THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Meta-Learning for Fast Adaptive Locomotion with Uncertainties in Environments and Robot Dynamics

Citation for published version:

Anne, T, Wilkinson, J & Li, Z 2021, Meta-Learning for Fast Adaptive Locomotion with Uncertainties in Environments and Robot Dynamics. in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Institute of Electrical and Electronics Engineers (IEEE), pp. 4568-4575, 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems , Prague, Czech Republic, 27/09/21. <https://doi.org/10.1109/IROS51168.2021.9635840>

Digital Object Identifier (DOI):

[10.1109/IROS51168.2021.9635840](https://doi.org/10.1109/IROS51168.2021.9635840)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2021 IEEE/RSJ International Conference on Intelligent Robots and Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Meta-Learning for Fast Adaptive Locomotion with Uncertainties in Environments and Robot Dynamics

Timothée Anne¹, Jack Wilkinson² and Zhibin Li²

Abstract—This work developed meta-learning control policies to achieve fast online adaptation to different changing conditions, which generate diverse and robust locomotion. The proposed method updates the interaction model constantly, samples feasible sequences of actions of estimated state-action trajectories, and then applies the optimal actions to maximize the reward. To achieve online model adaptation, our proposed method learns different latent vectors of each training condition, which is selected online based on newly collected data from the *past 10 samples within 0.2s*. Our work designs appropriate state space and reward functions, and optimizes feasible actions in an MPC fashion which are sampled directly in the joint space with constraints, hence requiring no prior design or training of specific gaits. We further demonstrated the robot’s capability of detecting unexpected changes during the interaction and adapting the control policy in *less than 0.2s*. The extensive validation on the SpotMicro robot in a physics simulation shows adaptive and robust locomotion skills under changing ground friction, external pushes, and different robot dynamics including motor failures and the whole leg amputation.

I. INTRODUCTION

Responsive control policies adapting to unforeseen environments are crucial to task success in robot motor control. These changes and unexpected situations can be intrinsic or extrinsic, such as robot damage, motor failures, changing friction, and external force disturbances. For robot locomotion, traditional approaches of planning and control require expert knowledge and accurate dynamics models and constraints of both the robot and the environment [1], [2], which are all subject to unforeseeable changes that are difficult to know beforehand. Moreover, even using data-efficient learning techniques such as Bayesian optimization to tune decision variables and control parameters, it can only achieve adaptation on a trial-by-trial basis [3] and also require extensive computation [4] which is not able to respond to changes on the fly.

Recent advances in Reinforcement-Learning (RL) lead to algorithms achieving human-like or animal-level performance in a range of difficult control tasks. Model-free RL can perform global search of control parameters and obtain globally optimal gaits while combined with walking pattern generation [5]. Also, an RL-based feedback policy can achieve human-like bipedal walking by imitating human motion capture data [6]. With simulation training including

Fig. 1: Adaptive and robust locomotion against uncertainties: (a) external force, (b) slippery ground, (c) jammed faulty motors, and (d) a complete leg amputation.

actuator properties, a model-free RL scheme can train different locomotion policies separately and deploy on a real quadrupedal robot [7]. By using a multi-expert learning, an hierarchical RL architecture can learn to fuse multiple motor skills and generate multimodal locomotion coherently on a real quadruped [8]. However, in general, model-free RL algorithms have limited sample efficiency, resulting in long training time to produce viable policies. For example, it took a model-free RL algorithm 83 hours to achieve human performance on the Atari game suite, compared to 15 minutes for a human [9]. Similarly, AlphaStar [10] used 200 years of equivalent real-time to reach expert human performance playing Starcraft II.

On the other hand, model-based approaches can achieve comparatively high performance while being more sample-efficient by several orders of magnitude, converging faster than model-free approaches for locomotion tasks [11]. To deploy robots in the real world, online adaptation to changes in the environment is required, since not all conditions can be considered by pre-trained policies, such as drastic changes of environments or robots by amputation. Hence, meta-learning, or *learning to learn* is a novel and promising approach for solving such generic adaptations. Model-based meta-RL has been used in the real world to adapt the control of a six-leg millirobot to different floor conditions [12]. A model-based meta-RL algorithm, FAMLE, was used in the real-world on a Minitaur quadruped, where a latent black-box context vector

¹Timothée Anne is with ENS Rennes, France; ²Jack Wilkinson and Zhibin Li are with the School of Informatics, the University of Edinburgh, UK. timothee.anne@ens-rennes.fr, jack.wilkinson, zhibin.li@ed.ac.uk. ©2021 IEEE. This is the author version. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses.

encoded different environment conditions [13].

Our proposed method has made new improvements that require no prior knowledge of specific gaits. For example, FAMLE relies on sinusoidal gaits and therefore needs to optimize the amplitudes and phases of sinusoidal patterns by the model predictive control at a low frequency of 0.5Hz. In our work, we also use random sampling as model predictive control using the learned forward model of the dynamics to estimate the samples rewards. Compared to their work, we directly sample in the joint space at a much higher frequency at 50Hz, and we further improve the sampling process by specifying constraints on velocity, acceleration and jerk of the desired joint trajectories. Our study extensively validated the capability of adaptation in simulated test scenarios with large variations in floor friction, external forces or unexpected damage to joints.

Based on the interaction model, our method allows changing the reward function online and therefore is able to modify the behavior of the robot. For example, the learned controller can track a variable forward velocity, even it has been trained on a fixed desired velocity. Likelihood estimation with condition latent vectors allows the meta-model to adapt to already seen conditions. Meta-training should allow "on the fly" optimization to better adapt to the current unknown condition.

In this paper, we present an improved model-based meta-RL approach to quadruped locomotion that is capable of online adaption to changing environments, as shown in Fig. 1. The main contributions of this work are:

- 1) The proposed algorithm is capable of learning from scratch and requires no prior knowledge of the type of gaits, such as periodic phases of leg movements.
- 2) Our methods introduces and applies hard constraints of velocity, acceleration and jerk on the sampled actions during the search process.
- 3) The capability and robustness of online adaption to changes in both the robot and the environment, such as external force disturbances, varying frictions, faulty motors and leg amputation.

The remainder of the paper is organized as follows. We outline the background in Section II and related work in Section III. In Section IV, we elaborate the methodology and technical details on the model-based RL algorithm and the improvements by meta-learning. Section V presents extensive simulation validations, results and analysis. Finally, we conclude and suggest future work in Section VI.

II. BACKGROUND

This section presents the preliminaries of RL, Model Predictive Control (MPC) and meta-learning.

A. Reinforcement Learning

In reinforcement learning, the agent learns to solve a task in an unknown environment E , defined by a Markov decision process (S, s_0, A, f_E, r) where S is the set of continuous states of the environment E , s_0 the initial state, A the set of continuous actions the agent can perform in the environment,

$f_E : S \times A \times S \rightarrow \mathbb{R}$ the probabilistic transition function and $r : S \times A \rightarrow \mathbb{R}$ the reward function.

The goal of the agent is to learn a policy $\Pi_\theta : S \rightarrow A$, parameterized by θ , which decides which action to perform given the current state s_t to *maximize the long-term reward* $R(t) = \sum_{i=t}^{t+H} \gamma^{i-t} r(s_i, \Pi(s_i))$, where H is the horizon and γ is the discount factor.

Model-free RL focuses in directly learning such a policy, whereas model-based RL focuses on learning a model of the transition function – the transition of states given the current state and actions – which can be used to train the policy with fictive transitions or with model predictive control.

B. Model Predictive Control

Given the current state s_t and a horizon H , Model Predictive Control (MPC) uses a forward model of dynamics to select an action sequence $a_{t:t+H}$ which maximizes the predicted cumulative reward $R(t)$. The agent performs the first action a_t from the action sequence and collects the resulting state s_{t+1} . The MPC then repeats such optimization and allows the agent to alleviate the possible error in the model prediction. Compared to model-free RL, we can change the reward *online* to control the agent's behavior using model-based RL in an MPC fashion.

C. Meta-Learning

We use meta-learning to train an agent to solve several tasks, where the neural network learns to adapt to several varying conditions, such as different floor frictions, the presence of external disturbances or having a damaged motor. For the neural network model, the initial set of weights θ^* must be found, such that only a small number of gradient descent steps with little collected data in a unknown environment can produce effective adaptations.

III. RELATED WORK

A. Model-Free Deep Reinforcement Learning

Proximal Policy Optimization was used to train a model-free controller first in simulation and then applied on the real robot with walking gaits [14]. Soft Actor-Critic (SAC) has also been used to train on the real Minotaur robot within two hours [15]. This limitation of sampling-efficiency motivates us to focus on model-based RL.

B. Model-Based Deep Reinforcement Learning

There are three types of model-based RL: learning to predict the expected return from a starting state distribution, for example using Bayesian Optimization [16]; learning to predict the outcome from a given starting state and given policy [17]; and learning to model the transition function using a forward dynamical model. Here, we use the third type of model.

Forward models of dynamics are either deterministic or probabilistic, where deterministic models can be linear models [18] or neural networks [19], and probabilistic models estimate uncertainty for modeling stochastic environments

Fig. 2: Schematic diagram of the proposed model-based meta-RL control framework for adaptive legged locomotion.

or estimating the long-term prediction uncertainty. Gaussian Processes [20], [21] or Bayesian neural networks [22] can be used to scale the abilities of Gaussian Processes models to higher dimensional environments.

For locomotion tasks, model-based RL with a forward model can have the same performance as model-free methods, while requiring at least an order of magnitude less samples [11]. An ensemble of feed-forward neural networks is used to model the forward dynamics of the environment with uncertainty estimation. MPC uses this uncertainty estimation to formulate a more robust control which alleviates early overfitting model-based RL. The same method has been used with meta-learning to adapt the control of a 6-leg real millirobot to different floors [12].

C. MPC and Meta-Learning

Several optimization methods have been used for model-based RL, for example, Model Predictive Path Integral [23], random shooting [15] or Cross-entropy method [24], [11]. We use random shooting for the simplicity, easy parallelism and proven performances on real robots [12].

There are two main methods: a meta-learner model outputs the set of initial weights θ^* of the learner [25]; or θ^* is optimized using a meta-loss, e.g., by gradient-descent [26]. Gaussian processes have been used in [27] but only for low dimension environments. For model-based RL, gradient based meta-learning was shown to be more data-efficient, resulting in a better and faster adaptation [12]. Hence, we use gradient-based meta-learning.

IV. METHODOLOGY

This section presents details of the model-based RL algorithm as discussed in Section IV-A and meta-learning algorithm Section IV-B. We highlight our improvements which results in new robot capability of robust and versatile walking without a predefined, parameterized gait.

A. Model-Based Reinforcement-Learning algorithm

The model-based RL algorithm runs at 50Hz, sending desired actions to PD controllers running at 250Hz to generate torques for physics simulation. The top level model-based RL runs at a lower frequency due to computation efficiency and because we empirically found that shorter timesteps decrease the MPC efficiency. The algorithm is composed of two main parts: the forward dynamics model, and MPC. Fig. 2 illustrates the schematics of the control framework.

1) *The Forward Model of Dynamics:* We use a fully-connected feed-forward neural network, with two hidden layers of 256 units using a ReLU (Rectified Linear Unit) activation function. It takes the concatenation of the current state and action (s_t, a_t) as input, and learns to predict the difference in the resulting state: $\Delta s_t = s_{t+1} - s_t$, which is a standard means to get the prediction $s_{t+1} = s_t + \Delta s_t$. The state is composed of the angular joints positions and velocities, the base orientation angles and angular rates, and the linear base velocities. The action corresponds to the desired angular joints positions.

The model parameter θ is the set of weights of the connections between the units. It is optimized using the gradient-based optimizer Adam [28] on a dataset of triplet $D \in S \times A \times S$ using mean squared error as loss function. We depict details of the model-based RL algorithm in Algorithm 1.

Compared to the work in [13] for the Minitaur, the addition of the angular orientation angles and angular rates of the base is the key of our success for controlling the robot at 50Hz. In contrast, only the linear base velocities in the horizontal plane were used in [13] to control the Minitaur gait parameters at a much lower frequency of 0.5Hz.

2) *Model Predictive Control:* The method of random shooting is implemented which is suitable for parallel computing, and the algorithm is detailed in Algorithm 2. At each time step, N_{pop} action sequences of length H are sampled. Each sequence is evaluated starting with the current state, using the model to estimate the corresponding state trajec-

Algorithm 1 Model-based reinforcement learning algorithm
Input: Environment E.

```

D = Ninit episodes in E with a random controller
θ = random weights
for Ntraining episodes do
  Train Mθ on D using Adam
  s0 = E.reset()
  for t = 0 to T − 1 steps do
    at = MPC(Mθ, st, at−3:t−1)
    st+1 = E.step(at)
    D = D ∪ {(st, at, st+1)}
  end for
end for

```

tory. From these trajectories, long term reward is computed and the action with the highest estimated reward is selected.

The reward is given by:

$$r(s, a) = e^{-|v_x - v_x^{desired}|} + e^{-|v_y|} + e^{-|v_z|} + e^{-|roll|} + e^{-|pitch|} + e^{-|yaw|}$$

One interesting feature of model-based RL compared to model-free is that at test time we can change the reward function to control the agent’s behavior. For example, we can change the value of $v_x^{desired}$ during an episode to change the agent’s tracked velocity, which results in actual change of forward velocity (Section V-A) although it has been trained with a fixed value of $v_x^{desired}$.

Real actuators have inherent limitations in velocity, acceleration and jerk. Instead of uniformly sampling desired joint angles within the limits, continuity constraints are used, where each desired joint state of the sequence is sampled using previous joint positions to ensure velocities, accelerations and jerks are smooth and bellow their respective limits.

As the improvement to the previous work [13], we enforced physical constraints during sampling of actions: $q_{min} \leq q \leq q_{max}$, $|A| \leq A_{max}$, $|V| \leq V_{max}$ and $|J| \leq J_{max}$, where q, V, A and J are the desired joint angle, velocity, acceleration and jerk, respectively. The limits of velocity, acceleration and jerk are the soft constraints for the smoothness and continuity of actions. For safety reasons, regarding the joint position limits, we further imposed hard constraints of sampled actions on q to avoid hitting the physical limit of joint movements.

This improvements on sampling enforces the MPC in a more suitable subspace. During training, it increased the distance traveled compared to the default condition during a 10s episode by an order of 2: from $1.64 \pm 0.12m$ (without), to $3.16 \pm 0.12m$ (with), where p-value ≤ 0.001 on 20 episodes. It also reduces the observed jerk by an order of 5: from $1.3 \times 10^4 \pm 260 \text{ rad/s}^3$ (without) to $2.4 \times 10^3 \pm 60 \text{ rad/s}^3$ (with), where p-value ≤ 0.001 on 20 episodes.

B. Meta-learning algorithm

Before meta-training, an expert is trained for each training condition using the proposed model-based RL algorithm to

Algorithm 2 MPC algorithm

Inputs: A model M_θ, initial state s₀, past actions a_{−3:−1}, horizon H and discount factor γ.

```

Sample a1:H1:Npop using a−3:−1 for continuity
for i = 1 to N do
  Ri = 0
end for
for t = 1 to H steps do
  for all i = 1 to Npop samples do
    st+1i = Mθ(ati, st)
    Ri = Ri + γt−1r(sti, ati)
  end for
end for
i* = argmaxi Ri
return a1i*

```

Algorithm 3 Meta-learning training, called once before adaptation

Inputs: N datasets D_{0:N−1} from N different conditions.

```

θ* = random weights
for i = 0 to N − 1 do
  Ci = random vector
end for
for nouter = 0 to Nouter − 1 do
  i = nouter mod N
  θ = θ*
  for ninner = 1 to Ninner do
    L =  $\frac{1}{|D_i|} \sum_{(s,a,s') \in D_i} (M_\theta(s,a,C_i) - s')^2$ 
    θ = θ − α∇θL
    Ci = Ci − α∇CiL
  end for
  θ* = θ* + (1 −  $\frac{n_{outer}}{N_{outer}}$ )β(θ − θ*)
end for
return θ*

```

collect its training data. To adapt the model to each condition i, a specific latent vector C_i is optimized during meta-learning using the regression loss on the data of the corresponding condition. This vector of fixed dimensions is then given to the input layer, alongside the current state and action when the condition i is selected. We use a first order meta-learning called Reptile [29], which is composed of two phases: meta-training (Algorithm 3) and meta-adaptation (Algorithm 4).

1) *Meta-training*: The initial set of weights θ* and each condition latent vector C_i, of dimension d_c, are optimized for adaptation. Meta-training is separated into two nested loops. In the inner loop, one training dataset D_i and its corresponding condition latent vector C_i are selected. The model weights θ are initialized to θ* and Adam [28] optimize both of them for the regression loss of the current dataset D_i.

In the outer loop, θ* is optimized by taking a small step, with a linearly decreasing schedule, towards the optimized weights of the inner-loop. This allows θ* to converge to a

Algorithm 4 Meta-learning adaptation, called at each step
Inputs: a meta-trained set of weights θ^* , a list of learned N condition latent vectors $\mathcal{C}_{1:N}$, a dataset $\mathcal{D}$ of the past K steps.

```

for  $i = 1$  to  $N$  do
   $\mathcal{L}_i = \frac{1}{|\mathcal{D}|} \sum_{(s,a,s') \in \mathcal{D}} (M_\theta^*(s,a,C_i) - s')^2$ 
end for
 $i^* = \operatorname{argmin}_i \mathcal{L}_i$ 
 $\theta = \theta^*_{i^*}$ 
for  $n_{inner} = 1$  to  $N_{inner}$  do
   $\mathcal{L} = \frac{1}{|\mathcal{D}|} \sum_{(s,a,s') \in \mathcal{D}} (M_\theta(s,a,C_{i^*}) - s')^2$ 
   $\theta = \theta - \alpha \nabla_\theta \mathcal{L}$ 
   $C_{i^*} = C_{i^*} - \alpha \nabla_{C_{i^*}} \mathcal{L}$ 
end for
return  $\theta, C_{i^*}$ 

```

nearby point (in the euclidean sense) to the optimal set of weights of each training condition. We detail the algorithm in Algorithm 3. For all experiments, we use the same hyper-parameters values as FAMLE [13].

2) *Meta-adaptation*: At each time step, we select the most likely training condition using the previous K time steps, each condition latent vector and the set of weights θ^* . We then optimize the corresponding latent vector and the set of model weights, starting from θ^* , using the same optimization procedure as the inner loop but with the past K steps. We detail the algorithm in Algorithm 4.

After the set of weights and the condition latent vector are optimized for the current condition, we use the MPC to select the optimal action to apply, and then new state information is collected and the whole meta-adaptation iterates. In this study, using only past 10 samples, i.e. $K = 10$ corresponding to 0.2s, the agent can detect any changes in the condition and adapt accordingly in less than 0.2s.

C. Limitations

Apart from the standard classical robot control of tuning PD gains and joints limits, the proposed method still requires fine-tuning of reward function, model architecture, hyper-parameters for the meta-learning and the adaptation. The use of MPC instead of a neural network policy has a trade-off between real-time computation and performance, i.e. MPC performs better in terms of adaption but requires more computation needed from the sampling procedure.

V. RESULTS

We used a custom version of the robot model (adapted from the open source SpotMicro robot [30]) in PyBullet simulation to validate our method. Here, we first present the learning capability of the model-based RL algorithm on SpotMicro with a first adaption to a sequence of different conditions (Section V-A). We further validate the adaptation capability of the proposed meta-learning under various fixed frictions (Section V-B) and time-varying, decreasing friction (Section V-C).

Fig. 3: Walking in presence of large uncertainties: (a) constant external force disturbance, (b) low-friction slippery ground, and (c) with one leg amputated.

Fig. 4: Walking with a fixed joint position (to emulated actuator failures): (a) rear-right hip, (b) rear-right knee, (c) front-right hip and (d) front-right knee.

A. Overview

We trained the expert model with a default condition of friction $\mu = 0.8$, and this resulted default controller for walking is robust to perturbations, withstanding several pushes of 10N for 0.2s. After 300 of 10s-episodes which produced training data in the given condition, the quadruped was able to walk on slippery ground (friction coefficient $\mu = 0.2$), against external forces, with a fixed joint (to emulated damaged motor) or a missing/amputated leg. Tab. I shows the comparison between experts and meta-trained models under different conditions. Snapshots of the walking gaits are shown in Fig. 3 and in Fig. 4.

Fig. 5: Measured joints trajectories from a 20s test scenario where the robot/environment changed every 5s: default friction $\mu = 0.8$ (green), blocked right hip pitch joint (pink), slippery ground (yellow) and constant external push (brown).

Expert\Condition	Default	Slippery	Lateral Force	Damaged Motor
Default	100%, 3.2 ± 0.2	20%, 0.7 ± 0.7	40%, 2.0 ± 1.0	100%, 0.7 ± 0.2
Slippery	30%, 1.3 ± 0.5	90%, 2.4 ± 0.4	10%, 0.7 ± 0.4	100%, 0.4 ± 0.1
Lateral Force	0%	0%	70%, 2.3 ± 0.7	90%, 0.4 ± 0.1
Damaged Motor	90%, 0.5 ± 0.2	70%, 0.3 ± 0.2	30%, 0.6 ± 0.3	100%, 2.5 ± 0.1
Meta-Trained	70%, 3.0 ± 0.6	80%, 2.4 ± 0.7	40%, 2.0 ± 1.3	100%, 2.7 ± 0.1

TABLE I: Success rate and average distance of travel for 10 episodes from different expert and the meta-trained model.

Fig. 6: Walking with a continuously changing velocity from 0.5m/s to -0.2m/s: (a) forward, (b) static, and (c) backward.

Using the proposed meta-learning method, the agent was able to adapt to four different conditions: default, fixed front-right hip motor, slippery ground and external forces. It traveled an average distance of 5.1 ± 1.5 m compared to a default expert which traveled 2.6 ± 0.8 m (averaged over 20 episodes). The joint trajectories from these test scenarios are shown in Fig. 5.

By changing the reward function online, we can command different desired velocities continuously as shown in Fig. 6 and Fig. 7. Moreover, the trained expert controller can also generate continuous control actions to track discrete, discontinuous commanded velocities (see Fig. 8). More details of

Fig. 7: Online tracking of continuous and variable walking velocity using the model trained only at the fixed velocity of 0.5m/s, while the desired velocity in the reward function changes from 0.5m/s to -0.2m/s during the test scenario.

walking performance can be seen in [the video here](#).

B. Ground with Constant Friction

We evaluated the adaptation capability using the meta-trained model and compared it to experts over the full range of different frictions (0.1 to 0.8 with 0.05 increments). We first trained 5 sets of experts for frictions 0.2, 0.4 and 0.6, using 300 10s-episodes, i.e., 50 minutes of data. Then we meta-trained 5 meta-models to adapt to these 3 frictions, each using one set of experts data, with the purpose that they could adapt to the full range afterwards.

Each set of these 5 models was evaluated for each friction with 4 10s-episodes, so this gives 20 evaluations per expert and 20 evaluations for the meta-learning. The meta-

Fig. 8: Online tracking of discrete and variable walking velocity using the model trained only at the fixed velocity of 0.5m/s, while every 2 s the desired velocity is randomly sampled within 0.5m/s and -0.2m/s during the test scenario.

Fig. 9: Distance traveled for the full range of frictions from 3 expert models and the meta-trained model.

trained models outperformed the experts on the full range of frictions (see Fig. 9). As expected, each expert had the best performance when the friction constant is around its trained value.

C. Ground with Decreasing Friction

As a comparison, we used the same experts and meta-model and benchmarked their adaptation capability on a ground with continuously decreasing friction. We evaluated each set of models with 4 10s-episodes where friction coefficient μ started at 0.8 and linearly decreased to 0.1. This gave 20 evaluations per expert and 20 evaluations for meta-learning.

The meta-trained models demonstrated better walking performance and traversed farther (3.38m) than the experts (3.07m), using a t-test with a p-value under 0.05, see Fig. 10. In Fig. 11, the curves and shaded areas are the means and standard deviations of the velocity, respectively.

Additionally, Fig. 12 depicts the estimated condition at each time step using the past 0.1s (i.e. 5 time steps). At the beginning of the episode, when friction was higher, the model estimated a friction of 0.6 to be more likely (from 0.8 down to 0.5, i.e., 0-4.5s), then switched to a friction of 0.4 (from

Fig. 10: Distance traveled with decreasing friction from the 3 expert models and the meta-trained model (p-value: * < 0.05, ** < 0.01 and *** < 0.001).

Fig. 11: Forward walking velocity with decreasing friction from 3 expert models and the meta-trained model.

0.5 down to 0.3, i.e., 4.5-7s), and finished by estimating a more likely friction of 0.2 (from 0.3 down to 0.1, i.e., 7-10s).

D. Discussion

The results validate the efficiency and effectiveness of the meta-learning method to detect the most probable current condition and adapt accordingly to different ground friction coefficients. The online adaptive walking using the proposed meta-learning outperformed the specialized experts which were specifically trained on specific frictions. This demonstrated the capability of meta-learning to incorporate knowledge from all training data.

We have also pushed the extreme test case in terms of unseen hardware failures. We specifically designed a case for meta-training where a motor of one leg of the quadruped was blocked at a fixed joint position (emulated actuator failures), and tested if meta-learning can adapt to the damage from a different leg faster than learning from scratch. Our investigation showed that the meta-adaptation is not able to adapt to such changes on a different leg, and we hypothesize that a second-order meta-learning algorithm, e.g. Model-Agnostic Meta-Learning [26], may be a better solution for such an extreme case.

Fig. 12: Likelihood of each pre-trained condition using online meta-adaptation in presence of decreasing friction.

VI. CONCLUSION AND FUTURE WORK

Based on the past work of model-based meta-RL [13], we have made contributions to improve the algorithm for adaptive and robust quadruped locomotion in changing robot dynamics (motor failure and amputation) and varying environmental constraints (time-varying friction, external pushes). In physics-based simulation, we have demonstrated this method can learn quadrupedal walking without using a periodic gait signal [13] or a phase vector [6]. Instead, by updating an interaction model of the robot and environment and applying the optimal control actions, a walking gait is naturally generated as the outcome of maximizing the task reward. We further validated the capability of our proposed framework in adapting to different conditions such as robot damage, changing friction, and external force disturbances.

Future work will apply this method on the real SpotMicro robot, and identify potential issues of sim2real transfer which will be addressed by new solutions for meta-model and more effective search of model predictive control procedure. We hypothesize the current meta-learning algorithm is efficient enough for multi-task learning, which can be further studied. Also, a second-order meta-learning algorithm [26] can be Incorporated could potentially achieve better adaptation to novel situations.

VII. ACKNOWLEDGEMENT

This work has been supported by EPSRC UK Robotics and Artificial Intelligence Hub for Offshore Energy Asset Integrity Management (EP/R026173/1).

APPENDIX

TABLE II: Parameters used in all experiments.

Episode duration: 10s	$N_{\text{training}}: 275$	$\alpha: 0.0001$
H: 25 (=0.5s)	$\gamma: 1.0$	$\beta: 0.3$
$N_{\text{pop}}: 1000$	$N_{\text{outer}}: 5000$	K: 10 (=0.2s)
$N_{\text{init}}: 25$	$N_{\text{inner}}: 10$	$d_c: 10$

REFERENCES

- [1] P. Fankhauser *et al.*, “Robust rough-terrain locomotion with a quadrupedal robot,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8.
- [2] I. Chatzinikolaïdis *et al.*, “Contact-implicit trajectory optimization using an analytically solvable contact model for locomotion on variable ground,” *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 6357–6364, 2020.
- [3] A. Rai *et al.*, “Bayesian optimization using domain knowledge on the atris biped,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1771–1778.
- [4] K. Yuan *et al.*, “Bayesian optimization for whole-body control of high-degree-of-freedom robots through reduction of dimensionality,” *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2268–2275, 2019.
- [5] H. Dallali *et al.*, “On global optimization of walking gaits for the compliant humanoid robot, coman using reinforcement learning,” *Cybernetics and Information Technologies*, vol. 12, no. 3, pp. 39–52, 2012.
- [6] C. Yang *et al.*, “Learning natural locomotion behaviors for humanoid robots using human bias,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2610–2617, 2020.
- [7] J. Hwangbo *et al.*, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, 2019.
- [8] C. Yang *et al.*, “Multi-expert learning of adaptive legged locomotion,” *Science Robotics*, vol. 5, no. 49, 2020.
- [9] M. Hessel *et al.*, “Rainbow: Combining improvements in deep reinforcement learning,” in *AAAI-18*, 2018, pp. 3215–3222.
- [10] O. Vinyals *et al.*, “AlphaStar: Mastering the Real-Time Strategy Game StarCraft II,” 2019.
- [11] K. Chua *et al.*, “Deep reinforcement learning in a handful of trials using probabilistic dynamics models,” in *NeurIPS 2018*, 2018, pp. 4759–4770.
- [12] A. Nagabandi *et al.*, “Learning to adapt in dynamic, real-world environments through meta-reinforcement learning,” in *ICLR 2019*, 2019.
- [13] R. Kaushik *et al.*, “Fast online adaptation in robotics through meta-learning embeddings of simulated priors,” 2020.
- [14] J. Tan *et al.*, “Sim-to-real: Learning agile locomotion for quadruped robots,” in *Robotics: Science and Systems XIV*, 2018.
- [15] T. Haarnoja *et al.*, “Learning to walk via deep reinforcement learning,” in *Robotics: Science and Systems XV*, 2019.
- [16] E. Brochu *et al.*, “A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning,” *CoRR*, vol. abs/1012.2599, 2010.
- [17] S. Forestier *et al.*, “Intrinsically motivated goal exploration processes with automatic curriculum learning,” *CoRR*, vol. abs/1708.02190, 2017.
- [18] S. Levine *et al.*, “End-to-end training of deep visuomotor policies,” *J. Mach. Learn. Res.*, vol. 17, pp. 39:1–39:40, 2016.
- [19] D. Ha and J. Schmidhuber, “Recurrent world models facilitate policy evolution,” in *NeurIPS 2018*, 2018, pp. 2455–2467.
- [20] M. Deisenroth and C. Rasmussen, “Pilco: A model-based and data-efficient approach to policy search,” in *ICML 2011*, 2011, pp. 465–472.
- [21] R. Kaushik *et al.*, “Multi-objective model-based policy search for data-efficient learning with sparse rewards,” in *CoRL 2018*, vol. 87, 2018, pp. 839–855.
- [22] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *ICML 2016*, vol. 48, 2016, pp. 1050–1059.
- [23] G. Williams *et al.*, “Model predictive path integral control using covariance variable importance sampling,” *CoRR*, vol. abs/1509.01149, 2015.
- [24] R. Y. Rubinstein and D. P. Kroese, *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*, 2004.
- [25] K. Li and J. Malik, “Learning to optimize,” in *ICLR*, 2017.
- [26] C. Finn *et al.*, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *International Conference on Machine Learning, ICML*, vol. 70. PMLR, 2017, pp. 1126–1135.
- [27] S. Sæmundsson *et al.*, “Meta reinforcement learning with latent variable gaussian processes,” *arXiv preprint arXiv:1803.07551*, 2018.
- [28] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR 2015*, 2015.
- [29] A. Nichol *et al.*, “On first-order meta-learning algorithms,” *CoRR*, 2018.
- [30] D.-Y. Kim, “Spotmicro - robot dog by kdy0523,” Feb 2019. [Online]. Available: <https://www.thingiverse.com/thing:3445283>