

# Learning Terrain-Aware Kinodynamic Model for Autonomous Off-Road Rally Driving With Model Predictive Path Integral Control

Hojin Lee, Taekyung Kim, Jungwi Mun, and Wonsuk Lee

**Abstract**—High-speed autonomous driving in off-road environments has immense potential for various applications, but it also presents challenges due to the complexity of vehicle-terrain interactions. In such environments, it is crucial for the vehicle to predict its motion and adjust its controls proactively in response to environmental changes, such as variations in terrain elevation. To this end, we propose a method for learning terrain-aware kinodynamic model which is conditioned on both proprioceptive and exteroceptive information. The proposed model generates reliable predictions of 6-degree-of-freedom motion and can even estimate contact interactions without requiring ground truth force data during training. This enables the design of a safe and robust model predictive controller through appropriate cost function design which penalizes sampled trajectories with unstable motion, unsafe interactions, and high levels of uncertainty derived from the model. We demonstrate the effectiveness of our approach through experiments on a simulated off-road track, showing that our proposed model-controller pair outperforms the baseline and ensures robust high-speed driving performance without control failure.<sup>1</sup>

**Index Terms**—Model learning for control, autonomous vehicle navigation, field robots.

## I. INTRODUCTION

**A**UTONOMOUS navigation in unstructured environments poses a significant challenge due to the complex and unpredictable nature of such surroundings. In various applications, including surveillance and rescue missions, high-speed driving capability is crucial for successful operation, particularly in off-road environments where unpaved roads and other challenging terrains may be encountered. Therefore, the control algorithm must be designed to be aware of environmental changes and to be capable of adapting promptly to them while ensuring safety.

Although extensive research efforts have addressed high-speed autonomous driving tasks, a major portion of these studies have focused on driving under flat terrain conditions [1]. Unfortunately, these approaches are not well-suited

Manuscript received: April 30, 2023; Revised: August 13, 2023; Accepted: September 11, 2023. This paper was recommended for publication by Editor J. Kober upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by the Agency for Defense Development Grant funded by the Korean Government in 2023. (Hojin Lee and Taekyung Kim are co-first authors.) (Corresponding author: Wonsuk Lee.)

Hojin Lee, Jungwi Mun, and Wonsuk Lee are with the AI Autonomy Technology Center, Agency for Defense Development, Daejeon, 34186, Republic of Korea {hojini1117, moonjw1109, wsblues82}@gmail.com

Taekyung Kim is with the Department of Robotics, University of Michigan, Ann Arbor, MI, 48109, USA taekyung@umich.edu

Digital Object Identifier (DOI): see top of this page.

<sup>1</sup>Project page: <https://sites.google.com/view/terrainawarekinodyn>

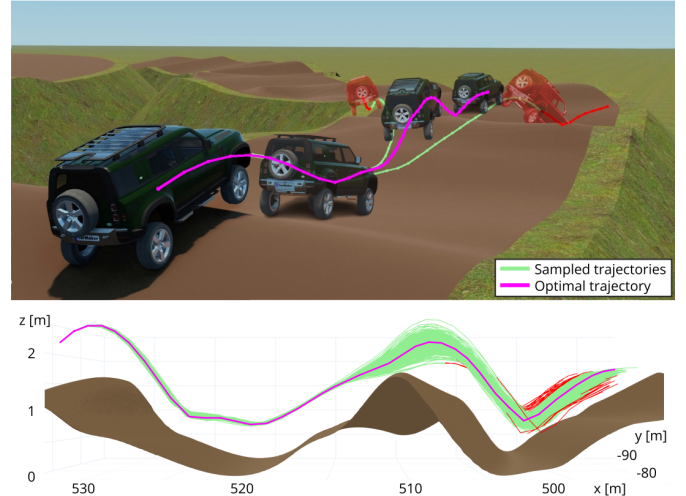


Fig. 1. Sampled trajectories and the optimal trajectory from the MPPI controller. Our terrain-aware kinodynamic model is capable of predicting the effect of terrain geometry on vehicle motion. The controller penalizes trajectories with unsafe vehicle motions and contact interactions, as well as high levels of uncertainty, allowing for the execution of aggressive maneuvers with enhanced safety. The penalized trajectories are depicted in red.

for off-road environments due to the challenges in accurately modeling the interactions between the vehicle and the terrain characterized by complex geometry. There have been attempts at analytical modeling of these interactions by using conventional methods such as with differential equations [2], [3]. However, these approaches rely heavily on assumptions that include constant contact between the wheels and the ground, single point contact, etc. These simplifying assumptions can be easily violated, leading to inaccurate predictions and degraded control performance, particularly in high-speed operations.

To effectively capture contact interactions in off-road environments, it is critical to model the vehicle dynamics by incorporating high-dimensional sensory data. Data-driven methods, such as Gaussian processes or online system identification, have emerged as tantalizing solutions for modeling nonlinear dynamics [4], [5]. However, most of these methods rely solely on proprioceptive information and are not scalable for handling vast amounts of exteroceptive data from various sensor modalities. In recent years, researchers have shifted their focus towards neural network-based approaches that can cope with large training datasets and guarantee constant-time inference [6]–[12]. By leveraging the advanced deep learning techniques, autonomous vehicles gain a greater understanding

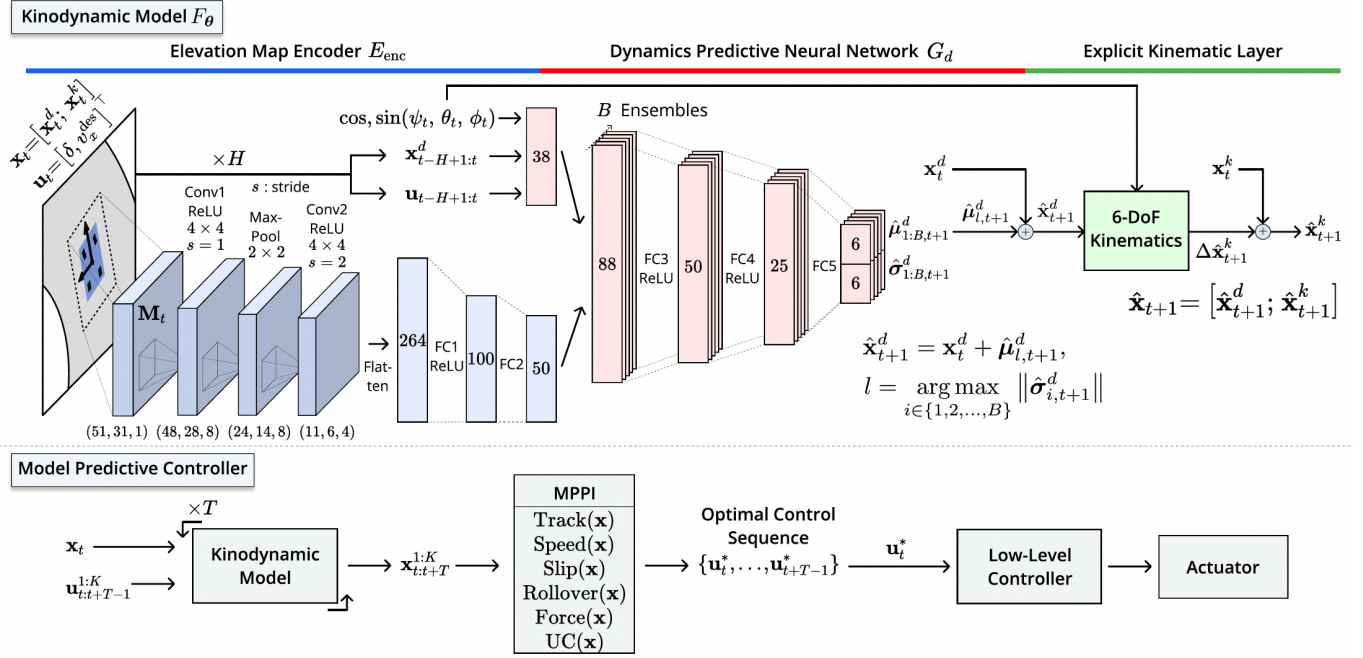


Fig. 2. An overview diagram of our method of learning terrain-aware kinodynamic model for autonomous high-speed off-road driving using MPPI controller.

of their surroundings and are able to proactively adjust their control strategy, paving the way for new navigational capabilities in challenging environments.

Despite the promising outcomes of using exteroceptive information to encode environmental context in neural network dynamics, there are still a few limitations that need to be resolved. Recent studies have demonstrated the effectiveness of utilizing learning-based models as dynamics of the systems combined with model-based controllers. Nonetheless, their applications are typically limited to path following tasks, where the path is computed offline without considering wheel-terrain interactions [6], [7], [11]. These fail to fully harness the model's capability to reliably predict interactive motions with the environments. Concretely, it is apparent that the model-based controller has the potential to generate kinodynamically optimal future trajectories on its own by utilizing neural network dynamics as a forward dynamics model [8], [9]. High-speed maneuvers in off-road environments including abrupt terrain profile changes may encompass highly discontinuous and impulsive contact forces to the systems. Pre-computed trajectories that ignore these effects can potentially degrade maneuverability, and in the worst case, lead to a loss of stability. Fig. 1 illustrates an example of such a scenario, where the vehicle exhibits significant roll and pitch motions due to the vertical forces exerted on the wheels. Therefore, the controller must be capable of optimizing future trajectories that are inherently feasible and robust to the environmental changes, respecting both the structural shape of the terrain and the dynamic characteristics of the controlled system.

In this paper, we propose an approach of learning terrain-aware kinodynamic model to achieve high-speed vehicle control in off-road environments while ensuring safety and

robustness. Our contributions are as follows:

- We develop a neural network model that can make reliable 6 Degrees-of-Freedom (DoF) motion predictions in terrains with rapid profile changes, by using a local elevation map as input and adopting a number of design considerations. In addition, our model can estimate the severity of ground impact by deriving contact forces and moments from its motion predictions, even without requiring ground truth force data during training.
- We describe a sampling-based model predictive control (MPC) scheme that employs our model as the forward dynamics to directly plan kinodynamically optimal trajectories on its own. This is achieved by a cost function design which penalizes trajectories that exhibit unstable motion, unsafe interactions, and high levels of uncertainty, all of which can be inferred from our model.
- We demonstrate the effectiveness of our idea by conducting experiments on a simulated off-road race track. Additionally, we validate that our method enhances vehicle maneuverability to a greater extent than the baseline, by presenting a scenario where our model encourages the MPC to make aggressive maneuvers rather than taking a conservative behavior.

## II. RELATED WORK

Early work on off-road autonomous driving mainly centered on analyzing the effect of vehicle-terrain interaction via traversability estimation [13]–[15]. They primarily used human supervision to assess the difficulty of the terrain to be traversed and to provide the vehicle with online maps of traversable regions. More recently, researchers have focused on self-supervised approaches that aim to learn traversability directly

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

from real driving experience [16], [17]. Unfortunately, all these methods may suffer from dynamics-related information loss when dealing with challenging terrain, because they tend to simplify traversability as a scalar value or a categorical label, and do not consider dynamic interactions which vary depending on the dynamic state of the vehicle.

On the other hand, a dynamics learning approach for off-road driving has been developed with the aim of capturing both the nonlinear vehicle dynamics and its interaction with the terrain, without relying on simplified traversability representations. Xiao et al. [18] developed an inverse kinodynamic model conditioned on onboard proprioceptive observations to sense the effect of vehicle-terrain interaction. Other previous studies demonstrated the benefits of leveraging exteroceptive information along with proprioceptive information to prepare the vehicle for upcoming interactions [6], [7], [11]. Building upon these strategies, Tremblay et al. [10] and Triest et al. [12] have further improved off-road dynamics prediction by expanding the model to accommodate multiple high-dimensional sensor modalities.

An off-road dynamics model can then be used to provide useful information to a model-based controller, which can facilitate the planning of optimal trajectories. Kahn et al. [8] demonstrated that the predictive model can be used to execute control inputs that prevent collisions and actively navigate the robot towards smoother terrain. Wang et al. [9] used constrained optimization to find trajectories with low divergence inferred from a probabilistic dynamics model. In this work, we extend this control strategy to high-speed off-road driving, by enabling the controller to perceive unstable motion and unsafe interactions, and by optimizing control inputs to prevent catastrophic failures when driving through challenging terrain.

### III. TERRAIN-AWARE KINODYNAMIC MODEL

In this section, we first define the state and action space to establish the kinodynamic model  $F_{\theta}$ . It predicts the change in the vehicle's state induced by control inputs and contact interactions during a discrete time period of  $\Delta t$ . Then, we describe the three main components of the model: 1) *Elevation Map Encoder*, 2) *Dynamics Predictive Neural Network*, and 3) *Explicit Kinematic Layer*.

First, we define the 6-DoF proprioceptive state  $\mathbf{x} = [\mathbf{x}^d; \mathbf{x}^k]$  at the center of mass (CoM) of the vehicle. The dynamic state  $\mathbf{x}^d$  includes linear velocity  $\mathbf{v}_b$  and angular velocity  $\boldsymbol{\omega}_b$ . The kinematic state  $\mathbf{x}^k$  comprises global position  $\mathbf{p}_s$  and orientation  $\mathbf{e}_s$ . The notations  $(\cdot)_s$  and  $(\cdot)_b$  denote vectors represented in a fixed reference frame and in the vehicle body frame attached to the CoM, respectively:

$$\begin{aligned} \mathbf{x}^d &= \begin{bmatrix} \mathbf{v}_b \\ \boldsymbol{\omega}_b \end{bmatrix} = [v_x, v_y, v_z, \omega_x, \omega_y, \omega_z]^\top, \\ \mathbf{x}^k &= \begin{bmatrix} \mathbf{p}_s \\ \mathbf{e}_s \end{bmatrix} = [x, y, z, \psi, \theta, \phi]^\top, \end{aligned} \quad (1)$$

where  $(\psi, \theta, \phi)$  are the yaw, pitch, and roll of the ZYX Euler angles, respectively.

Second, we augment the state space with exteroceptive information to incorporate the influence of terrain geometry.

In a recent study, exteroceptive information consisted of height samples obtained near each foot of the robot [19]. For bipedal or quadrupedal robots, contact points can be computed precisely using rigid body kinematics [20]. However, for full-size vehicles, deformable suspensions make the estimation of contact interaction unreliable using height samples around computed contact points.

Therefore, we use a vehicle-centric local elevation map cropped from the global map as exteroceptive information. We assume that the vehicle has full knowledge of the global map through existing techniques [21]. To ensure the model captures all possible contact interactions within  $\Delta t$ , we first specify the maximum operating speed of the vehicle, and set the size of the local map to be large enough to encompass the area that the vehicle can reach within  $\Delta t$  from its current position, as well as all four contact points. This strategy enables the neural network to implicitly predict contacts and interactions by considering the influence of terrain geometry around the vehicle as a whole, as opposed to using only the local geometry at each contact point. The local map is centered at the CoM and aligned with the yaw angle of the vehicle. We represent the elevation magnitudes of the local map, denoted as  $\mathbf{M}$ , as relative heights compared to the CoM. We set  $\Delta t$  to 0.1 s and the grid size of the local map to 10 cm.

In contrast to the electric vehicles used in recent learning-based autonomous driving studies [22], [23], our vehicle uses an automatic transmission for the gear-shifting mechanism, making throttle control impractical. Therefore, we select the desired speed as a control input instead of the throttle and use a Proportional-Integral (PI) low-level controller to track the desired speed. Hence, the control input  $\mathbf{u}$  consists of the steering angle  $\delta$  and the desired longitudinal speed  $v_x^{\text{des}}$ .

#### A. Elevation Map Encoder

The *Elevation Map Encoder* (EME)  $E_{\text{enc}}$  processes the local elevation map  $\mathbf{M}_t$  and outputs the latent terrain feature vector  $\mathbf{h}_t$ .  $\mathbf{M}_t$  is fed through a three-layer Convolutional Neural Network (CNN), whose output is flattened and then fed through two additional fully-connected layers:

$$\mathbf{h}_t = E_{\text{enc}}(\mathbf{M}_t). \quad (2)$$

#### B. Dynamics Predictive Neural Network

The *Dynamics Predictive Neural Network* (DPNN)  $G_d$  predicts the change in linear and angular velocities using proprioceptive information and a terrain feature vector  $\mathbf{h}_t$  as inputs. To enhance the model's prediction performance, we take a number of design considerations for the DPNN.

In real scenarios, exteroceptive sensors often provide noisy and imperfect data about the terrain. To account for this uncertainty, we build the DPNN as a probabilistic ensemble (PE) model [24], [25] capable of capturing both aleatoric and epistemic uncertainty, which can then be managed by the controller. Each ensemble member with the same network structure is randomly initialized and individually trained to capture epistemic uncertainty via disagreement in prediction.

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

Each member outputs the mean and standard deviation vector  $(\hat{\boldsymbol{\mu}}_{t+1}^d, \hat{\boldsymbol{\sigma}}_{t+1}^d)$  of the normal distribution to predict the change in velocities while incorporating aleatoric uncertainty.

In addition, rapid changes in the terrain profile cause highly time-varying and nonlinear vehicle motion that is difficult to model. Prior work suggested that adding history information to neural network models is effective for capturing time-varying behavior of the vehicle [26]. Therefore, we provide the DPNN with the history of velocities and control inputs  $(\mathbf{x}_{t-H+1:t}^d, \mathbf{u}_{t-H+1:t}) = (\mathbf{x}_{t-H+1}^d, \dots, \mathbf{x}_{t-1}^d, \mathbf{x}_t^d, \mathbf{u}_{t-H+1}, \dots, \mathbf{u}_{t-1}, \mathbf{u}_t)$ . In order to capture contact interactions, we also provide orientation angles  $(\psi_t, \theta_t, \phi_t)$  to characterize the spatial relationship between the vehicle and terrain.  $\cos$  and  $\sin$  of angular values are used to avoid the angle wrapping problem [6].

To sum up, the DPNN can be described as:

$$\begin{aligned} \hat{\boldsymbol{\mu}}_{i,t+1}^d, \hat{\boldsymbol{\sigma}}_{i,t+1}^d &= G_d(\mathbf{x}_{t-H+1:t}^d, \mathbf{u}_{t-H+1:t}, h_t, \\ &\quad c_\psi, s_\psi, c_\theta, s_\theta, c_\phi, s_\phi), \\ \hat{\mathbf{x}}_{i,t+1}^d &= \mathbf{x}_t^d + (\hat{\boldsymbol{\mu}}_{i,t+1}^d + \hat{\boldsymbol{\sigma}}_{i,t+1}^d \boldsymbol{\varepsilon}_i), \\ \boldsymbol{\varepsilon}_i &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \forall i \in \{1, 2, \dots, B\}. \end{aligned} \quad (3)$$

For notational simplicity, we abbreviate  $\cos(*)$  and  $\sin(*)$  as  $c_*$  and  $s_*$ , respectively.  $B$  is the number of models in the ensemble. In Section IV-C, we provide a slightly modified version of (3) for use in the controller.

### C. Explicit Kinematic Layer

The *Explicit Kinematic Layer* (EKL) analytically calculates the changes in global position and orientation. It transforms velocity outputs  $(\hat{\mathbf{v}}_b, \hat{\boldsymbol{\omega}}_b) \in \hat{\mathbf{x}}^d$  of the DPNN to vectors in a fixed reference frame. The change in position  $\Delta \hat{\mathbf{p}}_s$  can be obtained by sequentially applying the rotation matrices  $R$  corresponding to each Euler angle to the body frame displacement  $\hat{\mathbf{v}}_b \Delta t$ :

$$\begin{aligned} \Delta \hat{\mathbf{p}}_s &= R_z(\psi)R_y(\theta)R_x(\phi)\hat{\mathbf{v}}_b \Delta t \\ &= \begin{bmatrix} c_\psi c_\theta & c_\psi s_\theta s_\phi - s_\psi c_\phi & c_\psi s_\theta c_\phi + s_\psi s_\phi \\ s_\psi c_\theta & s_\psi s_\theta s_\phi + c_\psi c_\phi & s_\psi s_\theta c_\phi - c_\psi s_\phi \\ -s_\theta & c_\theta s_\phi & c_\theta c_\phi \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} \Delta t. \end{aligned} \quad (4)$$

The angular velocity vector  $\hat{\boldsymbol{\omega}}_b$  is the sum of Euler angle rate vectors transformed from their intermediate frames to the vehicle body frame [3]:

$$\begin{aligned} \hat{\boldsymbol{\omega}}_b &= \begin{bmatrix} \dot{\phi} \\ 0 \\ 0 \end{bmatrix} + R_x(-\phi) \begin{bmatrix} 0 \\ \dot{\theta} \\ 0 \end{bmatrix} + R_x(-\phi)R_y(-\theta) \begin{bmatrix} 0 \\ 0 \\ \dot{\psi} \end{bmatrix} \\ &= \begin{bmatrix} -s_\theta & 0 & 1 \\ c_\theta s_\phi & c_\phi & 0 \\ c_\theta c_\phi & -s_\phi & 0 \end{bmatrix} \begin{bmatrix} \dot{\psi} \\ \dot{\theta} \\ \dot{\phi} \end{bmatrix}. \end{aligned} \quad (5)$$

Then the change in orientation  $\Delta \hat{\mathbf{e}}_s$  can be derived via the inverse mapping of (5):

$$\Delta \hat{\mathbf{e}}_s = \begin{bmatrix} \dot{\psi} \\ \dot{\theta} \\ \dot{\phi} \end{bmatrix} \Delta t = \begin{bmatrix} 0 & s_\phi/c_\theta & c_\phi/c_\theta \\ 0 & c_\phi & -s_\phi \\ 1 & s_\theta s_\phi/c_\theta & s_\theta c_\phi/c_\theta \end{bmatrix} \hat{\boldsymbol{\omega}}_b \Delta t. \quad (6)$$

Finally, the 6-DoF kinematic state  $\hat{\mathbf{x}}_{t+1}^k$  is computed as:

$$\hat{\mathbf{x}}_{t+1}^k = \mathbf{x}_t^k + \begin{bmatrix} \Delta \hat{\mathbf{p}}_s \\ \Delta \hat{\mathbf{e}}_s \end{bmatrix}_{t+1}. \quad (7)$$

## IV. MODEL PREDICTIVE CONTROL USING THE TERRAIN-AWARE MODEL

### A. Model Predictive Path Integral Controller

Here, we provide a brief overview of our MPC method for off-road driving. Sampling-based MPC can handle both convex and non-convex optimization problems without requiring a gradient of the cost function. In recent years, the MPPI algorithm has been one of the most effective sampling-based approaches, aided by advancements in Graphics Processing Units (GPUs). The MPPI enables the design of arbitrary cost functions and the representation of the dynamics model by any nonlinear functions, such as neural networks [22].

In each timestep  $t$ ,  $K$  parallel samples of control sequences  $\mathbf{u}_{t:t+T-1}^{1:K} = (\mathbf{u}_{t:t+T-1}^1, \dots, \mathbf{u}_{t:t+T-1}^K)$  are drawn with a finite time horizon  $T$ . Each control sequence sample is transformed into the corresponding state trajectory by recursively applying  $F_\theta$ . Then the optimal sequence  $\mathbf{u}_{t:t+T-1}^*$  is obtained by taking the weighted average of samples, where the weight is determined by the evaluated cost of each trajectory using state-dependent cost function  $c(\mathbf{x})$ .

However, significant chattering of the optimal sequence may occur due to the stochastic process of this sampling-based algorithm. To resolve this problem, we utilize the Smooth MPPI (SMPPI) approach proposed in our prior work [27], which effectively attenuates chattering while adapting to a rapidly changing environment. Similar to prior work [22], we add the following three components to  $c(\mathbf{x})$ :

$$\begin{aligned} \text{Track}(\mathbf{x}) &= W_{\text{inf}} \mathbf{C}(x, y), \\ \text{Speed}(\mathbf{x}) &= (v_x - v_{\text{ref}})^2, \\ \text{Slip}(\mathbf{x}) &= \beta^2 + W_{\text{inf}} I(|\beta| > \beta_{th}). \end{aligned} \quad (8)$$

“ $th$ ” denotes an empirically defined threshold, and  $\beta = -\tan^{-1}\left(\frac{v_y}{|v_x|}\right)$  is the side-slip angle.  $\mathbf{C}(x, y)$  is the value in 2D cost-map at  $(x, y)$  which returns 1 if the vehicle is outside the track boundary and 0 otherwise.  $v_{\text{ref}}$  is the predefined reference speed, and  $I$  is an indicator function.  $W_{\text{inf}}$  is a sufficiently large constant to impose an impulse penalty.

### B. Off-Road Cost Function Design

We suggest a number of off-road cost function terms that can be incorporated into the state cost  $c(\mathbf{x})$  by employing predictions of the model. Instead of relying solely on (8), the inclusion of off-road costs allows the controller to make aggressive terrain-aware behavior while ensuring safety.

Firstly, we define the rollover cost to simply penalizing trajectories with large roll and pitch angles:

$$\begin{aligned} \text{Rollover}(\mathbf{x}) &= \theta^2 + \phi^2 + \\ &\quad W_{\text{inf}} (I(|\theta| > \theta_{th}) + I(|\phi| > \phi_{th})). \end{aligned} \quad (9)$$

Secondly, we penalize trajectories in which the vehicle is subjected to substantial interaction forces and moments

resulting from terrain geometry. When driving at high speeds on off-road terrain, unstable conditions of the vehicle not only include a high roll or pitch angle, but also a large vertical force  $F_z$  caused by shock impacts, a roll moment  $M_x$  and a pitch moment  $M_y$  resulting from uneven force distribution. This may severely damage the vehicle or compromise its stability. We can derive interaction forces  $\hat{\mathbf{f}}_b$  and moments  $\hat{\mathbf{m}}_b$  using velocity outputs from our model as:

$$\begin{aligned} \hat{\mathbf{f}}_b &= \begin{bmatrix} F_x \\ F_y \\ F_z \end{bmatrix}_b = m \left( \dot{\hat{\mathbf{v}}}_b + [\hat{\boldsymbol{\omega}}_b] \hat{\mathbf{v}}_b \right), \\ \hat{\mathbf{m}}_b &= \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix}_b = \mathcal{I}_b \dot{\hat{\boldsymbol{\omega}}}_b + [\hat{\boldsymbol{\omega}}_b] \mathcal{I}_b \hat{\boldsymbol{\omega}}_b. \end{aligned} \quad (10)$$

$[\hat{\boldsymbol{\omega}}_b]$  is a  $3 \times 3$  skew-symmetric matrix expression of  $\hat{\boldsymbol{\omega}}_b$  [28].  $\mathcal{I}_b$  is the rotational inertia matrix of the vehicle. In our model, we can approximate  $\dot{\hat{\mathbf{v}}}_b$  as  $\{\hat{\mathbf{v}}_b(t+1) - \hat{\mathbf{v}}_b(t)\} / \Delta t$ , and a similar approximation can be made for  $\dot{\hat{\boldsymbol{\omega}}}_b$ . Hence we can define the force cost as follows to explicitly penalize unsafe interactions without resorting to heuristics:

$$\begin{aligned} \text{Force}(\mathbf{x}) &= \sum_{j=1}^3 (\zeta_j^2 + W_{\text{inf}} I(|\zeta_j| > \zeta_{j,th})), \\ \zeta_1 &= \frac{F_z}{mg} - 1, \quad \zeta_2 = \frac{M_x}{mgd_x}, \quad \zeta_3 = \frac{M_y}{mgd_y}. \end{aligned} \quad (11)$$

$(d_x, d_y)$  denotes the distances from the CoM to the front-left wheel in the lateral and longitudinal directions, respectively.

Lastly, we incorporate the uncertainty cost into  $c(\mathbf{x})$ . This enables the controller to avoid state spaces with high uncertainty owing to noisy exteroceptive input or insufficient model learning.  $\text{UC}(\mathbf{x})$  is defined as the total amount of uncertainty inherent in the model prediction, which can be expressed as the sum of aleatoric uncertainty  $\text{AU}(\mathbf{x})$  and epistemic uncertainty  $\text{EU}(\mathbf{x})$  as follows [29]:

$$\text{UC}(\mathbf{x}) = \text{AU}(\mathbf{x}) + \text{EU}(\mathbf{x}) = \mathbb{E}_i \left[ \|\hat{\boldsymbol{\sigma}}_i^d\|^2 \right] + \|\text{Var}_i [\hat{\boldsymbol{\mu}}_i^d]\|. \quad (12)$$

To sum up, the final form of  $c(\mathbf{x})$  is:

$$\begin{aligned} c(\mathbf{x}) &= w_1 \text{Track}(\mathbf{x}) + w_2 \text{Speed}(\mathbf{x}) + w_3 \text{Slip}(\mathbf{x}) \\ &+ w_4 \text{Rollover}(\mathbf{x}) + w_5 \text{Force}(\mathbf{x}) + w_6 \text{UC}(\mathbf{x}). \end{aligned} \quad (13)$$

### C. Implementation Details

This subsection explains a few details for handling predictions of our model to reduce computation and enhance robustness of the MPPI. Since the PE model provides a distribution of trajectories instead of a single trajectory, multiple trajectories based on the Monte Carlo (MC) method are necessary to evaluate the cost of each control sequence. Such an approach requires an exponentially increasing number of state cost evaluations along sampling over each distribution, which is intractable to be implemented in real time.

To address this issue, we simply generate trajectories by propagating the mean prediction  $\hat{\boldsymbol{\mu}}^d$ . Prior work [24] indicated that the performance of a controller using a mean propagation is comparable to that of a controller utilizing an MC method.

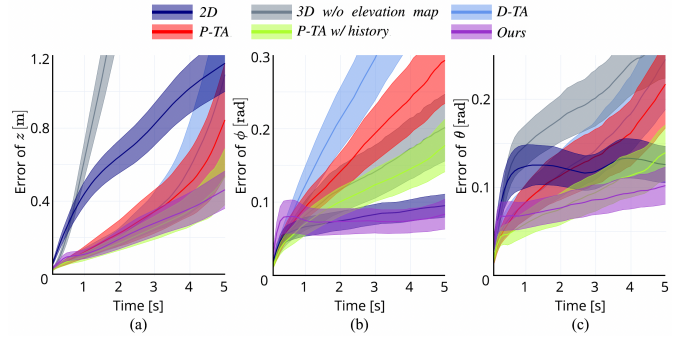


Fig. 3. Multi-step prediction errors for the values (a)  $z$ , (b)  $\phi$ , and (c)  $\theta$ . Each line represents the mean, and the shaded region represents  $1/5$  of the standard deviation for visual clarity.

Based on the above strategy, we select the ensemble model that has the greatest aleatoric uncertainty, and then use the mean prediction of the chosen model. This enhances the robustness of the MPPI since it considers the worst-case scenario among the ensembles.

In summary, (3) is modified for use in MPPI as:

$$\hat{\mathbf{x}}_{t+1}^d = \mathbf{x}_t^d + \hat{\boldsymbol{\mu}}_{l,t+1}^d, \quad l = \underset{i \in \{1,2,\dots,B\}}{\text{argmax}} \|\hat{\boldsymbol{\sigma}}_{i,t+1}^d\|. \quad (14)$$

## V. EXPERIMENTS

### A. Experimental Setup

We conducted experiments using the high-fidelity vehicle simulation software, the IPG CarMaker. We collected driving data using a variety of human-selected maneuvers [27]. By leveraging a simulator, we could even include risky maneuvers such as sudden acceleration or changes in steering commands, which broadens the data distribution. We added Gaussian noise to  $\delta$  and  $v_x^{\text{des}}$  in order to enhance the accuracy of prediction when highly noisy control sequences arising from the MPPI sampling procedure are given to the model.

Two completely different elevation profiles were generated on the same track; one for training data collection and the other for model validation. We collected 200 minutes of training data  $\mathcal{D}_{\text{train}}$  and 16 minutes of validation data  $\mathcal{D}_{\text{val}}$ .

### B. Training Procedure

The dataset was preprocessed using min-max normalization and standardization of proprioceptive and exteroceptive information, respectively. We applied Gaussian noise with zero means to the data to improve model robustness [30]. Then we trained each model in the ensemble by minimizing the Gaussian negative log likelihood (NLL) loss  $\mathcal{L}_d$ :

$$\mathcal{L}_d = \sum_{\mathbf{x}^d \in \mathcal{D}_{\text{train}}} \log \det \hat{\boldsymbol{\Sigma}} + \hat{\boldsymbol{\eta}}^\top \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\eta}}, \quad (15)$$

where  $\hat{\boldsymbol{\eta}} = \hat{\boldsymbol{\mu}}_{t+1}^d + \mathbf{x}_t^d - \mathbf{x}_{t+1}^d$ , and  $\hat{\boldsymbol{\Sigma}}$  is a diagonal matrix form of  $\hat{\boldsymbol{\sigma}}_{t+1}^d$ .

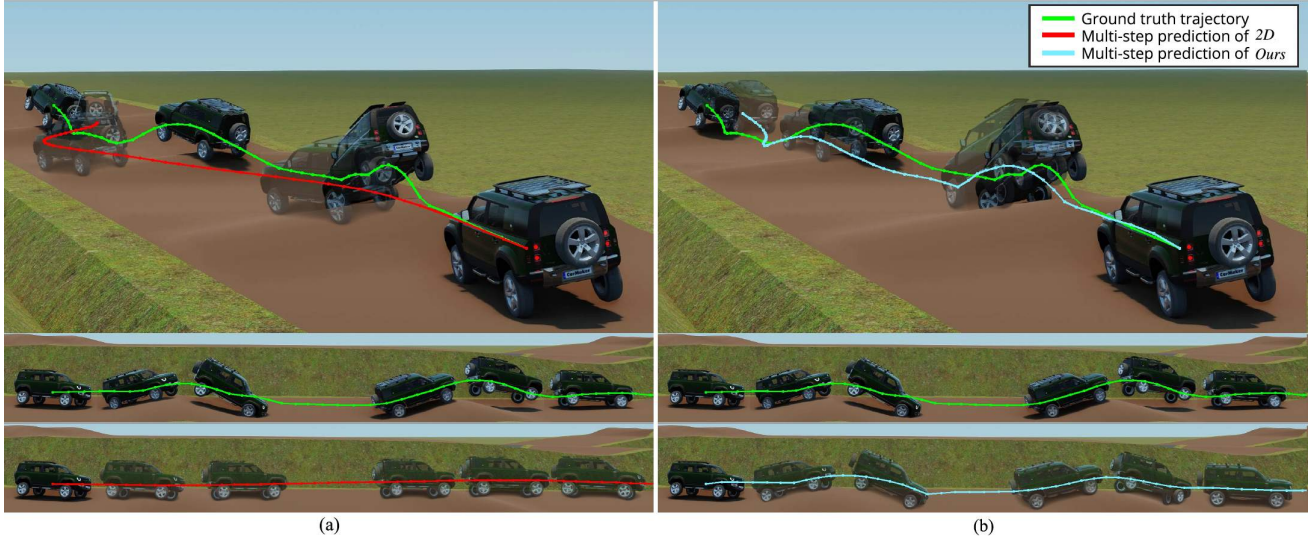


Fig. 4. Ground truth trajectory and predicted trajectories generated by multi-step prediction of (a) 2D and (b) *Ours* for one example trial.

### C. Analysis on The Kinodynamic Model

We compare the prediction performances of kinodynamic models based on a number of design considerations explained in Section III-B. Although we trained the model by minimizing the one-step loss, we chose  $T$ -step validation error as the evaluation metric since we would like to know how well the MPPI controller can generate a trajectory by recursively applying the learned model [30]. Given the initial state  $\mathbf{x}_t$  and the actual control sequence  $\mathbf{u}_{t:t+T-1}$  from  $\mathcal{D}_{\text{val}}$ , we compared the ground truth states  $\mathbf{x}_{t+1:t+T}$  and the multi-step predictions of the model  $\hat{\mathbf{x}}_{t+1:t+T}$ , which can be obtained by sequential one-step predictions:

$$\hat{\mathbf{x}}_{t+\tau} = \hat{\mathbf{x}}_{t+\tau-1} + F_{\theta}(\hat{\mathbf{x}}_{t+\tau-1}, \mathbf{u}_{t+\tau-1}). \quad (16)$$

To demonstrate the validity of our design considerations, we manually defined and trained the following ablation models using the same dataset: 1) “2D”: a 3-DoF baseline model; 2) “3D w/o elevation map”: a 6-DoF model without the knowledge of terrain; 3) “D-TA”: a Deterministic Terrain-Aware model; 4) “P-TA”: a Probabilistic Terrain-Aware model; 5) “P-TA w/ history”: a “P-TA” provided with history information; 6) “Ours”: a “P-TA w/ history” provided with orientation angles. In 1)-3), DPNN outputs the mean vector  $\hat{\boldsymbol{\mu}}^d$ , and the mean squared error (MSE) loss function is used for training. All models were ensemble models with a size of  $B = 5$ . We list the design specifications in Table I.

Fig. 3 shows the multi-step prediction errors in the values  $(z, \theta, \phi)$  over the prediction horizon  $T$ , which is highly dependent on terrain shape. 2D assumed constant values for these parameters, which resulted in discrepancies with ground truth values, and disregarded the influence of terrain geometry on the motion. 3D w/o elevation map exhibited more errors despite considering 6-DoF motion, due to the absence of terrain information, which led to divergence in a vertical direction. In contrast, D-TA took advantage of both 6-DoF motion and the terrain geometry, making it capable of fully estimating 3D vehicle-terrain interactions. However, the result

TABLE I  
THE DESIGN SPECIFICATIONS OF EACH MODEL. WE ABBREVIATED THE DETERMINISTIC/PROBABILISTIC MODEL AS D/P, RESPECTIVELY.

#	Model	Input of DPNN				
		DoF	$\mathbf{h}_t$	D / P	$H$	$(\psi, \theta, \phi)$
1	2D	3	✗	D	1	✗
2	3D w/o elevation map	6	✗	D	1	✗
3	D-TA	6	✓	D	1	✗
4	P-TA	6	✓	P	1	✗
5	P-TA w/ history	6	✓	P	3	✗
6	Ours	6	✓	P	3	✓

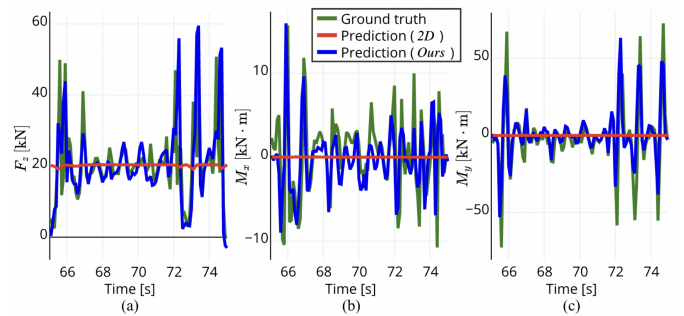


Fig. 5. Visualization of (a) vertical force  $F_z$ , (b) roll moment  $M_x$ , and (c) pitch moment  $M_y$  derived by using the predictions from 2D and *Ours*.

shows an eventual divergence on the open-loop prediction, due to the absence of uncertainty consideration.

The result of P-TA indicates that incorporating stochasticity into a model reduces the accumulation of prediction errors. Additionally, integrating history information of states and control inputs provides further enhancements by capturing the time-varying characteristics of vehicle-terrain interaction, as shown by the result of P-TA w/ history. Finally, our proposed model shows the most accurate multi-step predictions by using the orientation angles of the vehicle as input to the model. This

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

TABLE II

EXPERIMENTAL RESULTS ON THE RACE TRACK. TO ASSESS THE DIFFICULTY OF THE DRIVING TASK, WE CONDUCTED EXPERIMENTS USING A BASELINE MODEL AT BOTH  $v_{\text{REF}} = 30$  KM/H AND  $v_{\text{REF}} = 40$  KM/H. WE SET  $K$  AND  $T$  AS 2000 AND 20, RESPECTIVELY, WHICH ENABLES OUR ALGORITHM TO OPERATE AT 10HZ ON THE NVIDIA RTX 3090 GPU WITH CUDA AND PYTORCH. WE DISPLAYED THE MEAN AND STANDARD DEVIATION FOR LAP TIME AND  $F_z^{\text{PEAK}}$ .

#	Model	$v_{\text{ref}}$ (km/h)	Off-Road Cost Functions			Lap Time (s)	# of Failure	$ v _{\text{avg}}$ (km/h)	$ v _{\text{max}}$ (km/h)	$ \phi _{\text{max}}$ (deg)	$ \theta _{\text{max}}$ (deg)	$F_z^{\text{peak}}$ (kN)
			Rollover( $\mathbf{x}$ )	Force( $\mathbf{x}$ )	UC( $\mathbf{x}$ )							
1	2D	30	✗	✗	✗	127.62 ± 2.21	5	28.79	39.16	79.21	31.53	50.11 ± 20.17
2	2D	40	✗	✗	✗	116.59 ± 4.76	19	32.56	42.22	79.41	36.90	53.23 ± 24.27
3	Ours	40	✗	✗	✗	117.45 ± 3.69	17	32.09	41.85	78.09	33.91	53.86 ± 21.96
4	Ours	40	✓	✗	✗	120.86 ± 8.55	23	32.49	41.29	79.25	38.74	53.70 ± 26.44
5	Ours	40	✓	✓	✗	136.32 ± 2.07	3	28.04	41.37	78.71	76.42	44.50 ± 19.97
6	Ours	40	✓	✓	✓	134.69 ± 1.64	0	28.11	40.48	28.50	28.87	43.17 ± 16.22

is especially critical in terrains with severe curvatures or large altitude fluctuations, where changes in roll and pitch angles cannot be disregarded even if they are trivial in relatively flat terrain. Fig. 4 illustrates a multi-step prediction of 2D and Ours at a specific time step.

We also visualize  $(F_z, M_x, M_y)$  derived from (10), using the predictions obtained from 2D and Ours on  $\mathcal{D}_{\text{val}}$  (see Fig. 5). The result indicates that the estimated forces and moments of our model closely align with the ground truth values obtained from force sensors in the simulator, without any training process involving ground truth force data. Notably, Ours can reliably predict highly unsafe behavior in which the vehicle loses contact with the ground ( $F_z = 0$ ) or lands again, causing a strong collision with the ground ( $F_z \gg mg$ ). These findings validate the idea that integrating a force cost into the controller can help mitigate the risks resulting from severe impact with the ground.

#### D. Autonomous Rally Driving Task

The purpose of this section is to illustrate that both safe and aggressive driving over off-road terrain is achievable when the MPPI controller leverages information derived from our terrain-aware kinodynamic model. We evaluated the control performance of the MPPI using Ours with all the off-road costs, and manually defined other ablations by sequentially removing the following three off-road cost terms from  $c(\mathbf{x})$  for comparison: UC( $\mathbf{x}$ ), Force( $\mathbf{x}$ ), and Rollover( $\mathbf{x}$ ). These ablations were performed to assess the contribution of each term to the control performance. The reference speed  $v_{\text{ref}}$  was set to 40 km/h. Additionally, we evaluated the control performance of the MPPI using 2D at reference speeds of 30 km/h and 40 km/h. Note that it is infeasible to apply off-road costs when using 2D. This evaluation aimed to establish a lower bound baseline and assess the difficulty of high-speed driving on off-road terrain.

The objective of each controller was to complete 10 counterclockwise laps around the track. In these experiments, we employed the same race track that was used for validation data collection. If the vehicle deviated from the track or significantly lost its stability, we repositioned it at the centerline of the track and continued driving for the remaining laps. To evaluate the controllers, we measured the maximum roll and

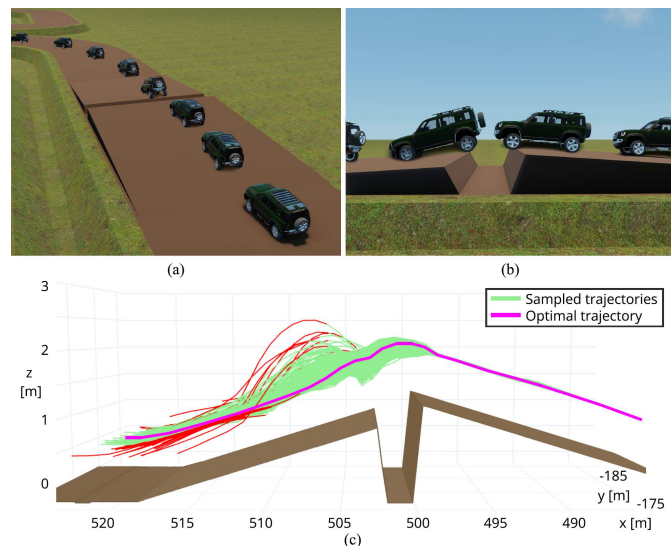


Fig. 6. (a) The back view, and (b) the side view of actual vehicle trajectory while crossing a cracked region. (c) The sampled trajectories of the MPPI right before the vehicle crosses the crack and the computed optimal trajectory in that time step. Intriguingly, it is estimated that trajectories with low speeds at which cannot be crossed are subjected to a significant impact on the vehicle's rear wheels by the terrain, resulting in a loss of stability (colored in red).

pitch angles of the vehicle and the range of impulsive vertical force caused by shock impacts with the ground.

Table II presents the experimental results. Results related to 2D demonstrates that ignoring the terrain geometry can lead to catastrophic failures, such as overturning, and frequent loss of ground contact, hence increasing both the frequency and magnitude of collisions with the ground. Similar tendencies are observed when using Ours without any off-road cost terms, or when using only the rollover cost in  $c(\mathbf{x})$ . This is because the orientation angles of the vehicle are insufficient to fully identify unsafe conditions such as strong ground impacts. We observe that the addition of the force cost considerably reduces the intensity of impulsive impact on the vehicle and lowers the frequency of failures. Nonetheless, failure still occurred because the estimated state cost based on model prediction does not consider the uncertainty. Finally, the inclusion of the uncertainty cost term guided the vehicle to a region with less uncertainty, resulting in the completion of 10 laps without any failure and a reduced magnitude of impulsive force.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

Lastly, we illustrate a specific scenario where our model facilitates the controller to execute more aggressive maneuvers in order to traverse a challenging terrain. Fig. 6 exhibits the control outcomes of the MPPI using *Ours* when crossing a cracked area on the terrain that cannot be traversed at low speeds. Our terrain-aware model is able to predict the impact by the cracked road, including a complete loss of stability at relatively low speeds. Hence, the controller slightly accelerates the vehicle to successfully cross the cracked region. In contrast, 2D baseline stops the vehicle in front of the crack, judging that it is a region that cannot be traversed. This result demonstrates that the use of our terrain-aware model in the controller improves the vehicle's maneuverability by considering the distinctive features of the terrain, instead of adopting a conservative approach by simply lowering the speed and cautiously driving over challenging terrain.

## VI. CONCLUSION

We presented a 6-DoF terrain-aware kinodynamic model for high-speed vehicle control on off-road terrain. Our model shows reliable predictions of vertical, roll, and pitch motions closely related to stability and physical feasibility, by using proprioceptive and exteroceptive information as inputs along with a few design considerations. Additionally, our model can predict contact interactions in terms of forces and moments, allowing for assessments of ground impact severity without any training using ground truth force data. Our MPC can utilize these predictions to execute safe and aggressive control inputs, leading to improved vehicle maneuverability. Our experiments demonstrate that our model-controller pair outperforms the baseline in terms of safety and robustness. Our future research will extend this work by applying this model to real-world vehicles via sim-to-real transfer.

## REFERENCES

- [1] J. Betz, H. Zheng, A. Liniger, U. Rosolia, P. Karle, M. Behl, V. Krovı, and R. Mangharam, "Autonomous vehicles on the edge: A survey on autonomous vehicle racing," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 458–488, 2022.
- [2] M. Tarokh and G. J. McDermott, "Kinematics modeling and analyses of articulated rovers," *IEEE Transactions on Robotics*, vol. 21, no. 4, pp. 539–553, 2005.
- [3] T. M. Howard and A. Kelly, "Optimal rough terrain trajectory generation for wheeled mobile robots," *The International Journal of Robotics Research*, vol. 26, no. 2, pp. 141–166, 2007.
- [4] W. Khalil and E. Dombre, *Modeling, identification and control of robots*. CRC Press, 2002.
- [5] M. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *Proc. International Conference on Machine Learning (ICML)*, 2011, pp. 465–472.
- [6] A. Nagabandi, G. Yang, T. Asmar, R. Pandya, G. Kahn, S. Levine, and R. S. Fearing, "Learning image-conditioned dynamics models for control of underactuated legged millirobots," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4606–4613.
- [7] R. Sonker and A. Dutta, "Adding terrain height to improve model learning for path tracking on uneven terrain by a four wheel robot," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 239–246, 2020.
- [8] G. Kahn, P. Abbeel, and S. Levine, "BADGR: An autonomous self-supervised learning-based navigation system," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.
- [9] S. J. Wang, S. Triest, W. Wang, S. Scherer, and A. Johnson, "Rough terrain navigation using divergence constrained model-based reinforcement learning," in *Proc. Conference on Robot Learning (CoRL)*, 2021, pp. 224–233.
- [10] J.-F. Tremblay, T. Manderson, A. Noca, G. Dudek, and D. Meger, "Multimodal dynamics modeling for off-road autonomous vehicles," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1796–1802.
- [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 *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 3294–3301.
- [12] S. Triest, M. Sivaprakasam, S. J. Wang, W. Wang, A. M. Johnson, and S. Scherer, "TartanDrive: A large-scale dataset for learning off-road dynamics models," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2546–2552.
- [13] P. Krüsi, P. Furgale, M. Bosse, and R. Siegwart, "Driving on point clouds: Motion planning, trajectory optimization, and terrain assessment in generic nonplanar environments," *Journal of Field Robotics*, vol. 34, no. 5, pp. 940–984, 2017.
- [14] R. O. Chavez-Garcia, J. Guzzi, L. M. Gambardella, and A. Giusti, "Learning ground traversability from simulations," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1695–1702, 2018.
- [15] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, "Real-time semantic mapping for autonomous off-road navigation," in *Proc. International Conference on Field and Service Robotics (FSR)*, 2017, pp. 335–350.
- [16] A. J. Sathyamoorthy, K. Weerakoon, T. Guan, J. Liang, and D. Manocha, "TerraPN: Unstructured terrain navigation using online self-supervised learning," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 7197–7204.
- [17] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, "ScaTE: A scalable framework for self-supervised traversability estimation in unstructured environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 888–895, 2023.
- [18] 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.
- [19] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [20] L. Wellhausen, A. Dosovitskiy, R. Ranftl, K. Walas, C. Cadena, and M. Hutter, "Where should I walk? predicting terrain properties from images via self-supervised learning," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1509–1516, 2019.
- [21] P. Fankhauser, M. Bloesch, and M. Hutter, "Probabilistic terrain mapping for mobile robots with uncertain localization," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3019–3026, 2018.
- [22] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [23] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [24] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- [25] T. Kim, J. Mun, J. Seo, B. Kim, and S. Hong, "Bridging active exploration and uncertainty-aware deployment using probabilistic ensemble neural network dynamics," in *Proc. Robotics: Science and Systems (RSS)*, 2023.
- [26] N. A. Spielberg, M. Brown, N. R. Kapania, J. C. Kegelmann, and J. C. Gerdes, "Neural network vehicle models for high-performance automated driving," *Science robotics*, vol. 4, no. 28, p. eaaw1975, 2019.
- [27] T. Kim, G. Park, K. Kwak, J. Bae, and W. Lee, "Smooth model predictive path integral control without smoothing," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10406–10413, 2022.
- [28] K. M. Lynch and F. C. Park, *Modern robotics*. Cambridge University Press, 2017.
- [29] R. Egele, R. Maulik, K. Raghavan, B. Lusch, I. Guyon, and P. Balaprakash, "AutoDEUQ: Automated deep ensemble with uncertainty quantification," in *Proc. International Conference on Pattern Recognition (ICPR)*, 2022, pp. 1908–1914.
- [30] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," in *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 7559–7566.