

Efficient Learning of Locomotion Skills through the Discovery of Diverse Environmental Trajectory Generator Priors

Shikha Surana^{*1}, Bryan Lim^{*1}, Antoine Cully¹

Abstract—Data-driven learning based methods have recently been particularly successful at learning robust locomotion controllers for a variety of unstructured terrains. Prior work has shown that incorporating good locomotion priors in the form of trajectory generators (TGs) is effective at efficiently learning complex locomotion skills. However, defining a good, single TG as tasks/environments become increasingly more complex remains a challenging problem as it requires extensive tuning and risks reducing the effectiveness of the prior. In this paper, we present Evolved Environmental Trajectory Generators (EETG), a method that learns a diverse set of specialised locomotion priors using Quality-Diversity algorithms while maintaining a single policy within the Policies Modulating TG (PMTG) architecture. The results demonstrate that EETG enables a quadruped robot to successfully traverse a wide range of environments, such as slopes, stairs, rough terrain, and balance beams. Our experiments show that learning a diverse set of specialized TG priors is significantly (5 times) more efficient than using a single, fixed prior when dealing with a wide range of environments.

I. INTRODUCTION

Legged robots [1], [2], [3] have tremendous potential for societal impact as they can be used for applications involving a wide range of environments such as rough, cluttered and unstructured terrain. From search and rescue and inspection work [4] to carrying heavy payloads, legged robots have the potential to undertake many of the physical activities humans and animals are capable of that are dangerous and unhealthy.

However, legged robots are also underactuated high-dimensional systems with many constraints making them challenging to control. Recently, reinforcement learning (RL) approaches [5], [6], [7], [8], [9], [10] have become competitive to more conventional model-based optimization methods [11], [12], [13], [14], demonstrating state-of-art locomotion abilities both in simulation and in the real-world [5], [6]. These learnt controllers are especially robust when evaluated across many different environments and perturbations. Despite these significant advances, learning based approaches in robotics are notoriously known for being sample inefficient and usually requires a large amount of data [15], [16]. Researchers have tried to address this problem in a number of different ways. For example, improving the sample efficiency of the underlying RL algorithm used [17] or using fast, highly parallel simulators [18], [8]. Another effective way is to incorporate useful priors in the learning process.

Policies Modulating Trajectory Generator (PMTG) [19] is one such method which incorporates a parameterized

Trajectory Generator (TG) as a prior, separate to the learnt policy. PMTG makes learning complex locomotion tasks easier and demonstrates that a good locomotion prior can significantly help the efficiency of reinforcement learning (RL) methods [19]. Lee et al. [6] also used this PMTG framework when demonstrating state-of-the-art locomotion across wide range of environments in the real-world, further demonstrating the effectiveness of the TG prior for locomotion.

However, there are still some questions surrounding the PMTG method. How are the parameters of the TG defined? What parameters make a good prior? The parameters of the TGs used in prior work are usually defined manually by engineers based on intuition of the locomotion task of interest. For instance, a forward swinging TG motion is useful when learning to walk forward [19]. On the other hand, for more complex tasks such as learning to walk across a diversity of difficult environments, a more generic and unbiased TG motion of stepping up and down in place had to be used [6]. While this TG choice proved effective, this could reduce the effectiveness of the prior in helping learning and could indicate that the policy still has to do the bulk of the work as it has to deal with the different environments. For example, a good TG prior for ascending steps would differ from that of descending steps. In this paper, we address this by learning good priors for tasks instead of manually crafting them. More importantly, we also learn a diverse set of specialized priors using Quality-Diversity (QD) algorithms rather than using just a single prior.

The main contribution of our work is a novel framework, Evolved Environmental Trajectory Generators (EETG) for discovering a diverse set of specialized Trajectory Generators (TGs) which act as priors for more efficient learning. We demonstrate in our experiments that our method enables a simulated A1 quadruped robot to learn dynamic locomotion behaviors over diverse environment types such as slopes, uneven terrain, and steps. Our experiments show that EETG is as good as learning individual TGs and policies across all environments while being significantly more efficient. Our work demonstrates that learning a diverse set of TG prior is more effective than a single fixed TG especially when dealing with many tasks and environments.

II. RELATED WORK

Legged Locomotion. Locomotion controllers have traditionally been designed using a modular control framework. This framework breaks down the difficult control problem into smaller sub-problems. Each sub-problem makes

^{*}Equal Contribution

¹Imperial College London, United Kingdom. {ss5721, bw1116, a.cully}@ic.ac.uk




Fig. 1. **A1 quadruped robot, trained with the presented approach, in a variety of challenging environments.** (A) Ascending slope. (B) Descending slope. (C) Ascending stairs. (D) Descending stairs. (E) Uneven/rough terrain. (F) Narrow balance beam.

approximations such as mass-less limbs and point mass dynamics [11], [13] and apply heuristics [14] which are used alongside trajectory optimization, footstep planning, model predictive control (MPC) methods [20], [11], [12].

Alternatively, data-driven learning based approaches are fast becoming a go-to option due to recent advances showing impressive robustness and performance [5], [6], [9], [10], while at the same time requiring less modeling and expert optimization knowledge. One of the key ideas that make learning based methods perform so well and give these controllers incredible robustness especially in the real world, is domain randomization (DR). DR of simulator physics and visual features enabled learning dexterous manipulation to solve a Rubik’s cube [15]. Hwangbo et al. [5] then demonstrated DR was also effective for locomotion in dynamic legged systems. This idea of DR can be also extended towards curriculum learning based methods and used beyond just physical parameters and dynamics of the robot. Both Lee et al. [6] and Kumar et al. [21] separately showed that learning in diverse environment types and terrains can result in learnt skills which can be robust and generalize to new environments. Similarly, our work utilizes environment diversity as strategy to learn robust and generalizable controllers.

As mentioned in the introduction, we also build on the PMTG control architecture [19] to utilize the effectiveness of priors for learning. However, instead of manually crafting parameters of the trajectory generator (TG), we learn a diverse and high-performing set of specialised TGs.

Quality-Diversity. Quality-Diversity (QD) [22], [23], [24] is a growing branch of optimization methods which aims to find a large set of diverse and locally optimal solution. This is in contrast to conventional optimization algorithms which find a single objective maximising solution. QD algorithms have demonstrated their effectiveness across a range of applications including robotics [25], [26], [27], reinforcement learning [28], [29], video-game design [30], [31], engineering design optimization [32] and more.

In the context of robotics, QD algorithms have commonly been used to learn a diverse repertoire of primitive controllers [33] which can then be used effectively for downstream applications such a damage recovery [25], [34]

or with planning algorithms to perform longer horizon tasks such as navigation to a goal [34], [35], [36]. In our work, we learn the parameters of a controller which is represented as an open loop TG. However, a critical difference in our method is that the descriptors/cells of the QD algorithm are not defined by the behavior of the controller but by the environment (i.e. variations of stairs, rough terrain etc.) in which the controller is evaluated in and is selected before evaluation. This is akin to multi-task MAP-Elites [37].

QD-like algorithms like POET [29], [38] have also been used to evolve environments of increasing complexity while learning specialised paired policies for each environment. Similarly, our work also maintains a diversity of environments while trying to discover specialized parameters for each environment. However, we learn specialised TGs instead of specialised policy networks for each environment. We then later learn a single policy over all the specialized TGs. Our method also differs in that we do not evolve environments but define them beforehand as we are interested in discovering specialized priors and not in auto-curricula of the environment. However, we expect this to further help despite requiring additional complexity of an additional optimization process and more compute. We leave this for future work.

III. METHODS

Our algorithm is composed of two parts: 1) Learning of a diverse set of environment specialized TG priors (Section III-A) and 2) policy optimization within the PMTG architecture (Section III-B). Figure 2 shows an overview of the two phases of the EETG algorithm. We describe these two parts in the following subsections.

A. Discovering Diverse Specialised Trajectory Generators

Our goal is to find a diverse set of priors in the form of TGs, each of which are specialised and high-performing for the corresponding task and environment. This problem can be formalised as a Quality-Diversity (QD) optimization problem. In this subsection, we first define the TG representation used and how the TGs are parameterised. The parameter space of the TGs is then used as the solution space which




Fig. 2. **EETG Overview.** (A) A diverse and high-performing set of environmental priors is learnt using QD optimization.. First, a solution and target environment is selected from the grid (1-2). The solution’s parameters are perturbed (3) and evaluated in the target environment (4). Finally, the solution is placed in the target environment cell, where the best solution in the cell is kept (5). (B) Policy optimisation. At each training episode, a target environment and its corresponding prior (i.e., TG parameters) are selected from the grid. The residual policy is optimised to modulate the TG actions with ARS [39].

the QD algorithm searches over. We then go through our method, EETG which applies some key modifications to the original MAP-Elites algorithm, a popular QD algorithm.

TG representation and parameterization. In our work, the TG generates an open-loop smooth nominal periodic trajectory for each foot similar to Lee et al. [6]. The trajectory for each foot is defined in Eqn 1.

$$\begin{aligned}
 \phi_i &= \phi_{i,0} + (f_0 + f_i)\tau \\
 \beta_1 &= \frac{\sin(\phi_i + \frac{\pi}{2}) - 1}{2} & \beta_2 &= \frac{\sin(2\phi_i + \frac{\pi}{2}) - 1}{2} \\
 x_i &= (s \times \beta_1) + s \\
 y_i &= \begin{cases} y_{\delta,i} & \phi_i \in [0, \pi] \\ y_{\delta,i} - (t \times \beta_2) & \phi_i \in [\pi, 2\pi] \end{cases} \\
 z_i &= \begin{cases} h & \phi_i \in [0, \pi] \\ h - (l \times \beta_2) & \phi_i \in [\pi, 2\pi] \end{cases}
 \end{aligned} \tag{1}$$

The TG determines the relative x, y, z coordinates of the foot positions for each leg i from the corresponding hip reference frame. The foot positions are computed based on the periodic phase variable ϕ and the residual TG frequencies f output by the residual policy network. The periodic phase ϕ which governs the frequency of the gait, ranges from $[0, 2\pi]$ and when $\phi \in [0, \pi]$ the leg is in swing movement, and when $\phi \in [\pi, 2\pi]$ the leg is in stance movement. We use a fixed base frequency f_0 of 1.25Hz. The motion of the TG is parametrized by 5 variables which make up the solution vector $\mathbf{x} = [s, t, l, y, \vec{\phi}_0]$ for EETG. All legs share the same parameters and these are defined as follows:

- Swing s : determines the maximum forwards and backwards range of the trajectory (Fig. 3A), and the parameter is continuous between 0m and 0.08m from the initial x position of each leg.
- Turn t : determines the maximum inwards (towards the center of the body) or outwards (Fig. 3B), and the range is continuous between 0m and 0.15m from the initial y position of each leg.
- Lift l : determines maximum height foot is raised above the ground (Fig. 3C) during the trajectory, and the range




Fig. 3. **Visualisations of the trajectory generator parameters.** (A) Swing - legs move forwards/backwards (along x -axis). (B) Turn - legs move leftwards/rightwards (along y -axis). (C) Lift - legs move upwards/downwards (along z -axis). (D) Y-position of the legs.

is continuous between 0m and 0.2m from the initial z position of each leg.

- y position of all legs: determines the initial y position of each leg (i.e., a negative y value means the legs are closer to the center of the body and a positive y value means the legs are further apart from the center of the body) (Fig. 3D), and the range is continuous between -0.05m to 0.12m which is offset to the initial y position of each leg.
- Phase offset $\vec{\phi}_0$: is a vector composed of 4 different phase offset values (i.e. one for each leg) which define the gait coordination between the legs. This parameter can be one of four different gaits: walk $[0, 0.25, 0.5, 0.75]$, trot $[0, 0.5, 0.5, 0]$, bound $[0, 0.5, 0, 0.5]$, and pronk $[0, 0, 0, 0]$ and thus, the parameter takes the form of a discrete integer between 0 and 3.

Evolved Environmental Trajectory Generators. As mentioned earlier, we formulate the problem of finding a diverse set of specialised priors as a QD problem. QD algorithms solve this problem by maintaining a population of solutions each defined by its performance when evaluated and a descriptor, a vector used to differentiate solutions to encourage diversity. We use MAP-Elites (ME), a popular QD algorithm, which tessellates the descriptor space into

a grid of cells [40], [25]. Because we are interested in solutions that work well in a diversity of environments, the cells of the descriptor space each represent an environment rather than a behavior of a solution. Additionally, these are defined beforehand like the multi-task ME setting [37]. In other words, we maintain container in the form of grid of diverse environments where each cell in the grid represents a task defined by the environment type and parameterisation (Fig. 2 A) and stores the best corresponding solution. Further details regarding environment parameterisation are provided in Section IV-A Fig. 2 and Alg. 1 shows the procedure of EETG. EETG first begins with random initialisation of solutions in 10% of the cells in the grid by uniformly sampling from the parameter space target environment in the grid. Solutions are added in to the corresponding cells even if they have a low fitness. EETG then follows a loop. At each iteration, a solution from a cell is first selected. It uniformly selects an existing solution from the grid. Next, a target cell environment is selected from the grid. To encourage effective *goal-switching* [41], [42], the target environment type is selected with a $p = 0.7$ probability to be same environment type as the solution selected and with $p = 0.1$ each of the three remaining environment types. The environment parameterisation is then sampled uniformly from the range of the corresponding selected environment type. Effective goal switching is key to EETG and is what makes QD algorithms work well. It allows solutions from different environments to act as useful *stepping stones* [29] in the optimization procedure.

The selected solution is mutated/perturbed via a variation operator. We use the iso-line directional variation operator [43]. The solution is then evaluated in the target environment by rolling out the open-loop TG motion to obtain the reward and performance of the solution. If the target environment cell is empty, the solution is inserted in the cell. However, if the cell already contains a solution, the fitness scores are compared, and the solution with the highest fitness score is kept in the cell. This procedure is repeated until we hit the evaluation budget defined.

B. Robustness via Policy Modulation

Given a diverse set of TGs, EETG builds on the PMTG [19] control architecture by learning a residual policy represented as a deep neural network to modulate and provide robustness to the base TG using RL.

Fig. 2B shows an overview of the PMTG control architecture used. At each timestep, the policy receives the robot state observation s_t , the TG foot frequencies, and the TG swing s_i , turn t_i and lift l_i parameters as input. The residual policy outputs the x_i, y_i, z_i residuals for each foot i and the four frequency multiplier f_i (one for each leg) to modulate the internal phase of the TGs (see Eqn1). The TG then takes TG parameters \mathbf{x} along with the residual frequencies f_i as input to compute the nominal x, y, z foot positions (see Eqn. 1). The four legs are decoupled for greater flexibility as the frequency f_i of each leg can be separately adjusted which leads to more robust locomotion. The nominal positions are then summed

with the residuals output by the policy to determine the target foot positions. Inverse kinematics is then applied to the target foot positions to compute the target joint angles which are given to a PD impedance controller used to determine the required torque for each joint. The policy and TG operate at 60 Hz while the tracking of the target foot positions by inverse kinematics and the PD controller operate at 240 Hz.

Algorithm 1 EETG

Require:

```

 $M \leftarrow$  number of iterations of MAP-Elites (ME)
 $G \leftarrow$  ME grid containing environments and solutions
 $G.init()$ 
for  $gen = 1, \dots, M$  do
   $\mathbf{x}, E \leftarrow SELECT(G)$ 
   $\tilde{\mathbf{x}} \leftarrow MUTATE(\mathbf{x})$ 
   $fitness \leftarrow EVALUATE(\tilde{\mathbf{x}}, E)$ 
   $G \leftarrow ADD(\tilde{\mathbf{x}}, E, fit, G)$ 

```

The residual policy is optimized using Evolutionary Strategies, more specifically Augmented Random Search (ARS) [39] as also done originally by Iscen et al. [19]. Conventionally, the policy is trained with only a single pre-defined TG. In our work, in order to maintain just a single residual policy to handle the diversity of learnt TGs generated, we also condition the residual policy on the TG parameters $\pi(\cdot|\mathbf{x}, s)$. At each iteration of ARS, we uniformly sample from an environment and its corresponding TG solution from a cell. To learn robust policies, we add gaussian noise ($\sigma \sim$ discretization of one cell) to the environment parameterization during policy optimization.

IV. EXPERIMENTS

A. Experimental Setup

We evaluate our method on the task of learning to walk forward over a variety of challenging terrains. We use the Unitree A1 quadruped robot and the PyBullet [44] simulator for all our experiments. The state of the robot consists of the CoM base orientation, CoM base velocity, CoM base angular velocity, joint positions, joint velocities and the foot positions in the base frame. We do not consider any privileged information, which are often used to train robust policies in simulation [6], [21]

Environment Setup. For our experiments, we consider 80 different environments which make up our environmental grid (as explained in Section III-A). The grid is 4x20 grid composed of four different environment types (Fig 1); slopes, stairs, uneven terrain and balance beams, where we consider a range of 20 environment encoding variations. The slope environment is parameterised by a steady decline and incline angle that ranges from -11.5° to 11.5° , where negative values represent walking down the slope and positive values walking up the incline. The stairs environment is parameterised by the step height of the stairs ranging from -0.1m to 0.1m where again both descending and ascending stair environments correspond to positive and negative values of the parameters

repectively. Next, the uneven terrain environment is parameterised by a maximum height variable which varies from -0.1m to 0.1m. The terrain is then constructed by a composition of squares with a fixed size and a random height that is sampled from a uniform distribution in the range from zero to this maximum height parameter. Negative values correspond to holes while positive values correspond to bumps. Lastly, the balance beam environments are parameterised simply by the width of the beam which varies from 0.25m to 0.75m.

Reward Function. We use the same reward function to evaluate the performance of the solution in the EETG phase and the policy training phase. The reward term r_t at each timestep is given by Eqn. 2 and is composed of five terms: target linear and angular velocity tracking reward r_{lv} and r_{av_t} , angular velocity penalty r_{av_p} , foot position smoothness reward r_s , and a, torque penalty r_{tp} . The performance (also known as the fitness) of a solution is simply the return R which is computed as the sum of the rewards $\sum_{t=0}^T r_t$ obtained at each timestep t of the episode length T .

$$r_t = r_{lv} + r_{av_t} + r_{av_p} + r_s + r_{tp} \quad (2)$$

Baselines. We compare our method against two baseline PMTG algorithms to extensively understand and demonstrate the benefits of learning a diverse set of specialized priors in EETG. The first baseline algorithm is the vanilla PMTG algorithm in which a residual policy operates on a single fixed pre-defined TG. Next, we compare against a baseline in which we first learn a good TG using the Covariance Matrix Adaptation Evolution Strategy (CMAES) [45] and consequently train the residual policy over this learnt TG. We refer to the this baseline as CMAES-PMTG. Both baselines use the same policy and PMTG architecture and use ARS for the policy optimization stage.

For a more complete comparison, we consider two different variations of the baseline vanilla PMTG and CMAES-PMTG algorithms:

- **Encoded (Enc.):** The first variant considers only a single TG and a single residual policy that is optimised on all the environment types and variations. To enable the policy to adapt to the environments, the policy receives the environment type and variation as a one-hot encoded vector input.
- **Independent (Ind.):** This variation is used as an upper baseline. We use independent TGs and policies for each environment (i.e. one TG and one residual policy per environment cell) in the grid. As a result, each TG and policy can specialise for each specific environment. This variant requires significantly more evaluations and is the least sample efficient as we have to independently train TGs (CMAES-PMTG baseline) and policies (both baselines) in each environment.

The total evaluation budget for training given to each variant is provided in Table I. To evaluate each algorithm, we take the final TG(s) and optimized residual policies provided at the end of each algorithm. This is deployed in each of the 80 environments in the grid. To ensure that we

evaluate robustness to unseen or variations in the environment, sampled gaussian noise is added to the environment parameterization during each evaluation. As a metric, we report the average total reward obtained for each environment over 20 replications of evaluations.

B. Results

Figure 4 and Table I shows the results of the performance of EETG and baselines for each environment type. Despite using the least number of samples, we can observe that EETG performs the best and is capable of traversing a variety of diverse terrains (see accompanying video). This can be attributed to a few factors explained by the performance of the baselines. First, we can observe that both vanilla PMTG variants obtained the lowest performance across all the environment types when compared to both EETG and CMAES-PMTG baselines. This result demonstrates the importance of a learnt TG prior over the constraints and biases of a fixed TG prior. This is because a TG that is learnt, already begins with a higher performance than a pre-defined TG. Thus, this shows that the samples used to learn a good TG prior on the task is more effective for learning than providing more samples to train policy with a fixed TG prior.

Next, we can also observe that the algorithms with a set of specialized TG's such as EETG and CMAES-PMTG (Ind.) tend to perform better than the variants with only a single TG. This observation can be explained by our earlier hypothesis where a single prior can struggle when learning on a diverse set of tasks/environments and reduces its effect in helping the learning algorithms. Learning a diverse set of TG priors can help to alleviate this when dealing with a diversity of environments and restore the effectiveness of the prior. Additionally, as observed previously, the process of learning a prior seems to be more sample efficient than learning a policy directly which makes learning a set of TG priors feasible.

Lastly, but also most importantly, we can observe that EETG performs slightly better or very similarly to CMAES-PMTG (Ind.) despite requiring five times less samples. This demonstrates that EETG can learn locomotion skills efficiently across a diversity of environments as good or better than a complete set of independently trained policies and TGs. We hypothesize that this is due to the goal switching phenomena and the availability of useful stepping stones in QD algorithms that we use in EETG. This also corresponds to observations in literature where it was shown that learning occurs more efficiently through the process of goal switching across a diversity of tasks in comparison to the more brute force goal directed optimization [46], [47].

One additional observation is that all algorithms (including EETG) perform badly on the balance beam environment. This can be explained by the fact that our robot is blind and only considers proprioception while this environment can easily result in catastrophic failure (i.e falling off) and would require vision information to correct behaviors. This is in contrast to the other three environments in which proprioception can be used to react effectively.

TABLE I
SUMMARY OF EETG , BASELINES, EVALUATION BUDGET AND PERFORMANCE.

Algorithm	Alg. Description		Evaluation Budget			Median Total Reward			
	TG	Policy	TG Opt.	Policy Opt.	Total	Slope	Stairs	Uneven Terrain	Balance Beam
EETG	Set, Learnt	Single	384,000	4,608,000	4,992,000	194.13	172.73	168.45	139.90
PMTG (Enc.)	Single, Fixed	Single	-	5,280,000	5,280,000	156.62	156.71	156.68	156.60
PMTG (Ind.)	Set, Fixed	Set	-	23,040,000	23,040,000	145.35	143.13	153.17	153.05
CMAES-PMTG (Enc.)	Single, Learnt	Single	384,000	4,608,000	4,992,000	186.85	164.19	167.79	158.60
CMAES-PMTG (Ind.)	Set, Learnt	Set	1,280,000	23,040,000	24,320,000	187.08	169.12	168.11	165.60




Fig. 4. The total reward obtained by EETG (blue) and the baselines: PMTG (Enc.) (orange), PMTG (Ind.) (green), CMAES-PMTG (Enc.) (red), and CMAES-PMTG (Ind.) (pink). The results are shown separately for each environment type.

C. Ablation

To further evaluate and find the optimal configuration for EETG, we aim to answer the following questions by running an ablation against a few variations of EETG: (1) Does adding repeated, sequential optimisation of the MAP-Elites grid and residual policy help performance? (2) Does including the policy during the MAP-Elites loop help performance?

We investigate this by comparing against two variants of EETG, both of which consists of multiple iteration loops between the two optimization phases: MAP-Elites and Policy Optimization. In EETG Itr., the policy is not used in the MAP-Elites TG learning stage where only the open loop TG is used. In EETG Itr.+ Policy, the latest policy from the policy optimization phase is also used in the controller during the MAP-Elites TG learning stage. In this experiment, all variants are given the same total evaluation budget as the default EETG but divided equally between the number of loop iterations.

Figure 5 shows the performances of EETG and its variations, EETG Itr. and EETG Itr.+ Policy, for each environment type. It can be seen that EETG obtains relatively better rewards across all the environments. However, the difference in total rewards between EETG and its variants only ranges from 0% to 4% of the maximum total reward that can be obtained in one episode. Furthermore, the differences in the observed behaviours of the robot trained with each algorithm are minimal. Therefore, it cannot be concluded that these




Fig. 5. The total rewards obtained by EETG (blue), and the two variants, EETG Itr. (orange) and EETG Itr.+ Policy (green). The results are shown separately for each different environment types.

variants lead to significant changes in the final performance when compared to EETG. Based on these results, it does not seem that (1) including multiple iterations of the ME and PMTG loop or (2) involving the residual policy in the ME evaluation stage lead to any significant changes in the final performance of the EETG algorithm.

V. DISCUSSION AND CONCLUSION

In summary, we present EETG, a novel method which evolves a set of high-performing TG priors for a diversity of environments and uses a single residual policy that regulates the learnt TGs. We show that this enables a quadruped robot to traverse a wide range of challenging environments. Our experiments highlight the need and effectiveness of a set of specialised TG priors over a single fixed manually defined TG prior when dealing with a variety of environments. We demonstrate that EETG is significantly more sample efficient and as effective as learning independent TGs and policies across diverse environments.

One of the limitations of this work is that we assume knowledge of what TG in the set to use during deployment. In future work, it would be interesting to eliminate this assumption in order to autonomously and efficiently select a suitable TG using Bayesian Optimisation through Intelligent Trial and Error [25], [34] via online data collected during deployment.

ACKNOWLEDGMENT

We thank Manon Flageat for reviewing earlier versions of this paper. This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) grant EP/V006673/1 project REcoVER.

REFERENCES

- [1] G. Bleedt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, and S. Kim, "Mit cheetah 3: Design and control of a robust, dynamic quadruped robot," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 2245–2252.
- [2] M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch *et al.*, "Anymal-a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 38–44.
- [3] B. Katz, J. Di Carlo, and S. Kim, "Mini cheetah: A platform for pushing the limits of dynamic quadruped control," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 6295–6301.
- [4] C. Gehring, P. Fankhauser, L. Isler, R. Diethelm, S. Bachmann, M. Potz, L. Gerstenberg, and M. Hutter, "Anymal in the field: Solving industrial inspection of an offshore hvdc platform with a quadrupedal robot," in *Field and Service Robotics*. Springer, 2021, pp. 247–260.
- [5] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaa5872, 2019.
- [6] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [7] J. Siekmann, Y. Godse, A. Fern, and J. Hurst, "Sim-to-real learning of all common bipedal gaits via periodic reward composition," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7309–7315.
- [8] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [9] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [10] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *arXiv preprint arXiv:2205.02824*, 2022.
- [11] J. Di Carlo, P. M. Wensing, B. Katz, G. Bleedt, and S. Kim, "Dynamic locomotion in the mit cheetah 3 through convex model-predictive control," in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2018, pp. 1–9.
- [12] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1560–1567, 2018.
- [13] D. Kim, J. Di Carlo, B. Katz, G. Bleedt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint arXiv:1909.06586*, 2019.
- [14] G. Bleedt and S. Kim, "Extracting legged locomotion heuristics with regularized predictive control," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 406–412.
- [15] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.
- [16] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [17] L. Smith, J. C. Kew, X. B. Peng, S. Ha, J. Tan, and S. Levine, "Legged robots that keep on learning: Fine-tuning locomotion policies in the real world," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1593–1599.
- [18] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [19] A. Iscen, K. Caluwaerts, J. Tan, T. Zhang, E. Coumans, V. Sindhwani, and V. Vanhoucke, "Policies modulating trajectory generators," in *Conference on Robot Learning*. PMLR, 2018, pp. 916–926.
- [20] M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, and S. Schaal, "Fast, robust quadruped locomotion over challenging terrain," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2665–2670.
- [21] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [22] J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, p. 40, 2016.
- [23] A. Cully and Y. Demiris, "Quality and diversity optimization: A unifying modular framework," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 2, pp. 245–259, 2017.
- [24] K. Chatzilygeroudis, A. Cully, V. Vassiliades, and J.-B. Mouret, "Quality-diversity optimization: a novel branch of stochastic optimization," in *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer, 2021, pp. 109–135.
- [25] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [26] A. Salehi, A. Coninx, and S. Doncieux, "Few-shot quality-diversity optimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4424–4431, 2022.
- [27] B. Lim, A. Reichenbach, and A. Cully, "Learning to walk autonomously via reset-free quality-diversity," *arXiv preprint arXiv:2204.03655*, 2022.
- [28] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, "First return, then explore," *Nature*, vol. 590, no. 7847, pp. 580–586, 2021.
- [29] R. Wang, J. Lehman, J. Clune, and K. O. Stanley, "Poet: open-ended coevolution of environments and their optimized solutions," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2019, pp. 142–151.
- [30] M. C. Fontaine, R. Liu, A. Khalifa, J. Modi, J. Togelius, A. K. Hoover, and S. Nikolaidis, "Illuminating mario scenes in the latent space of a generative adversarial network," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 7, 2021, pp. 5922–5930.
- [31] S. Earle, J. Snider, M. C. Fontaine, S. Nikolaidis, and J. Togelius, "Illuminating diverse neural cellular automata for level generation," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 68–76.
- [32] A. Gaier, A. Asteroth, and J.-B. Mouret, "Data-efficient design exploration through surrogate-assisted illumination," *Evolutionary computation*, vol. 26, no. 3, pp. 381–410, 2018.
- [33] A. Cully and J.-B. Mouret, "Behavioral repertoire learning in robotics," in *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, 2013, pp. 175–182.
- [34] K. Chatzilygeroudis, V. Vassiliades, and J.-B. Mouret, "Reset-free trial-and-error learning for robot damage recovery," *Robotics and Autonomous Systems*, vol. 100, pp. 236–250, 2018.
- [35] R. Kaushik, P. Desreumaux, and J.-B. Mouret, "Adaptive prior selection for repertoire-based online adaptation in robotics," *Frontiers in Robotics and AI*, p. 151, 2020.
- [36] B. Lim, L. Grillotti, L. Bernasconi, and A. Cully, "Dynamics-aware quality-diversity for efficient learning of skill repertoires," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 5360–5366.
- [37] J.-B. Mouret and G. Maguire, "Quality diversity for multi-task optimization," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, 2020, pp. 121–129.
- [38] R. Wang, J. Lehman, A. Rawal, J. Zhi, Y. Li, J. Clune, and K. Stanley, "Enhanced poet: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions," in *International Conference on Machine Learning*. PMLR, 2020, pp. 9940–9951.
- [39] H. Mania, A. Guy, and B. Recht, "Simple random search provides a competitive approach to reinforcement learning," *arXiv preprint arXiv:1803.07055*, 2018.
- [40] J.-B. Mouret and J. Clune, "Illuminating search spaces by mapping elites," *arXiv preprint arXiv:1504.04909*, 2015.
- [41] A. Nguyen, J. Yosinski, and J. Clune, "Understanding innovation engines: Automated creativity and improved stochastic optimization via deep learning," *Evolutionary computation*, vol. 24, no. 3, pp. 545–572, 2016.

- [42] J. Clune, "Ai-gas: Ai-generating algorithms, an alternate paradigm for producing general artificial intelligence," *arXiv preprint arXiv:1905.10985*, 2019.
- [43] V. Vassiliades and J.-B. Mouret, "Discovering the elite hypervolume by leveraging interspecies correlation," *Proceedings of the Genetic and Evolutionary Computation Conference*, 2018.
- [44] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [45] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)," *Evolutionary computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [46] J. Lehman and K. O. Stanley, "Efficiently evolving programs through the search for novelty," in *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, 2010, pp. 837–844.
- [47] B. G. Woolley and K. O. Stanley, "On the deleterious effects of a priori objectives on evolution and representation," in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, 2011, pp. 957–964.