



Combining Robot Experience with Human Expertise for State-Dependent Traversability Estimation

Andre Schreiber (andreams2@illinois.edu) and Katherine Driggs-Campbell (krdc@illinois.edu)
Human-Centered Autonomy Lab | University of Illinois Urbana-Champaign



HUMAN-CENTERED
AUTONOMY LAB

Introduction

Motivation

Mobile field robots must understand the traversability characteristics of their environments to safely navigate, but typical approaches do not include relevant aspects of robot state and often rely on only self-supervision or only strong human supervision

Prior Approaches

- Existing strongly-supervised approaches to traversability [1,2] typically rely on semantic segmentation
- Self-supervised methods quantify traversability via quantities like traction [3] but fail to explicitly label untraversable areas safely
- Most methods ignore crucial state components like rotation, but recent works [4] are starting to address this issue

Core Insights

- Traversability cannot be fully described using only a scalar quantity based on a robot's position! Traversability also depends on:



Speed



Direction of Travel



Rotation

- Self-supervision excels at labeling conditions that a robot can travel on, while human supervision excels at labeling lethal areas like obstacles





Our Solution: Human-Assisted Rotation-Aware Motion-Based Estimation (HARAMBE) for Off-Road Traversability

- Self-supervised labeling of traversable terrain, where traversability is predicted as linear/angular velocity expected to be attained as a function of position, rotation, and commanded linear/angular velocity
- Weak human supervision for safe, effective labeling of dangers that are predicted as a function of robot position

Architecture

Foundation model (DINOv3 [5] with features up-sampled using AnyUp [6]) computes visual features, which are projected to bird's-eye view (BEV) and fused over time to predict an obstacle map and BEV feature field that can be queried based on robot state.




Commands and BEV features at a given location are used by an MLP to predict candidate linear and angular velocities ($\tilde{v}$ and $\tilde{\omega}$), which are fed into affine transformations to predict final linear and angular velocities ($v_{\text{predicted}}$ and $\omega_{\text{predicted}}$). Each parameter for the affine transformations is computed via Fourier series whose argument is rotation (θ) and coefficients are BEV features.


Simulator Navigation Experiments

We perform navigation trials using our method for waypoint navigation of a differential-drive robot in a custom, photorealistic simulator. We use an MPPI controller like existing work [3,7] and compare against:

- Blind**: a naïve blind pursuit method
- Occupancy**: a simple occupancy mapping strategy that uses the depth camera for range measurements
- WayFASTER** [3]: a state-of-the-art self-supervised visual BEV traversability prediction method trained on the same dataset as our method
- WayFASTER*** [3]: same model as WayFASTER but trained on an expanded dataset that includes purposeful crashes to capture non-traversable conditions using self-supervision



Route for "semi-urban" sim environment. This environment is the same as used to collect training data. We run five trials per method. The start point is shown with a red circle, and the waypoints are numbered blue circles (i.e., 1 is first waypoint, 2 is second waypoint, etc.).



Route for "desert" sim environment. This environment differs significantly from that seen in training. We run five trials per method. The start is shown with a red circle, and the goal is shown as a blue circle. There are various possible obstacles (tall grass, trees, rocks) between the start and goal.

Method	Avg. Waypoints Achieved (↑)	Avg. Distance Traveled (↑)
Blind	0.0	30.9 m
Occupancy	0.6	124.3 m
WayFASTER	0.0	56.1 m
WayFASTER*	0.4	109.9 m
Obstacle-Only (w/o AnyUp)	2.4	236.3 m
HARAMBE (w/o AnyUp)	4.8	338.3 m
HARAMBE	5.0	344.0 m

Method	Success Rate (↑)	Avg. Distance for Successful Runs (↓)
Blind	0%	N/A
Occupancy	20%	20.2 m
WayFASTER	0%	N/A
WayFASTER*	60%	23.5 m
HARAMBE	100%	20.2 m

Real-World Navigation Experiments


We deploy our model on a differential-drive robot (Clearpath Jackal). We train our model using a dataset we collected on a university campus, and we test our method for navigation in environments of varying complexity (all real-world trials are in locations different from those used in training). Model inference and the MPPI controller are run in real-time onboard the robot using an Nvidia Jetson AGX Thor.



Method	Success Rate (↑)	Avg. Distance for Successful Runs (↓)
Occupancy	40%	26.4 m
WayFASTER	0%	N/A
WayFASTER*	80%	25.9 m
HARAMBE	100%	27.2 m


We perform additional comparisons between HARAMBE and our strongest baseline (WayFASTER*), running five trials per task:

- A simple task with few obstacles (2 trees), leaf litter, and tall grass in a tree grove (*left*).
- A hard task in which the robot must navigate out of a cluttered, densely-wooded area (*right*). The environmental complexity of this area is far greater than that seen in the training dataset.



Method	Success Rate (↑)	Method	Success Rate (↑)
WayFASTER*	100%	WayFASTER*	20%
HARAMBE	100%	HARAMBE	60%

Long-Distance Navigation



We run long-horizon (~550 m) navigation trials in a grove of trees using HARAMBE trained only on simulator data or only on real-world data. We ran each model (sim-trained or real-trained) 5 times and both models successfully completed the course in all trials.



We stress test HARAMBE in a forest with obstacles and significant ground clutter (roots, leaf litter, tall grass, logs/branches, etc.), and it successfully completed a ~600 m course between 5 waypoints with no interventions. Onboard images are shown at the bottom of the figure.

Conclusions

- We introduce a novel traversability prediction method (HARAMBE) that combines self-supervision and weak human supervision
- Our method predicts the anticipated achieved motion (linear/angular velocity) over traversable terrain as a function of commanded linear/angular velocity, position, and orientation
- We use human point annotation of obstacles to enable our method to safely and effectively predict untraversable conditions (e.g., obstacles and lethal terrain)
- HARAMBE outperformed an existing state-of-the-art visual BEV traversability prediction baseline in extensive navigation trials performed both in a photorealistic simulator and in real-world deployment
- We conduct long-horizon tests of our method, showing effective sim-to-real transfer and successful navigation for over 6 km of trials in challenging real-world outdoor settings

Acknowledgement

This work used Delta at the National Center for Supercomputing Applications through allocation CIS250944 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by U.S. National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296. Hardware development for this project was carried out in part in the Center for Autonomy Robotics Laboratories at the University of Illinois Urbana-Champaign (UIUC). We thank the UIUC Center for Autonomous Construction and Manufacturing at Scale for providing the robot platform used in the experiments. This material is based upon work supported by the National Science Foundation under Grant No. 2143435.

Maps throughout this poster were created using ArcGIS® software by Esri. ArcGIS® and ArcMap™ are the intellectual property of Esri and are used herein under license. Copyright © Esri. All rights reserved. For more information about Esri® software, please visit www.esri.com.

References

- Maturana *et al.*, "Real-Time Semantic Mapping for Autonomous Off-Road Navigation," in *International Conference on Field and Service Robotics*, 2017.
- Wigness *et al.*, "A RUGD Dataset for Autonomous Navigation and Visual Perception in Unstructured Outdoor Environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- Gasparino *et al.*, "WayFASTER: a Self-Supervised Traversability Prediction for Increased Navigational Awareness," in *IEEE International Conference on Robotics and Automation*, 2024.
- Dong *et al.*, "Learning Smooth State-Dependent Traversability from Dense Point Clouds," in *Conference on Robot Learning*, 2025.
- Siméoni *et al.*, "DINOv3," 2025, *arXiv:2508.10104*.
- Wimmer *et al.*, "AnyUp: Universal Feature Upsampling," in *International Conference on Learning Representations*, 2026.
- Schreiber *et al.*, "Do You Know the Way? Human-in-the-Loop Understanding for Fast Traversability Estimation in Mobile Robotics," *IEEE Robotics and Automation Letters*, 2025.