



Reinforcement Learning in the Wild with Maximum Likelihood-based Model Transfer

Downloaded from: <https://research.chalmers.se>, 2024-07-01 01:49 UTC

Citation for the original published paper (version of record):

Eriksson, H., Tram, T., Basu, D. et al (2024). Reinforcement Learning in the Wild with Maximum Likelihood-based Model Transfer. Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS, 2024: 516-524.
<http://dx.doi.org/10.48550/arXiv.2302.09273>

N.B. When citing this work, cite the original published paper.

Reinforcement Learning in the Wild with Maximum Likelihood-based Model Transfer

Hannes Eriksson

Zenseact

Chalmers University of Technology
Gothenburg, Sweden

hannes.eriksson@zenseact.com

Tommy Tram

Zenseact

Chalmers University of Technology
Gothenburg, Sweden

Debabrota Basu

Univ. Lille, Inria, CNRS

Centrale Lille, UMR 9189 – CRISTAL
Lille, France

Mina Alibeigi

Zenseact

Gothenburg, Sweden

Christos Dimitrakakis

University of Oslo

University of Neuchatel
Neuchatel, Switzerland

ABSTRACT

In this paper, we study the problem of transferring the available Markov Decision Process (MDP) models to learn and plan efficiently in an unknown but similar MDP. We refer to it as *Model Transfer Reinforcement Learning (MTRL)* problem. First, we formulate MTRL for discrete MDPs and Linear Quadratic Regulators (LQRs) with continuous state actions. Then, we propose a generic two-stage algorithm, MLEMTRL, to address the MTRL problem in discrete and continuous settings. In the first stage, MLEMTRL uses a *constrained Maximum Likelihood Estimation (MLE)*-based approach to estimate the target MDP model using a set of known MDP models. In the second stage, using the estimated target MDP model, MLEMTRL deploys a model-based planning algorithm appropriate for the MDP class. Theoretically, we prove worst-case regret bounds for MLEMTRL both in realisable and non-realisable settings. We empirically demonstrate that MLEMTRL allows faster learning in new MDPs than learning from scratch and achieves near-optimal performance depending on the similarity of the available MDPs and the target MDP.

KEYWORDS

Reinforcement Learning; Transfer Learning; Maximum Likelihood Estimation; Linear Quadratic Regulator

ACM Reference Format:

Hannes Eriksson, Tommy Tram, Debabrota Basu, Mina Alibeigi, and Christos Dimitrakakis. 2024. Reinforcement Learning in the Wild with Maximum Likelihood-based Model Transfer. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Deploying autonomous agents in the real world poses a wide variety of challenges. As in [11], we are often required to learn the real-world model with limited data, and use it to plan to achieve satisfactory performance in the real world. There might also be safety

and reproducibility constraints, which require us to track a model of the real-world environment [36]. In light of these challenges, we attempt to construct a framework that can aptly deal with optimal decision making for a novel task, by leveraging external knowledge. As the novel task is unknown, we adopt the Reinforcement Learning (RL) [38] framework to guide an agent’s learning process and to achieve near-optimal decisions.

An RL agent interacts directly with the environment to improve its performance. Specifically, in model-based RL, the agent tries to learn a model of the environment and then use it to improve performance [24]. In many applications, the depreciation in performance due to sub-optimal model learning can be paramount. For example, if the agent interacts with living things or expensive equipment, decision-making with an imprecise model might incur significant cost [30]. In such instances, boosting the model learning by leveraging external knowledge from the existing models, such as simulators [29], physics-driven engines, etc., can be of great value [41]. A model trained on simulated data may perform reasonably well when deployed in a new environment, given the novel environment is *similar enough* to the simulated model. Also, RL algorithms running on different environments yield data and models that can be used to plan in another similar enough real-life environment. In this work, we study the problem where we have access to multiple source models built using simulators or data from other environments, and we want to transfer the source models to perform efficient model-based RL in a different real-life environment.

EXAMPLE 1. *Let us consider that a company is designing autonomous driving agents for different countries in the world. The company has designed two RL agents that have learned to drive well in USA and UK. Now, the company wants to deploy a new RL agent in India. Though all the RL agents are concerned with the same task, i.e. driving, the models encompassing driver behaviors, traffic rules, signs, etc., can differ for each. For example, UK and India have left-handed traffic, while the USA has right-handed traffic. However, learning a new controller specifically for every new geographic location is computationally expensive and time-consuming, as both data collection and learning take time. Thus, the company might use the models learned for UK and USA, to estimate the model for India, and use it further to build a new autonomous driving agent (RL agent). Hence, being able to transfer the source models to the target environment allows the*



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

company to use existing knowledge to build an efficient agent faster and resource efficiently.

We address this problem of model transfer from source models to a target environment to plan efficiently. We observe that this problem falls at the juncture of *transfer learning* and *reinforcement learning* [19, 20, 42]. [20] enlists three approaches to transfer knowledge from the *source tasks* to a *target task*. (i) *Instance transfer*: data from the source tasks is used to guide decision-making in the novel task [41]. (ii) *Representation transfer*: a representation of the task, such as learned neural network features, are transferred to perform the new task [46]. (iii) *Parameter transfer*: the parameters of the RL algorithm or *policy* are transferred [34]. In our paper, the source tasks are equivalent to the source models, and the target task is the target environment. Moreover, we adopt the **model transfer reinforcement learning** approach (MTRL), which encompasses both (i) and (ii) (Section 4).

[18] describes three possible benefits of transfer learning. The first is **learning speed improvement**, i.e. decreasing the amount of data required to learn the solution. Secondly, **asymptotic improvement**, where the solution results in better asymptotic performance. Lastly, **jumpstart improvement**, where the initial proxy model serves as a better starting solution than that of one learning the true model from scratch. In this work, we propose a new algorithm to transfer RL that achieves both learning speed improvement and jumpstart improvement (Section 8). However, we might not find an asymptotic improvement in performance if compared with the best and unbiased algorithm in the true setting. Rather, we aim to achieve a model estimate that allows us to plan accurately in the target MDP (Section 6).

Contributions. We aim to answer two central questions:

- (1) *How can we accurately construct a model using a set of source models for an RL agent deployed in the wild?*
- (2) *Does the constructed model allow efficient planning and yield improvements over learning from scratch?*

In this paper, we address these questions as follows:

1. *A Taxonomy of MTRL*: First, we formulate the problem with the Markov Decision Processes (MDPs) setting of RL. We further provide a taxonomy of the problem depending on a discrete or continuous set of source models, and whether the target model is realisable by the source models (Section 4).
2. *Algorithm Design with MLE*: Following that, we design a two-stage algorithm MLEMTRL to plan in an unknown target MDP (Section 5). In the first stage, MLEMTRL uses a Maximum Likelihood Estimation (MLE) approach to estimate the target MDP using the source MDPs. In the second stage, MLEMTRL uses the estimated model to perform model-based planning. We instantiate MLEMTRL for discrete state-action (tabular) MDPs and Linear Quadratic Regulators (LQRs). We also derive a generic bound on the goodness of the policy computed using MLEMTRL (Section 6). We further provide a meta-algorithm, called Meta-MLEMTRL, to control the adaptation to the non-realizable setting (Section 7).
3. *Performance Analysis*: In Section 8, we empirically verify whether MLEMTRL improves the performance for unknown tabular MDPs and LQRs than learning from scratch. MLEMTRL exhibits learning speed improvement for tabular MDPs and LQRs. For LQRs, it incurs learning speed improvement and asymptotic improvement. We also

observe that the more similar the target and source models are, the better the performance of MLEMTRL, as indicated by the theoretical analysis. An ablation study of Meta-MLEMTRL under realisable and non-realizable settings further shows provable improvements yielded in the asymptotic and non-realizable regimes.

Before elaborating on the contributions, we posit this work in the existing literature (Section 2) and discuss the background knowledge of MDPs and MLEs (Section 3).

2 RELATED WORK

Our work on Model Transfer Reinforcement Learning is situated in the field of Transfer RL (TRL) and also is closely related to the multi-task RL and Bayesian multi-task RL literature. In this section, we elaborate on these connections.

TRL is widely studied in Deep Reinforcement Learning. [48] introduces different ways of transferring knowledge, such as *policy transfer*, where the set of source MDPs \mathcal{M}_s has a set of expert policies associated with them. The expert policies are used together with a new policy for the novel task by transferring knowledge from each policy. [34] uses this approach, where a student learner is combined with a set of teacher networks to guide learning in multi-task RL. [28] develops an actor-critic structure to learn ways to transfer its knowledge to new domains. [2] invokes generalisation across Q-functions by learning a master policy. Here, *we focus on model transfer instead of policy*.

Another seminal work in TRL, by [42] distinguishes between *multi-task learning* and *transfer learning*. Multi-task learning deals with problems where the agent aims to learn from a distribution over scenarios, whereas transfer learning makes no specific assumptions about the source and target tasks. Thus, in transfer learning, the tasks could involve different state and action spaces and different transition dynamics. Specifically, we focus on **model-transfer** [3] approach to TRL, where the state-action spaces and also dynamics can be different. [3] performs model transfer for a target task with an identical transition model. Thus, the main consideration is to transfer knowledge to tasks with the same dynamics but varying rewards. [19] assumes a context similar to that of [3], where the model dynamics are identical across environments. In our work, we rather assume that the reward function is the same, but the transition models are different. We believe this is an interesting question as the harder part of learning an MDP is learning the transition model. These works explicate a deep connection between the fields of *multi-task learning* and *TRL*. In general, TRL can be viewed as an extension of multi-task RL, where multiple tasks can either be learned simultaneously or have been learned *a priori*. This flexibility allows us to learn even in settings where the state-actions and transition dynamics are different among tasks. [33] describes a multi-task Maximum Likelihood Estimation procedure for optimal control of an aircraft. They identify a mixture of Gaussians, where the mixture is over each of the tasks. Here, we adopt an MLE approach to TRL to optimise performance for the target MDP (or a target task) rather than restricting to a mixture of Gaussians.

The Bayesian approach to multi-task RL [21, 44] tackles the problem of learning jointly how to act in multiple environments. [21] handles the *open-world assumption*, i.e. the number of tasks is unknown. This allows them to transfer knowledge from existing

tasks to a novel task, using value function transfer. However, this is significantly different from our setting, as we are considering model-based transfer. Further, we adopt an MLE-based framework in lieu of the full Bayesian procedure described in their work. In Bayesian RL, [39] also investigates a learning technique to generalise over multiple problem instances. By sampling a large number of instances, the method is expected to learn how to generalise from the existing tasks to a novel task. We do not assume access to such prior or posterior distributions to sample from.

There is another related line of work, namely multi-agent transfer RL [9]. For example, [22] develops a TRL framework for autonomous driving using federated learning. They accomplish this by aggregating knowledge for independent agents. This setting is different from general transfer learning but could be incorporated if the source tasks are learned simultaneously with the target task. This requires cooperation among agents, which is out of the scope.

3 BACKGROUND

Here, we introduce the important concepts on which this work is based upon. Firstly, we introduce the way we model the dynamics of the tasks. Secondly, we describe the Maximum Likelihood Estimation framework used in this work.

Markov Decision Process (MDP). We study sequential decision-making problems that can be represented as MDPs [31]. An MDP $\mu = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma)$ consists of a discrete or continuous state space denoted by \mathcal{S} , a discrete or continuous action-space \mathcal{A} , a reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ which determines the quality of taking action a in state s , and a transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ inducing a probability distribution over the successor states s' given a current state s and action a . Finally, in the infinite-horizon formulation, a discount factor $\gamma \in [0, 1)$ is assigned. The overarching objective for the agent is to compute a decision-making policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximises the expected sum of future discounted rewards up until the horizon T : $V_{\mu}^{\pi}(s) = \mathbb{E} \left[\sum_{t=0}^T \gamma^t \mathcal{R}(s_t, a_t) \right]$. $V_{\mu}^{\pi}(s)$ is called the value function of policy π for MDP μ . Let $V_{\mu}^* = V_{\mu}^{\pi^*}$ denote the optimal value function. The technique used to obtain the the optimal policy $\pi^* = \sup_{\pi} V_{\mu}^{\pi}$ depends on the MDP class. The MDPs with discrete state-action spaces are referred to as tabular MDPs. In this paper, we also study a class of MDPs with continuous state-action spaces, namely Linear Quadratic Regulators (LQRs) [15]. In tabular MDPs, we employ VALUEITERATION [31] for model-based planning, whereas in the LQR setting, we use RICCATITERATION [43].

The standard metric used to measure the performance of a policy π [6] for an MDP μ is *regret* $R(\mu, \pi)$. Regret is the difference between the optimal value function and the value function of π . In this work, we extend the definition of regret for MTRL, where the optimality is taken for a policy class in the target MDP.

Maximum Likelihood Estimation (MLE). One of the most popular methods of constructing point estimators is the *Maximum Likelihood Estimation* [7] framework. Given a density function $f(x | \theta_1, \dots, \theta_n)$ and associated i.i.d. data X_1, \dots, X_t , the goal of the MLE scheme is to maximise, $\ell(\theta | x) \triangleq \ell(\theta_1, \dots, \theta_n | x_1, \dots, x_t) \triangleq \log \prod_{i=1}^t f(x_i | \theta_1, \dots, \theta_n)$. $\ell(\cdot)$ is called the log-likelihood function. The set of parameters θ maximising $\ell(\theta | x)$ is called the *maximum likelihood estimator* of θ given the data X_1, \dots, X_t . MLE has many desirable properties that we leverage in this work. For example, the

MLE satisfies *consistency*, i.e. under certain conditions, it achieves optimality even for *constrained* MLE. An estimator being consistent means that if the data X_1, \dots, X_t is generated by $f(\cdot | \theta)$ and as $t \rightarrow \infty$, the estimate almost surely converges to the true parameter θ . [16] shows that MLE admits the consistency property given the following assumptions hold. The model is *identifiable*, i.e. the densities at two parameter values must be different unless the two parameter values are identical. Further, the parameter space is *compact* and *continuous*. Finally, if the log-density is *dominated*, one can establish that MLE converges to the true parameter almost surely [25]. For problems where the likelihood is unbounded, flat, or otherwise unstable, one may introduce a penalty term in the objective function. This approach is called *penalised maximum likelihood estimation* [8, 27]. As we in our work are mixing over known parameters, we do not need to add regularisation to our objective to guarantee convergence.

In this work, we iteratively collect data and compute new point estimates of the parameters and use them in our decision-making procedure. To carry out MLE, a likelihood function has to be chosen. In this work, we investigate two such likelihood functions in Section 5, one for each respective model class.

4 A TAXONOMY OF MODEL TRANSFER RL

Now, we formally define the Model Transfer RL problem and derive a taxonomy of settings encountered in MTRL.

4.1 MTRL: Problem Formulation

Let us assume that we have access to a set of source MDPs $\mathcal{M}_s \triangleq \{\mu_i\}_{i=1}^m$. The individual MDPs can belong to a finite or infinite but compact set depending on the setting. For example, for tabular MDPs with finite state-actions, this is always a finite set. Whereas for MDPs with continuous state-actions, the transitions can be parameterised by real-valued vectors/matrices, corresponding to an infinite but compact set. Given access to \mathcal{M}_s , we want to find an optimal policy for an unknown target MDP μ^* that we encounter while deploying RL in the wild. At each step t , we use \mathcal{M}_s and the data observed from the target MDP $D_{t-1} \triangleq \{s_0, a_0, s_1, \dots, s_{t-1}, a_{t-1}, s_t\}$ to construct an estimate of μ^* , say $\hat{\mu}^t$. Now, we use $\hat{\mu}^t$ to run a model-based planner, such as VALUEITERATION or RICCATITERATION, that leads to a policy π^t . After completing this planning step, we interact with the target MDP using π^t that yields an action a_t , and leads to observing s_{t+1}, r_{t+1} . We update the dataset with these observations: $D_t \triangleq D_{t-1} \cup \{a_t, s_t\}$. Here, we assume that all the source and target MDPs share the same reward function \mathcal{R} . We do not restrict the state-action space of target and source MDPs.

Our goal is to compute a policy π^t that performs as close as possible with respect to the optimal policy π^* for the target MDP as the number of interactions with the target MDP $t \rightarrow \infty$. This allows us to define a notion of regret for MTRL: $R(\mu^*, \pi_t) \triangleq V_{\mu^*}^* - V_{\mu^*}^{\pi_t}$. Here, π_t is a function of the source models \mathcal{M}_s , the data collected from target MDP D_t , and the underlying MTRL algorithm. The goal of an MTRL algorithm is to minimise $R(\mu^*, \pi_t)$. For the parametric policies π_{θ} with $\theta \in \Theta \subset \mathbb{R}^d$, we can specialise the regret further for this parametric family: $R(\mu^*, \pi_{\theta_t}) = V_{\mu^*}^{\pi_{\theta_t}} - V_{\mu^*}^{\pi_{\theta^*}}$. For example, for LQRs, we by default work with linear policies. We use this notion of regret in our theoretical and experimental analysis.




Figure 1: An illustration of the MTRL setting. The source models \mathcal{M}_s are the red boxes. The green area is the convex hull $C(\mathcal{M}_s)$ spanned by the source models. The target MDP μ^* is displayed in blue, and the best proxy model is contained in the convex hull μ in yellow. Finally, the estimator of the best proxy model $\hat{\mu}$ is shown in purple.

4.2 Three Classes of MTRL Problems

We begin by illustrating MTRL using Figure 1. In the figure, the source MDPs \mathcal{M}_s are depicted in red. This green area is the convex hull spanned by the source models $C(\mathcal{M}_s)$. The target MDP μ^* , the best representative within the convex hull of the source models μ , and the estimated MDP $\hat{\mu}$ are shown in blue, yellow, and purple, respectively. If the target model is *inside* the convex hull, we call it a **realisable** setting, whereas if the target model is *outside* (as in Figure 1), then we have a **non-realisable** setting.

Figure 1 also shows that the total deviation of the estimated model from the target model depends on two sources of errors: (i) **realisability**, i.e. how far is the target MDP μ^* from the convex hull of the source models $C(\mathcal{M}_s)$ available to us, and (ii) **estimation**, i.e. how close is the estimated MDP $\hat{\mu}$ to the best possible representation μ of the target MDP. In the realisable case, the realisability gap can be reduced to zero, but not otherwise. This approach allows us to decouple the effect of the expressibility of the source models and the goodness of the estimator.

Now, we further elaborate on these three classes and the corresponding implications of performing MLE.

I. Finite and Realisable Plausible Models. If the true model μ^* is one of the target models, i.e. $\hat{\mu} \in \mathcal{M}_s$, we have to identify the target MDP from a finite set of plausible MDPs. Thus, the corresponding MLE involves a finite set of parameters, i.e. the parameters of the source MDPs \mathcal{M}_s . We compute the MLE $\hat{\mu}$ by solving the optimisation problem:

$$\hat{\mu} \in \arg \max_{\mu' \in \mathcal{M}_s} \log \mathbb{P}(D_t | \mu'), D_t \sim \mu^*. \quad (1)$$

This method may serve as a reasonable heuristic for the TRL problem, where the target MDP is the same as or reasonably close to one of the source MDPs. However, this method will potentially be sub-optimal if the target MDP is too different from the source MDPs. Even if μ^* lies within the convex hull of the source MDPs (the green area in Figure 1), this setting restricts the selection of a model to one of the red boxes. Thus, this setting fails to leverage the

expressiveness of the source models as MLE allows us to accurately estimate models which are also in $C(\mathcal{M}_s)$. Thus, we focus on the two settings described below.

II. Infinite and Realisable Plausible Models. In this setting, the target MDP μ^* is in the convex hull $\mu^* \in C(\mathcal{M}_s)$ of the source MDPs. Thus, for Class I, we extend the parameter space considered in MLE to an infinite but compact parameter set.

Let us define the convex hull as $C(\mathcal{M}_s) \triangleq \{\mu_1 w_1 + \dots + \mu_m w_m \mid \mu_i \in \mathcal{M}_s, w_i \geq 0, i = 1, \dots, m, \sum_{i=1}^m w_i = 1\}$. Then, the corresponding MLE problem with the corresponding likelihood function is

$$\hat{\mu} \in \arg \max_{\mu' \in C(\mathcal{M}_s)} \log \mathbb{P}(D_t | \mu'), D_t \sim \mu^*. \quad (2)$$

Since $C(\mathcal{M}_s)$ induces a compact subset of model parameters $\mathcal{M}' \subset \mathcal{M}$, Equation (2) leads to a *constrained maximum likelihood estimation problem* [1]. It implies that if the parameter corresponding to the target MDP is in \mathcal{M}' , it can be correctly identified. In the case where the optimum lies inside, we can use constrained MLE to accurately identify the true parameters given enough experience from μ^* . This approach allows us to leverage the expressibility of the source models completely. However, μ^* might lie outside or on the boundary. Either of them may pose problems for the optimiser.

III. Infinite and Non-realisable Plausible Models. This class is similar to Class II with the important difference that the true parameter μ^* is outside the convex hull of source MDPs $C(\mathcal{M}_s)$, and thus, the corresponding parameter is not in the induced parameter subset \mathcal{M}' . This key difference means the true parameters cannot be correctly identified. Instead, the objective is to identify the best proxy model $\mu \in \mathcal{M}'$. The performance loss for using μ instead of μ^* is intimately related to the model dissimilarity $\|\mu^* - \mu\|_1$. This allows us to describe the limitation of expressivity of the source models by defining the *realisability gap*: $\epsilon_{\text{Realise}} \triangleq \min_{\mu \in C(\mathcal{M}_s)} \|\mu^* - \mu\|_1$. The realisability gap becomes important while dealing with continuous state-action MDPs with parameterised dynamics, such as LQRs.

5 MLEMTRL: MTRL WITH MAXIMUM LIKELIHOOD MODEL TRANSFER

Now, we present the proposed algorithm, MLEMTRL. The algorithm consists of two stages, a *model estimation* stage, and a *planning* stage. After having obtained a plan, then the agent will carry out its decision-making in the environment to acquire new experiences. We sketch an overview of MLEMTRL in Algorithm 1. For completeness, we also provide an extension to MLEMTRL called Meta-MLEMTRL. This extension combines the MLEMTRL estimated model with the empirical model of the target task. This allows us to identify the true model even in the non-realisable setting. The details of this algorithm are available in Section 7.

5.1 Stage 1: Model Estimation

The first stage of the proposed algorithm is *model estimation*. During this procedure, the likelihood of the data needs to be computed for the appropriate MDP class. In the tabular setting, we use a product of multinomial likelihoods, where the data likelihood is over the distribution of successor states s' for a given state-action pair (s, a) . In the LQR setting, we use a linear-Gaussian likelihood, which is also expressed as a product over data observed from the target MDP.

Algorithm 1 Maximum Likelihood Estimation for Model-based Transfer Reinforcement Learning (MLEMTRL)

```

1: Input: weights  $\mathbf{w}^0$ ,  $m$  source MDPs  $\mathcal{M}_s$ , data  $D_0$ , discount
   factor  $\gamma$ , iterations  $T$ .
2: for  $t = 0, \dots, T$  do
3:   // STAGE 1: MODEL ESTIMATION //
4:    $\mathbf{w}^{t+1} \leftarrow \text{OPTIMISER}(\log \mathbb{P}(D_t | \sum_{i=1}^m w_i \mu_i), \mathbf{w}^t)$ 
5:   Estimate the MDP:  $\mu^{t+1} = \sum_{i=1}^m w_i \mu_i$ 
6:   // STAGE 2: MODEL-BASED PLANNING //
7:   Compute the policy:  $\pi^{t+1} \in \arg \max_{\pi} V_{\mu^{t+1}}^{\pi}$ 
8:   // CONTROL //
9:   Observe  $s_{t+1}, r_{t+1} \sim \mu^*(s_t, a_t)$ ,  $a_t \sim \pi^{t+1}(s_t)$ 
10:  Update the dataset  $D_{t+1} = D_t \cup \{s_t, a_t, s_{t+1}, r_{t+1}\}$ 
11: end for
12: return An estimated MDP model  $\mu^T$  and a policy  $\pi^T$ 

```

Likelihood for Tabular MDPs. The log-likelihood that we attempt to maximise in tabular MDPs is a product over $|\mathcal{S}| \times |\mathcal{A}|$ of pairs of multinomials, where p_i is the probability of event i , $n^{s,a}$ is the number of times the state-action pairs (s, a) appear in the data D_t , and $x_i^{s,a}$ is the number of times the state-action pair (s, a, s_i) occurs in the data. That is, $\sum_{i=1}^{|\mathcal{S}|} x_i^{s,a} = n^{s,a}$. Specifically,

$$\log \mathbb{P}(D_t | \mathbf{p}) = \log \left(\prod_{s,a} n^{s,a}! \prod_{i=1}^{|\mathcal{S}|} \frac{p_i^{x_i^{s,a}}}{x_i^{s,a}!} \right) \quad (3)$$

Likelihood for Linear-Gaussian MDPs. For continuous state-action MDPs, we use a linear-Gaussian likelihood. In this context, let d_s be the dimensionality of the state-space, $\mathbf{s} \in \mathbb{R}^{d_s}$ and d_a be the dimensionality of the action-space. Then, the mean function \mathbf{M} is a $\mathbb{R}^{d_s} \times \mathbb{R}^{d_a+d_s}$ matrix. The mean visitation count to the successor state \mathbf{s}'_i when an action \mathbf{a}_t is taken at state \mathbf{s}_t is given by $\mathbf{M}(\mathbf{a}_t, \mathbf{s}_t)$. We denote the corresponding covariance matrix of size $\mathbb{R}^{d_s} \times \mathbb{R}^{d_s}$ by \mathbf{S} . Thus, we express the log-likelihood by

$$\log \mathbb{P}(D_t | \mathbf{M}, \mathbf{S}) = \log \prod_{i=1}^t \frac{\exp\left(-\frac{1}{2} \mathbf{v}^\top \mathbf{S}^{-1} \mathbf{v}\right)}{(2\pi)^{d_s/2} |\mathbf{S}|^{1/2}},$$

where $\mathbf{s}'_i - \mathbf{M}(\mathbf{a}_i, \mathbf{s}_i) = \mathbf{v}$.

Model Estimation as a Mixture of Models. As the optimisation problem involves weighing multiple source models together, we add a weight vector $\mathbf{w} \in [0, 1]^m$ with the usual property that \mathbf{w} sum to 1. This addition results in another outer product over the likelihoods shown above. Henceforth, μ will refer to either the parameters associated with the product-Multinomial likelihood or the linear-Gaussian likelihood, depending on the model class.

$$\begin{aligned} \min_{\mathbf{w}} \quad & \log \mathbb{P}(D_t | \sum_{i=1}^m w_i \mu_i), D_t \sim \mu^*, \mu_i \in \mathcal{M}_s, \\ \text{s.t.} \quad & \sum_{i=1}^m w_i = 1, w_i \geq 0. \end{aligned} \quad (4)$$

Because of the constraint on \mathbf{w} , this is a constrained nonlinear optimisation problem. We can use any optimiser algorithm, denoted by **OPTIMISER**, for this purpose.

OPTIMISER. In our implementations, we use Sequential Least-Squares Quadratic Programming (SLSQP) [17] as the **OPTIMISER**. SLSQP is

a quasi-Newton method solving a quadratic programming subproblem for the Lagrangian of the objective function and the constraints.

Specifically, in Line 4 of Algorithm 1, we compute the next weight vector \mathbf{w}^{t+1} by solving the optimisation problem in Eq. (4). Let $f(\mathbf{w}) = \log \mathbb{P}(D_t | \sum_{i=1}^m w_i \mu_i)$. Further, let $\lambda = \{\lambda_1, \dots, \lambda_m\}$ and κ be Lagrange multipliers. We then define the Lagrangian \mathcal{L} ,

$$\mathcal{L}(\mathbf{w}, \lambda, \kappa) = f(\mathbf{w}) - \lambda^\top \mathbf{w} - \kappa(1 - \mathbf{1}^\top \mathbf{w}). \quad (5)$$

Here, \mathbf{w}^k is the k -th iterate. Finally, taking the local approximation of Eq. (4), we define the optimisation problem as:

$$\begin{aligned} \min_{\mathbf{d}} \quad & \frac{1}{2} \mathbf{d}^\top \nabla^2 \mathcal{L}(\mathbf{w}, \lambda, \kappa) \mathbf{d} + \nabla f(\mathbf{w}^k) \mathbf{d} + f(\mathbf{w}^k) \\ \text{s.t.} \quad & \mathbf{d} + \mathbf{w}^k \geq 0, \mathbf{1}^\top \mathbf{w}^k = 1. \end{aligned} \quad (6)$$

This minimisation problem yields the search direction \mathbf{d}_k for the k -th iteration. Applying this iteratively and using the construction above ensures that the constraints posed in Eq. (4) are adhered to at every step of MLEMTRL. At convergence, the k -th iterate, \mathbf{w}^k is considered as the next \mathbf{w}^{t+1} in Line 1 of Algorithm 1.

5.2 Stage 2: Model-based Planning

When an appropriate model μ^t has been identified at time step t , the next stage of the algorithm involves model-based planning in the estimated MDP. We describe two model-based planning techniques, **VALUEITERATION** and **RICCATITERATION** for tabular MDPs and LQRs, respectively.

VALUEITERATION. Given the model, μ^t and the associated reward function \mathcal{R} , the optimal value function of μ^t can be computed iteratively as [38]:

$$V_{\mu^t}^*(s) = \max_a \sum_{s'} \mathcal{T}_{s,s'}^a \left(\mathcal{R}(s, a) + \gamma V_{\mu^t}^*(s') \right). \quad (7)$$

The fixed-point solution to Eq.7 is the optimal value function. When the optimal value function has been obtained, one can simply select the action maximising the action-value function. Let π^{t+1} be the policy selecting the maximising action for every state, then π^{t+1} is the policy the model-based planner will use at time step $t + 1$.

RICCATITERATION. A LQR-based control system, and thus, the corresponding MDP, is defined by four system matrices [15]: $\mathbf{A}, \mathbf{B}, \mathbf{Q}, \mathbf{R}$. The matrices \mathbf{A}, \mathbf{B} are associated with the transition model $\mathbf{s}_{t+1} - \mathbf{s}_t = \mathbf{A}\mathbf{s}_t + \mathbf{B}\mathbf{a}_t$. The matrices \mathbf{Q}, \mathbf{R} dictate the quadratic cost (or reward) of a policy π under an MDP μ is

$$V_{\mu}^{\pi} = \sum_{t=0}^T \mathbf{s}_t^\top \mathbf{Q} \mathbf{s}_t + \mathbf{a}_t^\top \mathbf{R} \mathbf{a}_t.$$

Optimal policy is identified following [43] that states $\mathbf{a}_t = -\mathbf{K}\mathbf{s}_t$ at time t , where \mathbf{K} is computed using $\mathbf{A}, \mathbf{B}, \mathbf{Q}, \mathbf{R}$. We refer to Appendix B for details.

6 THEORETICAL ANALYSIS OF MLEMTRL

In this section, we further justify the use of our framework by deriving worst-case performance degradation bounds relative to the optimal controller. The performance loss is shown to be related to the realisability of μ^* under $C(\mathcal{M}_s)$. In Figure 1, we visualise the model dissimilarities, where $\|\mu - \hat{\mu}\|_1$ is the model estimation error, $\|\mu^* - \mu\|_1$ is the realisability gap and $\|\mu^* - \hat{\mu}\|_1$ the total deviation

of the estimated model. Note that by the norm on MDP, we always refer to the L_1 norm over transition matrices.

THEOREM 6.1 (PERFORMANCE GAP FOR NON-REALISABLE MODELS). Let $\mu^* = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}^*, \gamma)$ be the true underlying MDP. Further, let $\mu = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma)$ be the maximum log-likelihood

$$\mu \in \arg \max_{\mu' \in C(\mathcal{M}_s)} -\log \mathbb{P}(D_\infty | \mu'), D_\infty \sim \mu^*$$

and $\hat{\mu} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \hat{\mathcal{T}}, \gamma)$ be a maximum log-likelihood estimator of μ . In addition, let $\pi^*, \pi, \hat{\pi}$ be the optimal policies for the respective MDPs. Then, if \mathcal{R} is a bounded reward function $\forall_{(s,a)} r(s,a) \in [0, 1]$ and with ϵ_{Estim} being the estimation error and

$$\epsilon_{\text{Realise}} \triangleq \min_{\mu \in C(\mathcal{M}_s)} \|\mu^* - \mu\|_1$$

the realisability gap. Then, the performance gap is given by,

$$\|V_{\mu^*}^* - V_{\mu^*}^{\hat{\pi}}\|_\infty \leq \frac{3(\epsilon_{\text{Estim}} + \epsilon_{\text{Realise}})}{(1 - \gamma)^2}. \quad (8)$$

For the full proof, see Appendix A.1. This result is comparable to recent results such as [47] but here with an explicit decomposition into model estimation error and realisability gap terms.

REMARK 1 (BOUND ON L_1 NORM DIFFERENCE IN THE REALISABLE SETTING). It is known [4, 32, 37] that in the realisable setting, it is possible to bound the model estimation error term ϵ_{Estim} via the following argument. Let μ^* be the true underlying MDP, and $\hat{\mu}$ be an MLE estimate of μ^* , as defined in Theorem 6.1. If \mathcal{R} is a bounded reward function, i.e. $r(s,a) \in [0, 1], \forall (s,a)$, and ϵ_{Estim} is upper bound on the L_1 norm between \mathcal{T}^* and $\hat{\mathcal{T}}$. If $n^{s,a}$ be the number of times (s,a) occur together, then with probability $1 - SA\delta$,

$$\|\mathcal{T}^* - \hat{\mathcal{T}}\|_1 \leq \epsilon_{\text{Estim}} \leq \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sqrt{\frac{2 \log((2^S - 2)/\delta)}{n^{s,a}}}.$$

From this, it can be said that the total L_1 norm then scales on the order of $\mathcal{O}(SA\sqrt{S + \log(1/\delta)}/\sqrt{T})$.

This result is specific to tabular MDPs. In tabular MDPs, the maximum likelihood estimate coincides with the empirical mean model. Further details are in Appendix A.2.

REMARK 2 (PERFORMANCE GAP IN THE REALISABLE SETTING). A trivial worst-case bound for the realisable case (Section 4.2) can be obtained by setting $\epsilon_{\text{Realise}} = 0$ because by definition of the realisable case $\mu^* \in C(\mathcal{M}_s)$.

7 A META-ALGORITHM FOR MLEMTRL UNDER NON-REALISABILITY

To guarantee good performance even in the non-realizable setting, we can proceed in two steps. First, we can add the target task to the set of source tasks. Second, we can construct a meta-algorithm, combining the model estimated by MLEMTRL and the empirical estimation of the target task. In this section, we propose a meta-algorithm based on the latter. We illustrate it in Algorithm 2.

The main change in Algorithm 2 compared to Algorithm 1 is internally keeping track of the empirical model, and in Line 8, computing a posterior probability distribution over the respective models by weighting the two likelihoods together with their respective priors. The resulting algorithm then trades off its bias to the

Algorithm 2 Meta-MLEMTRL

```

1: Input: prior  $p$ , weights  $\mathbf{w}^0$ ,  $m$  source MDPs  $\mathcal{M}_s$ , data  $D_0$ , discount factor  $\gamma$ , iterations  $T$ .
2: for  $t = 0, \dots, T$  do
3:   // STAGE 1: OBTAIN MODEL WEIGHTS //
4:    $\mathbf{w}^{t+1} \leftarrow \text{MLEMTRL}(\mathbf{w}^t, \mathcal{M}_s, \mathcal{D}_t, \gamma, 1)$ 
5:   Estimate the MDP:  $\mu^{t+1} = \sum_{i=1}^m w_i \mu_i$ 
6:   Compute log-likelihood  $\ell_{\text{MLEM}}^{t+1} = \log \mathbb{P}(\mathcal{D}_t | \mu^{t+1})$ 
7:   Compute log-likelihood of empirical model  $\ell_{\text{Empirical}}^{t+1} = \log \mathbb{P}(\mathcal{D}_t | \hat{\mu}^{t+1})$ 
8:   Sample  $\tilde{\mu}^{t+1}$  as  $\mu^{t+1}$  w.p.  $\propto p \exp\left(\frac{\ell_{\text{MLEM}}^{t+1}}{\pi}\right)$  and  $\hat{\mu}^{t+1}$  w.p.  $\propto (1-p) \exp\left(\frac{\ell_{\text{Empirical}}^{t+1}}{\pi}\right)$ .
9:   // STAGE 2: MODEL-BASED PLANNING //
10:  Compute the policy:  $\pi^{t+1} \in \arg \max_{\pi} V_{\tilde{\mu}^{t+1}}^{\pi}$ 
11:  // CONTROL //
12:  Observe  $s_{t+1}, r_{t+1} \sim \mu^*(s_t, a_t), a_t \sim \pi^{t+1}(s_t)$ 
13:  Update the dataset  $D_{t+1} = D_t \cup \{s_t, a_t, s_{t+1}, r_{t+1}\}$ 
14: end for
15: return An estimated MDP model  $\tilde{\mu}^T$  and a policy  $\pi^T$ 

```

prior based on the choice of prior hyperparameter p . Since asymptotically the likelihood of the data under the empirical model is greater than the likelihood of the data under the MLEM-estimated model, as more and more data is collected the Meta-MLEMTRL algorithm performs similarly to the optimal planning using the empirical estimate. This is intended behaviour and allows for the algorithm to asymptotically plan optimally even in the non-realizable setting.

We experimentally study the behaviour of Meta-MLEMTRL and its dependence on the prior parameter p in the next section.

8 EXPERIMENTAL ANALYSIS

To benchmark the performance of MLEMTRL, we compare ourselves to a posterior sampling method (PSRL) [26], equipped with a combination of product-Dirichlet and product-Normal Inverse Gamma priors for the tabular setting, and Bayesian Multivariate Regression prior [23] for the continuous setting. In PSRL, at every round, a new model is sampled from the prior, and it learns in the target MDP from scratch. Finally, for model-based planning, we use RICCATITERATIONS to obtain the optimal linear controller for the sampled model. In the continuous action setting, we compare the performance to the baseline algorithm multi-task soft-actor critic (MT-SAC) [14, 45] and a modified MT-SAC-TRL using data from the novel task during learning. In the tabular MDP setting, we compare against multi-task proximal policy optimisation (MT-PPO) [35, 45] and similarly MT-PPO-TRL.

The objectives of our empirical study are three-fold:

- (1) How does MLEMTRL impact performance in terms of **learning speed, jumpstart improvement** and **asymptotic convergence** compared to our baseline?
- (2) What is the performance loss of MLEMTRL in the **non-realisable setting**?
- (3) How does Meta-MLEMTRL perform in the **non-realisable setting**?




Figure 2: Depicted is the average cumulative reward at every time step computed over 10 novel tasks in the realisable/non-realisable setting. The shaded regions represent the standard error of the average cumulative reward at the time step.

We conduct two kinds of experiments to verify our hypotheses. Firstly, in the upper row of Figure 2, we consider the realisable setting, where the novel task μ^* is part of the convex hull $C(\mathcal{M}_s)$. In this case, we are looking to identify an improvement in some or all of the aforementioned qualities compared to the baselines. Further, in the bottom row of Figure 2, we investigate whether the algorithm can generalise to the case beyond what is supported by the theory in Section 4.2. We begin by recalling the goals of the transfer learning problem [18].

Learning Speed Improvement: A learning speed improvement would be indicated by the algorithm reaching its asymptotic convergence with less data.

Asymptotic Improvement: An asymptotic improvement would mean the algorithm converges asymptotically to a superior solution to that one of the baseline.

Jumpstart Improvement: A jumpstart improvement can be verified by the behaviour of the algorithm during the early learning process. In particular, if the algorithm starts at a better solution than the baseline, or has a simpler optimisation surface, it may more rapidly approach better solutions with much less data.

RL Environments. We test the algorithms in a tabular MDP, i.e. Chain [10], CartPole [5], and two LQR tasks in *Deepmind Control Suite* [40]: *dm_LQR_2_1* and *dm_LQR_6_2*. Further details on experimental setups are deferred to Appendix C.1.

(1) Impacts of Model Transfer with MLEMTRL. We begin by evaluating the proposed algorithm in the Chain environment. The results of the said experiment are available in the leftmost column

of Figure 2. In it, we evaluate the performance of MLEMTRL against PSRL, MT-PPO, MT-PPO-TRL. The experiments are done by varying the slippage parameter $p \in [0.00, 0.50]$ and the results are computed for each different setup of Chain from scratch. In this experiment, we can see the baseline algorithms MT-PPO and MT-PPO-TRL perform very well. This could partially be explained by PSRL and MLEMTRL not only having to learn the transition distribution but also the reward function. The value function transfer in the PPO-based baselines implicitly transfers not only the empirical transition model but also the reward function. We can see that MLEMTRL has improved learning speed compared to PSRL in both realisable and non-realisable settings. An additional experiment with a known reward function across tasks is shown in Figure 7 in Appendix.

In the centre and rightmost columns of Figure 2, we can see the results of running the algorithms in the LQR settings with the baseline algorithms PSRL, MT-SAC and MT-SAC-TRL. The variation over tasks is given by the randomness over the stiffness of the joints in the problem. In these experiments, we can see a clear advantage of MLEMTRL compared to all baselines in terms of learning speed improvements, and in some cases, asymptotic performance.

In Figure 2, the performance metric is the average cumulative reward at every time step, for 10^5 time steps and the shaded region represents the standard deviation, where the statistics are computed over 10 independent tasks.

(2) Impact of Realisability Gap on Regret. Now, we further illustrate the observed relation between model dissimilarity and degradation in performance. Figure 3 depicts the regret against the




Figure 3: A log-log plot of regret vs. KL-divergence between the true MDP and the best proxy model in CartPole. The thick blue line is a Gaussian Process regression model fitted on observed data (in purple).

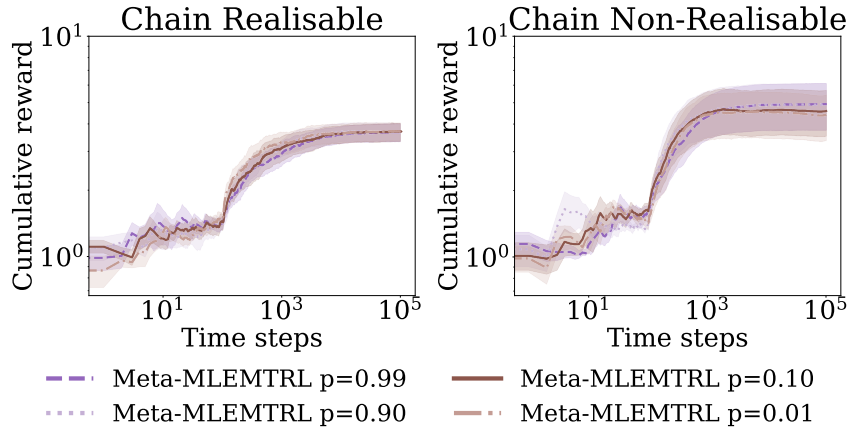


Figure 4: Figure depicting an ablation study of the prior parameter p in the Meta-MLEMTRL algorithm. The y-axis is the average cumulative reward at each time step computed over 10 novel tasks and the shaded region represents the standard error. When $p = 1$, the algorithm reduces to MLEMTRL and when $p = 0$ the algorithm reduces to standard maximum likelihood model estimation.

KL-divergence of the target model to the best proxy model in the convex set. We observe that model dissimilarity influences the performance gap in MLEMTRL. This is also justified in Section 6 where the bounds have an explicit dependency on the model difference. In this figure, only the non-zero regret experiments are shown. This is to have an idea of which models result in poor performance. As its shown, it is those models that are very dissimilar. Additional results in Figure 5 in Appendix further illustrate the dependency on model similarity.

(3) Performance of Meta-MLEMTRL under Non-realisability. In order to validate the performance of the proposed meta algorithm Meta-MLEMTRL, we perform an ablation study over the prior hyperparameter p . In Figure 4, we illustrate the results of running the Meta-MLEMTRL algorithm in the Chain environment for both the realisable and non-realisable settings. The choice of p determines how much the algorithm should be biased towards the model estimated using MLEMTRL and in the case when $p = 1$, Meta-MLEMTRL reduces to MLEMTRL. Similarly, if $p = 0$ then the algorithm will forego the MLEMTRL-estimate for the empirical estimate. Figure 4 shows that the performance of Meta-MLEMTRL is stable in long-run for different values of p . However, we identify that higher p values yield positive improvements in the cumulative reward over 10^5 steps, especially in the non-realizable setting. This indicates that the MLEMTRL-estimated model acts a good representation, while combined with the asymptotically converging empirical estimate obtained by Meta-MLEMTRL.

Summary of Results. In the experiments, we sought to identify whether the proposed algorithm shows superiority in terms of the transfer learning goals given by [18]. In the LQR-based environments, we can see a clear superiority of MLEMTRL in terms of learning speed compared to all baselines and in some cases, an asymptotic improvement. In the Chain environment, the proposed algorithm, MLEMTRL, outperforms PSRL in terms of learning

speed. Also, we perform an ablation study of Meta-MLEMTRL under realisable and non-realisable settings, demonstrating provable improvements in the asymptotic and non-realizable regimes.

9 DISCUSSIONS AND FUTURE WORK

In this work, we aim to answer two central questions.

- (1) *How can we accurately construct a model using a set of source models for an RL agent deployed in the wild?*
- (2) *Does the constructed model allow us to perform efficient planning and yield improvements over learning from scratch?*

Our answer to the first question is by adopting the *Model Transfer Reinforcement Learning* framework and weighting existing knowledge together with data from the novel task. We accomplished this by following a maximum likelihood-based approach. This has led to a novel algorithm, MLEMTRL, consisting of a model identification stage and a model-based planning stage. The second question is answered by the empirical results in Section 8 and the theoretical results in Section 6. Further, the model allows generalising to novel tasks, given that the tasks are similar enough to the existing task(s).

We motivate the use of our framework in settings where an agent is to be deployed in a new domain that is similar to existing, known, domains. We verify the quick, near-optimal performance of the algorithm in the case where the new domain is similar and we prove worst-case performance bounds of the algorithm in both the realisable and non-realizable settings. As a future work, it would be interesting to study the MTRL framework under Bayesian setting [39] and to deploy it with a risk-sensitive value function [12, 13].

ACKNOWLEDGMENTS

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. D. Basu acknowledges the Inria-Kyoto University Associate Team “RELIANT” and the ANR JCJC grant for the REPUBLIC project (ANR22-CE23-0003-01).

REFERENCES

- [1] John Aitchison and SD Silvey. 1958. Maximum-likelihood estimation of parameters subject to restraints. *The annals of mathematical Statistics* 29, 3 (1958), 813–828.
- [2] Isac Arnekvist, Danica Kragic, and Johannes A Stork. 2019. Vpe: Variational policy embedding for transfer reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 36–42.
- [3] Christopher G Atkeson and Juan Carlos Santamaria. 1997. A comparison of direct and model-based reinforcement learning. In *Proceedings of international conference on robotics and automation*, Vol. 4. IEEE, 3557–3564.
- [4] Peter Auer, Thomas Jaksch, and Ronald Ortner. 2008. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems* 21 (2008).
- [5] Andrew G Barto, Richard S Sutton, and Charles W Anderson. 1983. Neuron-like adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics* 5 (1983), 834–846.
- [6] David E Bell. 1982. Regret in decision making under uncertainty. *Operations research* 30, 5 (1982), 961–981.
- [7] George Casella and Roger L Berger. 2021. *Statistical inference*. Cengage Learning.
- [8] Gabriela Ciuperca, Andrea Ridolfi, and Jérôme Idier. 2003. Penalized maximum likelihood estimator for normal mixtures. *Scandinavian Journal of Statistics* 30, 1 (2003), 45–59.
- [9] Felipe Leno Da Silva and Anna Helena Reali Costa. 2019. A survey on transfer learning for multiagent reinforcement learning systems. *Journal of Artificial Intelligence Research* 64 (2019), 645–703.
- [10] Richard Dearden, Nir Friedman, and Stuart Russell. 1998. Bayesian Q-learning. *Aaai/iaai* 1998 (1998), 761–768.
- [11] Gabriel Dulac-Arnold, Nir Levine, Daniel J Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. 2021. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning* 110, 9 (2021), 2419–2468.
- [12] Hannes Eriksson, Debabrota Basu, Mina Alibeigi, and Christos Dimitrakakis. 2022. Sentinel: taming uncertainty with ensemble based distributional reinforcement learning. In *Uncertainty in Artificial Intelligence*. PMLR, 631–640.
- [13] Yannis Flet-Berliac and Debabrota Basu. 2022. SAAC: Safe Reinforcement Learning as an Adversarial Game of Actor-Critics. In *RLDM 2022-The Multi-disciplinary Conference on Reinforcement Learning and Decision Making*.
- [14] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*. PMLR, 1861–1870.
- [15] Rudolph Emil Kalman. 1960. A new approach to linear filtering and prediction problems. (1960).
- [16] Jack Kiefer and Jacob Wolfowitz. 1956. Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics* (1956), 887–906.
- [17] Dieter Kraft. 1988. A software package for sequential quadratic programming. *Forschungsbericht - Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt* (1988).
- [18] Pat Langley. 2006. Transfer of knowledge in cognitive systems. In *Talk, workshop on Structural Knowledge Transfer for Machine Learning at the Twenty-Third International Conference on Machine Learning*.
- [19] Romain Laroche and Merwan Barlier. 2017. Transfer reinforcement learning with shared dynamics. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [20] Alessandro Lazaric. 2012. Transfer in reinforcement learning: a framework and a survey. In *Reinforcement Learning*. Springer, 143–173.
- [21] Alessandro Lazaric and Mohammad Ghavamzadeh. 2010. Bayesian multi-task reinforcement learning. In *ICML-27th International Conference on Machine Learning*. Omnipress, 599–606.
- [22] Xinle Liang, Yang Liu, Tianjian Chen, Ming Liu, and Qiang Yang. 2023. Federated transfer reinforcement learning for autonomous driving. In *Federated and Transfer Learning*. Springer, 357–371.
- [23] Thomas Minka. 2000. *Bayesian linear regression*. Technical Report. Citeseer.
- [24] Thomas M Moerland, Joost Broekens, Aske Plaat, Catholijn M Jonker, et al. 2023. Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning* 16, 1 (2023), 1–118.
- [25] Whitney K Newey and James L Powell. 1987. Asymmetric least squares estimation and testing. *Econometrica: Journal of the Econometric Society* (1987), 819–847.
- [26] I. Osband, D. Russo, and B. Van Roy. 2013. (More) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*. 3003–3011.
- [27] Reda Ouahama, Debabrota Basu, and Odalric Maillard. 2023. Bilinear exponential family of MDPs: frequentist regret bound with tractable exploration & planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 9336–9344.
- [28] Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. 2015. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342* (2015).
- [29] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. 2018. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 3803–3810.
- [30] Athanasios S Polydoros and Lazaros Nalpantidis. 2017. Survey of model-based reinforcement learning: Applications on robotics. *Journal of Intelligent & Robotic Systems* 86, 2 (2017), 153–173.
- [31] Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- [32] Jian Qian, Ronan Fruit, Matteo Pirotta, and Alessandro Lazaric. 2020. Concentration inequalities for multinoulli random variables. *arXiv preprint arXiv:2001.11595* (2020).
- [33] Cédric Rommel, Joseph Frédéric Bonnans, Baptiste Gregorutti, and Pierre Martignon. 2017. Aircraft dynamics identification for optimal control. In *7th European Conference on Aeronautics and Space Sciences (EUCASS 2017)*.
- [34] Andrei A Rusu, Sergio Gomez Colmenarejo, Caglar Gulcehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2015. Policy distillation. *arXiv preprint arXiv:1511.06295* (2015).
- [35] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [36] Julian Skrzyński, Frederic Becker, and Falk Lieder. 2021. Automatic discovery of interpretable planning strategies. *Machine Learning* 110, 9 (2021), 2641–2683.
- [37] Alexander L Strehl and Michael L Littman. 2005. A theoretical analysis of model-based interval estimation. In *Proceedings of the 22nd international conference on Machine learning*. 856–863.
- [38] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.
- [39] Aviv Tamar, Daniel Soudry, and Ev Zisselman. 2022. Regularization Guarantees Generalization in Bayesian Reinforcement Learning through Algorithmic Stability. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 8423–8431.
- [40] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. 2018. Deepmind control suite. *arXiv preprint arXiv:1801.00690* (2018).
- [41] Matthew E Taylor, Nicholas K Jong, and Peter Stone. 2008. Transferring instances for model-based reinforcement learning. In *Joint European conference on machine learning and knowledge discovery in databases*. Springer, 488–505.
- [42] Matthew E Taylor and Peter Stone. 2009. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* 10, 7 (2009).
- [43] Jan Willems. 1971. Least squares stationary optimal control and the algebraic Riccati equation. *IEEE Transactions on automatic control* 16, 6 (1971), 621–634.
- [44] Aaron Wilson, Alan Fern, Soumya Ray, and Prasad Tadepalli. 2007. Multi-task reinforcement learning: a hierarchical bayesian approach. In *Proceedings of the 24th international conference on Machine learning*. 1015–1022.
- [45] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. 2020. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*. PMLR, 1094–1100.
- [46] Amy Zhang, Harsh Satija, and Joelle Pineau. 2018. Decoupling dynamics and reward for transfer learning. *arXiv preprint arXiv:1804.10689* (2018).
- [47] Amy Zhang, Shagun Sodhani, Khimya Khetarpal, and Joelle Pineau. 2020. Learning Robust State Abstractions for Hidden-Parameter Block MDPs. In *International Conference on Learning Representations*.
- [48] Zhuangdi Zhu, Kaixiang Lin, Anil K Jain, and Jiayu Zhou. 2023. Transfer learning in deep reinforcement learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023).