

ANIMAL LOCOMOTION

DTC: Deep Tracking Control

Fabian Jenelten*, Junzhe He, Farbod Farshidian†, Marco Hutter

Legged locomotion is a complex control problem that requires both accuracy and robustness to cope with real-world challenges. Legged systems have traditionally been controlled using trajectory optimization with inverse dynamics. Such hierarchical model-based methods are appealing because of intuitive cost function tuning, accurate planning, generalization, and, most importantly, the insightful understanding gained from more than one decade of extensive research. However, model mismatch and violation of assumptions are common sources of faulty operation. Simulation-based reinforcement learning, on the other hand, results in locomotion policies with unprecedented robustness and recovery skills. Yet, all learning algorithms struggle with sparse rewards emerging from environments where valid footholds are rare, such as gaps or stepping stones. In this work, we propose a hybrid control architecture that combines the advantages of both worlds to simultaneously achieve greater robustness, foot-placement accuracy, and terrain generalization. Our approach uses a model-based planner to roll out a reference motion during training. A deep neural network policy is trained in simulation, aiming to track the optimized footholds. We evaluated the accuracy of our locomotion pipeline on sparse terrains, where pure data-driven methods are prone to fail. Furthermore, we demonstrate superior robustness in the presence of slippery or deformable ground when compared with model-based counterparts. Last, we show that our proposed tracking controller generalizes across different trajectory optimization methods not seen during training. In conclusion, our work unites the predictive capabilities and optimality guarantees of online planning with the inherent robustness attributed to offline learning.

INTRODUCTION

Trajectory optimization (TO) is a commonly deployed instance of optimal control for designing motions of legged systems and has a long history of successful applications in rough environments since the early 2010s (1, 2). These methods require a model of the robot's kinematics and dynamics during run time, along with a parametrization of the terrain. Until recently, most approaches have used simple models, such as single rigid body (3) or inverted pendulum dynamics (4, 5), or have ignored the dynamic effects altogether (6). Research has shifted toward more complex formulations, including centroidal (7) or full-body dynamics (8). The resulting trajectories are tracked by a whole-body control (WBC) module, which operates at the control frequency and uses full-body dynamics (9). Despite the diversity and agility of the resulting motions, there remains a considerable gap between simulation and reality because of unrealistic assumptions. Most problematic assumptions include perfect state estimation, occlusion-free vision, known contact states, zero foot-slip, and perfect realization of the planned motions. Sophisticated hand-engineered state machines are required to detect and respond to various special cases not accounted for in the modeling process. Nevertheless, highly dynamic jumping maneuvers performed by Boston Dynamics' bipedal robot Atlas demonstrate the potential power of TO.

Reinforcement learning (RL) has emerged as a powerful tool in recent years for synthesizing robust legged locomotion. Unlike model-based control, RL does not rely on explicit models. Instead, behaviors are learned, most often in simulation, through random interactions of agents with the environment. The result is a closed-loop control policy, typically represented by a deep neural network, that maps raw observations to actions. Handcrafted state machines become obsolete because all relevant corner cases are eventually visited during training.

End-to-end policies, trained from user commands to joint target positions, have been deployed successfully on quadrupedal robots such as ANYmal (10, 11). More advanced teacher-student structures have substantially improved the robustness, enabling legged robots to overcome obstacles through touch (12) and perception (13). Although locomotion across gaps and stepping stones is theoretically possible, good exploration strategies are required to learn from the emerging sparse reward signals. So far, these terrains could only be handled by specialized policies, which intentionally overfit to one particular scenario (14) or a selection of similar terrain types (15–18). Despite promising results, distilling a unifying locomotion policy may be difficult and has only been shown with limited success (19).

Some of the shortcomings that appear in RL can be mitigated using optimization-based methods. Although the problem of sparse gradients still exists, two important advantages can be exploited: First, cost function and constraint gradients can be computed with a small number of samples. Second, poor local optima can be avoided by precomputing footholds (5, 8), presegmenting the terrain into steppable areas (7, 20), or smoothing out the entire gradient landscape (21). Another advantage of TO is its ability to plan actions ahead and predict future interactions with the environment. If model assumptions are generic enough, this allows for great generalization across diverse terrain geometries (7, 21).

The sparse gradient problem has been addressed extensively in the learning community. A notable line of research has focused on learning a specific task while imitating expert behavior. The expert provides a direct demonstration for solving the task (22, 23) or is used to impose a style while discovering the task (24–26). These approaches require collecting expert data, commonly done offline, through either retargeted motion capture data (24–26) or a TO technique (22, 23). The reward function can then be formulated to be dense, meaning that agents can collect nontrivial rewards even if they do not initially solve the task. Nonetheless, the goal is not to preserve the expert's accuracy but rather to lower the sample and reward complexity by leveraging existing knowledge.

Robotic Systems Lab, ETH Zurich, 8092 Zurich, Switzerland.

*Corresponding author. Email: fabian.jenelten@ethz.ch

†Present address: Boston Dynamics AI Institute, 145 Broadway, Cambridge, MA, USA.

To further decrease the gap between the expert and the policy performance, we speculate that the latter should have insight into the expert's intentions. This requires online generation of expert data, which can be conveniently achieved using any model-based controller. Unfortunately, rolling out trajectories is often orders of magnitude more expensive than a complete learning iteration. To circumvent this problem, one possible alternative is to approximate the expert with a generative model, for instance, by sampling footholds from a uniform distribution (15, 16) or from a neural network (17, 27, 28). However, for the former group, it might be challenging to capture the distribution of an actual model-based controller, whereas the latter group still does not solve the exploration problem itself.

In this work, we propose to guide exploration through the solution of TO. Because such data will be available both on- and offline, we refer to it as “reference” and not expert motion. We use a hierarchical structure introduced in deep loco (28), where a high-level planner proposes footholds at a lower rate and a low-level controller follows the footholds at a higher rate. Instead of using a neural network to generate the foothold plan, we leverage TO. Moreover, we use the target footholds not only as an indicator for a rough high-level direction but also as a demonstration of optimal foot placement.

The idea of combining model-based and model-free approaches is not new in the literature. For instance, supervised (29) and unsupervised (30, 31) learning has been used to warm-start nonlinear solvers. RL has been used to imitate (22, 23) or correct (32) motions obtained by solving TO problems. Conversely, model-based methods have been used to check the feasibility of learned high-level commands (27) or to track learned acceleration profiles (33). Compared with (32), we do not learn corrective joint torques around an existing WBC but, instead, learn the mapping from reference signals to joint positions in an end-to-end fashion.

To generate the reference data, we rely on an efficient TO method called terrain-aware motion generation for legged systems (TAMOLS) (21). It optimizes over footholds and base pose simultaneously, thereby enabling the robot to operate at its kinematic limits. We let the policy observe only a small subset of the solution, namely, planar footholds, desired joint positions, and the contact schedule. We found that these observations are more robust under the common pitfalls of model-based control while still providing enough information to solve the locomotion task. In addition, we limit computational costs arising from solving the optimization problems by using a variable update rate. During deployment, the optimizer runs at the fastest possible rate to account for model uncertainties and external disturbances.

Our approach incorporates elements introduced in (14), such as time-based rewards and position-based goal tracking. However, we reward desired foothold positions at planned touchdown instead of rewarding a desired base pose at an arbitrarily chosen time. Last, we use an asymmetric actor-critic structure similar to (22), where we provide privileged ground truth information to the value function and noisified measurements to the network policy.

We trained more than 4000 robots in parallel for 2 weeks on challenging ground covering a surface area of more than 76,000 m². Throughout the entire training process, we generated and learned from about 23 years of optimized trajectories. The combination of offline training and online replanning results in accurate, agile, and robust locomotion. As showcased in Fig. 1 and Movie 1, with our hybrid control pipeline, ANYmal (34) can skillfully traverse parkours with high precision and confidently overcome uncertain environments

with high robustness. Without the need for any post-training, the tracking policy can be deployed zero shot with different TO methods at different update rates. Moreover, Movie 2 demonstrates successful deployment in search-and-rescue scenarios, which demand both accurate foot placement and robust recovery skills. The contributions of our work are therefore twofold: First, we enable the deployment of model-based planners in rough and uncertain real-world environments. Second, we create a single unifying locomotion policy that generalizes beyond the limitations imposed by state-of-the-art RL methods.

RESULTS

To evaluate the effectiveness of our proposed pipeline, hereby referred to as Deep Tracking Control (DTC), we compared it with four different approaches: two model-based controllers, TAMOLS (21) and a nonlinear model predictive control (MPC) method presented in (7), and two data-driven methods, as introduced in (13) and (11). We refer to those as baseline-to-1 (TAMOLS), baseline-to-2 (MPC), baseline-rl-1 (teacher/student policy), and baseline-rl-2 (RL policy), respectively. These baselines mark the state of the art in MPC and RL before this work, and they have been tested and deployed under various conditions. If not noted differently, all experiments were conducted in the real world.

Evaluation of robustness

We conducted three experiments to evaluate the robustness of our hybrid control pipeline. The intent was to demonstrate survival skills on slippery ground and recovery reflexes when visual data are not consistent with proprioception or are absent altogether. We rebuilt harsh environments that are likely to be encountered on sites of natural disasters, where debris might further break down when stepped onto, and construction sites, where oil patches create slippery surfaces.

In the first experiment, we placed a rectangular cover plate with an area of 0.78 m by 1.19 m on top of a box with the same length and width as well as height (0.37 m) (Fig. 2A). The cover plate was shifted to the front by half of the box's length. ANYmal was then steered over the cover plate, which pitched down as soon as its center of mass passed beyond the edge of the box. Facing only forward and backward, the plate's movement was not detected through the depth cameras and could only be perceived through proprioceptive sensors. Despite the error between map and odometry reaching up to 0.4 m, the robot managed to successfully balance itself. This experiment was repeated three times with consistent outcomes.

In our second experiment (Fig. 2B), we created an obstacle parkour with challenging physical properties. A large wooden box with a sloped front face was placed next to a wet and slippery whiteboard. We increased the difficulty by placing a soft foam box in front and a rolling transport cart on top of the wooden box. The robot was commanded to walk over the objects with random reference velocities for approximately 45 s, after which the objects were redistributed to their original locations to account for any potential displacement. This experiment was repeated five times. Despite not being trained on movable or deforming obstacles, the robot demonstrated its recovery skills in all five trials without any falls.

The tracking policy was trained with perceptive feedback, meaning that the policy and the motion planner had partial or complete insight into the local geometrical landscape. Nevertheless, the locomotion policy was still capable of overcoming many




Fig. 1. Robust and precise locomotion in various indoor and outdoor environments. The marriage of model-free and model-based control allows legged robots to be deployed on terrains where steppable contact surfaces are sparse (bottom left) and environmental uncertainties are high (top right).

obstacles completely blind. To simulate a scenario with damaged depth sensors, we let ANYmal blindly walk over a stair with two treads, each 0.18 m high and 0.29 m wide (Fig. 2C). The experiment was repeated three times up and down, with a heading velocity (absolute value) of 0.5, 0.75, and 1 m/s. In some cases, a stair tread was higher than the swing motion of a foot. Thanks to a learned swing reflex, the stair set could be successfully cleared in all trials. We note that the same stair set was passed by a blind-folded version of baseline-rl-1 (13), which was trained in a complex teacher/student environment. In contrast, our method relies on an asymmetric actor/critics structure, achieving a similar level of robustness. Accompanying video clips can be found in movie S1.

Evaluation of accuracy

We demonstrate the precision of foothold tracking by devising a complex motion that requires the robot to perform a turn-in-place maneuver on a small surface area of 0.94 m by 0.44 m. The robot was commanded to walk up a slope onto a narrow table, execute a complete 360° turn, and descend onto a pallet. Some snapshots of the experiment are provided in Fig. 3A, whereas the full video is contained in movie S2.

To evaluate the quality of the foothold tracking, we collected data while ANYmal walked on flat ground. Each experiment lasted for approximately 20 s and was repeated with eight different heading velocities selected from ± 1.0 , ± 0.8 , ± 0.6 , or ± 0.4 m/s. We measured the tracking error as the smallest horizontal distance between a foot and its associated foothold during a stance phase. As shown in Fig. 3B, the footholds could be tracked with a very high precision of 2.3 cm and a SD of 0.48 cm when averaged over the broad spectrum of heading velocity commands.

Deployment with MPC

The maximum height that DTC in combination with TAMOLS can reliably overcome is about 0.40 m. The policy might hesitate to climb up taller objects because of the risk of potential knee joint collisions with the environment. This limitation is inherent to the chosen TO method, which only considers simplified kinematic constraints. We, therefore, deployed DTC with the planner of baseline-to-2, which takes into account the full kinematics of the system. To allow for zero-shot generalization, we implemented the same trotting gait as experienced during training. With this enhanced setup, ANYmal could climb up a box with a height of 0.48 m. This is 50% higher than what



Movie 1. The combination of model-based and model-free control techniques allows for accurate and robust robot legged locomotion.



Movie 2. Our method allows ANYmal to navigate complex, irregular, previously unexplored real-world environments.

baseline-rl-1 could climb and 380% more than what was reported for baseline-rl-2. The box climbing experiment was successfully repeated five times. The results are shown in movie S2 and for one selected trial in Fig. 3D. Furthermore, we measured the tracking error on flat ground. Despite the wider stance configuration of baseline-to-2, the error was found to be only 0.03 m on average (Fig. 3C).

The above two results seem to be unexpected at first glance but are easy to explain when considering the observation space and the training environment. Although the base-pose trajectory is considerably more detailed for baseline-to-2 because of frequency-loop shaping and increased system complexity, the foothold patterns are nevertheless quite similar. Thus, good generalization is facilitated by the specific choice of observations, which hides the optimized base pose from the policy. Furthermore, some terrains within the training environment can be seen as a combination of gaps and boxes, where each box is surrounded by a gap. During training, TAMOLS placed the footholds sufficiently far away from the box to avoid stepping into the gap. This allowed the policy to learn climbing maneuvers without knee joint collisions. Baseline-to-2, being aware of the spatial coordinates of the knees, naturally produces a similar foothold pattern even in the absence of the gap.

Benchmark against model-based control

TO was proven to be effective in solving complex locomotion tasks in simulation, such as the parkour shown in Fig. 4A. This parkour has been successfully traversed by ANYmal using baseline-to-1 and baseline-to-2, whereas it was found to be nontraversable for baseline-rl-1 and baseline-rl-2 (7). With our proposed approach, ANYmal could cross the same obstacle parkour in simulation back and forth at a speed of 1 m/s, which was 20% faster than baseline-to-1. The corresponding video clip can be found in movie S3.

Model-based controllers react sensitively to violation of model assumptions, which hinders applications in real-world scenarios, where, for instance, uncertainties in friction coefficients, contact states, and visual perception may be large. This issue is exemplified in Fig. 4B, where baseline-to-1 was used to guide ANYmal over a flat floor with an invisible gap. When the right front foot stepped on the trap, the planned and executed motions deviated from each other. This triggered a sequence of heuristic recovery strategies. For large mismatches, however, such scripted reflexes were not effective and resulted in failure. DTC uses the same high-level planner but incorporates

learned recovery and reflex skills. This allowed the robot to successfully navigate through the trap. The robustness is rooted in the ability to ignore both perception and reference motion while relying only on proprioception. Such behavior was learned in simulation by experiencing simulated map drift. The experiment was repeated five times with baseline-to-1, five times with baseline-to-2, and five times with our method, consistently leading to similar results. The video clips corresponding to the above experiments can be found in movie S3. The movie is further enriched with a comparison of baseline-to-2 against DTC on soft materials, which impose very similar challenges.

Benchmark against RL control

Although RL policies are known for their robustness, they may struggle in environments with limited interaction points. We demonstrate typical failure cases in two experiments using baseline-rl-1. In the first experiment (Fig. 5A), ANYmal was tasked to cross a small gap of 0.1 m with a reference heading velocity of 0.2 m/s. The model-free controller did not avoid the gap and thus could not reach the other side of the platform. In the second experiment, we connected two elevated boxes with a 1.0-m-long beam of height of 0.2 m (Fig. 5B). The robot was commanded to walk from the left to the right box but failed to make use of the beam.

In comparison, our hybrid policy achieved a 100% success rate for the same gap size over 10 repetitions. To further demonstrate the locomotion skills of DTC, we made the experiments more challenging. We replaced the small gap with four larger gaps, each 0.6 m wide and evenly distributed along the path (Fig. 5C). Similarly, we increased the length of the beam to a total of 1.8 m (Fig. 5D). Despite the increased difficulty, our approach maintained a 100% success rate across four repetitions of each experiment. Video clips of those experiments can be found in movie S4.

By using a specialized policy, ANYmal crossed a 0.6-m-wide gap within a premapped environment (14). Most notably, our locomotion controller, not being specialized or fine-tuned for this terrain type, crossed a sequence of four gaps with the same width while relying on online generated maps only.

The limitations of baseline-rl-1 were previously demonstrated (7) on the obstacle parkour of Fig. 4A, showing its inability to cross the stepping stones. We showcase the generality of our proposed control framework by conducting three experiments on stepping stones in the real world, each with an increased level of difficulty. The first



A moving plane**B** slippery, rolling, and deformable objects**C** walking blind upstairs

Fig. 2. Evaluation of robustness. (A) ANYmal walks along a loose cover plate that eventually pitches forward (left to right, top to bottom). The third row shows ANYmal's perception of the surroundings during the transition and recovery phase. (B) The snapshots were taken at critical time instances when walking on slippery ground, just before complete recovery. (C) ANYmal climbs up stairs with disabled perception (top to bottom). The collision of the right-front end effector with the stair tread triggers a swing reflex, visualized in orange.

experiment (Fig. 6A) required ANYmal to traverse a field of equally sized stepping stones, providing a contact surface of 0.2 m by 0.2 m each. The robot passed the 2.0-m-long field 10 times. Despite the varying heading velocity commands, the robot accurately hit the correct stepping stones, as indicated by the solution of the TO. For the second experiment (Fig. 6B), we increased the height of two randomly selected stones. The parkour was successfully crossed four out of four times. In the final experiment (Fig. 6C), we distributed three elevated platforms, *a*, *b*, and *c*, connected by loose wooden blocks of sizes 0.31 m by 0.2 m by 0.2 m and 0.51 m by 0.2 m by 0.2 m.

This environment posed considerable challenges because the blocks may move and flip over when stepped on. Following the path $a \rightarrow b \rightarrow a \rightarrow b \rightarrow c \rightarrow a$, the robot missed only one stepping stone, which, however, did not lead to failure. Video clips of the stepping stone experiments are provided in movie S5.

Simulation-based ablation study

During training, we computed a solution to the TO problem after variable time intervals but mainly after each foot touchdown. Although such a throttled rate greatly reduces computational costs, it



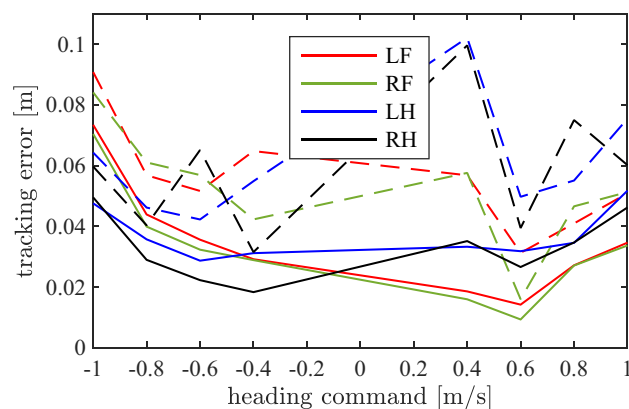
A artistic motion**B** foothold tracking error**C** foothold tracking error with baseline-to-2**D** Deployment with baseline-to-2

Fig. 3. Evaluation of tracking performance. (A) ANYmal climbs up a narrow table, turns, and descends back down to a box. The second image in the second row shows the robot's perception of the environment. (B) Euclidean norm of the planar foothold error, averaged over 20 s of operation using a constant heading velocity. The solid/dashed curves represent the average/maximum tracking errors. (C) Same representation as in (B), but the data were collected with baseline-to-2. (D) DTC deployed with baseline-to-2, enabling ANYmal to climb up a box of 0.48 m.

also leads to poor reactive behavior in the presence of quickly changing external disturbances, dynamic obstacles, or map occlusions. Moreover, the optimizer was updated using privileged observations, whereas, in reality, the optimizer is subject to elevation map drift, wrongly estimated friction coefficients, and unpredicted external forces. To compensate for such modeling errors, we deploy the optimizer in MPC fashion. We investigated the locomotion performance as a function of the optimizer update rate. Using the experimental setup outlined in Supplementary Methods (“Experimental setup for evaluation of optimizer rate” section), we collected a total of 6 days of

data in simulation. A robot was deemed successful if it could walk from the center to the border of its assigned terrain patch, failed if its torso made contact with the environment within its patch, and stuck otherwise. We report success and failure rates in Fig. 7A. Accordingly, when increasing the update rate from 1 to 50 Hz, the failure rate dropped by 7.11%, whereas the success rate increased by 4.25%.

In the second set of experiments, we compared our approach against baseline-rl-2 and against the same policy trained within our training environment. We refer to the emerging policy as baseline-rl-3. More details regarding the experimental setup can be found in


A obstacle parkour simulation (ours)**B** collapsing floor (baseline-to-1)**C** collapsing floor (ours)

Fig. 4. Benchmarking against model-based control. (A) DTC successfully traverses an obstacle parkour (left to right) in simulation with a heading velocity of 1 m/s. (B) Baseline-to-1 falls after stepping into a gap hidden from the perception (left to right). (C) ANYmal successfully overcomes a trapped floor using our hybrid control architecture (left to right).

Supplementary Methods (“Experimental setup for performance evaluation” section). As depicted in Fig. 7Bi, our approach exhibited a substantially higher success rate than baseline-rl-2. By learning on the same terrains, baseline-rl-3 could catch up but still did not match our performance in terms of success rate. The difference mainly originates from the fact that the retrained baseline still failed to solve sparse-structured terrains. To highlight this observation, we evaluated the performance on four terrain types with sparse (stepping stones, beams, gaps, and pallets) and four types with dense (stairs, pit, rough slope, and rings) stepping locations. On all considered terrain types, our approach outperformed baseline-rl-2 by a huge margin (Fig. 7B, ii), thereby demonstrating that learned locomotion generally does not extrapolate well to unseen scenarios. Our approach performed equally well as baseline-rl-3 on dense terrains but scored notably higher on sparse-structured terrains. This result suggests that the proposed approach itself was effective and that favorable locomotion skills were not encouraged by the specific training environment.

In an additional experiment, we investigated the effects of erroneous predictions of the high-level planner on locomotion performance. We did so by adding a drift value to the elevation map, sampled uniformly from the interval $\in(0, 0.5)$ m. Contrary to training, the motion was optimized over the perturbed height map. Other changes to the experimental setup are described in Supplementary Methods (“Experimental setup for performance evaluation under drift” section). As visualized in Fig. 7C, we collected tracking error, success, and failure rates with simulated drift on flat and rough ground. The tracking error grew mostly linearly with the drift value (Fig. 7C, i). On rough terrains, the success rate remained constant for drift values smaller than 0.1 m and decreased linearly for larger values (Fig. 7C, ii). On the

other hand, success and failure rates were not affected by drift on flat ground (Fig. 7C, iii).

We found that providing joint positions computed for the upcoming touchdown event greatly improved convergence time and foothold tracking performance. This signal encodes the foothold location in joint space, thus providing a useful hint for foothold tracking. It also simplifies the learning process because the network is no longer required to implicitly learn the solution of the inverse kinematics (IK). Evidence for our claims is given in Fig. 7D, showing two relevant learning curves. Tracking accuracy is represented by the foothold rewards, whereas technical skills are quantified using the average terrain level (11). Both scores are substantially higher if the footholds could be observed in both task and joint space.

DISCUSSION

This work demonstrates the potential of a hybrid locomotion pipeline that combines accurate foot placement and dynamic agility of state-of-the-art TO with the inherent robustness and reflex behaviors of novel RL control strategies. Our approach enables legged robots to overcome complex environments that either method alone would struggle with. Because such terrains are commonly found in construction sites, mines, and collapsed buildings, our work could help advance the deployment of autonomous legged machines in the fields of construction, maintenance, and search and rescue.

We have rigorously evaluated the performance in extensive real-world experiments over the course of about half a year. We included gaps, stepping stones, narrow beams, and tall boxes in our tests and demonstrated that our method outperformed the RL baseline

controller on every single terrain. Next, we evaluated the robustness on slippery and soft ground, each time outperforming two model-based controllers.

Furthermore, we have shown that the emerging policy can track the motion of two different planners using the same trotting gait. This was possible because the observed footholds seem to be mostly invariant under the choice of the optimizer. However, certain obstacles may encourage the deployed planner to produce footprint patterns that otherwise do not emerge during training. In this case, we would expect a degraded tracking performance.

In addition to our main contribution, we have demonstrated several other notable results. First, our policy, which was trained exclusively with visual perception, is still able to generalize to blind locomotion. Second, a simple multilayer perceptron (MLP) trained with an asymmetric actor/critic setup achieves similar robust behaviors as much more complex teacher/student trainings (12, 13).

Third, our locomotion policy can handle a lot of noise and drift in the visual data without relying on complicated gaited networks, which might be difficult to tune and train (13).

Contrary to our expectations, the proposed training environment was not found to be more sample efficient than similar unifying RL approaches (11, 13). The large number of epochs required for convergence suggests that foothold accuracy is something intrinsically complicated to learn.

In this work, we emphasized that TO and RL share complementary properties and that no single best method exists to address the open challenges in legged locomotion. The proposed control architecture leverages this observation by combining the planning capabilities of the former and the robustness properties of the latter. It does not, by any means, constitute a universal recipe to integrate the two approaches in an optimal way for a generic problem. Moreover, one could even extend the discussion with self- and unsupervised learning, indirect optimal control, dynamic programming, and stochastic optimal control. Nevertheless, our results may motivate future research to incorporate the aspect of planning into the concept RL.

We see several promising avenues for future research. Many successful data-driven controllers have the ability to alter the stride duration of the trotting gait. We expect a further increase in survival rate and technical skills if the network policy could suggest an arbitrary contact schedule to the motion optimizer. Moreover, a truly hybrid method, in which the policy can directly modify the cost function of the planner, may be able to generate more diversified motions. Our results indicate that IK is difficult to learn. To increase the sample efficiency and improve generalization across different platforms, a more sophisticated network structure could exploit prior knowledge of analytical IK. Another potential research direction may focus on leveraging the benefits of sampling trajectories from an offline buffer. This could substantially reduce the training time and allow for the substitution of TAMOLS with a more accurate TO method or even expert data gathered from real animals.

MATERIALS AND METHODS

Motivation

To motivate the specific architectural design, we first identify the strengths and weaknesses of the two most commonly used control paradigms in legged locomotion. TO amounts to open-loop control, which produces suboptimal solutions in the presence of stochasticity, modeling errors, and small prediction windows.

A small gap (baseline-rl-1)



B short beam (baseline-rl-1)



C sequence of large gaps (ours)



D long beam (ours)



Fig. 5. Benchmarking against reinforcement learning. (A) Baseline-rl-1 attempts to cross a small gap. ANYmal initially manages to recover from miss-stepping with its front legs but subsequently gets stuck when its hind legs fall inside the gap. (B) Using baseline-rl-1, the robot stumbles along a narrow beam. (C) With DTC, the robot can pass four consecutive large gaps (left to right) without getting stuck or falling. (D) ANYmal is crossing a long beam using the proposed control framework.

Unfortunately, these methods also introduce many assumptions, mostly to reduce computation time or to achieve favorable numerical properties. For instance, the feet are almost always preselected interaction points to prevent complex collision constraints, contact and actuator dynamics are usually omitted or smoothed out to circumvent stiff optimization problems, and the contact schedule is often pre-specified to avoid the combinatorial problem imposed by the switched system dynamics. Despite a large set of strong assumptions, real-time capable planners are not always truly real time. The reference trajectories are updated around 5 (31) to 100 (7) Hz and realized between 400 and 1000 Hz. In other words, these methods do not plan fast enough to catch up with the errors they are making. Although structural (2) or environmental (7, 20) decomposition may further contribute to the overall suboptimality, they were found to be useful for extracting good local solutions on sparse terrains. Because the concept of planning is not restricted to the tuning domain, model-based approaches tend to generalize well across different terrain geometries (7, 21).

Moreover, because numerical solvers perform very cheap and sparse operations on the elevation map, the map resolution can be arbitrarily small, facilitating accurate foothold planning.

RL, on the other hand, leads to policies that represent global closed-loop control strategies. Deep neural networks are large capacity models and, as such, can represent locomotion policies without introducing any assumption about the terrain or the system (except from being Markovian). They exhibit good interpolation between visited states but do not extrapolate well to unseen environments. Despite the large size of these networks, the inference time is usually relatively small. The integration of an actuator model has been demonstrated to improve sim-to-real transfer (10), and the stochasticity in the system dynamics and training environment can effectively be used to synthesize robust behaviors (12, 13). Contrary to TO, the elevation map is typically down-sampled (11, 13) to avoid immense memory consumption during training.

In summary, TO might be better suited if good generalization and high accuracy are required, whereas RL is the preferred method if robustness is of concern or onboard computational power is limited. Because locomotion combines challenges from both of these fields, we formulate the goal of this work as follows: RL shall be used to train a low-level tracking controller that provides markedly more robustness than classical inverse dynamics. The accuracy and planning capabilities of model-based TO shall be leveraged on a low level to synthesize a unifying locomotion strategy that supports diverse and generalizing motions.

Reference motions

Designing a TO problem for control always involves a compromise that trades off physical accuracy and generalization against good numerical conditioning, low computation time, convexity, smoothness, availability of derivatives, and the necessity of a high-quality initial guess. In our work, we generate the trajectories using TAMOLS (21). Unlike other similar methods, it does not require terrain segmentation nor precomputation of footholds, and its solutions are robust under varying initial guesses. The system dynamics and kinematics are simplified, allowing for fast updates. During deployment, we also compare against baseline-to-2, which builds on more complex kinodynamics. Because of the increased computation time and, particularly, the computationally demanding map-processing pipeline, this method is not well suited to be used directly within the learning process (the training time would be expected to be about eight times longer).

We added three crucial features to TAMOLS: First, we enabled parallelization on a central processing unit (CPU), which allows multiple optimization problems to be solved simultaneously. Second, we created a python interface using pybind11 (35), enabling it to run in a python-based environment. Last, we assumed that the measured contact state always matches the desired contact state. This renders the TO independent of contact estimation, which typically is the most fragile module in a model-based controller.

The optimizer requires a discretized 2.5-dimensional representation of its environment, a so-called elevation map, as input. We extract the map directly from the simulator by sampling the height across a fixed grid. For both training and deployment, we use a fixed trotting gait with a stride duration of 0.93 s and swing phase of 0.465 s and set the resolution of the grid map to 0.04 m by 0.04 m.

Overview of the training environment

The locomotion policy $\pi(\mathbf{a} | \mathbf{o})$ is a stochastic distribution of actions $\mathbf{a} \in \mathcal{A}$ that are conditioned on observations $\mathbf{o} \in \mathcal{O}$, parametrized by

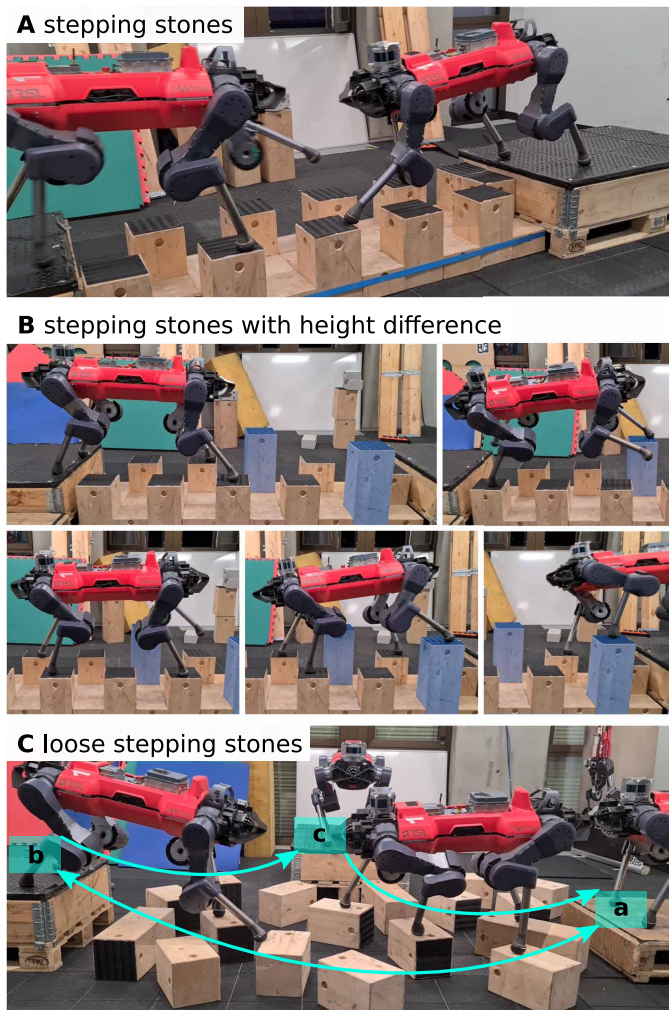


Fig. 6. Evaluation of the locomotion performance on stepping stones. (A) ANYmal reliably crosses a field of flat stepping stones (left to right). (B) The robot crosses stepping stones of varying heights (left to right). The two tall blocks are highlighted in blue. (C) ANYmal navigates through a field of loosely connected stepping stones, following the path $a \rightarrow b \rightarrow a \rightarrow b \rightarrow c \rightarrow a$.

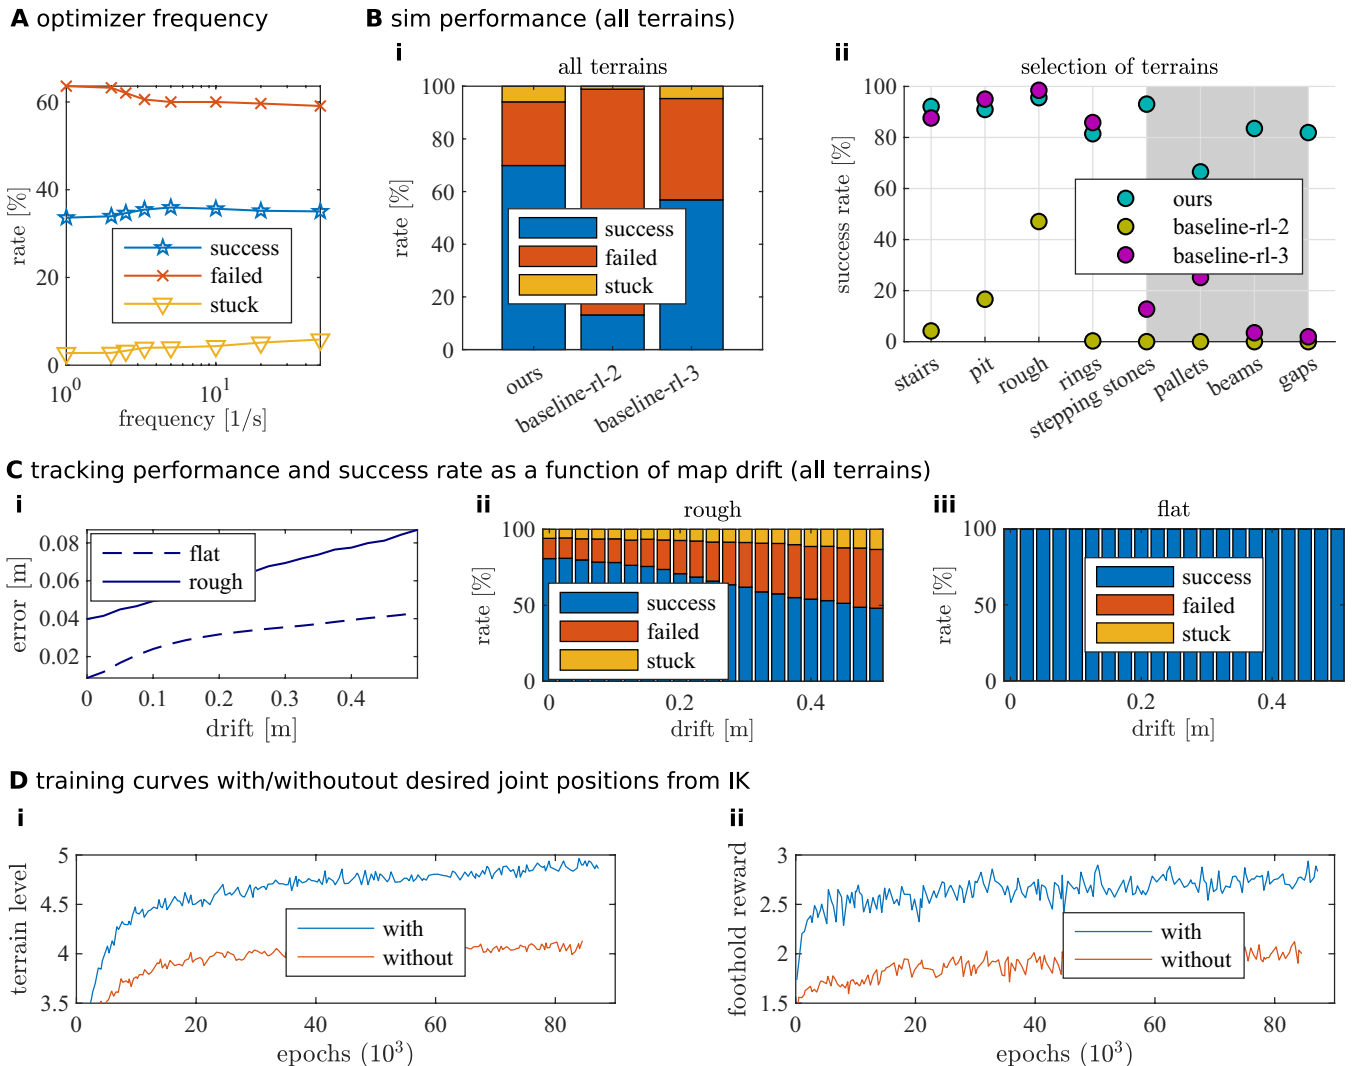


Fig. 7. Simulation results and ablation studies. (A) Success and failure rates of DTC, recorded for different update rates of the optimizer. The upper limit of 50 Hz is imposed by the policy frequency. (B) Comparison against baseline policies. (i) Evaluation on all 120 terrains. (ii) Evaluation on terrains where valid footholds are dense (white background) and sparse (gray background). (C) Influence of elevation map drift on the locomotion performance, quantified by tracking error (i), and success rate on rough (ii) and on flat ground (iii). (D) Average terrain level (i) and average foothold reward (ii) scored during training.

an MLP. The action space comprises target joint positions that are tracked using a proportional-derivative (PD) controller, following the approach in (10) and related works (12–14).

Given the state $\mathbf{s} \in \mathcal{S}$, we extract the solution at the next time step $\mathbf{x}'(\mathbf{s}) \in \mathcal{X} \subseteq \mathcal{S}$ from the optimizer, which includes four footholds $\mathbf{p}_{i=0,\dots,3}^*$, joint positions \mathbf{q}^* at touchdown time, and the base trajectory evaluated at the next time step. The base trajectory consists of base pose $\mathbf{b}^*(\Delta t)$, twist $\dot{\mathbf{b}}^*(\Delta t)$, and linear and angular acceleration $\ddot{\mathbf{b}}^*(\Delta t)$. More details can be found in Fig. 8A. We then sample an action from the policy. It is used to forward simulate the system dynamics, yielding a new state $\mathbf{s}' \in \mathcal{S}$, as illustrated in Fig. 8B.

To define a scalar reward $r(\mathbf{s}, \mathbf{s}', \mathbf{x}', \mathbf{a})$, we use a monotonically decreasing function of the error between the optimized and measured states, that is, $r \propto \mathbf{x}'(\mathbf{s}) \ominus \mathbf{x}(\mathbf{s}')$. The minus operator \ominus is defined on the set \mathcal{X} , the vector $\mathbf{x}'(\mathbf{s})$ is the optimized state, and $\mathbf{x}(\mathbf{s}')$

is the state of the simulator after extracting it on the corresponding subset. The policy network can also be understood as a learned model reference adaptive controller with the optimizer being the reference model.

In this work, we use an asymmetrical actor/critic method for training. The value function approximation $V(\mathbf{o}, \bar{\mathbf{o}})$ uses privileged $\bar{\mathbf{o}} \in \bar{\mathcal{O}}$ and policy observations \mathbf{o} .

Observation space

The value function is trained on policy observations and privileged observations, whereas the policy network is trained on the former only (22). All observations are given in the robot-centric base frame. The definition of the observation vector is given below, whereas noise distributions and dimensionalities of the observation vectors can be found in Supplementary Methods and table S2.

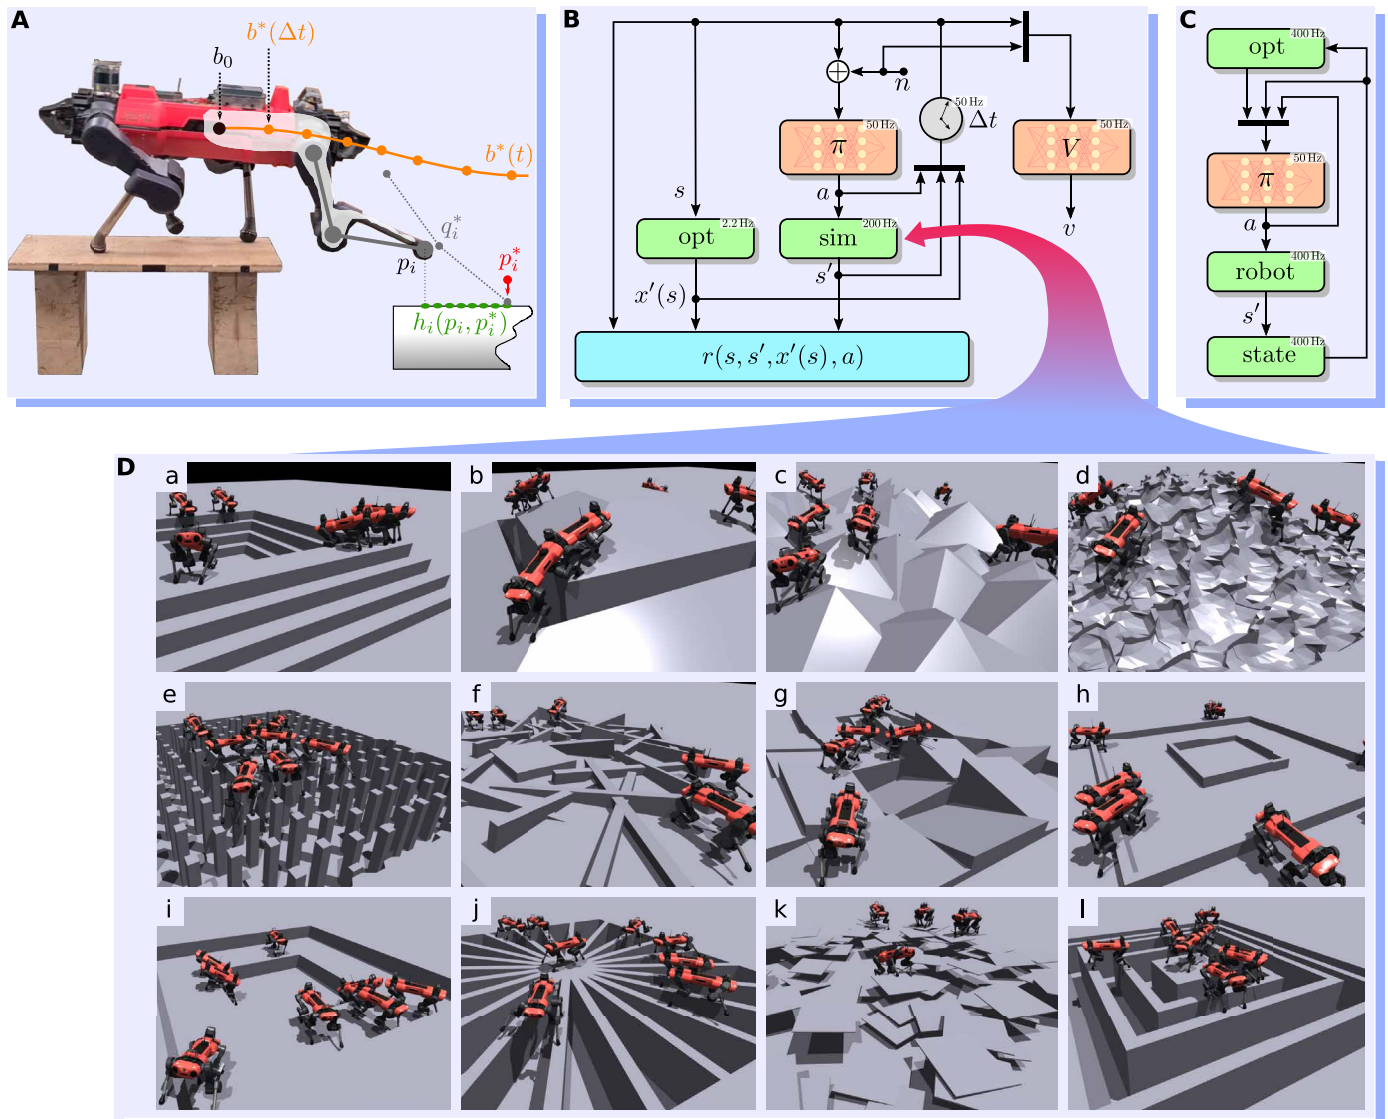


Fig. 8. Overview of the training method and deployment strategy. (A) The optimized solution provides footholds p_i^* , desired base pose b^* , twist b^* , and acceleration \ddot{b}^* (extracted one policy step Δt ahead), as well as desired joint positions q^* . In addition, a height scan h is sampled between the foot position p_i and the corresponding foothold. (B) Training environment: The optimizer runs in parallel to the simulation. At each leg touchdown, a new solution x' is generated. The policy π drives the system response s' toward the optimized solution $x'(s)$, which is encouraged using the reward function r . Actor observations are perturbed with the noise vector n , whereas critics and the TO receive ground truth data. (C) Deployment: Given the optimized footholds, the network computes target joint positions that are tracked using a PD control law. The state estimator (state) returns the estimated robot state, which is fed back into the policy and the optimizer. (D) The list of terrain types includes (a) stairs, (b) combinations of slopes and gaps, (c) pyramids, (d) sloped rough terrain, (e) stepping stones, (f) objects with randomized poses, (g) boxes with tilted surfaces, (h) rings, (i) pits, (j) beams, (k) hovering objects with randomized poses, and (l) pallets.

Policy observations

The policy observations comprise proprioceptive measurements such as base twist, gravity vector, joint positions, and joint velocities. The history only includes previous actions (11). Additional observations are extracted from the model-based planner, including planar coordinates of foothold positions (xy coordinates), desired joint positions at touchdown time, desired contact state, and time left in the current phase. The latter two are per-leg quantities that fully describe the gait pattern. Footholds only contain planar coordinates because the height can be extracted from the height scan.

The height scan, which is an additional part of the observation space, enables the network to anticipate a collision-free swing leg trajectory. In contrast to similar works, we do not construct a sparse elevation map around the base (11, 27) or the feet (13). Instead, we sample along a line connecting the current foot position with the desired foothold (Fig. 8A). This approach has several advantages: First, the samples are denser by only scanning terrain patches that are most relevant for the swing leg. Second, it prevents the network from extracting other information from the map, which is typically exposed to the most uncertainty (for instance, occlusion, reflection, odometry drift, discretization error, etc.). Third, it allows us to

conveniently model elevation map drift as a per-foot quantity, which means that each leg can have its own drift value.

We use analytical IK to compute the desired joint positions. Because the motion optimizer may not provide a swing trajectory, as is the case for TAMOLS, we completely skip the swing phase. This means that the IK is computed with the desired base pose and the measured foot position for a stance leg and the target foothold for a swing leg.

We do not provide the base pose reference as observation. This was found to reduce sensitivity to mapping errors and to render the policy independent of the used planner. Last, to allow the network to infer the desired walking direction, we add the reference twist (before optimization) to the observation space.

Privileged observations

The privileged observations contain the optimized base pose, base twist, and base linear and angular acceleration, extracted one time step ahead. In addition, the critics can observe signals confined to the simulator, such as the external base wrench, external foot forces, the measured contact forces, friction coefficients, and elevation map drift.

Reward functions

The total reward is computed as a weighted combination of several individual components, which can be categorized as follows: “tracking” of reference motions, encouraging “consistent” behavior, and other “regularization” terms necessary for successful sim-to-real transfer. The reward functions are explained below, whereas weights and parameters are reported in table S3.

Base pose tracking

To achieve tracking of the reference base pose trajectory, we use

$$r_{Bn} = e^{-\sigma_{Bn} \cdot \|\mathbf{b}^*(t+\Delta t) \ominus \mathbf{b}(t)\|^{2n}} \quad (1)$$

where $n = \{0,1,2\}$ is the derivative order, $\mathbf{b}(t)$ is the measured base pose, $\mathbf{b}^*(t + \Delta t)$ is the desired base pose sampled from the reference trajectory one policy step Δt ahead, and \ominus denotes the quaternion difference for base orientation and the vector difference otherwise. We refer to the above reward function as a “soft” tracking task because large values can be scored even if the tracking error does not perfectly vanish.

To further analyze the reward function, we decompose the base trajectory into three segments. The “head” starts at time zero, the “tail” stops at the prediction horizon, and the “middle” connects these two segments with each other. A logarithmic reward function would prioritize the tracking of the trajectory head, whereas a linear penalty would focus on making progress along the whole trajectory at once. Contrarily, the exponential shape of the reward function splits the tracking task into several steps. During the initial epochs, the tracking error of the trajectory middle and tail will likely be relatively large and thus does not contribute notably to the reward gradient. As a result, the network will minimize the tracking error of the trajectory head. Once its effect on the gradient diminishes, the errors corresponding to the trajectory middle will dominate the gradient landscape. In the final training stages, tracking is mostly improved around the trajectory tail.

Foothold tracking

We choose a logarithmic function

$$r_{pi} = -\ln(\|\mathbf{p}_i^* - \mathbf{p}_i\|^2 + \epsilon) \quad (2)$$

to learn foothold tracking, where \mathbf{p}_i is the current foot position of leg $i \in \{0, \dots, 3\}$, \mathbf{p}_i^* is the corresponding desired foothold, and $0 < \epsilon \ll 1$ is a small number ensuring the function is well defined. The above reward function may be termed a “hard” tracking task because the maximum value can only be scored if the error reaches zero. As the tracking improves, the gradients will become larger, resulting in even tighter tracking toward the later training stages.

A dense reward structure typically encourages a stance foot to be dragged along the ground to further minimize the tracking error. To prevent such drag motions from emerging, the above reward is given for each foot at most once during one complete gait cycle: More specifically, if and only if the leg is intended to be in contact and the norm of the contact force indicates a contact; that is, if $\|\mathbf{f}_i\| > 1 \text{ N}$, then the agent receives the reward.

Consistency

In RL for legged locomotion, hesitating to move over challenging terrains is a commonly observed phenomenon that prevents informative samples from being gathered and thus impedes the agent’s performance. This behavior can be explained by insufficient exploration: The majority of agents fail to solve a task, whereas a small number of agents achieve higher average rewards by refusing to act. To overcome this local optimum, we propose to encourage consistency by rewarding actions that are intended by previous actions. In our case, we measure consistency as the similarity between two consecutive motion optimizations. If the solutions are similar, then the agent is considered to be consistent. We measure similarity as the Euclidean distance between two adjacent solutions and write

$$r_c = \sum_{\delta t_j + t_0 \in (T_a \cap T_b)} -\delta t \|\mathbf{b}_a^*(\delta t_j + t_{0,a}) \ominus \mathbf{b}_b^*(\delta t_j + t_{0,b})\| - w_p \|\mathbf{p}_a^* - \mathbf{p}_b^*\| \quad (3)$$

Here, \mathbf{p}_i^* with $t = \{a, b\}$ is a vector of stacked footholds, $w_p > 0$ is a relative weight, $\delta t = 0.01 \text{ s}$ is the discretization time of the base trajectory, and t_0 is the time elapsed since the solution was retrieved. The index a refers to the most recent solution, whereas b refers to the previous solution. The two solution vectors \mathbf{x}_a and \mathbf{x}_b , from which we extract the base and footholds, are only defined on their respective time intervals given by the optimization horizon τ_h , that is, $t_a \in T_a = [0, \tau_{h,a}]$ and $t_b \in T_b = [0, \tau_{h,b}]$.

Regularization

To ensure that the robot walks smoothly, we use two different penalty terms enforcing complementary constraints. The first term, $r_{r1} = -\sum_i |\mathbf{v}_i^T \mathbf{f}_i|$, discourages foot scuffing and end-effector collisions by penalizing power measured at the feet. The second term, $r_{r2} = -\sum_i (\dot{\mathbf{q}}_i^T \boldsymbol{\tau}_i)^2$, penalizes joint power to prevent arbitrary motions, especially during the swing phase. Other regularization terms are stated in Supplementary Methods (“Implementation details” section).

Training environment

To train the locomotion policy, we use a custom version of Proximal Policy Optimization (PPO) (36) and a training environment that is mostly identical to that introduced in (11). It is explained in more detail in Supplementary Methods (“Training details” section) and table S1. Simulation and back propagation were performed on graphics processing unit (GPU), whereas the optimization problems were solved on CPU.

Termination

We use a simple termination condition where an episode is terminated if the base of the robot makes contact with the terrain.

Domain randomization

We inject noise into all observations except for those designated as privileged. At each policy step, a noise vector \mathbf{n} is sampled from a uniform distribution and added to the observation vector, with the only exceptions of the desired joint positions and the height scan.

For the elevation map, we add noise before extracting the height scan. The noise is sampled from an approximate Laplace distribution where large values are less common than small ones. We perturb the height scan with a constant offset, which is sampled from another approximate Laplace distribution for each foot separately. Both perturbations discourage the network to rely extensively on perceptive feedback and help to generalize to various perceptive uncertainties caused by odometry drift, occlusion, and soft ground.

All robots are artificially pushed by adding a twist offset to the measured twist at regular time instances. Friction coefficients are randomized per leg once at initialization time. To render the motion robust against disturbances, we perturb the base with an external wrench and the feet with external forces. The latter slightly stiffens up the swing motion but improves tracking performance in the presence of unmodeled joint frictions and link inertia. The reference twist is resampled in constant time intervals and then held constant.

The solutions for the TO problems are obtained using ground-truth data, which include the true friction coefficients, the true external base wrench, and noise-free height map. In the presence of simulated noise, drift, and external disturbances, the policy network is therefore trained to reconstruct a base trajectory that the optimizer would produce given the ground truth data. However, there is a risk that the network learns to remove the drift from the height scan by analyzing the desired joint positions. During hardware deployment, such a reconstruction will fail because the optimizer is subject to the same height drift. To mitigate this issue, we introduce noise to the desired joint position observations, sampled from a uniform distribution with boundaries proportional to the drift value.

Terrain curriculum

We use a terrain curriculum as introduced in (11). Before the training process, terrain patches of varying types and difficulties are generated, and each agent is assigned a terrain patch. As an agent acquires more skills and can navigate the current terrain, its level is upgraded, which means it will be respawned on the same terrain type but with a harder difficulty. We have observed that the variety of terrains encountered during training influences the sim-to-real transfer. We thus have included a total of 12 different terrain types with configurable parameters (Fig. 8D), leading to a total of 120 distinguishable terrain patches. The terrain types classify different locomotion behaviors, such as climbing (“stairs,” “pits,” “boxes,” and “pyramids”), reflexing (“rough,” “rings,” and “flying objects”), and walking with large steps (“gaps,” “pallets,” “stepping stones,” “beams,” and “objects with randomized poses”). Our terrain curriculum consists of 10 levels, where one of the configurable parameters is modulated to increase or decrease its difficulty. This results in a total of 1200 terrain patches, each with a size of 8 m by 8 m, summing to a total area of 76,800 m², which is approximately the size of 14 football fields or 10 soccer fields.

Training

Solving the TO problems at the policy frequency during training was found to provoke poor local optima. In such a case, the optimizer

adapts the solution after each policy step: If the agent is not able to follow the reference trajectory, the optimizer will adapt to the new state such that the tracking problem becomes feasible again. This means that the agent can exhibit “lazy” behavior and still collect some rewards. We prevent such a local optimum by updating the optimizer only at a leg touchdown (after 0.465 s). This also greatly reduces learning time because computational costs are reduced by a factor of 23. After a robot fell (on average, once every 18 s) or was pushed (after 10 s) or its twist commands changed (three times per episode), the optimized trajectories were no longer valid. To guarantee that the locomotion policy generalizes across different update rates, we additionally recompute the solution in all those scenarios.

We trained the policy with a massive parallelization of $64^2 = 4096$ robots for a total of 90,000 epochs. Each epoch consisted of 45 learning iterations, where each iteration covered a duration of 0.02 s. Considering the variable update rate explained previously, this resulted in a total of 8295 days (or 23 years) of optimized trajectories. The policy can be deployed after about 1 day of training (6000 epochs), reaches 90% of its peak performance after 3 days (20,000 epochs), and is fully converged after 2 weeks (90,000 epochs).

In comparison, the baseline-rl-1 policy was trained for 40,000 epochs with 1000 parallelized robots over 5 consecutive days. Each epoch lasted for 5 s, resulting in a throughput of 463 simulated seconds per second. Our policy was trained for 14 days, with each epoch lasting for 0.9 s, leading to a throughput of 274 simulated seconds per second. Thus, despite generating 1.6 years of desired motions per day, our approach has only a 1.7 times lower throughput than the baseline.

Deployment

We deploy the policy at a frequency of 50 Hz without any fine-tuning. The motion optimizer runs at the largest possible rate in a separate thread. For TAMOLS with a trotting gait, this is around 400 Hz, and for baseline-to-2, this is around 100 Hz (both are faster than the policy frequency). At each step, the policy queries the most recent solution from the thread pool and extracts it, $\Delta t = 0.02$ s ahead of the most recent time index.

For our experiments, we used three different types of ANYmal robots (34), two version C and one version D, for which we trained different policies. ANYmal C is by default equipped with four Intel RealSense D435 depth cameras, whereas ANYmal D has eight depth cameras of the same type. For the second version C, the depth cameras were replaced with two identical Robosense Bpearl dome LiDAR sensors. For the outdoor experiments, we mostly used this robot because the Bpearls tend to be more robust against lighting conditions. Motion optimization and the forward propagation of the network policy are done on a single Intel Core i7-8850H machine. Elevation mapping (37) runs on a dedicated onboard Nvidia Jetson.

Supplementary Materials

This PDF file includes:

Supplementary Methods
Tables S1 to S3

Other Supplementary Material for this manuscript includes the following:

Movies S1 to S5

REFERENCES AND NOTES

1. J. Z. Kolter, M. P. Rodgers, A. Y. Ng, A control architecture for quadruped locomotion over rough terrain, in *2008 IEEE International Conference on Robotics and Automation (IEEE, 2008)*, pp. 811–818.

2. M. Kalakrishnan, J. Buchli, P. Pastor, M. Mistry, S. Schaal, Fast, robust quadruped locomotion over challenging terrain, in *2010 IEEE International Conference on Robotics and Automation* (IEEE, 2010), pp. 2665–2670.
3. A. W. Winkler, C. D. Bellicoso, M. Hutter, J. Buchli, Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robot. Autom. Lett.* **3**, 1560–1567 (2018).
4. C. Mastalli, I. Havoutis, M. Focchi, D. G. Caldwell, C. Semini, Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping, and whole-body control. *IEEE Trans. Robot.* **3003464**, 1635–1648 (2020).
5. F. Jenelten, T. Miki, A. E. Vijayan, M. Bjelonic, M. Hutter, Perceptive locomotion in rough terrain – online foothold optimization. *IEEE Robot. Autom. Lett.* **5**, 5370–5376 (2020).
6. P. Fankhauser, M. Bjelonic, C. Dario Bellicoso, T. Miki, M. Hutter, Robust rough-terrain locomotion with a quadrupedal robot, in *2018 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2018), pp. 5761–5768.
7. R. Grandia, F. Jenelten, S. Yang, F. Farshidian, M. Hutter, Perceptive locomotion through nonlinear model-predictive control. *IEEE Trans. Robot.* **39**, 3402–3421 (2023).
8. C. Mastalli, W. Merkt, G. Xin, J. Shim, M. Mistry, I. Havoutis, S. Vijayakumar, Agile maneuvers in legged robots: a predictive control approach, arXiv:2203.07554v2 [cs.RO] (2022).
9. C. D. Bellicoso, F. Jenelten, C. Gehring, M. Hutter, Dynamic locomotion through online nonlinear motion optimization for quadrupedal robots. *IEEE Robot. Autom. Lett.* **3**, 2261–2268 (2018).
10. J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, M. Hutter, Learning agile and dynamic motor skills for legged robots. *Sci. Robot.* **4**, eaa5872 (2019).
11. N. Rudin, D. Hoeller, P. Reist, M. Hutter, Learning to walk in minutes using massively parallel deep reinforcement learning, in *5th Annual Conference on Robot Learning* (MLRPress, 2021), pp. 91–100.
12. J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, M. Hutter, Learning quadrupedal locomotion over challenging terrain. *Sci. Robot.* **5**, eabc5986 (2020).
13. T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, M. Hutter, Learning robust perceptive locomotion for quadrupedal robots in the wild. *Sci. Robot.* **7**, eabk2822 (2022).
14. N. Rudin, D. Hoeller, M. Bjelonic, M. Hutter, Advanced skills by learning locomotion and local navigation end-to-end, in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2022), pp. 2497–2503.
15. Z. Xie, H. Y. Ling, N. H. Kim, M. van de Panne, AllSteps: Curriculum-driven learning of stepping stone skills. *Comput. Graph. Forum* **39**, 213–224 (2020).
16. H. Duan, A. Malik, J. Dao, A. Saxena, K. Green, J. Siekmann, A. Fern, J. Hurst, Sim-to-real learning of footstep-constrained bipedal dynamic walking, in *2022 International Conference on Robotics and Automation (ICRA)* (IEEE, 2022), pp. 10428–10434.
17. W. Yu, D. Jain, A. Escontrela, A. Iscen, P. Xu, E. Coumans, S. Ha, J. Tan, T. Zhang, Visual-locomotion: Learning to walk on complex terrains with vision, in *Proceedings of the 5th Conference on Robot Learning*, A. Faust, D. Hsu, G. Neumann, Eds. (MLResearchPress, 2022), pp. 1291–1302.
18. A. Agarwal, A. Kumar, J. Malik, D. Pathak, Legged locomotion in challenging terrains using egocentric vision, in *6th Annual Conference on Robot Learning* (MLResearchPress, 2022), pp. 403–415.
19. K. Caluwaerts, A. Iscen, J. C. Kew, W. Yu, T. Zhang, D. Freeman, K.-H. Lee, L. Lee, S. Saliceti, V. Zhuang, N. Batchelor, S. Bohez, F. Casarini, J. E. Chen, O. Cortes, E. Coumans, A. Dostmohamed, G. Dulac-Arnold, A. Escontrela, E. Frey, R. Hafner, D. Jain, B. Jyenis, Y. Kuang, E. Lee, L. Luu, O. Nachum, K. Oslund, J. Powell, D. Reyes, F. Romano, F. Sadeghi, R. Sloat, B. Tabanpour, D. Zheng, M. Neunert, R. Hadsell, N. Heess, F. Nori, J. Seto, C. Parada, V. Sindhwani, V. Vanhoucke, J. Tan, Barkour: Benchmarking animal-level agility with quadruped robots arXiv:2305.14654 [cs.RO] (24 May 2023).
20. R. J. Griffin, G. Wiedebach, S. McCrory, S. Bertrand, I. Lee, J. Pratt, Footstep planning for autonomous walking over rough terrain, in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)* (IEEE, 2019), pp. 9–16.
21. F. Jenelten, R. Grandia, F. Farshidian, M. Hutter, Tamols: Terrain-aware motion optimization for legged systems, *IEEE Transactions on Robotics* (IEEE, 2022), pp.3395–3413.
22. P. Brakel, S. Bohez, L. Hasenclever, N. Heess, K. Bousmalis, Learning coordinated terrain-adaptive locomotion by imitating a centroidal dynamics planner, in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2022), pp. 10,335–10,342.
23. M. Bogdanovic, M. Khadiv, L. Righetti, Model-free reinforcement learning for robust locomotion using demonstrations from trajectory optimization. *Front. Robot. AI* **9**, 854212 (2022).
24. X. B. Peng, P. Abbeel, S. Levine, M. van de Panne, DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics* **37**, 1–14 (2018).
25. X. B. Peng, Z. Ma, P. Abbeel, S. Levine, A. Kanazawa, AMP: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics* **40**, 1–20 (2021).
26. S. Bohez, S. Tunyasuvunakool, P. Brakel, F. Sadeghi, L. Hasenclever, Y. Tassa, E. Parisotto, J. Humplik, T. Haarnoja, R. Hafner, M. Wulfmeier, M. Neunert, B. Moran, N. Siegel, A. Huber, F. Romano, N. Batchelor, F. Casarini, J. Merel, R. Hadsell, N. Heess, Imitate and repurpose: Learning reusable robot movement skills from human and animal behaviors arXiv:2203.17138 [cs.RO] (31 March 2022).
27. V. Tsounis, M. Alge, J. Lee, F. Farshidian, M. Hutter, Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning. *IEEE Robot. Autom. Lett.* **5**, 3699–3706 (2020).
28. X. B. Peng, G. Berseth, K. Yin, M. van de Panne, Deeploco: Dynamic locomotion skills using hierarchical deep reinforcement learning (ACM Transactions on Graphics, Proc. SIGGRAPH 2017) (2017), vol. 36.
29. O. Melon, M. Geisert, D. Surovik, I. Havoutis, M. Fallon, Reliable trajectories for dynamic quadrupeds using analytical costs and learned initializations, in *2020 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2020), pp. 1410–1416.
30. D. Surovik, O. Melon, M. Geisert, M. Fallon, I. Havoutis, Learning an expert skill-space for replanning dynamic quadruped locomotion over obstacles, in *Proceedings of the 2020 Conference on Robot Learning*, J. Kober, F. Ramos, C. Tomlin, eds. (MLResearchPress, 2021), pp. 1509–1518.
31. O. Melon, R. Orsolino, D. Surovik, M. Geisert, I. Havoutis, M. Fallon, Receding-horizon perceptive trajectory optimization for dynamic legged locomotion with learned initialization, in *2021 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2021), pp. 9805–9811.
32. S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, I. Havoutis, Rloc: Terrain-aware legged locomotion using reinforcement learning and optimal control. *IEEE Trans. Robot.* **38**, 2908–2927 (2022).
33. Z. Xie, X. Da, B. Babich, A. Garg, M. V. de Panne, Glide: Generalizable quadrupedal locomotion in diverse environments with a centroidal model, in *Algorithmic Foundations of Robotics XV*, S. M. LaValle, J. M. O’Kane, M. Otte, D. Sadigh, P. Tokekar, Eds. (Springer International Publishing, 2023), pp. 523–539.
34. M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, M. Hoepflinger, Hutter2016 – a highly mobile and dynamic quadrupedal robot, in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2016), pp. 38–44.
35. W. Jakob, J. Rhineland, D. Moldovan, pybind11–seamless operability between c++11 and python (2017); <https://github.com/pybind/pybind11>.
36. J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms, arXiv:1707.06347 [cs.LG] (28 August 2017).
37. T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, M. Hutter, Elevation mapping for locomotion and navigation using GPU, in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2022), pp. 2273–2280.
38. F. Jenelten, J. He, F. Farshidian, M. Hutter, Evaluation of tracking performance and robustness for a hybrid locomotion controller (Zenodo, 2023); <https://doi.org/10.5061/dryad.b5mkkwhqk>.

Funding: This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme grant agreement No 852044. This research was supported by the Swiss National Science Foundation (SNSF) as part of project no. 188596 and by the Swiss National Science Foundation through the National Centre of Competence in Research Robotics (NCCR Robotics). **Author contributions:** F.J. formulated the main ideas, trained and tested the policy using baseline-to-1, and conducted most of the experiments. J.H. interfaced baseline-to-2 with the tracking policy and conducted the box-climbing experiments. F.F. contributed to the theory and improved some of the original ideas. All authors helped to write, improve, and refine the paper. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** Datasets and code to generate all of our figures are made publicly available (38).

Submitted 21 April 2023
 Accepted 19 December 2023
 Published 17 January 2024
 10.1126/scirobotics.adh5401

DTC: Deep Tracking Control

Fabian Jenelten, Junzhe He, Farbod Farshidian, and Marco Hutter

Sci. Robot. **9** (86), eadh5401. DOI: 10.1126/scirobotics.adh5401

View the article online

<https://www.science.org/doi/10.1126/scirobotics.adh5401>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2024 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works