

Data-Efficient Task Generalization via Probabilistic Model-based Meta Reinforcement Learning

Arjun Bhardwaj¹, Jonas Rothfuss¹, Bhavya Sukhija¹, Yarden As¹,
 Marco Hutter¹, Stelian Coros¹, Andreas Krause¹

Abstract—We introduce PACOH-RL, a novel model-based Meta-Reinforcement Learning (Meta-RL) algorithm designed to efficiently adapt control policies to changing dynamics. PACOH-RL meta-learns priors for the dynamics model, allowing swift adaptation to new dynamics with minimal interaction data. Existing Meta-RL methods require abundant meta-learning data, limiting their applicability in settings such as robotics, where data is costly to obtain. To address this, PACOH-RL incorporates regularization and epistemic uncertainty quantification in both the meta-learning and task adaptation stages. When facing new dynamics, we use these uncertainty estimates to effectively guide exploration and data collection. Overall, this enables positive transfer, even when access to data from prior tasks or dynamic settings is severely limited. Our experiment results demonstrate that PACOH-RL outperforms model-based RL and model-based Meta-RL baselines in adapting to new dynamic conditions. Finally, on a real robotic car, we showcase the potential for efficient RL policy adaptation in diverse, data-scarce conditions.

I. INTRODUCTION

The field of Reinforcement Learning (RL) has seen remarkable advances in recent years [e.g., 1, 2]. Particularly, in gameplay and simulated robotic manipulation problems, RL agents can solve ever more complex tasks. These advances, however, largely rely on an abundance of agent-environment interactions. In contrast, in real-world robotic applications, obtaining interaction data is costly. Thus, a promising approach for real robotic platforms is *Model-based Reinforcement Learning (MBRL)*, which leverages a learned model of the environment dynamics to obtain a control policy more efficiently [e.g., 3–5].

Still, this requires a significant amount of data to learn a reliable dynamics model. This becomes even more problematic when generalization across multiple tasks with related but different dynamics is required, e.g., when we want to deploy the same robot on different surfaces or with different payloads/tools. The challenges above give rise to the following question: *How can we effectively transfer knowledge across tasks and adapt our model to changes in dynamics without the need for extensive data collection each time?*

Addressing this challenge, we propose PACOH-RL, a novel approach to model-based Meta-Reinforcement Learning (Meta-RL) which allows us to transfer prior experience

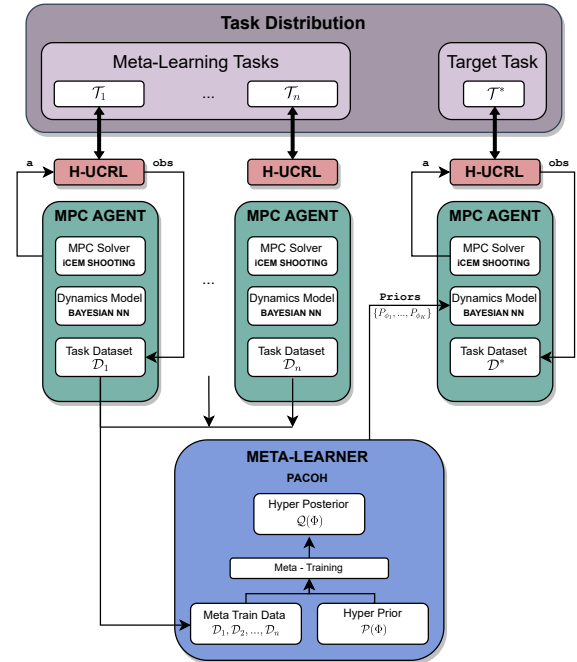


Fig. 1: The PACOH-RL framework uses datasets of transitions $\mathcal{D}_1, \dots, \mathcal{D}_n$ from previous RL tasks to meta-learn a BNN prior. Then, we equip our BNN dynamics model with the meta-learned prior. This significantly improves the sample efficiency of model-based RL on a new target task.

on the robotic platform to new dynamics conditions. In particular, we propose to *meta-learn a prior distribution over the dynamics model*, facilitating efficient (Bayesian) adaptation to new settings based on minimal interaction data. Unlike existing Meta-RL methods [e.g., 6, 7], our approach takes into account epistemic uncertainty, both throughout the meta-learning and task adaptation stages. This allows us to use an RL formulation that plans optimistically w.r.t. the epistemic uncertainty of the dynamics model [see 8]. As a result, PACOH-RL explores uncertain actions that plausibly lead to high rewards in a directed manner and, thus, efficiently adapts the model to the current dynamics conditions.

Crucially, our method is designed to operate *with only a handful of previous tasks* (i.e., dynamics settings), a regime where existing Meta-RL methods typically fail. Such efficiency is essential in real-world applications where collecting data across different tasks with different dynamics is limited.

Our experiments show that PACOH-RL can quickly adapt to new dynamics and consistently outperforms both model-based Meta-RL and standard model-based RL

*This research was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program grant agreement no. 815943 and the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40 180545. Jonas Rothfuss was supported by an Apple Scholars in AI/ML fellowship. Yarden As was supported by Hasler Foundation grant no. 21039.

¹ETH Zurich, correspondence to abhardwaj@ethz.ch

baselines. Thanks to the combination of uncertainty-aware meta-learning and directed exploration, PACOH-RL performs particularly favorably in RL environments with sparse rewards. Finally, we showcase how PACOH-RL facilitates successful transfer on a real robotic car with only a handful of meta-training datasets. This demonstrates the practicality of our approach and highlights its potential to efficiently adapt RL policies to changing conditions.

II. RELATED WORK

Model-Based Reinforcement Learning. MBRL methods are often considered for learning directly on hardware since they are generally more sample-efficient compared to model-free RL [3, 4]. However, asymptotically, model-free methods often outperform MBRL. One reason for this performance gap is the exploitation of model inaccuracies [4, 5]. Using dynamics models that are aware of epistemic uncertainty, such as ensembles or Bayesian Neural Networks (BNN), has been shown to alleviate the problem of model exploitation [4]. Building on this insight, we center our approach around BNN dynamics models. However, unlike learning a policy that is robust w.r.t. the model’s uncertainty, we use the optimistic RL objective of [8] that incentivizes exploration and data collection in areas where the model is uncertain. This allows us to reduce our model’s inaccuracies more efficiently.

Meta-Learning and Meta-RL. Meta-learning [9–11] aims to acquire useful inductive bias from a set of related learning tasks, allowing us to *quickly adapt* to a new, similar task. Some model-free Meta-RL approaches achieve this by training a sequence model to act as reinforcement learner [12, 6], meta-learning a policy initialization that can be quickly adapted to new tasks [10, 13] or conditioning the policy on a latent context variable that represents the different tasks [14, 15]. However, model-free Meta-RL methods require many tasks during meta-training. This renders them infeasible for real-world robotic problems.

Model-based meta-RL is more task and sample efficient as it focuses on meta-learning inductive bias for the dynamics model, and does not require online control of the robot during meta-training. A common model-based meta-RL approach is to condition the dynamics model on latent task variable [16–19]. Another method, proposed by [20], is to meta-learn an NN initialization for the dynamics model using MAML [10]. Our approach builds on PAC-Bayesian meta-learning [21–24]. In particular, we employ the PACOH-NN method [23] to meta-learn priors for our BNN dynamics models, which are adapted to the target task via (generalized) Bayesian inference. Our proposed method distinguishes itself from previous model-based meta-RL in the following two key characteristics. First, it features principled regularization and captures epistemic uncertainty both during the meta-learning and the task adaptation stage. Second, it uses the resulting uncertainty estimates to guide exploration and data collection to facilitate more efficient task adaptation. This allows PACOH-RL to perform positive transfer with only a handful of meta-learning datasets—a setting where previous approaches typically fail.

III. BACKGROUND

Reinforcement Learning. A discrete-time Markov decision process (MDP) is defined by the tuple $\mathcal{T} = (\mathcal{S}, \mathcal{A}, p, p_0, r, T)$. Here, $\mathcal{S} \subseteq \mathbb{R}^{d_s}$ is the state space, $\mathcal{A} \subseteq \mathbb{R}^{d_a}$ the action space, $p(s_{t+1}|s_t, a_t)$ the transition distribution, p_0 represents the initial state distribution, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a known reward function, and T the time horizon. For ease of notation, we exclude the discount factor γ in the ensuing discussions. We then define the return $R(\tau)$ as the cumulative sum of rewards along a trajectory $\tau := (s_0, a_0, \dots, s_{T-1}, a_{T-1}, s_T)$. The central objective of reinforcement learning is to find a policy $\pi(a|s)$ that optimizes the expected return $\mathbb{E}_{\tau \sim P_{\mathcal{T}}(\tau|\pi)} [\sum_{t=0}^{H-1} r(s_t, a_t)]$. Here, $P_{\mathcal{T}}(\tau|\pi) = p_0(s_0) \prod_{t=0}^{H-1} \pi(a_t|s_t) p(s_{t+1}|s_t, a_t)$ is the trajectory distribution under the MDP \mathcal{T} and the policy π .

Model-based RL. MBRL uses the collected state transition data $\mathcal{D} = \{(s_t, a_t, s_{t+1})\}$ to learn/estimate a model $\hat{p}(s_{t+1}|s_t, a_t)$ of the transition distribution, also referred to as a *dynamics model*. Often, function approximators such as neural networks are employed for this purpose. Then, the estimated dynamics model is either used to simulate trajectories to train a policy or to produce dynamics predictions/constraints for a controller. Generally, the dynamics model empowers the agent to simulate future states and rewards. Through that, the agent can anticipate the outcome of its actions without interacting with its MDP environment directly. Hence, MBRL methods are typically much more data-efficient than model-free methods.

Bayesian Neural Networks. Learning the dynamics model in model-based RL is a standard supervised learning problem. The inputs correspond to a concatenation of the current state and action, i.e., $x = [s, a]$, and the prediction target is the next state, i.e., $y = s'$, allowing us to write the training data set as $\mathcal{D} = \{(x_j, y_j)\}_{j=1}^m$. Let $h_{\theta} : \mathcal{X} \mapsto \mathcal{Y}$ be a function parameterized by a neural network (NN) with weights $\theta \in \Theta$. Using the NN mapping, we can define a conditional predictive distribution (a.k.a. likelihood function) $p(y|x, \theta) = \mathcal{N}(y|h_{\theta}(x), \sigma^2)$, where σ^2 is the variance corresponding to the aleatoric uncertainty.

Some MBRL methods (e.g., [25, 7]) simply fit a single NN h_{θ} by maximizing likelihood $p(\mathcal{D}|\theta) = \prod_{j=1}^m p(y_j|x_j, \theta)$ of the data. However, with a simple NN we cannot quantify *epistemic uncertainty*, which is crucial for directed exploration [8, 26]. In contrast, *Bayesian Neural Networks (BNNs)* maintain a distribution over NNs, allowing them to quantify uncertainty about $h_{\theta}(\cdot)$. In particular, BNNs presume a prior distribution $p(\theta)$ over the model parameters θ , which they combine with the data likelihood into a (generalized) posterior distribution $p(\theta|\mathcal{D}) \propto p(\mathcal{D}|\theta)^{\lambda} p(\theta)$. Note that in this paper, we resort to generalized Bayesian learning [27, 28] where the likelihood is tempered with $\lambda \in (0, 1)$, giving us additional robustness when the standard Bayesian assumptions are violated. To make probabilistic predictions, we typically form the predictive distribution as $p(y^*|x^*, \mathcal{D}) = \int p(y^*|x^*, \theta) p(\theta|\mathcal{D}) d\theta$ by marginalizing over the NN parameters θ .

Approximate Inference via SVGD. Since BNN posteriors are generally intractable, approximate inference techniques such as MCMC [29], variational inference (VI) [30] or particle VI methods [31, 32] are often applied. In this paper, we employ *Stein Variational Gradient Descent* (SVGD) [31] which approximates the posterior $p(\theta|\mathcal{D})$ by a set of L particles $\{\theta_1, \dots, \theta_L\}$. After initialization, SVGD iteratively transports the particles (here: NN parameters) to match $p(\theta|\mathcal{D})$ by applying a form of functional gradient descent that minimizes the KL divergence in the reproducing kernel Hilbert space induced by a kernel function $k(\cdot, \cdot)$. In particular, the update of a particle θ is computed as

$$\psi(\theta) = \frac{1}{L} \sum_{l'=1}^L \left[k(\theta_{l'}, \theta) \nabla_{\theta_{l'}} \log p(\theta_{l'}|\mathcal{D}) + \nabla_{\theta_{l'}} k(\theta_{l'}, \theta) \right] \quad (1)$$

and applied via $\theta_l \leftarrow \theta_l + \eta \psi(\theta_l)$ where η is the step size. While the first term in (1) moves the particles towards areas of higher probability, the second term, i.e., $\nabla_{\theta_{l'}} k(\theta_{l'}, \theta)$, acts as a repulsion force among the particles which ensures that they are well dispersed throughout the parameter space and do not collapse in the mode of the distribution.

IV. PACOH-RL: UNCERTAINTY-AWARE MODEL-BASED META-RL

A. Problem Statement: Meta-RL

We study the problem of Meta-RL where we face multiple MDPs that vary in their transition dynamics. While we could learn a policy that is robust to the varying dynamics, such policies are typically sub-optimal due to over-conservatism [33]. Instead, we want to swiftly adapt our agent’s behavior to the new dynamical conditions without requiring a large amount of agent-environment interactions.

Formally, we are given a sequence of MDP tasks $\mathcal{T}_1, \dots, \mathcal{T}_n$ sampled i.i.d. from the task distribution $p(\mathcal{T})$ with $\mathcal{T}_i = (\mathcal{S}, \mathcal{A}, p_i, p_0, r, H)$ where the transition probabilities $p_i(s_{t+1}|s_t, a_t)$ differ between the tasks. Our framework also seamlessly supports reward functions that may vary across tasks. However, for simplicity, we treat the reward as fixed throughout the remainder of the paper.

Suppose we have already collected (e.g., using RL) transition data $\mathcal{D}_i = \{(s, a, s')\}$ for n tasks corresponding to p_i , for $i = 1, \dots, n$. Now, we are facing a new target task $\mathcal{T}^* \sim p(\mathcal{T})$ for which we want to efficiently find an optimal policy. In particular, we focus on real-world robotic settings where n , i.e., the number of previous tasks, is small and the agent-environment interactions are costly. Hence, we want to explore the target task and find an optimal policy for it with as few interactions as possible. Out of this problem setting arise two key questions: 1) How can we transfer knowledge from the previously collected transition datasets $\mathcal{D}_1, \dots, \mathcal{D}_n$ to the new target task? 2) Which RL paradigm and exploration scheme should we use on the target task?

B. Our approach: Model-Based Meta-RL

Since data efficiency is one of our core concerns, we employ a model-based RL paradigm. To transfer knowledge from the previous MDP tasks, we use meta-learning to

extract inductive bias from the transitions of previous tasks $\mathcal{D}_1, \dots, \mathcal{D}_n$, which we then harness when estimating a dynamics model on the target task. Crucially, we choose a meta-learner and dynamics model that can reason about epistemic uncertainty. When performing RL on the target task, this allows us to explore in a directed manner towards areas of the state-action space in which we are more uncertain yet can plausibly obtain high rewards. As a result, our agent is able to swiftly collect transition data on the target task, making the dynamics model more accurate, and the resulting policy better. In the following, we explain the building blocks of our approach in more detail.

Meta-Learning a dynamics model prior. Due to its principled meta-level regularization, which allows successful meta-learning from only a handful of tasks as well as its principled treatment of uncertainty, we build on the *PAC-Bayesian meta-learning* framework [21–23]. In particular, we employ PACOH-NN [23] which meta-learns Bayesian Neural Network (BNN) priors from the meta-training data $\mathcal{D}_1, \dots, \mathcal{D}_n$. The PACOH framework uses a parametric family of priors $\{P_\phi | \phi \in \Phi\}$ over NN parameters θ . Due to computational convenience, we use Gaussian priors, i.e., $P_{\phi_k} = \mathcal{N}(\mu_{P_k}, \text{diag}(\sigma_{P_k}^2))$ with $\phi_k := (\mu_{P_k}, \ln \sigma_{P_k})$. The prior variance $\sigma_{P_k}^2$ is represented in the log-space to avoid additional positivity constraints. Employing a hyper-prior $\mathcal{P}(\phi)$ which acts as a regularizer on the meta-level, the meta-learner infers the hyper-posterior, a distribution over the prior parameters, in particular,

$$\mathcal{Q}(\phi) \propto \mathcal{P}(\phi) \exp \left(\frac{1}{\sqrt{nm} + 1} \sum_{i=1}^n \ln Z(\phi, \mathcal{D}_i) \right). \quad (2)$$

Here $\ln Z(\phi, \mathcal{D}_i) = \ln E_{\theta \sim P_\phi} \left[p(\mathcal{D}|\theta)^{1/\sqrt{n}} \right]$ denotes the generalized marginal log likelihood (MLL). Sampling from and determining the normalization constant of $\mathcal{Q}(\phi)$ is challenging. Thus, we follow [23] and approximate $\mathcal{Q}(\phi)$ by a set of K priors $P_{\phi_1}, \dots, P_{\phi_K}$ which are optimized via Stein Variational Gradient Descent (SVGD) [31] to closely resemble $\mathcal{Q}(\phi)$. By considering a distribution over priors rather than meta-learning a single prior, PACOH is able to quantify epistemic uncertainty on the meta-level. When the number of meta-learning tasks, i.e. n , is small, the sum of generalized MLLs in (2) is relatively small, and the hyper-prior keeps the uncertainty in \mathcal{Q} large. As we have more meta-learning tasks, the exponential term in (2) grows, and \mathcal{Q} becomes increasingly peaked in prior parameters that yield a large MLL across the tasks, reflecting reduced uncertainty on the meta-level.

Adapting the dynamics model to the target task. From the meta-learning stage, we have acquired a set of priors $P_{\phi_1}, \dots, P_{\phi_K}$ from the transitions of previous tasks, which give us good inductive bias towards the dynamics of our robot under varying conditions. Once we observe state transitions under the dynamical conditions of the target task \mathcal{T}^* , we can combine these empirical observations with our meta-learned prior knowledge into a BNN model. Let \mathcal{D}^* be the dataset of observed transitions on the target task. Then, the generalized BNN posterior corresponding to the meta-learned prior

Algorithm 1 PACOH-RL (MPC version)

Input: Transition datasets $\{\mathcal{D}_1, \dots, \mathcal{D}_n\}$ from previous tasks, test task \mathcal{T} , hyper-prior \mathcal{P}

- 1: $\{P_{\phi_1}, \dots, P_{\phi_K}\} \leftarrow \text{PACOH-NN}(\mathcal{D}_1, \dots, \mathcal{D}_n, \mathcal{P})$ ▷ Meta-learn set of priors to approx. \mathcal{Q}
- 2: $\mathcal{D}^* \leftarrow \emptyset$ ▷ Initialize empty transition dataset
- 3: **for** episode = 1, 2, ... **do**
- 4: **for** $k = 1, 2, \dots, K$ **do**
- 5: $\Theta_k \leftarrow \text{BNN-SVGD}(\mathcal{D}^*, P_{\phi_k})$ ▷ Train BNN with latest transition data
- 6: $\Theta := \{\Theta_1, \dots, \Theta_K\}$
- 7: $\hat{p}_\Theta \leftarrow \mathcal{N}(\hat{\mu}_\Theta(s, a), \hat{\sigma}_\Theta^2(s, a))$ ▷ Aggregate NN predictions into predictive distribution
- 8: $(s_0, a_0, \dots, a_{T-1}, s_T) \leftarrow \text{iCEM-MPC}(\hat{p}_\Theta, \mathcal{T})$ ▷ Perform rollout with MPC controller
- 9: $\mathcal{D}^* \leftarrow \mathcal{D}^* \cup \{(s_t, a_t, s_{t+1})\}_{t=0}^{T-1}$ ▷ Add transitions to dataset

P_{ϕ_k} follows as $Q_k(\theta; \mathcal{D}^*) \propto p(\mathcal{D}^* | \theta)^{1/\sqrt{n}} P_{\phi_k}(\theta)$. Since we have K priors, we also obtain K different posteriors, i.e., $Q_k(\theta; \mathcal{D}^*)$ $k = 1, \dots, K$. Similar to the meta-learning stage, we represent each posterior Q_k by a set of L NN parameters $\Theta_k = \{\theta_{k,1}, \dots, \theta_{k,L}\}$ which we optimize via SVGD to approximate the posterior density. This leaves us with $K \cdot L$ neural networks whose parameters we denote by $\Theta = \{\Theta_1, \dots, \Theta_K\}$. The K sets of NN particles represent the epistemic uncertainty on the meta-level whereas the particles within Θ_k correspond to the uncertainty on the target task. In all experiments, we use $K = L = 3$. To aggregate the individual neural networks’ predictions into a predictive distribution, we use a Gaussian approximation $\hat{p}_\Theta(s' | s, a) = \mathcal{N}(s'; \hat{\mu}_\Theta(s, a), \hat{\sigma}_\Theta^2(s, a))$ where $\hat{\mu}_\Theta(s, a) = \frac{1}{KL} \sum_{k=1}^K \sum_{l=1}^L h_{\theta_{k,l}}(s, a)$ is the predictive mean and $\hat{\sigma}_\Theta^2(s, a) = \frac{1}{KL} \sum_{k=1}^K \sum_{l=1}^L (h_{\theta_{k,l}}(s, a) - \hat{\mu}(s, a))^2$ the epistemic variance.

Model-based control and exploration. When performing model-based RL, we alternate between performing trajectory rollouts on the real environment and updating our dynamics model and policy with the newly collected data. This raises two important questions: *How to formulate and solve the model-based control/policy search problem? How to explore and collect informative transition data?*

A key feature of our BNN dynamics models is their ability to quantify epistemic uncertainty. We harness these uncertainty estimates to perform uncertainty-guided exploration. In particular, we employ *hallucinated upper-confidence reinforcement learning (H-UCRL)* [8] which explores by planning optimistically w.r.t. the dynamics model’s epistemic uncertainty. It hallucinates auxiliary controls $\eta(s, a) \in [-1, 1]^d$ that allow the policy to choose any state transition that is plausible within the (epistemic) confidence regions $[\hat{\mu}(s, a) \pm v\hat{\sigma}(s, a)]$ of the dynamics models. This results in the following optimistic H-UCRL RL objective:

$$\pi^* = \arg \max_{\pi} \max_{\eta} \mathbb{E}_{a_t \sim \pi(a_t | s_t)} \left[\sum_{t=0}^{T-1} r(s_t, a_t) \right] \quad (3)$$

s.t. $s_{t+1} = \hat{\mu}(s_t, a_t) + v\eta(s_t, a_t)\hat{\sigma}(s_t, a_t)$

The objective in (3) steers the policy to areas of the state-action space with high reward and high epistemic uncertainty. By collecting transition data from areas prone to

inaccurate predictions, the BNN dynamics becomes quickly more accurate, which facilitates efficient adaptation to the target task. As we collect more data, the epistemic variance $\hat{\sigma}^2(s, a)$ decreases, which leads to less and less exploration, ultimately ensuring convergence to an optimal policy.

To solve the control/policy-search problem in (3) we propose to use either a model-predictive control (MPC) solver, in particular, the improved cross-entropy method (iCEM) [34], or use policy search on the learned dynamics model. MPC-based approaches often perform favorably in the context of model-based RL since they are more robust w.r.t. model inaccuracies due to the constant re-planning [cf. 4]. However, compared to neural network-based policies, they are also much more computationally expensive at deployment time and tend to exhibit limitations when applied to high-dimensional action spaces. For scenarios with computational constraints or real-time requirements, e.g., on a real robot, we propose to learn an NN policy. Specifically, we employ Soft Actor-Critic (SAC) [35] with the BNN dynamics model as a swap-in for the real environment to learn a policy that optimizes the objective in (3). In this work, we evaluate our method with iCEM and SAC. However, our method is flexible enough to be used alongside most “off-the-shelf” MDP solvers and controllers.

Overview of the Approach. After we have introduced the building blocks, we provide an overview of our model-based meta-RL approach, which we refer to as PACOH-RL. The MPC version of our method is summarized in Alg. 1 and schematically illustrated in Fig. 1. Given the transition data $\mathcal{D}_1, \dots, \mathcal{D}_n$ from previous tasks, we form a particle approximation of the hyper-posterior $\mathcal{Q}(\phi)$ with SVGD, resulting in a set of BNN priors $\{P_{\phi_1}, \dots, P_{\phi_K}\}$. Each prior reflects our meta-learned prior knowledge about the general dynamics of our robot. The differences among the priors reflect the epistemic uncertainty due to the limited number of tasks and samples per task available for meta-learning.

Equipped with the meta-learned priors, we move on to model-based RL on the target task \mathcal{T}^* , which we aim to solve. Since our agent has not yet interacted with \mathcal{T}^* , we start with an empty dataset \mathcal{D}^* of transitions. Then, we iteratively alternate between fitting/updating our BNN dynamics models to the latest transition dataset \mathcal{D}^* and rolling out one episode with our control policy based on the updated dynamics model




Fig. 2: Returns on evaluation tasks averaged over five seeds. We compare PACOH-RL to its greedy counterpart, PACOH-RL (greedy), GrBAL [20], GrBAL-2x, H-UCRL [8], and PETS-DS [4]. For all the environments, PACOH-RL systematically outperforms the baselines in terms of sample efficiency and average return.

\hat{p}_Θ . At the end of each episode, we add the corresponding transition tuples to \mathcal{D}^* and repeat the process. Initially, when \mathcal{D}^* is empty, our BNN dynamics model reflects the meta-learned prior, constituting a much better model starting point than classical model-based RL methods without meta-learning. With every episode, our transition dataset grows in size, allowing the model to quickly adapt to the current dynamics conditions of the target task. As the dynamics model becomes more accurate, the performance of the model-based control policy also improves.

V. EXPERIMENTS

We evaluate PACOH-RL in simulation and on hardware, in particular, a highly dynamic remote-controlled (RC) race car (see Fig. 4). In our experiments, we investigate the following three aspects; (i) Does PACOH-RL improve sample efficiency on the target task?, (ii) does the uncertainty-aware, optimistic RL formulation in (3) improve exploration in environments with sparse rewards, and (iii) does PACOH-RL facilitate successful transfer on real-world hardware systems?

A. Simulation Experiments

For our simulation experiments, we consider the Pendulum, Cartpole, and Half-Cheetah environments from OpenAI gym [36], and a simulated model remote-controlled (RC) car. The RC car simulator is based on realistic race car dynamics used for autonomous racing [37]. We use a time-truncated version of the Half-Cheetah environment where the episode is terminated after 250 timesteps instead of 1000. Since the goal of PACOH-RL is to facilitate transfer and efficient adaptation to new dynamics settings, we vary the physical parameters of the simulation environment in our empirical studies. To showcase generalization in a realistic low-task regime, we sample 20 dynamical settings/tasks at random for the meta-training and 5 for evaluation. For all our experiments, we repeat the task generation and sampling procedure with five seeds and report the mean estimate along with the standard deviation of the achieved returns averaged over the 5 evaluation tasks.

Does PACOH-RL improve the efficiency of RL on the target task? To demonstrate the sample efficiency of

PACOH-RL, we compare it to two model-based RL algorithms; H-UCRL [8] and PETS with distribution sampling (PETS-DS [4]). Furthermore, we also compare PACOH-RL to the gradient-based adaptive learner (GrBAL) algorithm [20], a state-of-the-art model-based meta-RL method that is based on MAML [10] for meta-training. To our knowledge, GrBAL is the only prior model-based meta-RL algorithm that has been successfully implemented on hardware. As an ablation study, we also compare a greedy version of PACOH-RL, which does not use the optimistic planning objective in (3), and, instead, greedily maximizes the RL objective with PETS-DS while being robust w.r.t. the epistemic uncertainty in the dynamics model. We call this variant PACOH-RL (Greedy). To ensure comparability, we use the iCEM-MPC controller for all methods. Fig 2 reports the average returns for all the methods over the course of 25 episodes/trajectories.

We observe that thanks to its meta-learned BNN prior, PACOH-RL starts off with a significantly better policy than the baselines. Importantly, it still is able to improve performance quickly and maintain its advantage over the other methods. Overall, PACOH-RL’s ability to achieve higher rewards with fewer episodes demonstrates the effectiveness of the meta-learned BNN priors towards improving the sample efficiency of model-based RL. Unlike PACOH-RL, we observe that GrBAL [20] often stagnates in performance early on, or only improves slowly as it collects more trajectories. We hypothesize that the observed negative transfer of GrBAL is because MAML overfits the few meta-training tasks. Despite the very limited meta-learning data, PACOH-RL is able to achieve positive transfer which we attribute to the principled meta-level regularization and treatment of epistemic uncertainty of the approach.

Does the uncertainty-aware, optimistic RL formulation in (3) improve exploration in environments with sparse rewards? In Fig. 2, PACOH-RL with the optimistic exploration performs slightly better than the greedy version in the majority of environments. However, the difference between them is small because the environments have dense rewards and, thus, require little exploration.




Fig. 3: Returns after learning on evaluation tasks with sparse rewards for 10 episodes over five different seeds. We compare PACOH-RL to its greedy counterpart, PACOH-RL (greedy), H-UCRL, and its greedy version PETS-DS. In all environments, optimistic planning outperforms its greedy counterpart, with PACOH-RL performing the best.




Fig. 4: High torque motor RC car used in the hardware experiments. As depicted on the right, we have two different tire profiles for both the front and rear wheels. We can also add up to 400g of weight to the front of the car in a cylindrical box encircled in the image.

The authors of [8] show that in environments with sparse reward signals, principled exploration becomes much more crucial. To further investigate the influence of the optimistic planner from (3), we perform experiments on the Pendulum, Cartpole, and Pusher environments with sparse rewards. We compare PACOH-RL to its greedy variant and, as non-meta-learning baselines, report the performance of H-UCRL and its greedy counterpart PETS-DS. We train all agents for 25 episodes and report their average returns over the last ten episodes in Fig. 3. As we can observe, the optimistic, uncertainty-aware planning objective of PACOH-RL considerably improves the agent’s ability to achieve high returns in sparse reward environments. This is similarly true for both H-UCRL and PACOH-RL, and is in line with the results in [8].

B. Hardware Experiments

We use a high-torque motor remote-controlled (RC) car [see 38] for hardware experiments. The car can perform highly dynamic and nonlinear maneuvers such as drifting. The RL problem is to park reverse on a target position that is ca. 2 m away from the start position. This typically involves quickly rotating the car by 180° and then parking in reverse (see Fig. 6 or accompanying video¹). We represent the car with a six-dimensional state: the two-dimensional position and orientation of the car, and the corresponding velocities. The inputs to the car are the steering angle and throttle.

¹https://crl.ethz.ch/videos/pacoh_model_based_rl.mp4




Fig. 5: Trajectories of the RC car obtained under different dynamical settings. Starting at rest, we apply the same control sequence for 50 timesteps at 30 Hz. We repeat the experiment three times for each setting and plot the mean trajectory. The crosses along the trajectory correspond to the car’s mean position at an interval of 10 timesteps. The ellipses correspond to the empirical standard deviation in the car’s position. The first two digits in the legend labels denote the sets of wheels used in the front and rear, respectively, and the third digit denotes the added weight in hectograms.

To simulate different dynamical settings, we consider two different tire profiles with a varying amount of grip for both the front and rear wheels. Furthermore, we change the weight of the car from ca. 1.6kg to 1.8kg, and 2kg by adding weights and operating the car in slow and fast mode. In the slow mode, the motor consumes less power and applies smaller accelerations. In contrast, in the fast mode, the car accelerates considerably faster and performs highly dynamic maneuvers such as drifting. The hardware setup is illustrated in Fig. 4. In total, this gives us twenty-four settings, which result in considerable differences in the dynamic behavior of the car (see Fig. 5). The dynamicity and variability of the different settings make the RC car a compelling platform for applying PACOH-RL.

To collect meta-training datasets, we record trajectories under some of the different settings discussed above. In particular, for meta-training, we take only 5 tasks and use 4 minutes of recorded trajectories per task. After the meta-learning phase of PACOH-RL, we proceed with model-based



Fig. 6: The RC Car in the process of performing a highly dynamic parking maneuver.

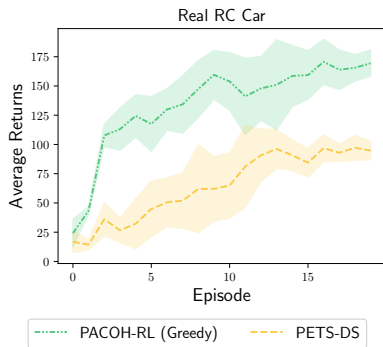


Fig. 7: Returns on the real car averaged over three different seeds. We compare PACOH-RL to its non-meta-learning counterpart. PACOH-RL systematically outperforms the baseline in terms of sample efficiency and average return.

RL on a new, highly dynamical target task. Instead of iCEM, we use SAC since the inference time of a SAC policy is much smaller than MPC, and, thus, can be run in real-time. This means that Line 8 in Alg. 1 is replaced by training a NN policy with SAC [35] on the current BNN dynamics model in a similar fashion to [39] and, then, running the trained policy on the real car to collect one trajectory.

Does PACOH-RL demonstrate sample efficiency on real-world hardware systems? To demonstrate the benefits of meta-learning, we compare PACOH-RL (Greedy) with PETS-DL [4]. We use the greedy versions since the reward signal is dense, and incentivizing additional exploration does not help. We chose one of the most dynamic settings of the RC car as the evaluation task, where we operate the car in fast mode with an added weight of 0.2 kg and the set of tires with the lowest friction. During the RL phase of both methods, we use SAC [35] to train policies which are then deployed on the car. Overall, we ran the experiment for 20 episodes, each consisting of updating the BNN dynamics model with the latest data, SAC training of the policy, and collecting one trajectory on the car. Fig. 7 displays the returns across the 20 episodes, averaged over 3 seeds. We observe that PACOH-RL significantly outperforms the non-meta-learning approach. The PACOH-RL agent achieves high returns within the first few rollouts and converges to an almost optimal policy within less than 2 min (20 episodes) of real-world data. Due to its principled meta-level regularization, PACOH-RL successfully meta-learns a prior specific to RC car dynamics from only 5 meta-learning tasks. Equipped with this prior, the BNN dynamic model adapts quickly to the target task, resulting in substantial efficiency improvements compared to PETS.

VI. CONCLUSION

We have proposed PACOH-RL, a novel approach to model-based Meta-RL. PACOH-RL meta-learns a dynamics model prior from previous experience on the robotic platform, and, thereby, facilitates efficient adaptation under new dynamical conditions. The key characteristics of our method are its principled regularization and its holistic treatment of epistemic uncertainty, which guides exploration and data collection during the task adaptation stage. This allows PACOH-RL to achieve positive transfer and efficient task adaptation, even with only a handful of meta-learning datasets. Hence, PACOH-RL is one of the first Meta-RL approaches that are applicable to real robotic hardware where data is scarce.

We have focused on harnessing the epistemic uncertainty quantification of our approach towards exploration. Many other problems, such as RL under safety constraints [40, 41] or without rewards [42, 43], and off-policy evaluation [44, 45] also benefit from good uncertainty estimates. We leave extending our approach to these settings as future work.

REFERENCES

- [1] V. Mnih, K. Kavukcuoglu, D. Silver *et al.*, “Human-level control through deep reinforcement learning,” *nature*, 2015.
- [2] A. Rajeswaran, V. Kumar, A. Gupta *et al.*, “Learning complex dexterous manipulation with deep reinforcement learning and demonstrations,” *arXiv preprint arXiv:1709.10087*, 2017.
- [3] M. P. Deisenroth and C. E. Rasmussen, “Pilco: A model-based and data-efficient approach to policy search,” in *ICML*, 2011.
- [4] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models,” in *NeurIPS*, 2018.
- [5] I. Clavera, J. Rothfuss, J. Schulman, Y. Fujita, T. Asfour, and P. Abbeel, “Model-based reinforcement learning via meta-policy optimization,” in *CoRL*, 2018.
- [6] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, “RI2: Fast reinforcement learning via slow reinforcement learning,” in *ICLR*, 2017.
- [7] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, “Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning,” in *ICRA*, 2018.

- [8] S. Curi, F. Berkenkamp, and A. Krause, “Efficient model-based reinforcement learning through optimistic policy search and planning,” *NeurIPS*, 2020.
- [9] S. Thrun and L. Pratt, *Learning to Learn: Introduction and Overview*. Springer US, 1998.
- [10] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *ICML*, 2017.
- [11] K. L. Pavasovic, J. Rothfuss, and A. Krause, “Mars: Meta-learning as score matching in the function space,” in *ICLR*, 2022.
- [12] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick, “Learning to reinforcement learn,” *arXiv preprint arXiv:1611.05763*, 2016.
- [13] J. Rothfuss, D. Lee, I. Clavera, T. Asfour, and P. Abbeel, “ProMP: Proximal Meta-Policy Search,” in *ICLR*, 2019.
- [14] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen, “Efficient off-policy meta-reinforcement learning via probabilistic context variables,” in *ICML*, 2019.
- [15] F. Luo, S. Jiang, Y. Yu, Z. Zhang, and Y.-F. Zhang, “Adapt to environment sudden changes by learning a context sensitive policy,” in *AAAI*, 2022.
- [16] S. Sæmundsson, K. Hofmann, and M. P. Deisenroth, “Meta reinforcement learning with latent variable gaussian processes,” in *UAI*, 2018.
- [17] C. Perez, F. P. Such, and T. Karaletsos, “Generalized hidden parameter MDPs: Transferable model-based RL in a handful of trials,” in *AAAI*, 2020.
- [18] K. Lee, Y. Seo, S. Lee, H. Lee, and J. Shin, “Context-aware dynamics model for generalization in model-based reinforcement learning,” in *ICML*, 2020.
- [19] T. Hiraoka, T. Imagawa, V. Tangkaratt, T. Osa, T. Onishi, and Y. Tsuruoka, “Meta-model-based meta-policy optimization,” in *ACML*, 2021.
- [20] I. Clavera, A. Nagabandi, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, “Learning to adapt in dynamic, real-world environments through meta-reinforcement learning,” in *ICLR*, 2019.
- [21] A. Pentina and C. Lampert, “A PAC-Bayesian bound for lifelong learning,” in *ICML*, 2014.
- [22] R. Amit and R. Meir, “Meta-learning by adjusting priors based on extended PAC-Bayes theory,” in *ICML*, 2018.
- [23] J. Rothfuss, V. Fortuin, M. Josifoski, and A. Krause, “PACOH: Bayes-optimal meta-learning with PAC-guarantees,” in *ICML*, 2021.
- [24] J. Rothfuss, C. Koenig, A. Rupenyan, and A. Krause, “Meta-learning priors for safe Bayesian optimization,” in *CoRL*, 2023.
- [25] N. Heess, G. Wayne, D. Silver, T. Lillicrap, T. Erez, and Y. Tassa, “Learning continuous control policies by stochastic value gradients,” in *NeurIPS*, 2015.
- [26] R. Houthoofd, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel, “Curiosity-driven exploration in deep reinforcement learning via bayesian neural networks,” in *NeurIPS*, 2016.
- [27] P. Grünwald, “The safe Bayesian: learning the learning rate via the mixability gap,” in *ALT*, 2012.
- [28] B. Guedj, “A primer on pac-bayesian learning,” *arXiv preprint arXiv:1901.05353*, 2019.
- [29] M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *ICML*, 2011.
- [30] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *JASA*, 2017.
- [31] Q. Liu and D. Wang, “Stein variational gradient descent: A general purpose bayesian inference algorithm,” in *NeurIPS*, 2016.
- [32] C. Chen, R. Zhang, W. Wang, B. Li, and L. Chen, “A unified particle-optimization framework for scalable bayesian sampling,” in *UAI*, 2018.
- [33] S. H. Lim, H. Xu, and S. Mannor, “Reinforcement learning in robust markov decision processes,” *NeurIPS*, 2013.
- [34] C. Pinneri, S. Sawant, S. Blaes, J. Achterhold, J. Stueckler, M. Rolinek, and G. Martius, “Sample-efficient cross-entropy method for real-time planning,” in *CORL*, 2020.
- [35] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *ICML*, 2018.
- [36] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.
- [37] J. Kabzan, M. I. Valls, V. J. Reijgwart, H. F. Hendrikx, C. Ehmke, M. Prajapat, A. Bühler, N. Gosala, M. Gupta *et al.*, “Amz driverless: The full autonomous racing system,” *Journal of Field Robotics*, 2020.
- [38] B. Sukhija, N. Köhler, M. Zamora, S. Zimmermann, S. Curi, A. Krause, and S. Coros, “Gradient-based trajectory optimization with learned dynamics,” in *ICRA*, 2023.
- [39] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” in *NeurIPS*, 2019.
- [40] J. Garcia and F. Fernández, “A comprehensive survey on safe reinforcement learning,” *JMLR*, no. 1, 2015.
- [41] Y. As, I. Usmanova, S. Curi, and A. Krause, “Constrained policy optimization via bayesian world models,” in *ICLR*, 2022.
- [42] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak, “Planning to explore via self-supervised world models,” in *ICLR*, 2020.
- [43] B. Sukhija, L. Treven, C. Sankar, S. Blaes, S. Coros, and A. Krause, “Optimistic active exploration of dynamical systems,” *NeurIPS*, 2024.
- [44] J. Rothfuss, B. Sukhija, T. Birchler, P. Kassraie, and A. Krause, “Hallucinated adversarial control for conservative offline policy evaluation,” in *UAI*, 2023.
- [45] J. P. Hanna, P. Stone, and S. Niekum, “Bootstrapping with models: Confidence intervals for off-policy evaluation,” in *AAAI*, 2017.