

# Toward Wheeled Mobility on Vertically Challenging Terrain: Platforms, Datasets, and Algorithms

Aniket Datar\*, Chenhui Pan\*, Mohammad Nazeri, and Xuesu Xiao

**Abstract**—Most conventional wheeled robots can only move in flat environments and simply divide their planar workspaces into free spaces and obstacles. Deeming obstacles as non-traversable significantly limits wheeled robots’ mobility in real-world, extremely rugged, off-road environments, where part of the terrain (e.g., irregular boulders and fallen trees) will be treated as non-traversable obstacles. To improve wheeled mobility in those environments with *vertically challenging terrain*, we present two wheeled platforms with little hardware modification compared to conventional wheeled robots; we collect datasets of our wheeled robots crawling over previously non-traversable, vertically challenging terrain to facilitate data-driven mobility; we also present algorithms and their experimental results to show that conventional wheeled robots have previously unrealized potential of moving through vertically challenging terrain. We make our platforms, datasets, and algorithms publicly available to facilitate future research on wheeled mobility.<sup>1</sup>

## I. INTRODUCTION

Building mobile robots that are capable of reaching as many places as possible has long been a dream for many robotics researchers. Indeed, autonomous mobile robots have ventured into remote deserts for scientific exploration [1], explored extraterrestrial planets to look for signs of life [2], [3], and assisted with search and rescue missions in hazardous or difficult-to-reach environments [4]–[7].

One particular thrust in this area of research is the development of ground robots capable of navigating vertically challenging terrain (e.g., steep slopes, rocky outcroppings, and uneven surfaces) [8]–[10]. Achieving reliable and robust mobility in these environments is challenging due to the intricate nature of the terrain, the complex vehicle-terrain interactions, the adverse impact caused by gravity, the potential deformation of the vehicle chassis, and the varying traction between the wheels and the terrain. Despite these difficulties, such mobility has been made possible mainly through advancement in hardware, including the development of specialized robot platforms [3], [4], [10]–[12] and the use of new materials [13]. These robots are capable of climbing walls [13], scaling cliffs [10], and traversing rough terrain with ease [11], [12], making them suitable for a wide range of real-world applications.

However, despite these hardware advances, the vast majority of currently available ground robots are still conventional wheeled platforms, whose mobility is mostly limited to flat surfaces. Most existing autonomous navigation systems for wheeled robots merely divide their presumably planar



Fig. 1: The Verti-Wheelers: Conventional Wheeled Vehicles Moving through Vertically Challenging Terrain.

workspaces into free spaces (traversable) or obstacles (non-traversable), significantly reducing these robots’ reachability in the real world, especially outdoor off-road environments where vertical protrusions from the ground are not uncommon. Our work is motivated by such limitations and aims at expanding the mobility of these widely available wheeled robot platforms so that they can venture into vertically challenging environments, which would otherwise be deemed as obstacles (non-traversable) or require specialized hardware. Note that the *vertically challenging* terrain we are interested in conquering requires driving wheeled robots over irregular and complex obstacles and is therefore much more difficult compared to simply driving on *non-flat* environments.

To this end, we first present an open-source design of two wheeled robot platforms, the Verti-Wheelers, which are representative of the majority of existing conventional ground mobile robot platforms, and hypothesize that conventional wheeled robots can also navigate many vertically challenging terrain, which are normally considered as non-traversable obstacles (Fig. 1). Second, considering the difficulty in creating accurate analytical models to solve this problem, we collect and make publicly available datasets to facilitate data-driven approaches. Third, we present three algorithms to autonomously drive wheeled robots over vertically challenging terrain: an open-loop, a classical rule-based, and an end-to-end learning-based approach, which can be used as baselines to benchmark future research. We provide extensive experimental results of all three algorithms in both an indoor controlled testbed and outdoor natural environments, confirm our hypothesis, and point out future research directions for wheeled mobility on vertically challenging terrain. To the

\*Equally contributing authors. All authors are with the Department of Computer Science, George Mason University {adatar, cpan7, mnazerir, xiao}@gmu.edu

<sup>1</sup>Website: <https://cs.gmu.edu/~xiao/Research/Verti-Wheelers/>

best of our knowledge, no existing work has tackled vertically challenging terrain with conventional wheeled robots in the real world.

## II. RELATED WORK

This section discusses conventional wheeled robot mobility, reviews novel hardware to expand ground mobility beyond the capability of conventional wheeled platforms, and surveys data-driven approaches to improve robot mobility.

### A. Conventional Wheeled Mobility

Due to their simplicity and efficiency, wheeled robots have been widely used in various applications such as scientific exploration [1], autonomous delivery [14], and search and rescue [8]. Equipped with differential-drive mechanism [15], Ackermann steering [16], or omnidirectional wheels [17], these robots can move through their planar workspaces and reach their goals.

One basic capability for wheeled robots is obstacle avoidance. Roboticists have developed autonomous navigation systems which first divide their planar workspaces into free spaces and obstacles and then move robots to their goals without collisions [18]–[22]. Researchers have also investigated off-road navigation [23]–[26], in which the planar workspace includes a variety of terrain types, such as gravel, grass, mud, sand, and snow. Instead of the binary free/obstacle depiction of the planner workspace, off-road navigation systems usually build semantic or traversability maps [27], [28], e.g., gravel is better to drive on than grass, which is better than mud. However, for vertical protrusions from the ground, e.g., large boulders or fallen tree trunks, these systems still treat them as non-traversable obstacles due to mobility limitations caused by the wheels.

However, in real-world unstructured environments, vertical protrusions from the ground are not necessarily non-traversable, even for wheeled robots. Consider the motor-sport of rock crawling, in which human drivers are able to drive off-road vehicles over obstacles of similar size to the vehicles themselves. Deeming all these vertically challenging terrain as completely non-traversable largely limits the mobility of wheeled robots in the real world. In this work, we aim at expanding wheeled mobility to these previously impossible, vertically challenging terrain.

### B. Novel Hardware for Vertically Challenging Terrain

One way to improve ground mobility is through novel hardware design. To overcome the limitations of wheels, researchers have developed robots with alternative actuation: wheels can be replaced by legs when facing extremely rugged terrain [29], [30]; active suspensions are widely used on planetary rovers to achieve better maneuverability, reconfigurability, and therefore mobility by allowing the chassis to actively conform to different underlying terrain [31], [32]; tracked robots are not particularly novel, but the increased track surface contact compared to wheels makes them more appropriate on vertically challenging terrain, while compromising efficiency on flat surfaces. Another thrust is to




Fig. 2: Components of the Verti-Wheelers.

develop new adhesive materials so that robots can adhere to steep or even vertical surfaces [13].

Despite their superior mobility on vertically challenging environments, these costly, specialized hardware require extra engineering effort and may adversely affect vehicles' mobility and efficiency on flat environments. Considering that most ground robots are wheeled with no or passive suspension systems, our research aims at equipping those conventional wheeled platforms with enhanced mobility to autonomously move through vertically challenging terrain.

### C. Data-Driven Robot Mobility

Thanks to the recent advancement in machine learning, data-driven approaches have been leveraged to improve robot mobility [33]. Researchers have investigated using learning to achieve adaptive navigation in a variety of environments [34]–[40], agile navigation in highly constrained spaces [41]–[43], high-speed off-road navigation [23]–[26], visual-only navigation [44]–[47], and socially compliant navigation [34], [48]–[54]. Learning from data using either imitation learning [36], [44], [55] or reinforcement learning [56]–[58] removes the necessity of building analytical models of the environments, such as vehicle-terrain or human-robot interactions, and alleviates the burden of crafting delicate cost functions [34], [59] or tuning unintuitive parameters [35]–[39].

The problem of autonomously driving wheeled robots through vertically challenging terrain shares many aforementioned difficulties. Therefore, we hypothesize that data-driven approaches are one avenue toward enabling enhanced wheeled mobility on previously impossible, vertically challenging terrain.

### III. PLATFORMS AND DATASETS

Considering that most existing mobile robots are wheeled, our goal is to equip these easily available wheeled platforms with no or little specialized hardware modification. To assure it is physically feasible for wheeled platforms to move through vertically challenging terrain, we identify the following seven desiderata for their hardware:

**All-Wheel Drive (D1):** Due to the high possibility of losing wheel contact with the irregular terrain, our wheeled platforms need to be all-wheel drive to maximize traction.

**Independent Suspensions (D2):** Considering the non-flatness of vertically challenging terrain, independent suspensions allow the chassis to conform to and most wheels to contact with the ground to maintain stability and traction.

**Differential Lock (D3):** While traditional differentials aim at increasing efficiency during vehicle turning on flat surfaces, losing left wheel contact will cause the right wheel to lose traction, or vice versa. Therefore, our wheeled platforms need to be able to lock the differential(s) when necessary.

**Low/High Gear (D4):** Although high gear improves energy efficiency on flat surfaces, to mitigate downward force caused by gravity, our platforms require high torque and low gear on vertically challenging terrain.

**Wheel Speed / RPM Sensing (D5):** Precise end-point wheel speed measurement (instead of motor encoder readings that will be skewed by the drivetrain and differential) allows the vehicle to sense wheel slippage before losing wheel contact.

**Ground Speed Sensing (D6):** Our platforms need an easy way to identify when the vehicle is getting stuck, which is very likely to happen on vertically challenging terrain.

**Actuated Perception (D7):** Due to the increased vehicle pose variability (i.e., pitch and roll in addition to yaw) when crawling over vertically challenging terrain, traditional fixed onboard sensors may lose sight of the terrain. Therefore, our platforms need to be able to point their sensors toward the terrain, regardless of how the vehicle pose varies.

#### A. Verti-6-Wheeler (V6W)

For the mechanical components in D1 to D4, we base our platform on an off-the-shelf, three-axle, six-wheel, all-wheel-drive, off-road vehicle chassis from Traxxas. The length of the chassis is 0.863m, with 0.471m front-to-middle and 0.603m front-to-rear axle wheelbase. The total height and width of V6W after outfitting all mechanical and electrical components is 0.200m and 0.249m. D1 and D2 are therefore achieved. We use an Arduino Mega micro-controller to lock/unlock the front and rear differential (D3) and switch between low and high gear (D4) through three servos. For D5, we install four magnetic sensors on the front and middle axles, and eight magnets per wheel to sense the wheel rotation. We choose not to install magnetic sensors on the rear axle considering the rear differential lock is shared by both middle and rear axles. For D6, we install a Crazyflie Flow deck v2 sensor on the chassis facing downward, providing not only 2D ground speed ( $x$  and  $y$ ) but also distance between the sensor and the ground ( $z$ ).

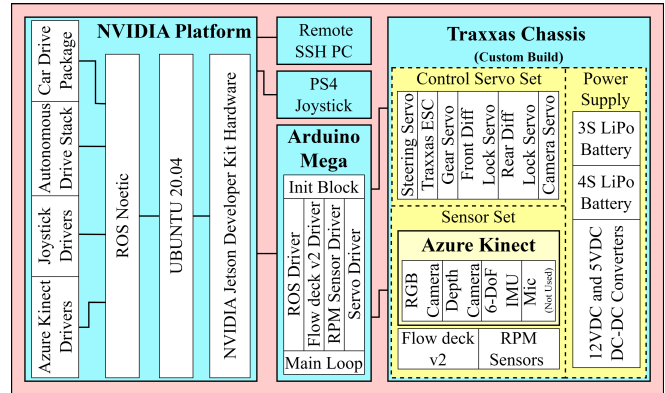


Fig. 3: System Architecture of the Verti-Wheelers.

We choose an Azure Kinect RGB-D camera due to its high-resolution depth perception at close range. For D7, we add a tilt joint for the camera actuated by a servo. We use a complementary filter and camera Inertia Measurement Unit (IMU) readings to estimate the camera orientation and a PID controller to regulate the camera pitch angle. For the core computation unit, we use an NVIDIA Jetson AGX Orin to provide both onboard CPU and GPU computation. To interface all low-level sensors and actuators, we use an Arduino Mega micro-controller.

#### B. Verti-4-Wheeler (V4W)

Considering that most wheeled robots have four wheels, we also build a four-wheeled platform based on an off-the-shelf, two axle, four-wheel-drive, off-road vehicle from Traxxas. The length of the V4W chassis is 0.523m, with a 0.312m wheelbase. The total height and width of V4W is 0.200m and 0.249m. Most components remain the same as the six-wheeler, but to accommodate the small footprint and payload capacity of the four-wheeler, we replace the NVIDIA Jetson AGX Orin with a Xavier NX. The mechanical and electrical components and system architecture for both V4W and V6W are shown in Fig. 2 and 3 respectively.<sup>2</sup>

#### C. Testbed

In addition to outdoor field experiments, we construct a custom-built indoor testbed for vertically challenging terrain for controlled and repeatable experiments: hundreds of rocks and boulders of an average size of 30cm (at the same scale of the V6W and V4W) are randomly laid out and stacked up on a  $3.1 \times 1.3$ m test course. The highest elevation of the test course can reach up to 0.5m, more than twice the height of both vehicles, as shown in Fig. 4. The testbed is highly reconfigurable by shuffling the pose of each rock/boulder.

#### D. Datasets

Considering the difficulty in representing surface topography and modeling complex vehicle dynamics and the recent success in data-driven mobility [33], we collect two datasets

<sup>2</sup><https://github.com/RobotiXX/Verti-Wheelers>



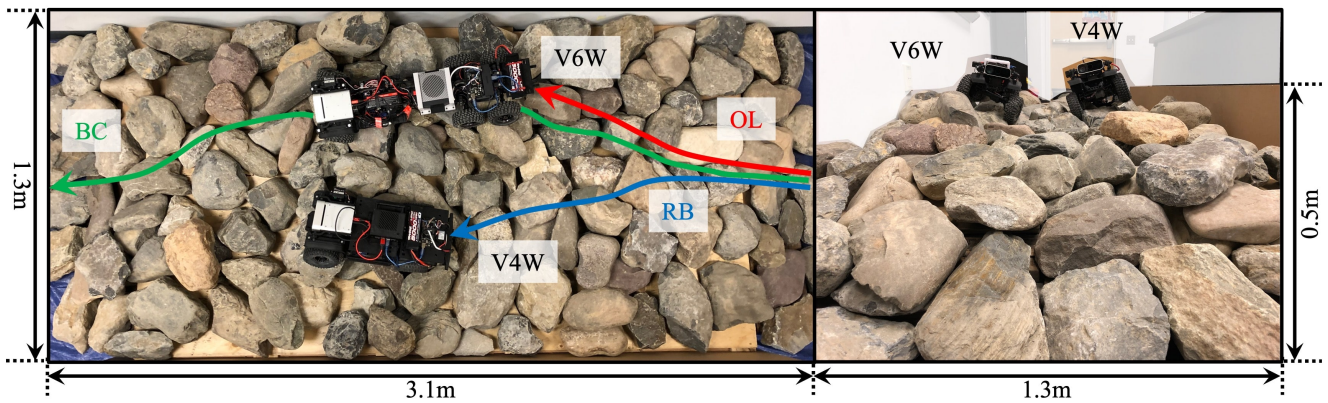


Fig. 4: Custom-Built Testbed with V6W and V4W and Example Traversals by the Three Algorithms (OL, RB, and BC).

with the two wheeled robots on our custom-built testbed.<sup>3</sup> We reconfigure our testbed multiple times and both robots are manually driven through different vertically challenging terrain. We collect the following data streams from the onboard sensors and human teleoperation commands: RGB ( $1280 \times 720 \times 3$ ) and depth ( $512 \times 512$ ) images  $i$ , wheel speed  $w$  (4D float vector for four wheels), ground speed  $g$  (relative movement indicators along  $\Delta x$  and  $\Delta y$  and displacement along  $z$ , along with two binary reliability indicators for speeds and displacement), differential release/lock  $d$  (2D binary vector for both front and rear differentials), low/high gear switch  $s$  (1D binary vector), linear velocity  $v$  (scalar float number), and steering angle  $\omega$  (scalar float number). Each dataset  $\mathcal{D}$  is therefore  $\mathcal{D} = \{i_t, w_t, g_t, d_t, s_t, v_t, \omega_t\}_{t=1}^N$ , where  $N$  indicates the total number of data frames.

The initial release of the V6W and V4W datasets include 50 and 64 teleoperated trials of 46667 and 70143 data frames of the 6-wheeler and 4-wheeler crawling over different vertically challenging rock/boulder courses respectively. To assure the human demonstrator has only access to the same perception as the robot, teleoperation is conducted in a first-person-view from the onboard camera, rather than a third-person-view. Fig. 5 shows an example data frame of the RGB and depth images  $i$  of the V4W dataset.

With a more mechanically capable chassis and two more wheels, the demonstrator can demonstrate crawling behaviors on V6W at ease, i.e., mostly driving forward, slowing down when approaching an elevated terrain patch, and only using the steering to circumvent very difficult obstacles or ditches in front of V6W. On the other hand, the demonstration of V4W takes much more effort, and the demonstrator needs to carefully control both the linear velocity and steering angle at fine resolution to negotiate through a variety of vertically challenging terrain.

#### IV. ALGORITHMS

Requiring no or little specialized hardware in addition to conventional wheeled platforms, we develop three algorithms

<sup>3</sup><https://dataverse.orc.gmu.edu/dataset.xhtml?persistentId=doi:10.13021/orc2020/QSN50Q>



Fig. 5: Example RGB (Left) and Depth (Right) Images in the Verti-Wheelers Dataset.

to drive our robots through vertically challenging terrain: an open-loop, a classical rule-based, and an end-to-end learning-based controller.

##### A. Open-Loop Controller

As a baseline, we implement an open-loop controller that drives the robots toward vertically challenging terrain previously deemed as non-traversal obstacles, simply treating them as free spaces. Our open-loop controller locks the differentials and uses the low gear all the time. We set a constant linear velocity to drive the robot forward. No onboard perception is used for the open-loop controller.

##### B. Classical Rule-Based Controller

We design a classical rule-based controller based on our heuristics on off-road driving: we lock the corresponding differential when we sense wheel slippage; we use the low gear when ascending steep slopes; when getting stuck on rugged terrain, we first increase the wheel speed and attempt to move the robot forward beyond the stuck point; if unsuccessful, we then back up the robot to get unstuck, and subsequently try a slightly different route. With the hardware described in Sec. III, our robots are able to perceive all aforementioned information with the onboard sensors and initiate corresponding actions. The finite state machine of the rule-based controller is shown in Fig. 6.

##### C. End-to-End Learning-Based Controller

We also develop an end-to-end learning-based controller to enable data-driven mobility. We aim at learning a motion




Fig. 6: Finite State Machine of the Rule-Based Controller

policy that maps from the robots’ onboard perception to raw motor commands to drive the robots over vertically challenging terrain. Utilizing the datasets we collect, we adopt an imitation learning approach, i.e., Behavior Cloning [60], to regress from perceived vehicle state information to demonstrated actions.

Our motion policy  $\pi_\theta : \mathcal{X} \rightarrow \mathcal{A}$ , is a mapping from vehicle state space  $\mathcal{X}$  to vehicle action space  $\mathcal{A}$ , parameterized by a set of parameters  $\theta$ . To instantiate a vehicle state  $x \in \mathcal{X}$ , we use all (or a subset of) onboard perception, including RGB and depth image  $i$ , wheel speed  $w$ , and ground speed  $g$ , while the vehicle action  $a \in \mathcal{A}$  includes differential release/lock  $d$ , low/high gear switch  $s$ , and most importantly, linear velocity  $v$  and steering angle  $\omega$ :

$$x = (i, w, g), \quad a = (d, s, v, \omega). \quad (1)$$

To train  $\pi_\theta$ , we minimize an end-to-end behavior cloning loss to seek for the optimal parameter set  $\theta^*$ :

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \sum_{(x,a) \in \mathcal{D}} \|a - \pi_\theta(x)\|_H, \quad (2)$$

where  $\|v\|_H = v^T H v$  is the induced norm by a diagonal matrix  $H$  with positive real entries to weigh each dimension of the action. We instantiate  $\pi_\theta$  as a neural network and use backpropagation to find the appropriate set of parameters  $\theta^*$ .

## V. EXPERIMENTS

All three algorithms in Sec. IV are implemented on both V6W and V4W in Sec. III, with the end-to-end learning-based controller learned from the collected datasets.

### A. Implementations

Considering the extremely rugged testbed (Fig. 4), all three methods lock both differentials and use low gear. The open-loop controller (OL) uses a constant linear velocity of 0.5m/s. When the rule-based (RB) controller gets stuck for 2 seconds, the wheel speed is gradually increased from 0.5m/s to 0.75m/s. If unsuccessful for 3 seconds, the robots back up for 2 seconds at 0.5m/s, and then continue forward at 0.75m/s steering  $18^\circ$  left for 2 seconds and right for 2 seconds.

For the end-to-end Behavior Cloning (BC) controller, we use both datasets on both vehicles. Specifically, in addition to using their own corresponding dataset, we also cross-deploy the learned model on the other platform, i.e., learned model from V6W data deployed on V4W and vice versa. Such a cross-deployment aims at revealing the relationship between different mechanical capabilities and qualities of the training data. Similar to OL and RB, BC maintains locked differentials and low gear. Our BC implementation takes vehicle state  $x = (i)$  (depth image only) and action

$a = (v, \omega)$  (Eqn. 1) and seeks optimal parameter  $\theta^*$  (Eqn. 2) in a neural network. BC uses a RESNET18 to process the  $224 \times 224$  depth images (down sampled from  $512 \times 512$ ) before feeding the 512-dimensional embedding into four fully connected layers with 256, 128, 64, and 2 neurons with ReLU activation except the last layer.

### B. Results

Three different test courses are built by reconfiguring the rocks/boulders on the testbed (Fig. 4), whose difficulty levels range from easy, medium, to difficult. For the difficult level, we add wooden blocks to the rock/boulder course, which do not exist in the datasets, to test BC’s generalizability. For each test course, we run four different approaches, i.e., OL, RB, BC, and BC with cross-deployment (we denote the model trained with the V6W and V4W dataset as BC6 and BC4 respectively), each ten trials, running from both directions of the course. We report both number of successful trials (out of 10 attempts) and mean traversal time (for the successful trials in seconds) with variance of all 240 experiment trials in Tab. I. A failure trial can either be the vehicle getting stuck or tipping over on the test course. For all four approaches, we start the vehicles at the same starting location and orientation facing the test course.

For a certain vehicle on a certain difficulty level, in general BC achieves higher success rate than both OL and RB, with OL most frequently getting stuck or tipping over. Among all successful trials, BC mostly achieves the shortest traversal time, but not always, because BC learns to slow down to smoothly go through rugged terrain while OL and RB may drive aggressively. With increasing difficulty level, success rate decreases and traversal time increases, although V6W’s performance is not very sensitive to the Easy and Medium levels, considering its superior mechanical capability.

One very interesting finding is the cross-deployment results. While V4W with its own dataset (BC4) outperforms V4W cross-trained with V6W’s dataset (BC6), V6W’s performance is opposite: V6W performs better when cross-trained with V4W’s dataset (BC4) compared to when being trained with its own dataset (BC6). Such a discrepancy is caused by the different demonstrations in the V6W and V4W datasets, which are further caused by the difference in the two vehicles’ mechanical capabilities: as mentioned above, the demonstrator can drive the more mechanically capable V6W through different test courses at ease without the need to constantly adjust the speed and steering, while the limited mechanical capability of V4W requires fine-grained control, causing the quality of the BC4’s dataset to exceed that of BC6’s. Therefore, even V6W, a different vehicle, can achieve improved mobility with BC4, compared to BC6 with its own dataset. On the other hand, learning from the less careful driving demonstrations with BC6 jeopardizes the learned mobility of V4W.

### C. Outdoor Mobility Demonstration

We also deploy both Verti-Wheelers in outdoor vertically challenging environments beyond the controlled indoor

TABLE I: Number of Successful Trials (Out of 10) and Mean Traversal Time (of Successful Trials in Seconds) with Variance

	V6W				V4W			
	OL	RB	BC6	BC4	OL	RB	BC6	BC4
Easy	5 (20.7 ± 1.7)	8 (19.2 ± 3.9)	9 (13.8 ± 8.2)	<b>10</b> (11.6 ± 1.9)	6 (17.7 ± 3.8)	6 (13.4 ± 2.5)	7 (17.2 ± 6.7)	<b>9</b> (14.1 ± 7.7)
Medium	6 (15.4 ± 0.9)	9 (14.8 ± 2.2)	9 (14.6 ± 11.2)	<b>10</b> (13.6 ± 2.3)	4 (15.6 ± 14.2)	6 (12.9 ± 1.8)	3 (19.2 ± 10.6)	<b>8</b> (13.7 ± 1.6)
Difficult	3 (24.1 ± 2.6)	6 (14.3 ± 1.9)	6 (15.7 ± 18.5)	<b>9</b> (14.9 ± 2.9)	3 (19.7 ± 29.4)	5 (16.8 ± 20.5)	3 (23.3 ± 43.4)	<b>7</b> (14.9 ± 8.2)



Fig. 7: Outdoor Demonstration: V4W’s right front wheel suspends in the air (left) and V6W climbs a steeper rock pile than the indoor testbed.

testbed (Fig. 7). The outdoor environments contain different rock sizes, including both smaller pebbles and larger boulders compared to the indoor testbed. Furthermore, the outdoor environments are more jagged, including larger ditches between rocks that may easily get the vehicles stuck (Fig. 7 left) and steeper slopes that require stable traction between all wheels and the terrain (Fig. 7 right). While OL and RB can drive smoothly on small pebbles, they often get stuck on larger and more jagged boulders. BC is able to generalize well to rocks of similar sizes as in the indoor testbed, but may occasionally get stuck on large boulders, wide ditches, and steep slopes as well, showing future research is still required to improve wheeled mobility on vertically challenging terrain in the real world.

## VI. CONCLUSIONS AND FUTURE WORK

This paper presents two vehicle platforms, demonstration datasets, and three algorithms to unlock the previously unrealized potential of wheeled mobility on vertically challenging terrain. The two platforms, with minimal hardware modification to conventional wheeled robots, achieve seven desiderata to tackle such challenging environments. We open-source our hardware design, software implementation, and demonstration datasets to facilitate future research. The experiment results of the three algorithms confirm our hypothesis that conventional wheeled robots have the mechanical capability of navigating many vertically challenging terrain, which are normally considered as non-traversable obstacles, especially with the help of data-driven approaches.

This paper opens up a new direction of research to achieve superior robot mobility beyond their original design with only limited mechanical capability. Despite the fact that the simple Behavior Cloning algorithm works well on many vertically challenging terrain, it still fails in many more challenging scenarios, such as in more difficult test courses than the Difficult level in Tab. I or in the natural

outdoor environments (Fig. 7). However, these environments can be conquered by human teleoperation and are therefore within the vehicles’ mechanical limit. One promising future research direction is to explicitly model terrain traversability instead of implicitly learning it in a motion policy. It is also possible to design or learn a vehicle dynamics model on vertically challenging terrain. Another interesting direction is to combine autonomous crawling and goal-oriented navigation with obstacle avoidance, i.e., real obstacles that are absolutely beyond the robots’ mechanical capability. Last but not least, considering the interesting findings in the BC cross-deployment experiments, Verti-Wheelers provide ideal physical robot platforms for research in transfer learning for cross-platform mobility, learning with embodiment mismatch, and meta learning for general vehicle mobility and maneuverability.

## ACKNOWLEDGEMENTS

This work has taken place in the RobotiXX Laboratory at George Mason University. RobotiXX research is supported by Army Research Office (ARO, W911NF2220242, W911NF2320004, W911NF2420027), US Air Forces Central (AFCENT), Google DeepMind (GDM), Clearpath Robotics, and Raytheon Technologies (RTX).

## REFERENCES

- [1] D. S. Wettergreen and D. R. Thompson, “Science on the fly: Enabling science autonomy during robotic traverse,” in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 1110–1111.
- [2] M. Maurette, “Mars rover autonomous navigation,” *Autonomous Robots*, vol. 14, no. 2-3, pp. 199–208, 2003.
- [3] M. R. Islam, F. H. Chowdhury, S. Rezwana, M. J. Ishaque, J. U. Akanda, A. S. Tuhel, and B. B. Riddhe, “Novel design and performance analysis of a mars exploration robot: Mars rover mongol pothik,” in *2017 Third International Conference on Research in Computational Intelligence and Communication Networks (ICRCICN)*. IEEE, 2017, pp. 132–136.
- [4] X. Xiao, E. Cappelletti, W. Zhen, J. Dai, K. Sun, C. Gong, M. J. Travers, and H. Choset, “Locomotive reduction for snake robots,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3735–3740.
- [5] R. Murphy, J. Dufek, T. Sarmiento, G. Wilde, X. Xiao, J. Braun, L. Mullen, R. Smith, S. Allred, J. Adams *et al.*, “Two case studies and gaps analysis of flood assessment for emergency management with small unmanned aerial systems,” in *2016 IEEE international symposium on safety, security, and rescue robotics (SSRR)*. IEEE, 2016, pp. 54–61.
- [6] X. Xiao, J. Dufek, T. Woodbury, and R. Murphy, “Uav assisted usv visual navigation for marine mass casualty incident response,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6105–6110.
- [7] X. Xiao, J. Dufek, and R. R. Murphy, “Autonomous visual assistance for robot operations using a tethered uav,” in *Field and Service Robotics: Results of the 12th International Conference*. Springer, 2021, pp. 15–29.
- [8] R. R. Murphy, *Disaster robotics*. MIT press, 2014.



- [9] X. Xiao and R. Murphy, "A review on snake robot testbeds in granular and restricted maneuverability spaces," *Robotics and Autonomous Systems*, vol. 110, pp. 160–172, 2018.
- [10] P. McGarey, F. Pomerleau, and T. D. Barfoot, "System design of a tethered robotic explorer (trex) for 3d mapping of steep terrain and harsh environments," in *Field and Service Robotics: Results of the 10th International Conference*. Springer, 2016, pp. 267–281.
- [11] C. Zheng, S. Sane, K. Lee, V. Kalyanram, and K. Lee, " $\alpha$ -waltr: Adaptive wheel-and-leg transformable robot for versatile multiterrain locomotion," *IEEE Transactions on Robotics*, 2022.
- [12] K. Xu, Y. Lu, L. Shi, J. Li, S. Wang, and T. Lei, "Whole-body stability control with high contact redundancy for wheel-legged hexapod robot driving over rough terrain," *Mechanism and Machine Theory*, vol. 181, p. 105199, 2023.
- [13] Y. Liu and T. Seo, "Anyclimb-ii: Dry-adhesive linkage-type climbing robot for uneven vertical surfaces," *Mechanism and Machine Theory*, vol. 124, pp. 197–210, 2018.
- [14] Starship, "Starship," <https://www.starship.xyz/>, 2023.
- [15] S. K. Malu and J. Majumdar, "Kinematics, localization and control of differential drive mobile robot," *Global Journals of Research in Engineering*, vol. 14, no. H1, pp. 1–7, 2014.
- [16] P. Atreya, H. Karnan, K. S. Sikand, X. Xiao, S. Rabiee, and J. Biswas, "High-speed accurate robot control using learned forward kinodynamics and non-linear least squares optimization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 11 789–11 795.
- [17] J. Lu, X. Tuo, H. Liu, and F. Shi, "Design and implementation of a spherical robot for radiation monitoring," in *2022 5th International Conference on Pattern Recognition and Artificial Intelligence (PRAI)*. IEEE, 2022, pp. 1261–1266.
- [18] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [19] S. Quinlan and O. Khatib, "Elastic bands: Connecting path planning and control," in *[1993] Proceedings IEEE International Conference on Robotics and Automation*. IEEE, 1993, pp. 802–807.
- [20] X. Xiao, Z. Xu, Z. Wang, Y. Song, G. Warnell, P. Stone, T. Zhang, S. Ravi, G. Wang, H. Karnan *et al.*, "Autonomous ground navigation in highly constrained spaces: Lessons learned from the benchmark autonomous robot navigation challenge at icra 2022 [competitions]," *IEEE Robotics & Automation Magazine*, vol. 29, no. 4, pp. 148–156, 2022.
- [21] D. Perille, A. Truong, X. Xiao, and P. Stone, "Benchmarking metric ground navigation," in *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2020, pp. 116–121.
- [22] A. Nair, F. Jiang, K. Hou, Z. Xu, S. Li, X. Xiao, and P. Stone, "Dynabarn: Benchmarking metric ground navigation in dynamic environments," in *2022 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2022, pp. 347–352.
- [23] Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. A. Theodorou, and B. Boots, "Imitation learning for agile autonomous driving," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 286–302, 2020.
- [24] X. Xiao, J. Biswas, and P. Stone, "Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 6054–6060, 2021.
- [25] M. Sivaprakasam, S. Triest, W. Wang, P. Yin, and S. Scherer, "Improving off-road planning techniques with learned costs from physical interactions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4844–4850.
- [26] H. Karnan, K. S. Sikand, P. Atreya, S. Rabiee, X. Xiao, G. Warnell, P. Stone, and J. Biswas, "Vi-ikd: High-speed accurate off-road navigation using learned visual-inertial inverse kinodynamics," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3294–3301.
- [27] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, "Real-time semantic mapping for autonomous off-road navigation," in *Field and Service Robotics*. Springer, 2018, pp. 335–350.
- [28] A. Shaban, X. Meng, J. Lee, B. Boots, and D. Fox, "Semantic terrain classification for off-road autonomous driving," in *Conference on Robot Learning*. PMLR, 2022, pp. 619–629.
- [29] M. D. Kutzer, M. S. Moses, C. Y. Brown, M. Armand, D. H. Scheidt, and G. S. Chirikjian, "Design of a new independently-mobile reconfigurable modular robot," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2758–2764.
- [30] Y. Bu, Y. Bu, H. Li, S. Mao, and H. Zhu, "Development of wheel-leg hybrid climbing robot with switchable permanent magnetic omni-wheels as feet," in *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2022, pp. 1–6.
- [31] F. Cordes, C. Oekermann, A. Babu, D. Kuehn, T. Stark, F. Kirchner, and D. Bremen, "An active suspension system for a planetary rover," in *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2014, pp. 17–19.
- [32] H. Jiang, G. Xu, W. Zeng, F. Gao, and K. Chong, "Lateral stability of a mobile robot utilizing an active adjustable suspension," *Applied Sciences*, vol. 9, no. 20, p. 4410, 2019.
- [33] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Autonomous Robots*, vol. 46, no. 5, pp. 569–597, 2022.
- [34] X. Xiao, T. Zhang, K. M. Choromanski, T.-W. E. Lee, A. Francis, J. Varley, S. Tu, S. Singh, P. Xu, F. Xia, S. M. Persson, L. Takayama, R. Frostig, J. Tan, C. Parada, and V. Sindhwani, "Learning model predictive controllers with real-time attention for real-world navigation," in *Conference on robot learning*. PMLR, 2022.
- [35] X. Xiao, Z. Wang, Z. Xu, B. Liu, G. Warnell, G. Dhamankar, A. Nair, and P. Stone, "Appl: Adaptive planner parameter learning," *Robotics and Autonomous Systems*, vol. 154, p. 104132, 2022.
- [36] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [37] Z. Wang, X. Xiao, B. Liu, G. Warnell, and P. Stone, "Appli: Adaptive planner parameter learning from interventions," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 6079–6085.
- [38] Z. Wang, X. Xiao, G. Warnell, and P. Stone, "Apple: Adaptive planner parameter learning from evaluative feedback," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7744–7749, 2021.
- [39] Z. Xu, G. Dhamankar, A. Nair, X. Xiao, G. Warnell, B. Liu, Z. Wang, and P. Stone, "Applr: Adaptive planner parameter learning from reinforcement," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 6086–6092.
- [40] B. Liu, X. Xiao, and P. Stone, "A lifelong learning approach to mobile robot navigation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1090–1096, 2021.
- [41] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1503–1510, 2021.
- [42] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 7316–7322.
- [43] Z. Wang, X. Xiao, A. J. Nettekoven, K. Umasankar, A. Singh, S. Bommakanti, U. Topcu, and P. Stone, "From agile ground to aerial navigation: Learning from learned hallucination," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 148–153.
- [44] H. Karnan, G. Warnell, X. Xiao, and P. Stone, "Voila: Visual-observation-only imitation learning for autonomous navigation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2497–2503.
- [45] G. Kahn, P. Abbeel, and S. Levine, "Badgr: An autonomous self-supervised learning-based navigation system," *arXiv preprint arXiv:2002.05700*, 2020.
- [46] D. Shah, A. Sridhar, N. Dashora, K. Stachowicz, K. Black, N. Hirose, and S. Levine, "Vint: A foundation model for visual navigation," *arXiv preprint arXiv:2306.14846*, 2023.
- [47] K. Stachowicz, D. Shah, A. Bhorkar, I. Kostrikov, and S. Levine, "Fastrlap: A system for learning high-speed driving via deep rl and autonomous practicing," *arXiv preprint arXiv:2304.09831*, 2023.
- [48] R. Mirsky, X. Xiao, J. Hart, and P. Stone, "Conflict avoidance in social navigation—a survey," *arXiv preprint arXiv:2106.12113*, 2021.
- [49] A. Francis, C. Pérez-d'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra *et al.*, "Principles and guidelines for evaluating social robot navigation algorithms," *arXiv preprint arXiv:2306.16740*, 2023.
- [50] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, "Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 807–11 814, 2022.

- [51] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1343–1350.
- [52] D. M. Nguyen, M. Nazeri, A. Payandeh, A. Datar, and X. Xiao, "Toward human-like social robot navigation: A large-scale, multi-modal, social human navigation dataset," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023.
- [53] J.-S. Park, X. Xiao, G. Warnell, H. Yedidsion, and P. Stone, "Learning perceptual hallucination for multi-robot navigation in narrow hallways," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 10 033–10 039.
- [54] J. Hart, R. Mirsky, X. Xiao, S. Tejada, B. Mahajan, J. Goo, K. Baldauf, S. Owen, and P. Stone, "Using human-inspired signals to disambiguate navigational intentions," in *International Conference on Social Robotics*. Springer, 2020, pp. 320–331.
- [55] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, and C. Cadena, "From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots," in *IEEE International Conference on Robotics and Automation*. IEEE, 2017.
- [56] Z. Xu, B. Liu, X. Xiao, A. Nair, and P. Stone, "Benchmarking reinforcement learning techniques for autonomous navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9224–9230.
- [57] Z. Xu, X. Xiao, G. Warnell, A. Nair, and P. Stone, "Machine learning methods for local motion planning: A study of end-to-end vs. parameter learning," in *2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE, 2021, pp. 217–222.
- [58] A. Faust, K. Oslund, O. Ramirez, A. Francis, L. Tapia, M. Fiser, and J. Davidson, "Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 5113–5120.
- [59] K. S. Sikand, S. Rabiee, A. Uccello, X. Xiao, G. Warnell, and J. Biswas, "Visual representation learning for preference-aware path planning," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 303–11 309.
- [60] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.