

# Learning to walk in confined spaces using 3D representation

Takahiro Miki, Joonho Lee, Lorenz Wellhausen and Marco Hutter

**Abstract**—Legged robots have the potential to traverse complex terrain and access confined spaces beyond the reach of traditional platforms thanks to their ability to carefully select footholds and flexibly adapt their body posture while walking. However, robust deployment in real-world applications is still an open challenge. In this paper, we present a method for legged locomotion control using reinforcement learning and 3D volumetric representations to enable robust and versatile locomotion in confined and unstructured environments. By employing a two-layer hierarchical policy structure, we exploit the capabilities of a highly robust low-level policy to follow 6D commands and a high-level policy to enable three-dimensional spatial awareness for navigating under overhanging obstacles. Our study includes the development of a procedural terrain generator to create diverse training environments. We present a series of experimental evaluations in both simulation and real-world settings, demonstrating the effectiveness of our approach in controlling a quadruped robot in confined, rough terrain. By achieving this, our work extends the applicability of legged robots to a broader range of scenarios.

## I. INTRODUCTION

Robots have attracted attention for their ability to perform a wide range of tasks, including exploring and accessing areas that are dangerous or inaccessible to humans. Among the different robotic platforms, legged robots are well suited for this type of task due to their ability to move on uneven, complex terrain, and to remain stable on unstable surfaces. In recent years, legged robots have been successfully deployed in a variety of challenging environments, from mountains, forests or simulated space environments [1]–[3] to underground spaces such as tunnels and caves [4]–[7]. However, their full potential has yet to be realized. One unique feature of legged robots is their ability to adjust their posture using its high degree of freedom, such as crouching. This provides capabilities to access confined area not easily replicated by other robotic platforms with a similar size. Despite this, the challenge of robustly deploying legged robots in confined or restricted spaces is still an open problem.

Some studies have shown that legged robots can navigate under low obstacles using a special crouching gait. However, these tests were conducted in simple environments such as on flat ground with a simple overhanging obstacle [8]–[11]. Buchanan et al. addressed more complex confined spaces including both steps and sloped overhanging obstacles

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 780883.

This research was supported by the Swiss National Science Foundation (SNSF) as part of project No.188596.

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101016970.

All authors are with Robotic Systems Lab, ETH Zurich



Fig. 1. Real-world experiment: Successful confined space traversal by the quadruped robot including a simulated collapsed building environment. The terrain consists of loose gravel or unstable steps, while the overhead structures have tilted configurations with narrow openings. The robot could adapt its posture to traverse these challenging conditions.

using trajectory optimization techniques for motion planning and a two-layer elevation map to perceive the environment, representing both the ground and the ceiling [12], [13]. However, the gait was limited to a static gait and it assumed that the map is accurate and the robot would always have clear space to move in the map, which can be a problem if the sensors make errors or provide unclear data.

To achieve a fast dynamic walking on various terrains while being robust against degraded perception, learning-based methods have recently demonstrated their effectiveness by integrating proprioceptive data with noisy exteroceptive data [2]. However, until now, the deployments are mostly limited to environments without overhanging obstacles [2], [14]–[17] or on a flat terrains with single obstacle [9], [10]. Our research aims to extend the operational range of legged robots to more complex environments in the wild including confined spaces while keeping the robust and smooth locomotion capability of the state of the art methods.

In this paper, we propose a two-layer hierarchical framework to address the problem. At the lower level, the policy negotiates ground-level obstacles while complying with extended high-level commands such as desired velocity, body height, and body orientation, while the higher level policy uses 3D geometric data to guide the actions of the lower level policy. This architecture allows us to benefit from the robustness of established locomotion methods, while introducing a new layer of spatial awareness through high-level policy formulation. To achieve spatial awareness of the confined environment, on-board sensors such as cameras, LiDARs, or depth cameras can be used to measure the surrounding geometry. However, they often exhibit degraded performance

in the complex environments due to factors such as physical obstructions, blind spots, and limited sensor ranges. In addition, the variety of sensor configurations makes it difficult to create a generalizable policy. To circumvent these challenges, we use a teacher-student training setup where we distill the teacher policy into a robust student policy which takes noisy occupancy voxels as our geometric representation. The voxel representation can integrate different sensor measurements into one data format, allowing a flexibility in different sensor configurations. In addition, the same policy can be used without retraining for the different sensor setup unlike the case of depth images representation. Furthermore, it naturally allows the incorporation of memory-based strategies, such as mapping to preserve information.

The contributions are as follows: First, we introduce a hierarchical policy framework that decomposes the confined space navigation into low-level and high-level tasks where the low-level policy inherits the performance of the robust perceptive locomotion on rough terrain while adding a new capability of entering confined space through high-level policy's posture adaptation. Second, we train the policy in a wide variety of environments by incorporating procedural terrain generator in simulation. Finally, we demonstrate the real-world applicability of our approach through deployments in complex scenarios in the field.

## II. RELATED WORKS

Environmental perception is essential for legged robots to smoothly navigate complex terrains. Early methods relied on pre-generated elevation maps for offline planning [18], [19]. This approach evolved to use on-board sensors to dynamically generate elevation maps in real time [20]–[22], which were subsequently used in optimization-based controllers [23]–[29]. A basic assumption of these methods is accurate terrain information. Therefore, their robustness degrades in the face of mapping inaccuracies, a problem often encountered in unstructured outdoor environments.

On the other hand, learning-based methods have gained momentum for locomotion tasks. Reinforcement learning is used to train a neural network policy in simulation and deployed on the hardware [1], [2], [30]–[34]. Recently exteroceptive information has been integrated into the locomotion control through elevation map as an intermediate representation to keep past information [2], [35], [36]. They sampled the height scan from the map to feed into a Recurrent Neural Network (RNN) based policy, however this 2.5D representation can not handle overhanging structures. As another intermediate representation, voxels were used for locomotion task [10], [37]. This can represent full 3D structures, however it was demonstrated only in a structured terrain with simple geometry. Other works have used depth images as input [14]–[16], [38]. In these approaches, depth images are first processed by a Convolutional Neural Network (CNN) and either combined with historical features or fed into a Recurrent Neural Network (RNN) to preserve past information. To extend the memory function of RNN-based methods using depth images, implicit voxel representations

have been introduced [17]. In this method, latent voxels from the previous frame are transformed to the current frame using a predicted transformation obtained from two depth images. However, the verification of these locomotion capabilities did not include diverse and confined spaces.

## III. METHODS

### A. Overview

Our approach uses a hierarchical framework as shown in Fig. 2. The training of the control policy is structured into four progressive stages: the low-level teacher policy, the low-level student policy, the high-level teacher policy, and the high-level student policy.

First, the low-level teacher policy is trained using reinforcement learning to follow high-level commands over varied, rough terrain. Using 6D inputs - including  $x$ ,  $y$ , and yaw velocities, roll, pitch, and body height - the policy acquires the ability to navigate smoothly over uneven surfaces while following given commands. Then the low-level student policy is distilled from the low-level teacher policy to handle noisy observations following [2].

We then train the high-level teacher policy, whose primary task is to generate 6 DOF commands for the low-level policy while going over rough terrain and avoiding overhanging obstacles in a procedurally generated terrains. This policy also employs reinforcement learning and uses spherical scans as shown in Fig. 2 to capture localized geometric data. These scans provide an omnidirectional view of the environment and inform the high-level policy's decision-making process.

Finally, we train a high-level student policy, which is intended for the robot's final deployment in the field. Unlike the high-level teacher policy, the student policy is trained using noisy observations simulated to resemble data from the robot's on-board sensors. Specifically, it employs noisy voxel data as a three-dimensional representation of the environment.

### B. Low-level locomotion policy

We extend the perceptive locomotion pipeline presented in [2] by adding three more command inputs: desired roll  $\xi$ , pitch  $\theta$ , and body height  $h$  to allow the robot to have more control over its motion. The low-level teacher policy is trained using Proximal Policy Optimization (PPO) algorithm [39], whereby the robot interacts with the environment and receives rewards based on its performance in following the given 6D commands on the rough terrain.

The command is randomly sampled at the beginning of each episode, ensuring that the robot is exposed to a variety of command scenarios during training. We sample roll and pitch command from a normal distribution with mean zero and standard deviation 0.25 rad and the desired body height from the ground is randomly sampled from an interval ranging from 0.1 to 0.6 above ground level. To avoid spikes in ground height values near edges, we use the average of five sampled ground heights to represent the body's height above the ground.

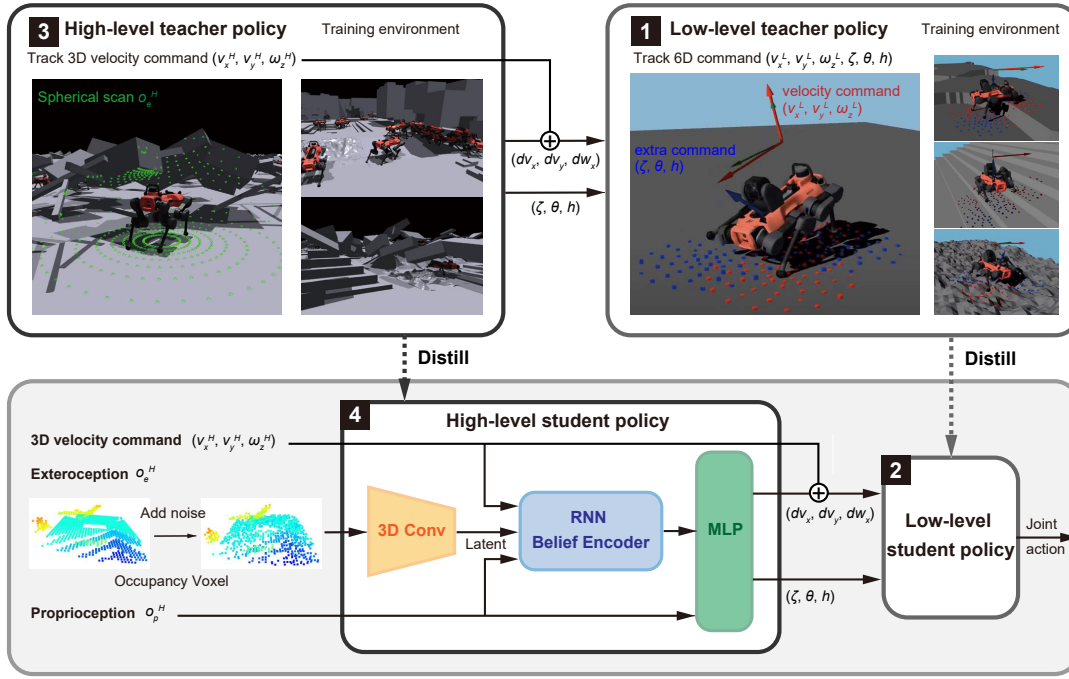


Fig. 2. Overview of our method. We use a two-layer policy setup. The low-level policy learns to walk over rough terrain while following 6D commands consisting of x,y, yaw rate and roll pitch and body height. The high-level policy is trained in a procedurally generated confined environment to guide the robot by giving commands to the low-level policy. We first train a low-level teacher policy and then distill it into a low-level student policy. We follow this by training a high-level teacher policy using spherical scans for exteroceptive perception. Finally, we distill this into a high-level student policy using noisy voxel grids as exteroceptive input.

The low-level teacher observation includes a 6D command, proprioceptive observation  $o_p^L$ , exteroceptive observation  $o_e^L$ , and privileged state  $s_p^L$ . The proprioceptive observation ( $o_p^L$ ) consists of the body velocity, orientation, history of joint position error and velocity that stack few frames between each control loop, action history, as well as the phase of each leg. The exteroceptive observation ( $o_e^L$ ) is represented by height samples taken around each foot with five distinct radii [2]. The privileged state  $s_p^L$  encompasses various factors such as contact states, contact forces, contact normals, friction coefficient, thigh and shank contact states, external forces and torques applied to the body, and swing phase duration. We adopt the same action space as that described in previous research [2] which consists of the phase difference  $\Delta\phi_l$  of the periodic motion generator and the residual joint position target  $\Delta q_i$  that is added on top of the periodic motion. To represent our teacher policy  $\pi_\theta$  we utilize a Multi-Layer Perceptron model.

We add two additional reward terms on top of the original rewards which consists of velocity tracking rewards and penalties such as joint torque, velocity, acceleration or slip rewards. Please refer to the previous research [2] for the detailed reward definitions for locomotion. The additional reward terms are employed for roll, pitch and body height tracking.

The orientation tracking reward is:

$$r_{\text{orientation}} = \exp(-\alpha \times e_{rp}), \quad (1)$$

where  $e_{rp}$  denotes the error in either roll or pitch and  $\alpha$  is a

scaling factor that determines the sensitivity of the reward to orientation errors. Similarly, the body height tracking reward used to encourage the robot to maintain a desired body height is described as:

$$r_{\text{base height}} = \exp(-\alpha \times e_h), \quad (2)$$

where the base height error,  $e_h$ , represents the absolute distance between the desired and actual body heights above the ground.

The student policy for the low-level policy is trained following the same manner of the previous work [2] to remove the privileged observation dependency and apply noise on the exteroception. User It was trained using extensive domain randomization applied to exteroceptive measurements, employing behavior cloning to minimize the action difference between the teacher policy.

### C. High-level teacher policy

After we get the low-level policy, a high-level teacher policy is trained with PPO to command the low-level policy in navigating diverse terrains that include overhanging obstacles. While the low-level policy focuses on training specifically for rough terrain, the high-level policy is trained in confined spaces with overhanging obstacles. This separation in training allows us to focus the high-level policy training on the specific challenges of navigating confined spaces.

The high-level policy's task is to follow a desired velocity target by commanding the low-level policy. It takes proprioception  $o_p^H$  as well as exteroceptive observations  $o_e^H$  as

input.  $o_p^H$  includes velocity commands, body velocities, joint positions, joint velocities, body orientation, and previous actions.  $o_e^H$  consists with the same height samples as the low-level policy ( $o_e^L$ ), and a spherical scans for omnidirectional observations. We used the sparse ray tracing pattern, because using more complicated data structure, such as depth images or voxel grids, could slow down the reinforcement learning process due to their computational requirements. As the action space, a residual approach was utilized to increase training efficiency where it outputs the residual for  $v_x$ ,  $v_y$  and  $\omega_z$  instead of outputting the velocity commands directly. This can skip the process of learning to output the same input velocity commands to track them. For commands related to roll ( $a_\xi$ ), pitch ( $a_\theta$ ), and base height ( $a_h$ ), they are directly outputted.

We employ rewards consisting of task rewards that encourage the desired behavior and penalty rewards to avoid undesirable behavior. For task rewards, we used velocity tracking reward and base distance reward, designed to maintain a safe distance from obstacles and thus enhance safety during navigation. For base distance reward we used

$$r_{base} = \exp(-\alpha \times (d_{max} - \min(d, d_{max}))), \quad (3)$$

where  $d_{max}$  is a parameter to set a threshold and  $\alpha$  defines the coefficient. To measure the distances  $d$ , we used ray casting in a spherical pattern. For penalties, we discourage the robot from hitting objects with a collision penalty by simply giving a large negative reward when the body hits obstacles. We also use penalties for joint speed, acceleration and torque for smoother motion and vertical velocity and orientation to avoid unnecessary movements.

#### D. High-level student policy training

After training a teacher policy that can perform the task using ground truth data, we distill this policy into a student policy that can be deployed within the sensory constraints of a physical robot. We convert the spatial information into the occupancy voxel representation as an abstraction. This allows to combine different sensor configurations and also allows using voxel mapping pipeline [37], [40], [41], or transforming the past observations into the current frame [17] to accumulate information.

To compress the voxel representation to a smaller latent space before feeding to the belief encoder, we leverage a three layer voxel encoder to compress the data into a more compact latent space. The voxel grid dimensions are  $32 \times 32 \times 32$ , with a resolution of 0.08 along each axis. The voxel encoder consists of 3D convolutions coupled with layer normalization and ELU activation following the methodology described in [37].

We use a belief encoder based on a Gated Recurrent Unit (GRU) to integrate the proprioceptive data and the external latent representation into a coherent belief state similar to [2]. We flattened the voxel latent representation before putting into the belief encoder. The belief encoder has a hidden size of 128 and consists of two layers. Additionally, we incorporate a belief decoder that decodes privileged information

from the latent feature to help extracting information that is useful for student policy. Finally, the output of the belief encoder is fed into an MLP network, consistent with existing architectures, to produce the final action commands. The loss is the action loss which is the mean square error between the teacher policy and the reconstruction loss of the privileged states.

#### E. Procedural Terrain Generation




Fig. 3. Procedural terrain generation. We generate the terrain mesh by tiling mesh parts. First, the tiles are connected procedurally based on connectivity calculated from terrain height array. Then, overhanging obstacles are added on top of the mesh.

To simulate different environments, we developed a procedural terrain generator that builds mesh representations of different terrains<sup>1</sup>. This generator uses the Wave Function Collapse (WFC) method [42], widely used in game development, to ensure smooth transitions between neighboring tiles based on their connectivity. We used terrain heights to calculate the connectivity between each tile.

We defined 1570 types of tiles, including terrains such as steps, stairs, ramps and rough in different heights. After the mesh was generated, boxes of different dimensions were added to the terrain to simulate confined spaces. This method allows for the generation of a wide variety of terrain combinations, reflecting the complex environments a robot may encounter, as shown in Figure 3.

## IV. EXPERIMENTAL RESULTS

In this section, we describe the setup of our experiments, which include both simulations and real-world tests. In the simulation phase, we evaluate the performance of the policy on evaluation terrains. After that, we test the trained policy on actual hardware to validate the performance in real-world.

#### A. Evaluation in simulation

To evaluate the high-level student policy, we created a test environment with a variety of obstacle heights and overhanging obstacles. In this setup, the robot had to move from a starting position to a goal position that is 6 meters away. The criterion for a successful trial was the robot's ability to reach the target within 30 seconds. Note that we applied these test conditions to student policies that had been

<sup>1</sup>We open-sourced the software <https://github.com/leggedrobotics/terrain-generator>

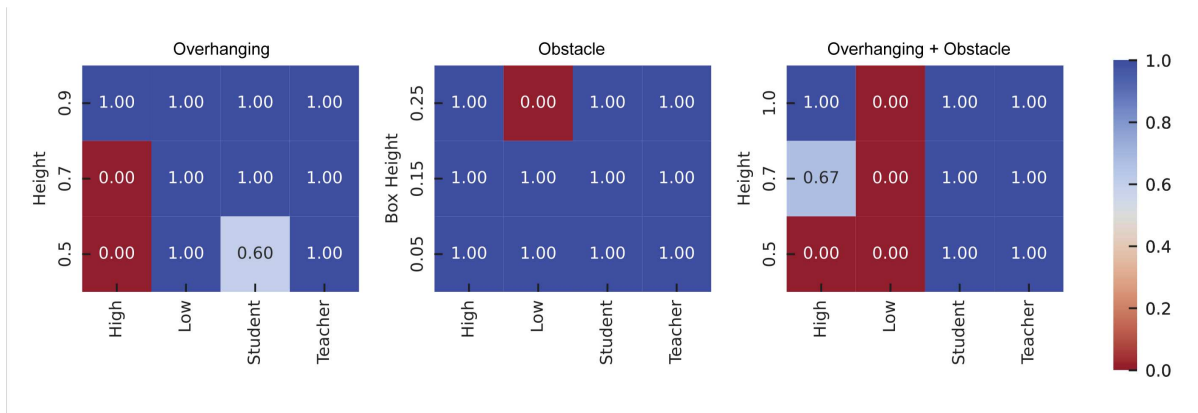


Fig. 4. Success rate of different overhanging obstacle height and obstacle box height from the ground. The  $x$  axis show different methods and  $y$  axis shows the different parameters of the obstacles. For the obstacle + overhanging, we used the obstacle height of  $0.25m$  and varied the height of the overhanging box. The baseline methods which always walk at normal height (High) and always walk with crouching (Low) was compared against our method. The results show that the combination of overhanging and rough terrain needs an adaptive body height control.




Fig. 5. Sequence on evaluation terrains. We have evaluated the policy’s performance on the evaluation terrains where it has three terrain types with different parameters. Overhanging has an overhanging box in the middle, and Obstacle has a box on the ground, Overhanging + obstacle is a combination of both obstacle and overhanging boxes.

trained on procedurally generated terrains, not to policies specifically trained on the evaluation setup. The results are detailed in Figure 4. We compared the performance against two baselines both employ simple strategy without the high-level policy. One is the policy which always walk at normal height (High), and the another one always crouching on the ground (Low). The result shows that the student policy could handle three types of the terrains fairly well and closely matched the performance of the teacher policy. On the other hand, the Low baseline could traverse overhanging obstacles but failed to overcome obstacle on the ground and High baseline could traverse obstacles on the ground but failed to go through overhanging obstacle.

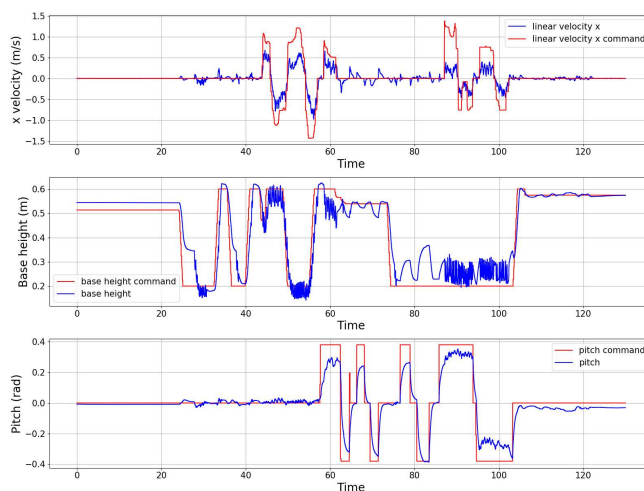


Fig. 6. Tracking performance of a low-level policy when responding to simultaneous commands for body height, pitch, and linear  $x$  velocity. The results confirm that the policy can effectively manage combined actions such as crouching, tilting, and walking simultaneously.

## B. Real-World Experiments

Our trained policies were empirically tested using ANYmal-C and ANYmal-D robots equipped with either dome LiDARs or depth cameras. Depth sensor data were transformed into voxel structures - achieved by voxel mapping for depth cameras and straightforward voxelization in the LiDAR configuration.

1) *Low-level policy*: First, we evaluate the low-level policy’s capabilities with the operator giving commands to guide the robot’s behavior. The robot effectively followed the combinations of different commands as shown in Figure 6. Even under complex conditions, such as maintaining a crouched position while traversing stairs, steps and uneven terrain, the robot was able to adjust its body height and orientation according to the operator’s instructions. These empirical results, illustrated in Figure 7, show the robustness and adaptability of our low-level policy.

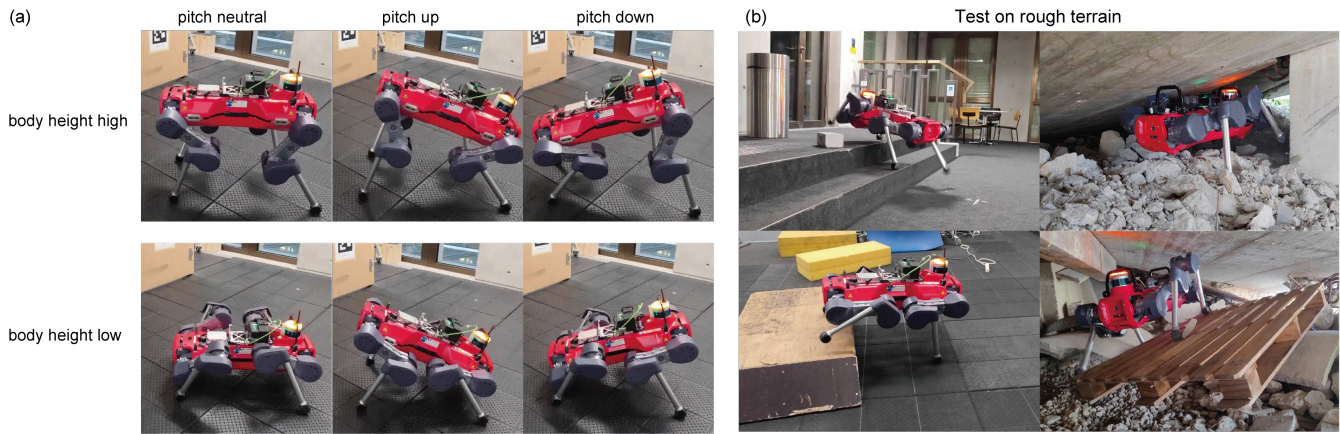


Fig. 7. Tests of the low-level policy on the hardware. (a) The robot could follow additional commands such as body height or pitch. (b) The policy showed its robustness on different rough terrains such as stairs, step, loose ground or wooden steps while following low body height command.




Fig. 8. Real-world deployment of the integrated high-level and low-level policies. The sequence of images illustrates the robot's ability to autonomously adapt its body height in the presence of overhanging obstacles.

2) *Combined policy in the field:* In the final part of our experiments, we tested a combined high-level and low-level policy in a real-world environment, as shown in Figure 1 and 8. The test environment consisted of a simulated collapsed building and some objects such as tables and benches.

In the tests, the operator provided commands for the robot's  $x$ ,  $y$ , and yaw velocities and the robot autonomously determined its own body height, roll, and pitch. The robot faced several challenges, such as navigating confined spaces with ceilings at varying angles. The terrain conditions were also difficult, with loose and rough gravels and unstable steps. However the robot could handle all challenges robustly and showed that the robot could reliably enter and traverse confined spaces under challenging conditions. This confirms the robustness of our combined policy approach in real-world scenarios.

### C. Simulation details

For low-level policy training, we used RaiSim [43] to simulate the robot and the environment, while for high-level policy, Isaac Gym [44] was used to simulate confined space environment with depth sensor measurements such as

spherical scan, depth camera and voxel map. We used 1000 parallel environments for the low-level and high-level teacher policy training and 300 for the student policy training.

## V. CONCLUSIONS

In this paper, we present a legged robot system designed for navigating confined spaces, broadening the operational environments for robotics. Our framework adopts a two-layer hierarchical policy structure to decompose the complexity of navigation into more manageable components. The low-level policy ensures robust traversal over uneven terrain based on 6D commands, including body velocity and orientation, while the high-level policy strategically directs these commands using 3D spatial understanding. Our validation involved training on diverse terrains and deploying in a challenging collapsed building scenario, demonstrating the system's real-world applicability. However, limitations exist in highly dynamic environments and performing advanced maneuvers like using body contacts in tight spaces, pointing to future research directions.

## REFERENCES

- [1] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, 2020.
- [2] 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. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abk2822>
- [3] P. Arm, G. Waibel, J. Preisig, T. Tuna, R. Zhou, V. Bickel, G. Ligeza, T. Miki, F. Kehl, H. Kolvenbach *et al.*, "Scientific exploration of challenging planetary analog environments with a team of legged robots," *Science robotics*, vol. 8, no. 80, p. eade9548, 2023.
- [4] M. Tranzatto, M. Dharmadhikari, L. Bernreiter, M. Camurri, S. Khatatak, F. Mascari, P. Pfreundschuh, D. Wisth, S. Zimmermann, M. Kulkarni, V. Reijgwart, B. Casseau, T. Homberger, P. De Petris, L. Ott, W. Tubby, G. Waibel, H. Nguyen, C. Cadena, R. Buchanan, L. Wellhausen, N. Khedekar, O. Andersson, L. Zhang, T. Miki, T. Dang, M. Mattamala, M. Montenegro, K. Meyer, X. Wu, A. Briod, M. Mueller, M. Fallon, R. Siegwart, M. Hutter, and K. Alexis, "Team CERBERUS wins the DARPA subterranean challenge: Technical overview and lessons learned," Jul. 2022.
- [5] M. Tranzatto, T. Miki, M. Dharmadhikari, L. Bernreiter, M. Kulkarni, F. Mascari, O. Andersson, S. Khatatak, M. Hutter, R. Siegwart, and K. Alexis, "CERBERUS in the DARPA subterranean challenge," *Sci Robot*, vol. 7, no. 66, p. eabp9742, May 2022.
- [6] A. Bouman, M. F. Ginting, N. Alatur, M. Palieri, D. D. Fan, T. Touma, T. Pailevanian, S.-K. Kim, K. Otsu, J. Burdick, and A.-A. Aghamohammadi, "Autonomous spot: Long-Range autonomous exploration of extreme environments with legged locomotion," Oct. 2020.
- [7] N. Kottege, J. Williams, B. Tidd, F. Talbot, R. Steindl, M. Cox, D. Frousheger, T. Hines, A. Pitt, B. Tam, B. Wood, L. Hanson, K. Lo Surdo, T. Molnar, M. Wildie, K. Stepanas, G. Catt, L. Tyche-Smith, D. Penfold, L. Overs, M. Ramezani, K. Khosoussi, F. Kendoul, G. Wagner, D. Palmer, J. Manderson, C. Medek, M. O'Brien, S. Chen, and R. C. Arkin, "Heterogeneous robot teams with unified perception and autonomy: How team CSIRO data61 tied for the top score at the DARPA subterranean challenge," Feb. 2023.
- [8] L. Han, Q. Zhu, J. Sheng, C. Zhang, T. Li, Y. Zhang, H. Zhang, Y. Liu, C. Zhou, R. Zhao, J. Li, Y. Zhang, R. Wang, W. Chi, X. Li, Y. Zhu, L. Xiang, X. Teng, and Z. Zhang, "Lifelike agility and play on quadrupedal robots using reinforcement learning and generative pre-trained models," Aug. 2023.
- [9] Z. Zhuang, Z. Fu, J. Wang, C. G. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, "Robot parkour learning," in *7th Annual Conference on Robot Learning*, 2023.
- [10] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *arXiv preprint arXiv:2306.14874*, 2023.
- [11] Y. Kim, H. Oh, J. Lee, J. Choi, G. Ji, M. Jung, D. Youm, and J. Hwangbo, "Not only rewards but also constraints: Applications on legged robot locomotion," Aug. 2023.
- [12] R. Buchanan, L. Wellhausen, M. Bjelonic, T. Bandyopadhyay, N. Kottege, and M. Hutter, "Perceptive whole-body planning for multilegged robots in confined spaces," *Journal of Field Robotics*, vol. 38, no. 1, pp. 68–84, 2021.
- [13] R. Buchanan, T. Bandyopadhyay, M. Bjelonic, L. Wellhausen, M. Hutter, and N. Kottege, "Walking posture adaptation for legged robot navigation in confined spaces," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2148–2155, 2019.
- [14] R. Yang, M. Zhang, N. Hansen, H. Xu, and X. Wang, "Learning vision-guided quadrupedal locomotion end-to-end with cross-modal transformers," in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=nhnJ3oo6AB>
- [15] W. Yu, D. Jain, A. Escontrela, A. Iscen, P. Xu, E. Coumans, S. Ha, J. Tan, and T. Zhang, "Visual-LoComotion: Learning to walk on complex terrains with vision," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 2022, pp. 1291–1302.
- [16] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," Sep. 2022.
- [17] R. Yang, G. Yang, and X. Wang, "Neural volumetric memory for visual locomotion control," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1430–1440.
- [18] M. Zucker, J. A. Bagnell, C. G. Atkeson, and J. Kuffner, "An optimization approach to rough terrain locomotion," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 3589–3595.
- [19] P. D. Neuhuis, J. E. Pratt, and M. J. Johnson, "Comprehensive summary of the institute for human and machine cognition's experience with LittleDog," *The International Journal of Robotics Research*, vol. 30, no. 2, pp. 216–235, 2011.
- [20] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [21] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and 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.
- [22] G. Erni, J. Frey, T. Miki, M. Mattamala, and M. Hutter, "Mem: Multi-modal elevation mapping for robotics and learning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 11 011–11 018.
- [23] J. Z. Kolter, Y. Kim, and A. Y. Ng, "Stereo vision and terrain modeling for quadruped robots," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 1557–1564.
- [24] I. Havoutis, J. Ortiz, S. Bazeille, V. Barasoul, C. Semini, and D. G. Caldwell, "Onboard perception-based trotting and crawling with the hydraulic quadruped robot (HyQ)," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 6052–6057.
- [25] D. Belter, P. Łabecki, and P. Skrzypczyński, "Adaptive motion planning for autonomous rough terrain traversal with a walking robot," *Journal of Field Robotics*, vol. 33, no. 3, pp. 337–370, 2016.
- [26] D. Kim, D. Carballo, J. Di Carlo, B. Katz, G. Bleedt, B. Lim, and S. Kim, "Vision aided dynamic exploration of unstructured terrain with a small-scale quadruped robot," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2464–2470.
- [27] F. Jenelten, T. Miki, A. E. Vijayan, M. Bjelonic, and M. Hutter, "Perceptive locomotion in rough terrain—online foothold optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5370–5376, 2020.
- [28] F. Jenelten, R. Grandia, F. Farshidian, and others, "TAMOLS: Terrain-aware motion optimization for legged systems," *IEEE Transactions on*, 2022.
- [29] R. Grandia, F. Jenelten, S. Yang, and others, "Perceptive locomotion through nonlinear Model-Predictive control," *IEEE Transactions*, 2023.
- [30] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-real: Learning agile locomotion for quadruped robots," in *Robotics: Science and Systems*, 2018.
- [31] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Sci Robot*, vol. 4, no. 26, Jan. 2019.
- [32] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," *Robotics: Science and Systems*, 2020.
- [33] 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.
- [34] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, "Blind bipedal stair traversal via sim-to-real reinforcement learning," in *Robotics: Science and Systems*, 2021.
- [35] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "RLOC: Terrain-aware legged locomotion using reinforcement learning and optimal control," *arXiv preprint arXiv:2012.03094*, 2020.
- [36] K. Caluwaerts, A. Iscen, J. C. Kew, W. Yu, T. Zhang, D. Freeman, K.-H. Lee, L. Lee, S. Saliceti, V. Zhuang *et al.*, "Barkour: Benchmarking animal-level agility with quadruped robots," *arXiv preprint arXiv:2305.14654*, 2023.
- [37] D. Hoeller, N. Rudin, C. Choy, and others, "Neural scene representation for locomotion on structured terrain," *IEEE Robotics and*, 2022.
- [38] C. S. Imai, M. Zhang, Y. Zhang, M. Kierebiński, R. Yang, Y. Qin, and X. Wang, "Vision-Guided quadrupedal locomotion in the wild with Multi-Modal delay randomization," pp. 5556–5563, Oct. 2022.

- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [40] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [41] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: an efficient probabilistic 3d mapping framework based on octrees," *Auton. Robots*, vol. 34, no. 3, pp. 189–206, 2013. [Online]. Available: <http://dblp.uni-trier.de/db/journals/arobots/arobots34.html#HornungWBSB13>
- [42] M. Gumin, "WaveFunctionCollapse," <https://github.com/mxgmn/WaveFunctionCollapse>.
- [43] J. Hwangbo, J. Lee, and M. Hutter, "Per-contact iteration method for solving contact dynamics," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 895–902, 2018. [Online]. Available: [www.raisim.com](http://www.raisim.com)
- [44] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and Gavriel State, "Isaac gym: High performance GPU-Based physics simulation for robot learning," Aug. 2021.