision. Petrik & Scherrer (2008) has shown that the approximation error bound incurred in an inaccurate model can be tightened by using a low discount factor. Jiang et al. (2015) has shown the similar result, but they have analyzed in terms of the complexity of the class of policies. They show that there is a trade-off between the complexity of the policy space and the error due to the approximated model according to the effective planning horizon, which can be controlled by a discount factor. Recently, Hu et al. (2022) analyzed the role of discount factor from the perspective of offline RL. In addition to the trade-off relationship between optimality and sample efficiency, they show theoretically and empirically that the low discount factor can also be seen as a model-based pessimism. On the other hand, there has not been a study on the effect of discount factor on offline imitation learning, up to our knowledge.

Offline imitation learning by leveraging the duality of RL Imitation learning (IL) aims to mimic the expert policy by using the expert demonstrations and the online interactions, but offline supplementary dataset is given instead in offline case. Recently, there have been a large amount of literature that leverages the duality in RL to develop a novel algorithm estimating quantities related to visitation distributions. In a field of offline IL, DemoDICE (Kim et al., 2021) proposed to solve it by deriving the dual of visitation distribution matching problem with an explicit regularizer minimizing the f -divergence between the visitation distribution induced by policy and the supplement dataset distribution. SMODICE (Ma et al., 2022) considers the IfO problem, which solves IL with state-only demonstrations by expert, and proposes a versatile offline IL algorithm by leveraging f -divergence and Fenchel duality instead of KL used in (Kim et al., 2021).

6 EXPERIMENTS

We evaluate the performance of our algorithm in both discrete and continuous MDP. We first show the trade-off effect by the discount factor in finite-discrete MDP, then evaluate the imitation performance of our algorithm using IGI in continuous MDP.

6.1 FINITE AND DISCRETE MDP

In finite-discrete MDP, we empirically show the trade-off effect by the discount factor in the Random MDP. This environment generates the finite and discrete MDP randomly. We follow the environment configuration of Lee et al. (2021). In this experiment setting, we can see that the performance can be optimized by choosing a lower discount factor. Details of the experiment in finite-discrete MDP and its results with various settings are shown in Appendix E.1 due to the lack of space.

6.2 CONTINUOUS MDP

For continuous MDP, we evaluate the imitation performance of our algorithm using IGI in the MuJoCo continuous control environment (Todorov et al., 2012): HalfCheetah-v2, Hopper-v2, Walker2d-v2 and Ant-v2. The dataset used for learning is obtained from the D4RL datasets (Fu et al., 2020). We sampled expert trajectories and random trajectories from the expert dataset and random dataset of D4RL benchmark respectively. We use a single expert trajectory for \mathcal{D}^E , and our suboptimal dataset \mathcal{D}^O is constructed as a union of expert trajectories and random-policy trajectories with specific ratios.

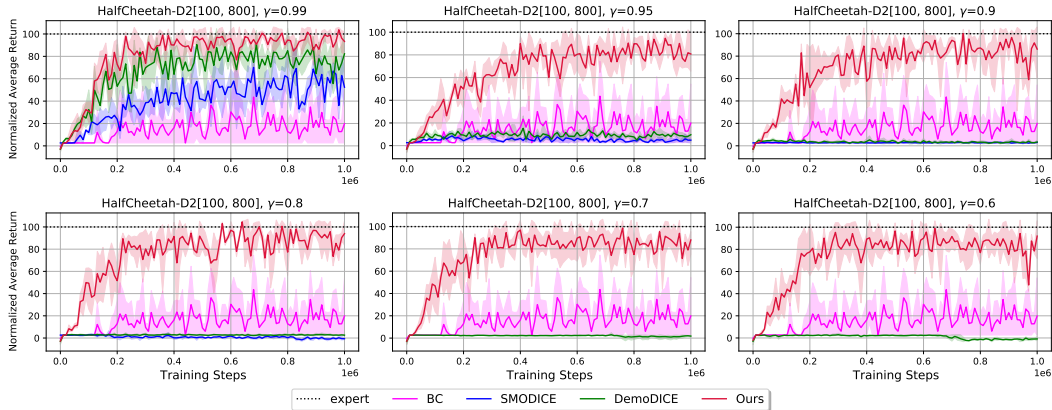


Figure 4: Performance of our algorithm (IGI) and baselines according to different γ s on HalfCheetah-v2 with $D2$ dataset. Here, shaded area shows the standard error of the normalized evaluation over 3 seeds.

We conduct experiments by changing the ratio between the expert data and the random-policy data of the suboptimal dataset. We call suboptimal dataset consisting of 400 expert trajectories and 800 random-policy trajectories as D1, 100:800 as D2, 50:800 as D3. We compare our algorithm using IGI with other offline IL algorithms, DemoDICE (Kim et al., 2021), SMODICE (Ma et al., 2022), and BC. Performance of other algorithms is measured based on the official hyperparameters made public by authors without any modification. For BC, we show the result of learning with \mathcal{D}^D , which has the best performance among the various combination of the data to be cloned. All experiment results in this paper are averaged over 3 seeds and normalized so that 0 corresponds to the average score of the random-policy dataset, and 100 corresponds to the average score of the expert policy dataset.

In order to demonstrate the robustness of our algorithm over different γ s, we show the results by applying various discount factor. We use discount factor for $\gamma \in \{0.99, 0.95, 0.9, 0.8, 0.7, 0.6\}$. We also tested every algorithm with γ below 0.6, but due to exploding gradient, it was not possible to run DemoDICE and SMODICE for lower γ s below 0.6. Therefore, we report only the result of γ s mentioned above.

The results of baselines on different γ s are shown in Figure 4. It can be seen that the introduced method IGI makes the algorithm perform well regardless of γ . Additional results on other environments and D1, D3-ratio datasets are shown in the Appendix E.2. In addition, the experiment in Figure 4 has too large enough size of the dataset to make ϵ_P in the second term of Theorem 1 negligible, and the effect of controlling the discount factor is not clearly visible. Therefore, we make additional experiments by lowering the amount of the dataset about 10 times less (40 expert trajectories and 80 random-policy trajectories) than D1 setting to show the effect of controlling the discount factor clearly. Relevant results are shown in Appendix F.1. Also, the pseudocode and the hyperparameters used to demonstrate IGI can be found in Appendix D.1.

7 CONCLUSION

In this paper, we analyze the effect of controlling the discount factor on offline IL and motivate that the discount factor can take a role of a regularizer to prevent the sampling error of the supplementary dataset from hurting the performance. We show that the previously suggested imitation learning algorithms that utilize discriminators and a visitation distribution matching objective suffer from the discrepancy between the visitation distribution and the empirical distribution when low discount factor γ is applied. To this end, we proposed Inverse Geometric Initial state sampling (IGI), which uses the whole dataset with the weighting inversely proportional to the geometric distribution, to alleviate the problem that we cannot recover the expert policy π^E . We show that our approach shows stable and competitive performance regardless of the discount factor compared to other visitation distribution matching algorithms with explicit regularization.

REFERENCES

- Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- Hao Hu, Yiqin Yang, Qianchuan Zhao, and Chongjie Zhang. On the role of discount factor in offline reinforcement learning. *arXiv preprint arXiv:2206.03383*, 2022.
- Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. *Advances in Neural Information Processing Systems*, 32, 2019.
- Nan Jiang, Alex Kulesza, Satinder Singh, and Richard Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1181–1189. Citeseer, 2015.
- Liyiming Ke, Sanjiban Choudhury, Matt Barnes, Wen Sun, Gilwoo Lee, and Siddhartha Srinivasa. Imitation learning as f-divergence minimization. In *International Workshop on the Algorithmic Foundations of Robotics*, pp. 313–329. Springer, 2020.
- Geon-Hyeong Kim, Seokin Seo, Jongmin Lee, Wonseok Jeon, HyeongJoo Hwang, Hongseok Yang, and Kee-Eung Kim. Demodice: Offline imitation learning with supplementary imperfect demonstrations. In *International Conference on Learning Representations*, 2021.
- Geon-Hyeong Kim, Jongmin Lee, Youngsoo Jang, Hongseok Yang, and Kee-Eung Kim. Lobsdice: Offline imitation learning from observation via stationary distribution correction estimation. *arXiv preprint arXiv:2202.13536*, 2022.
- Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. *arXiv preprint arXiv:1809.02925*, 2018.
- Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. Imitation learning via off-policy distribution matching. *arXiv preprint arXiv:1912.05032*, 2019.
- Byungjun Lee, Jongmin Lee, Peter Vrancx, Dongho Kim, and Kee-Eung Kim. Batch reinforcement learning with hyperparameter gradients. In *International Conference on Machine Learning*, pp. 5725–5735. PMLR, 2020.
- Jongmin Lee, Wonseok Jeon, Byungjun Lee, Joelle Pineau, and Kee-Eung Kim. Optidice: Offline policy optimization via stationary distribution correction estimation. In *International Conference on Machine Learning*, pp. 6120–6130. PMLR, 2021.
- Ziniu Li, Tian Xu, Yang Yu, and Zhi-Quan Luo. Rethinking valuedice: Does it really improve performance? *arXiv preprint arXiv:2202.02468*, 2022.
- Yecheng Jason Ma, Andrew Shen, Dinesh Jayaraman, and Osbert Bastani. Smoalice: Versatile offline imitation learning via state occupancy matching. *arXiv preprint arXiv:2202.02433*, 2022.
- Andrew Y Ng, Stuart Russell, et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, pp. 2, 2000.
- Marek Petrik and Bruno Scherrer. Biasing approximate dynamic programming with a lower discount factor. *Advances in neural information processing systems*, 21, 2008.
- Dean A Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1):88–97, 1991.

- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033. IEEE, 2012.
- Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. *arXiv preprint arXiv:1807.06158*, 2018.
- Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *CoRR*, abs/1911.11361, 2019. URL <http://arxiv.org/abs/1911.11361>.
- Tian Xu, Ziniu Li, and Yang Yu. Error bounds of imitating policies and environments. *Advances in Neural Information Processing Systems*, 33:15737–15749, 2020.

A PROOF OF THEOREM 1

The proof of Theorem 1 is based on the analysis in previous works (Petrik & Scherrer, 2008; Jiang et al., 2015; Lee et al., 2020; Xu et al., 2020). We provide three Lemmas first and use them to prove the Theorem 1. To prove three Lemmas, we define several notations. \mathbf{P} and $\hat{\mathbf{P}}$ denote a matrix of the underlying transition dynamics and estimated transition dynamics, respectively. $\mathbb{E}_{d_{\hat{\mathbf{P}},\gamma}^\pi}[r(s,a)]$ denotes the expected total reward discounted by γ for the policy π under transition dynamics $\hat{\mathbf{P}}$. \mathbf{d}_0 and \mathbf{r} denote a vector of the initial state probability and the reward, respectively. For \mathbf{r} and \mathbf{P} , if the π is at superscript, it means ‘‘following policy π ’’.

Lemma 1. For any MDP with bounded rewards $|r(s,a)| \leq R_{\max}$, for all $\pi : S \rightarrow A$ and $\hat{\gamma} \leq \gamma$,

$$\left| \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] - \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] \right| \leq R_{\max} \frac{\gamma - \hat{\gamma}}{(1-\gamma)(1-\hat{\gamma})}.$$

Proof.

$$\begin{aligned} \left| \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] - \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] \right| &= \left| \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] - \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] \right| \\ &\leq \left\| \sum_{t=0}^{\infty} (\gamma^t - \hat{\gamma}^t) \mathbf{P}^{\pi t} \mathbf{r}^\pi \right\|_{\infty} \\ &\leq \sum_{t=0}^{\infty} (\gamma^t - \hat{\gamma}^t) R_{\max} \\ &= \left(\frac{1}{1-\gamma} - \frac{1}{1-\hat{\gamma}} \right) R_{\max} \\ &= \frac{\gamma - \hat{\gamma}}{(1-\gamma)(1-\hat{\gamma})} R_{\max} \end{aligned}$$

□

We define one more notation for Lemma 2 and 3. $\mathbf{d}_{\mathbf{P},\gamma}^\pi$ denotes a vector of the marginal state probability induced by π under transition dynamics \mathbf{P} using γ .

Lemma 2. We can bound the difference of the evaluations of policy π on two different MDPs with bounded rewards $|r(s,a)| \leq R_{\max}$ as

$$\left| \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] - \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] \right| \leq \frac{2\gamma R_{\max}}{(1-\gamma)^2} \mathbb{E}_{s \sim d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi, a \sim \pi} \left[\text{TV} \left(\hat{P}(s'|s,a) \| P(s'|s,a) \right) \right].$$

Proof.

$$\left| \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] - \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] \right| = \left| \mathbb{E}_{d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi}[r(s,a)] - \mathbb{E}_{d_{\mathbf{P},\gamma}^\pi}[r(s,a)] \right| \quad (17)$$

$$= \left| \mathbf{r}^\top \left[(I - \gamma \hat{\mathbf{P}})^{-1} - (I - \gamma \mathbf{P})^{-1} \right] \mathbf{d}_0 \right| \quad (18)$$

$$= \left| \mathbf{r}^\top (I - \gamma \mathbf{P})^{-1} \left[\gamma \hat{\mathbf{P}} - \gamma \mathbf{P} \right] (I - \gamma \hat{\mathbf{P}})^{-1} \mathbf{d}_0 \right| \quad (19)$$

$$= \gamma \left| \mathbf{r}^\top (I - \gamma \mathbf{P})^{-1} \left[\hat{\mathbf{P}} - \mathbf{P} \right] \frac{\mathbf{d}_{\hat{\mathbf{P}}}^\pi}{1-\hat{\gamma}} \right| \quad (20)$$

$$\leq \frac{\gamma}{1-\hat{\gamma}} \left\| \mathbf{r}^\top (I - \gamma \mathbf{P})^{-1} \right\|_{\infty} \left\| \left[\hat{\mathbf{P}} - \mathbf{P} \right] \mathbf{d}_{\hat{\mathbf{P}}}^\pi \right\|_1 \quad (21)$$

$$\leq \frac{\gamma R_{\max}}{(1-\hat{\gamma})^2} \left\| \left[\hat{\mathbf{P}} - \mathbf{P} \right] \mathbf{d}_{\hat{\mathbf{P}}}^\pi \right\|_1 \quad (22)$$

$$\leq \frac{\gamma R_{\max}}{(1-\hat{\gamma})^2} \sum_{s',a,s} \left| \hat{P}(s'|s,a) - P(s'|s,a) \right| \pi(a|s) d_{\hat{\mathbf{P}}}^\pi(s) \quad (23)$$

$$= \frac{2\gamma R_{\max}}{(1-\hat{\gamma})^2} \mathbb{E}_{s \sim d_{\hat{\mathbf{P}},\hat{\gamma}}^\pi, a \sim \pi} \left[\text{TV} \left(\hat{P}(s'|s,a) \| P(s'|s,a) \right) \right] \quad (24)$$

□

Lemma 3. We can bound the difference of the evaluations of two policies π and μ with bounded rewards $|r(s, a)| \leq R_{\max}$ as

$$\left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^\mu} [r(s, a)] \right| \leq \frac{2R_{\max}}{(1-\gamma)^2} \mathbb{E}_{s \sim d_P^\pi} [\mathbb{T}\mathbb{V}(\pi(a|s) \|\mu(a|s))].$$

Proof. We can expand the difference of the evaluations of two policies like

$$\begin{aligned} \left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^\mu} [r(s, a)] \right| &= \left| [(\mathbf{r}^\pi)^\top (I - \gamma \mathbf{P}^\pi)^{-1} - (\mathbf{r}^\mu)^\top (I - \gamma \mathbf{P}^\mu)^{-1}] \mathbf{d}_0 \right| \\ &\leq \left| (\mathbf{r}^\pi - \mathbf{r}^\mu)^\top \frac{\mathbf{d}_P^\pi}{1-\gamma} \right| + \left| (\mathbf{r}^\mu)^\top [(I - \gamma \mathbf{P}^\pi)^{-1} - (I - \gamma \mathbf{P}^\mu)^{-1}] \mathbf{d}_0 \right|. \end{aligned}$$

The first term is bounded as

$$\left| (\mathbf{r}^\pi - \mathbf{r}^\mu)^\top \frac{\mathbf{d}_P^\pi}{1-\gamma} \right| \leq \frac{2R_{\max}}{1-\gamma} \mathbb{E}_{s \sim d_{P,\gamma}^\pi} [\mathbb{T}\mathbb{V}(\pi(a|s) \|\mu(a|s))]. \quad (25)$$

Using the process (18-22), the second term is bounded as

$$\left| (\mathbf{r}^\mu)^\top [(I - \gamma \mathbf{P}^\pi)^{-1} - (I - \gamma \mathbf{P}^\mu)^{-1}] \mathbf{d}_0 \right| \leq \frac{\gamma R_{\max}}{(1-\gamma)^2} \|\mathbf{P}^\pi - \mathbf{P}^\mu\|_1 \mathbf{d}_P^\pi \quad (26)$$

$$= \frac{\gamma R_{\max}}{(1-\gamma)^2} \sum_{s'} \left| \sum_{s,a} P(s'|s, a) [\pi(a|s) - \mu(a|s)] d_P^\pi(s) \right| \quad (27)$$

$$\leq \frac{\gamma R_{\max}}{(1-\gamma)^2} \sum_{s,a} |\pi(a|s) - \mu(a|s)| d_P^\pi(s) \quad (28)$$

$$= \frac{2\gamma R_{\max}}{(1-\gamma)^2} \mathbb{E}_{s \sim d_{P,\gamma}^\pi} [\mathbb{T}\mathbb{V}(\pi(a|s) \|\mu(a|s))]. \quad (29)$$

Combining (25) and (29), we have the proof of Lemma 3. \square

Theorem 1. Let P an underlying transition dynamics and \hat{P} an estimated transition dynamics. γ is a discount factor used for evaluating the policy and $\hat{\gamma}$ is a discount factor used for training the policy where $\hat{\gamma} \leq \gamma$. Let $d_{P,\gamma}^\pi$ a state-action visitation distribution induced by π under the dynamics P using γ , and $d_{P,\gamma}^E$ a state-action visitation distribution induced by π^E under the dynamics P using γ . Assume that the reward is bounded in $[0, R_{\max}]$. Then, the error bound for imitated policy π is

$$\left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^E} [r(s, a)] \right| \leq \frac{2R_{\max}}{1-\hat{\gamma}} \left(\frac{\gamma - \hat{\gamma}}{1-\hat{\gamma}} + \frac{\hat{\gamma}\epsilon_P}{1-\hat{\gamma}} + \frac{\epsilon_\pi}{1-\hat{\gamma}} \right), \quad (14)$$

where $\epsilon_P = \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [D_{TV}(\hat{P} \| P)] + \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^E} [D_{TV}(\hat{P} \| P)]$, and $\epsilon_\pi = \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [D_{TV}(\pi \| \pi^E)]$.

Proof. The error bound for imitated policy π can be expanded as

$$\begin{aligned} &\left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^E} [r(s, a)] \right| \\ &\leq \left| \mathbb{E}_{d_{P,\gamma}^\pi} [r(s, a)] - \mathbb{E}_{d_{P,\hat{\gamma}}^\pi} [r(s, a)] \right| + \left| \mathbb{E}_{d_{P,\hat{\gamma}}^\pi} [r(s, a)] - \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [r(s, a)] \right| \\ &+ \left| \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^\pi} [r(s, a)] - \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^E} [r(s, a)] \right| + \left| \mathbb{E}_{d_{\hat{P},\hat{\gamma}}^E} [r(s, a)] - \mathbb{E}_{d_{P,\gamma}^E} [r(s, a)] \right| \\ &+ \left| \mathbb{E}_{d_{P,\gamma}^E} [r(s, a)] - \mathbb{E}_{d_{P,\hat{\gamma}}^E} [r(s, a)] \right|. \end{aligned}$$

The first and last terms are bounded by Lemma 1, and the second and fourth terms are bounded by Lemma 2. Lastly, the third term is bounded by Lemma 3. This completes the proof. \square

B PROOF FOR SIMPLE MDP

In Section 4.1, we define expert policy as

$$\pi^E(a_1|s_0) = 0.6, \pi^E(a_2|s_0) = 0.4, \pi^E(a_1|s_1) = 1, \pi^E(a_1|g) = 1.$$

we said if N trajectories are sampled, the number of τ_1 and τ_2 is $0.6N$ and $0.4N$. Since each trajectories have three (s, a) pairs, total number of (s, a) pairs is $3N$. Calculating the previously defined $E(s, a)$ is as follows

$$\begin{aligned} E(s_0, a_1) &= \frac{0.6N}{3N} = \frac{1}{5}, & E(s_0, a_2) &= \frac{0.4N}{3N} = \frac{2}{15} \\ E(s_1, a_1) &= \frac{0.6N}{3N} = \frac{1}{5}, & E(g, a_1) &= \frac{1.4N}{3N} = \frac{7}{15} \end{aligned}$$

Next, imitation policy is defined as

$$\pi(a_1|s_0) = \theta, \pi(a_2|s_0) = 1 - \theta, \pi(a_1|s_1) = 1, \pi(a_1|g) = 1.$$

Then, $d^\pi(s, a)$ has the following values

$$\begin{aligned} d^\pi(s_0, a_1) &= \pi(a_1|s_0) \left[(1 - \gamma)p_0(s_0) + \gamma \sum_{\bar{s}, \bar{a}} P(s_0|\bar{s}, \bar{a}) d^\pi(\bar{s}, \bar{a}) \right] = \theta(1 - \gamma) \\ d^\pi(s_0, a_2) &= \pi(a_2|s_0) \left[(1 - \gamma)p_0(s_0) + \gamma \sum_{\bar{s}, \bar{a}} P(s_0|\bar{s}, \bar{a}) d^\pi(\bar{s}, \bar{a}) \right] = (1 - \theta)(1 - \gamma) \\ d^\pi(s_1, a_1) &= \pi(a_1|s_1) \left[(1 - \gamma)p_0(s_1) + \gamma \sum_{\bar{s}, \bar{a}} P(s_1|\bar{s}, \bar{a}) d^\pi(\bar{s}, \bar{a}) \right] \\ &= 1 \cdot [(1 - \gamma) \cdot 0 + \gamma \cdot 1 \cdot d^\pi(s_0, a_1)] = \theta\gamma(1 - \gamma) \\ d^\pi(g, a_1) &= \pi(a_1|g) \left[(1 - \gamma)p_0(g) + \gamma \sum_{\bar{s}, \bar{a}} P(g|\bar{s}, \bar{a}) d^\pi(\bar{s}, \bar{a}) \right] \\ &= 1 \cdot [0 + \gamma(1 \cdot d^\pi(s_0, a_2) + 1 \cdot d^\pi(s_1, a_1) + 1 \cdot d^\pi(g, a_1))] \\ &= (1 - \theta)\gamma(1 - \gamma) + \theta\gamma^2(1 - \gamma) + \gamma d^\pi(g, a_1) \\ &\rightarrow d^\pi(g, a_1) = (1 - \theta)\gamma + \theta\gamma^2 \end{aligned}$$

Using $E(s, a)$ and $d^\pi(s, a)$, $D_{KL}(d^\pi \| E)$ is expressed as

$$\begin{aligned} \min_{\pi} D_{KL}(d^\pi(s, a) \| E(s, a)) &= d^\pi(s_0, a_1) \log \frac{d^\pi(s_0, a_1)}{E(s_0, a_1)} + d^\pi(s_0, a_2) \log \frac{d^\pi(s_0, a_2)}{E(s_0, a_2)} \\ &\quad + d^\pi(s_1, a_1) \log \frac{d^\pi(s_1, a_1)}{E(s_1, a_1)} + d^\pi(g, a_1) \log \frac{d^\pi(g, a_1)}{E(g, a_1)} \\ \min_{\theta} D_{KL}(d^\pi(s, a) \| E(s, a)) &= \theta(1 - \gamma) \log \frac{\theta(1 - \gamma)}{1/5} + (1 - \theta)(1 - \gamma) \log \frac{(1 - \theta)(1 - \gamma)}{2/15} \\ &\quad + \theta\gamma(1 - \gamma) \log \frac{\theta\gamma(1 - \gamma)}{1/5} + (1 - \theta)\gamma + \theta\gamma^2 \log \frac{(1 - \theta)\gamma + \theta\gamma^2}{7/15} \end{aligned}$$

To find a minimizer θ^* of above optimization, we use first-order optimality condition. That is,

$$\frac{\partial D_{KL}(d^\pi(s, a) \| E(s, a))}{\partial \theta} = 0.$$

We derive first derivative of $D_{KL}(d^\pi(s, a) \| E(s, a))$. Result is

$$\begin{aligned} \frac{\partial D_{KL}(d^\pi(s, a) \| E(s, a))}{\partial \theta} &= \\ (1 - \gamma) \left[(1 + \gamma) \log \theta - \log(1 - \theta) + \log \frac{2}{3} + \gamma \left[\log \frac{7\gamma(1 - \gamma)}{3} - \log((1 - \theta)\gamma + \theta\gamma^2) \right] \right] &= 0 \end{aligned}$$

It is hard to get the exact θ^* value because equation is very complicated. Therefore, we substitute the probability value $\frac{3}{5}$, which is probability of expert policy, into θ to check whether the derivative becomes 0.

$$\begin{aligned} \frac{\partial D_{KL}(d^\pi(s, a) \| E(s, a))}{\partial \theta} &= \\ (1 - \gamma) \left[(1 + \gamma) \log \frac{3}{5} - \log(1 - \frac{3}{5}) + \log \frac{2}{3} + \gamma \left[\log \frac{7\gamma(1 - \gamma)}{3} - \log((1 - \frac{3}{5})\gamma + \frac{3}{5}\gamma^2) \right] \right] &= 0 \\ \rightarrow \gamma \log \frac{3}{5} + \gamma \left[\log \frac{7\gamma(1 - \gamma)}{3} - \log \left(\frac{2}{5}\gamma + \frac{3}{5}\gamma^2 \right) \right] &= 0 \\ \rightarrow \log \frac{3}{5} + \log \frac{7\gamma(1 - \gamma)}{3} - \log \left(\frac{2}{5}\gamma + \frac{3}{5}\gamma^2 \right) &= \log \left(\frac{7(1 - \gamma)}{2 + 3\gamma} \right) = 0 \\ \rightarrow 7(1 - \gamma) &= 2 + 3\gamma \end{aligned}$$

Since $\gamma \in [0, 1)$, we exclude $\gamma = 1$ case. Solving last equation yields $\gamma = 0.5$. In summary, if $\gamma = 0.5$, the expert policy can be recovered through optimization $\min_{\pi} D_{KL}(d^\pi(s, a) \| E(s, a))$, but if other γ values are applied, it can be interpreted that the optimal policy is not the same as the expert policy.

C FINDING IGI DISTRIBUTION

Our goal is to find the initial distribution \tilde{p}_0 that makes the distribution made by $t_0 + t_{geom}$ uniform, where $t_0 \sim \tilde{p}_0(t)$ and $t_{geom} \sim \text{Geom}_{1-\gamma}(t|t_0)$. Note that $\text{Geom}_{1-\gamma}(t|t_0)$ is conditional distribution. Because, in general, there is a maximum timestep, $|T|$, for episode in the learning environment, so in order to prevent the sum of timestep exceed over $|T|$, $\text{Geom}_{1-\gamma}(t|t_0)$ must be conditional distribution which considers the sampled initial timestep. Let $T = t_0 + t_{geom}$, Then $P_T(T)$ is defined as

$$\begin{aligned} P_T(T) &= \sum_{t_0=0}^T \text{Geom}_{1-\gamma}(T - t_0|t_0) \tilde{p}_0(t_0) \\ &= \sum_{t_0=0}^T \frac{(1 - \gamma)\gamma^{T-t_0}}{\sum_{i=0}^{|T|-t_0} (1 - \gamma)\gamma^i} \tilde{p}_0(t_0) \\ &= \sum_{t_0=0}^T \frac{(1 - \gamma)\gamma^{T-t_0}}{\text{Sum}(|T| - t_0)} \tilde{p}_0(t_0). \end{aligned} \quad (30)$$

For the convenience of notation, we rewrite $\sum_{i=0}^{|T|-t_0} (1 - \gamma)\gamma^i$ as $\text{Sum}(|T| - t_0)$. We can express Equation (30) in linear problem form. To this end, it is represented as follows.

$$\underbrace{\begin{bmatrix} P_T(0) \\ \vdots \\ P_T(|T|) \end{bmatrix}}_{P_T} = \underbrace{\begin{bmatrix} \frac{1-\gamma}{\text{Sum}(|T|)} & 0 & \cdots & 0 \\ \frac{(1-\gamma)\gamma}{\text{Sum}(|T|)} & \frac{1-\gamma}{\text{Sum}(|T|-1)} & 0 & \cdots & 0 \\ \vdots & \frac{(1-\gamma)\gamma}{\text{Sum}(|T|-1)} & \frac{1-\gamma}{\text{Sum}(|T|-2)} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{(1-\gamma)\gamma^{|T|}}{\text{Sum}(|T|)} & \frac{(1-\gamma)\gamma^{|T|-1}}{\text{Sum}(|T|-1)} & \frac{(1-\gamma)\gamma^{|T|-2}}{\text{Sum}(|T|-2)} & \cdots & 1 - \gamma \end{bmatrix}}_{P_{t_{geom}}} \underbrace{\begin{bmatrix} \tilde{p}_0(0) \\ \vdots \\ \tilde{p}_0(|T|) \end{bmatrix}}_{\tilde{p}_0}$$

When we actually obtain P_T , we obtain it through $P_T(T) = \frac{\text{The number of data with timestep T in dataset}}{\text{Total number of data in dataset}}$. Finally, solving $\tilde{p}_0 = P_{t_{geom}}^{-1} \cdot P_T$ gives us the $\tilde{p}_0(t)$ we want.

D EXPERIMENTAL DETAILS

D.1 ALGORITHM DETAILS

In this subsection, we provide a simple pseudocode for our algorithm using IGI.

Algorithm 1 Our algorithm using IGI

Require: dataset $\mathcal{D}^D = \mathcal{D}^E \cup \mathcal{D}^O$, **IGI distribution** \tilde{p}_0 from Appendix C, policy network π_θ with parameter θ , neural networks ν_ϕ and c_ψ with parameter ϕ and ψ , discount factor γ , batch size B , learning rate α

Ensure: P_{T_D} is a distribution P_T from Appendix C using dataset \mathcal{D} , and \mathcal{D}_t is an dataset \mathcal{D} consisted of data with timestep t

Make IGI distribution \tilde{p}_0 using γ and $P_{T_{\mathcal{D}^D}}$

while total iterations **do**

 Sample initial timestep: $\{t_0^{(i)}\}_{i=1}^B \sim \tilde{p}_0$ ▷ using IGI distribution

 Sample uniform timestep: $\{t^{(i)}\}_{i=1}^B \sim P_{T_{\mathcal{D}^D}}$

 Sample expert timestep: $\{t_e^{(i)}\}_{i=1}^B \sim P_{T_{\mathcal{D}^E}}$

 Sample initial state: $\{s_0^{(i)}\} \sim \mathcal{D}_{t_0^{(i)}}^D$, for $i = 1, \dots, B$

 Sample total data: $\{(s^{(i)}, a^{(i)}, s'^{(i)})\} \sim \mathcal{D}_{t^{(i)}}^D$, for $i = 1, \dots, B$

 Sample expert data: $\{(s_e^{(i)}, a_e^{(i)}, s'_e^{(i)})\} \sim \mathcal{D}_{t_e^{(i)}}^E$, for $i = 1, \dots, B$

 Compute discriminator loss J_c (11) using $\{(s^{(i)}, a^{(i)})\}_{i=1}^B$ and $\{(s_e^{(i)}, a_e^{(i)})\}_{i=1}^B$

 Compute critic loss J_ν (12) using $\{s_0^{(i)}\}_{i=1}^B$ and $\{(s^{(i)}, a^{(i)}, s'^{(i)})\}_{i=1}^B$

 Compute policy loss J_π (13) using $\{(s^{(i)}, a^{(i)}, s'^{(i)})\}_{i=1}^B$

 Update $\psi \leftarrow \psi - \alpha \nabla_\psi J_c$

 Update $\phi \leftarrow \phi - \alpha \nabla_\phi J_\nu$

 Update $\theta \leftarrow \theta - \alpha \nabla_\theta J_\pi$

end while

To compute three loss functions in Algorithm 1, we need sampled data from d^E , D , and p_0 . This begins by sampling the timesteps for each distribution. First of all, compute the IGI distribution using the geometric distribution $\text{Geom}(1 - \gamma)$ given a discount factor γ as we described in Appendix C. At this time, uniform distribution P_T for the total dataset D^D is also used for practical performance as shown in Figure 3 of Section 4.2. P_T is defined as $\frac{\text{The number of data with timestep } T \text{ in dataset}}{\text{Total number of data in dataset}}$ as in Appendix C. Since IGI enables sampling for the discounted distribution uniformly regardless of the timestep, sampling from d^E is replaced with sampling from $P_{T_{\mathcal{D}^E}}$. After three types of timesteps are sampled, we need actual data corresponding to each type of timestep to use in loss functions. We uniformly sample actual data of corresponding timesteps in each three dataset. To compute the loss function (11) of the discriminator, uniformly sampled expert data is used to compute the first term and the total data is used to compute the second term. In the case of the critic network, the initial data from IGI is used to compute the first term in (12) and the total data is used to compute the second term. When we compute $A_\nu(s, a)$ in the second term, use $\gamma \nu(s'^{(i)})$ instead of $\gamma \mathcal{P} \nu(s^{(i)})$ and use the discriminator to compute $\log \frac{d^E(s, a)}{D(s, a)}$ as notified in the sentence below (11). For the policy network, total data is used to compute (13). Lastly, update the parameters of each network with a learning rate α . Repeating this process for the total number of iterations.

D.2 IMPLEMENTATION DETAIL

In this subsection, we provide the hyperparameter settings of our algorithm. We use an absorbing state for practical implementation (Kostrikov et al., 2018). For discriminator learning, we use WGAN-GP (Gulrajani et al., 2017) to achieve more robust learning.

Table 1: Hyperparameter settings

Hyperparameters	Our setting
Actor learning rate	3×10^{-4}
Actor network size	[256, 256]
Critic learning rate	3×10^{-4}
Critic network size	[256, 256]
Critic gradient L2-norm coefficient	10^{-4}
Discriminator learning rate	3×10^{-4}
Discriminator network size	[256, 256]
Discriminator gradient penalty coefficient	10
Batch size	256
Number of total iteration	10^6
Random seeds	1, 2, 3
Kernel initializer	He normal initializer

E EXPERIMENT RESULTS

E.1 FINITE AND DISCRETE MDP

In this subsection, we provide experiment details in finite-discrete MDP and its results.




Figure 5: The trend of evaluation result with various discount factors in different settings. We plot the normalized average evaluation over 1000 random seeds.

In finite and discrete MDP, we empirically show the trade-off effect by the discount factor in the Random MDP. This environment generates the finite and discrete MDP randomly. We follow the environment configuration of Lee et al. (2021).

We construct the expert dataset \mathcal{D}^E and suboptimal dataset \mathcal{D}^O by rolling expert policy and suboptimal policy, respectively. We characterized the expert policy as the stochastic policy based on the optimal Q (state-action) value of the randomly generated MDP. The suboptimal policy has performance between optimal and uniformly random policy by controlling the hyperparameter $\omega \in [0, 1]$ as $\omega V^*(s_0) + (1 - \omega)V^{\pi_{\text{unif}}}(s_0)$ where V^* and $V^{\pi_{\text{unif}}}$ is the value function of optimal policy and uniformly random policy, respectively. Since the MDP is finite and discrete, we can compute the discounted visitation distribution of the policy directly, which means learning the discriminator is unnecessary. Thus, there is no need to use IGI in this environment. We use Maximum likelihood estimation (MLE) transition dynamics \hat{P} and discount factor $\hat{\gamma}$ for training the policy as we discussed in Theorem 1. We use discount factor for $\hat{\gamma} \in \{0.99, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1\}$ with varying the number of suboptimal trajectories for $\{5, 20, 60, 100\}$. The number of expert trajectory is 1 for all settings. Figure 5 shows the policy evaluation result on the true MDP with true discount factor γ . It shows evaluation trends so that the trade-off effect by discount factor suggested in Section 3.2 can be confirmed. We can see that the performance can be optimized by choosing lower γ . Furthermore, the figure shows that the performance of using large discount factors becomes better as the number of suboptimal trajectories gets larger.

E.2 CONTINUOUS MDP

We provide the supplementary experiment results of Figure 4 here. We plot the learning curve according to γ and dataset ratio. Figure 6 represents the result in HalfCheetah-v2 environment, and Figure 7, 8, 9 is the result on the Hopper-v2, Walker2d-v2 and Ant-v2, respectively.

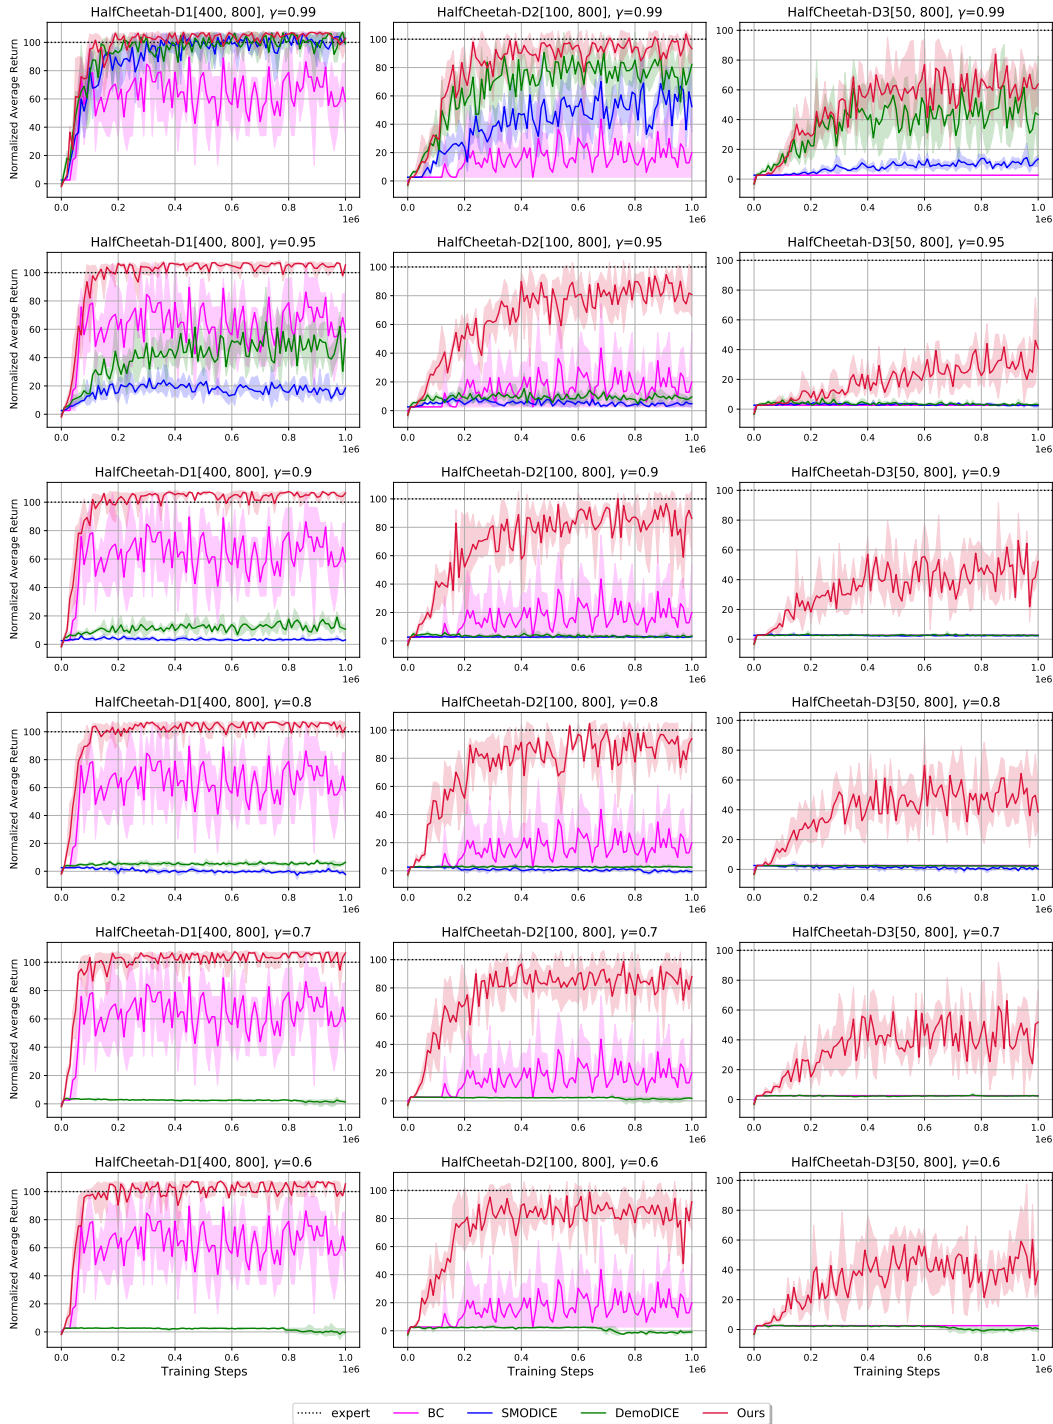


Figure 6: Performance of our algorithm using IGI and other baseline algorithms with $\gamma \in \{0.99, 0.95, 0.9, 0.8, 0.7, 0.6\}$ and dataset ratio D1, D2 and D3 in HalfCheetah-v2 environment.

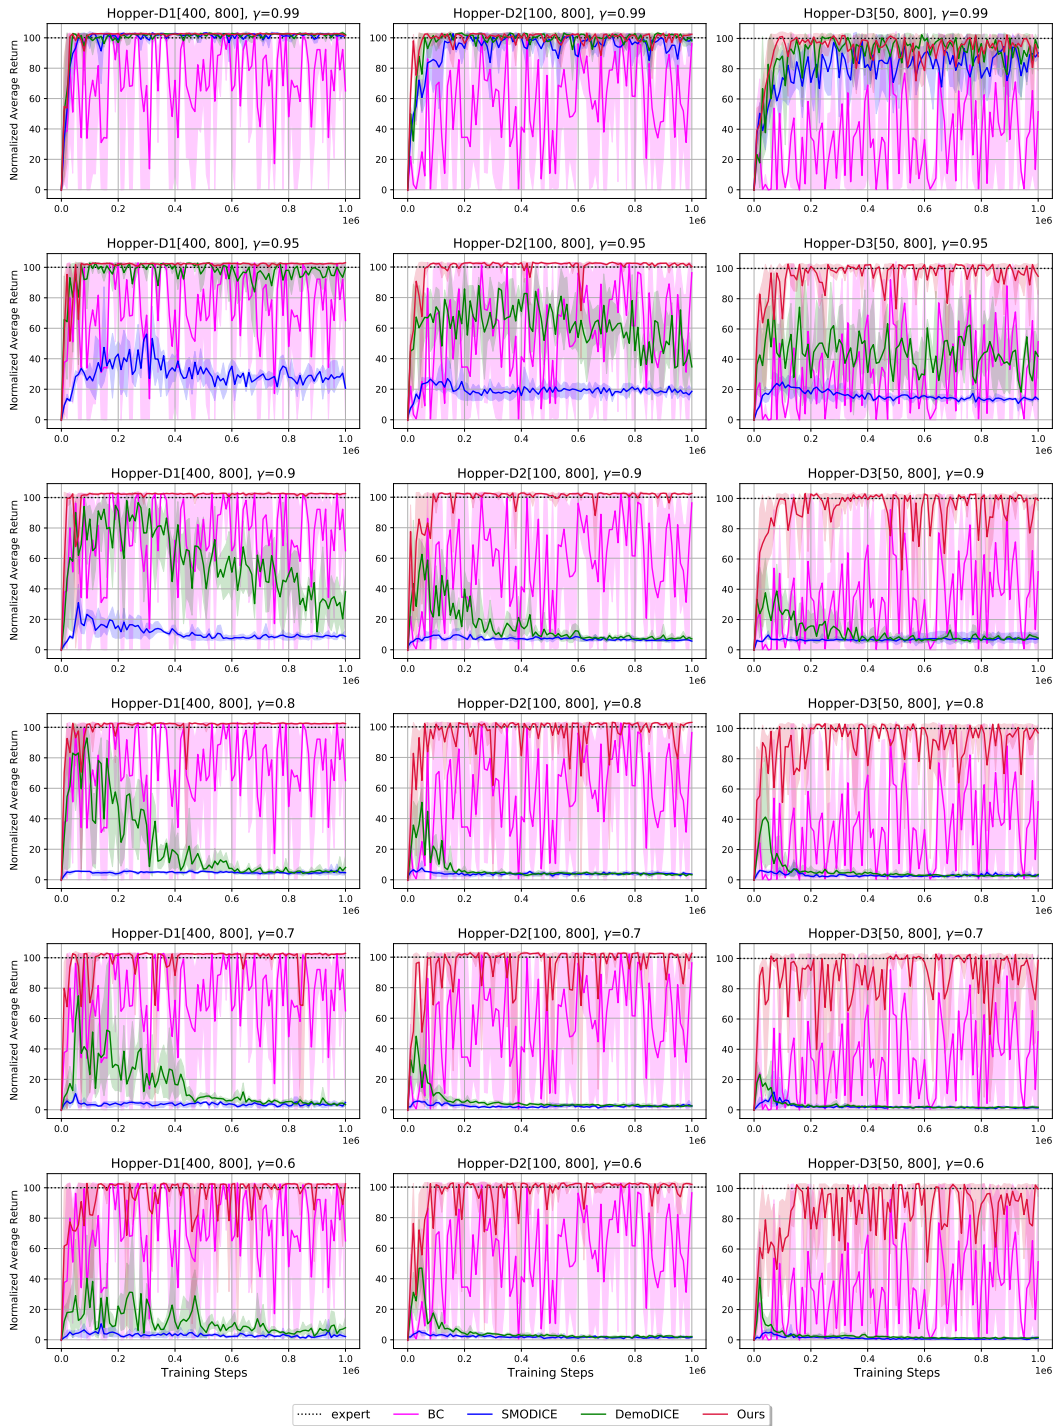


Figure 7: Performance of our algorithm using IGI and other baseline algorithms with $\gamma \in \{0.99, 0.95, 0.9, 0.8, 0.7, 0.6\}$ and dataset ratio D1, D2 and D3 in Hopper-v2 environment.

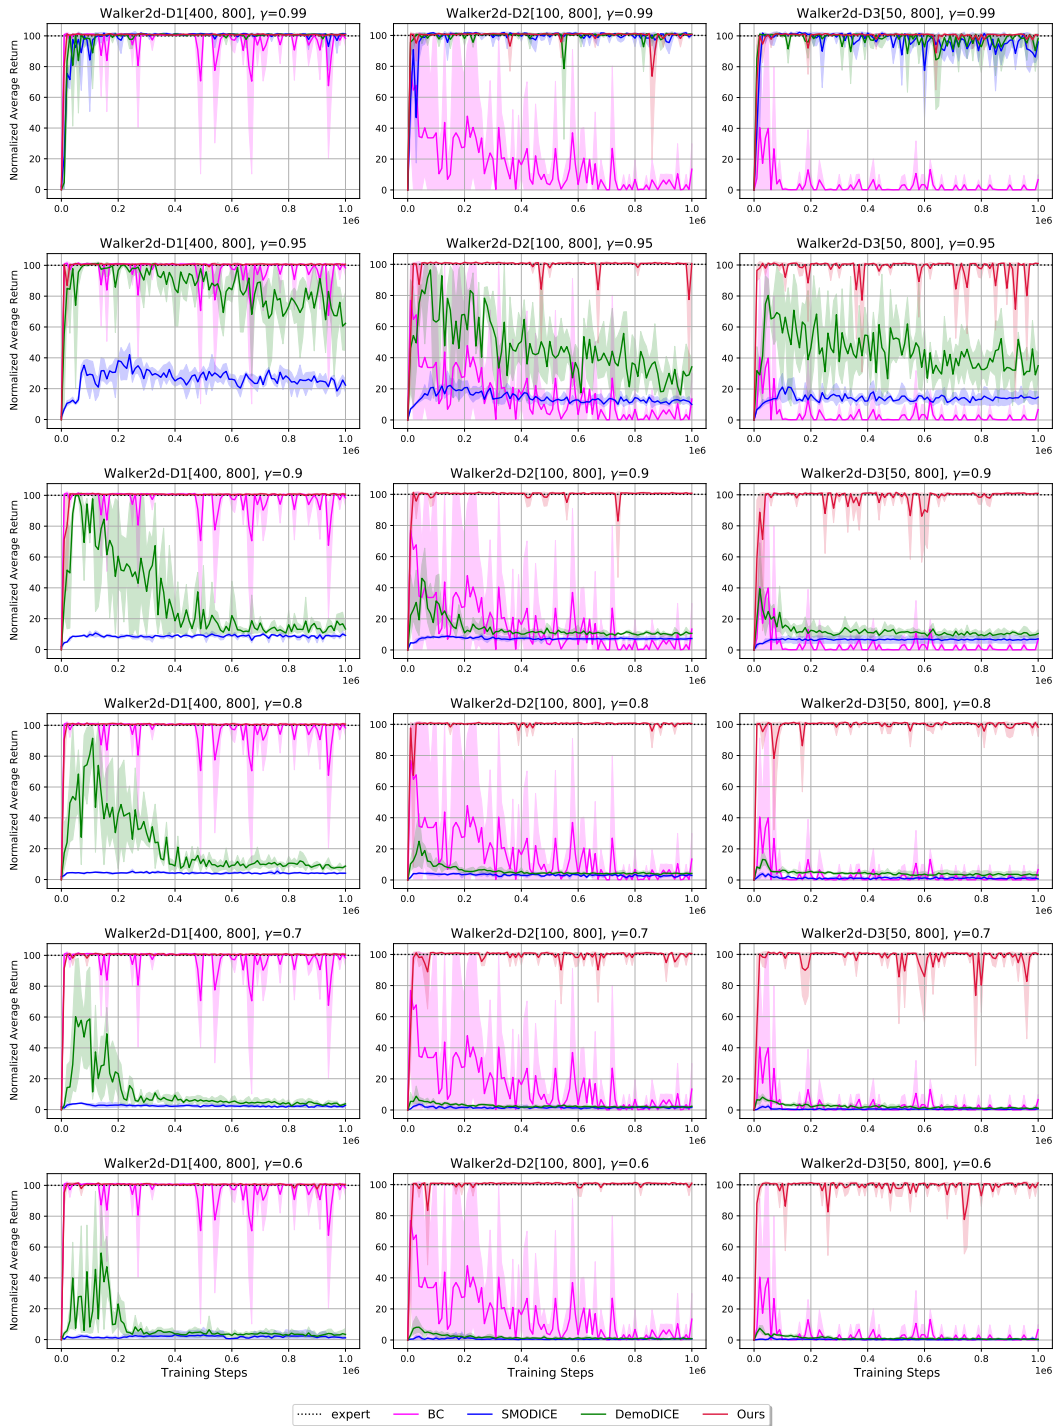


Figure 8: Performance of our algorithm using IGI and other baseline algorithms with $\gamma \in \{0.99, 0.95, 0.9, 0.8, 0.7, 0.6\}$ and dataset ratio D1, D2 and D3 in Walker2d-v2 environment.

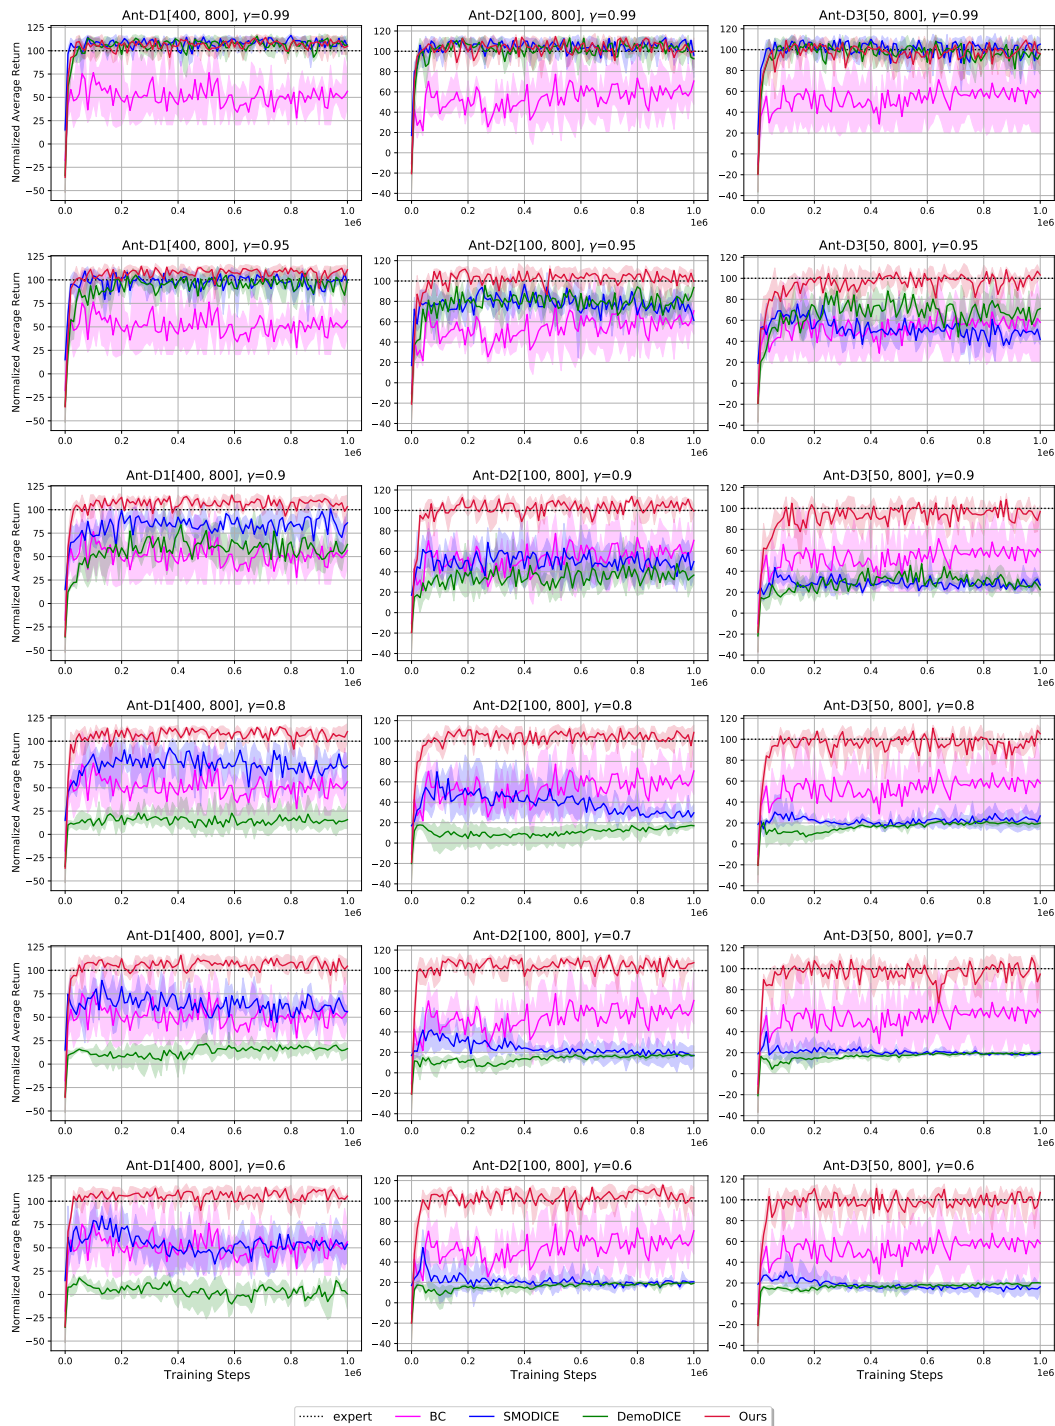


Figure 9: Performance of our algorithm using IGI and other baseline algorithms with $\gamma \in \{0.99, 0.95, 0.9, 0.8, 0.7, 0.6\}$ and dataset ratio D1, D2 and D3 in Ant-v2 environment.

F ADDITIONAL EXPERIMENT RESULT

F.1 EXPERIMENT ON SMALL DATASET IN CONTINUOUS MDP




Figure 10: The trend of the final evaluation over 3 random seeds with various discount factors in different environment. For comparison, we plot green line which represents the normalized average returns of DemoDICE($\gamma = 0.99$).

We make additional experiments by lowering the amount of dataset about 10 times less (40 expert trajectories and 80 random-policy trajectories) than D1 / D2 / D3 settings. Figure 10 shows the last evaluation of IGI with applying aforementioned setting in 4 MuJoCo environments over 3 random seeds, and green line shown in the Figure 10 is the maximum last evaluation over 3 seed of DemoDICE among different discount factors for each environment (Note that best discount factor of DemoDICE is 0.99 for all environments).

As shown in the attached figure, we confirmed that the optimal discount factor appeared at lower than 0.99, except for the HalfCheetah-v2 environment. To explain the reason why a different trend came out only in HalfCheetah-v2, we inform that unlike the other environments, HalfCheetah-v2 environment has the peculiarity that learning is stable in the offline RL problem, even when the policy is unregularized. For example, refer to the experiment result by varying the policy regularization factor in Figure 9 of Wu et al. (2019), HalfCheetah-v2 environment shows good performance compared to other environments even when a low regularization factor α is applied. Combining these experimental results with the regularization effect on the discount factor in the offline RL referred to Hu et al. (2022) (Section 3), we can expect that the performance of the learned policy can be improved when a high discount factor is used in the HalfCheetah-v2 environment. Eventually, the reason why Figure 10 shows a different trend only in the HalfCheetah-v2 environment is interpreted as reflecting this environmental disposition, and the error bound that we proposed in Theorem 1 is valid under general situations.

F.2 TEST ON VALIDITY OF THEOREM 1

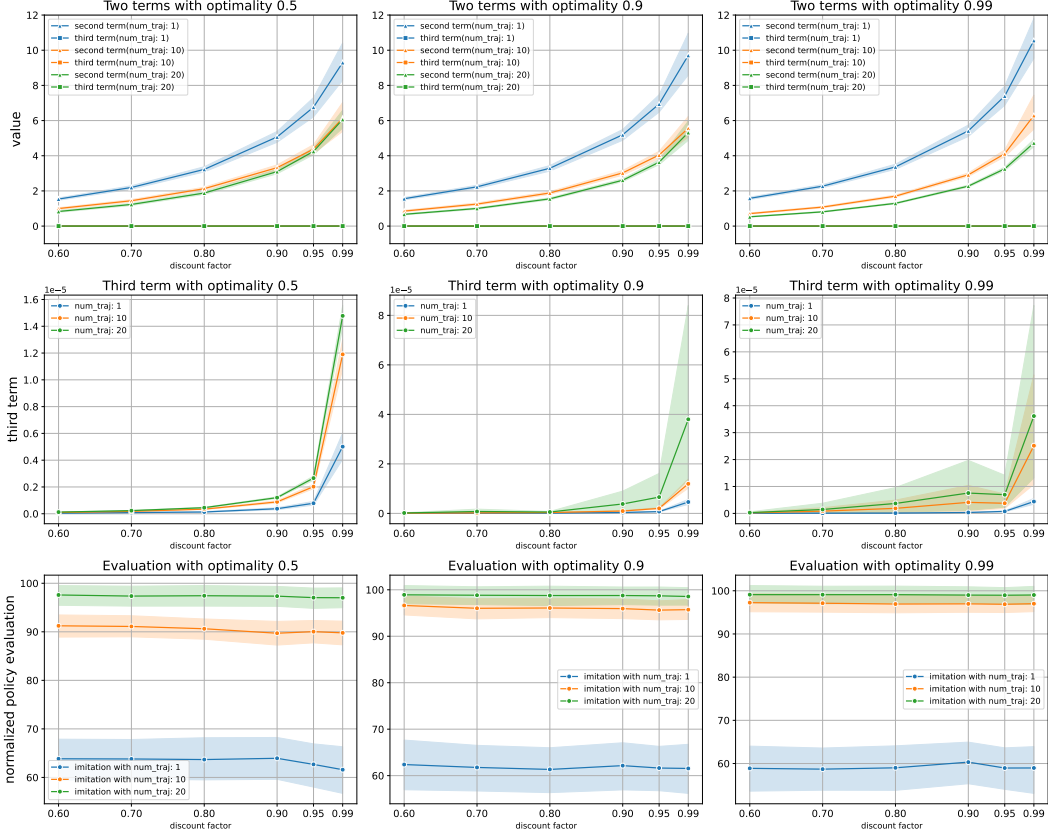


Figure 11: Figures in the first row show the magnitude of the second and third terms averaged over 100 random seeds with different discount factors and the number of suboptimal trajectories (num_traj) for each optimality setting. The second row shows only the magnitude of the third term, and the last row represents the performance of imitated policies.

To check the validity of Theorem 1, we measured the magnitude of the second and third terms in finite-discrete MDP described in E.1 where ϵ_P and ϵ_π can be exactly computed. Figure 11 shows the trend of two terms in (14) with different discount factors at each setting. Results show that the third term has a much lower value than the second term. As a result, Theorem 1 is mainly dependent on the first and second terms.

The peculiar part is that the third term increases as the suboptimal data increases. In this situation, the proportion of expert data in the total dataset is reduced, and in that case, the state-marginal distribution, $d_{\hat{P}, \hat{\gamma}}^\pi(s)$ generated by MLE MDP will give unnecessary amount of probability for the suboptimal data. Therefore, it is interpreted that the third term is increased because samples with high $D_{TV}(\pi || \pi^E)$ are more frequently reflected in ϵ_π as suboptimal data increases.

F.3 PERFORMANCE ON DIFFERENT SUBOPTIMAL DATASET




Figure 12: Last evaluation of IGI and best among DemoDICE when suboptimal dataset is consisted of 850 medium-trajectory. Gray line represents best of DemoDICE ($\gamma = 0.99$).

We additionally conducted an experiment using medium trajectory as suboptimal data. Figure 12 shows the evaluations of learning 1 million steps according to the discount factor when the suboptimal dataset is consisted of 850 medium trajectories, and the gray line represents DemoDICE’s best performance. As a result, in an environment other than `Hopper-v2` it was possible to obtain a higher return than DemoDICE. We also included the learning curve of this experiment in next page, Figure 13. Here, we can see that the IGI still shows robust learning on the discount factor.

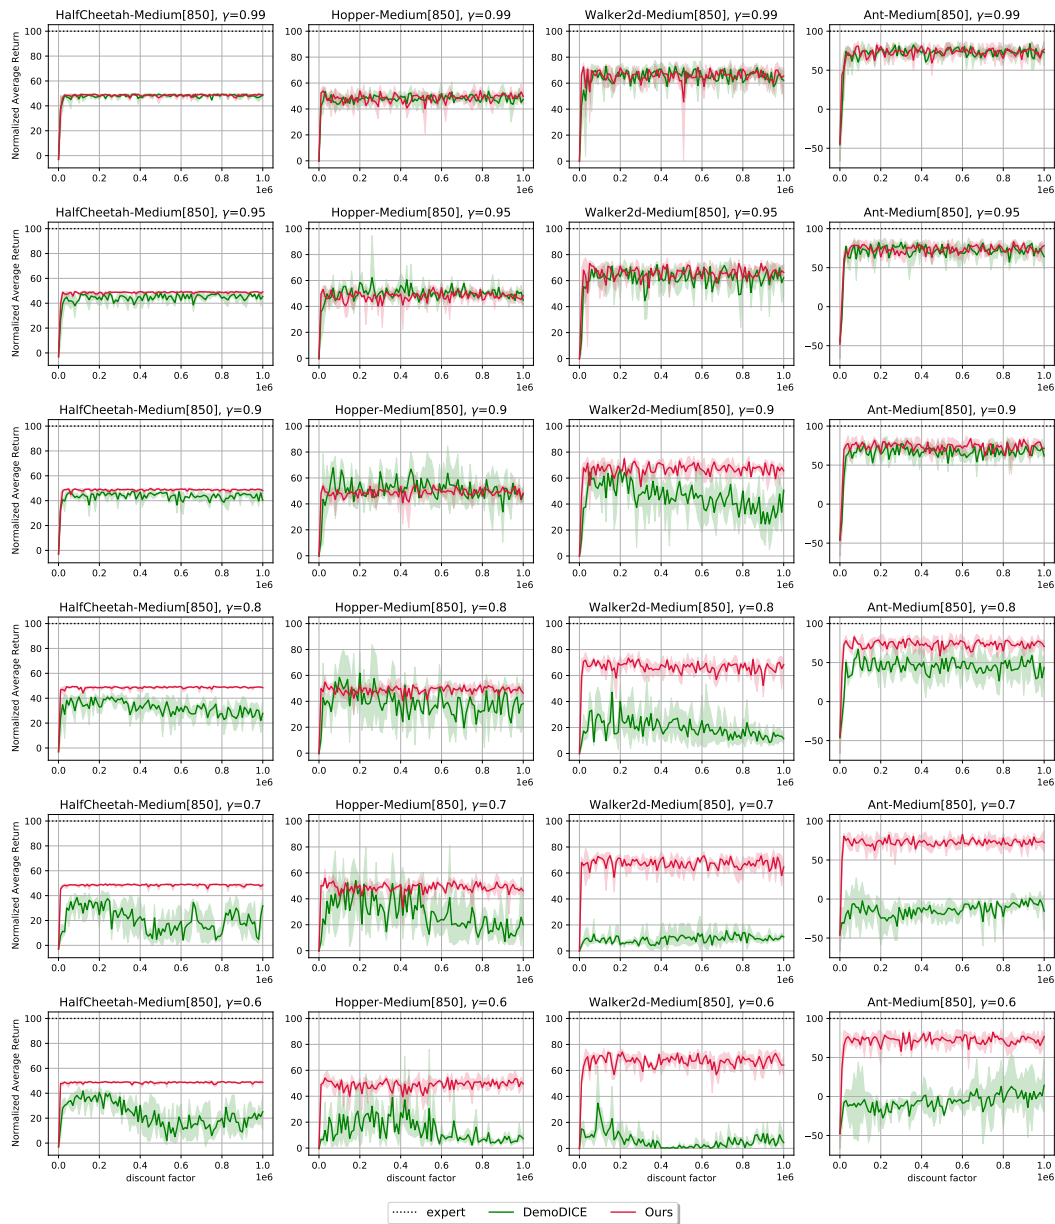


Figure 13: Performance of IGI and DemoDICE when suboptimal dataset is made by 850 medium-trajectories. Each column shows the normalized return of two algorithms according to the discount factor in the same environment.