

# Bridge the Gap: Enhancing Quadruped Locomotion with Vertical Ground Perturbations

Maximilian Stasica<sup>\*,1</sup>, Arne Bick<sup>\*,1,2</sup>, Nico Bohlinger<sup>\*,2</sup>, Omid Mohseni<sup>1,3</sup>, Max Johannes Alois Fritzsche<sup>4</sup>, Clemens Hübler<sup>4</sup>, Jan Peters<sup>2,5</sup>, and André Seyfarth<sup>1</sup>

**Abstract**—Legged robots, particularly quadrupeds, excel at navigating rough terrains, yet their performance under vertical ground perturbations, such as those from oscillating surfaces, remains underexplored. This study introduces a novel approach to enhance quadruped locomotion robustness by training the Unitree Go2 robot on an oscillating bridge—a 13.24-meter steel-and-concrete structure with a 2 Hz eigenfrequency designed to perturb locomotion. Using Reinforcement Learning (RL) with the Proximal Policy Optimization (PPO) algorithm in a MuJoCo simulation, we trained 15 distinct locomotion policies, combining five gaits (trot, pace, bound, free, default) with three training conditions: rigid bridge and two oscillating bridge setups with differing height regulation strategies (relative to bridge surface or ground). Domain randomization ensured zero-shot transfer to the real-world bridge. Our results demonstrate that policies trained on the oscillating bridge exhibit superior stability and adaptability compared to those trained on rigid surfaces. Our framework enables robust gait patterns even without prior bridge exposure. These findings highlight the potential of simulation-based RL to improve quadruped locomotion during dynamic ground perturbations, offering insights for designing robots capable of traversing vibrating environments.

## I. INTRODUCTION

The development of legged robots, particularly quadrupeds, has surged in popularity due to their ability to traverse challenging terrains such as mountainsides [1] and obstacle-rich environments [2], [3]. While these approaches can handle rigid uneven surfaces, their performance under active ground perturbations—encompassing moving obstacles, and both horizontal and vertical ground movements—remains under-examined [4].

This work was funded by German Research Foundation (DFG) under the *HUMAN-structure interaction and gait adaptation during locomotion on Vibrating structures* (HUMVIB) project (grant number: 446124066) and the INTENTION project (grant number: 506123304). This project was also partially funded by National Science Centre, Poland under the OPUS call in the Weave program UMO-2021/43/I/ST6/02711.

\*Equal contribution.

<sup>1</sup>M. Stasica, A. Bick, O. Mohseni and A. Seyfarth are with the Laulabor Locomotion Laboratory, Institute of Sports Science and Centre of Cognitive Science, Technical University of Darmstadt, 64289 Darmstadt, Germany [stasica@sport.tu-darmstadt.de](mailto:stasica@sport.tu-darmstadt.de)

<sup>2</sup>A. Bick, N. Bohlinger and J. Peters are with the Intelligent Autonomous Systems group, Department of Computer Science, Technical University of Darmstadt, 64289 Darmstadt, Germany

<sup>3</sup>O. Mohseni is with the Measurement and Sensor Technology Group, Department of Electrical Engineering and Information Technology, Technical University of Darmstadt, Darmstadt, Germany

<sup>4</sup>M. J. A. Fritzsche and C. Hübler are with the Institute of Structural Mechanics and Design, Technical University of Darmstadt, 64289 Darmstadt, Germany

<sup>5</sup>J. Peters is with the German Research Center for AI (DFKI), Research Department: Systems AI for Robot Learning; Hessian.AI; Robotics Institute Germany (RIG)



Fig. 1: We train the Unitree Go2 quadruped in simulation and zero-shot transfer the learned policies to the oscillating real-world HUMVIB bridge to investigate the impact of vertical ground perturbations on locomotion robustness.

Although many robot designs and controllers have been evaluated on terrain with different stiffness levels [5] or friction coefficients [6], [7], vertical ground movements are rarely studied, despite their relevance to real-world scenarios such as disaster zones, industrial sites, or hazardous fields. This gap, which also appears in human locomotion research [8], arises partly from the challenge of replicating such dynamic conditions in controlled environments—leaving a critical blind spot in our understanding of robotic adaptability to unstable surfaces.

Traditional robotic controllers often struggle with dynamic perturbations, as they are typically designed for predictable or rigid terrains. While some systems manage uneven ground or obstacles effectively [2], [3], their reliance on pre-tuned parameters limits resilience to sudden vertical shifts. In human studies, vertical perturbations like varying ground stiffness significantly alter performance. Human runners adapt instantly to these perturbations by adjusting leg mechanics [9], [10], but the robotic equivalent remains under-explored. The logistical challenge of testing such scenarios has hindered progress, underscoring the need for both innovative platforms and advanced control strategies to enhance locomotion robustness in unpredictable settings.

In recent years, Deep Reinforcement Learning (DRL) has become a popular approach for learning agile locomotion controllers for legged robots [4]. On-policy DRL algorithms, such as Proximal Policy Optimization (PPO) [11], can be scaled up through many parallel simulated physics environments and big batch sizes to learn complex and dynamic




Fig. 2: Schematic of the Unitree Go2 quadruped crossing the 13.24-meter long HUMVIB bridge.

locomotion behaviors for any legged robot form factor [12], [6], [13]. There are many examples of learning highly agile and robust locomotion for quadruped robots, such as fast-paced running and turning [7], [14], dynamic jumping on parkour courses [15], [16], [2], [17], climbing on ladders or over tall obstacles [18], [3], hiking on mountainous terrain [1] and even performing handstands, backflips and other acrobatic maneuvers [2], [19]. To transfer the learned policies to real-world robots, a wide range of domain randomization is necessary to bridge the sim-to-real gap [20], [12], [6], [15]. This includes randomizing the robot’s properties, such as mass, inertia, and actuator dynamics, as well as the environment’s properties, such as ground friction and roughness of the terrain. To further improve the robustness and to teach the policy to adapt to unforeseen external disturbances, the training has to be augmented with adversarial perturbations that can occur at any moment in a learning episode. However, the majority of works only consider sensor noise and pushes on the robot’s trunk as active disturbances [12], [18], [21], while the effects of active ground perturbations, like vertical ground movements, on the robot’s locomotion are widely unexplored.

Understanding how Reinforcement Learning (RL) policies shape locomotion requires examining quadruped gaits, which are defined by the coordination of leg movements. Common quadruped gaits include the trot (diagonal legs move together), pace (lateral legs move in unison), bound (front and back legs pair for a leaping motion), and more flexible styles like free (unconstrained coordination) [22]. The interplay between gait selection and dynamic ground conditions, such as vertical oscillations, still have to be fully understood in robotics. Investigating these gait types under such perturbations is crucial for assessing how robots maintain stability and robustness, complementing the focus of RL on policy optimization with insights into physical locomotion dynamics.

To tackle these challenges, we introduce a novel approach using the Unitree Go2 quadruped (Unitree, Hangzhou, China) on a purpose-built oscillating bridge (Figure 2). As

part of the HUMVIB project (*HUMAN*-structure interaction and gait adaptation during locomotion on *VIB*rating structures), this structure, with a span of 13.24 m, composed of two steel beams, and thirteen concrete slabs, features an eigenfrequency of approximately 2 Hz, making it highly susceptible to locomotion-induced oscillations [23], [24]. Equipped with a 2.5 m wide track and sensors like force plates and accelerometers, the bridge provides a controlled yet dynamic testbed for vertical perturbations. By training the robot in simulation with state-of-the-art RL techniques and evaluating its gaits in the real-world setting, our study bridges the gap between rigid terrain research and the demands of unstable environments, aiming to enhance robotic locomotion stability and adaptability.

## II. LEARNING LOCOMOTION ON AN OSCILLATING BRIDGE

We propose learning locomotion policies for quadruped robots on an oscillating bridge to investigate the impact of vertical ground perturbations on the robot’s locomotion and achieve higher robustness of the learned policies. We first describe the RL training setup and the modeling of the bridge in simulation, followed by the necessary components for learning different gaits on the bridge, and the evaluation setup in the real world.

### A. Training setup

We trained the Unitree Go2 quadruped in simulation using the CPU-based MuJoCo physics engine [25] with 48 parallel environments for fast data collection. Like previous works [12], [6], [13], we used the PPO algorithm [11] to learn different locomotion policies. We built on the DRL library RL-X [26] to implement and integrate the algorithm with the simulation environment. The policies were trained to control the robot at 50 Hz with target joint positions, and to walk with a given x-y-yaw-command velocity  $\bar{v} \in [-1.0, 1.0]^3$  with respect to the robot’s trunk. To enable zero-shot transfer of the learned policies to the real robot, we applied a wide range of domain randomization during training, including randomizing the robot’s mass, inertia, Center of Mass (CoM) (CoM), actuator dynamics and delays, the ground properties, such as friction and compliance, sensor noise, and pushes on the robot’s trunk. To investigate the robustness of the policies through the impact of vertical ground perturbations with the HUMVIB bridge, we trained the robot to walk on either a rigid or harmonically oscillating ground with varying eigenfrequencies and amplitudes.

### B. Bridge model

The HUMVIB bridge is modeled in MuJoCo as a harmonic oscillator emulating the dynamics of the real bridge. We fixed the surface of the bridge at 1.05 m over the ground—the peak of the oscillation—while the equilibrium position can be adjusted to modify the oscillation amplitude. The stiffness is tuned such that the bridge exhibits an eigenfrequency of 2 Hz with an oscillation amplitude of  $\pm 0.1$  m, similar to the real bridge. During training, we varied the eigenfrequency

of the bridge between 0.75 Hz and 7.5 Hz and the amplitude between zero and a constrained maximum value that ensures the bridge’s acceleration remains below  $9.81 \text{ m/s}^2$  given its mass and adjusted stiffness. This constraint is necessary to prevent the robot from experiencing too much acceleration to become airborne.

### C. Learning different gaits

We used the same reward terms, coefficients and curriculum as in [13], and denoted the learned gait style as *default*. To encourage the emergence of the different *trot*, *pace*, and *bound* gaits, we modified the existing symmetry reward term to penalize deviations from the characteristic stance phases of the respective gait. Finally, removing the symmetry constraints results in the gait style we refer to as *free*.

Beyond the gait-specific reward terms, the environments in which the policies are learned can significantly shape the resulting gaits. For gaits learned on the oscillating bridge, two different policies were trained by altering the base height reward term, which promotes the robot to maintain a constant height of 0.325 m—the trunk height in the nominal standing position—with respect to the surface it is walking on. In one case, the robot is rewarded to maintain a constant height relative to the oscillating bridge surface, called *equidistant bridge* (*eb*), while in the other, it is encouraged to maintain a constant height relative to the ground, called *equidistant ground* (*eg*). When training on the rigid bridge, the base height reward term only has a single interpretation as the robot’s height with respect to the ground, called *no oscillation* (*nos*). Overall, this approach results in 15 distinct policies, derived from the five different gaits (*trot*, *pace*, *bound*, *free*, *default*), each trained under three conditions, one on the rigid bridge and two on the oscillating bridge with the different height regulation strategies (*nos*, *eb*, *eg*) respectively.

### D. Real-world setup

The real world experiments took place on the HUMVIB bridge. To validate the presented algorithm, the structure was equipped with six Delsys Trigno Wireless sensors (Delsys, Natick, US) which were used as an IMU to track the bridges acceleration in different locations. For each combination of gait style and training setting, the robot had to complete eight passes over the pre-oscillated and the idle bridge (Figure 1) respectively. The command velocity  $\bar{v}$  was controlled by a human operator. This ensured a constant speed of 0.5 m/s in the x-direction and the operator could adjust the y- and yaw-velocity to keep the robot on track in the case of lateral drift.

## III. RESULTS

We first evaluated the learning dynamics of the different gaits and training conditions. To ensure a fair comparison, we measured the performance of the policies by calculating the episode return without the gait-specific reward terms. Figure 3 shows that the *free* and *default* gaits are the easiest




Fig. 3: Top: Episode return for the different gaits (left) and training conditions (right) during training. Bottom: Episode return over the command velocity on the oscillating bridge for the different gaits (left) and training conditions (right) of the final policies during evaluation.

to learn, as they impose no or fewer restrictions on the policy compared to the *trot*, *pace*, and *bound* gaits that rely on strong gait-specific reward terms to enforce their desired footfall patterns. We show all of the learned footfall patterns in Figure 4. Of the three gait styles, the *trot* gait is visually and reward-wise the closest to the *default* gait, while the *pace* and *bound* gaits are more distinct and achieve significantly lower returns. This indicates that these policies are worse at tracking the command velocity and more susceptible to disturbances from the training environment. When comparing the training conditions, the *nos* policy performs best, as it is trained on the rigid surface and is less disturbed by the oscillations, while the *eb* and *eg* policy perform similarly to each other.

Next, we evaluated the performance of the final policies on the oscillating bridge with varying command velocities. Figure 3 shows that the *free* and *default* gaits perform best with smaller command velocities, where most of the probability mass for the sampled commands is located. The *trot* gait performs well at higher command velocities, and even outperforms the *free* and *default* gaits. The *pace* and *bound* gaits perform worse than the other gaits in most cases. When comparing the training conditions, the *nos* policy performs significantly worse on the oscillating bridge than the *eb* and *eg* policies, highlighting the importance of training locomotion policies with vertical ground perturbations to achieve higher robustness in such environments.

To investigate how the different policies cope with the oscillating bridge, we evaluated the movement of the robot’s CoM compared to the bridge’s surface. Figure 5 shows the *x*, *y* and *z* components of the CoM movement of the

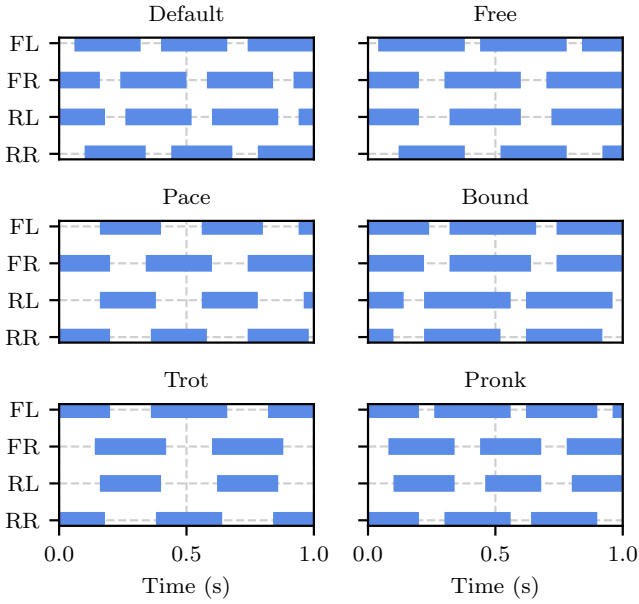


Fig. 4: Footfall pattern of the different gait styles on the rigid bridge with a target speed  $v_x = 0.5$  m/s. The feet of the robot are denoted with front left (*FL*), front right (*FR*), rear left (*RL*), and rear right (*RR*). In our study, we omit the *pronk* gait, as we were unable to train a policy that sufficiently learned the desired footfall pattern, due to the agent resorting to a more stable trot-like gait during learning.

robot with a fixed velocity command  $v_x = 0.5$  m/s for the *default* gait style under the three training conditions. When the bridge is oscillating with a frequency of 2.0 Hz and an amplitude  $\pm 0.1$  m, all policies clearly show adaptation in their  $z$ -movement to the bridge’s oscillation with a small phase shift of around  $0.3\pi$ , due to the robot’s inertia. The *nos* policy, however, struggles to keep the robot’s CoM constant, exhibits more chaotic behavior, and struggles to walk at the commanded velocity, progressing slower over the bridge than the *eb* and *eg* policies. Interestingly, the *eb* and *eg* policies show a lateral drift in the  $y$ -direction, when the bridge is not oscillating, which is not present in the *nos* policy, as it was trained in this scenario. On the oscillating bridge, the *eb* and *eg* policies are able to move mostly straight ahead, while now the *nos* policy shows a strong drift in the  $y$ -direction. This indicates that the *eb* and *eg* policies are more robust to the oscillating bridge but are slightly worse adapted to the non-oscillating surface. Additionally, we provide snapshots of the simulated gait cycles on the oscillating bridge for two policies, *bound nos* and *default eb*, in Figure 6.

In the real-world validation experiments, the different policies were evaluated on the HUMVIB structure, with the exception of the *pace* gait. This exclusion was due to the instability of the *pace* policies, increasing the risk of the robot falling off the structure and potentially damaging itself. All other policies could be evaluated and showed a decrease in movement speed while passing over the middle section of the structure, where the oscillation were highest. The *eb* an




Fig. 5: Movement of the robot’s CoM with a target speed  $v_x = 0.5$  m/s for the *default* gait pattern under three conditions: no oscillation (*nos*), equidistant bridge (*eb*), and equidistant ground (*eg*). Subfigures on the first column show  $x$ ,  $y$ , and  $z$  directions of robot’s CoM on an idle bridge, while the subfigures on the second column show the same quantities, but on an oscillating bridge.

*eg* policies showed a more stable performance with regard to vertical and lateral movements of the CoM. Interestingly, although the *bound* gait performed poor in simulation, it was capable of stable locomotion on the oscillating bridge in the real world, hinting at inherent robustness and transferability of the gait to the real system. While the bridge was pre-oscillated by human operators, the *trot* policies were able to excite the structure, since their step frequency of around 4.0 Hz fell close to the 2nd harmonic of the structure.

To assess the interaction between the robot and the bridge, we measured the mean forces in the feet of the robot. When oscillating of the bridge, the *nos* policies showed the highest forces, while the *eb* policies recorded the lowest ones (Figure 7), indicating the learned adaptation to the oscillations of the real bridge when trained with the simulated model. Considering the gait styles, *trot* and *default* recorded similar forces, while the *free* policies produced more force, with *bound* topping the scale (Figure 7), due to its inherent leaping motion.






Fig. 6: Snapshots of simulated gait cycles on the oscillating bridge for *bound nos* (left) and *default eb* (right). Each set is arranged left to right, row-wise, spanning one full oscillation cycle of the bridge.




Fig. 7: Means and standard deviations of the force readouts of the foot sensors of the robot on the oscillating HUMVIB bridge in the real world.

#### IV. DISCUSSION

During the learning phase, the *free* gait was the easiest to learn for the RL agents. This is expected, since it imposes no restrictions on the footfall patterns compared to the gait styles found in quadrupeds in nature. Also, the worse performance of the *bound* and *pace* policies is expected, since they include relatively longer contact times, which were penalized by the learner, and are inherently less stable compared to the *trot* and *default* gaits. While this study evaluated the performance of 15 different policies in the given scenario, each policy was only trained with one seed, limiting the statistical power. This could also be part of the reason, why the *prongk* policy reversed to a *trot* behavior during training. Further optimization of the policies is required to provide an exhaustive comparison of the different gait styles. Since currently the real-world bridge can only exhibit an eigenfrequency of 2.0 Hz or be rigid, further frequencies could not be tested in our study. When looking at the different height regulation strategies, the *nos* policies superiority during learning is a product of the simpler training environment, as the robot is not disturbed by oscillations of the bridge. The performance differences of the policies with respect to the CoM movement in simulation can

also be explained by their relative training scenarios. Since the *eb* and *eg* policies trained on oscillating surfaces, they had a harder time to adapt to the rigid bridge. However, the oscillation seems to impose so much of a disturbance that the drift of the *nos* policy becomes excessive to a point where it would fall off the bridge after less than 10 s.

While most policies cope well with the setting they used for learning, we can observe that the current black-box modelling approach does not generalize perfectly to situations that have not been learned. A mixed black-box and white-box model approach could be used to overcome this limitation. White-box models are structural models that implement physiological details about the sensorimotor system (e.g., neural circuitries, muscle properties [27], sensor-mechanical couplings [28]). This combination of structure models and learning may help to better understand the versatility of humans and animals who are able to traverse over moving grounds benefiting of the underlying neural controller dynamics [29].

Future work will extend our research to explore the combination of ground perturbations with moving obstacles, as well as the integration of high-level planning systems to navigate complex multi-layered terrains with dynamic and adversarial perturbations.

Lastly, while no quadrupedal animals have been tested on the HUMVIB bridge yet, future research is planned to understand the biological coping strategies in this scenario, covering the full picture of quadrupedal gaits during vertical ground oscillations [30].

#### V. CONCLUSION

This study demonstrates that locomotion policies trained in simulation on an oscillating surface significantly outperform those developed on rigid terrain when tested on the Unitree Go2 quadruped navigating the HUMVIB bridge. By employing RL with the PPO algorithm across 15 distinct policies—spanning five gaits and three training conditions—we found that exposure to vertical ground perturbations during training enhances stability and adaptability, a result validated through zero-shot transfer to the physical bridge. Notably, this work marks the first effort to both simulate and experimentally evaluate a robot’s response to such dynamic

bridge-induced disturbances, revealing robust gait patterns that withstand real-world instabilities without prior exposure to the testbed. Our findings underscore the advantage of incorporating dynamic perturbations into simulation environments, advancing the understanding of quadruped locomotion under challenging conditions and paving the way for designing more robust robots capable of operating effectively in dynamic and unpredictable environments.

#### ACKNOWLEDGMENT

We sincerely thank the German Research Center for AI (DFKI), Research Department: Systems AI for Robot Learning, for lending us the Unitree Go2 quadruped.

#### REFERENCES

- [1] 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.
- [2] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, "Extreme parkour with legged robots," in *RoboLetics: Workshop on Robot Learning in Athletics@ CoRL 2023*, 2023.
- [3] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024.
- [4] S. Ha, J. Lee, M. van de Panne, Z. Xie, W. Yu, and M. Khadiv, "Learning-based legged locomotion: State of the art and future perspectives," *The International Journal of Robotics Research*, p. 02783649241312698, 2024.
- [5] S. Choi, G. Ji, J. Park, H. Kim, J. Mun, J. H. Lee, and J. Hwangbo, "Learning quadrupedal locomotion on deformable terrain," *Science Robotics*, vol. 8, no. 74, p. eade2256, 2023.
- [6] G. Ji, J. Mun, H. Kim, and J. Hwangbo, "Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4630–4637, 2022.
- [7] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *The International Journal of Robotics Research*, vol. 43, no. 4, pp. 572–587, 2024.
- [8] D. Tokur, M. Grimmer, and A. Seyfarth, "Review of balance recovery in response to external perturbations during daily activities," *Human movement science*, vol. 69, p. 102546, 2020.
- [9] T. A. McMahon and P. R. Greene, "Fast running tracks," *Scientific American*, vol. 239, no. 6, p. 148, 1978.
- [10] D. P. Ferris, K. Liang, and C. T. Farley, "Runners adjust leg stiffness for their first step on a new running surface," *Journal of biomechanics*, vol. 32, no. 8, pp. 787–794, 1999.
- [11] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [12] 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.
- [13] N. Bohlinger, G. Czechmanowski, M. Krupka, P. Kicki, K. Walas, J. Peters, and D. Tateo, "One policy to run them all: an end-to-end learning approach to multi-embodiment locomotion," *Conference on Robot Learning*, 2024.
- [14] S. Mahankali, C.-C. Lee, G. B. Margolis, Z.-W. Hong, and P. Agrawal, "Maximizing quadruped velocity by minimizing energy," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 467–11 473.
- [15] 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.
- [16] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, "Robot parkour learning," in *Conference on Robot Learning (CoRL)*, 2023.
- [17] E. Chane-Sane, J. Amigo, T. Flayols, L. Righetti, and N. Mansard, "Soloparkour: Constrained reinforcement learning for visual locomotion from privileged experience," in *8th Annual Conference on Robot Learning*, 2024.
- [18] D. Vogel, R. Baines, J. Church, J. Lotzer, K. Werner, and M. Hutter, "Robust ladder climbing with a quadrupedal robot," *arXiv preprint arXiv:2409.17731*, 2024.
- [19] D. Kim, H. Kwon, J. Kim, G. Lee, and S. Oh, "Stage-wise reward shaping for acrobatic robots: A constrained multi-objective reinforcement learning approach," *arXiv preprint arXiv:2409.15755*, 2024.
- [20] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [21] F. Shi, C. Zhang, T. Miki, J. Lee, M. Hutter, and S. Coros, "Rethinking robustness assessment: Adversarial attacks on learning-based quadrupedal locomotion controllers," *arXiv preprint arXiv:2405.12424*, 2024.
- [22] Y. Fukuoka, Y. Habu, and T. Fukui, "A simple rule for quadrupedal gait generation determined by leg loading feedback: a modeling study," *Scientific reports*, vol. 5, no. 1, p. 8169, 2015.
- [23] A. Firus, J. Schneider, W. Francke, and B. Kunkel, "Tu darmstadt humvib-bridge: Investigation of the human-structure interaction on a full-scale experimental pedestrian bridge," *Bulletin of the Transilvania University of Brasov. Series I-Engineering Sciences*, pp. 63–72, 2016.
- [24] M. Fritzsche, H. Berthold, S. Lorenzen, J. Schneider, A. Firus, M. Stastica, G. Zhao, and A. Seyfarth, "Integrated measurement concept for identification of human-structure interaction of flexible structures for natural gait," in *Bridge Safety, Maintenance, Management, Life-Cycle, Resilience and Sustainability*. CRC Press, 2022, pp. 713–720.
- [25] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.
- [26] N. Bohlinger and K. Dorer, "RI-x: A deep reinforcement learning library (not only) for robocup," in *Robot World Cup*. Springer, 2023, pp. 228–239.
- [27] C. Schumacher and A. Seyfarth, "Sensor-motor maps for describing linear reflex composition in hopping," *Frontiers in computational neuroscience*, vol. 11, p. 108, 2017.
- [28] U. B. Rongala, A. Seyfarth, V. Hayward, and H. Jörntell, "The import of skin tissue dynamics in tactile sensing," *Cell Reports Physical Science*, vol. 5, no. 5, 2024.
- [29] M. M. Van Der Krogt, W. W. De Graaf, C. T. Farley, C. T. Moritz, L. Richard Casius, and M. F. Bobbert, "Robust passive dynamics of the musculoskeletal system compensate for unexpected surface changes during human hopping," *Journal of applied physiology*, vol. 107, no. 3, pp. 801–808, 2009.
- [30] K. T. Kalveram and A. Seyfarth, "Inverse biomimetics: How robots can help to verify concepts concerning sensorimotor control of human arm and leg movements," *Journal of Physiology-Paris*, vol. 103, no. 3-5, pp. 232–243, 2009.