

Learning Inverse Kinodynamics for Accurate High-Speed Off-Road Navigation on Unstructured Terrain

Xuesu Xiao , Joydeep Biswas , and Peter Stone 

Abstract—This letter presents a learning-based approach to consider the effect of unobservable world states in kinodynamic motion planning in order to enable accurate high-speed off-road navigation on unstructured terrain. Existing kinodynamic motion planners either operate in structured and homogeneous environments and thus do not need to explicitly account for terrain-vehicle interaction, or assume a set of discrete terrain classes. However, when operating on unstructured terrain, especially at high speeds, even small variations in the environment will be magnified and cause inaccurate plan execution. In this letter, to capture the complex kinodynamic model and mathematically unknown world state, we learn a kinodynamic planner in a data-driven manner with onboard inertial observations. Our approach is tested on a physical robot in different indoor and outdoor environments, enables fast and accurate off-road navigation, and outperforms environment-independent alternatives, demonstrating 52.4% to 86.9% improvement in terms of plan execution success rate while traveling at high speeds.

Index Terms—Autonomous vehicle navigation, motion and path planning, machine learning for robot control.

I. INTRODUCTION

CURRENT mobile robot navigation methods can navigate a robot from one point to another safely and reliably in structured and homogeneous environments [1], [2], such as indoor hallways or outdoor paved surfaces. These consistent environments allow the robots to use simple kinodynamic motion

Manuscript received February 24, 2021; accepted June 5, 2021. Date of publication June 16, 2021; date of current version July 2, 2021. This letter was recommended for publication by Associate Editor R. Liu and Editor D. Popa upon evaluation of the reviewers' comments. This work has taken place in the Learning Agents Research Group (LARG) and Autonomous Mobile Robotics Laboratory (AMRL) at UT Austin. LARG research is supported in part by NSF (CPS-1739964, IIS-1724157, NRI-1925082), ONR (N00014-18-2243), FLI (RFP2-000), ARO (W911NF-19-2-0333), DARPA, Lockheed Martin, GM, and Bosch. AMRL research is supported in part by NSF (CAREER-2046955, IIS-1954778, SHF-2006404), ARO (W911NF-19-2-0333), DARPA (HR001120C0031), Amazon, JP Morgan, and Northrop Grumman Mission Systems. Peter Stone serves as the Executive Director of Sony AI America and receives financial compensation for this work. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research. (*Corresponding author: Xuesu Xiao*)

Xuesu Xiao and Joydeep Biswas are with the Department of Computer Science, The University of Texas at Austin, Austin TX 78712. (e-mail: xiao@cs.utexas.edu; joydeepb@cs.utexas.edu).

Peter Stone is with the Department of Computer Science, The University of Texas at Austin, Austin TX 78712, and also with Sony AI. (e-mail: pstone@cs.utexas.edu).

Digital Object Identifier 10.1109/LRA.2021.3090023

planners independent of the environment, thanks to the limited environment disturbances and stochasticity.

Dating back at least to DARPA's Grand Challenge [3] and LAGR (Learning Applied to Ground Vehicles) [4] program, researchers have also looked into applying autonomous navigation in unstructured outdoor environments. Challenges arise from multiple fronts in those natural spaces, but most off-road navigation work focused on perception, e.g., detecting natural obstacles [4], classifying underlying terrain types [5]–[7], or building semantic maps [8]–[10]. For motion control, most off-road robots simply travel at low speeds to minimize uncertainty and to maximize safety [11]–[14].

While recent advances in deep learning provide roboticists with a different avenue to investigate those perception problems in off-road navigation, researchers also started to combine perception, planning, and motion control using end-to-end learning in unstructured environments [15], [16]. These systems do not require a heavily-engineered navigation pipeline, and can react to natural environments in a data-driven approach. Although these methods can enable successful navigation, they are data-intensive and generally do not lead to better navigation than their classical counterparts.

Focusing on the motion control side of off-road navigation on unstructured terrain, the contribution of this letter is to use learning to capture the effect of complex and unknown environmental factors on the robot's low-level kinodynamic model. We use onboard inertial observation to encode environmental factors and learn an inverse kinodynamic model to produce fast and accurate low-level control inputs. Using our method, even with extensive disturbances caused by high-speed terrain-vehicle interaction, the robot is still able to perform fast and accurate off-road navigation (Fig. 1), compared to a kinodynamic model based on ideal assumptions and a learned model that does not consider environmental factors.

II. RELATED WORK

In this section, we review related off-road navigation work in terms of perception and motion control.

A. Off-Road Perception

The first challenge arises from off-road navigation is perception. In unstructured off-road environments, perception is no longer simply in the geometric sense (e.g., free vs. obstacle),




Fig. 1. The UT Automata scale 1/10th autonomous vehicle drives an eight-turn (T1–T8 on the red path) outdoor race track with unstructured terrain. Close-ups of some terrain (black inserts) are shown, including cement, mud, grass, which are covered by leaves, stalks, and/or twigs at different densities. For high-speed terrain-aware navigation, the robot has to interact with these unstructured terrain (red insert).

but also requires semantic information (e.g., grass vs. mud). A plethora of research in terrain classification has leveraged vibration-based signals to classify terrain types [5]–[7]. Vision-based sensors, e.g., LiDAR and camera, combined with current deep learning methods, have also been used to build semantic maps [8]–[10]. These perception methods assign costs to discrete terrain classes for planning, but do not consider robot’s kinodynamic model when moving on these terrain. Our work does not distinguish among discrete terrain classes and uses observations collected during interactions with different terrain to enable fast and accurate kinodynamic planning.

B. Off-Road Motion Control

Although research thrust for off-road navigation has been primarily focused on perception, roboticists have also investigated off-road navigation from the motion control side. Many wheel slip models [11]–[14] have been developed and used to design controllers to minimize slip. Most of these models treat slip only as a function of the vehicle kinematics. But to achieve high-speed off-road navigation, slip is inevitable and also highly dependent on the underlying terrain.

Researchers have also used machine learning for motion control in off-road navigation. A recent survey [17] pointed out that learning is most efficient when targeting at navigation components, e.g., learning local planners [18]–[23], costmaps [24], [25], or planner parameters [26]–[30]. Research on using learning for motion control in off-road scenarios is scarce. Pan, *et al.* [31] enabled high-speed navigation with end-to-end imitation learning from RGB input in a closed circular dirt track. The expert demonstrator is a model predictive controller with access to high-precision sensors including GPS-RTK. The end-to-end learning approach most likely does not generalize well to other terrain and tracks. Aiming at a variety of terrain, Siva, *et al.* [32] used imitation learning from human demonstration to navigate five discrete terrain types (concrete, grass, mud, pebble, rock). In contrast, our method only targets at learning a kinodynamic

model and can navigate any global path. We also do not intentionally separate terrain into discrete types, and treat different terrain characteristics in a continuous manner. Reinforcement learning has also shown potential in learning motion control for navigation, but at the cost of extensive training overhead: Brunnbauer, *et al.* [33] took 8 million time steps (37 hours of interactions) to learn a simple *kinematic* model. In comparison, our approach only requires 30 minutes to learn a *kinodynamic* model, which would require far more training by such reinforcement learning approaches.

III. LEARNING INVERSE KINODYNAMIC MODELS

Most navigation systems either assume kindodynamic models to be independent of the environment, or that there exists a discrete set of environment classes [32], e.g., one model for paved terrain and another for grass. In this work, we relax these assumptions by learning a single continuous inverse kinodynamic model that accounts for environmental factors across different terrain types without having to perform discrete terrain classification, or analytic modelling. The learned model takes as input inertial observations that make the impact of environmental factor on kinodynamic motion observable (e.g., how bumpiness from gravel will result in understeer at high speeds). This inverse kinodynamic model is learned in a data-driven manner.

A. Problem Formulation

Given vehicle state x , control input u , and world state w , the state dynamics and observation y are given by

$$\dot{x} = f(x, u, w), \quad y = g(x, w), \quad (1)$$

where $f(\cdot, \cdot, \cdot)$ is the system’s forward kinodynamic function, while $g(\cdot, \cdot)$ is the observation function. Note that in most cases, w is not directly observable and cannot be easily modeled. A navigation planner generates a global plan $\Pi: [0, 1] \rightarrow X$ mapping from a unitless progress variable $s \in [0, 1]$ to planned vehicle state $x \in X$, incorporating both global (e.g., traversable map) and local (e.g., sensed obstacles) information to take the robot from the start state $\Pi(0)$ to the goal state $\Pi(1)$. A projection operator $\rho: X \rightarrow [0, 1]$ maps the robot state x (e.g., from localization) to infer the progress variable s (i.e., the robot’s progress along the global plan so far), such that the closest state in the plan to a robot state x is $\Pi(\rho(x))$. For simplicity of notation, we represent the projected state at any time as $x_\Pi = \Pi(\rho(x))$. We also omit the explicit time-dependence of variables $x(t)$, $u(t)$, and $y(t)$, denoting them simply as x , u , and y . The objective of our controller u is thus to minimize the total navigation time T while following the plan precisely, as represented by the joint cost function

$$J = T + \gamma \int_0^T \|x(t) - x_\Pi(t)\|^2 dt. \quad (2)$$

Here, γ is a hyperparameter that trades total navigation time for execution accuracy.

We formulate the solution to this optimal control problem as a receding-horizon controller u^* over a unitless progress horizon (along the global plan) Δ and corresponding time-horizon Δt

such that the control input drives the robot state from x to the receding horizon plan state $\Pi(\rho(x) + \Delta)$ over time-period Δt :

$$u^* = \arg \min_u \left(\Delta t + \gamma \left\| \Delta x_{\Pi} - \int_0^{\Delta t} f(x, u, w) dt \right\|^2 \right),$$

$$\Delta x_{\Pi} = \Pi(\rho(x) + \Delta) - x, \quad (3)$$

where Δx_{Π} is the state change between the receding horizon plan and current vehicle state.¹ The optimal control u^* can be solved using the receding horizon inverse kinodynamic model f^{-1} as

$$u^* = f^{-1}(\Delta x, x, w), \quad (4)$$

that takes as input the desired relative state change Δx , the current robot state x , and world state w . Unfortunately, it is hard to express f^{-1} accurately via analytical models, and even if it could be expressed accurately, computing u^* is error-prone since the world state w is not directly observable.

B. Learning Inverse Kinodynamics

In this work, in order to enable fast and accurate navigation under the influence of different terrain interactions from the world state w , we adopt a data-driven approach to capture the effect of w . Specifically, we introduce the function

$$f_{\theta}^+(\Delta x, x, y) \approx f^{-1}(\Delta x, x, w), \quad (5)$$

parameterized by θ , as an approximation for the original receding horizon inverse kinodynamic function (superscript $+$ denotes pseudo inverse).

The key insight in this approximation is that the impact of w on u^* becomes predictable given observations y related to high-speed terrain-aware navigation—in our case we use onboard inertial sensing to capture speed-dependent terrain interaction. To learn f_{θ}^+ , a training dataset \mathcal{T} with N samples

$$\mathcal{T} = \{ \{ \Delta x^i, x^i, u^i, y^i \}_{i=1}^N \}$$

is desired, using the optimal but unknown receding horizon inverse kinodynamic function $u^i = f^{-1}(\Delta x^i, x^i, w^i)$ and observation function $y^i = g(x^i, w^i)$ from Eqns. 1 and 4. Unfortunately, we neither know f^{-1} , nor do we know the world states w^i . However, we do have access to f as a black-box function via real-world execution: we can simply pick arbitrary sample controls u^i at corresponding starting states x^i , and observe the resulting state change Δx^i after the chosen receding horizon Δt , including the impact of the unknown w^i :

$$\Delta x^i = \int_0^{\Delta t} f(x^i, u^i, w^i) dt.$$

Thus, the original chosen control u^i is the control² for the resulting state change Δx^i from the original state x^i , for the corresponding but unknown (and hence unrecorded) world state w^i . Along with the corresponding sensor observation y^i , we

¹In general, the minus signs in Eqn. 3 is the generalized difference operator \ominus over Special Euclidean Group $SE(n)$ and the corresponding Lie Algebra.

²Here, we assume the optimal control that follows the global path (second term in Eqn. 2) also implicitly minimizes navigation time (first term).

generate each sample i for dataset \mathcal{T} . To ensure that the learned parameters θ of f_{θ}^+ approximate f^{-1} accurately at states that the robot will encounter during execution, \mathcal{T} must include representative samples for x , u , and y . With the collected dataset \mathcal{T} , we formulate deriving $f_{\theta}^+(\cdot, \cdot, \cdot)$ as a learning problem by minimizing a supervised loss:

$$\theta^* = \arg \min_{\theta} \sum_{(\Delta x^i, x^i, y^i) \in \mathcal{T}} \|f^{-1}(\cdot, \cdot, \cdot) - f_{\theta}^+(\Delta x^i, x^i, y^i)\|_H$$

$$= \arg \min_{\theta} \sum_{(u^i, \Delta x^i, x^i, y^i) \in \mathcal{T}} \|u^i - f_{\theta}^+(\Delta x^i, x^i, y^i)\|_H, \quad (6)$$

where $\|v\|_H = v^T H v$ is the norm induced by positive definite matrix H , used to weigh the learning loss between the different dimensions of the control input u^i . We represent $f_{\theta}^+(\cdot, \cdot, \cdot)$ as a neural network and can therefore use gradient decent to find an approximately optimal θ^* . In this work, we collect raw 6-DoF readings from an onboard IMU, and construct y by feeding inertial data through an autoencoder to encode relevant terrain-vehicle interaction at different driving speeds. More details are provided in Section IV.

C. Online Execution

The learned inverse kinodynamic model $f_{\theta^*}^+(\cdot, \cdot, \cdot)$ provides a means to approximately account for w using the onboard observation y for fast, terrain-aware, and precise navigation. At each time step t during online execution, we compute the desired change of state $\Delta x_{\Pi} = \Pi(\rho(x) + \Delta) - x$ with x from localization, the projection operator $\rho(\cdot)$, projection horizon Δ , and global plan $\Pi(\cdot)$. Along with onboard observation y and current vehicle state x , we use the learned inverse kinodynamic model $f_{\theta^*}^+$ to produce system control input:

$$u(t) = f_{\theta^*}^+(\Delta x_{\Pi}, x, y), \quad (7)$$

and repeat this process for every time step.

IV. EXPERIMENTS

In this section, we present experimental results using a learned inverse kinodynamic model $f_{\theta^*}^+(\Delta x, x, y)$, which considers unobservable world state w by taking y as input, and can precisely track different global plans by different Π .

We denote the baseline forward and inverse kinodynamic functions, which do not consider world state w , as $f_B(x, u)$ and $f_B^+(\Delta x, x)$, respectively. As an ablation study to test the effectiveness of capturing the world state w with y , we also learn an ablated inverse kinodynamic model $f_{A\phi^*}^+(\Delta x, x)$, which is parameterized by ϕ^* and does not take observation y as input to represent w . Note that the baseline represents classical model-based local planners [2], while the ablation is equivalent to existing learning-based local planners, e.g., Imitation Learning [34]. We show that the learned $f_{A\phi^*}^+$ can outperform the baseline f_B^+ . Adding the learned observation y as another input, $f_{\theta^*}^+(\Delta x, x, y)$ can outperform both $f_{A\phi^*}^+(\Delta x, x)$ and $f_B^+(\Delta x, x)$.

The learned inverse kinodynamic model for online execution is also agnostic to different online plans and unseen terrain. We show that the model learned through a global planner Π_1 (a




Fig. 2. (a): A UT Automata robot, scale 1/10th autonomous vehicle platform used in the experiments. (b): The sampling-based baseline approximates Δx_{Π} by rolling out the optimal u_B among 100 samples with f_B .

randomly exploring policy driven by a human operator) can also generalize well to a complicated global planner Π_2 on an outdoor race track, and a simple global planner Π_3 on an indoor track. We also test the generalizability with respect to unseen world states through the encoded y and $f_{\theta^*}^+(\Delta x, x, y)$ on unseen terrain.

A. Implementation

1) *Robot Platform*: Our learned inverse kinodynamic model is implemented on a UT Automata robot, a scale 1/10th autonomous vehicle platform (Fig. 2 a). The Ackermann-steering four-wheel drive vehicle is equipped with a 2D Hokuyo UST-10LX LiDAR for localization, a Vectornav VN-100 IMU for inertial sensing (200 Hz), a Flipsky VESC 4.12 50 A motor controller, and a Traxxas Titan 550 Motor. Although the platform has individual suspensions, the relatively short-travel suspensions are not specifically designed for off-road navigation. We specifically pick this platform because its small size and weight and the lack of designed off-road capability can maximize the difference in accuracy between navigating with and without the learned terrain-aware inverse kinodynamic model. The robot has a NVIDIA Jetson onboard, but only the CPU is used during deployment.

2) *Environment*: The environment comprises of cement, grass, and mud; and some patches are covered by different artifacts, such as leaves, stalks, and/or twigs, with different densities (Fig. 1). For a small vehicle like the UT Automata robot, these artifacts can cause significantly different world states (Fig. 1 red insert). Note that the terrain also changes due to environmental factors such as sunlight, wind, and moisture, and is also affected by the robot's wheel and chassis. To minimize the effect of artifacts being pushed off the course during extended experiments, we frequently shuffle and redistribute the artifacts. We do not specify discrete terrain classes and treat the terrain characteristics in a continuous manner.

3) *Model Implementation*: During autonomous navigation, the robot uses Episodic Non-Markov Localization (ENML) [35] with a pre-built map of the environment to derive vehicle state x . A global planner Π includes a pre-generated global path for the robot to follow and uses line-of-sight control (similar to [36], [37]) to generate desired receding horizon plan state $\Pi(\rho(x) + \Delta)$ on the global path 1 m away from the robot. In a model predictive control manner, the robot uses the baseline




Fig. 3. Neural network architecture. input: blue IMU encoder and orange desired state change (desired velocity and curvature in practice); Output: purple learned function approximator $f_{\theta^*}^+$ as the inverse kinodynamic model.

forward kinodynamic function f_B and samples candidate velocity and curvature control inputs $u \in U$ evenly distributed within a physically-feasible window to jointly find the desired state change Δx_{Π} and control input u_B (shown in Fig. 2 b). More specifically, to compute control input, the baseline inverse kinodynamic model f_B^+ produces the curvature input, which results in the desired Δx_{Π} and drives the robot as close to $\Pi(\rho(x) + \Delta)$ as possible:

$$u_B = f_B^+(\cdot, \cdot) = \arg \min_u \|\Pi(\rho(x) + \Delta) - \int_{t=0}^{\Delta t} f_B(x, u) dt\|, \quad (8)$$

for the second term in Eqn. 2, and selects the fastest possible velocity for the first term T , considering the robot's acceleration limit and a safety distance to decelerate in case of obstacles.

For our learned ablated and final model, $f_{A\phi^*}^+$ and $f_{\theta^*}^+$, we utilize the Δx_{Π} from the baseline kinodynamic model (corresponds to u_B), but instead of using the baseline's control input, we query our learned models to produce $u = f_{A\phi^*}^+(\Delta x_{\Pi}, x)$ or $u = f_{\theta^*}^+(\Delta x_{\Pi}, x, y)$. In practice, we use the baseline control input $u_B = \{v, c\}$ (linear velocity and steering curvature) to represent the desired state change rate Δx_{Π} .

4) *Data Collection*: To collect training data, the robot is teleoperated with a joystick in an open environment with linear velocity $v \in [0, 3]$ m/s and steering curvature $c \in [-1.35, 1.35]$ m⁻¹ for 30 minutes (24 418 data points). The teleoperator randomly varies both linear velocity and steering curvature (Π_1). In our specific implementation, the robot reasons in the robot frame and therefore the state x^i in the training trajectory $\mathcal{T} = \{\langle \Delta x^i, x^i, u^i, y^i \rangle_{i=1}^N\}$ becomes the origin in the robot frame. The ground truth Δx^i is represented as real $\{v_r^i, c_r^i\}$, where v_r^i is from vehicle odometry and $c_r^i = \omega_r^i / v_r^i$ (ω_r^i is the sensed angular velocity around the vertical z axis from the IMU). Currently, we take v_r^i from wheel odometry only, which can be further improved by adding visual, point cloud, and/or inertial information in future work. The commanded control input $u^i = \{v_c^i, c_c^i\}$ is recorded from joystick input. For y , we collect the 6-DoF raw IMU signal, including 3-DoF accelerometer and 3-DoF gyroscope, as a sliding history window.

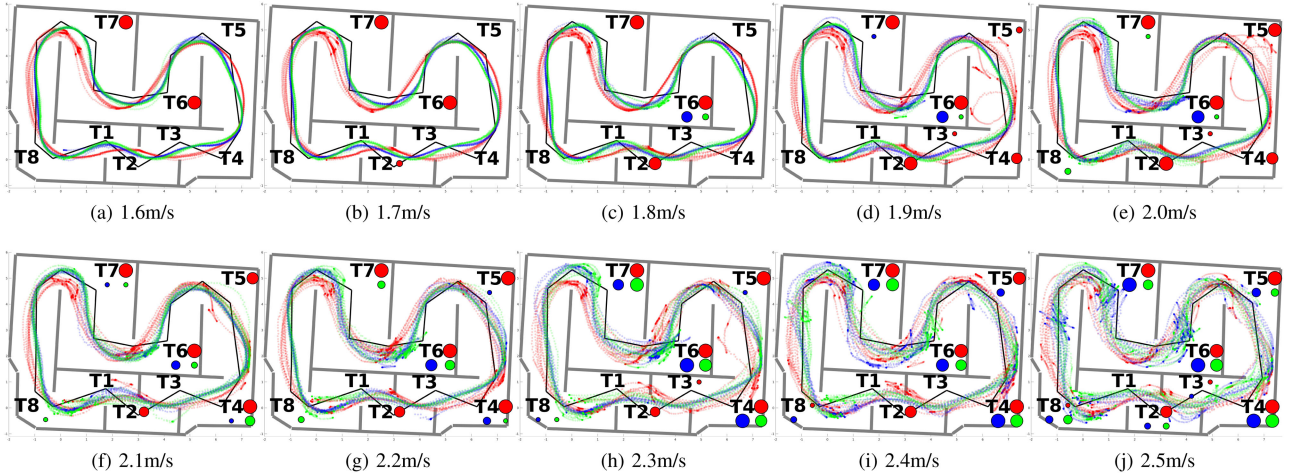


Fig. 4. Results of outdoor experiments on seen terrain: Localized robot positions of all 300 laps are plotted around the pre-defined global path (black line segments) on the map (grey lines). The size of the circles at each turn denotes the number of failures at that turn. Red: baseline $f_B^+(\Delta x, x)$, blue: ablated model $f_{A\phi^+}^+(\Delta x, x)$, green: learned model $f_{\theta^+}^+(\Delta x, x, y)$.

5) *Network Architecture*: As a function approximator for $f_{\theta^+}^+$, we use a two-layer neural network with 32 neurons each layer (shown in purple in Fig. 3) and we empirically show that such a small network is very efficient to train and suffices to improve off-road navigation. The neural network takes realistic/desired $\{v_r^i, c_r^i\}$ (as a proxy for Δx^i , Fig. 3, orange) and observation y as input, and outputs to-be-commanded control input $u^i = \{v_c^i, c_c^i\}$. For observation y , we concatenate the last 100 IMU readings (0.5 s) into a 600-dimensional vector, and feed it into two 256-neuron layers as an autoencoder (Fig. 3, blue). The final embedding for y is a two dimensional vector, then concatenated with $\{v_r^i, c_r^i\}$, and finally trained in an end-to-end fashion. The entire network architecture is shown in Fig. 3. For the ablated model $f_{A\phi^+}^+$, the two-dimensional y embedding is removed (only the orange and purple components remain). Training both models takes less than five minutes on a NVIDIA GeForce GTX 1650 laptop GPU. During runtime, the trained model is used onboard the robot's Jetson CPU with `libtorch`.

B. Navigation on Seen Terrain

We first test the inverse kinodynamic model's performance on the same terrain where the training data is collected, but with a different global planner Π_2 . After collecting the training data with Π_1 , an outdoor race track is constructed using plastic panels and wooden posts (Fig. 1). Starting from the origin (robot location in Fig. 1), eight turns are created (T1–T8, Π_2). While Turn 1, 2, and 3 are relatively gentle left-, right-, and left-hand turns, Turn 4 and Turn 8 are roughly 90° left-hand turns. Turn 5, 6, and 7 are sharp 180° left-, right-, and left-hand turns.

Ten different target speeds (the maximum speed the robot targets at reaching while maintaining safety tolerance to decelerate and avoid potential collisions) are tested, ranging from 1.6 m/s to 2.5 m/s with 0.1 m/s intervals. For the three models, the baseline $f_B^+(\Delta x, x)$, the ablated $f_{A\phi^+}^+(\Delta x, x)$, and learned $f_{\theta^+}^+(\Delta x, x, y)$, we repeat ten trials/laps each for statistical significance. A total 300 laps are executed. The localized robot position from ENML




Fig. 5. Failure rate per target speed.

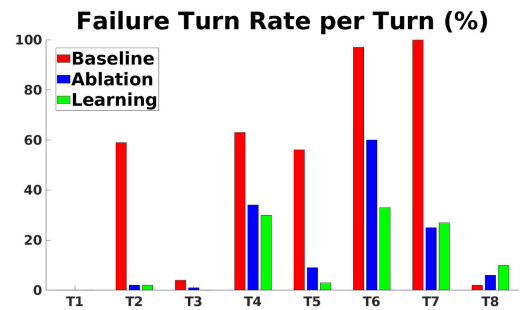


Fig. 6. Failure rate per turn.

TABLE I
OVERALL SUCCESS RATES OF ALL SPEEDS AND TURNS

	Baseline	Ablation	Learning
Seen	52.4%	82.9%	86.9%
Unseen	49.5%	–	87.0%

are shown in Fig. 4 in the subplots corresponding to the target speeds.

At lower target speeds, the green trajectories by the learned model $f_{\theta^+}^+(\Delta x, x, y)$ are much closer to the pre-defined global path, compared to the baseline $f_B^+(\Delta x, x)$, because the latter

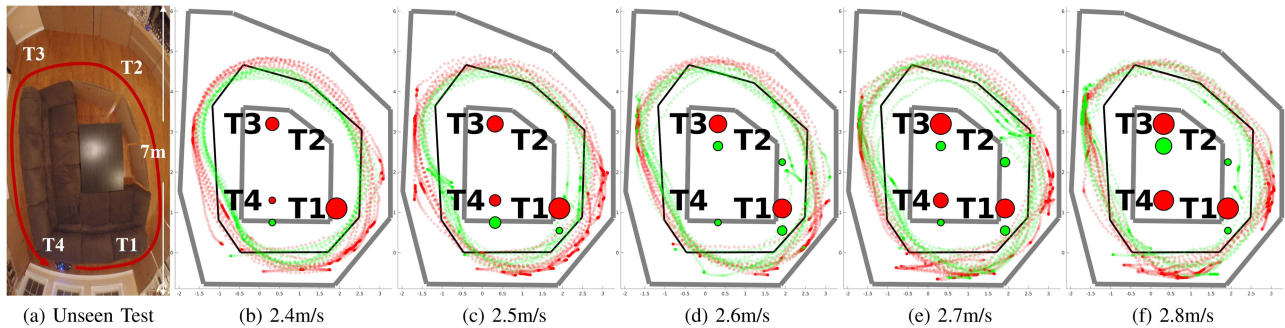


Fig. 7. Experiment results in unseen environment.

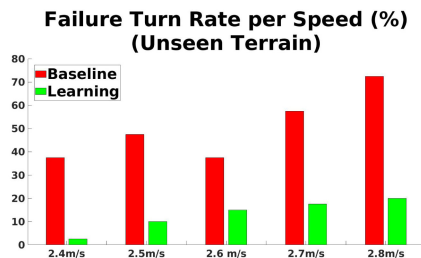


Fig. 8. Failure rate per target speed on unseen terrain.




Fig. 9. Failure rate per turn on unseen terrain.

model fails to consider the world state caused by the unstructured terrain. With increasing speed, the robot trajectory becomes more scattered around the global path due to increased stochasticity from vehicle-terrain interaction. But overall speaking, the green trajectories are always closer to the global path than other alternatives. The blue trajectory is generated by the ablated model $f_{A\phi^+}^+(\Delta x, x)$. Like $f_{\theta^+}^+$, it learns from actual terrain interactions, but it does not consider the current observation y . So $f_{A\phi^+}^+$ is roughly an averaged model over the continuous spectrum of terrain. Therefore, $f_{A\phi^+}^+$ outperforms the baseline f_B^+ , but underperforms $f_{\theta^+}^+$ because it fails to consider the current world state.

At each turn, the size of the red, blue, and green circles represents the number of failed turns (collision or getting stuck) in the ten attempted turns. Turn 6 and 7 cause a lot of trouble for the baseline even at lower speeds. With increasing speed, more turns cause failure for other models as well, but in general, the baseline fails more frequently at most turns than the ablated and learned models. Fig. 5 and Fig. 6 show the percentage of failed turns per target speed and per turn, respectively. In

Fig. 5, failure rate increases with faster speed, while within each speed, the learned model achieves the lowest failure rate, while the baseline fails most frequently. In Fig. 6, Turn 6 is the most difficult for all three alternatives, and at most turns, the learned model outperforms the ablation and the baseline. Since Turn 8 is immediately after the terrain change from grass to cement, the robot sometimes oversteers (to compensate for slip on grass) and does not react quickly enough to understeer (for higher friction on cement), causing it to get stuck in a few laps with the learned model. This problem can be addressed by adding forward looking camera to predict future wheel-terrain interaction in future work. The overall success rates of the three models for all turns are shown in the first row in Table I.

C. Navigation on Unseen Terrain

To test that the learned model generalizes to different global planners Π and also to unseen terrain, we further conduct an indoor experiment with a different track and global path (Π_3) on an unseen wooden floor (Fig. 7 a). Note that $f_{\theta^+}^+$ has only seen training data from random exploration on the outdoor unstructured terrain,³ we increase the navigation target speed to 2.4 m/s - 2.8 m/s, also with 0.1 m/s intervals. The baseline and the learned model are applied with these five different target speeds, ten repetitions each. Fig. 7 shows the results from the 100 laps on the unseen terrain with a different global path. Similar to the results on seen terrain, the learned model produces more concentrated and also closer robot trajectories to the global path to be tracked. As shown in Fig. 8 and 9, the learned model also outperforms the baseline in terms of failure rate at all target speeds and in most turns (except Turn 2). The overall success rates of the baseline and learned model are shown in the second row in Table I.

V. CONCLUSION

In this letter, we present a data-driven approach to learn an inverse kinodynamic model for accurate high-speed navigation on unstructured terrain. To capture the elusive and stochastic

³We speculate that the generalization would not be as good were the model trained indoors (on easy terrain) but applied outdoors.

world state caused by vehicle-terrain interaction at different high speeds, we use an inertia-based observation embedding as an input to the learned inverse kinodynamic function. This approach is tested on a physical robot on seen and unseen terrain with different global plans at different high speeds. The experimental results show that the learned model can significantly outperform an ideal baseline model without consideration of world state. Our ablation study also shows our observation embedding is useful to enable fast and accurate off-road navigation on unstructured terrain. For future work, better ground truth linear velocity estimation needs to be investigated: in addition to wheel odometry alone, other sources of perception, e.g., vision, point cloud, and/or inertia, can be leveraged. Better linear velocity estimation can account for significant wheel slippage on more challenging terrain, e.g., on ice, and enable even faster navigation. Adding vision-based observation also has the potential to enable the robot to prepare for future interactions, e.g., to reduce the failures at Turn 8. Another interesting direction to investigate in the future is generalization from easier to harder environments.

REFERENCES

- [1] S. Quinlan and O. Khatib, "Elastic bands: Connecting path planning and control," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1993, pp. 802–807.
- [2] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Automat. Mag.*, vol. 4, no. 1, pp. 23–33, Mar. 1997.
- [3] G. Seetharaman, A. Lakhota, and E. P. Blasch, "Unmanned vehicles come of age: The DARPA grand challenge," *Computer*, vol. 39, no. 12, pp. 26–29, 2006.
- [4] L. D. Jackel, E. Krotkov, M. Perschbacher, J. Pippine, and C. Sullivan, "The DARPA LAGR program: Goals, challenges, methodology, and phase i results," *J. Field Robot.*, vol. 23, no. 11-12, pp. 945–973, 2006.
- [5] C. Bai, J. Guo, and H. Zheng, "Three-dimensional vibration-based terrain classification for mobile robots," *IEEE Access*, vol. 7, pp. 63 485–63 492, 2019.
- [6] W. Shi, Z. Li, W. Lv, Y. Wu, J. Chang, and X. Li, "Laplacian support vector machine for vibration-based robotic terrain classification," *Electron.*, vol. 9, no. 3, 2020, Art. no. 513.
- [7] M. Mei, J. Chang, Y. Li, Z. Li, X. Li, and W. Lv, "Comparative study of different methods in vibration-based terrain classification for wheeled robots with shock absorbers," *Sensors*, vol. 19, no. 5, 2019, Art. no. 1137.
- [8] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, "Real-time semantic mapping for autonomous off-road navigation," *Field Service Robot.*. Springer, 2018, pp. 335–350.
- [9] P. Wolf, A. Vierling, T. Ropertz, and K. Berns, "Advanced Scene Aware Navigation for the Heavy Duty Off-Road Vehicle Unimog," in *Proc. IOP Conf. Series: Materials Sci. and Eng.*, IOP Publishing, 2020, vol. 997, no. 1, Art. no. 012093.
- [10] S. Sharma, J. E. Ball, B. Tang, D. W. Carruth, M. Doude, and M. A. Islam, "Semantic segmentation with transfer learning for off-road autonomous driving," *Sensors*, vol. 19, no. 11, 2019, Art. no. 2577.
- [11] S. Rabiee and J. Biswas, "A friction-based kinematic model for skid-steer wheeled mobile robots," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 8563–8569.
- [12] Y. Tian and N. Sarkar, "Control of a mobile robot subject to wheel slip," *J. Intell. Robot. Syst.*, vol. 74, no. 3, pp. 915–929, 2014.
- [13] F. Rogers-Marcovitz, M. George, N. Seegmiller, and A. Kelly, "Aiding off-road inertial navigation with high performance models of wheel slip," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 215–222.
- [14] M. Seyr and S. Jakubek, "Proprioceptive navigation, slip estimation and slip control for autonomous wheeled mobile robots," in *Proc. IEEE Conf. Robot. Automat. Mechatronics.*, 2006, pp. 1–6.
- [15] U. Muller, J. Ben, E. Cosatto, B. Flepp, and Y. L. Cun, "Off-road obstacle avoidance through end-to-end learning," *Adv. Neural Inf. Process. Syst.*. Citeseer, 2006, pp. 739–746.
- [16] A. Giusti *et al.* "A machine learning approach to visual perception of forest trails for mobile robots," *IEEE Robot. Automat. Lett.*, vol. 1, no. 2, pp. 661–667, Jul. 2016.
- [17] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion control for mobile robot navigation using machine learning: A survey," 2020, *arXiv:2011.13112*.
- [18] Z. Wang *et al.*, "Fromagile Ground to Aerial Navigation: Learning From Learned Hallucination," Tech. Rep., The University of Texas at Austin, 2021.
- [19] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021.
- [20] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1503–1510, Apr. 2021.
- [21] B. Liu, X. Xiao, and P. Stone, "A lifelong learning approach to mobile robot navigation," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1090–1096, Apr. 2021.
- [22] H.-T. L. Chiang, J. Hsu, M. Fiser, L. Tapia, and A. Faust, "RL-RRT: Kinodynamic motion planning via learning reachability estimators from RL policies," *IEEE Robot. Automat. Lett.*, vol. 4, no. 4, pp. 4298–4305, Oct. 2019.
- [23] G. P. Kontoudis and K. G. Vamvoudakis, "Kinodynamic motion planning with continuous-time q-learning: An online, model-free, and safe navigation framework," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 12, pp. 3803–3817, Dec. 2019.
- [24] M. Wigness, J. G. Rogers, and L. E. Navarro-Serment, "Robot navigation from human demonstration: Learning control behaviors," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 1150–1157.
- [25] C. Richter and N. Roy, "Safe Visual Navigation Via Deep Learning and Novelty Detection," *Robot.: Sci. Syst. XIII* Jul. 2017.
- [26] X. Xiao *et al.*, "APPL: Adaptive planner parameter learning," *arXiv:2105.07620*, 2021.
- [27] Z. Xu *et al.*, "APPLR: Adaptive planner parameter learning from reinforcement," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021.
- [28] Z. Wang, X. Xiao, G. Warnell, and P. Stone, "APPL: Adaptive Planner Parameter Learning From Evaluative Feedback," Tech. Rep., The University of Texas at Austin, 2021.
- [29] Z. Wang, X. Xiao, B. Liu, G. Warnell, and P. Stone, "APPLI: Adaptive planner parameter learning from interventions," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021.
- [30] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "APPLD: Adaptive planner parameter learning from demonstration," *IEEE Robot. Automat. Lett.*, vol. 5, no. 3, pp. 4541–4547, Jul. 2020.
- [31] Y. Pan *et al.*, "Imitation learning for agile autonomous driving," *Int. J. Robot. Res.*, vol. 39, no. 2-3, pp. 286–302, 2020.
- [32] S. Siva, M. Wigness, J. Rogers, and H. Zhang, "Robot adaptation to unstructured terrains by joint representation and apprenticeship learning," in *Robot.: Sci. Syst.*, 2019.
- [33] A. Brunnbauer *et al.*, "Model-based versus model-free deep reinforcement learning for autonomous racing cars," *arXiv:2103.04909*, 2021.
- [34] 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 *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 1527–1533.
- [35] J. Biswas and M. M. Veloso, "Episodic non-markov localization," *Robot. Auton. Syst.*, vol. 87, pp. 162–176, 2017.
- [36] X. Xiao, J. Dufek, T. Woodbury, and R. Murphy, "UAV assisted USV visual navigation for marine mass casualty incident response," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 6105–6110.
- [37] X. Xiao *et al.*, "Locomotive reduction for snake robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 3735–3740.