Meta-Inverse Reinforcement Learning for Mean Field Games via Probabilistic Context Variables

Yang Chen^{1,2}, Xiao Lin³, Bo Yan³, Libo Zhang², Jiamou Liu², Neset Özkan Tan^{1,2}, Michael Witbrock^{1,2}

¹ NAOInstitute, University of Auckland, New Zealand

² School of Computer Science, University of Auckland, New Zealand

³ School of Computer Science, Beijing Institute of Technology, Beijing, China

{yang.chen,jiamou.liu,neset.tan,m.witbrock}@auckland.ac.nz, lzha797@aucklanduni.ac.nz, linxiao_zj@foxmail.com,

vanbo@bit.edu.cn

Abstract

Designing suitable reward functions for numerous interacting intelligent agents is challenging in real-world applications. Inverse reinforcement learning (IRL) in mean field games (MFGs) offers a practical framework to infer reward functions from expert demonstrations. While promising, the assumption of agent homogeneity limits the capability of existing methods to handle demonstrations with heterogeneous and unknown objectives, which are common in practice. To this end, we propose a deep latent variable MFG model and an associated IRL method. Critically, our method can infer rewards from different yet structurally similar tasks without prior knowledge about underlying contexts or modifying the MFG model itself. Our experiments, conducted on simulated scenarios and a real-world spatial taxi-ride pricing problem, demonstrate the superiority of our approach over state-of-theart IRL methods in MFGs.

Introduction

Understanding incentives among interacting agents in realworld decision-making and control tasks is a fundamental challenge in multi-agent systems. Inverse reinforcement learning (IRL) (Ng, Harada, and Russell 1999; Ng and Russell 2000) addresses this issue by inferring reward functions from expert demonstrations, offering a succinct representation of tasks. IRL in multi-agent systems serves two primary purposes. First, it aids in comprehending and predicting the objectives of interacting agents, such as determining the destinations of autonomous vehicles (You et al. 2019). Second, it enables the design of agent environments with known reward signals to guide their behaviour as desired, akin to mechanism design (Fu et al. 2021).

A prominent challenge with IRL is the "curse of the agent number", *i.e.*, an increase in agent number leads to exponential interaction complexities, resulting in impractical time and memory expenses. Fortunately, recent advancements in IRL within mean field games (MFGs) (Yang et al. 2018a; Chen et al. 2022, 2023) offer a solution. By leveraging mean-field approximation, these methods simplify interactions among many agents to only two agents (individualpopulation), mitigating the computational burden. These MFG-based approaches have demonstrated compelling outcomes across diverse large-scale multi-agent tasks. Applications include modelling and predicting social media population behaviour (Yang et al. 2018a), product pricing in expansive markets, virus propagation modelling and explaining emerging social norms (Chen et al. 2022, 2023).

While appealing, these existing methods typically rely on the theoretically powerful yet practically unideal assumption in MFGs: all agents are homogeneous, *i.e.*, they are identical in the reward function, state-action space and dynamics. Real-world scenarios, however, frequently involve demonstrated behaviour with distinct and unknown rewards. For example, in spatial pricing for taxi rides, drivers may have different preferences based on factors such as distance, origins, and passenger destinations (Ata, Barjesteh, and Kumar 2019). To address this issue, efforts have been made to generalise mean-field approximation by introducing additional type variables for each agent to differentiate reward functions (Subramanian and Mahajan 2019; Ganapathi Subramanian et al. 2020; Ghosh and Aggarwal 2020). However, this approach requires prior knowledge of these variables, making it unsuitable for handling agents with unknown types or contexts, as observed in taxi trajectories.

Additionally, incorporating types into mean-field approximation makes the model more complex, necessitating the reevaluation of important theoretical properties, such as the existence and uniqueness of an equilibrium (Ghosh and Aggarwal 2020), to ensure a well-defined corresponding IRL problem. Although such generalisations are theoretically valued, from a machine learning viewpoint, simpler models with fewer theoretical restrictions are often preferred. Given the considerations above, the question arises: *can we enable IRL to handle numerous agents with unknown reward functions without altering the mean-field approximation*?

On the other hand, an emerging branch of IRL called meta-IRL (Seyed Ghasemipour, Gu, and Zemel 2019; Yu et al. 2019; Xu et al. 2019) combines IRL and meta-learning to address similar reward structures among demonstrations from different tasks. It introduces a latent probabilistic context variable that influences the reward function and performs context-conditioned reward inference without prior knowledge of the contexts. Although currently limited only to single-agent scenarios, this approach inspires the idea

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: An overview of the mechanism of PEMMFIRL. The context variable is unknown to the framework.

of assigning probabilistic context variables (representing types) externally to a family of MFGs instead of internally to agents within an MFG. By doing so, each MFG in this family retains its original theoretical properties, enabling the development of a meta-IRL method on top of it.

This paper implements this idea into a novel IRL framework called Probabilistic Embedding for Meta Mean Field IRL (PEMMFIRL) that answers the question we raised above. It integrates meta-IRL, mean-field approximation, and latent variable models into a unified framework, as illustrated in Fig. 1. Our contributions are threefold:

- We extend the notion of MFG by introducing a probabilistic contextual variable, significantly enhancing its ability to handle heterogeneous agents without introducing additional assumptions to mean-field approximation.
- 2. We develop PEMMFIRL, an associated IRL framework for this generalised MFG model, capable of inferring reward functions from demonstrations with different context variables. Importantly, PEMMFIRL can infer rewards from structurally similar tasks without prior knowledge about underlying contexts.
- 3. Experimental results on simulated tasks and a real-world spatial taxi-ride pricing problem demonstrate the effectiveness of our approach, achieving an outstanding increase in drivers' average profit when applied to real-world passenger demand data.

Related Work

Our work is closely related to the literature on MFGs, initially introduced in the continuous setting by (Huang et al. 2006) and (Lasry and Lions 2007). Later, MFGs were extended to the discrete model by (Gomes, Mohr, and Souza 2010), which is commonly adopted in agent learning. Recently, learning MFGs has garnered significant attention (Cardaliaguet and Hadikhanloo 2017), with existing methods relying on reinforcement learning techniques (Yang et al. 2018b; Guo et al. 2019; Subramanian and Mahajan 2019; Cui and Koeppl 2021). Our method stands out by directly recovering the reward function from observed behaviour, eliminating the need for manual reward design.

IRL was introduced by (Ng and Russell 2000) for singleagent settings. Early IRL methods based on margin optimization (Ratliff, Bagnell, and Zinkevich 2006) were illdefined. To resolve this, Maximum Entropy IRL (MaxEnt IRL) was proposed by (Ziebart et al. 2008; Ziebart, Bagnell, and Dey 2010). However, MaxEnt IRL is limited to small, discrete problems due to its iterative nature in reward function tuning. To extend MaxEnt IRL to high-dimensional or continuous domains, (Fu, Luo, and Levine 2018) introduced Adversarial IRL (AIRL). It utilises a sampling-based approximation by relating MaxEnt IRL to generative adversarial networks (Goodfellow et al. 2014), enabling effective reward tuning for complex scenarios. Some recent work extends AIRL to the multi-agent setting (Yu, Song, and Ermon 2019; Fu et al. 2021) and the mean-field setting (Yang et al. 2018a; Chen et al. 2023). However, they are limited to either the countable-agent or the homogeneous many-agent cases.

We build our approach on the problem of meta-learning, also known as "learning to learn" (Thrun and Pratt 2012), which aims to train models to adapt quickly to new tasks. Many methods have been proposed including the memorybased methods (Duan et al. 2016; Santoro et al. 2016; Wang et al. 2016; Mishra et al. 2017), methods that learn an optimiser and/or initialisation (Andrychowicz et al. 2016; Ravi and Larochelle 2016; Finn, Abbeel, and Levine 2017; Sun et al. 2018). Our method resembles (Rakelly et al. 2019) that differentiates different tasks through a deep latent variable. Notably, meta-IRL (Xu et al. 2019; Yu et al. 2019) incorporates meta-learning into IRL in order to rapidly adapt to new tasks or domains. While these methods show promise in both tabular and high-dimensional/continuous tasks, they can only deal with a handful of agents. Unlike these prior work, with the help of MFG, our approach is capable of handling super large-scale tasks. In this sense, our work can be seen as a significant extension of IRL to address many-agent problems with multi-task demonstrations.

Preliminaries

Mean Field Games

Consider a population of homogeneous agents sharing the same finite (local) state space S and action space A. Let $\{1, 2, ..., N\}$ be an enumeration of agents. MFGs reduce the all-agent interactions to two-party interactions between a single representative agent and the population, wherein the population is embodied by an empirical distribution of states of the system, called a *mean field*, given by

$$\mu \in \Delta(\mathcal{S})$$
 such that $\mu(s) \triangleq \lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\{s_i = s\}},$

where 1 denotes the indicator function, *i.e.*, $\mathbb{1}_x = 1$ if x is true and 0 otherwise. The reward function $r : S \times A \times \Delta(S) \to \mathbb{R}$ and state transition dynamics $P: S \times A \times \Delta(S) \times S \to [0, 1]$ of the representative agent thus depend on the current state, action and the additional mean field. Let T > 0 be a finite horizon and the initial mean field μ^0 be given, the sequence of mean fields $\mu = \mu^0, \mu^1, \dots, \mu^T$ is called a *mean-field flow*. Likewise, a *policy follow* $\pi = \pi^0, \pi^1, \dots, \pi^t$ determines the agent's strategy, where $\pi^t : S \to \Delta(A)$ maps from states to a distribution over actions.

A mean-field flow μ is said to be *consistent* with a policy flow π if for all t < T, μ^{t+1} matches an individual's state marginal distribution when executing π^t . This can be formally written by the discrete-time McKean-Vlasov (MKV) equation (Carmona, Delarue, and Lachapelle 2013):

$$\mu^{t+1}(s') = \sum_{s \in \mathcal{S}} \mu^t(s) \sum_{a \in \mathcal{A}} \pi^t(a|s) \ P(s'|s, a, \mu^t).$$
(1)

Let $\tau = s^0, a^0, \dots, s^T, a^T$ denote a state-action trajectory of an individual. The trajectory distribution induced by a pair of mean field flow and policy flow can be written as:

$$p_{\mu,\pi}(\tau) = \mu^0(s^0) \prod_{t=0}^T \pi^t(a^t|s^t) P(s^{t+1}|s^t, a^t, \mu^t)$$

In particular, if μ is consistent with π , we can rewrite $p_{\mu,\pi}(\tau)$ as $p_{\mu,\pi}(\tau) = \prod_{t=0}^{T} \mu^t(s^t)\pi^t(a^t|s^t)$. A policy flow π is said to be *optimal* to a given μ if it max-

A policy flow π is said to be *optimal* to a given μ if it maximises the expected return $\mathbb{E}_{\tau \sim p_{\mu,\pi}(\tau)} [\sum_{t=0}^{T} r(s^t, a^t, \mu^t)]^{.1}$ While, an optimal policy flow may not be unique. Entropy-regularised MFG (Cui and Koeppl 2021) resolves this ambiguity by augmenting the reward with the policy entropy $\mathbb{E}_{a \sim \pi}[-\log \pi(a|s)]$, resulting in the following objective:

$$\max_{\boldsymbol{\pi}} \mathbb{E}_{\tau \sim p_{\boldsymbol{\mu}, \boldsymbol{\pi}}(\tau)} \left[\sum_{t=0}^{T} r(s^t, a^t, \mu^t) - \log \pi^t(a^t | s^t) \right].$$
(2)

The solution concept is called the *entropy-regularised mean* field Nash equilibrium (ERMFNE) which is a pair of mean field flow and policy flow (μ, π) such that π is optimal to μ and, in turn, μ is consistent with π , *i.e.*, it fulfils the condition in Eq. (1) and maximises the objective in Eq. (2). Shown by (Cui and Koeppl 2021, Theorem 3), an ERMFNE exists uniquely under certain conditions.

Inverse Reinforcement Learning for MFGs

Suppose we have no access to the reward function $r(s, a, \mu)$ but have a set of expert demonstrated trajectories $\mathcal{D} = \{\tau_j\}_{j=1}^M$ sampled from an ERMFNE $(\boldsymbol{\mu}_E, \boldsymbol{\pi}_E)$ via $s^0 \sim \mu^0, a^t \sim \boldsymbol{\pi}_E^t(a|s^t), s^{t+1} \sim P(s|s^t, a^t, \boldsymbol{\mu}_E^t)$. Mean field Adversarial IRL (MF-AIRL) (Chen et al. 2023) aims to recover the underlying reward function from demonstrations, which can be interpreted as the following optimisation problem:

$$\min D_{\mathrm{KL}}\left(p_{\boldsymbol{\mu}_{E},\boldsymbol{\pi}_{E}}(\tau) \parallel p_{\omega}(\tau)\right) \tag{3}$$

$$p_{\omega}(\tau) = \frac{1}{Z(\omega)} \left[\prod_{t=0}^{T} \mu_{\omega}^{t}(s^{t}) \right] \cdot \exp\left(\sum_{t=0}^{T} r_{\omega} \left(s^{t}, a^{t}, \mu_{\omega}^{t} \right) \right)$$

Here, r_{ω} is the ω -parameterised reward function, p_{ω} denotes the probability that a trajectory is generated under the r_{ω} induced ERMFNE denoted by $(\mu_{\omega}, \pi_{\omega})$. The summation $Z(\omega)$ denotes the partition function, *i.e.*, the sum over all trajectories. Directly optimising the objective above is intractable as we have no access to the analytical form of μ_{ω}^t , which is a result of the entanglement between the mean field and the policy in MFGs. Fortunately, we can bypass such entanglement by establishing an unbiased estimate of μ_E^t :

$$\hat{\mu}_{E}^{t}(s) = \frac{1}{M} \sum_{j=1}^{M} \mathbb{1}_{s_{j}^{t}=s}.$$
(4)

Proven in (Chen et al. 2023, Theorem 2), with μ_{ω}^{t} being substituted with $\hat{\mu}_{E}^{t}$, the solution to the above optimisation problem approaches the optimal reward parameter if the number of demonstrated trajectories is sufficiently large.

However, computing Z_{ω} is intractable if the state-action space is large. To address this issue, MF-AIRL takes the mechanism of adversarial IRL (Fu, Luo, and Levine 2018) and reframes Eq. (3) as optimising a generative adversarial network (Goodfellow et al. 2014). It uses a discriminator D_{ω} (a binary classifier) and a sequence of adaptive samplers π_{θ} (a policy flow) whose update is equivalent to improving a sampling-based approximation to Z_{ω} . Particularly, D_{ω} takes the form of $D_{\omega}(s, a, \hat{\mu}_E^t) =$ $\exp\left(f_{\omega}(s,a,\hat{\mu}_{E}^{t})\right)/(\exp\left(f_{\omega}(s,a,\hat{\mu}_{E}^{t})\right)+\pi_{\theta}^{t}(a|s)),$ where f_{ω} serves as the parameterised reward function. The update of D_{ω} is interleaved with the update of π_{θ} : D_{ω} is trained to update the reward function by distinguishing between the trajectories sampled from the expert and the adaptive samplers, *i.e.*, to maximise $\mathbb{E}_{\mu_E, \pi_E}[\log D_{\omega}] + \mathbb{E}_{\pi_{\theta}}[\log(1-D_{\omega})];$ while $\pi_{ heta}$ is trained to maximise the entropy-augmented cumulative rewards $\mathbb{E}_{\pi_{\theta}} \left[\log D_{\omega} - \log (1 - D_{\omega}) \right] = \mathbb{E}_{\pi_{\theta}} \left[f_{\omega}(s^{t}, a^{t}, \hat{\mu}_{E}^{t}) - \log \pi_{\theta}^{t}(a^{t}|s^{t}) \right]$. Under certain conditions, f_{ω} will recover the underlying reward function $r(s, a, \mu)$ (Chen et al. 2023).

Meta-Mean Field IRL with Probabilistic Context Variables

MFGs with Probabilistic Context Variables

We extend the concept of MFGs by introducing a probabilistic contextual variable $m \in \mathcal{M}$, which follows a prior distribution p(m). Here, \mathcal{M} represents a discrete value space. MFGs with different values of m can be viewed as multiple large-scale multi-agent tasks with a shared structure. A realworld example is taxi hailing, where each taxi is associated with a specific m representing the driver's trip preferences. As a result, the policy, mean field, and reward function now have dependencies on m, denoted as $\pi^t : S \times \mathcal{M} \to \Delta(\mathcal{A})$, $\mu^t : \mathcal{M} \to \Delta(\mathcal{S}), \text{ and } r : \mathcal{S} \times \mathcal{A} \times \Delta(\mathcal{S}) \times \mathcal{M} \to$ \mathbb{R} , respectively. We assume independence among the stateaction space, transition function (Finn, Abbeel, and Levine 2017; Rakelly et al. 2019), and initial mean field w.r.t. m_{1} , which leads to the following expression for the trajectory distribution conditioned on the additional context variable: $p_{\mu,\pi}(\tau|m) =$

$$\mu^{0}(s^{0})\prod_{t=0}^{T}\pi(a^{t}|s^{t},m)P(s^{t+1}|s^{t},a^{t},\mu^{t}(\cdot|m)).$$

Given the insights above and following the spirit of entropy-regularised MFG, our objective is to maximise the expected return over the additional probabilistic contextual variable while maintaining the consistency between the

¹The reward at the last step is zero (Elie et al. 2020).



Figure 2: The graphic model of the meta-MFIRL problem. Note that the mean field (population's state density) matches an individual's state marginal when executing π_E under μ_E .

mean field flow and policy flow. This can be formulated as the following constrained optimisation problem:

$$\max_{\boldsymbol{\mu},\boldsymbol{\pi}} \mathbb{E}_{m,\tau} \left[\sum_{t=0}^{T} r(s^{t}, a^{t}, \mu^{t}, m) - \log \pi^{t}(a^{t} | s^{t}, m) \right]$$
where $m \sim p(m), \tau \sim p_{\boldsymbol{\mu},\boldsymbol{\pi}}(\tau | m)$
s.t. $\boldsymbol{\mu}(s | m)$ is consistent with $\boldsymbol{\pi}(a | s, m), \forall m \in \mathcal{M}.$
(5)

Intuitively, for all $m \in \mathcal{M}$, $(\mu(s|m), \pi(a|s, m))$ constitutes the ERMFNE of the MFG specified by m.

Problem Setup

We now introduce the problem of meta-mean field IRL (meta-MFIRL) with multi-task demonstrations. Let us consider the existence of a ground-truth reward function $r(s, a, \mu, m)$, a prior distribution p(m), and a pair of *m*-conditioned mean field flow and policy flow $(\boldsymbol{\mu}_{E}(\cdot|m), \boldsymbol{\pi}_{E}(\cdot|m))$ that solves the constrained optimisation problem in Eq. (5). Given a collection of demonstrated trajectories that are *i.i.d.* samples drawn from the resulting marginal distribution $p_{\mu_E, \pi_E}(\tau) =$ $\sum_{m \in \mathcal{M}} p(m) p_{\mu_E, \pi_E}(\tau | m)$, our objective is to metalearn an inference model $q(m|\tau)$ and a reward function $f(s, a, \mu, m)$. The aim is to ensure that when presented with a new trajectory τ_E generated by sampling $m' \sim p(m)$ and $\tau_E \sim p_{\mu_E, \pi_E}(\tau | m')$, and with \hat{m} being inferred as $\hat{m} \sim q(m | \tau_E)$, both $r(s, a, \mu, \hat{m})$ and $f(s, a, \mu, \hat{m})$ yield the same solution to the problem in Eq. (5). To illustrate dependencies between variables, we depict the graphic model underlying the meta-MFIRL problem in Fig. 2.

It is important to note that we assume no access to the prior task distribution p(m), the underlying m value for each trajectory, nor the transition dynamics $P(s'|s, a, \mu)$. Additionally, we suppose that the entire supervision comes solely from the demonstrated data, meaning that requesting additional demonstrations is not allowed.

Mutual Information Regularisation over Context Variables and the Reward Function

Under MF-AIRL, we use a ψ -parameterised context variable inference model $q_{\psi}(m|\tau)$, an ω -parameterised reward function $f_{\omega}(s, a, \mu, m)$. With m being inferred by $q_{\psi}(m|\tau)$, let $(\boldsymbol{\mu}_{\omega}(\cdot|m), \boldsymbol{\pi}_{\omega}(\cdot|s, m))$ denote the ERMFNE induced by f_{ω} . The trajectory distribution under $(\mu_{\omega}, \pi_{\omega})$ is given by:

$$p_{\omega}(\tau|m) = \frac{1}{Z(\omega)} \left[\prod_{t=0}^{T} \mu_{\omega}^{t}(s^{t}) \right] \cdot e^{\sum_{t=0}^{T} f_{\omega}\left(s^{t}, a^{t}, \mu_{\omega}^{t}, m\right)},$$
(6)

where μ_{ω} denotes the induced mean field flow consistent with π_{ω} , and $Z(\omega)$ is the partition function.

Directly applying MF-AIRL without further constraints on the context variable m would result in the pre-inferred m by q_{ψ} being treated as a constant in each component of MF-AIRL. As a result, maximising the likelihood with respect to the conditional distribution defined in Eq. (6) would lead to m being ignored, making the learned reward function f_{ω} independent of m. Consequently, the learned reward function would be indistinguishable among different tasks. To establish a connection between the reward function and the context variable, it is essential to consider the mutual information between m and the trajectories τ sampled from the reward-induced distribution. This measure quantifies the amount of information gained about the context variable by observing the trajectories, serving as an ideal measure, as done in (Zhao, Song, and Ermon 2018; Yu et al. 2019) from the perspective of information theory. The mutual information between m and τ under the joint distribution $p_{\omega}(m, \tau) = p(m)p_{\omega}(\tau|m)$ is given by:

$$I_{p_{\omega}}(m;\tau) = \mathbb{E}_{m \sim p(m), \tau \sim p_{\omega}(\tau|m)} \left[\log p_{\omega}(m|\tau) - \log p(m)\right]$$

where $p_{\omega}(m|\tau)$ is the posterior distribution corresponding to the conditional distribution $p_{\omega}(\tau|m)$ as defined in Eq. (6).

Optimising the mutual information is intractable as we have no access to the prior distribution p(m), the posterior $p(m|\tau)$, nor the ω -induced conditional mean field $\mu_{\omega}^{t}(\cdot|m)$. To address this, we employ several approximations based on sampling and estimation. First, we replace $\mu_{\omega}^{t}(\cdot|m)$ with an empirical value estimated from demonstrations using $q_{\psi}(m|\tau)$, shifting its dependency from ω to ψ . Second, we use $q_{\psi}(m|\tau)$ as a variational approximation to $p_{\omega}(m|\tau)$, allowing for approximate sampling from p(m). These approximations will be further explained in the next section. Now, our final goal in meta-MFIRL is to learn a reward function and an inference model for the task type. Formally, this can be interpreted as the pursuit of the following two desiderata:

1. Reward desideratum. Match conditional distributions:

$$\mathbb{E}_{m \sim p(m)}[D_{\mathrm{KL}}(p_{\boldsymbol{\mu}_E, \boldsymbol{\pi}_E}(\tau|m)||p_{\omega}(\tau|m))] = 0$$

2. Context desideratum. Match posterior distributions:

$$\mathbb{E}_{\tau \sim p_{\omega}(\tau)}[D_{\mathrm{KL}}\left(p_{\omega}(m|\tau)||q_{\psi}(m|\tau)\right)] = 0.$$

The reward desideratum aligns with the objective in MF-AIRL, encouraging the trajectory distribution induced by the learned reward function to match the distribution of demonstrations. The context desideratum ensures a match between the context variable inference model and the posterior distribution induced by the learned reward function, facilitating $q_{\psi}(m|\tau)$ to serve as a suitable variational approximation to the unknown $p_{\omega}(m|\tau)$. This allows correct inference of the context variable from a new demonstrated trajectory. Combining mutual information as the optimisation objective and the two aforementioned desiderata as constraints, we arrive at the final formula for our target optimization problem, with its Lagrangian dual function being:

$$\min_{\omega,\psi} \alpha \cdot \mathbb{E}_{m \sim p(m)} [D_{\mathrm{KL}} (p_{\mu_E, \pi_E}(\tau | m) || p_{\omega}(\tau | m))] + \beta \cdot \mathbb{E}_{\tau \sim p_{\omega}(\tau)} [D_{\mathrm{KL}} (p_{\omega}(m | \tau) || q_{\psi}(m | \tau))] - I_{p_{\omega}(m; \tau)}.$$

By fixing Lagrangian multipliers to specific values ($\alpha =$ $\beta = 1$) that confirm the incentive to maximise the mutual information (Zhao, Song, and Ermon 2018; Yu et al. 2019), we can rewrite the above Lagrangian dual function as

$$\min_{\omega,\psi} \mathbb{E}_{m \sim p(m)} [D_{\mathrm{KL}} \left(p_{\boldsymbol{\mu}_{E},\boldsymbol{\pi}_{E}}(\tau|m) || p_{\omega}(\tau|m) \right)] - \mathbb{E}_{m \sim p(m),\tau \sim p_{\omega}(\tau|m)} [\log q_{\psi}(m|\tau)].$$
(7)

The detailed derivation is given in Appendix A.

Probabilistic Embeddings for Meta-MFIRL

To optimise the objective in Eq. (7) tractably, we need to replace the unknown $\mu_{\omega}^{t}(\cdot|m)$ (part of $p_{\omega}(\tau|m)$) with a known value that retains optimality in ideal conditions. One approach is to imitate the operation in MF-AIRL, using an estimate $\hat{\mu}_{E}^{t}(\cdot|m)$ and substituting it for $\mu_{\omega}^{t}(\cdot|m)$ (Eq.(4)). However, the dependency of $\mu_{\omega}^t(\cdot|m)$ on m prevents this substitution, as we lack access to the posterior distribution $p(m|\tau)$. Fortunately, we can use $q_{\psi}(m|\tau)$ as a variational approximation to $p(m|\tau)$ to account for the uncertainty over tasks. Using this approximation, we construct an estimate of $\mu_E^t(\cdot|m)$ according to the following rule:

$$\hat{\mu}^t_{\psi}(s|m) = \mathbb{E}_{\tau_E \sim p_{\boldsymbol{\mu}_E, \boldsymbol{\pi}_E}(\tau)} \left[q_{\psi}(m|\tau_E) \cdot \mathbb{1}_{s^t = s} \right], \qquad (8)$$

which is unbiased when ψ is trained to the optimality. Note that the replacement of $\mu_{\omega}^{t}(\cdot|m)$ with $\hat{\mu}_{\psi}^{t}(s|m)$ will make $p_{\omega}(\tau|m)$ additionally depend on ψ . We rewrite the resulting conditional distribution as

$$p_{\omega,\psi}(\tau|m) = \frac{1}{Z(\omega,\psi)} \left[\prod_{t=0}^{T} \hat{\mu}_{\psi}^{t}(s^{t}) \right] \cdot e^{\sum_{t=0}^{T} f_{\omega}\left(s^{t}, a^{t}, \hat{\mu}_{\psi}^{t}, m\right)}$$

Accordingly, our original target problem in Eq. (7) now takes the following form:

$$\min_{\omega,\psi} \mathcal{K}(\omega,\psi) - \mathcal{L}(\omega,\psi) \tag{9}$$

$$\mathcal{K}(\omega,\psi) = \mathbb{E}_{m\sim p(m)} [D_{\mathrm{KL}} \left(p_{\boldsymbol{\mu}_{E},\boldsymbol{\pi}_{E}}(\tau|m) || p_{\omega,\psi}(\tau|m) \right)],$$

$$\mathcal{L}(\omega,\psi) = \mathbb{E}_{m\sim p(m),\tau\sim p_{\omega,\psi}(\tau|m)} [\log q_{\psi}(m|\tau)].$$

We next introduce how to approximately optimise the objective in Eq. (9) with sampling-based gradient estimation. First, we observe that $\mathcal{K}(\omega, \psi)$ with fixed ψ can be maximised using the adversarial framework in MF-AIRL, where the adaptive sampler (policy) takes the forms of $\pi_{\theta}(a|s, m)$. This observation can be formally stated in the following lemma that will be useful to derive the gradients of $\mathcal{L}(\omega, \psi)$. Lemma 1. Let the adversarial framework of MF-AIRL take the adaptive samplers $\pi^t_{\theta}(a|s,m)$ and the discriminator $D_{\omega}(s,a,\hat{\mu}_{\psi}^{t},m) = \frac{\exp(f_{\omega}(s,a,\hat{\mu}_{\psi}^{t}))}{\exp(f_{\omega}(s,a,\hat{\mu}_{\psi}^{t}) + \pi_{\theta}(a|s,m))}.$ If the adaptive samplers are trained to the optimility π_{θ}^{*} w.r.t. the reward signal $\log D_{\omega} - \log(1 - D_{\omega})$, then the trajectory distribution induced by $(\hat{\mu}_{\psi}, \pi_{\theta}^*)$ matches the conditional distribution $p_{\omega,\psi}(\tau|m)$, i.e., $p_{\hat{\mu}_{\psi},\pi_{\alpha}^{*}}(\tau|m) = p_{\omega,\psi}(\tau|m)$.

Algorithm 1: PEMMFIRL Meta-Training

- 1: Input: Expert trajectories $\mathcal{D}_E = \{\tau_j\}_{j=1}^M$.
- 2: Initialisation: Parameters f_{ω} , q_{ψ} , π_{θ} .
- 3: for each iteration do
- Sample two set of trajectories $\tau_E, \tau'_E \sim \mathcal{D}_E$. 4:
- Infer a batch of context variables $\tilde{m} \sim q_{\psi}(m|\tau_E)$. 5:
- Estimate a batch of mean fields $\hat{\mu}_{\psi}(s|\tilde{m})$ with Eq. (8). 6:
- 7: Generate a set of trajectories \mathcal{D} using $\pi_{\theta}(a|s, \tilde{m})$ and $\hat{\mu}_{\psi}(s|\tilde{m})$ with the fixed \tilde{m} for each trajectory.
- 8: Update ψ to decrease $\mathcal{K}(\omega, \psi) - \mathcal{L}(\omega, \psi)$ with gradients $\frac{\partial \mathcal{K}}{\partial \psi} - \frac{\partial \mathcal{L}}{\partial \psi}$ estimated on \mathcal{D} .
- Update ω to increase $\mathcal{L}(\omega, \psi)$ with $\frac{\partial \mathcal{L}}{\partial \omega}$ estimated on 9:
- 10: Update ω to increase following objective:
- $\mathbb{E}_{\tilde{m},\tau'_{E}}[\sum_{t=0}^{T} \log D_{\omega}] + \mathbb{E}_{\tilde{m},\tau \sim \mathcal{D}}[\sum_{t=0}^{T} \log(1-D_{\omega})].$ Update θ with RL to increase the following objective: $\mathbb{E}_{\tilde{m},\tau'_{E}}[\sum_{t=0}^{T} f_{\omega} \log \pi^{t}_{\theta}(a^{t}|s^{t},\tilde{m})].$ 11:

13: **Output:** Reward function f_{ω} and inference model q_{ψ} .

Lemma 1 tells us we can instead use π^*_{θ} to generate trajectory samples from the energy-based model $p_{\omega,\psi}(\tau|m)$ that is difficult to directly sample. Now we are ready to estimate the gradients of $\mathcal{L}(\omega, \psi)$ w.r.t. ω and ψ .

Proposition 1. With $m \sim p(m)$, $\hat{\tau}$ and $\hat{\tau}' \sim p_{\hat{\mu}_{\psi}, \pi_{a}^{*}}(\tau|m)$ and f_{ω} denoting $f_{\omega}(s^t, a^t, \hat{\mu}^t_{\psi}, m)$ for short, the gradients of $\mathcal{L}(\omega, \psi)$ w.r.t. ω and ψ can be estimated with:

$$\mathbb{E}_{m,\hat{\tau}}\left[\log q_{\psi}(m|\hat{\tau})\left[\sum_{t=0}^{T}\frac{\partial f_{\omega}}{\partial \omega} - \mathbb{E}_{\hat{\tau}'}\left[\sum_{t=0}^{T}\frac{\partial f_{\omega}}{\partial \omega}\right]\right]\right] and \mathbb{E}_{m,\hat{\tau}}\left[\log q_{\psi}(m|\hat{\tau})(\kappa(\hat{\tau},m) - \mathbb{E}_{\hat{\tau}'}[\kappa(\hat{\tau}',m)]) + \frac{\partial \log q_{\psi}(m|\hat{\tau})}{\partial \psi}\right],$$

where
$$\kappa(\tau, m) = \sum_{t=0}^{T} \left[\left(\frac{\partial f_{\omega}}{\partial \hat{\mu}_{\psi}^{t}} + \frac{1}{\hat{\mu}_{\psi}^{t}(s^{t}|m)} \right) \frac{\partial \hat{\mu}_{\psi}^{t}(s^{t}|m)}{\partial \psi} \right]$$

Proof. See Appendix B.

Proof. See Appendix B.

Note that $\frac{\partial f_{\omega}}{\partial \hat{a}^t}$ is a function of ω and depends on the model

(e.g., neural networks) for f_{ω} , and calculating $\frac{\partial \hat{\mu}^{t}_{\psi}(s^{t}|m)}{\partial \psi}$ is straightforward according to its definition in Eq. (8). Also note that the expectations in both estimates above are taken over the prior p(m). Since we have no access to p(m) but have a set of demonstrated trajectories, we can generate synthetic samples for m using the following generative process

$$\tau_E \sim p_{\boldsymbol{\mu}_E, \boldsymbol{\pi}_E}(\tau), \tilde{m} \sim q_{\psi}(m|\tau_E), \tag{10}$$

which matches p(m) when ω and ψ are trained to optimality. We are left to estimate the gradient of $\mathcal{K}(\omega, \psi)$ w.r.t. ψ .

Proposition 2. With $\tilde{\tau} \sim p_{\hat{\mu}_{\psi}, \pi^*_{\theta}}(\tau | \tilde{m})$ and \tilde{m} by Eq. (10), the gradient of $\mathcal{K}(\omega, \psi)$ w.r.t. ψ can be estimated with:

$$\mathbb{E}_{\tau_E,\tilde{m}}[\mathbb{E}_{\tilde{\tau}}[\kappa(\tilde{\tau},\tilde{m})] - \kappa(\tau_E,\tilde{m})],$$

Proof. See Appendix C.

Now we are ready to introduce our proposed framework termed Probabilistic Embeddings for Meta-Mean Field IRL (PEMMFIRL) where we alternatively update an ω parameterised reward function $f_{\omega}(s, a, \mu, m)$ and a ψ parameterised context variable inference model $q_{\psi}(m|\tau)$. We update ω by $\frac{\partial}{\partial \omega} \mathcal{L}(\omega, \psi)$ and invoking MF-AIRL to solve $\mathcal{K}(\omega, \psi)$ given the current ψ , where we train a discriminator D_{ω} given in Lemma 1 and adaptive samplers π_{θ} as

$$\max_{\omega} \mathbb{E}_{\tau_{E},\tilde{m}} \left[\sum_{t=0}^{T} \log D_{\omega} \right] + \mathbb{E}_{\tilde{m},\hat{\tau}} \left[\sum_{t=0}^{T} \log(1 - D_{\omega}) \right]$$

and
$$\max_{\theta} \mathbb{E}_{\tilde{m},\hat{\tau}} \left[\sum_{t=0}^{T} f_{\omega} - \log \pi_{\theta}^{t}(a^{t} | s^{t}, \tilde{m}) \right].$$

The update of ψ is according to $\frac{\partial}{\partial \psi} \mathcal{K}(\omega, \psi)$ and $\frac{\partial}{\partial \psi} \mathcal{L}(\omega, \psi)$. As a summary, we present the pseudocode of PEMMIRL meta-training in Alg. 1. The pseudocode of the meta-test is deferred until Appendix D.

Experiments

We experimentally seek answers to two critical questions: (1) Can PEMMFIRL accurately recover task type and reward function from demonstrated trajectories of multiple tasks? (2) Can PEMMFIRL effectively learn good policies for each task type? We evaluate our algorithm on simulated MFG environments and a real-world spatial pricing problem for taxi rides. Detailed task descriptions are provided in Appendix E, while specific settings for dataset preprocessing, hyper-parameters, network architectures, and hardware environments can be found in Appendix F.

Simulated Tasks

Task Descriptions. There are three simulated MFG environments ordered in increasing complexity: (1) Virus infection (VIRUS) (Cui and Koeppl 2021): Agents choose between "social distancing" or "going out", affecting susceptible, infection and recovery probabilities. (2) Malware spread (MALWARE) (Huang and Ma 2017): This model is representative of several problems with positive externalities, such as flu vaccination and economic models involving the entry and exit of firms. (3) Investment in product quality (INVEST) (Subramanian and Mahajan 2019): This model captures investment decisions in a fragmented market with multiple competing firms producing the same product.

Settings. For each task, we define different reward functions $r(s, a, \mu, m)$ based on context variables m. As in (Chen et al. 2022), for each m, the expert training process involves iteratively updating $\mu_E(\cdot|m)$ and $\pi_E(\cdot|m)$ until convergence: π_E is updated using backward induction, while μ_E is updated using the MKV equation (Eq. (1)). Demonstrated trajectories (each with a length of 50) are generated using $(\mu_E(\cdot|m), \pi_E(\cdot|m))$ with m drawn from a uniform distribution p(m). The IRL algorithms have no access to $r(s, a, \mu, m), m$, or p(m).

Baselines. Two state-of-the-art methods are compared: (1) **Population-level IRL (PLIRL)** (Yang et al. 2018a): A centralised method that converts MFG to MDP and performs



Figure 3: Results for simulated tasks. The curves and bars are the median and variance over ten independent runs.

IRL on this MDP. It assumes cooperation among agents to maximise the population's average rewards. (2) Mean Field Adversarial IRL (MF-AIRL) (Chen et al. 2023): This method leverages adversarial learning to efficiently infer rewards for entropy-regularised MFGs.

Performance Metrics. The quality of a learned reward function f_{ω} is evaluated by comparing its induced π_{ω} with the expert policy π_E . Two metrics are used to measure this difference – statistical distance and cumulative rewards:

- 1. Policy deviation. The expectation over KL-divergence: $\mathbb{E}_{m \sim p(m)} [\sum_{t=0}^{T} \sum_{s \in \mathcal{S}} D_{\mathrm{KL}}(\pi_{E}^{t}(\cdot|s,m) \parallel \pi_{\omega}^{t}(\cdot|s,m))]$ measures the statistical distance between two policies.
- 2. *Expected return.* The difference between two expected returns of $(\mu_{\omega}, \pi_{\omega})$ and (μ_E, π_E) under the ground-truth reward function $r(s, a, \mu, m)$ and the prior p(m).

Results. As depicted in Fig. 3, the learned policy by PEMM-FIRL shows a minor deviation from the expert policy, leading to a slight difference in expected return under the ground-truth reward function. Since PLIRL and MF-AIRL cannot handle different context variables, they both exhibit large deviations. Moreover, PLIRL shows the largest deviation as it is not suitable for non-fully cooperative environments. PEMMFIRL experiences a slight decrease in time performance due to the additional input parameters required for the context inference model $q(m|\tau)$. However, this performance decrease remains acceptable. Overall, PEMMFIRL performs well in various numerical experimental tasks, producing accurate results similar to expert performance. It demonstrates good stability during iterations, with a significantly lower average variance (32%) in expected returns and policy deviations than other algorithms, indicating low fluctuations during the iterations.

Spatial Pricing for Taxi Rides

Task Description. Spatial pricing for taxi rides aims to implement different pricing strategies in diverse areas to address the mismatched supply and demand in taxi services. By applying higher pricing in areas with high demand and low supply, taxis are encouraged to move to those regions when empty, increasing availability. The problem can be addressed using mean-field games, where drivers make deci-



Figure 4: Illustration of the initial setup. From left to right: distribution of passengers, distribution of taxis, and distribution of passengers' travelling distances.



Figure 5: Illustration for the underlying price multipliers and the learned policies (with different context variables) at the initial step for drivers at the starred location. A driver takes passengers in a darker area with a higher probability.

sions based on passenger density (Ata, Barjesteh, and Kumar 2019). By taking the value of PEMMIRL in the sense of pricing by differentiating drivers' personal preferences, we seek to improve drivers' profits further.

Settings. Real-world taxi operation records, including passenger travel time, departure locations, and destinations, are used as experimental data for passenger travel demand. This data is extracted from the New York Yellow Taxi Dataset (Ata, Barjesteh, and Kumar 2019). To ensure sufficient data in each segmented node, we select a subset of the data with one-month records that fall within part of Queens borough of New York City (see Fig. 4) and divide it into a 10×10 grid using a boundary of 0.01° for longitude and latitude.

Ata, Barjesteh, and Kumar (2019) developed an empirical MFG model that fits the dataset. The driver's reward function calculates revenue considering destinations and deducts movement-related losses, such as fuel consumption and vehicle depreciation. It applies a surcharge based on the driver's origin. The origin-only reward function reported in (Ata, Barjesteh, and Kumar 2019) is expressed as follows:

$$r(s=i,a=j,\mu) = (\eta^{0.5265} - \eta^{0.5265}_s) \cdot f(i,j,\mu),$$

where *i* and *j* represent the driver's origin area and the passenger's location, respectively; μ denotes the current density of empty taxis; *f* is a fixed function to calculate the standard profit based on the passenger demand and passengers' travelling distances; η is a constant upper limit for the price passengers are not willing to pay more than; η_s is referred to as the *price multiplier*, varying across regions to determine the surcharge scale of each area. As in (Ata, Barjesteh, and Kumar 2019), we set $\eta = 2.33$ and adopt η_s shown in Fig. 5.

To additionally consider the dependency of the reward function on the context variable, we assume a context-conditioned profit $f(i, j, \mu, m)$ (modelled by a neural net-

| Value of η | Decay rate of served passengers | Increase rate of average profit | Increased fare per ride |
|-----------------|---------------------------------|---------------------------------|----------------------------|
| 5 | - 0.4% | +2.8% | +\$0.1308 |
| 10 | - 0.5% | +2.3% | +\$0.1074 |
| 15 | - 0.6% | +3.4% | +\$0.1589 |
| 20 | - 0.7% | +3.1% | +\$0.1448 |

Table 1: Increase rate of the average profit and increased fare per ride under different decay rates of served passengers.

work) with a discrete context variable $m \in \{1, 2\}$. Intuitively, we believe m captures the impact of drivers' preferences, *e.g.*, trip distances, over profits. A trajectory has a total of 120 steps, and each step simulates five minutes. The initial setup for trajectory sampling and reward (profit) calculation in terms of distributions of the passengers, the taxis and passengers' travelling distances is illustrated in Fig. 4.

Baseline. We compare all drivers' average profit induced by the learned policy of PEMMFIRL with that induced by the optimal policy of the empirical MFG model (Ata, Barjesteh, and Kumar 2019), under the context-conditioned profit.

Results. Fig. 5 demonstrates the learned policy by PEMM-FIRL based on the inferred rewards functions. Applying this new pricing strategy and the corresponding policy to actual passenger demand data decreases served passengers compared to the original data. This is expected because no additional passenger demands are introduced beyond the dataset. By adjusting the upper limit of price η , different decay rates of passengers and increased rates (compared to the profits under the empirical MFG model in (Ata, Barjesteh, and Kumar 2019)) of drivers' average profit are obtained. Note that changing η does not affect the optimal policy, as it is a constant. The results in Tab. 1 indicate that for a relatively small decay proportion, the model increases the average profit of all drivers by 2.8% (\$0.1308 per ride) with a 0.4% reduction in passenger numbers. Further analysis reveals that most reduced passengers are those on short trips. For a higher decay proportion, the model increases the average profit of all drivers by 3.1% (\$0.1448 per ride) with a 0.7% reduction in passenger numbers. Overall, this suggests that the learned policy effectively raises the profit of taxi drivers.

Conclusions

This paper addresses the incapacity of IRL in the face of a large number of agents with distinct and unknown rewards. By introducing probabilistic contextual variables, we extend MFGs to handle heterogeneous agents without altering mean-field approximation. Based on this generalisation, we develop the Probabilistic Embeddings for Meta-Mean Field IRL (PEMMFIRL), allowing inferring reward functions from mixed-type demonstrations without prior knowledge of types. Experiments on simulated tasks and a realworld taxi-ride pricing problem reveal improvements over state-of-the-art IRL methods in MFGs.

Acknowledgments

This work was supported by a grant from the New Zealand Tertiary Education Commission and by the Strong AI Lab at the University of Auckland.

References

Andrychowicz, M.; Denil, M.; Gomez, S.; Hoffman, M. W.; Pfau, D.; Schaul, T.; Shillingford, B.; and De Freitas, N. 2016. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29.

Ata, B.; Barjesteh, N.; and Kumar, S. 2019. Spatial pricing: An empirical analysis of taxi rides in New York City. In *The University of Chicago Booth School of Business Chicago, IL Working paper*.

Cardaliaguet, P.; and Hadikhanloo, S. 2017. Learning in mean field games: the fictitious play. *ESAIM: Control, Optimisation and Calculus of Variations*, 23(2): 569–591.

Carmona, R.; Delarue, F.; and Lachapelle, A. 2013. Control of McKean–Vlasov dynamics versus mean field games. *Mathematics and Financial Economics*, 7(2): 131–166.

Chen, Y.; Zhang, L.; Liu, J.; and Hu, S. 2022. Individuallevel inverse reinforcement learning for mean field games. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-agent Systems*.

Chen, Y.; Zhang, L.; Liu, J.; and Witbrock, M. 2023. adversarial inverse reinforcement learning for mean field games. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-agent Systems*.

Cui, K.; and Koeppl, H. 2021. Approximately Solving Mean Field Games via Entropy-Regularized Deep Reinforcement Learning. In *International Conference on Artificial Intelligence and Statistics*, 1909–1917. PMLR.

Duan, Y.; Schulman, J.; Chen, X.; Bartlett, P. L.; Sutskever, I.; and Abbeel, P. 2016. Rl²: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*.

Elie, R.; Pérolat, J.; Laurière, M.; Geist, M.; and Pietquin, O. 2020. On the Convergence of Model Free Learning in Mean Field Games. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 7143–7150.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.

Fu, J.; Luo, K.; and Levine, S. 2018. Learning Robust Rewards with Adverserial Inverse Reinforcement Learning. In *International Conference on Learning Representations*.

Fu, J.; Tacchetti, A.; Perolat, J.; and Bachrach, Y. 2021. Evaluating strategic structures in multi-agent inverse reinforcement learning. *Journal of Artificial Intelligence Research*, 71: 925–951.

Ganapathi Subramanian, S.; Poupart, P.; Taylor, M. E.; and Hegde, N. 2020. Multi Type Mean Field Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, 411– 419. Ghosh, A.; and Aggarwal, V. 2020. Model free reinforcement learning algorithm for stationary mean field equilibrium for multiple types of agents. *arXiv preprint arXiv:2012.15377*.

Gomes, D. A.; Mohr, J.; and Souza, R. R. 2010. Discrete time, finite state space mean field games. *Journal de Mathématiques Pures et Appliquées*, 93(3): 308–328.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.

Guo, X.; Hu, A.; Xu, R.; and Zhang, J. 2019. Learning mean-field games. In *Advances in Neural Information Processing Systems*, 4967–4977.

Huang, M.; and Ma, Y. 2016. Mean field stochastic games: Monotone costs and threshold policies. In 2016 IEEE 55th Conference on Decision and Control (CDC), 7105–7110. IEEE.

Huang, M.; and Ma, Y. 2017. Mean field stochastic games with binary actions: Stationary threshold policies. In 2017 *IEEE 56th Annual Conference on Decision and Control (CDC)*, 27–32. IEEE.

Huang, M.; Malhamé, R. P.; Caines, P. E.; et al. 2006. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3): 221–252.

Lasry, J.-M.; and Lions, P.-L. 2007. Mean field games. *Japanese Journal of Mathematics*, 2(1): 229–260.

Mishra, N.; Rohaninejad, M.; Chen, X.; and Abbeel, P. 2017. Meta-learning with temporal convolutions. *arXiv preprint arXiv:1707.03141*, 2(7): 23.

Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, 278–287.

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.

Rakelly, K.; Zhou, A.; Finn, C.; Levine, S.; and Quillen, D. 2019. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, 5331–5340. PMLR.

Ratliff, N. D.; Bagnell, J. A.; and Zinkevich, M. A. 2006. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning*, 729–736.

Ravi, S.; and Larochelle, H. 2016. Optimization as a model for few-shot learning. In *International conference on learning representations*.

Santoro, A.; Bartunov, S.; Botvinick, M.; Wierstra, D.; and Lillicrap, T. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, 1842–1850. PMLR.

Seyed Ghasemipour, S. K.; Gu, S. S.; and Zemel, R. 2019. Smile: Scalable meta inverse reinforcement learning through context-conditional policies. *Advances in Neural Information Processing Systems*, 32. Subramanian, J.; and Mahajan, A. 2019. Reinforcement learning in stationary mean-field games. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-agent Systems*, 251–259.

Sun, H.; Chen, X.; Shi, Q.; Hong, M.; Fu, X.; and Sidiropoulos, N. D. 2018. Learning to optimize: Training deep neural networks for interference management. *IEEE Transactions on Signal Processing*, 66(20): 5438–5453.

Thrun, S.; and Pratt, L. 2012. *Learning to learn*. Springer Science & Business Media.

Wang, J. X.; Kurth-Nelson, Z.; Tirumala, D.; Soyer, H.; Leibo, J. Z.; Munos, R.; Blundell, C.; Kumaran, D.; and Botvinick, M. 2016. Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.

Weintraub, G. Y.; Benkard, C. L.; and Van Roy, B. 2010. Computational methods for oblivious equilibrium. *Operations research*, 58(4-part-2): 1247–1265.

Xu, K.; Ratner, E.; Dragan, A.; Levine, S.; and Finn, C. 2019. Learning a prior over intent via meta-inverse reinforcement learning. In *International conference on machine learning*, 6952–6962. PMLR.

Yang, J.; Ye, X.; Trivedi, R.; Xu, H.; and Zha, H. 2018a. Learning Deep Mean Field Games for Modeling Large Population Behavior. In *International Conference on Learning Representations*.

Yang, Y.; Luo, R.; Li, M.; Zhou, M.; Zhang, W.; and Wang, J. 2018b. Mean Field Multi-Agent Reinforcement Learning. In *35th International Conference on Machine Learning*, volume 80, 5571–5580. PMLR.

You, C.; Lu, J.; Filev, D.; and Tsiotras, P. 2019. Advanced planning for autonomous vehicles using reinforcement learning and deep inverse reinforcement learning. *Robotics and Autonomous Systems*, 114: 1–18.

Yu, L.; Song, J.; and Ermon, S. 2019. Multi-Agent Adversarial Inverse Reinforcement Learning. In *International Conference on Machine Learning*, 7194–7201.

Yu, L.; Yu, T.; Finn, C.; and Ermon, S. 2019. Meta-inverse reinforcement learning with probabilistic context variables. *Advances in Neural Information Processing Systems*, 32.

Zhao, S.; Song, J.; and Ermon, S. 2018. The information autoencoding family: A lagrangian perspective on latent variable generative models. *arXiv preprint arXiv:1806.06514*.

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 Machine Learning*, 1255–1262.

Ziebart, B. D.; Maas, A.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, 1433–1438.