
Inverse Reinforcement Learning from a Gradient-based Learner

Giorgia Ramponi
Politecnico Di Milano
Milan, Italy
giorgia.ramponi@polimi.it

Gianluca Drappo
Politecnico Di Milano
Milan, Italy
gianluca.drappo@mail.polimi.it

Marcello Restelli
Politecnico Di Milano
Milan, Italy
marcello.restelli@polimi.it

Abstract

Inverse Reinforcement Learning addresses the problem of inferring an expert’s reward function from demonstrations. However, in many applications, we not only have access to the expert’s near-optimal behaviour, but we also observe part of her learning process. In this paper, we propose a new algorithm for this setting, in which the goal is to recover the reward function being optimized by an agent, given a sequence of policies produced during learning. Our approach is based on the assumption that the observed agent is updating her policy parameters along the gradient direction. Then we extend our method to deal with the more realistic scenario where we only have access to a dataset of learning trajectories. For both settings, we provide theoretical insights into our algorithms’ performance. Finally, we evaluate the approach in a simulated GridWorld environment and on the MuJoCo environments, comparing it with the state-of-the-art baseline.

1 Introduction

Inverse Reinforcement Learning (IRL) [20] aims to infer an expert’s reward function from her demonstrations [21]. In the standard setting, an expert shows behaviour by repeatedly interacting with the environment. This behaviour, encoded by its policy, is optimizing an *unknown* reward function. The goal of IRL consists of finding a reward function that makes the expert’s behaviour optimal [20]. Compared to other imitation learning approaches [3, 15], which output an imitating policy (e.g. Behavioral Cloning [3]), IRL explicitly provides a succinct representation of the expert’s *intention*. For this reason, it provides a generalization of the expert’s policy to unobserved situations.

However, in some cases, it is not possible to wait for the convergence of the demonstrator’s learning process. For instance, in multi-agent environments, an agent has to infer the unknown reward functions that the other agents are learning, before actually becoming “experts”; so that she can either cooperate or compete with them. On the other hand, in many situations, we can learn something useful by observing the learning process of an agent. These observations contain important information about the agent’s intentions and can be used to infer her interests. Imagine a driver who is learning a new circuit. During her training, we can observe how she behaves in a variety of situations (even dangerous ones) and this is useful for understanding which states are good and which should be avoided. Instead, when expert behaviour is observed, only a small sub-region of the state space could be explored, thus leaving the observer unaware of what to do in situations that are unlikely under the expert policy.

Inverse Reinforcement Learning from not expert agents, called Learning from a Learner (LfL), was recently proposed by Jacq et al. in [17]. LfL involves two agents: a *learner* who is currently learning a task and an *observer* who wants to infer the learner’s intentions. In [17] the authors assume that the learner is learning under an entropy-regularized framework, motivated by the assumption that the learner is showing a sequence of constantly improving policies. However many Reinforcement Learning (RL) algorithms [35] do not satisfy this and also human learning is characterized by mistakes that may lead to a non-monotonic learning process.

In this paper we propose a new algorithm for the LfL setting called Learning Observing a Gradient not-Expert Learner (LOGEL), which is not affected by the violation of the constantly improving assumption. Given that many successful RL algorithms are gradient-based [22] and there is some evidence that the human learning process is similar to a gradient-based method [32], we assume that the learner is following the gradient direction of her expected discounted return. The algorithm learns the reward function that minimizes the distance between the actual policy parameters of the learner and the policy parameters that should be obtained if she were following the policy gradient using that reward function.

After a formal introduction of the LfL setting in Section 3, we provide in Section 4 a first solution of the LfL problem when the observer has full access to the learner’s policy parameters and learning rates. Then, in Section 5 we extend the algorithm to the more realistic case in which the observer can identify the optimized reward function only by analyzing the learner’s trajectories. For each problem setting, we provide a finite sample analysis to give to the reader an intuition on the correctness of the recovered weights. Finally, we consider discrete and continuous simulated domains to empirically compare the proposed algorithm with state-of-the-art baselines in this setting [17, 7]. The proofs of all the results are reported in Appendix A. In the appendix we report preliminary results on a simulated autonomous driving task B.3.

2 Preliminaries

A **Markov Decision Process** (MDP) [27, 35] is a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, \gamma, \mu, R)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is the transition function, which defines the density $P(s'|s, a)$ of state $s' \in \mathcal{S}$ when taking action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$, $\gamma \in [0, 1)$ is the discount factor, $\mu : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is the initial state distribution and $R : \mathcal{S} \rightarrow \mathbb{R}$ is the reward function. An RL agent follows a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$, where $\pi(\cdot|s)$ specifies for each state s a distribution over the action space \mathcal{A} , i.e., the probability of taking action a in state s . We consider stochastic differentiable policies belonging to a parametric space $\Pi_{\Theta} = \{\pi_{\theta} : \theta \in \Theta \subseteq \mathbb{R}^d\}$. We evaluate the performance of a policy π_{θ} as its expected cumulative discounted return:

$$J(\theta) = \mathbb{E}_{\substack{S_0 \sim \mu, \\ A_t \sim \pi_{\theta}(\cdot|S_t), \\ S_{t+1} \sim P(\cdot|S_t, A_t)}} \left[\sum_{t=0}^{+\infty} \gamma^t R(S_t, A_t) \right].$$

To solve an MDP, we must find a policy π_{θ^*} that maximizes the performance $\theta^* \in \arg \max_{\theta} J(\theta)$.

Inverse Reinforcement Learning [21, 20, 2] addresses the problem of recovering the *unknown* reward function optimized by an expert given demonstrations of her behavior. The expert plays a policy π^E which is (nearly) optimal for some unknown reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. We are given a dataset $D = \{\tau_1, \dots, \tau_n\}$ of trajectories from π^E , where we define a *trajectory* as a sequence of states and actions $\tau = (s_0, a_0, \dots, s_{T-1}, a_{T-1}, s_T)$, where T is the trajectory length. The goal of an IRL agent is to find a reward function that explains the expert’s behavior. As commonly done in the Inverse Reinforcement Learning literature [24, 41, 2], we assume that the expert’s reward function can be represented by a linear combination with weights ω of q basis functions ϕ :

$$R_{\omega}(s, a) = \omega^T \phi(s, a), \quad \omega \in \mathbb{R}^q, \quad (1)$$

where $\phi : \mathcal{S} \times \mathcal{A} \rightarrow [-M_r, M_r]^q$ is a bounded feature vector function.

We define the **feature expectations** of a policy π_{θ} as:

$$\psi(\theta) = \mathbb{E}_{\substack{S_0 \sim \mu, \\ A_t \sim \pi_{\theta}(\cdot|S_t), \\ S_{t+1} \sim P(\cdot|S_t, A_t)}} \left[\sum_{t=0}^{+\infty} \gamma^t \phi(S_t, A_t) \right].$$

The **expected discounted return**, under the linear reward model, is defined as:

$$J(\boldsymbol{\theta}, \boldsymbol{\omega}) = \mathbb{E}_{\substack{S_0 \sim \mu, \\ A_t \sim \pi_{\boldsymbol{\theta}}(\cdot | S_t), \\ S_{t+1} \sim P(\cdot | S_t, A_t)}} \left[\sum_{t=0}^{+\infty} \gamma^t R_{\boldsymbol{\omega}}(S_t, A_t) \right] = \boldsymbol{\omega}^T \boldsymbol{\psi}(\boldsymbol{\theta}). \quad (2)$$

3 Inverse Reinforcement Learning from learning agents

The Learning from a Learner Inverse Reinforcement Learning setting (LfL), proposed in [17], involves two agents:

- a *learner* which is learning a task defined by the reward function $R_{\boldsymbol{\omega}^L}$,
- and an *observer* which wants to infer the learner’s reward function.

More formally, the learner is an RL agent which is learning a policy $\pi_{\boldsymbol{\theta}} \in \Pi_{\Theta}$ in order to maximize its *discounted expected return* $J(\boldsymbol{\theta}, \boldsymbol{\omega}^L)$. The learner is improving its own policy by an update function $f(\boldsymbol{\theta}, \boldsymbol{\omega}) : \mathbb{R}^d \times \mathbb{R}^q \rightarrow \mathbb{R}^d$, i.e., at time t , $\boldsymbol{\theta}_{t+1} = f(\boldsymbol{\theta}_t, \boldsymbol{\omega})$. The observer, instead, perceives a sequence of learner’s policy parameters $\{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{m+1}\}$ and a dataset of trajectories for each policy $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_{m+1}\}$, where $\mathcal{D}_i = \{\tau_1^i, \dots, \tau_n^i\}$. Her goal is to recover the reward function $R_{\boldsymbol{\omega}^L}$ that explains $\pi_{\boldsymbol{\theta}_i} \rightarrow \pi_{\boldsymbol{\theta}_{i+1}}$ for all $1 \leq i \leq m$, i.e the updates of the learner’s policy.

Remark 3.1. *It is easy to notice that this problem has the same intention as Inverse Reinforcement Learning since the demonstrating agent is motivated by some reward function. On the other hand, in classical IRL the learner agent is an expert, and not a non-stationary agent. For this reason, we cannot simply apply standard IRL algorithms to this problem or use Behavioral Cloning [26, 3, 21] algorithms, which mimic a suboptimal behavior.*

4 Learning from a learner following the gradient

Many algorithms that are the state of the art of reinforcement learning are policy-gradient methods [22, 36], i.e. approaches which optimize the expected discounted return with gradient updates of the policy parameters. Recently it has been proved that even standard RL algorithms such as Value Iteration or Q-learning have strict connections with policy gradient methods [13, 30]. For the above reasons, we assume that the learner is optimizing the expected discounted return using gradient descent.

For the sake of presentation, we start by considering the simplified case in which we assume that the observer can perceive the sequence of the learner’s policy parameters $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{m+1})$, the associated gradients of the feature expectations $(\nabla_{\boldsymbol{\theta}} \boldsymbol{\psi}(\boldsymbol{\theta}_1), \dots, \nabla_{\boldsymbol{\theta}} \boldsymbol{\psi}(\boldsymbol{\theta}_m))$, and the learning rates $(\alpha_1, \dots, \alpha_m)$. Then, we will replace the exact knowledge of the gradients with estimates built on a set of demonstrations \mathcal{D}_i for each learner’s policy $\pi_{\boldsymbol{\theta}_i}$ (Section 4.2). Finally, we introduce our algorithm LOGEL, which, using behavioral cloning and an alternate block-coordinate optimization [37], is able to estimate the reward’s parameters without requiring as input the policy parameters and the learning rates (Section 5).

4.1 Exact gradient

We express the gradient of the expected return as [36, 23]:

$$\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}, \boldsymbol{\omega}) = \mathbb{E}_{\substack{S_0 \sim \mu, \\ A_t \sim \pi_{\boldsymbol{\theta}}(\cdot | S_t), \\ S_{t+1} \sim P(\cdot | S_t, A_t)}} \left[\sum_{t=0}^{+\infty} \gamma^t R_{\boldsymbol{\omega}}(S_t, A_t) \sum_{l=0}^t \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(A_l | S_l) \right] = \nabla_{\boldsymbol{\theta}} \boldsymbol{\psi}(\boldsymbol{\theta}) \boldsymbol{\omega},$$

where $\nabla_{\boldsymbol{\theta}} \boldsymbol{\psi}(\boldsymbol{\theta}) = (\nabla_{\boldsymbol{\theta}} \psi_1(\boldsymbol{\theta}) | \dots | \nabla_{\boldsymbol{\theta}} \psi_q(\boldsymbol{\theta})) \in \mathbb{R}^{d \times q}$ is the Jacobian matrix of the feature expectations $\boldsymbol{\psi}(\boldsymbol{\theta})$ w.r.t the policy parameters $\boldsymbol{\theta}$. In the rest of the paper, with some abuse of notation, we will indicate $\boldsymbol{\psi}(\boldsymbol{\theta}_t)$ with $\boldsymbol{\psi}_t$.

We define the gradient-based learner updating rule at time t as:

$$\boldsymbol{\theta}_{t+1}^L = \boldsymbol{\theta}_t^L + \alpha_t \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}_t^L, \boldsymbol{\omega}) = \boldsymbol{\theta}_t^L + \alpha_t \nabla_{\boldsymbol{\theta}} \boldsymbol{\psi}_t^L \boldsymbol{\omega}^L, \quad (3)$$

where α_t is the learning rate. Given a sequence of consecutive policy parameters $(\theta_1^L, \dots, \theta_{m+1}^L)$, and of learning rates $(\alpha_1, \dots, \alpha_m)$ the observer has to find the reward function R_ω such that the improvements are explainable by the update rule in Eq. (3). This implies that the observer has to solve the following minimization problem:

$$\min_{\omega \in \mathbb{R}^q} \sum_{t=1}^m \|\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega\|_2^2, \quad (4)$$

where $\Delta_t = \theta_{t+1} - \theta_t$. This optimization problem can be easily solved in closed form under the assumption that $(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t)^{-1}$ is invertible.

Lemma 4.1. *If the matrix $(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t)^{-1}$ is full-rank than optimization problem (4) is solved in closed form by*

$$\hat{\omega} = \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right)^{-1} \left(\sum_{t=1}^m \alpha_t \nabla_{\theta} \psi_t^T \Delta_t \right). \quad (5)$$

When problem (4) has no unique solution or when the matrix to be inverted is nearly singular, in order to avoid numerical issues, we can resort to a regularized version of the optimization problem. In the case we add an L2-norm penalty term over weights ω we can still compute a closed-form solution (see Lemma A.5 in Appendix A).

4.2 Approximate gradient

In practice, we do not have access to the Jacobian matrix $\nabla_{\theta} \psi$, but the observer has to estimate it using the dataset D and some unbiased policy gradient estimator, such as REINFORCE [40] or G(PO)MDP [5]. The estimation of the Jacobian will introduce errors on the optimization problem (4). Obviously the estimation of the reward weights ω becomes more accurate when more data are available [25]. On the other hand during the learning process, the learner will produce more than one policy improvement, and the observer can use these improvements to get better estimates of the reward weights.

In order to have an insight on the relationship between the amount of data needed to estimate the gradient and the number of learning steps, we provide a finite sample analysis on the norm of the difference between the learner's weights ω^L and the recovered weights $\hat{\omega}$. The analysis takes into account the learning steps data and the gradient estimation data, without having any assumption on the policy of the learner. We denote with $\Psi = [\nabla_{\theta} \psi_1, \dots, \nabla_{\theta} \psi_m]^T$ the concatenation of the Jacobians and $\hat{\Psi} = [\widehat{\nabla_{\theta} \psi_1}, \dots, \widehat{\nabla_{\theta} \psi_m}]^T$ the concatenation of the estimated Jacobians.

Theorem 4.1. *Let Ψ be the real Jacobians and $\hat{\Psi}$ the estimated Jacobian from n trajectories $\{\tau_1, \dots, \tau_n\}$. Assume that Ψ is bounded by a constant M and $\lambda_{\min}(\hat{\Psi}^T \hat{\Psi}) \geq \lambda > 0$. Then w.h.p.:*

$$\|\omega^L - \hat{\omega}\|_2 \leq O \left(\frac{1}{\lambda} M \sqrt{\frac{dq \log(\frac{2}{\delta})}{2n}} \left(\sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right) \right).$$

We have to underline that a finite sample analysis is quite important for this problem. In fact, the number of policy improvement steps of the learner is finite as the learner will eventually achieve an optimal policy. So, knowing the finite number of learning improvements m , we can estimate how much data we need for each policy to get an estimate with a certain accuracy. More information about the proof of the theorem can be found in appendix A.

Remark 4.1. *Another important aspect to take into account is that there is an intrinsic bias [18] due to the gradient estimation error that cannot be solved by increasing the number of learning steps, but only with a more accurate estimation of the gradient. However, we show in Section 7 that, experimentally, the component of the bound that does not depend on the number of learning steps does not influence the recovered weights.*

5 Learning from improvement trajectories

In a realistic scenario, the observer has access only to a dataset $\mathcal{D} = (\mathcal{D}_1, \dots, \mathcal{D}_{m+1})$ of trajectories generated by each policy, such that $\mathcal{D}_i = \{\tau_1, \dots, \tau_n\} \sim \pi_{\theta_i}$. Furthermore, the learning rates are unknown and possibly the learner applies an update rule other than (3). The observer has to infer the policy parameters $\Theta = (\theta_1, \dots, \theta_{m+1})$, the learning rates $A = (\alpha_1, \dots, \alpha_m)$, and the reward weights ω . If we suppose that the learner is updating its policy parameters with gradient ascent on the discounted expected return, the natural way to see this problem is to maximize the log-likelihood of $p(\theta_1, \omega, A | \mathcal{D})$:

$$\max_{\theta_1, \omega, A} \sum_{(s,a) \in \mathcal{D}_1} \log \pi_{\theta_1}(a|s) + \sum_{t=2}^{m+1} \sum_{(s,a) \in \mathcal{D}_t} \log \pi_{\theta_t}(a|s),$$

where $\theta_t = \theta_{t-1} + \alpha_{t-1} \nabla_{\theta} \psi_{t-1}$. Unfortunately, solving this problem directly is not practical as it involves evaluating gradients of the discounted expected return up to the m -th order. To deal with this, we break down the inference problem into two steps: the first one consists in recovering the policy parameters Θ of the learner and the second in estimating the learning rates A and the reward weights ω (see Algorithm 1).

5.1 Recovering learner policies

Since we assume that the learner's policy belongs to a parametric policy space Π_{Θ} made of differentiable policies, as explained in [24], we can recover an approximation of the learner's parameters Θ through behavioural cloning, exploiting the trajectories in $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_{m+1}\}$. For each dataset $\mathcal{D}_i \in \mathcal{D}$ of trajectories, we cast the problem of finding the parameter θ_i to a maximum-likelihood estimation. Solving the following optimization problem we obtain an estimate $\hat{\theta}_i$ of θ_i :

$$\max_{\theta_i \in \Theta} \frac{1}{n} \sum_{l=1}^n \sum_{t=0}^{T-1} \log \pi_{\theta_i}(a_{l,t} | s_{l,t}). \quad (6)$$

It is known that the maximum-likelihood estimator is consistent under mild regularity conditions on the policy space Π_{Θ} and assuming the identifiability property [8]. Some finite-sample guarantees on the concentration of distance $\|\hat{\theta}_i - \theta_i\|_p$ were also derived under stronger assumptions, e.g., in [34].

5.2 Recovering learning rates and reward weights

Given the parameters $(\hat{\theta}_1, \dots, \hat{\theta}_{m+1})$, if the learner is updating her policy with a constant learning rate we can simply apply Eq. (4). On the other hand, with an unknown learner, we cannot make this assumption and it is necessary to estimate also the learning rates $A = (\alpha_1, \dots, \alpha_m)$. The optimization problem in Eq. (4) becomes:

$$\min_{\omega \in \mathbb{R}^q, A \in \mathbb{R}^m} \sum_{t=1}^m \left\| \hat{\Delta}_t - \alpha_t \widehat{\nabla_{\theta} \psi}_t \omega \right\|_2^2 \quad (7)$$

$$\text{s.t. } \alpha_t \geq \epsilon \quad 1 \leq t \leq m. \quad (8)$$

where $\hat{\Delta}_t = \hat{\theta}_{t+1} - \hat{\theta}_t$ and ϵ is a small constant. To optimize this function we use alternate block-coordinate descent [37]. We alternate the optimization of parameters A and the optimization of parameters ω . Furthermore, we can notice that these two steps can be solved in closed form. When we optimize on ω , the optimization can be done using Lemma 4.1. When we optimize on A we can solve for each parameter α_t , with $1 \leq t \leq m$, in closed form.

Lemma 5.1. *The minimum of (7) with respect to α_t is equal to:*

$$\hat{\alpha}_t = \max \left(\epsilon, \left((\widehat{\nabla_{\theta} \psi}_t \omega)^T (\widehat{\nabla_{\theta} \psi}_t \omega) \right)^{-1} (\widehat{\nabla_{\theta} \psi}_t \omega)^T \hat{\Delta}_t \right). \quad (9)$$

The inner matrix cannot be inverted only if the vector $\widehat{\nabla_{\theta} \psi} \omega$ is equal to $\mathbf{0}$. This would happen only if the expert is at a stationary point, so $\widehat{\nabla_{\theta} \psi}$ is $\mathbf{0}$. The optimization converges under the assumption that there exists a unique minimum for each variable A and ω [37].

Algorithm 1 LOGEL

Require: Dataset $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_{m+1}\}$ with $\mathcal{D}_j = \{(\tau_1, \dots, \tau_{n_j}) \mid \tau_i \sim \pi_{\theta_j}\}$ **Ensure:** Reward weights $\omega \in \mathbb{R}^q$

- 1: Estimate policy parameters $(\hat{\theta}_1, \dots, \hat{\theta}_{m+1})$ with Eq. (6)
 - 2: Initialize A and ω
 - 3: Compute learning rates A and reward weights ω by alternating (9) and (5) up to convergence
-

5.3 Theoretical analysis

In this section, we provide a finite-sample analysis of LOGEL when only one learning step is observed, assuming that the Jacobian matrix $\widehat{\nabla_{\theta} \psi}$ is bounded and the learner’s policy is a Gaussian policy $\pi \sim \mathcal{N}(\theta^T \varphi(s), \sigma^2)$. The analysis evaluates the norm of the difference between the learner’s weights ω^L and the recovered weights $\hat{\omega}$. Without loss of generality, we consider the case where the learning rate $\alpha = 1$. The analysis takes into account the bias introduced by the behavioural cloning and the gradient estimation.

Theorem 5.1. *Let $\pi_{\theta_1}, \pi_{\theta_2}$ be two Gaussian policies $\pi_{\theta_i}(\cdot|s) \sim \mathcal{N}(\theta_i^T \varphi(s), \sigma^2)$ with $i \in \{1, 2\}$, such that π_{θ_2} is the improvement of π_{θ_1} . Let $i \in [1, 2]$. Given datasets $D_i = \{\tau_1^i, \dots, \tau_n^i\}$ of trajectories generated by π_i , such that $S_i \in \mathbb{R}^{n \times t \times d}$ is the matrix of corresponding states features, let the minimum singular value of $\sigma_{\min}(S_i^T S_i) \geq \eta > 0$, $\widehat{\nabla_{\theta_1} \psi}$ uniformly bounded by M , the state features bounded by M_S , and the reward features bounded by M_R . Then with probability $1 - \delta$:*

$$\|\omega^L - \hat{\omega}\|_2 \leq O\left(\frac{(M + M_S^2 M_R)}{\sigma_{\min}(\nabla_{\theta_1} \psi)} \sqrt{\frac{\log(\frac{2}{\delta})}{n\eta}}\right)$$

where ω^L are the real reward parameters and $\hat{\omega}$ are the parameters recovered using Lemma 4.1.

The theorem, that relies on perturbation analysis [39] and least squares with fixed design [29], underlines how LOGEL, with a sufficient number of samples to estimate the policy parameters and the gradients, succeeds in recovering the correct reward parameters.

6 Related Works

The problem of estimating the reward function of an agent who is learning is quite new. This setting was proposed by Jacq et al. [17] and, to the best of our knowledge, it is studied only in that work. In [17] the authors proposed a method based on entropy-regularized reinforcement learning, in which they assumed that the learner is performing soft policy improvements. In order to derive their algorithm, the authors also assume that the learner respects the policy improvement condition. We do not make this assumption as our formulation assumes only that the learner is changing its policy parameters along the gradient direction (which can result in a performance loss).

The problem, as we underlined in Section 3, is close to the Inverse Reinforcement Learning problem [21, 20], since they share the intention of acquiring the unknown reward function from the observation of an agent’s demonstrations. LOGEL relies on the assumption that the learner is improving her policy through gradient ascent updates. A similar approach, but in the expert case, was taken in [25, 19] where the authors use the null gradient assumption to learn the reward from expert’s demonstrations.

In another line of works, sub-optimal demonstrations are used in the preference-based IRL [11, 16] and ranking-based IRL [7, 9]. Some of these works require that the algorithm asks a human to compare possible agent’s trajectories to learn the underlying reward function of the task. We can imagine that LOGEL can be used in a similar way to learn from humans who are learning a new task. Instead, in [4] was proposed an Imitation Learning setting where the observer tries to imitate the behavior of a supervisor that demonstrates a converging sequence of policies.

In works on *theory of minds* [28, 33], the authors propose an algorithm that uses meta-learning to build a system that learns how to model other agents. In these works it is not required that agents are experts but they must be stationary. Instead, in the setting considered by LOGEL, the observed agent is non-stationary.




Figure 1: Gridworld experiment with known policy parameters. The learner is using G(PO)MDP. From left the expected discounted return and the norm difference between the real weights and the recovered ones with one learning step; the same measures with fixed batch size (5 trajectories with length 20). The performance of the observers are evaluated on the learner’s reward weights. Results are averaged over 20 runs. 98% c.i as shaded area.

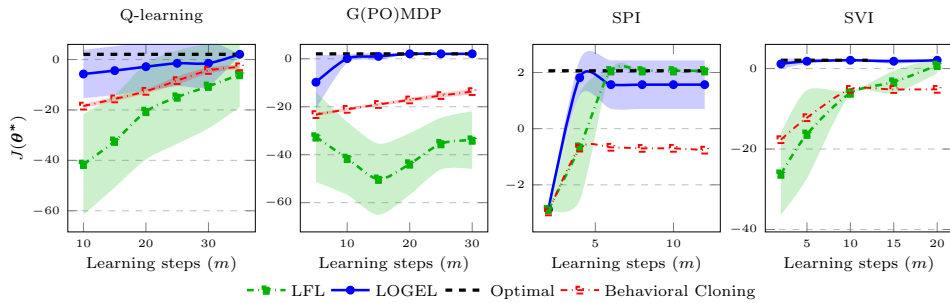


Figure 3: Gridworld experiment with estimated policy parameters and four learners: from left Q-learning, G(PO)MDP, SPI, SVI. The green line is the LfL observer, the blue one is the LOGEL observer and the red one Behavioral Cloning. The performance of the observers are evaluated on the learner’s reward weights. Results are averaged over 20 runs. 98% c.i. as shaded area.

7 Experiments

This section is devoted to the experimental evaluation of LOGEL. The algorithm LOGEL is compared to the state-of-the-art baseline Learner From a Learner (LfL) [17] and T-REX [7] in a gridworld navigation task and in two MuJoCo environments. In these experiments the assumption that the learner is gradient-based is violated and in the MuJoCo task the reward features are constructed by states and actions features. Therefore we can argue that in this experiment the reward linearity assumption is violated, since we use a different reward space for the recovered reward function. More details on the experiments are in Appendix B.1.

7.1 Gridworld

The first set of experiments aims at evaluating the performance of LOGEL in a discrete Gridworld environment. The Gridworld, represented in Figure 2, is composed of five regions with a different reward for each area. The agent starts from the cell top left and when she reaches the green state, then returns to the starting state. The reward feature space is composed of the one-hot encoding of five areas: the orange, the light grey, the dark grey, the blue, and the green. The learner weights for the areas are $(-3, -1, -5, 7, 0)$ respectively. As a first experiment, we want to verify in practice the theoretical finding exposed in Section 4. In this experiment, the learner uses a Boltzmann policy and she is learning with the G(PO)MDP policy gradient algorithm. The observer has access to the true policy parameters of the learner. Figure 1 shows the performance of LOGEL in two settings: a single learning step and increasing batch size (5, 10, 20, 30, 40, 50); a fixed batch size (batch size 5 and trajectory length 20) and an increasing number




Figure 2: Gridworld environment: every area has a different reward weight. In the green area the agent is reset to the starting state.

of learning steps (2, 4, 6, 8, 10). The figure shows the expected discounted return (evaluated in closed form) and the difference in norm between the learner’s weights and the recovered weights ¹. We note that, as explained in Theorem 4.1, with a more accurate gradient estimate, the observer succeeds in recovering the reward weights by observing even just one learning step. On the other hand, as we can deduce from Theorem 4.1, if we have a noisy estimation of the gradient, with multiple learning steps, the observer succeeds in recovering the learner’s weights. It is interesting to notice that, from this experiment, it seems that the bias component, which does not vanish as the learning steps increase (see Theorem 4.1), does not affect the correctness of the recovered weights.

In the second experiment we consider four different learners using: Q-learning [35], G(PO)MDP, Soft policy improvement (SPI) [17] and Soft Value Iteration (SVI) [14] ². For this experiment, we compare the performance of LOGEL, LfL [17] and Behavioral Cloning. In Figure 3 we can notice as LOGEL succeeds in recovering the learner’s reward weights even with learner algorithms other than gradient-based ones. Instead, LfL does not recover the reward weights of the G(PO)MDP learner and needs more learning steps than LOGEL to learn the reward weights when Q-learning learner and SVI are observed. Behavioral Cloning only mimics the last seen policy which can be suboptimal.

7.2 MuJoCo environments

In the second set of experiments, we show the ability of LOGEL to infer the reward weights in more complex and continuous environments. We use two environments from the MuJoCo control suite [6]: Hopper and Reacher. As in [17], the learner is trained using Policy Proximal Optimization (PPO) [31]³, with 16 parallel agents for each learning step. For each step, the length of the trajectories is 2000. Then we use LOGEL, LfL or T-REX [7] to recover the reward parameters. In the end, the observer is trained with the recovered weights using PPO and the performances are evaluated on the learner’s weights, starting from the same initial policy of the learner for a fair comparison. The scores are normalized by setting to 1 the score of the last observed learner policy and to 0 the score of the initial one (as done in [17]). In both environments, the observer learns using the learning steps from 10 to 20 as the first learning steps are too noisy. The reward function of LfL is the same as the one used in the original paper, where the reward function is a neural network equal to the one used for the learner’s policy.

Instead, for LOGEL we used linear reward functions derived only from state and action features. The reward for the Reacher environment is a 26-grid radial basis function that describes the distance between the agent and the goal, plus the 2-norm squared of the action. In the Hopper environment, instead, the reward features are the distance between the previous and the current position and the 2-norm squared of the action. The T-REX algorithm aims to recover a reward function from ranked trajectories, where the rank is given by an oracle and is based on the expected discounted return. We use the algorithm in the LfL setting, where we approximate the ranking with the temporal updates of the policies, as was done in an example in the original paper. We implement the reward function as in the original paper with a three layer neural network with 256 as hidden size.

The results are shown in Figure 4, where we reported results averaged over 10 runs. For Behavioral Cloning we report the performance of the policy learnt at the 20th step; in fact Behavioral Cloning cannot improve its performance with learning steps. We can notice that LOGEL succeeds in




Figure 4: From the left, the Reacher and the Hopper MuJoCo environments. The red line is the performance of the learner during 20 learning steps. The observers, LfL, LOGEL, T-REX and Behavioral Cloning (BC) observe the trajectories of the learning steps from 10 to 20. The performance of the observers are evaluated on the learner’s reward weights. Scores are normalized setting to 0 the first return of the learner and to 1 the last one. The results are averaged over 10 runs. 98% c.i. are shown as shaded areas.

¹To perform this comparison, we normalize the recovered weights and the learner’s weights

²In Appendix B.1 the learning process of each learning agent is shown.

³It is important to notice that PPO violates the gradient learning assumption of LOGEL.

identifying a good reward function in both environments, although in the Reacher environment the recovered reward function causes slower learning. Instead, LfL fails to recover an effective reward function for the Hopper environment [17]. The T-REX algorithm, as in the original paper, succeeds in recovering a good approximation of the reward weights in the Hopper domain; instead, it does not succeed into recovering the reward function of the Reacher environment.

8 Conclusions

In this paper we propose a novel algorithm, LOGEL, for the “Learning from a Learner Inverse Reinforcement Learning” setting. The proposed method relies on the assumption that the learner updates her policy along the direction of the gradient of the expected discounted return. We provide some finite-sample bounds on the algorithm performance in recovering the reward weights when the observer observes the learner’s policy parameters and when the observer observes only the learner’s trajectories. Finally, we tested LOGEL on a discrete gridworld environment and on two MuJoCo continuous environments, comparing the algorithm with the state-of-the-art baselines [17, 7]. As future work, we plan to extend the algorithm to account for the uncertainty in estimating both the policy parameters and the gradient.

Broader impact

In this paper, we focus on the Inverse Reinforcement Learning [2, 15, 3, 21] task from a Learning Agent [17]. The first motivation to study Inverse Reinforcement Learning algorithms is to overcome the difficulties that can arise in specifying the reward function from human and animal behaviour. Sometimes, in fact, it is easier to infer human intentions by observing their behaviours than to design a reward function by hand. An example is helicopter flight control [1], in which we can observe a helicopter operator and through IRL a reward function is inferred to teach a physical remote-controlled helicopter. Another example is to predict the behavior of a real agent as route prediction tasks of taxis [41, 42] or anticipation of pedestrian interactions [12] or energy-efficient driving [38]. However, in many cases, the agents are not really experts and on the other hand, only expert demonstrations can not show their intention to avoid dangerous situations. We want to point out that learning what the agent wants to avoid because harmful is as important as learning his intentions.

The possible outcomes of this research are the same as those of Inverse Reinforcement Learning mentioned above, avoiding the constraint that the agent has to be an expert. In future work, we will study how to apply the proposed algorithm in order to infer the pilot’s intentions when they learn a new circuit.

A relevant possible complication of using IRL is the error on the reward feature engineering which can lead to errors in understanding the agent’s intentions. In an application such as autonomous driving, errors in the reward function can cause dangerous situations. For this reason, verification through the simulated environment of the effectiveness of the retrieve rewards is quite important.

Funding Transparency Statement

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Market”.

References

- [1] Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y Ng. An application of reinforcement learning to aerobatic helicopter flight. In *Advances in neural information processing systems*, pages 1–8, 2007.
- [2] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [3] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.

- [4] Ashwin Balakrishna, Brijen Thananjeyan, Jonathan Lee, Felix Li, Arsh Zahed, Joseph E. Gonzalez, and Ken Goldberg. On-policy robot imitation learning from a converging supervisor. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, volume 100 of *Proceedings of Machine Learning Research*, pages 24–41. PMLR, 2019.
- [5] Jonathan Baxter and Peter L Bartlett. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350, 2001.
- [6] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *CoRR*, abs/1606.01540, 2016.
- [7] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. Extrapolating beyond sub-optimal demonstrations via inverse reinforcement learning from observations. In *International Conference on Machine Learning*, pages 783–792, 2019.
- [8] George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA, 2002.
- [9] Pablo Samuel Castro, Shijian Li, and Daqing Zhang. Inverse reinforcement learning with multiple ranked experts. *arXiv preprint arXiv:1907.13411*, 2019.
- [10] Yudong Chen and Constantine Caramanis. Noisy and missing data regression: Distribution-oblivious support recovery. In *International Conference on Machine Learning*, pages 383–391, 2013.
- [11] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pages 4299–4307, 2017.
- [12] Shu-Yun Chung and Han-Pang Huang. A mobile robot that understands pedestrian spatial behaviors. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5861–5866. IEEE, 2010.
- [13] Vineet Goyal and Julien Grand-Clement. A first-order approach to accelerated value iteration. *arXiv preprint arXiv:1905.09963*, 2019.
- [14] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [15] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):21, 2017.
- [16] Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward learning from human preferences and demonstrations in atari. In *Advances in Neural Information Processing Systems*, pages 8011–8023, 2018.
- [17] Alexis Jacq, Matthieu Geist, Ana Paiva, and Olivier Pietquin. Learning from a learner. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2990–2999, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [18] Brian McWilliams, Gabriel Krummenacher, Mario Lucic, and Joachim M Buhmann. Fast and robust least squares estimation in corrupted linear models. In *Advances in Neural Information Processing Systems*, pages 415–423, 2014.
- [19] Alberto Maria Metelli, Matteo Pirota, and Marcello Restelli. Compatible reward inverse reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2050–2059, 2017.
- [20] Andrew Y. Ng and Stuart J. Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford University, Stanford, CA, USA, June 29 - July 2, 2000, pages 663–670. Morgan Kaufmann, 2000.
- [21] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 7(1-2):1–179, 2018.

- [22] Jan Peters and Stefan Schaal. Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2219–2225. IEEE, 2006.
- [23] Jan Peters and Stefan Schaal. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4):682–697, 2008.
- [24] Matteo Pirotta and Marcello Restelli. Inverse reinforcement learning through policy gradient minimization. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 1993–1999. AAAI Press, 2016.
- [25] Matteo Pirotta, Marcello Restelli, and Luca Bascetta. Adaptive step-size for policy gradient methods. In *Advances in Neural Information Processing Systems*, pages 1394–1402, 2013.
- [26] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.
- [27] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1994.
- [28] Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, SM Ali Eslami, and Matthew Botvinick. Machine theory of mind. In *International Conference on Machine Learning*, pages 4218–4227, 2018.
- [29] Philippe Rigollet. High-dimensional statistics. spring 2015. *Massachusetts Institute of Technology: MIT OpenCourseWare*, 2015.
- [30] John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.
- [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [32] Hanan Shteingart and Yonatan Loewenstein. Reinforcement learning and human behavior. *Current Opinion in Neurobiology*, 25:93–98, 2014.
- [33] Michael Shum, Max Kleiman-Weiner, Michael L Littman, and Joshua B Tenenbaum. Theory of minds: Understanding behavior in groups through inverse planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6163–6170, 2019.
- [34] Vladimir Spokoiny et al. Parametric estimation. finite sample theory. *The Annals of Statistics*, 40(6):2877–2909, 2012.
- [35] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford book. Bradford Book, 1998.
- [36] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, pages 1057–1063, 2000.
- [37] Paul Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494, 2001.
- [38] Adam Vogel, Deepak Ramachandran, Rakesh Gupta, and Antoine Raux. Improving hybrid vehicle fuel efficiency using inverse reinforcement learning. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.
- [39] Perke Wedin. Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics*, 13(2):217–232, 1973.
- [40] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [41] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In Dieter Fox and Carla P. Gomes, editors, *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*, pages 1433–1438. AAAI Press, 2008.
- [42] Brian D Ziebart, Andrew L Maas, Anind K Dey, and J Andrew Bagnell. Navigate like a cabbie: Probabilistic reasoning from observed context-aware behavior. In *Proceedings of the 10th international conference on Ubiquitous computing*, pages 322–331, 2008.

A Proofs and derivations

We start the proofs given some introduction on Perturbation on Least Square problems and on Least Square problems with fixed design. Then we report the proofs and derivations for the results of Sections 4 and 5. For the rest of this section we assume that:

- $\omega^L, \hat{\omega} \in \mathbb{R}^q$,
- $\theta^L, \hat{\theta} \in \mathbb{R}^d$.

We define with ω^L (θ^L) the reward (policy) parameters of the learner, and with $\hat{\omega}$ ($\hat{\theta}$) the reward (policy) parameters recovered by the observer.

A.1 Preliminaries

Definition A.1. *The condition number of a matrix $\mathbf{A} \in \mathbb{R}^{m \times q}$ $\mathbf{A} \neq 0$ is:*

$$\kappa = \|\mathbf{A}\|_2 \|\mathbf{A}^+\|_2 = \frac{\sigma_1}{\sigma_r},$$

where $0 < r = \text{rank}(\mathbf{A}) \leq \min(m, q)$, and $\sigma_1 \geq \dots \geq \sigma_r > 0$ are the nonzero singular values of \mathbf{A} .

A least-squares problem is defined as:

$$\min_x \|\mathbf{A}x - b\|_2, \quad (10)$$

where the solution is $x = \mathbf{A}^+b$. We denote with \mathbf{A}^+ the pseudoinverse of \mathbf{A} , the perturbed \mathbf{A} as $\hat{\mathbf{A}} = \mathbf{A} + \delta\mathbf{A}$ and the perturbed b as $\hat{b} = b + \delta b$ and the perturbed solution $\hat{x} = \hat{\mathbf{A}}^+\hat{b} = x + \delta x$. Finally, we denote with \mathbf{A}^H the adjoint of the matrix \mathbf{A} .

We define as $\chi = \frac{\|\delta\mathbf{A}\|_2}{\|\mathbf{A}\|_2}$ and $y = \mathbf{A}^H x$.

Lemma A.1 (Perturbation on Least Square Problems [39]). *Assume that $\text{rank}(\mathbf{A} + \delta\mathbf{A}) = \text{rank}(\mathbf{A})$ and $\chi\kappa < 1$ then:*

$$\|x - \hat{x}\|_2 \leq \frac{\kappa}{(1 - \chi\kappa)} (\chi \|x\|_2 \|\mathbf{A}\|_2 + \chi\kappa \|r\|_2 + \|\delta b\|_2) + \chi \|y\|_2 \|\mathbf{A}\|_2. \quad (11)$$

Proof. The proof can be find in [39]. □

We adapt the lemma 6 in [10] to our context where $\hat{\omega}$ are the reward weights recovered with lemma 4.1.

Lemma A.2 (From lemma 6 in [10]). *Let $\Sigma = (\widehat{\nabla_{\theta}\psi}^T \widehat{\nabla_{\theta}\psi})$ and suppose the following strong convexity condition holds: $\lambda_{\min}(\Sigma) \geq \lambda > 0$. Then the estimation error satisfies:*

$$\|\hat{\omega} - \omega^L\|_2 \leq O\left(\frac{1}{\lambda} \|\nabla_{\theta}\psi^T \Delta - \Sigma\omega^L\|_2\right).$$

Lemma A.3 (Revised from lemma 11 in [18]). *Suppose $X \in \mathbb{R}^{m \times q}$ and $W \in \mathbb{R}^{n \times M}$ are zero-mean sub-gaussian matrices with parameters $(\frac{1}{n}\Sigma_x, \frac{1}{n}\sigma_x^2)$, $(\frac{1}{n}\Sigma_w, \frac{1}{n}\sigma_w^2)$ respectively. Then for any fixed vectors v_1, v_2 , we have:*

$$P[|v_1^T (W^T X - \mathbb{E}[W^T X]) v_2| \geq t \|v_1\|_2 \|v_2\|_2] \leq 3 \exp\left(-cn \min\left\{\frac{t^2}{\sigma_x^2 \sigma_w^2}, \frac{t}{\sigma_x \sigma_w}\right\}\right),$$

in particular if $n \gtrsim \log p$ we have that:

$$|v_1^T (W^T X - \mathbb{E}[W^T X]) v_2| \leq \sigma_x \sigma_w \|v_1\|_2 \|v_2\|_2 \sqrt{\frac{\log p}{n}}.$$

Setting v_1 to be the first standard basis vector and using a union bound over $j = 1, \dots, p$ we have:

$$\|(W^T X - \mathbb{E}[W^T X])v\|_{\infty} \leq \sigma_x \sigma_w \|v\|_2 \sqrt{\frac{\log p}{n}},$$

with probability $1 - c_1 \exp(-c_2 \log p)$ where c_1, c_2 are positive constants which are independent from σ_x, σ_w, n, p .

Theorem A.1 (from Chapter 2 [29]). Assume that the least-squares model:

$$\min_x \|\mathbf{A}x - b + \epsilon\|$$

holds where $\epsilon \sim \text{subGn}(\sigma^2)$. Then, for any $\delta > 0$, with probability $1 - \delta$ it holds:

$$\|x - \hat{x}\|_2 \leq \sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\sigma_{\min}}},$$

where $\sigma_{\min} = \frac{\mathbf{A}^T \mathbf{A}}{n}$ is the minimum singular value of $\mathbf{A}^T \mathbf{A}$ and r is the rank($\mathbf{A}^T \mathbf{A}$).

A.2 Additional Results

In this section we give the proofs and derivations of the theorems in Section 4.

First, we will provide a finite sample analysis on the difference in norm between the reward vector of the learner ω^L and the reward vector recovered using (5), with a single learning step. This result was omitted in the main paper as we can see this as a special case of Theorem 4.1, but with a different technique. We add it here as it provides a first insight on how, having enough demonstrations, we can recover the correct weights. In the demonstration, without loss of generality, we assume that the learning rate is 1.

Lemma A.4. Let $\nabla_{\theta}\psi$ be the real Jacobian and $\widehat{\nabla_{\theta}\psi}$ the estimated Jacobian from n trajectories $\{\tau_1, \dots, \tau_n\}$. Assume that $\widehat{\nabla_{\theta}\psi}$ is uniformly bounded by M . Then with probability $1 - \delta$

$$\|\widehat{\nabla_{\theta}\psi} - \nabla_{\theta}\psi\|_2 \leq M\sqrt{qd} \sqrt{\frac{\log(\frac{2}{\delta})}{2n}}.$$

Proof. We use Hoeffding's inequality:

$$\mathbb{P} \left[\|\widehat{\nabla_{\theta}\psi} - \nabla_{\theta}\psi\|_2 \geq t \right] \leq \mathbb{P} \left[\sqrt{qd} \|\widehat{\nabla_{\theta}\psi} - \nabla_{\theta}\psi\|_{\infty} \geq t \right] \leq 2 \exp \left(\frac{-2t^2 n}{dqM^2} \right)$$

The result follows by setting $\delta = 2 \exp \left(\frac{-2t^2 n}{dqM^2} \right)$. \square

Theorem A.1. Let $\nabla_{\theta}\psi$ be the real Jacobian and $\widehat{\nabla_{\theta}\psi}$ the estimated Jacobian from n trajectories $\{\tau_1, \dots, \tau_n\}$. Assume that $\widehat{\nabla_{\theta}\psi}$ is uniformly bounded by M , $\text{rank}(\widehat{\nabla_{\theta}\psi}) = \text{rank}(\nabla_{\theta}\psi)$ and $\|\widehat{\nabla_{\theta}\psi} - \nabla_{\theta}\psi\|_2 \cdot \kappa_{\nabla_{\theta}\psi} < \|\nabla_{\theta}\psi\|_2$. Then with probability $1 - \delta$:

$$\|\omega^L - \hat{\omega}\|_2 \leq M\sqrt{qd} \sqrt{\frac{\log(\frac{2}{\delta})}{2n}} \left(\frac{\kappa_{\nabla_{\theta}\psi} \|\omega^L\|_2}{c \|\nabla_{\theta}\psi\|_2} + \|y\|_2 \right), \quad (12)$$

where ω^L are the real reward parameters and $\hat{\omega}$ are the parameters recovered with Equation (5), $c = 1 - \frac{\|\widehat{\nabla_{\theta}\psi} - \nabla_{\theta}\psi\|_2}{\|\nabla_{\theta}\psi\|_2} \kappa_{\nabla_{\theta}\psi} > 0$, and $y = \nabla_{\theta}\psi^{+H} \omega$.

Proof. We need to bound the difference in norm between ω^L and $\hat{\omega}$ that are the true parameters and the parameters that we recovered solving the minimization problem (4).

$$\|\omega^L - \hat{\omega}\|_2 \quad (13)$$

$$\leq \frac{\kappa}{\left(1 - \kappa \frac{\|\delta \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2}\right) \|\nabla_{\theta} \psi\|_2} \left(\frac{\|\delta \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2} \|\omega^L\|_2 \|\nabla_{\theta} \psi\|_2 \right) + \frac{\|\delta \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2} \|y\|_2 \|\nabla_{\theta} \psi\|_2 \quad (14)$$

$$\leq \frac{\kappa}{c \|\nabla_{\theta} \psi\|_2} \left(\frac{\|\delta \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2} \|\omega^L\|_2 \|\nabla_{\theta} \psi\|_2 \right) + \frac{\|\delta \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2} \|y\|_2 \|\nabla_{\theta} \psi\|_2 \quad (15)$$

$$= \|\delta \nabla_{\theta} \psi\|_2 \left(\frac{\kappa \|\omega^L\|_2}{c \|\nabla_{\theta} \psi\|_2} + \|y\|_2 \right) \quad (16)$$

$$\leq M \sqrt{qd} \sqrt{\frac{\log(\frac{2}{\delta})}{2n}} \left(\frac{\kappa \|\omega^L\|_2}{c \|\nabla_{\theta} \psi\|_2} + \|y\|_2 \right), \quad (17)$$

where line (14) is obtained by using Lemma A.1, lines (15), (16) by rearranging the terms, and line (17) by using Lemma A.4. We can observe that the last term vanishes when the rank($\nabla_{\theta} \psi$) = q (see [39]). \square

A.3 Proofs and derivation of Section 4

Now we will give the proofs and derivations of Lemmas 4.1 and Theorem 4.1.

Lemma 4.1. *If the matrix $(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t)^{-1}$ is full-rank then optimization problem (4) is solved in closed form by*

$$\hat{\omega} = \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right)^{-1} \left(\sum_{t=1}^m \alpha_t \nabla_{\theta} \psi_t^T \Delta_t \right). \quad (5)$$

Proof. Taking the derivative of (4) with respect to ω :

$$\begin{aligned} \nabla_{\omega} \sum_{t=1}^m \|\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega\|_2^2 &= \sum_{t=1}^m \nabla_{\omega} (\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega)^T (\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega) \\ &= \sum_{t=1}^m \nabla_{\omega} (\Delta_t^T \Delta_t + (\alpha_t \nabla_{\theta} \psi_t \omega)^T (\alpha_t \nabla_{\theta} \psi_t \omega) - 2\alpha_t \nabla_{\theta} \psi_t \omega)^T \Delta_t \\ &= 2 \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right) \omega - 2 \sum_{t=1}^m (\alpha_t \nabla_{\theta} \psi_t^T \Delta_t). \end{aligned}$$

Taking it equal to zero:

$$\begin{aligned} \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right) \omega - \sum_{t=1}^m (\alpha_t \nabla_{\theta} \psi_t^T \Delta_t) &= 0 \\ \omega &= \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right)^{-1} \left(\sum_{t=1}^m \alpha_t \nabla_{\theta} \psi_t^T \Delta_t \right) \end{aligned}$$

\square

Lemma A.5. *The regularized version of (4) is equal to:*

$$\min_{\omega} \sum_{t=1}^m \|\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega\|_2^2 + \lambda \|\omega\|_2^2,$$

where $\lambda > 0$. We can solve the regularized problem in closed form:

$$\omega = \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t + \lambda \mathbf{I}_d \right)^{-1} \left(\sum_{t=1}^m \alpha_t \nabla_{\theta} \psi_t^T \Delta_t \right).$$

Proof. Taking the derivative respect to ω :

$$\begin{aligned} \nabla_{\omega} \sum_{t=1}^m \|\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega\|_2^2 + \lambda \|\omega\|_2^2 &= \sum_{t=1}^m \nabla_{\omega} (\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega)^T (\Delta_t - \alpha_t \nabla_{\theta} \psi_t \omega) + \nabla_{\omega} \lambda \omega^T \omega \\ &= \sum_{t=1}^m \nabla_{\omega} (\Delta_t^T \Delta_t + (\alpha_t \nabla_{\theta} \psi_t \omega)^T (\alpha_t \nabla_{\theta} \psi_t \omega) - 2\alpha_t \nabla_{\theta} \psi_t \omega^T \Delta_t) + 2\lambda \omega \\ &= 2 \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t \right) \omega - 2 \sum_{t=1}^m (\alpha_t \nabla_{\theta} \psi_t^T \Delta_t) + 2\lambda \omega. \end{aligned}$$

Taking it equal to zero:

$$\begin{aligned} \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t + \lambda \mathbf{I}_d \right) \omega - \sum_{t=1}^m (\alpha_t \nabla_{\theta} \psi_t^T \Delta_t) &= 0 \\ \omega &= \left(\sum_{t=1}^m \alpha_t^2 \nabla_{\theta} \psi_t^T \nabla_{\theta} \psi_t + \lambda \mathbf{I}_d \right)^{-1} \left(\sum_{t=1}^m \alpha_t \nabla_{\theta} \psi_t^T \Delta_t \right). \end{aligned}$$

□

Theorem 4.1. Let Ψ be the real Jacobians and $\widehat{\Psi}$ the estimated Jacobian from n trajectories $\{\tau_1, \dots, \tau_n\}$. Assume that Ψ is bounded by a constant M and $\lambda_{\min}(\widehat{\Psi}^T \widehat{\Psi}) \geq \lambda > 0$. Then w.h.p.:

$$\|\omega^L - \widehat{\omega}\|_2 \leq O \left(\frac{1}{\lambda} M \sqrt{\frac{dq \log(\frac{2}{\delta})}{2n}} \left(\sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right) \right).$$

Proof. We decompose the estimated Jacobian $\Psi = \Psi + E$, where E is the random variable component caused by the estimation of the $\nabla_{\theta} \psi$. Since we estimate the jacobians with an unbiased estimator the mean of E is 0. We reshape the Ψ and E as $\Psi \in \mathbb{R}^{m \times dq}$ and $E \in \mathbb{R}^{m \times dq}$. Now E , since its mean is 0 and all lines are independent of each other, is a sub-Gaussian matrix with parameters $(\frac{1}{m} \Sigma_E, \frac{1}{m} \sigma_E)$. The proof is similar to the proof of Theorem 1 in [18].

$$\begin{aligned} &\left\| (\Psi + E)^T (\Psi^T \omega^L) - (\Psi + E)^T (\Psi + E) \omega^L \right\|_2 \\ &= \left\| \Psi^T \nabla_{\theta} \psi^T \omega^* + E^T \Psi^T \omega^L - \Psi^T \Psi \omega^L - \Psi^T E \omega^* - E \Psi^T \omega^* - E^T E \omega^L \right\|_2 \\ &= \left\| -\Psi^T E \omega^L - E^T E \omega^L \right\|_2. \end{aligned}$$

Now we bound separately these two terms, using Lemma A.3 as in [18]:

$$\begin{aligned} \left\| \Psi^T E \omega^L \right\|_2 &\leq \|\Psi\|_2 \sigma_E \|\omega^L\|_2 \sqrt{\frac{\log dq}{m}} \\ \left\| E^T E \omega^L \right\|_2 &= \left\| (E^T E + \sigma_E^2 \mathbf{I}_{qd} - \sigma_E^2 \mathbf{I}_{qd}) \omega^L \right\|_2 \leq \sigma_E^2 \left(C \sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right) \|\omega^L\|_2 \end{aligned}$$

with probability $1 - c_1 \exp(-c_2 \log q)$ where c_1, c_2 are positive constants that do not depend on σ_E, n, q . So now applying Lemma A.2:

$$\|\omega^L - \widehat{\omega}\|_2 \leq \frac{1}{\lambda} \left(\|\Psi\|_2 \sigma_E \|\omega^L\|_2 \sqrt{\frac{\log dq}{m}} + \sigma_E^2 \left(C \sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right) \|\omega^L\|_2 \right)$$

w.h.p..

Now we need to bound the random variable σ_E . Remember that $E_i = \nabla_{\theta} \psi_i - \widehat{\nabla_{\theta} \psi_i}$. Since $\widehat{\nabla_{\theta} \psi}$ are assumed to be bounded by M , by applying Hoeffding's inequality, with probability $1 - \delta_1$:

$$\|E_i\|_2 = \left\| \widehat{\nabla_{\theta} \psi_i} - \nabla_{\theta} \psi_i \right\|_2 \leq M \sqrt{\frac{dq \log(\frac{2}{\delta_1})}{2n}}.$$

So E is a subgaussian random variable where each component is bounded by $M\sqrt{\frac{dq \log(\frac{2}{\delta_1})}{2n}}$.

Then:

$$\begin{aligned} \mathbb{P} \left[\|\boldsymbol{\omega}^L - \widehat{\boldsymbol{\omega}}\|_2 \geq \frac{1}{\lambda} M \sqrt{\frac{dq \log(\frac{2}{\delta_1})}{2n}} \|\boldsymbol{\omega}^L\|_2 \|\Psi\|_2 \sqrt{\frac{\log dq}{m}} + M \frac{dq \log(\frac{2}{\delta_1})}{2n} C \sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right] \\ \leq \mathbb{P} \left[\|\boldsymbol{\omega}^L - \widehat{\boldsymbol{\omega}}\|_2 \geq \frac{1}{\lambda} M \sqrt{\frac{dq \log(\frac{2}{\delta_1})}{2n}} \|\boldsymbol{\omega}^L\|_2 \left(\|\Psi\|_2 \sqrt{\frac{\log dq}{m}} + C \sqrt{\frac{\log dq}{m}} + \sqrt{dq} \right) \right] \\ \leq \delta_1 + c1 \exp(-c2 \log dq) \end{aligned}$$

So the result follows where with w.h.p. we mean with probability $1 - (\delta_1 + c1 \exp(-c2 \log dq))$ as in [18]. \square

A.4 Proofs of Section 5

In this section we provide the proofs and derivations of the theorems in Section 5.

Lemma A.6. *Given a dataset $D = \{(s_1, a_1), \dots, (s_n, a_n)\}$ of state-action couples sampled from a Gaussian linear policy $\pi_{\boldsymbol{\theta}}(\cdot|s) \sim \mathcal{N}(\boldsymbol{\theta}^T \boldsymbol{\varphi}(s), \sigma^2)$ such that $S \in \mathbb{R}^{n \times p}$ is the matrix of states features and let the minimum singular value of $(S^T S)$ $\sigma_{\min} \geq \eta$, then the error between the maximum likelihood estimator $\boldsymbol{\theta}^{MLE}$ and the mean $\boldsymbol{\theta}$ is, with probability $1 - \delta$:*

$$\|\boldsymbol{\theta}^{MLE} - \boldsymbol{\theta}\|_2 \leq \sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}},$$

where r is the rank($S^T S$).

Proof. We start by stating that the maximum likelihood for linear Gaussian policies can be recast as an ordinary least-squares problem. We write the Likelihood $L(\boldsymbol{\theta})$

$$\begin{aligned} \log L(\boldsymbol{\theta}) &= \log \left(\prod_{i=1}^n \pi(a_i|s_i) \right) \\ &= \sum_{i=1}^n \log \left(\frac{1}{\sqrt{2\pi\sigma^2}} \exp \left(-\frac{(a_i - \boldsymbol{\theta}^T \boldsymbol{\varphi}(s_i))^2}{2\sigma^2} \right) \right) \\ &= n \log \left(\frac{1}{\sqrt{2\pi\sigma^2}} \right) - \sum_{i=1}^n \frac{(a_i - \boldsymbol{\theta}^T \boldsymbol{\varphi}(s_i))^2}{2\sigma^2} \end{aligned}$$

The resulting maximum likelihood problem is given by:

$$\max_{\boldsymbol{\theta}} \log L(\boldsymbol{\theta}) = \min_{\boldsymbol{\theta}} \sum_{i=1}^n (a_i - \boldsymbol{\theta}^T \boldsymbol{\varphi}(s_i))^2$$

So we have the following linear least-squares problem:

$$\min_{\boldsymbol{\theta}} \|S\boldsymbol{\theta} - A + \boldsymbol{\epsilon}\|_2,$$

where $\boldsymbol{\epsilon}$ is an error with mean 0 and variance σ^2 , $S \in \mathbb{R}^{n \times d}$ is the matrix of states features and $A \in \mathbb{R}^n$ is the vector of actions. Using Theorem A.1, we can say that with probability $1 - \delta$:

$$\|\boldsymbol{\theta}^{MLE} - \boldsymbol{\theta}\|_2 \leq \sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}},$$

where r is the rank($S^T S$). \square

Lemma A.7. *Given two Gaussian policies $\pi_{\boldsymbol{\theta}_1}(\cdot|s) \sim \mathcal{N}(\boldsymbol{\theta}_1^T \boldsymbol{\varphi}(s), \sigma^2)$ and $\pi_{\boldsymbol{\theta}_2}(\cdot|s) \sim \mathcal{N}(\boldsymbol{\theta}_2^T \boldsymbol{\varphi}(s), \sigma^2)$ with same variance and the state features are bounded by M_S :*

$$\|\nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}_1}(a|s) - \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}_2}(a|s)\|_2 \leq \frac{M_S^2}{\sigma^2} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2.$$

Proof. The gradient of the log policy of a general policy $\pi_{\theta}(a|s)$ is:

$$\nabla_{\theta} \log \pi(a|s) = \frac{\varphi(s)^T (a - \theta^T \varphi(s))}{\sigma^2}.$$

Now we apply this result to the difference in norm between two Gaussian log policies:

$$\|\nabla_{\theta} \log \pi_{\theta_1}(a|s) - \nabla_{\theta} \log \pi_{\theta_2}(a|s)\|_2 = \left\| \frac{\varphi(s)^T (a - \theta_1^T \varphi(s))}{\sigma^2} - \frac{\varphi(s)^T (a - \theta_2^T \varphi(s))}{\sigma^2} \right\|_2 \quad (18)$$

$$= \left\| \frac{\varphi(s)}{\sigma^2} (\theta_1^T \varphi(s) - \theta_2^T \varphi(s)) \right\|_2 \quad (19)$$

$$\leq \left\| \frac{\varphi(s)}{\sigma^2} \right\|_2 \|\theta_1 - \theta_2\|_2 \|\varphi(s)\|_2 \quad (20)$$

$$\leq \frac{M_S^2}{\sigma^2} \|\theta_1 - \theta_2\|_2. \quad (21)$$

In line (19) we use the Cauchy-Schwartz inequality, and in line (20) the assumption that the state features are bounded by M_S . \square

Lemma A.8. *Given a dataset $D = \{\tau_1, \dots, \tau_n\}$ of trajectories such that every trajectory $\tau_i = \{(s_1, a_1), \dots, (s_T, a_T)\}$ is sampled from a Gaussian linear policy $\pi_{\theta}(\cdot|s) \sim \mathcal{N}(\theta^T \varphi(s), \sigma)$, the maximum likelihood estimator θ^{MLE} estimated on D , the condition of Lemma A.6 holds, the $\widehat{\nabla_{\theta} \psi}$ uniformly bounded by M , the state features bounded by M_S , the reward features bounded by M_R . Let $S \in \mathbb{R}^{n \times p}$ be the matrix of state features and let $\sigma_{\min}(S^T S) \geq \eta$. Then with probability $1 - \delta$:*

$$\left\| \widehat{\nabla_{\theta} \psi}(\theta^{MLE}) - \nabla_{\theta} \psi(\theta) \right\|_2 \leq M \sqrt{qd} \sqrt{\frac{\log(\frac{2}{\delta})}{2n}} + \frac{T M_S^2 M_R}{(1-\gamma)\sigma} \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}},$$

where γ is the discount factor and r is the rank of $S^T S$.

Proof. We start by decomposing the norm of the difference in two components, using triangular inequality:

$$\left\| \widehat{\nabla_{\theta} \psi}(\hat{\theta}) - \nabla_{\theta} \psi(\theta) \right\|_2 \leq \left\| \widehat{\nabla_{\theta} \psi}(\theta) - \nabla_{\theta} \psi(\theta) \right\|_2 + \left\| \widehat{\nabla_{\theta} \psi}(\theta) - \widehat{\nabla_{\theta} \psi}(\hat{\theta}) \right\|_2.$$

The first component is bounded by Lemma A.4. We will bound now the second component, using Reinforce estimator for the gradient:

$$\left\| \widehat{\nabla_{\theta} \psi}(\theta) - \widehat{\nabla_{\theta} \psi}(\hat{\theta}) \right\|_2 = \quad (22)$$

$$= \left\| \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \nabla_{\theta} \log \pi_{\theta}(a_{i,t}|s_{i,t}) R_{i,t} \gamma^t - \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \nabla_{\theta} \log \pi_{\hat{\theta}}(a_{i,t}|s_{i,t}) R_{i,t} \gamma^t \right\|_2 \quad (23)$$

$$= \frac{1}{n} \left\| \sum_{i=1}^n \sum_{t=1}^T (\nabla_{\theta} \log \pi_{\theta}(a_{i,t}|s_{i,t}) - \nabla_{\theta} \log \pi_{\hat{\theta}}(a_{i,t}|s_{i,t})) R_{i,t} \gamma^t \right\|_2 \quad (24)$$

$$\leq \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left\| (\nabla_{\theta} \log \pi_{\theta}(a_{i,t}|s_{i,t}) - \nabla_{\theta} \log \pi_{\hat{\theta}}(a_{i,t}|s_{i,t})) \right\|_2 \|R_{i,t} \gamma^t\|_2 \quad (25)$$

$$\leq \frac{1}{n} \frac{M_R}{(1-\gamma)} \sum_{i=1}^n \sum_{t=1}^T \frac{M_S^2}{\sigma^2} \|\theta - \hat{\theta}\|_2 \quad (26)$$

$$\leq \frac{T M_S^2 M_R}{\sigma^2 (1-\gamma)} \sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}}. \quad (27)$$

In line (26) we apply the Cauchy-Schwartz inequality. In line (27) we apply lemma A.4 and in line (28) we apply lemma A.4. Merging the two results the proof follows. \square

Theorem 5.1. Let $\pi_{\theta_1}, \pi_{\theta_2}$ be two Gaussian policies $\pi_{\theta_i}(\cdot|s) \sim \mathcal{N}(\theta_i^T \varphi(s), \sigma^2)$ with $i \in \{1, 2\}$, such that π_{θ_2} is the improvement of π_{θ_1} . Let $i \in [1, 2]$. Given datasets $D_i = \{\tau_1^i, \dots, \tau_n^i\}$ of trajectories generated by π_i , such that $S_i \in \mathbb{R}^{n \times t \times d}$ is the matrix of corresponding states features, let the minimum singular value of $\sigma_{\min}(S_i^T S_i) \geq \eta > 0$, $\widehat{\nabla_{\theta_1} \psi}$ uniformly bounded by M , the state features bounded by M_S , and the reward features bounded by M_R . Then with probability $1 - \delta$:

$$\|\omega^L - \widehat{\omega}\|_2 \leq O\left(\frac{(M + M_S^2 M_R)}{\sigma_{\min}(\nabla_{\theta_1} \psi)} \sqrt{\frac{\log(\frac{2}{\delta})}{n\eta}}\right)$$

where ω^L are the real reward parameters and $\widehat{\omega}$ are the parameters recovered using Lemma 4.1.

Proof. First we have to bound the error on Δ created by the behavioral cloning. Given $\Delta = \theta_2 - \theta_1$ and $\widehat{\Delta} = \widehat{\theta}_2 - \widehat{\theta}_1$:

$$\|\Delta - \widehat{\Delta}\| = \|\theta_2 - \theta_1 - \widehat{\theta}_2 + \widehat{\theta}_1\| \leq \|\theta_1 - \widehat{\theta}_1\| + \|\theta_2 - \widehat{\theta}_2\| \leq 2\sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}}. \quad (28)$$

So we can bound the difference in norm between the real weights ω^L and the estimated weights $\widehat{\omega}$.

We indicate with κ the condition number of $\nabla_{\theta} \psi$, with $\chi = \frac{\|\widehat{\nabla_{\theta} \psi} - \nabla_{\theta} \psi\|_2}{\|\nabla_{\theta} \psi\|_2}$ and $y = \nabla_{\theta} \psi +^H \omega$. We apply the perturbation Lemma A.1.

$$\begin{aligned} \|\omega^L - \widehat{\omega}\|_2 &\leq \frac{\kappa}{(1 - \kappa\chi) \|\nabla_{\theta} \psi\|_2} (\chi \|\omega^L\|_2 \|\nabla_{\theta} \psi\|_2 + \|\Delta - \widehat{\Delta}\|_2) + \chi \|y\|_2 \|\nabla_{\theta} \psi\|_2 \quad (29) \\ &\leq \frac{\kappa \|\widehat{\nabla_{\theta} \psi} - \nabla_{\theta} \psi\|_2}{c \|\nabla_{\theta} \psi\|_2} \|\omega^L\|_2 + \frac{\kappa \|\Delta - \widehat{\Delta}\|_2}{c \|\nabla_{\theta} \psi\|_2} + \|\nabla_{\theta} \psi\|_2 \|y\|_2 \\ &= \|\nabla_{\theta} \psi - \widehat{\nabla_{\theta} \psi}\|_2 \left(\frac{\kappa \|\omega^L\|_2}{c \|\nabla_{\theta} \psi\|_2} + \|y\|_2 \right) + \|\Delta - \widehat{\Delta}\|_2 \frac{\kappa}{c \|\nabla_{\theta} \psi\|_2} \\ &= \|\nabla_{\theta} \psi - \widehat{\nabla_{\theta} \psi}\|_2 \left(\frac{\|\omega^L\|_2}{c \sigma_{\min}(\nabla_{\theta} \psi)} + \|y\|_2 \right) + \|\Delta - \widehat{\Delta}\|_2 \frac{1}{c \sigma_{\min}(\nabla_{\theta} \psi)} \\ &\quad + 2\sigma \sqrt{\frac{r + \log(\frac{1}{\delta})}{n\eta}} \frac{1}{c \sigma_{\min}(\nabla_{\theta} \psi)} \quad (30) \end{aligned}$$

$$\leq O\left(\frac{(M + M_S^2 M_R)}{\sigma_{\min}(\nabla_{\theta} \psi)} \sqrt{\frac{\log(\frac{2}{\delta})}{n\eta}}\right), \quad (31)$$

where in line (29) we apply Lemma A.1 and in line (30) we apply Equation (28) and Lemma A.8. \square

B Experiments

In this appendix, we report some experimental details together with some additional experiments.

B.1 Gridworld

In the Gridworld experiment, we select different learning steps for different learners. The number of learning steps depends on the number of policy updates that the learner takes to become an expert. In the following plots, we report the expected discounted return for each learner: Q-Learning (Figure 6), G(PO)MDP (Figure 5), SPI (Figure 7), SVI (Figure 8). In these plots, the expected discounted return is estimating using a batch of 50 trajectories for each learner. The discount factor used in all experiments is 0.96.




Figure 5: Learning performance of G(PO)MDP. 20 runs, 98% c.i.




Figure 6: Learning performance of Q-Learning. 20 runs, 98% c.i.




Figure 7: Learning performance of SPI. 20 runs, 98% c.i.




Figure 8: Learning performance of SVI. 20 runs, 98% c.i.

B.2 MuJoCo additional experiments

For the MuJoCo experiments, we use the same hyperparameters as in [17], apart from that we use 16 parallel agents for PPO, due to resource constraints. The number of forward steps are settled to 2000. As in [17], we select a subset of the learner’s trajectories and we do not use the first 10 trajectories because the first phase of learning is too noisy. We evaluate the algorithms on the first 1 million environment steps.

	Recovered Weights	Real Weights
Time	0.0401	0.0017
Jerk	0.0174	0.0003
Slow	0.0001	0.0000
Crash	0.9424	0.9980

Table 1: Reward weights for the autonomous simulate driving scenario.

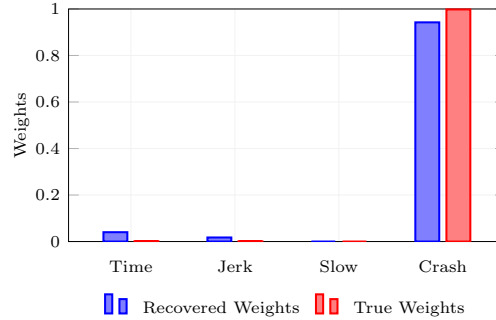


Figure 9: Reward weights for the autonomous simulate driving scenario.

B.3 Autonomous driving scenario

In this section, we report an additional, preliminary experiment, that we perform on a simulator driving scenario. We employ SUMO simulator, an open-source, continuous road traffic simulation package designed to handle large road networks. SUMO focuses on the high-level control of the car, integrating an internal system that controls the vehicle dynamics. During the simulation, SUMO provides information on the other vehicles around the ego vehicle.

We consider a crossroad scenario which consists of an intersection with an arbitrary number of roads. The vehicle coming from the source road has to reach a target road that has a higher priority. The goal of the agent is to drive the ego car and enter the target road, avoiding dangerous manoeuvres.

The reward features consists of four components: *Time*, a constant feature at each decision step; *Jerk*, the absolute value of the instantaneous jerk, i.e., the finite- difference derivative of the acceleration; *Harsh Slow Down*, a binary feature, which activates whenever the velocity is lower than a threshold; *Crash*, a binary feature which activates when the vehicle violates the safety constraints or performs a crash.

The agent’s policy is a rule-based policy, i.e., a set of parametrized rules, which is learned using Policy Gradients with Parameter-based Exploration (PGPE). It is important to notice that the agent’s policy is not differentiable.

We perform 10 PGPE updates of the agents and then we use the learning trajectories with LOGEL. In the behavioural cloning step, we use a linear layer to approximate the policy of the learner. Table 1 shows the normalized weights recovered by LOGEL and the normalized real weights. As shown in Table 1 and in Figure 9, the reward weights recovered are quite similar to the real ones.

We thank Amarildo Likmeta to provide us the data to perform this experiment.