

Inverse Reinforcement Learning by Estimating Expertise of Demonstrators

Mark Beliaev, Ramtin Pedarsani

University of California, Santa Barbara
mbeliaev@ucsb.edu, ramtin@ucsb.edu

Abstract

In Imitation Learning (IL), utilizing suboptimal and heterogeneous demonstrations presents a substantial challenge due to the varied nature of real-world data. However, standard IL algorithms consider these datasets as homogeneous, thereby inheriting the deficiencies of suboptimal demonstrators. Previous approaches to this issue rely on impractical assumptions like high-quality data subsets, confidence rankings, or explicit environmental knowledge. This paper introduces IRLEED, *Inverse Reinforcement Learning by Estimating Expertise of Demonstrators*, a novel framework that overcomes these hurdles without prior knowledge of demonstrator expertise. IRLEED enhances existing Inverse Reinforcement Learning (IRL) algorithms by combining a general model for demonstrator suboptimality to address reward bias and action variance, with a Maximum Entropy IRL framework to efficiently derive the optimal policy from diverse, suboptimal demonstrations. Experiments in both online and offline IL settings, with simulated and human-generated data, demonstrate IRLEED’s adaptability and effectiveness, making it a versatile solution for learning from suboptimal demonstrations.

Code — <https://github.com/mbeliaev1/IRLEED>

Extended version — <https://arxiv.org/abs/2402.01886>

1 Introduction

Reinforcement Learning (RL) has proven to be a powerful tool across a wide range of applications, from controlling robotic systems, to playing complex games like Go and Chess. The efficacy of RL algorithms is largely attributed to their ability to optimize hand-crafted reward functions (Brockman et al. 2016). However, designing these functions is often challenging and impractical in complex environments (Hadfield-Menell et al. 2017). A prevalent alternative to this dilemma is Imitation Learning (IL), which bypasses the need for cumbersome reward engineering by leveraging expert demonstrations to instill desired behaviors (Argall et al. 2009).

The two main approaches utilized in IL are: behavioral cloning (Pomerleau 1991), which acquires a policy through a supervised learning approach using state-action pairs provided by the expert, and Inverse Reinforcement Learning (IRL) (Ng and Russell 2000), which solves for the reward

function that makes the expert’s behavior optimal, subsequently facilitating the training of an IL policy. Despite the effectiveness of behavior cloning in simple environments with large amounts of data, learning a policy to fit single time step decisions leads to compounding errors due to covariance shift (Ross and Bagnell 2010). On the other hand, IRL learns a reward function that prioritizes entire trajectories over others, considering the sequential nature of the decision making problem at hand (Abbeel and Ng 2004). Consequently, the success of IRL has spurred the development of IL techniques that either explicitly or implicitly incorporate environment dynamics (Ho and Ermon 2016; Fu, Luo, and Levine 2018; Kostrikov, Nachum, and Tompson 2019; Garg et al. 2021).

However, a critical assumption in these methodologies is the availability of high-quality demonstration data. In many practical situations, especially with crowd-sourced or varied data, the quality of demonstrations is inconsistent (Mandlekar et al. 2021; Belkhale, Cui, and Sadigh 2023). For robotics, data curation becomes vital, since utilizing smaller real world datasets with noise or biases can lead to hazardous situations. Furthermore, assuming uniformity in demonstration quality overlooks the unique intentions of individual demonstrators, potentially leading to suboptimal learning outcomes (Eysenbach et al. 2018). Therefore, it is critical to develop IL methods that can account for the heterogeneity and suboptimality of demonstration data.

Related Work Addressing the challenges posed by learning from suboptimal and heterogeneous demonstrations is complex. One line of research focuses on heterogeneous demonstrations without considering their quality (Chen et al. 2020, 2023), assuming that all demonstrations are optimal. Other works consider suboptimal demonstrations, but require unrealistic prerequisites such as explicit knowledge of environment dynamics (Brown, Goo, and Niekum 2020; Chen, Paleja, and Gombolay 2021), a set of confidence scores over the demonstrations (Zhang et al. 2021; Cao and Sadigh 2021; Wu et al. 2019; Kuhar et al. 2023), or direct access to a subset of expert demonstrations (Shiarlis, Messias, and Whiteson 2016; Xu et al. 2022; Yang, Levine, and Nachum 2021).

A recent approach that avoids these assumptions is ILEED (Beliaev et al. 2022), which directly models the expertise of demonstrators within a supervised learning framework, thus enabling one to learn from the optimal behavior while filtering out the suboptimal behavior of each demonstrator.

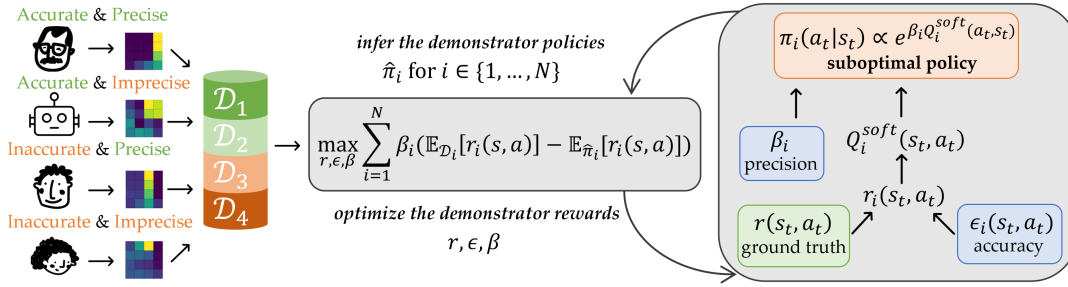


Figure 1: IRLEED is applied in the suboptimal setting to estimate the ground truth reward r , which is used to find the optimal policy. *Left*: A heterogeneous dataset is collected from multiple sources with varying optimality. We categorize this optimality by using accuracy to represent the reward bias, and precision to represent the variance in action choices. *Right*: We infer the demonstrator policies using a model for behavior based on the Boltzmann rationality principle, which captures both the accuracy ϵ_i , and the precision β_i , as compared to the ground truth reward r . *Middle*: Using estimates of the demonstrator policies $\hat{\pi}_i$ along with the demonstrations \mathcal{D}_i , we can optimize for the true reward r , and the parameters that capture accuracy ϵ and precision β .

While this method was proven effective, it is based on a behavioral cloning formulation, which overlooks the underlying environment dynamics, and its representation of suboptimal behavior is limited, presuming demonstrators to be noisy approximations of experts.

Overview To overcome these challenges, this paper introduces IRLEED, *Inverse Reinforcement Learning by Estimating Expertise of Demonstrators*. As illustrated in Figure 1, IRLEED is a novel framework for IRL that accounts for demonstrator suboptimality without prior expertise knowledge. It comprises two core components: (1) a general model of demonstrator suboptimality based on the Boltzmann rationality principle, which captures both the reward bias and the variance in action choices of each demonstrator as compared to the optimal policy (2) a Maximum Entropy IRL framework, which is used to efficiently find the optimal policy when provided with a set of suboptimal and heterogeneous demonstrations. These elements enable IRLEED to effectively recover the ground truth policy from suboptimal and heterogeneous demonstrations, surpassing the limitations of previous models. Furthermore, IRLEED’s simplicity and compatibility with existing IRL techniques make it a versatile and powerful tool in the field of IL.

The main contributions of our work are:

- We propose a novel framework designed to enhance existing IRL algorithms, addressing the challenges of learning from suboptimal and heterogeneous data.
- We provide comparative insights against standard IRL and the behavior cloning approach, ILEED, showing how IRLEED generalizes both methods.
- We empirically validate the success of our method against relevant baselines in both online and offline IL scenarios.

2 Preliminary

Notations We use the Markov decision process (MDP) setting, defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, p_0, r, \gamma)$, where \mathcal{S} , \mathcal{A} represent state and action spaces, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ represents the dynamics, $p_0 : \mathcal{S} \rightarrow [0, 1]$ represents the initial state distribution, $r \in \mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ represents the reward function, and $\gamma \in (0, 1)$ represents the discount factor.

In Sections 3 and 4, we describe our method using finite state and action spaces, \mathcal{S} and \mathcal{A} , but our experiments later utilize continuous environments. We use $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ to denote a policy that assigns probabilities to actions in \mathcal{A} given states in \mathcal{S} . Considering the γ -discounted infinite time horizon setting, we define the expected value under a policy π in terms of the trajectory it produces. Specifically, this expectation is represented as $\bar{f}_\pi = \mathbb{E}_\pi[f(s, a)]$, which equals $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t f(s_t, a_t)]$, where $s_0 \sim p_0, a_t \sim \pi(\cdot | s_t)$, and $s_{t+1} \sim \mathcal{T}(\cdot | s_t, a_t)$. Additionally, we will use $\bar{f}_{\mathcal{D}} = \mathbb{E}_{\mathcal{D}}[f(s, a)]$ to denote the empirical expectation with respect to the trajectories $\tau_j = (s_0, a_0, \dots, s_T, a_T)$ in demonstration set $\mathcal{D} = \{\tau_j\}_{j=1}^M$, where we leave out denoting trajectories of varying lengths for simplicity.

Maximum Entropy IRL The goal of IRL is to recover a reward $r \in \mathcal{R}$ that rationalizes the demonstrated behavior in dataset \mathcal{D} , where the provided dataset consists of trajectories τ sampled from an expert policy π_E . This idea was first formulated as a feature-expectation matching problem (Abbeel and Ng 2004): assuming that the feature vector $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^k$ components quantify the expert’s behavior, to find a policy π that performs equal to or better than the expert π_E , it suffices that their feature expectations match $\bar{f}_\pi = \bar{f}_{\pi_E}$. Unfortunately, this leads to an ill-posed problem as many policies can lead to the same feature counts. Maximum Entropy IRL resolves this ambiguity by choosing the policy which does not exhibit any additional preferences beyond matching feature expectations (Ziebart et al. 2008):

$$\arg \max_{\pi} H(\pi), \text{ such that: } \bar{f}_\pi = \bar{f}_{\mathcal{D}}, \quad (1)$$

where in the infinite-time horizon setting, this equates to maximizing the discounted causal entropy $H(\pi)$ under the feature matching constraint (Bloem and Bambos 2014).

In the MDP setting, the above problem is equivalent to finding the maximum likelihood estimate of θ :

$$\hat{\theta} = \arg \max_{\theta} \log \prod_{\tau_j \in \mathcal{D}} p_0(s_0) \prod_{t=0}^T \pi_{\theta}(a_t | s_t) \mathcal{T}(s_{t+1} | s_t, a_t), \quad (2)$$

with the parameterized policy π_θ defined recursively using Q_θ^{soft} and V_θ^{soft} :

$$\pi_\theta(a_t|s_t) = \exp(Q_\theta^{soft}(s_t, a_t) - V_\theta^{soft}(s_t)), \quad (3)$$

$$Q_\theta^{soft} = \theta^\top f(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \mathcal{T}(\cdot|s_t, a_t)}[V_\theta^{soft}(s_{t+1})], \quad (4)$$

$$V_\theta^{soft} = \log \sum_{a_t \in \mathcal{A}} \exp(Q_\theta^{soft}(s_t, a_t)). \quad (5)$$

The parameters $\theta \in \mathbb{R}^k$ specifying the soft Bellman policy, $\pi_\theta(a_t|s_t)$, correspond to the dual variables of the feature matching constraint in Eq. (1). Without loss of generality, if we define the true reward signal as $r(s, a) = \theta^{*\top} f(s, a)$ for some reward parameter $\theta^* \in \mathbb{R}^k$ and features f , then the soft Bellman policy parameterized by θ^* achieves the maximum possible return (for more details see (Ziebart 2010), Theorem 6.2 and Corollary 6.11).¹ This formulates an iterative approach to finding the true reward parameter θ^* : given demonstration set \mathcal{D} and your estimate θ , alternate between *inferring* the policy $\pi_\theta(a_t|s_t)$ based on Eq. (3), and *optimizing* the parameter θ based on Eq. (2).

This line of work provides many insights that have been used to derive modern IRL (Fu, Luo, and Levine 2018; Reddy, Dragan, and Levine 2019; Gao et al. 2018; Garg et al. 2021) and RL (Hao et al. 2017, 2018) algorithms. Unlike these prior works, we consider the setting where we no longer have access to the expert policy π_E , and instead are given a heterogeneous set of potentially suboptimal demonstrations from which we want to learn a well performing policy. To this end, we first describe the model we use to capture the suboptimality of demonstrators in Section 3, following which we frame this problem within the IRL framework and derive our own method IRLEED in Section 4. Finally, we describe our experimental results in Section 5, and conclude our work in Section 6.

3 Suboptimal Demonstrator Model

This work considers the case where dataset $\mathcal{D} = \{(i, \mathcal{D}_i)\}_{i=1}^N$ consists of a mixture of N demonstrations with varying quality. Each demonstration $\mathcal{D}_i = \{\tau_j\}_{j=1}^{M_i}$ contains a set of M_i trajectories sampled from a distinct, fixed policy π_i . Rather than considering \mathcal{D} as a homogeneous set derived from one expert policy π_E , our approach focuses on leveraging the specific source of each demonstration to enhance the IRL framework. By identifying the varying qualities within this mixed collection of demonstrations, we aim to more accurately deduce the ground truth policy π_E , as opposed to a simplistic method that averages behaviors. To achieve this, it is essential to develop a model for the demonstrators π_i that accurately reflects their suboptimal behavior relative to the optimal policy π_E .

A popular framework for modeling decision making is the Boltzmann rationality model, which is widely used in psychology (Baker, Saxe, and Tenenbaum 2009), economics (Luce 1959) and RL (Laidlaw and Dragan 2021). In

¹To allow an arbitrary reward specification, the IRL objective defined in Eq. 1 can be expressed in terms of occupancy measures instead of feature expectations (Ho and Ermon 2016).

this model, the likelihood of an agent choosing an action is proportional to the exponential of the reward associated with that action: $\pi(a|s; r, \beta) \propto \exp(\beta r(s, a))$, where the inverse temperature parameter β controls the randomness of the choice. While this formulation provides an intuitive way to capture one form of suboptimality through β , it traditionally assumes that agents have unbiased knowledge about the true reward $r(s, a)$. To broaden this model for suboptimal demonstrations, we propose considering agents that follow a biased reward $r(s, a) + \epsilon(s, a)$, introducing ϵ as the deviation from the true reward. This modification enables us to represent an agent’s *accuracy* through the deviation ϵ , adding a bias to their perceived rewards, and their *precision* through parameter β , adding variance to their action choices.

Within the context of the IRL framework described in Section 2, we recall that the expert policy π_E aligns with the true reward, formulated as $r(s, a) = \theta^{*\top} f(s, a)$. Applying the aforementioned concept, we can model demonstrator i ’s perceived reward as a deviation from the true reward: $r_i(s, a) = (\theta^* + \epsilon_i)^\top f(s, a)$, where $\epsilon_i \in \mathbb{R}^k$. By incorporating this altered reward r_i into the *soft* Bellman policy defined in Eq. (3), we derive a recursive parameterization for demonstrator i ’s policy:

$$\pi_{\theta^*, \epsilon_i, \beta_i}(a_t|s_t) \propto \exp(\beta_i(Q_{\theta^* + \epsilon_i}^{soft}(s_t, a_t))), \quad (6)$$

where $\beta_i \in [0, \infty)$, and the normalizing factor not shown above should be modified to account for β_i : $V_{\theta^* + \epsilon_i}^{soft}(s_t) = \frac{1}{\beta_i} \log \sum_{a_t \in \mathcal{A}} \exp(\beta_i Q_{\theta^* + \epsilon_i}^{soft}(s_t, a_t))$.

Our model captures two distinct aspects of suboptimal behavior: (1) ϵ_i quantifies the demonstrator’s *accuracy* in estimating the true reward r , where $\epsilon_i^\top f(\tau)$ represents the estimation error. (2) β_i quantifies the demonstrator’s *precision* in action selection, where $\beta_i \rightarrow \infty$ interpolates from the soft Bellman policy, which samples actions according to their Q-values, to the standard Bellman policy, which chooses actions that maximize the Q-value. This model offers a versatile and mathematically convenient way to depict suboptimal demonstrator behaviors within the soft Bellman policy domain, as well as extend the results of Maximum Causal Entropy IRL to the suboptimal setting.

Remark 3.1. *It is important to note that this model extends outside the scope of problems where the true reward r is an affine transformation of some feature vector f . In the general setting, we can parameterize demonstrator i ’s reward as $r_i = r + \epsilon_i$, a combination of the true reward $r \in \mathcal{R}$ and their deviation $\epsilon_i \in \mathcal{R}$, where r and ϵ_i can be represented by neural networks to allow for general function approximation. We opt to formulate our proposed model within the classic feature matching scope, as this facilitates direct comparison with the seminal works used to derive our results in the following section. However, our proposed model’s applicability extends beyond this traditional framework, as we elaborate in Section 4, and demonstrate through experiments in Section 5. Until then, we utilize the feature matching framework described, referencing it as standard IRL.*

4 Learning from Suboptimal Data with IRL

Up to this point, we have formulated a model for each demonstrator i based on the parameter θ^* , which defines the true reward, and parameters ϵ_i and β_i , which characterize demonstrator suboptimality. However, we do not have access to these parameters, and are instead provided with a dataset of demonstrations \mathcal{D} . Recall that standard IRL assumes that all demonstrations come from the optimal policy π_E , and follows the iterative approach described in Section 2 to find it. However, applying this approach to the heterogeneous regime will yield an averaged policy, leading to subpar performance if a significant portion of the demonstrations are suboptimal. To refine this, we introduce IRLEED, an innovative extension of the IRL framework tailored for the suboptimal setting.

IRLEED

Utilizing the dataset $\mathcal{D} = \{(i, \mathcal{D}_i)\}_{i=1}^N$, we can jointly estimate the maximum likelihood parameters:

$$\mathcal{L}(\theta, \epsilon, \beta) = \log \prod_{\mathcal{D}_i \in \mathcal{D}} \prod_{\tau_j \in \mathcal{D}_i} P(\tau_j | \theta, \epsilon_i, \beta_i), \quad (7)$$

where $\epsilon = \{\epsilon_i\}_{i=1}^N$, $\beta = \{\beta_i\}_{i=1}^N$, and the probability $P(\tau_j | \theta, \epsilon_i, \beta_i)$ is given by: $P(\tau | \theta, \epsilon_i, \beta_i) = p_0(s_0) \prod_{t=0}^T \pi_{\theta, \epsilon_i, \beta_i}(a_t | s_t) \mathcal{T}(s_{t+1} | s_t, a_t)$.

Under the specified demonstrator model, this likelihood serves as a suitable loss function. Specifically, rewriting the joint likelihood equivalently only over θ : $\mathcal{L}(\theta) = \max_{\epsilon, \beta} \mathcal{L}(\theta, \epsilon, \beta)$, it is easy to show that $\mathcal{L}(\theta)$ is a proper objective function.

Proposition 4.1. θ^* is the (non-unique) maximizer of $\mathcal{L}(\theta)$.

This provides us with a convenient way to derive the gradient of likelihood function $\mathcal{L}(\theta, \epsilon, \beta)$.

Lemma 4.2. Given the likelihood \mathcal{L} defined according to Eq. (7) and demonstrator policies $\pi_{\theta, \epsilon_i, \beta_i}$ parameterized by Eq. (6), we can compute the gradients of \mathcal{L} as follows:

$$\nabla_{\theta} \mathcal{L} = \sum_{i=1}^N \beta_i (\tilde{f}_{\mathcal{D}_i} - \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}), \quad (8)$$

$$\nabla_{\epsilon_i} \mathcal{L} = \beta_i (\tilde{f}_{\mathcal{D}_i} - \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}), \quad (9)$$

$$\frac{\partial \mathcal{L}}{\partial \beta_i} = (\theta + \epsilon_i)^{\top} (\tilde{f}_{\mathcal{D}_i} - \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}). \quad (10)$$

Proof. First, the parameter β_i can be folded into the modified soft Bellman policy of Eq. (6) as the reward estimate: $\beta_i(\theta + \epsilon_i)^{\top} f(s_t, a_t)$. Second, we can express the likelihood as $\mathcal{L} = \sum_{i=1}^N \mathcal{L}_i$, where $\mathcal{L}_i = \log \prod_{\tau_j \in \mathcal{D}_i} P(\tau_j | \theta, \epsilon_i, \beta_i)$. Setting $u_i = \beta_i(\theta + \epsilon_i)$, we know that $\nabla_{u_i} \mathcal{L}_i = \tilde{f}_{\mathcal{D}_i} - \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}} \forall i \in \{1, \dots, N\}$ (see (Ziebart 2010), Theorem 6.2, Lemma A.2). Using the chain rule completes the proof. \square

This result has an intuitive interpretation. Recall that for each demonstrator i , $\tilde{f}_{\mathcal{D}_i} - \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}$ represents the difference between the empirical feature vector given samples from their policy, $\tilde{f}_{\mathcal{D}_i}$, and the expected feature vector under the probabilistic model, $\bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}$. With this in mind, Eq. (9) shows

that we update ϵ_i to match feature expectation with demonstrator π_i , just like standard IRL. On the other hand, Eq. (8) shows that we update θ to match a weighted average over feature expectations provided by all demonstrators, where their respective precision β_i determines their contribution. Lastly, Eq. (10) shows that we update this precision β_i to balance the expected returns of the demonstrator, π_i , and our probabilistic model, $\pi_{\theta, \epsilon_i, \beta_i}$, under our estimate of the perceived reward, $r_i(s, a) = (\theta + \epsilon_i)^{\top} f(s, a)$. This way, when the probabilistic model outperforms the demonstration set under the estimated reward r_i , we decrease the precision β_i to lower its relative performance, and vice versa.

This gives us a general algorithm that extends the IRL framework. Specifically, IRLEED follows an iterative approach to finding the true reward parameter θ^* : given demonstration set $\mathcal{D} = \{(i, \mathcal{D}_i)\}_{i=1}^N$ and estimates θ, ϵ, β , for each demonstrator $i \in \{1, \dots, N\}$, alternate between *inferring* the policy $\pi_{\theta, \epsilon_i, \beta_i}(a_t | s_t)$ based on Eq. (6), and *optimizing* the parameters $\theta, \epsilon_i, \beta_i$ based on Eqs. (8), (9), (10), respectively.

As aforementioned, many techniques can be used for both the inference and optimization procedures. To infer the policy $\pi_{\theta, \epsilon_i, \beta_i}(a_t | s_t)$ given $\theta, \epsilon_i, \beta_i$, one can use soft versions of the standard value iteration and Q-learning algorithms, depending on whether the dynamics \mathcal{T} are provided or not (Bloem and Bambos 2014). To compute feature expectations $\bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}$ based on the inferred policy, one utilize the standard dynamic programming operator when the dynamics \mathcal{T} are known, or estimate the expectations by utilizing Monte Carlo simulations of the MDP when the dynamics \mathcal{T} are unknown.

Comparisons to Standard IRL and ILEED

This likelihood maximization problem is closely related to the maximum causal entropy IRL framework described in Eqs. (1)- (5). Specifically, IRLEED makes two modifications: (1) For each $i \in \{1, \dots, N\}$, it defines individual dual variables $\theta_i = \theta + \epsilon_i$ that correspond to feature matching constraint $\tilde{f}_{\mathcal{D}_i} = \bar{f}_{\pi_{\theta, \epsilon_i, \beta_i}}$, tying the demonstrator policies together. (2) It uses a learnable parameter β_i to tune the magnitude of uncertainty related to demonstrations provided in \mathcal{D}_i . Moreover, IRLEED generalizes standard IRL.

Remark 4.3. IRLEED recovers the IRL framework when we set $\epsilon_i = [0]^k$ and $\beta_i = 1$ as constants.

To show the advantage of IRLEED within the suboptimal demonstration setting, we note that the policy recovered by standard IRL can only perform as well as the average demonstration \mathcal{D}_i provided. This is a direct result of the feature matching constraint imposed by the IRL objective.

Proposition 4.4. Given $\mathcal{D} = \{(i, \mathcal{D}_i)\}_{i=1}^N$ produced by sub-optimal policies π_i according to Eq. (6), let θ_{IRL} denote the naive solution to the IRL problem defined by Eq. 2, which treats demonstration set \mathcal{D} as one homogeneous set produced by π_E . The performance of IRL is bounded by:

$$\mathbb{E}_{\pi_{\theta_{IRL}}} [r(s, a)] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\pi_i} [r(s, a)] \leq \mathbb{E}_{\pi_{\theta^*}} [r(s, a)], \quad (11)$$

where we assume that all demonstrators provide the same number of trajectories $M_i = M$.

As mentioned in Proposition 4.1, IRLEED can recover the true reward parameter θ^* , which outperforms θ_{IRL} unless the demonstrations provided are optimal. Unfortunately, since θ^* is not the unique solution to IRLEED, we can not make guarantees on recovering the true reward parameter θ^* . Nonetheless, our experimental results in Section 5 demonstrate that IRLEED improves performance over IRL frameworks in both the feature matching setting described here, as well as the general IRL setting, which can be applied to continuous control tasks with high dimensional input spaces.

Finally, we note that when the MDP is deterministic, we can express the policy without recursion: $\pi_{\theta, \epsilon_i, \beta_i}(\tau) \propto \exp(\beta_i(\theta + \epsilon_i)^\top f(\tau))$, where $f(\tau) = \sum_{(s,a) \in \tau} f(s,a)$, and $\pi(\tau)$ can directly map from trajectories due to the deterministic dynamics \mathcal{T} . Replacing the probability $P(\tau|\theta, \epsilon_i, \beta_i)$ with $\pi_{\theta, \epsilon_i, \beta_i}(\tau)$, maximizing the likelihood defined by Eq. (7) corresponds to the supervised learning approach employed by ILEED (Beliaev et al. 2022). Hence unlike ILEED, our framework provides a dynamics-aware solution to the problem of learning from mixtures of suboptimal demonstrations.

Remark 4.5. *IRLEED recovers the ILEED framework if we assume all demonstrators i have knowledge of the true reward, $\epsilon_i = [0]^k$, and ignore the dynamics of the MDP.*

Practical Algorithm

While the results in the previous sections provide insight on how to learn from suboptimal demonstrations using IRL, there are two main concerns: (1) Relying on hand-crafted features f can hinder our ability to express complex behaviors. (2) Nothing in the loss prevents the model from overfitting by learning zero shared reward θ , and a completely different reward deviation ϵ_i , for each demonstrator i . To address these concerns, we can extend our approach to generalized IRL:

$$\max_{r \in \mathcal{R}} \min_{\pi \in \Pi} \mathbb{E}_{\pi_E} [r(s,a)] - \mathbb{E}_{\pi} [r(s,a)] - H(\pi) - \psi(r), \quad (12)$$

where the reward r belongs to a non restrictive set of functions $\mathcal{R} = \{r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}\}$, and ψ is a convex regularizer that is used to prevent overfitting (Ho and Ermon 2016).

In brief, we extend Eq. (12) to the suboptimal setting by using neural networks to parameterize r and $\epsilon = \{\epsilon_i\}_{i=1}^N$, and implement ℓ_2 -regularization on the outputs of ϵ to mitigate overfitting. Our experiment results in Section 5 demonstrate that IRLEED performs well in this generalized setting, and can even be trained offline when paired with Inverse soft-Q learning (IQ), an approach that avoids the iterative process defined by learning a single Q-function, implicitly representing both reward and policy (Garg et al. 2021). We leave the details of this section to our extended version, covering how IRLEED can be utilized alongside generalized IRL algorithms to address both of the aforementioned concerns.

5 Experiments

In this section we evaluate how IRLEED performs when learning from suboptimal demonstrations, using experiments in both online and offline IL settings, with simulated and human-generated data. Throughout this section we compared IRLEED to maximum entropy IRL (Ziebart, Bagnell, and

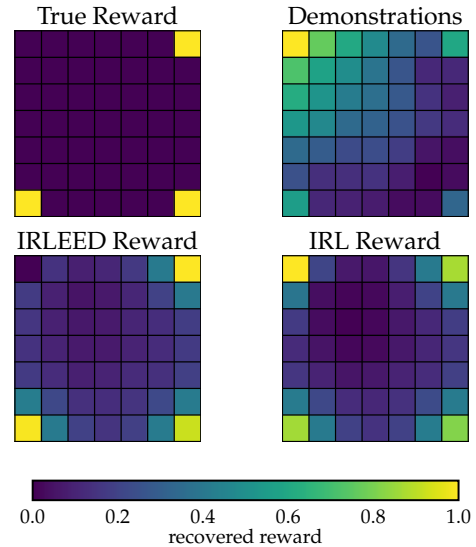


Figure 2: We visualize the reward recovered by IRLEED and IRL when trained using suboptimal demonstrations. Top left shows the true reward, where the three yellow corners are terminal states. Top right shows the normalized state visitation frequency over the entire dataset. Bottom left and right show the normalized rewards recovered by IRLEED and IRL respectively. We can see that the provided demonstrations are misaligned with the ground truth reward: the state visitation frequency for the top left corner is higher due to demonstrator suboptimalities. As expected, the feature matching constraint of IRL absorbs this suboptimality. Although the reward recovered by IRL contains information about the true reward, it is incorrectly biasing the top left corner. On the other hand, IRLEED is able to remove this bias, providing a better estimate of the ground truth reward.

Dey 2010), Inverse soft-Q learning (IQ) (Garg et al. 2021), and ILEED (Beliaev et al. 2022). In total, we performed four sets of experiments: (1) Using a custom Gridworld environment, we simulated suboptimal demonstrations with varying levels of *precision* and *accuracy*, and compared the performance and recovered rewards of IRLEED and maximum entropy IRL. (2) Using the control tasks, Cartpole, and Lunar Lander, we simulated suboptimal demonstrations with varying levels of *expertise* and *accuracy*, and compared the performance of IRLEED, IQ, and ILEED in both the online and offline settings. (3) Using the Atari environments, Space Invaders, and Qbert, we combined both simulated and human demonstrations, and compared the performance of IRLEED, IQ, and ILEED in the offline setting. (4) Using the continuous control Mujoco task Hopper, we collected suboptimal demonstrations from pretrained policies, and compared the performance of IRLEED and IQ.

Implementation We utilized the codebases provided by the authors to implement ILEED and IQ, creating IRLEED on top of the IQ algorithm by adding a learnable parameter β_i , and an additional critic network ϵ_i , for each demonstrator i .

Setting		Cartpole						Lunar Lander					
Prec.	Acc.	Mean	IQ^*	$IRLEED^*$	ILEED	IQ	IRLEED	Mean	IQ^*	$IRLEED^*$	ILEED	IQ	IRLEED
<i>L</i>	<i>L</i>	0.58	0.97	1.55	1.38	0.96	1.56	0.46	0.68	1.28	0.84	0.81	1.34
<i>L</i>	<i>H</i>	0.70	0.92	1.00	1.00	0.96	1.00	0.54	0.84	1.14	1.06	1.05	1.18
<i>H</i>	<i>L</i>	0.64	1.03	1.09	1.05	0.95	1.05	0.57	0.94	1.30	0.85	0.99	1.14
<i>H</i>	<i>H</i>	1.00	0.99	1.00	0.92	1.00	1.00	0.86	0.93	0.98	1.01	0.99	0.98

Table 1: Simulated Control Tasks. We list the average return of the recovered policies (IQ, IRLEED, ILEED), relative to the best demonstrator’s return, where IQ^* and $IRLEED^*$ are used to denote *online* algorithms. This is provided for 4 dataset settings (Prec., Acc.), where we list the mean performance (Mean) of all demonstrators relative to the best. The last row (H,H) corresponds to the clean dataset. The return of the best demonstrator for each setting (top to bottom) is 260, 500, 411, 500 (Cartpole) and 171, 213, 166, 277 (Lunar Lander).

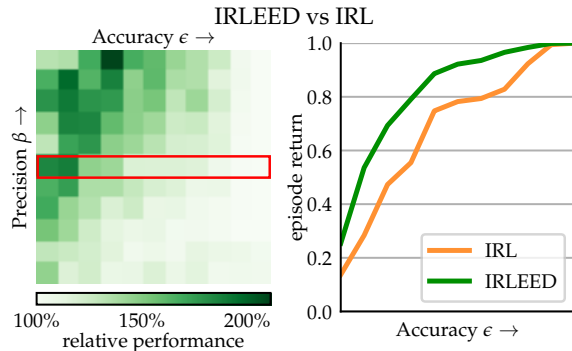


Figure 3: We compare the performance of the policies recovered by IRLEED and IRL. The left plot shows the relative performance of IRLEED over IRL under varying dataset settings, where the top right corner corresponds to expert data. The right plot shows the performance of both policies as we increase the accuracy ϵ of demonstrators in the dataset, corresponding to the data outlined in red in the left plot. On average, IRLEED provided a 30.3% improvement over IRL.

Since IQ learns a reward and policy directly by using a single Q function as a critic, we used our additional critic network ϵ_i to directly add bias to the state action function instead of the reward. For further implementation details refer to the Appendix found in our extended version.

1. Gridworld For the Gridworld experiments, we implemented IRL and IRLEED by utilizing stochastic value iteration on the reward estimate θ to infer the policy π , and Monte Carlo simulations of the MDP to estimate the feature expectations required to optimize θ , repeating both steps until convergence. For IRLEED, we updated ϵ and β only after θ converged, and repeated this for two iterations.

The reward in this setting was linearly parameterized as $r(s) = \theta^\top f(s)$, where $f(s)$ is a one hot encoding for each gridworld state. For each level of *precision* and *accuracy*, we generated 5 suboptimal policies according to Equation (6), where β_i was sampled from a uniform distribution starting at 0, with mean equal to the *precision* level, and ϵ_i was sampled from a multivariate normal distribution $\mathcal{N}([0]^k, \mathbf{I}_k/\lambda^2)$, with λ equal to the *accuracy* level. We tested 121 dataset settings, collecting 40 trajectories from each policy, and us-

ing 100 seeds for each setting. Note that we visualized the recovered reward $r(s, a)$ for both IRLEED and IRL under one dataset setting in Figure 2. Given demonstrations that are misaligned with the ground truth reward, we saw that the feature matching constraint of IRL absorbed this suboptimality, while IRLEED was able to remove this bias, providing us with a better estimate of the ground truth reward.

To see the effect of dataset quality on performance, we use Figure 3 to show the relative improvement of the policy recovered by IRLEED over maximum entropy IRL for each dataset setting. We can see from the left plot that IRLEED provides improvement over IRL as we decrease *accuracy* ϵ (by decreasing λ), especially when *precision* β is high. This signifies the importance of modeling the accuracy ϵ , or reward bias, in addition to the precision β , or action variance, when learning from suboptimal demonstrations. This point is further highlighted in the right plot, which shows that while both methods can achieve good performance under low precision settings, IRLEED can tolerate lower demonstrator accuracy compared to IRL.

2. Simulated Control Tasks For this experiment, each level of *precision* and *accuracy* corresponded to 3 suboptimal policies generated according to $\pi_i(a|s) \propto \exp(\beta_i(Q(s, a) + \alpha\epsilon_i(s, a)))$, where $Q(s, a)$ is a pretrained critic network, $\epsilon_i(s, a)$ is a randomly initialized critic network, and α is a scaling factor which is set to zero for the high *accuracy* setting. For the low *precision* setting, β_i was sampled from a uniform distribution, whereas for the high *precision* setting, the stochasticity of the policy was removed by following the maximum state action value. We utilized both the online and offline setting, using 30 seeds for each dataset setting, and displaying our results in Table 1. Note that with near binary performance of the recovered policies, the variance is due to demonstration quality, hence we report averages relative to the best demonstrator computed over all initializations.

We can see that in the online setting, IRLEED outperforms its ablated counterpart, IQ, under all 3 suboptimal dataset settings (not including the clean dataset), for both environments. Furthermore, we see that in the more complex Lunar Lander environment, IQ can not perform as well as the best demonstrator, whereas IRLEED consistently performs on par or better than the best demonstrator. To have a fair comparison between IRLEED and its behavior cloning counterpart, ILEED, we repeat these experiments without online access

to environment interactions. We can see that offline IRLEED outperforms ILEED, under all 3 suboptimal dataset settings, for both environments. Furthermore, although ILEED has comparable performance in the simpler Cartpole environment, it performs poorly in the more complex Lunar Lander environment. Moreover, we can see that ILEED shows diminished performance when demonstrators have low accuracy ϵ , signifying that it can not account for the reward biases present in these settings, as was highlighted earlier in Remark 4.5.

3. Atari with Human Data Our demonstrations comprised two unique sources, dataset *A*: simulated data using a pre-trained expert policy and dataset *B*: collected data using adept human players (Kurin et al. 2017). Since the human dataset *B* was acquired using screenshots of web-browsers, there is a mismatch in sampling frequencies and colors compared to the simulated environment used in dataset *A*. This provides a challenge akin to crowd-sourcing data from varying sources, where we test if IRLEED can improve IQ when training on the combined dataset *AB* as opposed to dataset *A* alone.

The results shown in Table 2 list the return of the recovered policies for both environments (Space Invaders and Qbert), averaged over 5 seeds. While using dataset *A* is enough to reach decent performance, all methods do poorly when trained exclusively on dataset *B* due to the mismatch between the simulation environment and the data. However, we can see that IRLEED significantly improves its performance when utilizing the combined dataset *AB*. Overall we see IRLEED shows a 17% improvement over IQ, and a 41% improvement over ILEED in the combined *AB* setting.

4. Mujoco For our results on the Mujoco task Hopper, we collect a single trajectory from three sources, dataset *A*: a pretrained expert policy (return of 3457), dataset *B*: a partially trained policy (return of 2121) and a barely trained policy *C* (return of 216). Due to the complex nature of this continuous control tasks, the partially trained policies can not provide adequate demonstrations for learning. Hence our goal was to see if IRLEED could improve the capability of IQ when training on the combined dataset *ABC* by filtering the suboptimal demonstrations.

The results shown in Figure 4 plot the mean episode return of the policies during training, averaged over 15 initializations. As expected, both methods do poorly when training on dataset *C*, while achieving improved performance when training on dataset *B*. However, we see IRLEED outperforms

Dataset	Space Invaders			Qbert		
	ILEED	IQ	IRLEED	ILEED	IQ	IRLEED
<i>AB</i>	802	755	910	5511	8182	9336
<i>A</i>	783	754	768	5926	8092	8049
<i>B</i>	110	179	123	0	809	1087

Table 2: Utilizing Demonstrations from Varying Sources. We list the average return of the recovered policies, using the simulated dataset *A*, human dataset *B* (Kurin et al. 2017), and their combination *AB*. We note that IRLEED maintains similar per-seed variance as IQ on *AB* - Space Invaders: 61 vs 54 (IQ), Qbert: 315 vs 376 (IQ).

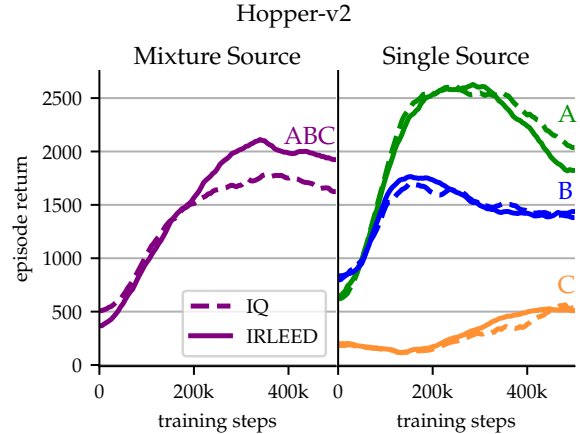


Figure 4: We plot the mean episode return of the policies learned by IRLEED (solid) and IQ (dashed) as they train on data from a single source (right) vs a mixture source (left).

IQ when training on the combined dataset *ABC* since it is more robust to the suboptimal demonstrations present in the data. However, we note that there is still a gap between the desired performance we see when training on dataset *A* alone, which demonstrates the difficulty of the task at hand.

6 Conclusion

Summary This paper addresses the challenges of leveraging suboptimal and heterogeneous demonstrations in IL by introducing IRLEED, a novel IRL framework. By integrating a general model of demonstrator suboptimality that accounts for reward bias and action variance, with a maximum entropy IRL framework, IRLEED mitigates the limitations inherent in standard IL algorithms, which often fail to account for the diversity and imperfections in real-world data. Our key insight is that: *supervised learning frameworks should account for the heterogeneity in crowd-sourced datasets by leveraging information about the source identity during training.*

Limitations Although we have shown that IRLEED is a general framework that improves over IRL in a multitude of settings, design choices such as the demonstrator model and reward regularization can be further analyzed. While we do try to make our framework easy to interpret, there is inherent complexity added as we need to simultaneously learn multiple demonstrator models. We also note that apart from the Gridworld experiments, we did not analyze reward recovery as it is hard to interpret in complex environments. However, our work addresses a problem setting where achieving desirable performance is a challenge in itself, and many recovered policies are below optimal. Lastly, we have not provided guarantees that state whether IRLEED will recover a better reward estimate compared to IRL. While we believe this is a good direction for future work, our goal was to present a novel formulation easily integrated with existing IRL methods, where further theoretical analysis is only possible under additional assumptions, and not crucial given the nonconvexity of the general setting.

Acknowledgements

This research was supported by NSF ECCS grant #2419982.

References

- Abbeel, P.; and Ng, A. Y. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, 1.
- Argall, B. D.; Chernova, S.; Veloso, M.; and Browning, B. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5): 469–483.
- Baker, C. L.; Saxe, R.; and Tenenbaum, J. B. 2009. Action understanding as inverse planning. *Cognition*, 113(3): 329–349.
- Beliaev, M.; Shih, A.; Ermon, S.; Sadigh, D.; and Pedarsani, R. 2022. Imitation learning by estimating expertise of demonstrators. In *International Conference on Machine Learning*, 1732–1748. PMLR.
- Belkhal, S.; Cui, Y.; and Sadigh, D. 2023. Data Quality in Imitation Learning. *arXiv preprint arXiv:2306.02437*.
- Bloem, M.; and Bambos, N. 2014. Infinite time horizon maximum causal entropy inverse reinforcement learning. In *53rd IEEE Conference on Decision and Control*, 4911–4916.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *arXiv:1606.01540*.
- Brown, D. S.; Goo, W.; and Niekum, S. 2020. Better-than-demonstrator imitation learning via automatically-ranked demonstrations. In *Conference on robot learning*, 330–359. PMLR.
- Cao, Z.; and Sadigh, D. 2021. Learning from imperfect demonstrations from agents with varying dynamics. *IEEE Robotics and Automation Letters*, 6(3): 5231–5238.
- Chen, L.; Jayanthi, S.; Paleja, R. R.; Martin, D.; Zakharov, V.; and Gombolay, M. 2023. Fast lifelong adaptive inverse reinforcement learning from demonstrations. In *Conference on Robot Learning*, 2083–2094. PMLR.
- Chen, L.; Paleja, R.; Ghuy, M.; and Gombolay, M. 2020. Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*, 659–668.
- Chen, L.; Paleja, R.; and Gombolay, M. 2021. Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*, 1262–1277. PMLR.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is All You Need: Learning Skills without a Reward Function. In *International Conference on Learning Representations*.
- Fu, J.; Luo, K.; and Levine, S. 2018. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.
- Gao, Y.; Xu, H.; Lin, J.; Yu, F.; Levine, S.; and Darrell, T. 2018. Reinforcement learning from imperfect demonstrations. *arXiv preprint arXiv:1802.05313*.
- Garg, D.; Chakraborty, S.; Cundy, C.; Song, J.; and Ermon, S. 2021. Iq-learn: Inverse soft-q learning for imitation. *Advances in Neural Information Processing Systems*, 34: 4028–4039.
- Haarnoja, T.; Tang, H.; Abbeel, P.; and Levine, S. 2017. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, 1352–1361. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Hadfield-Menell, D.; Milli, S.; Abbeel, P.; Russell, S. J.; and Dragan, A. 2017. Inverse reward design. *Advances in neural information processing systems*, 30.
- Ho, J.; and Ermon, S. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29.
- Kostrikov, I.; Nachum, O.; and Tompson, J. 2019. Imitation Learning via Off-Policy Distribution Matching. In *International Conference on Learning Representations*.
- Kuhar, S.; Cheng, S.; Chopra, S.; Bronars, M.; and Xu, D. 2023. Learning to Discern: Imitating Heterogeneous Human Demonstrations with Preference and Representation Learning. In Tan, J.; Toussaint, M.; and Darvish, K., eds., *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, 1437–1449. PMLR.
- Kurin, V.; Nowozin, S.; Hofmann, K.; Beyer, L.; and Leibe, B. 2017. The atari grand challenge dataset. *arXiv preprint arXiv:1705.10998*.
- Laidlaw, C.; and Dragan, A. 2021. The Boltzmann Policy Distribution: Accounting for Systematic Suboptimality in Human Models. In *International Conference on Learning Representations*.
- Luce, R. D. 1959. *Individual choice behavior*. John Wiley.
- Mandlekar, A.; Xu, D.; Wong, J.; Nasiriany, S.; Wang, C.; Kulkarni, R.; Fei-Fei, L.; Savarese, S.; Zhu, Y.; and Martín-Martín, R. 2021. What Matters in Learning from Offline Human Demonstrations for Robot Manipulation. In *Conference on robot learning*.
- Ng, A. Y.; and Russell, S. J. 2000. Algorithms for Inverse Reinforcement Learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, 663–670.
- Pomerleau, D. A. 1991. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1): 88–97.
- Reddy, S.; Dragan, A. D.; and Levine, S. 2019. SQIL: Imitation Learning via Reinforcement Learning with Sparse Rewards. In *International Conference on Learning Representations*.
- Ross, S.; and Bagnell, D. 2010. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 661–668. JMLR Workshop and Conference Proceedings.

- Shiarlis, K.; Messias, J.; and Whiteson, S. 2016. Inverse reinforcement learning from failure. *International Foundation for Autonomous Agents and Multiagent Systems*.
- Wu, Y.-H.; Charoenphakdee, N.; Bao, H.; Tangkaratt, V.; and Sugiyama, M. 2019. Imitation learning from imperfect demonstration. In *International Conference on Machine Learning*, 6818–6827. PMLR.
- Xu, H.; Zhan, X.; Yin, H.; and Qin, H. 2022. Discriminator-weighted offline imitation learning from suboptimal demonstrations. In *International Conference on Machine Learning*, 24725–24742. PMLR.
- Yang, M.; Levine, S.; and Nachum, O. 2021. TRAIL: Near-Optimal Imitation Learning with Suboptimal Data. In *International Conference on Learning Representations*.
- Zhang, S.; Cao, Z.; Sadigh, D.; and Sui, Y. 2021. Confidence-aware imitation learning from demonstrations with varying optimality. *Advances in Neural Information Processing Systems*, 34: 12340–12350.
- Ziebart, B. D. 2010. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University.
- Ziebart, B. D.; Bagnell, J. A.; and Dey, A. K. 2010. Modeling interaction via the principle of maximum causal entropy. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 1255–1262. ISBN 9781605589077.
- Ziebart, B. D.; Maas, A.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd national conference on Artificial intelligence-Volume 3*, 1433–1438.