End-to-End Learning for Wheeled Mobility on Vertically Challenging Terrain

Aniket Datar*, Chenhui Pan*, Mohammad Nazeri, and Xuesu Xiao
Project Website: https://cs.gmu.edu/~xiao/Research/Verti-Wheelers/

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, non-flat, off-road environments, where part of the terrain (e.g., steep slopes or rugged boulders) will be treated as non-traversable obstacles. 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.

Thanks to the recent advancement in machine learning, data-driven approaches have been leveraged to improve robot mobility [1]. Researchers have investigated using learning to achieve adaptive navigation in a variety of environments [2]–[7], high-speed off-road navigation [8]–[11], and socially compliant navigation [2], [12]–[15]. Learning from data 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 [2] or tuning unintuitive parameters [3]. Therefore, we hypothesize that data-driven approaches are one avenue toward enabling enhanced wheeled mobility on previously impossible, vertically challenging terrain.

To this end, We develop three algorithms to autonomously drive wheeled robots over vertically challenging terrain based on our open-source design of two wheeled robot platforms, the Verti-Wheelers (Fig. 1): an Open-Loop (OL), a classical Rule-Based (RB), and an end-to-end learning-based approach, Behavior Cloning (BC).

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 vehicle differentials and uses their low gear all the time. We set a constant linear velocity to drive the robots 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 robots forward beyond the stuck point; if unsuccessful, we then back up the robots to get unstuck, and subsequently try a slightly different route.

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

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 policy that maps from the robots' onboard perception to raw motor commands to drive the robots over vertically challenging terrain. Utilizing a dataset we collect, we adopt an imitation learning approach, BC, to regress from perceived vehicle state information to demonstrated actions.

Three different test courses are built by reconfiguring the rocks/boulders on the testbed, whose difficulty levels range from easy, medium, to difficult. We report both number of successful trials and mean traversal time (for the successful trials in seconds) with variance of all experiment trials. A failure trial can either be the vehicle getting stuck or tipping over on the test course. For all three 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. For full results, we refer the readers to our full paper [16].

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

References

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] R. Mirsky, X. Xiao, J. Hart, and P. Stone, "Prevention and resolution of conflicts in social navigation–a survey," *arXiv preprint arXiv:2106.12113*, 2021.

[13] 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.

[14] 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.

[15] 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.

[16] A. Datar, C. Pan, M. Nazeri, and X. Xiao, "Toward wheeled mobility on vertically challenging terrain: Platforms, datasets, and algorithms," *arXiv preprint arXiv:2303.00998*, 2023.