

# METAVerse: Meta-Learning Traversability Cost Map for Off-Road Navigation

Junwon Seo<sup>1</sup>, Taekyung Kim<sup>2</sup>, Seongyong Ahn<sup>3</sup>, and Kiho Kwak<sup>3</sup>

**Abstract**—Autonomous navigation in off-road conditions requires an accurate estimation of terrain traversability. However, traversability estimation in unstructured environments is subject to high uncertainty due to the variability of numerous factors that influence vehicle-terrain interaction. Consequently, it is challenging to obtain a generalizable model that can accurately predict traversability in a variety of environments. This paper presents *METAVerse*, a meta-learning framework for learning a global model that accurately and reliably predicts terrain traversability across diverse environments. We train the traversability prediction network to generate a dense and continuous-valued cost map from a sparse LiDAR point cloud, leveraging vehicle-terrain interaction feedback in a self-supervised manner. Meta-learning is utilized to train a global model with driving data collected from multiple environments, effectively minimizing estimation uncertainty. During deployment, online adaptation is performed to rapidly adapt the network to the local environment by exploiting recent interaction experiences. To conduct a comprehensive evaluation, we collect driving data from various terrains and demonstrate that our method can obtain a global model that minimizes uncertainty. Moreover, by integrating our model with a model predictive controller, we demonstrate that the reduced uncertainty results in safe and stable navigation in unstructured and unknown terrains.

## I. INTRODUCTION

During off-road navigation, autonomous vehicles encounter diverse unstructured terrains with distinct characteristics. To ensure safe and stable navigation in off-road environments, it is necessary to predict the interaction of a vehicle with terrains in an upcoming trajectory [1]. The predicted difficulty of interaction, which represents the traversability of the terrain, plays critical roles in various navigational strategies such as local path planning and control [2], [3]. However, estimating terrain traversability accurately in off-road environments with limited sensors is challenging. Even though human-annotated datasets can be utilized to train a semantic classifier, off-road environments are fraught with unseen and ambiguous terrains, resulting in inaccurate predictions of traversability. Furthermore, terrains of the same terrain class may have varying degrees of traversability due to their complex and variable traversability-related properties, which cannot be captured through manually labeled data.

This work was supported by the Agency For Defense Development Grant funded by the Korean Government in 2024.

<sup>1</sup>Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA. junwonseo@andrew.cmu.edu

<sup>2</sup>Department of Robotics, University of Michigan, Ann Arbor, MI, USA. taekyung@umich.edu

<sup>3</sup>Agency for Defense Development, Daejeon, Republic of Korea. {seongyong.ahn, kkwak.add}@gmail.com

Our video can be found at <https://youtu.be/4rIAMM1ZKMo>

Recently, off-road navigation has benefited significantly from self-supervised approaches that utilize vehicle-terrain interaction of actual navigation experiences to learn terrain traversability [4]–[6]. These methods learn a mapping from exteroceptive data (e.g., RGB and LiDAR point clouds) to the traversability cost defined by the vehicle-terrain interaction measured with proprioceptive sensors (e.g., Inertial Measurement Unit (IMU)). The resultant traversability cost maps with continuous-valued scores can precisely represent the difficulty of navigation. Moreover, these cost maps reflect the navigation capabilities of a vehicle and the terrain characteristics based on actual experiences, resulting in enhanced navigational performance.

However, these methods cannot be generalized to various environments since it is challenging to acquire a global model that operates reliably in numerous environments. As the interaction data only provides supervision for regions the vehicle has interacted with, the model’s predictions for unexplored terrains are subject to high epistemic uncertainty [7]. Even if interaction data are obtained on a variety of terrains to reduce epistemic uncertainty, the estimation of the model is still subject to a substantial amount of aleatoric uncertainty. In real-world off-road environments, the traversability of the terrain is intricately influenced by a multitude of interconnected and complex factors (e.g., platform, geometry, deformability, bumpiness, friction, and roughness) [8]. Nonetheless, such subtle variations cannot be captured precisely with a limited sensor configuration. While the geometric characteristics of the terrain captured by a sensor would be comparable, its ground-truth traversability would vary considerably, resulting in a high level of aleatoric uncertainty. These uncertainties result in less precise predictions and make it impossible to estimate traversability costs that are more nuanced than simple terrain classification, resulting in suboptimal off-road navigational performances.

This paper presents a meta-learning-based framework for learning terrain traversability (*METAVerse*), which is capable of learning a global model that accurately predicts terrain traversability in various off-road environments. During training, the traversability cost prediction network is learned to generate a dense, continuous-valued cost map in real-time from a single sweep LiDAR point clouds. To minimize the uncertainty of the network trained with data collected from various terrains, meta-learning is employed to discover initial network parameters that enable effective generalization with a few gradient descent steps. During deployment, the model quickly adapts to the local context with gradient descents based on the recent history of vehicle-terrain interactions,

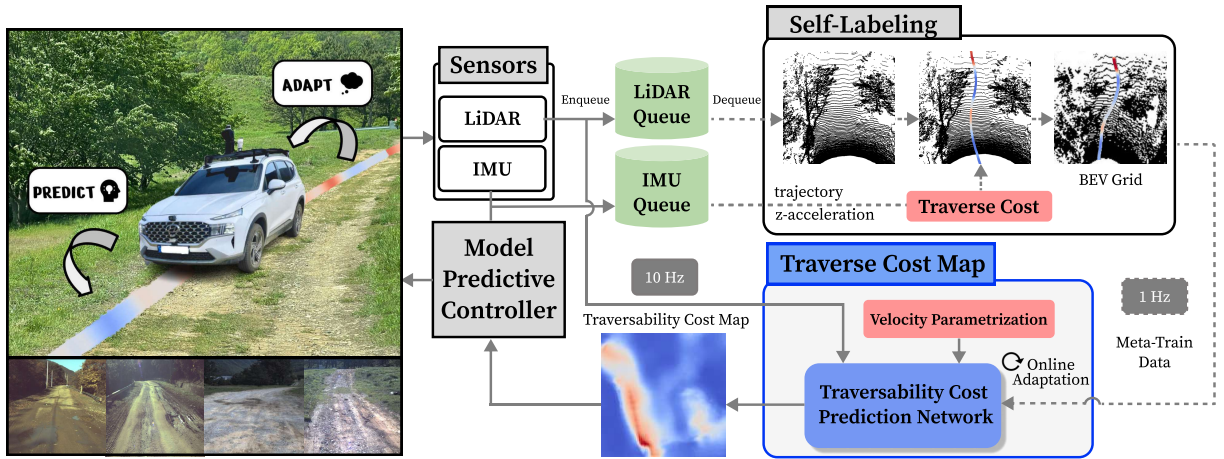


Fig. 1: Diagram of our method for off-road navigation that can be implemented in various environments. The traversability cost prediction network produces a dense and continuous-valued cost map in BEV. Using self-labeled data, the network minimizes uncertainty through online adaptation during deployment. A model predictive controller utilizes the resultant accurate and reliable traversability cost map for safe and stable off-road navigation.

allowing for effective off-road navigation. We demonstrate that our method can learn a global model that reduces prediction uncertainty using real-world driving data collected on unstructured terrains with varying properties. In addition, by integrating our framework with a sampling-based model predictive controller, we demonstrate that our method facilitates stable and safe navigation in unstructured environments. An overview of our framework is presented in Fig. 1.

In summary, the main contributions of our work are:

- We introduce a deep meta-learning framework for learning a global terrain traversability prediction network that can reliably generate a dense cost map from a sparse LiDAR point cloud in various environments.
- We propose a method to minimize uncertainty in estimation by online adapting the network based on the vehicle’s navigation experience.
- We demonstrate that our method can reduce uncertainty in prediction using data collected in various terrains.
- We validate that our method leads to stable navigation by integrating our traversability cost map with a sampling-based model predictive controller.

## II. RELATED WORKS

### A. Traversability Estimation in Off-Road

Earlier works estimate traversability based on rule-based features derived from geometric and visual appearances, such as the terrain’s roughness and slope [9]–[11]. With the advent of deep neural networks, semantic segmentation has been widely utilized to classify off-road terrains according to their navigability levels [12], [13]. Recent work [14] has utilized semantic scene completion to generate a dense terrain classification map in bird’s eye view (BEV) from a sparse point cloud [15].

Unfortunately, these human-supervised methods cannot provide adequate information for effective navigation in complex and unstructured off-road environments. The cost assigned to a predetermined terrain class would be irrelevant to a given environment or inaccurately depict navigability [16]. Self-supervised approaches use terrain interaction

feedback to circumvent such limitations [5], [6], [17]. Using information about the terrain traversed by a vehicle, they identify traversable regions or classify terrains into multiple classes to designate differential costs [4], [18], [19]. Nevertheless, these approaches abstract away the subtle variations in traversability within the same class [8].

Recent research has shifted toward predicting a continuous-valued cost for more effective navigation in unstructured environments [20]–[24]. To define the continuous costs, inertial information pertinent to navigational stability is processed [8], [20], [23], or a reinforcement learning framework is leveraged [3], [25], [26]. By incorporating the local cost map into path planning and control, these methods improve navigational performance in terms of stability and safety. Nevertheless, these methods prioritize distinguishing between terrain types over evaluating subtle differences in traversability among terrains of the same class. Also, these approaches cannot account for terrains with unknown properties [18]. While recent works detect and avoid unobserved terrains [7], [27], they cannot obtain a global model that adapts or generalizes to numerous environments.

### B. Meta-Learning

The goal of meta-learning is to train a model that can rapidly adapt to new tasks [28]. Model-agnostic meta-learning (MAML) [29] obtains the initial parameters of a neural network such that taking a few gradient descent steps from this initialization results in effective generalization. By considering environments or situations with distinct characteristics as different tasks, meta-learning can learn a global initial parameter from heterogeneous datasets without confusion. During inference, the global model is generalizable to numerous tasks, including a novel one, through a few gradient updates [30].

Various works in robotics literature have adopted meta-learning to learn a global model that can reduce uncertainty through rapid adaptation [31]. Nagabandi et al. [32] trained the dynamics model of a legged robot, which rapidly adapts

to its local environment. Recently, Visca et al. [33] proposed a meta-adaptive energy predictor for path planning in unknown terrains.

### III. METHODS

This section details our proposed framework for learning traversability. First, we present our traversability cost prediction network, which predicts the continuous traversability cost derived from vehicle-terrain interaction and generates a dense cost map in BEV using a single sweep LiDAR point cloud. Then, we describe our meta-learning method for training the network (*METAverse*) to acquire a global model that is generalizable in various environments.

#### A. Dense Traversability Cost Map

Off-road environments are fraught with bumps and obstacles of varying shapes, despite being in the same terrain class. For safe and effective navigation in off-road terrain, estimating the nuanced traverse cost of the terrain is necessary. A path planner can optimize a trajectory that minimizes disturbances during navigation with the predicted cost. Therefore, we generate dense and continuous traverse cost maps in BEV from a single LiDAR point cloud.

The z-axis linear acceleration measured by an IMU can be utilized to define traversability cost derived from vehicle-terrain interaction. This component effectively captures the terrain’s properties related to the stability of the vehicle in off-road navigation [34]. In addition, the definition of traversability based on vertical acceleration can be advantageous for control performance because model-based controllers frequently employ a vehicle dynamics model ignorant of vertical motions for computational simplicity [35], [36].

Motivated by recent work [8], traversability cost is defined using the spectral analysis of z-axis linear acceleration. The wavelet power spectrum is used to precisely characterize the costs of a time series signal, as it eliminates the need to segment signals and apply the Fourier transform to each segment. A continuous wavelet transformation with the Morlet wavelet is performed on the z-acceleration  $a_z(t)$  to generate the wavelet coefficient  $w_z(f_n, t)$  for each frequency scale  $f_n = 2^n \cdot f_0$  and time step  $t$ . Then, the traversability cost  $c_t$  is defined by the wavelet power spectrum as follows:

$$c_t = \sum_{n=0}^j \frac{\|w_z(f_n, t)\|^2}{f_n}, \quad (1)$$

where squares of the coefficient are divided by frequency scale to rectify the power spectrum, as suggested by Liu et al. [37], and summed over a certain frequency scale range of  $f_0 = 0.16$  to  $f_j = 5.12$ . The calculated ground-truth costs are then assigned to BEV grids along the positions of the trajectory and used for training the traversability cost prediction network. Note that traversability cost can be defined in multiple other ways using proprioceptive signals.

The traversability cost prediction network is trained to produce a dense top-view cost map using a sparse single-sweep LiDAR point cloud, as shown in Fig. 2. Following PointPillars [38], each point is discretized into sparse pillars.

The point in each pillar is encoded as a 4-dimensional feature consisting of offset from the pillar center and distance from origin  $(\Delta x, \Delta y, \Delta z, d)$ . Using a simplified PointNet [39] that consists of a linear layer, BatchNorm, and ReLU, each pillar of size  $(N, 4)$  is converted into sparse pillar features of size  $C$ , where  $N$  is the maximum number of points per pillar. Each pillar feature is scattered back to the pillar locations to create a BEV sparse feature representation of size  $(H, W, C)$ , where  $H$  and  $W$  denote the width and height of the grid, respectively. The empty pillars are zero-initialized. A U-Net [40] structured network, which has an encoder-decoder architecture with skip connections, is employed to generate a dense pillar feature map of size  $D$ . It progressively reduces the spatial size of features and captures higher-level semantic information while the decoder upsamples feature maps to recover spatial information.

The dense pillar features are concatenated with parameterized velocity to produce velocity-conditioned cost maps. Fourier feature mapping [41] is used to incorporate the vehicle’s velocity into the cost prediction [8]. The velocity vector is mapped into a higher dimensional representation:

$$\gamma(v) = [\cos(2\pi b_1 v), \sin(2\pi b_1 v), \dots, \cos(2\pi b_P v), \sin(2\pi b_P v)], \quad (2)$$

where  $v$  is the norm of the velocity vector,  $b_i \sim \mathcal{N}(0, 5^2)$  are sampled from a Gaussian distribution, and  $P = 10$  is the number of samples. Finally, the MLP head predicts the mean  $\mu_i$  and standard deviation  $\sigma_i$  of the traversability for each pillar  $i$ . The network is trained to minimize the Gaussian log-likelihood:

$$\mathcal{L}^{\text{traverse}}(\tau, \theta) = \frac{1}{2} \sum_i \left( \log(\sigma_i) + \frac{(\mu_i - c_i)^2}{\sigma_i} \right), \quad (3)$$

where  $\theta$ ,  $\tau$ , and  $c_i$  represent the model parameter, driving data along a trajectory segment, and the ground truth cost associated with the pillar  $i$ , respectively. The loss calculation is restricted to pillars assigned with ground truth, that is, the vehicle has traversed. Multiple data augmentations, such as random flip, rotation, and translation, are implemented during training to prevent overfitting and produce a dense cost map.

#### B. METAverse: Meta-Learning Traversability Cost Map

Learning a global traversability model using a large dataset  $\mathcal{D}$  of multiple environments leads to high aleatoric uncertainty. To address this problem, we propose a method that can effectively learn a global traversability model capable of rapidly adapting to a new environment based on its recent experiences. Meta-learning can be used to learn a global model of predicting self-supervised traversability cost because it can handle the variability of interaction data obtained from distinct environments. In addition, by performing online adaptation with self-labeled data, it would be able to accurately predict traversability costs during deployment in a variety of environments, including unknown ones.

MAML [29] is used to learn the global traversability model. MAML aims to find the initial parameters of the

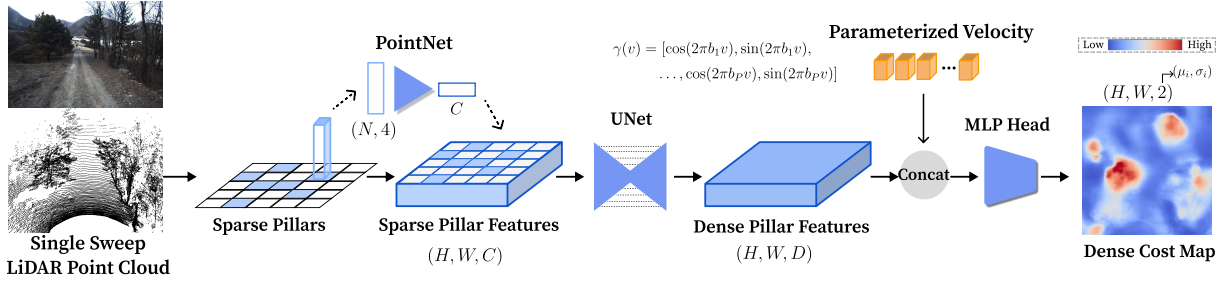


Fig. 2: Overview pipeline of the traversability cost prediction network. From a single sweep LiDAR point cloud, the network generates a dense and continuous-valued traversability cost map in BEV.

network so that adaptation with a few gradient descent steps from this initialization leads to effective generalization to the current circumstances. This meta-objective enables the model for predicting traversability cost to incorporate driving data collected in various environments without confusion caused by aleatoric uncertainty in the training phase. During the deployment phase, the network is updated based on recent vehicle-terrain interaction experiences to adapt to dynamic environments and generate an accurate cost map.

While terrain properties vary significantly in different environments, we assume the environment is locally consistent. Consequently, vehicle-terrain interaction of each local trajectory segment, denoted as  $\tau$ , is regarded as a separate *task*. Instead of considering the entire dataset with distinct properties as a single task, the network is trained with the meta-objective that the recent experiences can provide information about the current task. The past  $M$  timesteps provide insight into how to adapt the model to predict the traversability costs of future trajectories precisely. The network is trained to adapt using the *meta-train* data of the past  $M$  timesteps,  $\tau(t-M, t)$ , to predict the traversability cost of *meta-eval* data from the next  $K$  timesteps,  $\tau(t, t+K)$ , as follows:

$$\begin{aligned} \underset{\theta}{\operatorname{argmin}} \quad & \mathbb{E}_{\tau(t-M, t+K) \sim \mathcal{D}} [\mathcal{L}^{\text{traverse}}(\tau(t, t+K), \theta')] \\ \text{s.t.} \quad & \theta' = \theta - \alpha \nabla_{\theta} \mathcal{L}^{\text{traverse}}(\tau(t-M, t), \theta). \end{aligned} \quad (4)$$

Algorithm 1 outlines the meta-learning-based training procedure of *METAverse* for obtaining the global traversability model. The inner loops adapt the model with meta-train data by taking  $N_A$  adaptation steps via gradient descent. The outer loop updates the initial parameters with losses calculated with  $N$  trajectories within a minibatch.

During deployment, the traversability prediction network is online adapted utilizing meta-train data, as illustrated in Fig. 1. Meta-train data is automatically generated by self-labeling sensor data from LiDAR and IMU that are stored in queues. The online adaptation of the network is performed asynchronously, and only inner loops are executed to accomplish rapid adaptation. The updated network generates an accurate representation of environments, which is then employed for navigation by a model predictive controller. The online adaptation enables the trained global model to be applicable in various environments and even adapt to terrains with unknown properties.

---

### Algorithm 1: Meta Learning of Traversability Cost

---

**Given:**  $\mathcal{D}$ : Traversability data from various environments;  
 $M, K$ : Number of past and future timesteps;  
 $N$ : Number of sampled trajectories within a batch;  
 $N_A$ : Number of the inner loops;  
 $\alpha, \beta$ : Learning rates for the inner and outer loops;  
 Randomly initialize  $\theta$ ;  
**for**  $i \leftarrow 0$  **to** maximum iterations **do**  
   **for**  $j \leftarrow 0$  **to**  $N - 1$  **do**  
     Sample  $\tau(t - M, t), \tau(t, t + K) \sim \mathcal{D}$ ;  
     Self-Label  $\tau(t - M, t)$  and  $\tau(t, t + K)$ ;  
      $\theta' \leftarrow \theta$ ;  
     **for**  $k \leftarrow 0$  **to**  $N_A - 1$  **do**  
        $\theta' \leftarrow \theta' - \alpha \nabla_{\theta'} \mathcal{L}^{\text{traverse}}(\tau(t - M, t), \theta')$ ;  
        $\mathcal{L}_j \leftarrow \mathcal{L}^{\text{traverse}}(\tau(t, t + K), \theta')$   
      $\theta \leftarrow \theta - \beta \nabla_{\theta} \frac{1}{N} \sum_{j=1}^N \mathcal{L}_j$

---

## IV. EXPERIMENTS

In this section, we validate the efficacy of our proposed framework. We first demonstrate that our method is capable of learning a global model that minimizes uncertainty in prediction, resulting in a more accurate prediction of traversability costs in diverse environments. We then verify that this accurate and dependable traversability cost prediction leads to stable and effective off-road navigation.

Our experiments address the following key questions: **(Q1)** Can our method learn a global model that minimizes uncertainty for learning traversability in various environments? **(Q2)** Does our method facilitate safe and effective navigation in unstructured environments? **(Q3)** Can our traversability prediction network adapt effectively to unknown terrains during navigation?

### A. Implementation Details

For all experiments, the input point cloud is cropped at  $[(0, 51.2), (-25.6, 25.6), (-5, 10)]$  meters along the  $x, y, z$  axes, and a pillar grid size of  $0.2m \times 0.2m$  is used. For the traversability prediction network, the maximum number of points per pillar is set to  $N = 20$ , and the channels of sparse and dense pillar features are set to  $C = 128$  and  $D = 64$ , respectively. Each encoder and decoder has five layers, each consisting of a max pooling or transposed convolution layer and two  $3 \times 3$  convolution layers, with a ReLU and a BatchNorm layer in the middle.

TABLE I: The details of the evaluation dataset. The sequence length in time (*Total Len*) and the standard deviations of the ground-truth traversability costs (*STDEV of Cost*) are displayed for each category.

	Unpaved	Grassland	Profiled Road	Simulation
<i>Total Len</i> (s)	531.5	390.2	1056.9	545.9
<i>STDEV of Cost</i>	0.4751	0.5166	0.4638	5.5651




Fig. 3: Real-world driving data. Based on the terrain characteristics, the evaluation data is divided into three distinct categories: (a) Unpaved, (b) Grassland, and (c) Profiled Road. RGB images and the traversability cost map generated by LiDAR points of these scenes are displayed. More results are available in the multimedia material.

Our model is trained for 60 epochs using the Adam optimizer with the outer loop learning rate of  $\beta = 3e^{-4}$  and a batch size of 16. Each trajectory data within a batch comprises *meta-train* data for the inner loop and *meta-eval* data for the outer loop. Each meta-train data is composed of eight LiDAR point clouds and the ground-truth traversability costs, which are generated from the trajectory of the previous  $M = 8$  seconds from a reference time. The meta-eval data also consists of eight point clouds and the ground truth, but they are generated based on the trajectory of the upcoming  $K = 8$  seconds. All network parameters are subject to adaptation and are updated  $N_A = 3$  times at a learning rate for inner loops of  $\alpha = 1e^{-4}$ . During training, random horizontal flipping is applied with a probability of 50%, random rotation along the z-axis is applied between  $(-\frac{\pi}{4}, \frac{\pi}{4})$  radians, and random translation is applied  $(-5, 5)$  meters in the  $x$  and  $y$  axes.

### B. Learning Global Traversability Model

We validate the efficacy of our meta-learning method for traversability cost prediction (Q1) with real-world off-road driving data. Specifically, the validity of the global traversability model trained with meta-objective (4) is evaluated in terms of prediction accuracy.

1) *Experimental Setup*: Driving data is collected in off-road environments with various types of terrain using our platform equipped with OS1-128 LiDAR and IMU (See Fig. 1). Also, the driving data is obtained in a simulation environment consisting of randomly patterned rough terrain and bumps, which is hazardous to interact with in

TABLE II: Validation error of the experiment with real-world driving data. Our method shows a significant margin compared to the baseline, and it can predict traversability more accurately through adaptations during inference.

Method	Adaptation	Evaluation Dataset Category			
		Unpaved	Grassland	Profiled Road	Simulation
<i>Baseline</i>	✗	0.1222	0.2460	0.2039	0.7263
<i>METAverse</i>	✗	0.0713	0.1961	0.1907	0.6668
<i>METAverse</i>	✓	<b>0.0114</b>	<b>0.1767</b>	<b>0.1523</b>	<b>0.5725</b>

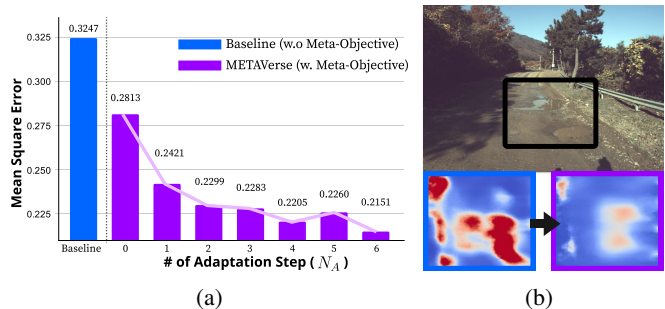


Fig. 4: (a) Mean square error for the evaluation dataset. The error decreases as the number of adaptation steps increases. Our model reduces prediction uncertainty significantly compared to a model trained without a meta-objective (*Baseline*). (b) The initially inaccurate cost map is updated through adaptation, incorporating the recent experience of vehicles traversing similar bumps and puddles to predict the traversability cost more accurately.

real-world environments [7]. Approximately three hours of driving data are utilized to train the network. Then, the trained traversability prediction network is evaluated using an evaluation dataset consisting of separate trajectory sequences not included in the training dataset.

According to the terrain characteristics, we divide our real-world evaluation dataset into 3 categories, namely, *unpaved*, *grassland*, and *profiled road*. Fig. 3 shows example images of these terrains along with their corresponding visualizations of traversability cost maps. The *unpaved* consists of rough and unpaved dirt tracks with numerous bumps and puddles. The *grassland* represents grassy roads with vegetation, bushes, and cobblestones. The *profiled road* is composed of roads with five different types of artificial road profiles of varying sizes that are used for assessing driving stability. Lastly, *simulation* category is added for the evaluation, which refers to data collected while navigating the off-road evaluation track in the simulation setup (See Fig. 5). The details of our evaluation dataset are presented in Table. I.

For comparison, a network is trained using all train data without meta-objective and adaptation (*Baseline*). Also, the network trained with meta-objective is evaluated with zero adaptation steps ( $N_A=0$ ) during inference as well as with varying numbers of adaptation steps. Mean Square Error (MSE) between the ground-truth traversability and the predicted traversability cost of corresponding grids is calculated for the evaluation.

2) *Experimental Result*: The quantitative results are presented in Table II. In every category, the model trained with the meta-objective outperforms the baseline model trained without the meta-objective. It indicates that the non-meta-learned baseline failed to converge well due to high aleatoric uncertainty in ground-truth traversability collected across

TABLE III: Quantitative results for the navigation experiments (Q2). The average and maximum motions of the vehicle across all trials are shown. A trial is considered unsuccessful if the vehicle deviates off the track or flips over before completing the lap.

Method	Self Supervised	Online Adaptation	Success Rate	Vertical Vel. [m/s]		Vertical Acc. [m/s <sup>2</sup> ]		Roll Rate [rad/s]		Pitch Rate [rad/s]		Roll Acc. [rad/s <sup>2</sup> ]		Pitch Acc. [rad/s <sup>2</sup> ]	
				Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max
<i>Elevation Based</i>	✗	✗	0.46	0.1522	2.7664	1.2543	40.1081	0.1433	0.1516	1.4675	2.3831	1.2877	54.7836	1.1770	34.0290
<i>Slope Based</i>	✗	✗	0.53	0.1286	1.8467	0.9994	30.4818	0.1288	1.6243	0.1273	1.3686	1.0184	38.9109	0.9531	26.8223
<i>Point-wise</i>	✓	✗	0.80	0.1091	1.7433	0.8862	26.3765	<b>0.1007</b>	1.4877	0.1180	1.4961	<b>0.8528</b>	36.2358	0.8516	23.7459
<i>METAverse</i>	✓	✗	0.93	0.1094	1.6901	0.8684	27.9480	0.1071	1.6461	0.1180	1.5383	0.8990	36.0222	0.8780	27.0468
<i>METAverse</i>	✓	✓	<b>1.00</b>	<b>0.1064</b>	<b>1.5913</b>	<b>0.8542</b>	<b>25.6867</b>	0.1038	<b>1.4582</b>	<b>0.1134</b>	<b>1.4349</b>	0.8926	<b>32.2179</b>	<b>0.8495</b>	<b>22.0796</b>

various terrains. In contrast, our meta-learned model can converge well by incorporating such uncertainty during training and adapting the model during inference. Even without adaptation in the evaluation phase, our method outperforms the baseline, implying that utilizing the meta-objective leads to a finding of a better initial parameter [29]. In addition, the model’s performance improves as it adapts during inference using recent interaction experiences.

Fig. 4a shows the experimental results on the whole evaluation data with varying numbers of adaptation steps during inference. The accuracy improves as the number of adaptation steps increases. Also, our method with meta-objective shows a significant margin over the baseline that does not conduct adaptation during both training and inference. After adaptation, the initially erroneous cost map is adjusted to more accurately represent the cost of terrains by incorporating recent experiences, as illustrated in Fig. 4b.

### C. Autonomous Navigation in Unstructured Environments

We demonstrate that our method can lead to stable navigation (Q2-Q3) by integrating our traversability cost map with a sampling-based model-predictive controller. A high-fidelity vehicle dynamics simulator - IPG CarMaker is utilized for the experiments, which enables navigation in a controlled setup and, thus, enables more comprehensive comparisons and analyses of the results. In all experiments, navigation is performed only using local traversability maps generated in real-time from LiDAR point clouds, with no prior knowledge of the environments or global path planning.

TABLE IV: Control parameters of SMPPI. Note that the definitions of the variables and the remaining parameters not specified in this paper are provided in the original paper [35].

Control Frequency	Target Speed	Trajectories	Horizon	Sampling Variance
10 Hz	30 km/h	5,000	4 s	Diag(1.6, 0.4)

1) *Experimental Setup*: For navigation, we employ the Smooth Model Predictive Path Integral (SMPPI) [35] controller, a sampling-based model predictive controller that can generate smooth actions during deployment. Table IV lists the controller’s parameters. Based on our previous work [35], we formulate a simple state-dependent running cost function  $q(\mathbf{x}_t)$  of the controller:

$$q(\mathbf{x}_t) = \alpha_1 \text{Track}(\mathbf{x}_t) + \alpha_2 \text{Stable}(\mathbf{x}_t) + \alpha_3 \text{Speed}(\mathbf{x}_t), \quad (5)$$

where  $\mathbf{x}_t$  represents the vehicle state and  $\text{Stable}(\mathbf{x}_t)$  is the predicted traversability cost. Based on our previous work for identifying non-traversable regions [7],  $\text{Track}(\mathbf{x}_t)$  imposes a significant penalty on regions with high uncertainty to prevent collisions.  $\text{Speed}(\mathbf{x}_t)$  is a simple quadratic cost that



Fig. 5: (Left) The off-road race track for the navigation experiment (Q2). (Right) Vehicle trajectories taken by employing various traversability maps. We observe that *METAverse* can navigate through relatively safer trajectories by accurately identifying nuanced traversability.

penalizes the difference between the target speed and the vehicle’s current speed. Each cost is normalized into  $[0, 1]$ , and the weight coefficients are set as  $\alpha_1 = 10000$ ,  $\alpha_2 = 10$ , and  $\alpha_3 = 1$ , namely  $\alpha_1 > \alpha_2 > \alpha_3$ . This setting ensures that the vehicle navigates only through traversable regions, and within those regions, it optimizes the trajectory to minimize traversability costs while maintaining the target speed as much as possible.

Based on our previous work [36], we use a probabilistic ensemble neural network as the vehicle dynamics. The state and action inputs follow the simplified bicycle model, and it also leverages the history of state action pairs to extract contextual information. For real-time implementation, we use a four-layer MLP and five ensembles.

2) *Safe Navigation*: An off-road environment (see Fig. 5) is designed to conduct navigation (Q2) where challenging unstructured terrains with large and irregularly patterned bumps make navigation solely with classification-based traversability maps inadequate. The control vehicle is a Volvo XC90, which is distinct from the vehicle used for training data collection. To evaluate the efficacy of our global model in enhancing navigation performance, navigation is performed with our method (*METAverse*) with and without online adaptation. Experiments are also conducted with a self-supervised method that predicts point-wise traversability (*Point-wise*) [7] instead of generating a dense map. In addition, rule-based methods are compared, including an elevation-map based method (*Elevation Based*) [42], [43] and a slope-based method (*Slope Based*) [9], [11]. We conduct navigation 15 times for each method.

The trajectories taken during navigation are shown in Fig. 5, and the statistics about angular and vertical motions of the vehicle and success rates of navigations are shown in Table III. The navigational performance of *METAverse* employing online adaptation is superior to that of our model without the adaptation. The traversability network is adjusted online to the novel environment and vehicle of deployment, resulting in improved navigational performance.




Fig. 6: (a) The off-road race track designed for online adaptation (Q3). The vehicle trajectories are visualized, and the colors of the lines illustrate the rotational impacts exerted on the vehicle. (b) Vertical acceleration of the vehicle during navigation. By conducting online adaptation, the vehicle can plan paths that can minimize impacts exerted on the vehicle, leading to stable navigation in off-road.

In addition, our method with voxelization to efficiently construct a dense cost map improves navigation performance compared to the point-wise self-supervised method, which predicts traversability cost using a point-wise network and generates local traversability maps via interpolation [7]. Due to real-time constraints, the network for point-wise prediction becomes shallower than the voxel-based network, which is sufficiently deep to generate a continuous cost map in BEV. The efficient network structure that directly generates maps in BEV enables our model to embed richer information to predict traversability more precisely in a limited amount of time, thereby enhancing real-time navigational performance. Lastly, the inability of the rule-based methods to reason about nuanced interactions with unstructured terrain leads to instability and an increase in navigational failure rates. On the other hand, the vehicle utilizing a self-supervised traversability map effectively navigates along paths that minimize disturbances. Our method can effectively identify risky regions where terrain adversely impacts vehicle stability.

3) *Effect of Online Adaptation:* To validate the efficacy of online adaptation to unknown terrains (Q3), which ultimately leads to safe navigation, an additional off-road environment is designed with several unknown types of bumps (see Fig. 6). The vehicle begins adaptation after experiencing three types of unknown bumps. This allows the vehicle to generate trajectories that minimize disturbance when re-encountering these bumps. The navigation is conducted for five trials using our trained traversability prediction networks, with and without online adaptation during inference.

Fig. 6a illustrates the navigation results. The vehicle experiences unseen bumps and effectively adapts the model to accommodate the experience. Initially, the predicted traversability of all unobserved bumps is not relatively differentiated because the model has no information about the bumps. As the network begins to be online-adapted using the experience, the predicted costs for bump #2 become lower than the costs for bump #1 and #3. Therefore, the controller chooses to circumvent the challenging bumps, resulting in successful navigation. In contrast, the vehicle without adaptation fails to discern the difficulties and avoid them, eventually resulting in a rollover.

Fig. 6b shows the vertical acceleration experienced by the vehicle during navigation. By beginning to online-adapt the network, the vehicle can reduce the impact exerted on it by adjusting its trajectory based on recent experiences, whereas the vehicle that does not conduct adaptation continues to ex-

TABLE V: Navigation results in the scene for evaluating online adaptation (Q3). The average vehicle motions across 5 trials are shown.

Online Adaptation	Vertical Vel. [m/s]	Vertical Acc. [m/s <sup>2</sup> ]	Roll Rate [rad/s]	Pitch Rate [rad/s]	Roll Acc. [rad/s <sup>2</sup> ]	Pitch Acc. [rad/s <sup>2</sup> ]
✗	0.3319	2.7385	0.2563	2.1254	2.9504	2.7333
✓	<b>0.1910</b>	<b>1.4347</b>	<b>0.1409</b>	<b>1.9994</b>	<b>1.4919</b>	<b>1.6075</b>

perience enormous impacts due to the inability to overcome uncertainty. The navigational stability measured by averaging the vertical and angular motions of the vehicle is presented in Table V. It verifies that adaptation can induce stable vehicle motions by adjusting the traversability prediction model to incorporate experiences with unknown terrains.

## V. CONCLUSION

This paper proposes a meta-learning framework for off-road traversability estimation. Our traversability prediction network predicts terrain traversability derived from vehicle-terrain interactions and generates a dense and continuous-valued cost map from a single-sweep LiDAR point cloud. Meta-learning is used to train a global model that can accurately predict terrain traversability in a variety of environments by minimizing uncertainty. During deployment, the network performs online adaptation utilizing recent interaction experiences to improve the accuracy of predictions. Extensive experiments demonstrate that the proposed method can reduce the uncertainty of the global model, resulting in stable off-road navigation in unstructured and unknown terrains. We believe this concept can be used for the broader deployment of autonomous robots in unstructured environments and improve the reliability and generalizability of off-road navigation systems that utilize self-supervision for learning traversability.

In future work, we intend to extend this framework using multiple sensor fusions to reduce uncertainty, such as RGB-LiDAR fusion. In addition, the integration of this framework for obtaining accurate terrain representation with vehicle dynamics learning would enhance the effectiveness of off-road high-speed navigation.

## REFERENCES

- [1] P. Borges, T. Peynot, S. Liang, B. Arain, M. Wildie, M. Minareci, S. Lichman, G. Samvedi, I. Sa, N. Hudson *et al.*, "A survey on terrain traversability analysis for autonomous ground vehicles: Methods, sensors, and challenges," *Field Robotics*, 2022.
- [2] D. D. Fan, A.-A. Agha-Mohammadi, and E. A. Theodorou, "Learning risk-aware costmaps for traversability in challenging environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 279–286, 2021.

- [3] J. Frey, D. Hoeller, S. Khattak, and M. Hutter, "Locomotion policy guided traversability learning using volumetric representations of complex environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 5722–5729.
- [4] D. Kim, J. Sun, S. M. Oh, J. M. Reh, and A. F. Bobick, "Traversability classification using unsupervised on-line visual learning for outdoor robot navigation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2006, pp. 518–525.
- [5] J. Zürn, W. Burgard, and A. Valada, "Self-supervised visual terrain classification from unsupervised acoustic feature learning," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 466–481, 2020.
- [6] 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.
- [7] 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.
- [8] M. Guaman Castro, S. Triest, W. Wang, J. M. Gregory, F. Sanchez, J. G. Rogers III, and S. Scherer, "How does it feel? self-supervised costmap learning for off-road vehicle traversability," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [9] J. Sock, J. Kim, J. Min, and K. Kwak, "Probabilistic traversability map generation using 3d-lidar and camera," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 5631–5637.
- [10] J. Ahtiainen, T. Stoyanov, and J. Saarinen, "Normal distributions transform traversability maps: Lidar-only approach for traversability mapping in outdoor environments," *Journal of Field Robotics*, vol. 34, no. 3, pp. 600–621, 2017.
- [11] J. Kim, J. Min, K. Kwak, and K. Bae, "Traversable region detection based on a lateral slope feature for autonomous driving of ugv's," *Journal of Institute of Control, Robotics and Systems*, vol. 23, no. 2, pp. 67–75, 2017.
- [12] B. Gao, S. Hu, X. Zhao, and H. Zhao, "Fine-grained off-road semantic segmentation and mapping via contrastive learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5950–5957.
- [13] T. Guan, D. Kothandaraman, R. Chandra, A. J. Sathiamoorthy, K. Weerakoon, and D. Manocha, "Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8138–8145, 2022.
- [14] A. Shaban, X. Meng, J. Lee, B. Boots, and D. Fox, "Semantic terrain classification for off-road autonomous driving," in *Conference on Robot Learning (CoRL)*, 2022, pp. 619–629.
- [15] J. Fei, K. Peng, P. Heidenreich, F. Bieder, and C. Stiller, "Pillarsegnet: Pillar-based semantic grid map estimation using sparse lidar data," in *IEEE Intelligent Vehicles Symposium (IV)*, 2021, pp. 838–844.
- [16] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, "Fast traversability estimation for wild visual navigation," in *Robotics: Science and Systems (RSS)*, 2023.
- [17] M. V. Gasparino, A. N. Sivakumar, Y. Liu, A. E. Velasquez, V. A. Higuity, J. Rogers, H. Tran, and G. Chowdhary, "Wayfast: Navigation with predictive traversability in the field," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10651–10658, 2022.
- [18] J. Seo, S. Sim, and I. Shim, "Learning off-road terrain traversability with self-supervisions only," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4617–4624, 2023.
- [19] T. Guan, R. Song, Z. Ye, and L. Zhang, "Vinet: Visual and inertial-based terrain classification and adaptive navigation over unknown terrain," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 4106–4112.
- [20] X. Yao, J. Zhang, and J. Oh, "Rca: Ride comfort-aware visual navigation via self-supervised learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 7847–7852.
- [21] A. J. Sathiamoorthy, K. Weerakoon, T. Guan, J. Liang, and D. Manocha, "Terrapn: Unstructured terrain navigation using online self-supervised learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 7197–7204.
- [22] G. G. Waibel, T. Löw, M. Nass, D. Howard, T. Bandyopadhyay, and P. V. K. Borges, "How rough is the path? terrain traversability estimation for local and global path planning," *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [23] H. Karnan, E. Yang, D. Farkash, G. Warnell, J. Biswas, and P. Stone, "STERLING: Self-supervised terrain representation learning from unconstrained robot experience," in *Conference on Robot Learning (CoRL)*, 2023.
- [24] E. Chen, C. Ho, M. Maulimov, C. Wang, and S. Scherer, "Learning-on-the-drive: Self-supervised adaptation of visual offroad traversability models," *arXiv preprint arXiv:2306.15226*, 2023.
- [25] Z. Zhu, N. Li, R. Sun, D. Xu, and H. Zhao, "Off-road autonomous vehicles traversability analysis and trajectory planning based on deep inverse reinforcement learning," in *IEEE Intelligent Vehicles Symposium (IV)*, 2020, pp. 971–977.
- [26] K. Weerakoon, A. J. Sathiamoorthy, U. Patel, and D. Manocha, "Terp: Reliable planning in uneven outdoor environments using deep reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 9447–9453.
- [27] X. Cai, M. Everett, L. Sharma, P. R. Osteen, and J. P. How, "Probabilistic traversability model for risk-aware motion planning in off-road environments," *arXiv preprint arXiv:2210.00153*, 2022.
- [28] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, "Meta-learning in neural networks: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 9, pp. 5149–5169, 2021.
- [29] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International Conference on Learning Representations (ICLR)*, 2017, pp. 1126–1135.
- [30] D. Li, Y. Yang, Y.-Z. Song, and T. Hospedales, "Learning to generalize: Meta-learning for domain generalization," in *AAAI Conference on Artificial Intelligence (AAAI)*, vol. 32, no. 1, 2018.
- [31] M. Wortsman, K. Ehsani, M. Rastegari, A. Farhadi, and R. Mottaghi, "Learning to learn how to learn: Self-adaptive visual navigation using meta-learning," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [32] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," in *International Conference on Learning Representations (ICLR)*, 2019.
- [33] M. Visca, R. Powell, Y. Gao, and S. Fallah, "Deep meta-learning energy-aware path planner for unmanned ground vehicles in unknown terrains," *IEEE Access*, vol. 10, pp. 30055–30068, 2022.
- [34] M. A. Bekhti, Y. Kobayashi, and K. Matsumura, "Terrain traversability analysis using multi-sensor data correlation by a mobile robot," in *IEEE/SICE International Symposium on System Integration*, 2014, pp. 615–620.
- [35] 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.
- [36] T. Kim, J. Mun, J. Seo, B. Kim, and S. Hong, "Bridging active exploration and uncertainty-aware deployment using probabilistic ensemble neural network dynamics," *Robotics: Science and Systems (RSS)*, 2023.
- [37] Y. Liu, X. San Liang, and R. H. Weisberg, "Rectification of the bias in the wavelet power spectrum," *Journal of Atmospheric and Oceanic Technology*, vol. 24, no. 12, pp. 2093–2102, 2007.
- [38] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 12697–12705.
- [39] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 652–660.
- [40] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2015, pp. 234–241.
- [41] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, "Fourier features let networks learn high frequency functions in low dimensional domains," *Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 7537–7547, 2020.
- [42] P. Fankhauser, M. Bloesch, C. Gehring, M. Hutter, and R. Siegwart, "Robot-centric elevation mapping with uncertainty estimates," in *Mobile Service Robotics*, 2014, pp. 433–440.
- [43] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, "Elevation mapping for locomotion and navigation using gpu," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 2273–2280.