

PIETRA: Physics-Informed Evidential Learning for Traversing Out-of-Distribution Terrain

Cai, Xiaoyi; Queeney, James; Xu, Tong; Datar, Aniket; Pan, Chenhui; Miller, Max; Flather, Ashton; Osteen, Philip R.; Roy, Nicholas; Xiao, Xuesu; How, Jonathan P.

TR2025-022 March 01, 2025

Abstract

Self-supervised learning is a powerful approach for developing traversability models for off-road navigation, but these models often struggle with inputs unseen during training. Existing methods utilize techniques like evidential deep learning to quantify model uncertainty, helping to identify and avoid out-of-distribution terrain. However, always avoiding out-of-distribution terrain can be overly conservative, e.g., when novel terrain can be effectively analyzed using a physics-based model. To overcome this challenge, we introduce Physics-Informed Evidential Traversability (PIETRA), a self-supervised learning framework that integrates physics priors directly into the mathematical formulation of evidential neural networks and introduces physics knowledge implicitly through an uncertainty-aware, physics-informed training loss. Our evidential network seamlessly transitions between learned and physics-based predictions for out-of-distribution inputs. Additionally, the physics-informed loss regularizes the learned model, ensuring better alignment with the physics model. Extensive simulations and hardware experiments demonstrate that PIETRA improves both learning accuracy and navigation performance in environments with significant distribution shifts.

IEEE Robotics and Automation Letters (RA-L) 2025

PIETRA: Physics-Informed Evidential Learning for Traversing Out-of-Distribution Terrain

Xiaoyi Cai¹, James Queeney^{1,2}, Tong Xu³, Aniket Datar³, Chenhui Pan³,
Max Miller¹, Ashton Flather¹, Philip R. Osteen⁴, Nicholas Roy¹, Xuesu Xiao³, and Jonathan P. How¹

Abstract—Self-supervised learning is a powerful approach for developing traversability models for off-road navigation, but these models often struggle with inputs unseen during training. Existing methods utilize techniques like evidential deep learning to quantify model uncertainty, helping to identify and avoid out-of-distribution terrain. However, always avoiding out-of-distribution terrain can be overly conservative, e.g., when novel terrain can be effectively analyzed using a physics-based model. To overcome this challenge, we introduce Physics-Informed Evidential Traversability (PIETRA), a self-supervised learning framework that integrates physics priors directly into the mathematical formulation of evidential neural networks and introduces physics knowledge implicitly through an uncertainty-aware, physics-informed training loss. Our evidential network seamlessly transitions between learned and physics-based predictions for out-of-distribution inputs. Additionally, the physics-informed loss regularizes the learned model, ensuring better alignment with the physics model. Extensive simulations and hardware experiments demonstrate that PIETRA improves both learning accuracy and navigation performance in environments with significant distribution shifts.

I. INTRODUCTION

Recent advancements in perception and mobility have accelerated the deployment of autonomous robots in challenging real-world environments such as office spaces, construction sites, forests, deserts and Mars [1]–[6], where both geometric and semantic comprehension of the terrain is crucial for reliable navigation. In these settings, self-supervised traversability learning has emerged as a powerful tool to train neural networks (NNs) to predict terrain models from navigation data without manual labeling [7]–[10], where the learned representations can work directly with model-based motion planners, providing better interpretability and flexibility compared to fully learned navigation policies. However, the lack of abundant and diverse training data limits the reliability of learned traversability models in novel environments. This is a well-known issue in the learning literature that arises from out-of-distribution (OOD) inputs at test time due to the distribution shift between training and test data [11], [12].

This research was partially sponsored by ARL grant W911NF-21-2-0150 and by ONR grant N00014-18-1-2832. JQ was exclusively supported by MERL. Supplementary Video: <https://youtu.be/OTnNZ96oJRK>

¹Massachusetts Institute of Technology, Cambridge, MA 02139, USA. {xyc, queeney, maxm25, ashton_f, nickroy, jhow}@mit.edu.

²Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA. queeney@merl.com.

³George Mason University, Fairfax, VA 22030, USA. {txu25, adatar, cpan7, xiao}@gmu.edu.

⁴DEVCOM Army Research Laboratory, Adelphi, MD 20783, USA. philip.r.osteen.civ@army.mil. Distribution Statement A: Approved for public release; distribution is unlimited.




Fig. 1: Real-world navigation scenario where the robot, trained on flat ground and turf, encounters unseen tall vegetation and ramps while navigating to a goal. Unlike prior works that avoid OOD terrain, this work successfully navigates across OOD terrain to reach the goal by integrating physics knowledge into the traversability model.

Many recent works mitigate the risk of encountering OOD scenarios by quantifying the *epistemic uncertainty*, which is the model uncertainty due to distribution shift [12]. For example, OOD terrain can be detected by learning a density estimator for the training data to identify test input with low density [13], [14], or a terrain auto-encoder for detecting poorly reconstructed terrain [8], [15]. While avoiding OOD terrain has been shown to improve mission success rate in our prior work [10], doing so can be too conservative, such as the situation considered in this work (see Fig. 1). To this end, we propose **Physics-Informed Evidential TRAversability** (PIETRA), a self-supervised learning framework that seamlessly combines learning-based and physics-based traversability analysis methods such that the downstream planner relies on the learned model for in-distribution (ID) terrain and the physics-based model for OOD terrain. Improving upon our prior work [10], we exploit the mathematical structure of evidential learning to embed a physics-based prior that automatically gets invoked when epistemic uncertainty is high. Moreover, we propose an uncertainty-aware physics-informed loss function inspired by existing works to regularize the learned model to improve generalization [16], [17].

In summary, the contributions of this work are threefold:

- 1) An evidential traversability learning framework with an explicit physics-based prior that gets invoked when encountering OOD features at test time.
- 2) An uncertainty-aware physics-inspired loss function that implicitly injects physics knowledge at training time to improve learning accuracy for both ID and OOD features.
- 3) Extensive simulation and hardware experiments showing that our approach improves both the learning accuracy and the downstream navigation performance in test environments with significant distribution shift.

II. RELATED WORK

The field of traversability analysis studies how to infer suitability of terrain for navigation (see survey [18]). Compared to hand-crafting planning costs based on terrain features, directly learning traversability models from data requires less manual labeling and results in a more accurate assessment of vehicle-terrain interaction. Based on navigation data, visual and/or geometric features of the visited terrain can be used to train a traversability predictor based on estimated traversability values [8], [19]–[22]. Building upon this basic idea, visual and visual-inertial representation learning [9], [23], self-training with pseudo-labels for unvisited terrain [24], temporal fusion of robot states and sensor measurements [7], and data augmentation via vision foundation models [25] can all improve learning accuracy. Alternatively, when a hand-crafted traversability model such as [26] is available, it can be used to provide supervisory signals for training NNs that are faster [2], [3]. One concern is that, because learning-based methods rely on likely limited real-world data, the learned models may not generalize to environments unseen during training. While our work also uses self-supervised learning to obtain a traversability model, we incorporate physics knowledge into the model to improve generalization to OOD environments.

OOD detection is well studied and closely related to uncertainty quantification (see surveys [11], [12]). At a high level, input features that are not well-represented in the training data can lead to high *epistemic uncertainty*, which can be estimated with techniques such as Bayesian dropout [27], model ensembles [28], and evidential methods [29]. Epistemic uncertainty can also be estimated via a terrain auto-encoder for detecting high reconstruction error [15], a density estimator fit to the training data distribution [14], or Gaussian Process regression [30]. Our prior work [10] adopts the evidential method proposed by [31] to efficiently identify OOD terrain and shows that avoiding OOD terrain during planning can improve mission success rate. However, always avoiding OOD terrain can be too conservative. To address this limitation, this new work exploits the mathematical formulation of evidential learning to explicitly embed a physics prior that is activated when encountering OOD features. While informative priors have recently been combined with evidential learning in other problem settings, such as the use of a rule-based prior for trajectory prediction in autonomous driving [32], our work focuses on off-road navigation and additionally introduces a physics-informed loss to further improve OOD generalization.

Incorporating physics and expert knowledge for navigating challenging terrain is crucial for both performance and safety, which can be achieved *explicitly* or *implicitly* (see survey [33]). For example, explicit safety constraints based on terrain geometry and robot states can be imposed during planning [23], [34]–[36]. Physics laws can be explicitly incorporated into NNs via differentiable physics engines or neuro-symbolic methods [37], [38]. In addition, custom models can be directly used as priors in an evidential framework [32]. In contrast to previously mentioned explicit approaches, physics knowledge can also be infused into NN models implicitly by learning to reduce prediction errors with respect to both the training data

and the physics model [16], [17]. This work described herein uses both explicit and implicit methods to infuse physics knowledge into the learned traversability model. As will be shown in Sec. V, compared to NNs trained with a physics-based loss that suffer from distribution shifts in far-OOD regime, our method gracefully falls back to the explicit physics prior. Furthermore, compared to an evidential network that only has a physics prior, our method uses a physics-based loss to further improve generalization.

III. PROBLEM FORMULATION

We consider the problem of motion planning over uneven terrain for a ground vehicle. This section introduces the robot model with uncertain traversability parameters caused by rough terrain and sensing uncertainty in Sec. III-A and the risk-aware planning formulation in Sec. III-B. Compared to our prior work [10] that only considers the linear and angular traction, this new work additionally accounts for the roll and pitch angles important for navigating uneven terrain.

A. Dynamical Models with Traversability Parameters

Consider the discrete time system:

$$\mathbf{x}_{t+1} = F(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\psi}_t), \quad (1)$$

where $\mathbf{x}_t \in \mathbf{X} \subseteq \mathbb{R}^n$ is the state vector such as the position and heading of the ground robot, $\mathbf{u}_t \in \mathbb{R}^m$ is the control input, and $\boldsymbol{\psi}_t \in \boldsymbol{\Psi} \subseteq \mathbb{R}^r$ is the parameter vector that captures traversability of the terrain. In this work, we focus on the bicycle model which is applicable for Ackermann-steering robots used in our simulation and hardware experiments:

$$\begin{bmatrix} p_{t+1}^x \\ p_{t+1}^y \\ \theta_{t+1} \end{bmatrix} = \begin{bmatrix} p_t^x \\ p_t^y \\ \theta_t \end{bmatrix} + \Delta \cdot \begin{bmatrix} \psi_{1,t} \cdot v_t \cdot \cos(\theta_t) \\ \psi_{1,t} \cdot v_t \cdot \sin(\theta_t) \\ \psi_{2,t} \cdot v_t \cdot \tan(\delta_t)/L \end{bmatrix}, \quad (2)$$

where $\mathbf{x}_t = [p_t^x, p_t^y, \theta_t]^\top$ contains the X, Y positions and yaw angle; $\mathbf{u}_t = [v_t, \delta_t]^\top$ contains the commanded speed and steering angle; $0 \leq \psi_{1,t}, \psi_{2,t} \leq 1$ are the linear and angular traction; $\Delta > 0$ is the time interval; and $L > 0$ is the wheelbase. We additionally consider the absolute roll and pitch angles of the robot $\psi_{3,t}, \psi_{4,t} \geq 0$ that do not appear in (2). Therefore, the traversability parameter vector is $\boldsymbol{\psi}_t = [\psi_{1,t}, \psi_{2,t}, \psi_{3,t}, \psi_{4,t}]^\top$. Intuitively, traction captures the “slip” or the ratio between achieved and commanded velocities, which is important for fast navigation, and the roll and pitch values are important for rollover prevention.

For rough terrain with vegetation, the traversability values are often unknown but can be empirically learned. Additionally, due to the noisy nature of the empirical data, we model traversability $\boldsymbol{\psi}_t$ as random variables and mitigate the risk of encountering poor traversability during planning.

B. CVaR-Based Risk-Aware Navigation

Given the initial state \mathbf{x}_0 and maximum roll and pitch angles $\psi_3^{\max}, \psi_4^{\max} \geq 0$, we want to find a control sequence $\mathbf{u}_{0:T-1}$ that minimizes the time to reach the goal using the objective proposed in [10]. We use Conditional Value at Risk (CVaR, as visualized in Fig. 2) to quantify the risk of obtaining low




Fig. 2: Conditional Value at Risk (CVaR) is the expected value of the worst-case $\alpha \in (0, 1]$ portion of total probability and Value at Risk (VaR) is the worst-case α -quantile for some random variable Z . We use both the left-tail and right-tail definitions proposed in [10].

traction and large roll and pitch angles. We achieve risk-aware planning by simulating the state trajectory using the left-tail CVaR_α^- of traction and imposing the maximum attitude constraints over the right-tail CVaR_α^+ of roll and pitch angles:

$$\min_{\mathbf{u}_{0:T-1}} C(\mathbf{x}_{0:T}) \quad (3)$$

$$\text{s.t. } \mathbf{x}_{t+1} = F(\mathbf{x}_t, \mathbf{u}_t, \bar{\psi}_t) \quad (4)$$

$$\bar{\psi}_{3,t} \leq \psi_3^{\max}, \quad \bar{\psi}_{4,t} \leq \psi_4^{\max} \quad (5)$$

$$\bar{\psi}_t = \begin{bmatrix} \text{CVaR}_\alpha^-(\psi_{1,t}) \\ \text{CVaR}_\alpha^-(\psi_{2,t}) \\ \text{CVaR}_\alpha^+(\psi_{3,t}) \\ \text{CVaR}_\alpha^+(\psi_{4,t}) \end{bmatrix}, \quad \begin{bmatrix} \psi_{1,t} \\ \psi_{2,t} \\ \psi_{3,t} \\ \psi_{4,t} \end{bmatrix} \sim p(\mathbf{o}_t) \quad (6)$$

$$\mathbf{o}_t \text{ is the terrain feature at } \mathbf{x}_t \quad (7)$$

$$\forall t \in \{0, \dots, T-1\}, \quad (8)$$

where $\bar{\psi}_t$ contains the worst-case expected traversability values, $\alpha \in (0, 1]$ is the risk tolerance, and $p(\mathbf{o}_t)$ is the traversability distribution after observing the terrain feature at state \mathbf{x}_t . We use Model Predictive Path Integral control (MPPI [39]) to solve (3–8) because MPPI is gradient-free and parallelizable on GPU. Note that a similar formulation of (3–8) has been shown by [10] to outperform methods that assume no slip and state-of-the-art methods such as [21], [40], but this new work introduces additional constraints over roll and pitch.

IV. PHYSICS-INFORMED EVIDENTIAL LEARNING

In this section, we present the proposed method PIETRA shown in Fig. 3. At a high level, the traversability predictor p_ϕ parameterized by ϕ outputs a categorical distribution over discretized traversability values to capture aleatoric uncertainty (inherent, irreducible data uncertainty) for each traversability parameter. Moreover, the normalizing flow network [41] p_λ parameterized by λ estimates the densities of latent features as proxies for epistemic uncertainty (model uncertainty due to distribution shift). PIETRA improves upon our prior work EVORA [10] by having an explicit physics prior that is invoked when encountering OOD inputs (Sec. IV-A) and an uncertainty-aware physics-informed loss that implicitly injects physics knowledge to further improve model accuracy in OOD terrain (Sec. IV-B). Lastly, we design custom physics priors used in our experiments in Sec. IV-C and discuss the implementation details in Sec. IV-D.

A. Dirichlet Distribution with Physics Prior

The Dirichlet distribution $q = \text{Dir}(\beta)$ with concentration parameters $\beta = [\beta_1, \dots, \beta_B]^T \in \mathbb{R}_{>0}^B$ is a hierarchical

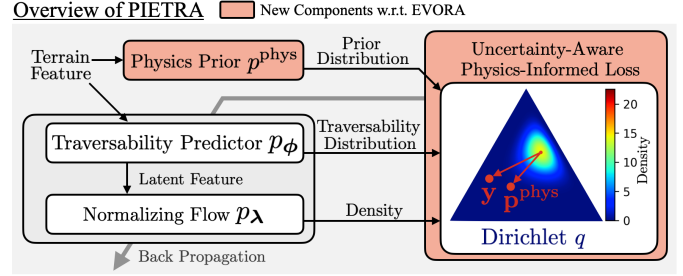


Fig. 3: Overview of the proposed physics-informed evidential learning framework. In contrast to our prior work EVORA [10], this work explicitly embeds a physics prior in the Dirichlet posterior update and implicitly infuses physics knowledge via an uncertainty-aware, physics-informed loss during training.

distribution over categorical distributions $\text{Cat}(\mathbf{p})$, where $\mathbf{p} \in \mathbb{R}_{>0}^B$ is a normalized probability mass function (PMF) over $B > 0$ discretized traversability values, i.e., $\sum_{b=1}^B p_b = 1$. The parameters \mathbf{p} of the lower-level categorical distribution $\text{Cat}(\mathbf{p})$ are sampled from the higher-level Dirichlet distribution, i.e., $\mathbf{p} \sim \text{Dir}(\beta)$. The mean (also called the *expected PMF*) of the Dirichlet distribution is given by $\mathbb{E}_{\mathbf{p} \sim q}[\mathbf{p}] = \beta / \sum_{b=1}^B \beta_b$. The sum $n = \sum_{b=1}^B \beta_b$ reflects how concentrated the Dirichlet distribution is around its mean and corresponds to the “total evidence” of a data point observed during training.

Given an input feature \mathbf{o} and a physics model p^{phys} that maps its input to a traversability PMF with the evidence $n^{\text{phys}} > 0$, the NN performs an input-dependent posterior update:

$$\beta_{\phi, \lambda}^{\mathbf{o}} = n^{\text{phys}} p^{\text{phys}}(\mathbf{o}) + n_{\lambda}^{\mathbf{o}} p_{\phi}(\mathbf{o}), \quad (9)$$

$$n_{\lambda}^{\mathbf{o}} = N p_{\lambda}(\mathbf{z}_{\mathbf{o}}), \quad (10)$$

where the posterior Dirichlet distribution $q_{\phi, \lambda}^{\mathbf{o}} = \text{Dir}(\beta_{\phi, \lambda}^{\mathbf{o}})$ depends on the physics prior and its evidence, the predicted traversability PMF $p_{\phi}(\mathbf{o})$ and the predicted evidence $n_{\lambda}^{\mathbf{o}}$ that is proportional to the density $p_{\lambda}(\mathbf{z}_{\mathbf{o}})$ for the latent feature $\mathbf{z}_{\mathbf{o}}$ weighted by a fixed certainty budget $N > 0$. The posterior Dirichlet distribution leads to the expected traversability PMF:

$$\mathbf{p}_{\phi, \lambda}^{\mathbf{o}} = \frac{n^{\text{phys}} p^{\text{phys}}(\mathbf{o}) + n_{\lambda}^{\mathbf{o}} p_{\phi}(\mathbf{o})}{n^{\text{phys}} + n_{\lambda}^{\mathbf{o}}}, \quad (11)$$

which is used by the risk-aware planner (3–8). The prior evidence is set to a small number (e.g., $n^{\text{phys}} = B$) such that the predicted evidence $n_{\lambda}^{\mathbf{o}}$ is much larger than n^{phys} for ID features. As a result, the learned traversability PMF $p_{\phi}(\mathbf{o})$ is used for planning on ID terrain. However, the new evidence $n_{\lambda}^{\mathbf{o}}$ diminishes if the features are OOD, so the physics model $p^{\text{phys}}(\mathbf{o})$ is used for planning instead. The graceful transition between learned and physics-based traversability estimates eliminates the need to avoid OOD terrain, as is required in our prior work EVORA [10] that uses an uninformative prior.

B. Uncertainty-Aware Physics-Informed Loss

In addition to explicitly embedding the physics prior in the posterior update (11), physics knowledge can also be implicitly infused into the NN by adding a physics-based loss term during training. In contrast to existing methods [16], [17] that are not uncertainty-aware, we adapt the physics-informed loss to the context of evidential learning.

We design the training loss based on the squared Earth Mover’s distance (EMD² [42]) which is a better measure of




Fig. 4: Terrain features used by the physics priors and how multiple potential traversability values are fused, as discussed in Sec. IV-C. (a) Terrain slope. (b) Terrain height. (c) Inter-wheel distances of the front wheel pair and back wheel pair. (d) Inter-wheel distances of the left wheel pair and right wheel pair. (e) Multiple candidates are fused into a single estimated traversability PMF and combined with the uniform distribution to account for uncertainty.

error compared to KL divergence that treats prediction errors independently across discrete bins. EMD^2 has a closed-form expression after dropping the constant multiplicative factor:

$$\text{EMD}^2(\mathbf{p}, \mathbf{y}) := \|\text{cs}(\mathbf{p}) - \text{cs}(\mathbf{y})\|^2, \quad (12)$$

where $\text{cs} : \mathbb{R}^B \rightarrow \mathbb{R}^B$ is the cumulative sum operator, and \mathbf{p} and \mathbf{y} are the predicted and target PMFs. To train an evidential NN, our prior work proposes the uncertainty-aware EMD^2 (UEMD^2) loss, which is the expected EMD^2 under the predicted Dirichlet q (see the closed form in [10, Thm. 1]):

$$L^{\text{UEMD}^2}(q, \mathbf{y}) := \mathbb{E}_{\mathbf{p} \sim q} [\text{EMD}^2(\mathbf{p}, \mathbf{y})]. \quad (13)$$

This uncertainty-aware loss has been shown to improve both learning accuracy and navigation performance compared to using (12) or cross-entropy-based losses for training.

To implicitly infuse physics knowledge, we propose the uncertainty-aware physics-informed (UPI) loss, defined as the expectation of the weighted sum of data loss and physics loss given the predicted Dirichlet distribution:

$$\begin{aligned} L^{\text{UPI}}(q, \mathbf{y}) &:= \mathbb{E}_{\mathbf{p} \sim q} [\text{EMD}^2(\mathbf{p}, \mathbf{y}) + \kappa \cdot \text{EMD}^2(\mathbf{p}, \mathbf{p}^{\text{phys}})] \\ &= \text{UEMD}^2(q, \mathbf{y}) + \kappa \cdot \text{UEMD}^2(q, \mathbf{p}^{\text{phys}}), \end{aligned} \quad (14)$$

where $\mathbf{p}^{\text{phys}} = p^{\text{phys}}(\mathbf{o})$ is the PMF predicted by the physics model given some feature \mathbf{o} , and $\kappa \geq 0$ is a hyperparameter. Intuitively, the physics-based term ensures that NN predictions stay close to the physics prior, thus improving generalization in near-OOD regimes. As our ablation study in Sec. V-C later shows, both the UPI loss (14) and the physics prior (11) are important for achieving the best learning performance.

C. Uncertainty-Aware Physics Prior

We now present our custom physics priors for uneven terrain with two semantic types: the rigid *dirt* terrain and the soft *vegetation* terrain. Note that alternative prior designs could be used provided they are based on physics principles that hold in both ID and OOD scenarios.

The traction priors are inferred from the terrain properties under the wheels (also called the “footprint”) assuming the robot has 0° roll and pitch. Over dirt terrain, as the tire traction depends on friction that reduces with greater slopes,

we design a simple slope-based model that predicts lower traction for more sloped terrain. In our hardware experiment, taller vegetation deforms less and slows down the robot more, so we propose a traction model that predicts lower traction for taller vegetation. To factor in the uncertainty due to outcomes from the four wheels, our physics-based traction models for the traversability parameter $\psi' \in \{\psi_1, \psi_2\}$ are

$$\mathbf{p}_{\text{dirt}}^{\psi'} = \frac{1}{4} \sum_{i=1}^4 \mathbb{1}_{\text{dirt}}^{\psi'} \left(\text{clip} \left(\frac{s^{\text{max}} - s_i}{s^{\text{max}}}, 0, 1 \right) \right), \quad (15)$$

$$\mathbf{p}_{\text{veg}}^{\psi'} = \frac{1}{4} \sum_{i=1}^4 \mathbb{1}_{\text{veg}}^{\psi'} \left(\text{clip} \left(\frac{h^{\text{max}} - h_i}{h^{\text{max}}}, 0, 1 \right) \right), \quad (16)$$

where s_i, h_i are the terrain slope and the height of vegetation under the wheel i , as shown in Fig. 4(a-b); the clip function restricts the traction values between 0 and 1; and the operator $\mathbb{1}_s^{\psi'}$ maps a scalar to a PMF where the estimated traversability value for ψ' has probability 1 for the semantic type $s \in \mathcal{S} := \{\text{dirt}, \text{veg}\}$. Intuitively, the priors merge the estimated traction outcomes from the four wheels into a common PMF. Note that the terrain slope is estimated in the direction of the robot’s heading, and we use the absolute slope to produce low traction values for both uphill and downhill terrain.

The roll and pitch angles are estimated based on the terrain height under the wheels using trigonometry. The physics model for the traversability parameter $\psi' \in \{\psi_3, \psi_4\}$ and the semantic type $s \in \mathcal{S}$ is

$$\mathbf{p}_s^{\psi'} = \frac{1}{2} \sum_{(i,j) \in \mathcal{W}^{\psi'}} \mathbb{1}_s^{\psi'} \left(\arctan \left(\frac{h_i - h_j}{d_{i,j}} \right) \right), \quad (17)$$

where $d_{i,j}$ is the distance between the wheels i, j . Roll and pitch priors use different wheel index pairs defined in $\mathcal{W}^{\psi_3} = \{(1, 4), (2, 3)\}$ and $\mathcal{W}^{\psi_4} = \{(1, 2), (4, 3)\}$ to account for uncertain terrain contact points as shown in Fig. 4(c-d).

To handle different semantic types within the footprint feature \mathbf{o} , the proposed physics prior for the traversability parameter $\psi' \in \{\psi_1, \dots, \psi_4\}$ combines the prior predictions based on different semantic types via:

$$p^{\text{phys}, \psi'}(\mathbf{o}) = w^{\text{unif}} \mathbf{p}^{\text{unif}} + (1 - w^{\text{unif}}) \sum_{s \in \mathcal{S}} r_s \mathbf{p}_s^{\psi'}, \quad (18)$$

where r_s is the ratio of the semantic types within the footprint, $\mathbf{p}_s^{\psi'}$ is the physics prior based on (15–17), \mathbf{p}^{unif} is a uniform PMF, and $w^{\text{unif}} \in [0, 1]$ determines how much uniform PMF is incorporated into the final physics prior to account for inaccurate prior predictions, as shown in Fig. 4(e). Note that (18) is used as the physics prior in (11).

D. Implementation Details

The traversability predictor p_ϕ and the normalizing flow p_λ are trained jointly using the proposed UPI loss (14) on an empirically collected dataset $\{(\mathbf{o}, \psi)_k\}_{k=1}^K$ where $K > 0$ and every estimated traversability parameter in ψ is converted to a target PMF \mathbf{y} using one-hot encoding with $B = 12$ discrete bins. Terrain features can be extracted from a semantic elevation map (e.g., [43]), while traversability values can be computed using the pose and velocity estimates obtained from onboard odometry, along with the commanded velocities.

For simplicity, we use a multi-layer perceptron (MLP) as the shared encoder to process the flattened and concatenated semantic and elevation patches of the terrain. The shared encoder is followed by separate MLP decoders with soft-max outputs for predictions of each traversability parameter. To reduce computation, we train a single flow network using the encoder’s outputs. As recommended by [31], we scale the constant certainty budget N exponentially with the dimension of the encoder’s outputs for numerical purposes. As the accuracy of physics priors differs across traversability parameters, we use a standalone MLP to map the shared encoder’s output to scalars between 0 and 1 to downscale the predicted evidence in the posterior update (11) for each traversability parameter.

In contrast to our prior work [10], this work uses yaw-aligned features because roll, pitch and traction are directional over uneven terrain. At deployment time, feature patches are obtained in a sliding-window fashion with a fixed number of yaw angles. The predicted traversability distributions are then converted to CVaR values for look-up during planning.

V. SIMULATION RESULTS

We use the Chrono Engine [44] to simulate navigation over rocky terrain and collect benchmark data. When compared against several baselines with or without physics knowledge, PIETRA achieves the best prediction accuracy (Sec. V-B). Furthermore, we conduct an ablation study to validate the proposed improvements (Sec. V-C). When used for navigation, PIETRA leads to the best success rate and time to goal (Sec. V-D). While only the dirt terrain is simulated, the real-world experiments have both dirt and vegetation (Sec. VI).

A. Simulation Setup

An overview of the simulation environment is shown in Fig. 5. We generate 30 distinct synthetic elevation maps based on the real-world rock testbed proposed in [45], which are split evenly for training, validation and testing. Every map has sides of 50 m and we use an Ackermann-steering robot of dimension 3.4 m by 1.8 m to collect 20 navigation trials in each map. For simplicity, the robot executes sinusoidal steering and 2 m/s speed commands, and the episode ends if the robot rolls over, gets stuck, or exits the map. Note that terrain features and traversability values are obtained directly from the simulator.

To emphasize the testing of OOD generalization, the test elevation maps are scaled more than the training and validation maps, inducing significant distribution shift as visualized in Fig. 5, where we use the standard deviations of the elevation features to measure terrain unevenness. To clearly mark the ID/OOD boundary, we consider a terrain feature as ID (or OOD) if its unevenness falls below (or above) the 50th percentile of the training dataset. We only use the ID data for training and validation, but the accuracy of trained NNs is evaluated on the entire test dataset. Note that the test maps are also used for evaluating navigation performance.

B. Learning Benchmark

The proposed method PIETRA is compared to several state-of-the-art methods, including our prior work EVORA [10], an

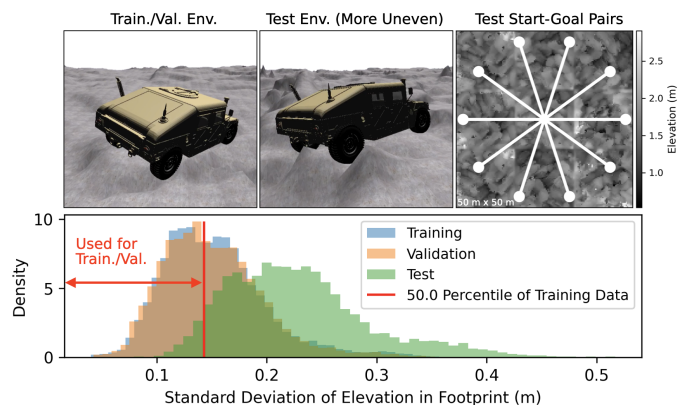


Fig. 5: Overview of the simulation setup with synthetic rocky terrain, where the terrain unevenness is measured by the standard deviation of elevation in the robot footprint. The test maps are more uneven than the training and validation ones to induce distribution shift. During training and validation, only the data with terrain unevenness below the 50th percentile of the training dataset is used. Note that test maps are also used to evaluate navigation performance.

encoder-decoder NN trained with the EMD² loss (Vanilla), and an encoder-decoder NN trained with the physics-informed loss (PI) adapted from [16] to use the EMD² loss. Based on the training data and vehicle characteristics, we tune the max slopes (15) for linear and angular traction to correspond to 30° and 15° with $w^{\text{unif}} = 0.2$ for mixing in the uniform PMF (18). Hyperparameter sweeps are conducted over learning rates in $\{1e-3, 1e-4, 1e-5\}$ for the Adam optimizer, and weights for the physics loss term in $\{0.1, 0.5, 1.0\}$. To encourage smoothness, we penalize Dirichlet entropy [31] with weights in $\{1e-3, 1e-4, 1e-5\}$ for evidential NNs. We identify the best parameter set for each method based on the average validation EMD² errors over 5 training seeds and report the mean and standard deviations of the prediction errors over the seeds.

The learning benchmark results in Table I show that PIETRA achieves the best overall, ID and OOD performance. To gain more intuition, the test errors are binned based on terrain unevenness in Fig. 6. Notably, while PI outperforms Vanilla and performs similarly to PIETRA on ID features, the improvement thanks to the physics-informed loss degrades as features become more OOD. On the other hand, PIETRA and EVORA fall back to their corresponding priors (physics prior and uniform prior) due to decreasing predicted evidence (11).

C. Ablation Study

An ablation study for the proposed improvements with respect to EVORA is summarized in Table II. For completeness, we include a variant where EVORA uses the physics model

TABLE I: Learning benchmark results.

Method	Prediction Error (EMD ² ↓)		
	Overall	ID	OOD
PIETRA (Proposed)	2.77 ± 0.01	2.68 ± 0.01	2.78 ± 0.01
Physics-Informed (PI)	2.99 ± 0.01	2.68 ± 0.01	3.02 ± 0.01
Vanilla	3.05 ± 0.01	2.74 ± 0.01	3.08 ± 0.01
EVORA	3.30 ± 0.02	2.72 ± 0.01	3.35 ± 0.02
Physics Prior	3.09	2.93	3.10
Uniform Prior	4.35	4.34	4.35




Fig. 6: Prediction errors binned by the standard deviation of elevation in robot footprint. While all trained models perform similarly on ID data, PIETRA generalizes the best to OOD data. Unlike PIETRA that falls back to the physics prior, EVORA falls back to the uniform prior which has higher error. Without explicit priors, PI and Vanilla perform poorly in far-OOD regime.

for OOD features with latent densities lower than the densities observed during training. The takeaway is that both the physics prior and physics-informed loss are important for achieving the best overall accuracy. While the UPI loss (14) alone leads to the best ID accuracy, it offers limited improvement on OOD accuracy. Interestingly, EVORA equipped with a physics prior achieves better OOD accuracy than OOD-based switching, verifying the benefit of the explicitly embedded physics prior.

TABLE II: Ablation study on infusing physics knowledge into EVORA [10].

Changes w.r.t. EVORA	Prediction Error (EMD ² ↓)		
	Overall	ID	OOD
PP & UPI (Proposed)	2.77 ± 0.01	2.68 ± 0.01	2.78 ± 0.01
UPI	3.27 ± 0.02	2.67 ± 0.01	3.32 ± 0.02
PP	2.86 ± 0.01	2.73 ± 0.01	2.87 ± 0.01
Use phys. model if OOD	2.93 ± 0.01	2.72 ± 0.01	2.95 ± 0.01
EVORA	3.30 ± 0.02	2.72 ± 0.01	3.35 ± 0.02

PP: Physics Prior UPI: Uncertainty-Aware Physics-Informed Loss (14)

D. Navigation Benchmark

We deploy the NNs from Sec. V-B trained with the best hyperparameters in the test environments. As discussed in [10], EVORA detects OOD terrain that is avoided during planning via auxiliary costs. As an ablation of the physics prior, we also consider PIETRA with OOD avoidance. We consider 10 test maps and 10 start-goal pairs that are 40 m apart as visualized in Fig. 5. The robot aims to reach the goal without exceeding 30° of roll and pitch angles or getting immobilized. The robot has a maximum steering angle of 30° and a maximum speed of 1 m/s to ensure stable crawling. Note that MPPI runs at 10 Hz on GPU with a 5 s planning horizon and 1024 rollouts. We consider 3 risk tolerances α in $\{0.4, 0.6, 0.8\}$ and report the success rate and the time to goal of successful trials.

The navigation results in Fig. 7 show that PIETRA achieves the best success rate and time to goal. In comparison, PIETRA with OOD avoidance is more conservative but still outperforms EVORA thanks to the proposed UPI loss (14). While PI outperforms Vanilla, it performs worse than the physics prior due to poor OOD generalization.

VI. HARDWARE EXPERIMENTS

While the simulation results in Sec. V have shown improved learning and navigation performance achieved by PIETRA




Fig. 7: Success rate and time to goal from the simulated navigation benchmark. The figure shows the mean and standard deviation values for each method. The proposed PIETRA performs the best because it seamlessly transitions between the learned and physics models on OOD terrain. When PIETRA avoids OOD terrain, it becomes more conservative but still outperforms EVORA thanks to the physics-informed training loss. While PI outperforms Vanilla, it performs worse than the physics prior due to unreliable traversability predictions.

with respect to the state-of-the-art, this section presents further evidence of PIETRA’s practical utility via real-world experiments on a computationally constrained platform in the presence of multiple semantic types in the environment.

A. Data Collection, Training and Deployment

The indoor 9.6 m by 8 m arena contains turf and fake bushes to mimic outdoor vegetation. The *dirt* semantic type includes the concrete floor and a skateboard ramp. A 0.33 m by 0.25 m RC car carries a RealSense D455 depth camera, and a computer with an Intel Core i7 CPU and Nvidia RTX 2060 GPU. The robot runs onboard traversability prediction, motion planning, and elevation mapping with 0.1 m resolution, but Vicon is used for estimating the pose and velocity of the robot. For simplicity, terrain elevation and semantics are estimated separately based on depth and color measurements. Each map cell contains a single semantic type initially set to dirt and is later updated to vegetation if the corresponding pixels in the image space are green.

Traversability models are trained on 10 min of manual driving data over *flat* concrete floor and turf, making the tall bushes and skateboard ramp OOD at test time. For the vegetation and dirt traction priors, we empirically tune the maximum dirt slope s^{\max} (15) to correspond to 30° and the maximum vegetation height h^{\max} (16) to be 0.2 m which is slightly greater than the 0.15 m wheel diameter. To prevent damaging the hardware, we do not consider PI and Vanilla because their predictions for OOD terrain are unreliable. Therefore, we only consider PIETRA, the physics prior, and EVORA that avoids OOD terrain. All models are trained with the learning rate of $1e-4$, entropy weight of $1e-4$ and physics loss weight of 0.1 when applicable. We set a fixed risk tolerance of $\alpha = 0.5$ for the risk-aware planner. Note that MPPI runs at 10 Hz on GPU with a 5 s planning horizon and 1024 rollouts, and the generation of CVaR maps runs at 5 Hz with the map dimension of 81×96 with 13 yaw discretizations.

B. Navigation Results

The robot is tasked to navigate to a fixed goal without leaving the bounding box of the arena or violating the roll and

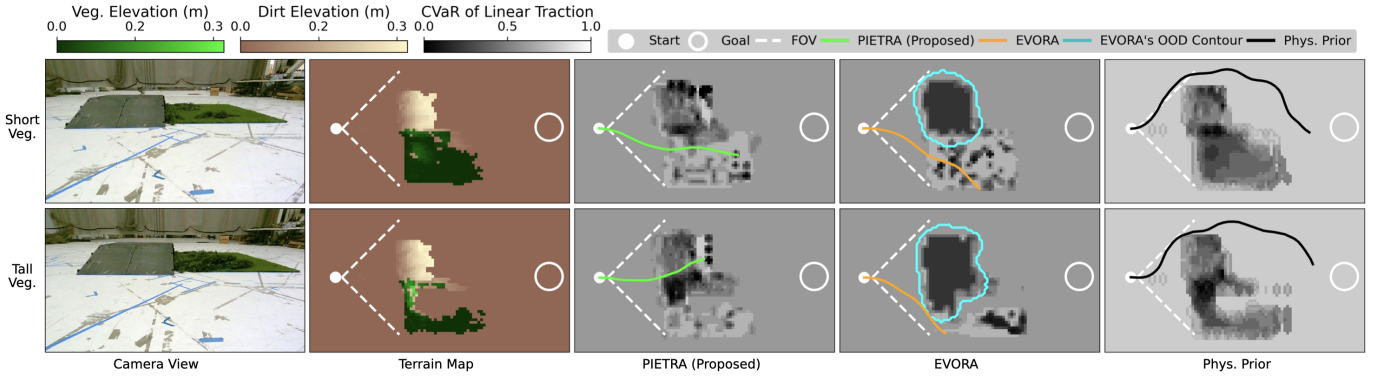


Fig. 8: Snapshots of the start of the two missions with short and tall vegetation. From the onboard camera view, the robot builds semantic elevation map online, which is processed to generate CVaR of the traversability values. Due to limited space, we only show the CVaR of the predicted linear traction aligned with the robot heading for each method and overlay the initial best trajectories. For EVORA, we additionally show the OOD contour that the robot avoids.

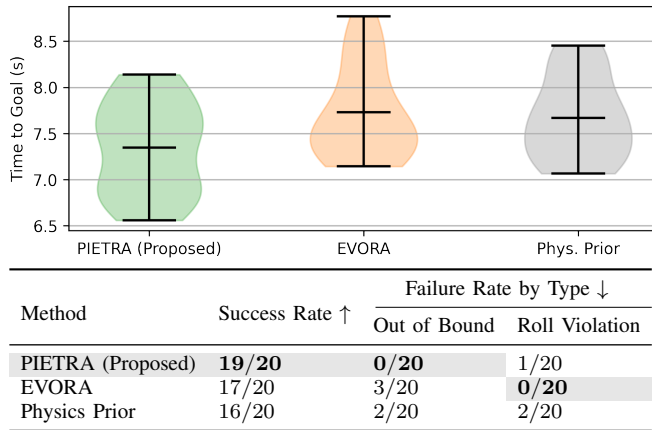


Fig. 9: Hardware experiment results showing that PIETRA achieves the best time to goal and success rate. To avoid OOD terrain, EVORA takes wide turns that increase the time to goal and the odds of going out of bound. While the physics prior achieves better time to goal than EVORA, it suffers from more failures due to inaccurate traction estimates on the flat terrain.

pitch constraints of 30° and 45° respectively, while deciding whether to drive on vegetation, climb over the ramp, or stay on the concrete floor. We vary the elevation and density of bushes and repeat each method 20 times. Some snapshots of the start of the missions are visualized in Fig. 8 for scenarios with short and tall vegetation. The success rates and the time to goal of the successful trials are reported in Fig. 9, showing that PIETRA achieves the best performance. PIETRA violates the roll constraint once because the robot drives off the down-ramp prematurely. EVORA has 3 out-of-bound failures because the robot frequently takes wide turns to avoid OOD terrain, which makes the robot more likely to miss the goal and leave the bounding box. The physics prior leads to 2 out-of-bound failures because it misses the goal region slightly due to over-optimistic traction estimates. The physics prior also leads to 2 roll violations due to driving off the down-ramp prematurely.

Representative trials in Fig. 10 show that PIETRA goes over vegetation when the bushes are short, but it goes over the ramp when the bushes are tall. In comparison, the large OOD regions force EVORA to take wide turns. Interestingly, the robot using the physics prior favors the concrete floor. To explain this phenomenon, Fig. 8 shows the predicted CVaR of linear traction for each method, suggesting that the physics prior is over-confident about the traction of the concrete floor.




Fig. 10: Full trajectories, terrain maps and trial time of representative trials. PIETRA chooses between vegetation and the ramp based on vegetation height. While the physics prior and PIETRA share similar traversability assessment for tall vegetation and the ramp, the physics prior overestimates the terrain traction of the concrete floor, making the robot frequently go around the ramp and vegetation. EVORA detects and avoids the OOD terrain unseen during training, causing undesirable wide turns due to limited steering authority.

This further shows that PIETRA combines the best of EVORA and physics prior in that it relies on the learned model in ID terrain, but falls back to the physics model in OOD terrain.

VII. CONCLUSION AND FUTURE WORK

We presented PIETRA, a self-supervised traversability learning method that incorporates physics knowledge explicitly via the custom physics prior and implicitly via the physics-informed training loss. Extensive simulations and hardware experiments show that PIETRA improves learning accuracy and navigation performance under significant distribution shifts.

One of the limitations of PIETRA is its large memory footprint due to discretization, so improvements are needed for higher-dimensional systems. Moreover, PIETRA relies on accurate terrain features and state estimation, but the risk due to uncertain sensing and estimation may need to be addressed for some real-world environments. Lastly, PIETRA can be further improved by leveraging data augmentation [17], differentiable physics simulators [37] and neuro-symbolic networks [38].

REFERENCES

- [1] M. F. Ginting, S.-K. Kim, D. D. Fan, M. Palieri, M. J. Kochenderfer, and A.-a. Agha-Mohammadi, "SEEK: Semantic Reasoning for Object Goal Navigation in Real World Inspection Tasks," *arXiv:2405.09822*, 2024.
- [2] S. Triest, D. D. Fan, S. Scherer, and A.-A. Agha-Mohammadi, "UN-RealNet: Learning Uncertainty-Aware Navigation Features from High-Fidelity Scans of Real Environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024, pp. 12 627–12 634.
- [3] J. Frey, S. Khattak, M. Patel, D. Atha, J. Nubert, C. Padgett, M. Hutter, and P. Spieler, "RoadRunner–Learning Traversability Estimation for Autonomous Off-road Driving," *arXiv:2402.19341*, 2024.
- [4] R. Chung, G. Georgakis, P. Spieler, C. Padgett, A. Agha, and S. Khattak, "Pixel to Elevation: Learning to Predict Elevation Maps at Long Range using Images for Autonomous Offroad Navigation," *IEEE Robot. Autom. Lett.*, 2024.
- [5] X. Meng, N. Hatch, A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Q. Chen, S. Deng, G. Okopal, D. Fox, B. Boots, and A. Shaban, "TerrainNet: Visual Modeling of Complex Terrain for High-speed, Off-road Navigation," in *Proc. Robot. Sci. Syst.*, 2023.
- [6] R. Takemura and G. Ishigami, "Uncertainty-Aware Trajectory Planning: Using Uncertainty Quantification and Propagation in Traversability Prediction of Planetary Rovers," *IEEE Robot. Autom. Mag.*, 2024.
- [7] M. V. Gasparino, A. N. Sivakumar, and G. Chowdhary, "WayFASTER: A Self-Supervised Traversability Prediction for Increased Navigation Awareness," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024, pp. 8486–8492.
- [8] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, "Fast Traversability Estimation for Wild Visual Navigation," in *Proc. Robot. Sci. Syst.*, 2023.
- [9] J. Seo, S. Sim, and I. Shim, "Learning Off-Road Terrain Traversability With Self-Supervisions Only," *IEEE Robot. Autom. Lett.*, vol. 8, no. 8, pp. 4617–4624, 2023.
- [10] X. Cai, S. Ancha, L. Sharma, P. R. Osteen, B. Bucher, S. Phillips, J. Wang, M. Everett, N. Roy, and J. P. How, "EVORA: Deep Evidential Traversability Learning for Risk-Aware Off-Road Autonomy," *IEEE Trans. Robot.*, vol. 40, pp. 3756–3777, 2024.
- [11] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized Out-of-Distribution Detection: A Survey," *arXiv:2110.11334*, 2021.
- [12] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, M. Shahzad, W. Yang, R. Bamler, and X. X. Zhu, "A Survey of Uncertainty in Deep Neural Networks," *Artif. Intell. Rev.*, vol. 56, no. 1, pp. 1513–1589, 2023.
- [13] S. Ancha, P. R. Osteen, and N. Roy, "Deep Evidential Uncertainty Estimation for Semantic Segmentation under Out-Of-Distribution Obstacles," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024.
- [14] 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," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 11 297–11 304.
- [15] R. Schmid, D. Atha, F. Schöller, S. Dey, S. Fakoorian, K. Otsu, B. Ridge, M. Bjelonic, L. Wellhausen, M. Hutter, *et al.*, "Self-Supervised Traversability Prediction by Learning to Reconstruct Safe Terrain," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 12 419–12 425.
- [16] A. Saviolo, G. Li, and G. Loianno, "Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 10 256–10 263, 2022.
- [17] P. Maheshwari, W. Wang, S. Triest, M. Sivaprakasam, S. Aich, J. G. Rogers III, J. M. Gregory, and S. Scherer, "PIAug–Physics Informed Augmentation for Learning Vehicle Dynamics for Off-Road Navigation," *arXiv:2311.00815*, 2023.
- [18] P. Papadakis, "Terrain Traversability Analysis Methods for Unmanned Ground Vehicles: A Survey," *Eng. Appl. Artif. Intell.*, vol. 26, no. 4, pp. 1373–1385, 2013.
- [19] A. Datar, C. Pan, M. Nazeri, A. Pokhrel, and X. Xiao, "Terrain-Attentive Learning for Efficient 6-DoF Kinodynamic Modeling on Vertically Challenging Terrain," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024.
- [20] A. Datar, C. Pan, and X. Xiao, "Learning to Model and Plan for Wheeled Mobility on Vertically Challenging Terrain," *arXiv:2306.11611*, 2023.
- [21] X. Cai, M. Everett, J. Fink, and J. P. How, "Risk-Aware Off-Road Navigation via a Learned Speed Distribution Map," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 2931–2937.
- [22] K. S. Sikand, S. Rabiee, A. Uccello, X. Xiao, G. Warnell, and J. Biswas, "Visual representation learning for preference-aware path planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 11 303–11 309.
- [23] A. Pokhrel, A. Datar, M. Nazeri, and X. Xiao, "CAHSOR: Competence-Aware High-Speed Off-Road Ground Navigation in SE (3)," *arXiv:2402.07065*, 2024.
- [24] I. Cho and W. Chung, "Learning Self-Supervised Traversability With Navigation Experiences of Mobile Robots: A Risk-Aware Self-Training Approach," *IEEE Robot. Autom. Lett.*, 2024.
- [25] S. Jung, J. Lee, X. Meng, B. Boots, and A. Lambert, "V-STRONG: Visual Self-Supervised Traversability Learning for Off-road Navigation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024, pp. 1766–1773.
- [26] D. D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick, and A. akbar Agha-mohammadi, "STEP: Stochastic Traversability Evaluation and Planning for Risk-Aware Off-road Navigation," in *Proc. Robot. Sci. Syst.*, 2021.
- [27] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning," in *Proc. Int. Conf. Mach. Learn.*, vol. 48, 2016, pp. 1050–1059.
- [28] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy, "Deep Exploration via Bootstrapped DQN," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016.
- [29] D. T. Ulmer, C. Hardmeier, and J. Frellsen, "Prior and Posterior Networks: A Survey on Evidential Deep Learning Methods For Uncertainty Estimation," *Trans. Mach. Learn. Res.*, 2023.
- [30] M. Endo, T. Tani, R. Yonetani, and G. Ishigami, "Risk-Aware Path Planning via Probabilistic Fusion of Traversability Prediction for Planetary Rovers on Heterogeneous Terrains," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2023, pp. 11 852–11 858.
- [31] B. Charpentier, O. Borchert, D. Zügner, S. Geisler, and S. Günnemann, "Natural Posterior Network: Deep Bayesian Predictive Uncertainty for Exponential Family Distributions," in *Proc. Int. Conf. Learn. Representations*, 2022.
- [32] J. Patrikar, S. Veer, A. Sharma, M. Pavone, and S. Scherer, "RuleFuser: Injecting Rules in Evidential Networks for Robust Out-of-Distribution Trajectory Prediction," *arXiv:2405.11139*, 2024.
- [33] Z. Hao, S. Liu, Y. Zhang, C. Ying, Y. Feng, H. Su, and J. Zhu, "Physics-Informed Machine Learning: A Survey on Problems, Methods and Applications," *arXiv:2211.08064*, 2022.
- [34] T. Han, A. Liu, A. Li, A. Spitzer, G. Shi, and B. Boots, "Model Predictive Control for Aggressive Driving Over Uneven Terrain," *arXiv:2311.12284*, 2023.
- [35] L. Sharma, M. Everett, D. Lee, X. Cai, P. Osteen, and J. P. How, "RAMP: A Risk-Aware Mapping and Planning Pipeline for Fast Off-Road Ground Robot Navigation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2023, pp. 5730–5736.
- [36] S. Talia, M. Schmittle, A. Lambert, A. Spitzer, C. Mavrogiannis, and S. S. Srinivasa, "Demonstrating HOUND: A Low-cost Research Platform for High-speed Off-road Underactuated Nonholonomic Driving," in *Proc. Robot. Sci. Syst.*, 2024.
- [37] R. Agishev, K. Zimmermann, M. Pecka, and T. Svoboda, "MonoForce: Self-Supervised Learning of Physics-Aware Grey-Box Model for Predicting the Robot-Terrain Interaction," *arXiv:2309.09007*, 2023.
- [38] Z. Zhao, B. Li, Y. Du, T. Fu, and C. Wang, "PhysORD: A Neuro-Symbolic Approach for Physics-infused Motion Prediction in Off-road Driving," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024.
- [39] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Reh, B. Boots, and E. A. Theodorou, "Information Theoretic MPC for Model-Based Reinforcement Learning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2017, pp. 1714–1721.
- [40] Z. Wang, O. So, K. Lee, and E. A. Theodorou, "Adaptive Risk Sensitive Model Predictive Control with Stochastic Search," in *Proc. Learn. Dyn. Control Conf.*, 2021, pp. 510–522.
- [41] I. Kobyzev, S. J. Prince, and M. A. Brubaker, "Normalizing Flows: An Introduction and Review of Current Methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 3964–3979, 2021.
- [42] L. Hou, C.-P. Yu, and D. Samaras, "Squared Earth Mover's Distance-Based Loss for Training Deep Neural Networks," *arXiv:1611.05916*, 2016.
- [43] G. Erni, J. Frey, T. Miki, M. Mattamala, and M. Hutter, "MEM: Multi-Modal Elevation Mapping for Robotics and Learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 11 011–11 018.
- [44] A. Tasora, R. Serban, H. Mazhar, A. Pazouki, D. Melanz, J. Fleischmann, M. Taylor, H. Sugiyama, and D. Negrut, "Chrono: An Open Source Multi-physics Dynamics Engine," in *Proc. Springer High Perf. Comput. Sci. Eng.*, 2016, pp. 19–49.
- [45] A. Datar, C. Pan, M. Nazeri, and X. Xiao, "Toward Wheeled Mobility on Vertically Challenging Terrain: Platforms, Datasets, and Algorithms," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2024, pp. 16 322–16 329.