# Learning Smooth State-Dependent Traversability from Dense Point Clouds

Zihao Dong<sup>1</sup>, Alan Papalia<sup>1</sup>, Leonard Jung<sup>1</sup>, Alenna Spiro<sup>1</sup>, Philip R. Osteen<sup>2\*</sup>, Christa S. Robison<sup>2\*</sup>, Michael Everett<sup>1</sup>

<sup>2</sup>DEVCOM Army Research Laboratory (ARL) <sup>1</sup>Northeastern University <sup>\*</sup>Equal Contribution

Abstract—A key open challenge in off-road autonomy is that the traversability of terrain often depends on the vehicle's state. In particular, some obstacles are only traversable from some orientations. However, learning this interaction by encoding the angle of approach as a model input demands a large and diverse training dataset and is computationally inefficient during planning due to repeated model inference. To address these challenges, we present SPARTA, a method for estimating approach angle conditioned traversability from point clouds. Specifically, we impose geometric structure into our network by outputting a smooth analytical function over the 1-Sphere that predicts risk distribution for any angle of approach with minimal overhead and can be reused for subsequent queries. The function is composed of Fourier basis functions, which has important advantages for generalization due to their periodic nature and smoothness. We demonstrate SPARTA both in a highfidelity simulation platform, where our model achieves a 91% success rate crossing a 40m boulder field (compared to 73% for the baseline), and on hardware, illustrating the generalization ability of the model to real-world settings.

#### I. INTRODUCTION

Many applications for autonomous mobile robots involve traversing rough, off-road environments [1]-[3]. One key to reliably navigating in these environments is the ability to understand the interaction between the robot and the terrain, to assess relevant risk metrics (e.g., potential vehicle damage). Despite recent success in identifying traversable terrain [4]-[7], motion planning in geometrically complex environments (e.g., dense forests, boulder fields) remains a challenge. For example, the risk associated with traversing a given terrain is often a complicated function of the robot's state, especially its speed and angle of approach, and the terrain geometry<sup>1</sup>. While high-speed traversal through risky terrain can be explicitly penalized within the planner, it is less clear how to account for the angle of approach (e.g., using hand-crafted risk heuristics based on terrain elevation [8]-[10]), as this requires both detailed terrain information and nuanced consideration of the underlying geometry.

In contrast, learning-based approaches have promise in being able to flexibly model the complex relationship between terrain and traversal risk. In particular, point cloud-based approaches are suitable for the problem setting, because point clouds are capable of capturing fine geometric details of the environment that would be missed in other representations like an elevation map. As shown in Fig. 1, in the point cloud, the





Fig. 1. Dense point cloud (background) captures geometric detail of the environment (Top right), while discretizing dense point cloud into elevation map results in information loss. The underlying terrain is clearly distinguishable from noise (leaves, sensor noise, etc.) in the point cloud, but doing so is hard in an elevation map, leading to wrong risk estimation (Viridis Circle where Green denotes low risk).

terrain (ground) is distinguishable from noise (e.g. leaves). As a result, a point cloud-based method would prefer the safe approach angle (Red) over the risky one (Magenta) due to a relatively smooth elevation change. However, noise and terrain cannot be easily separated in the elevation map despite its high resolution (5cm). The terrain appears smoother from the risky angle (Magenta) thus an elevation map-based method would (incorrectly) select it instead.

In this paper, we propose a technique to learn risk variables of interest, based on a point cloud representation of a given terrain patch. Specifically, to model system noise implicitly captured in the dataset, we represent risk as a categorical distribution due to its ease of use and capability of approximating multimodal distributions, as demonstrated in [4]. However, a point cloud-based network that naively encodes the angle of approach would be hard to train in practice due to the sparsity of data in robotics applications. As a result, our model takes in the point cloud and predicts an analytical function (represented using Fourier basis functions) that maps angle of approach to the target risk variable distribution. The choice of Fourier basis ensures that the resulting function is periodic, thus avoiding the wrap-around discontinuity [11], and has important advantages for generalization in the training process due to its upperbounded Lipschitz constant by construction. The function, once computed for a terrain patch, can be queried efficiently with minimal computation overhead for any angle of approach at any timestep when integrated with modern planners like MPPI [12], [13]. To summarize, in this paper we propose

SPARTA (Smooth Point-cloud Approach-angle Reasoning for Terrain Assessment) with the following key contributions:

- A novel learning-based approach leveraging Fourier basis functions, which estimates angle-of-approach dependent traversal risk for complex terrains represented as point clouds.
- Theoretical analysis of the advantage of imposing smoothness by using Fourier basis functions, with connection to Lipschitz continuity and model generalization capability.
- Evaluations of the proposed pipeline through learning to predict potential vehicle damage, a risk variable prior works handle using handcrafted rules, with extensive results both in a high-fidelity simulator [14] and on hardware, with integration into a MPPI planner [12].

## II. RELATED WORK

a) Traversability Estimation: Traversability estimation algorithms aim to identify terrain suitable for robot navigation. Traversability estimates typically rely on either manually designed features [8], [10] or are derived from learned vehicle movement patterns [4]-[7], [15]. Both [8] and [10] use local elevation measurements to estimate stepping difficulty using heuristics such as step-wise elevation change. On the other hand, learned methods usually rely on multiple sensing modalities, especially elevation maps and semantic maps [4]-[7], [15]. WVN [5] uses camera input and trains a traversable terrain classifier online leveraging pretrained features from DINO-ViT [16]. Meanwhile, [6], [7] fuse semantic and elevation measurements, and use separate decoder heads to predict dense semantic and elevation maps, which are used to score planner rollouts using various heuristics. [4], [15] use semantic and elevation maps to predict traversability distributions represented as linear/angular traction and incorporate the worstcase traction into the vehicle dynamics model for risk-aware planning. However, none of these approaches models the angle-dependent nature of traversability. As a result, they can be overly conservative in geometrically complex environments, while underestimate traversal risk when the terrain is only dangerous from a few orientations.

b) Learning Harmonic Functions: Geometric deep learning has found its place in numerous machine learning applications. For example, FBM [17] uses Fourier basis functions to composite long-time series data and achieved state-of-the-art performance on multiple benchmarks. NeuRBF [18] replaces traditional grid-based structures with radial basis functions for constructing implicit neural representation of scenes, allowing the model to better adapt to the underlying signal structures. In robotics, OrbitGrasp [19] leverages spherical harmonics functions to represent a grasping quality function. Off-road autonomy, especially traversability estimation, is a setting with rich geometric structure, and as a result we propose to leverage Fourier basis functions to exploit this intrinsic prior.

c) Placement of This Work: Unlike previous traversability estimation algorithms, SPARTA models the angle dependency of risk variables, which existing algorithms failed to capture, by exploiting the geometric structure of angle of approach (the 1-Sphere). Because we impose few assumptions on the risk variable of interest (i.e., it has clear lower and upper bounds and does not change abruptly with a small change in angle of approach), our approach could be used to extend (most) existing traversability estimation methods for more accurate and granular prediction. Besides tire deformation or potential vehicle damage demonstrated in our experiments (Section IV), we show a few other variables in off-road autonomy with a similar geometric structure in Section G.

### III. METHOD

In this section we introduce SPARTA, a method for predicting the risk associated with traversing a terrain patch at a given angle based on prior data traversing a wide variety of terrains (represented as dense point cloud). One challenge in traversability estimation is handling aleatoric uncertainty, the inherent uncertainty in a system which cannot be reduced with more data [20]. SPARTA captures this uncertainty by treating risk as a random variable – a categorical distribution with *B* discrete bins – whose concentration parameters are predicted by our model (Section III-A and Section III-B). The tail of the distribution is used to estimate risk for risk-aware planning (Section III-D).

The primary novelty in SPARTA is in representing this risk distribution. SPARTA predicts an analytical function for each bin, and each function smoothly maps any angle of approach to the corresponding risk concentration parameter (Section III-B). We characterize the smoothness of these risk functions (Section III-C), which we hypothesize to be useful for generalization across approach angles (empirically supported in Section IV-B). Once computed, the function can be re-used for subsequent queries for the same terrain patch, thus avoiding repeated neural network inference during planning. We include an experiment demonstrating the query efficiency of our model in Section C. An overview of the risk estimation pipeline and its application to planning is shown in Fig. 2.

## A. Modeling Aleatoric Uncertainty with Categorical Distributions

In this work, we are interested in predicting a risk variable conditioned on the angle of approach. We assume we have a dataset recording the risk variable while traversing various terrains. However, the aleatoric uncertainty of this data will likely prevent us from reliably estimating risk as a single scalar value. To develop robust real-world systems, it is important to properly account for aleatoric uncertainty [4], [15]. SPARTA captures aleatoric uncertainty by modeling risk as a random variable whose distribution reflects this uncertainty. Without assuming a specific parametric form, we approximate this distribution with a categorical distribution [4].

More formally, given a terrain patch represented as a point cloud  $\mathbf{Q} \in \mathbb{R}^{m \times 3}$  where m > 0 is the number of points in the point cloud and the angle of approach  $\phi \in [0, 2\pi]$ , we are interested in obtaining a function  $\Gamma: (\mathbf{Q}, \phi) \mapsto \mathbf{p}(\gamma_{\mathbf{Q}, \phi})$ ,



Fig. 2. Overall pipeline of the proposed algorithm. For all terrain patches we are interested in, if its Fourier coefficients are not in the database, we compute them and store in the database. Otherwise we retrieve the Fourier coefficients from the database and construct the analytical functions, which are queried to compute the risk variable distribution, whose tail distribution is considered for risk aware planning.

where  $\gamma_{\mathbf{Q},\phi} \in \mathbb{R}$  denotes a risk related variable<sup>2</sup> and  $\mathbf{p}(\gamma_{\mathbf{Q},\phi})$ is the probability mass function (PMF) of a categorical distribution with *B* bins.<sup>3</sup> The categorical distribution  $\mathbf{p}(\gamma_{\phi})$ is parameterized by non-negative concentration parameters  $\bar{\gamma}_{\phi} = [\gamma_{\phi}^1, \ldots, \gamma_{\phi}^B] \in \mathbb{R}^B_{\geq 0}$  and computed by normalizing  $\bar{\gamma}_{\phi}$ to sum to 1:

$$\mathbf{p}(\gamma_{\phi}) = \frac{\bar{\gamma}_{\phi}}{\sum_{b=1}^{B} \gamma_{\phi}^{b}}.$$
(1)

To obtain this risk distribution  $\mathbf{p}(\gamma_{\phi})$ , we can train a neural network on a dataset of traversals on various terrains,  $\{(\mathbf{Q}_k, \mathbf{y}_k, \phi_k)\}_{k=1}^K$ , where  $\mathbf{y}_k$  denotes the ground truth (empirical) risk PMF. The network is trained to minimize the squared Earth Mover's Distance (EMD<sup>2</sup>):

$$L^{\text{EMD}^2}(\mathbf{p}(\phi_k), \mathbf{y}_k) = ||\mathbf{cs}(\mathbf{p}(\phi_k)) - \mathbf{cs}(\mathbf{y}_k)||_2^2, \qquad (2)$$

where  $cs(\cdot)$  is the cumulative sum operator. However, training such a network for risk-aware planning poses two major difficulties. The first issue is due to the structure of angles of approach  $\phi$ , i.e., there is a periodicity of  $2\pi$  and we expect that small changes in  $\phi$  should lead to small changes in the risk distribution. While a neural network can approximate this structure, it typically requires a large and diverse dataset to do so. The second issue is the model's query efficiency, which becomes a major concern for real-time planning. In subsequent sections, we discuss how SPARTA addresses these issues by leveraging Fourier basis functions to construct a smooth and periodic analytical function.

## B. Leveraging $S^1$ as a Continuous Representation of Approach Angle

Here we introduce how SPARTA represents the risk concentration parameters  $\gamma_{\phi}$  as a function of the angle of approach  $\phi$  via Fourier basis functions. This representation stems from two main insights: (i) the set of angles  $[0, 2\pi]$  can be mapped to the 1-sphere  $S^1$  (i.e., the unit circle that lies in the terrain ground plane and is centered at the intersection of the terrain normal vector and the plane) and (ii) Fourier basis function is a fundamental way of representing arbitrary functions over  $S^1$ .



Fig. 3. (a): Naive model (orange) suffers from wrap-around discontinuity, whereas our model (blue) predicts continuous signal. (b): The phenomenon propagates to the CVaR of the predicted distribution.

Given a representation of angles on  $S^1$ , we can consider the risk concentration parameters  $\gamma_{\phi}$  as functions over  $S^1$ . It follows from harmonic analysis and representation theory [21] that any square-integrable real-valued function over the 1-sphere  $S^1$  can be uniquely decomposed into a (possibly infinite) linear combination of the Fourier basis functions  $\{1\} \cup \{\cos(k\phi), \sin(k\phi) : k \in [1, n]\}$ . This suggests these basis functions as the natural choice for representing arbitrary functions on the 1-sphere. Altogether this allows us to compute risk concentration parameters at angle  $\phi$  by taking a linear combination of the Fourier basis functions and applying the sigmoid function<sup>4</sup>  $\sigma(\cdot)$  to ensure non-negativity (so the result is a valid concentration parameter):

$$\gamma_{\phi}^{i} = \sigma(\frac{a_{i}^{0}}{2} + \sum_{k=1}^{n} [a_{i}^{k} \cos(k\phi) + b_{i}^{k} \sin(k\phi)]).$$
(3)

Concretely, our neural network  $\Gamma_{\theta}$  takes an input point cloud  $\mathbf{Q}$  and, for each bin  $i \in [1, B]$ , predicts the corresponding coefficients  $\mathbf{F}^i \triangleq [a_i^0, a_i^1, b_i^1, \cdots a_i^n, b_i^n] \in \mathbb{R}^{2n+1}$ . Eq. (3) is then used to compute the concentration parameters  $\bar{\gamma}_{\phi}$  during planning. Eq. (3) is fast to compute as it requires a single (batched) dot product and a pass through the sigmoid function, and its runtime scales linearly with the maximum frequency n, a small number in practice<sup>5</sup>. We note that the choice of Fourier basis function naturally avoids wrap-around discontinuities at  $\phi = 2\pi$ . We demonstrate this in Fig. 3a, where we compare the

<sup>&</sup>lt;sup>2</sup>For example, risk related variables may include vehicle damage, terrain traction, or rollover risk.

<sup>&</sup>lt;sup>3</sup>For conciseness, we omit the subscript  $\mathbf{Q}$  in the rest of this paper.

 $<sup>^{4}</sup>$ While we used sigmoid function, other non-negative functions could theoretically be used (e.g., softplus).

<sup>&</sup>lt;sup>5</sup>we found n = 3 to suffice in our experiments. We include model runtime comparison in Section C.



Fig. 4. Demonstration of model in overfitting (red) and smoothinterpolation regime (blue).

predictions of a baseline multilayer perceptron model (orange) that encodes the angle of approach  $\phi$  directly as a model input, and our model (blue). It can be seen that the baseline model (orange) is discontinuous at the wrap-around point, resulting in a large jump in the predicted concentration parameters. This discontinuity is known to be harmful to model training when working with angle-related variables [11]. In contrast, our model (blue) naturally respects the periodicity and generates a globally smooth signal over  $S^1$ . Besides avoiding the wraparound discontinuity, the Fourier basis functions allow better generalization when trained with limited data, as we detail through analyzing its Lipschitz constant in the next subsection.

### C. Imposed Smoothness Improves Prediction Generalization

In off-road settings we may expect that small changes in approach angle  $\phi$  induce relatively small changes in the risk variable  $\gamma_{\phi}$ . In this subsection we show that  $\gamma_{\phi}$  is smooth with respect to  $\phi$  and provide a simple upper bound on the first derivative (i.e., an upper bound on the Lipschitz constant). We connect this Lipschitz constant analysis to results in learning theory [22], [23] to motivate the hypothesis that our representation improves generalization (see Section IV-B for empirical evidence).

Formally, if  $L^i$  denotes the Lipschitz constant of the mapping from approach angle to risk concentration parameter  $\phi \to \gamma_{\phi}^i$ , then adding a small perturbation  $\delta$  to  $\phi$  bounds the change in  $\gamma_{\phi}^i$ :  $|\gamma_{\phi+\delta}^i - \gamma_{\phi}^i| \leq L^i |\delta|$ . Our representation of the concentration parameters  $\gamma_{\phi}^i$  guarantees that each  $\gamma_{\phi}^i$  is smooth with respect to  $\phi$ , as each  $\gamma_{\phi}^i$  is a linear combination of sines and cosines. Furthermore, the Fourier coefficients  $\mathbf{F}^i$  provide an upper bound on the Lipschitz constant  $L^i$ .

**Theorem 1.** The concentration parameter  $\gamma_{\phi}^{i}$  defined by the Fourier series representation in Eq. (3) is Lipschitz continuous with respect to  $\phi$  with an upper bound on the Lipschitz constant  $L^{i}$  given by:

$$L^{i} \leq \frac{1}{4} \sum_{k=1}^{n} k \sqrt{(a_{i}^{k})^{2} + (b_{i}^{k})^{2}}.$$
(4)

Proof: See Section A.

Having a bounded Lipschitz constant is desirable, as small Lipschitz constants have been connected to improved generalization to unseen data in deep learning [22]–[24]. With this in mind, we consider Theorem 1 to suggest that our representation improves the model's ability to generalize to unseen angles of approach  $\phi$ . As SPARTA attempts to estimate

risk based on both point cloud data and the angle of approach  $(\mathbf{Q}, \phi)$ , we argue that an improved generalization over  $\phi$  allows the learning process to focus on extracting meaningful information from the point cloud  $\mathbf{Q}$ , rather than having to simultaneously learn the relevant structure of the angle of approach  $\phi$ . In other words, we believe that the smoothness and bounded Lipschitz constants represent informative geometric priors that allows the model to more efficiently learn relevant features. We provide empirical evidence for this in Section IV-B.

## D. Incorporating the Learned Risk Model into a Motion Planner

Let  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^s$  denote the state of the robot,  $\mathbf{u} \in \mathcal{U} \subseteq$  $\mathbb{R}^{u}$  the control input, and subscript t > 0 the timestep. We assume the robot's motion evolves as a discrete time system,  $\mathbf{x}_{t+1} = F(\mathbf{x}_t, \mathbf{u}_t)$ . The planner optimizes for a control input sequence  $\mathbf{u}_{0:T}$  given an initial state  $\mathbf{x}_0$  that minimizes the worst-case planning objective. Specifically, we break the cost function into two parts: A task-related term  $C(\mathbf{x}_{0:T})$  (e.g., goal reaching cost) and the worst-case risk computed from the risk distribution  $\mathbf{p}(\gamma_{\phi_t})$ . We adopt the Conditional Value at Risk (CVaR) as the worst-case risk because it is coherent and thus suitable for trustworthy risk assessment in robotics [25]. Specifically, the CVaR at level  $\alpha \in [0,1]$  of the risk variable  $\gamma_{\phi}$  is  $\text{CVaR}_{\alpha}(\gamma_{\phi}) := \frac{1}{\alpha} \int_{0}^{\alpha} \min\{\gamma \mid p(\gamma_{\phi} > \gamma) \leq \tau\} d\tau$ , where the probabilities can be estimated from  $\mathbf{p}(\gamma_{\phi})$ . Intuitively,  $\text{CVaR}_{\alpha}(\gamma_{\phi})$  represents the expectation of the last (right tail)  $\alpha$  portion of the distribution  $\mathbf{p}(\gamma_{\phi})$ . Notice that while the categorical distribution is discrete, so its CVaR is not necessarily continuous w.r.t.  $\phi$ , our formulation guarantees the CVaR does not suffer from a wrap-around discontinuity, as shown in Fig. 3b. To summarize, we adopt MPPI [12], [13] with the following planning problem:

$$\min_{\mathbf{u}_{0:T}} \quad C(\mathbf{x}_{0:T}) + \sum_{t=0}^{T} \operatorname{CVaR}_{\alpha}(\gamma_{\phi_{t}})$$
s.t.  $\mathbf{x}_{t+1} = F(\mathbf{x}_{t}, \mathbf{u}_{t})$ 
 $\mathbf{F}_{t+1}^{i} = \Gamma_{\theta}(\mathbf{Q}_{t+1})[i], \ \forall \ i \in [1, B]$ 
 $\gamma_{\phi_{t+1}}^{i} \leftarrow Eq. \ (3), \ \forall \ i \in [1, B]$ 
 $\mathbf{p}_{t+1}(\gamma_{\phi_{t+1}}) \leftarrow Eq. \ (1)$ 
IV. EXPERIMENTS

In this section, we demonstrate the application of the proposed framework to assess potential vehicle damage (e.g., flat tires, broken axle). Although vehicle damage frequently occurs during real off-road autonomous operations, previous research has largely overlooked this risk metric or relied on relatively simple handcrafted heuristics (such as stepwise elevation change) to model it [8]–[10].

## A. Tire Deformation As A Surrogate Metric

In the following experiments, we use tire deformation extent as a surrogate metric for the probability of vehicle damage. While our formulation can handle other risk-related variables



Fig. 5. Our Model (Red and Blue) achieve lower test EMD<sup>2</sup> loss compared to *AngleInput* (Green and Orange) without overfitting.

like wheel force and acceleration, which are also related to terrain contact, we use tire deformation because the tire is the part that directly contacts the terrain and thus is less susceptible to noise accumulated from simulating other vehicle parts. We collect training data using a high-fidelity simulation platform BeamNG.tech [14]. Specifically, for a tire node w, let  $r_w$  denotes the distance from the (deformed) tire node to wheel center,  $r_{\text{inner}}$  and  $r_{\text{outer}}$  the radius of the wheel frame and non-deformed tire, respectively, and we compute the deformation extent as  $d_w = \frac{r_w - r_{\text{inner}}}{r_{\text{outer}} - r_{\text{inner}}}$ . Section D includes further details of the data collection process, and Section G includes examples of other risk-related variables that our formulation can handle.

## B. Model Performance

To demonstrate the generalization capability of our method, we compare its test set performance against a model that encodes the angle of approach as an input (AngleInput), where our model outperforms AngleInput. Our model and the AngleInput share the same architecture and have approximately the same number of trainable parameters, except that the AngleInput has an additional encoder head for the angle of approach, which is fused downstream with the point cloud feature to predict the concentration parameters. Details of the model architecture can be found in Section B. In Fig. 5, we show the test  $EMD^2$  loss achieved by our model and the AngleInput with varying model sizes. The AngleInput baselines (Green and Orange) quickly overfit to the training data, with a best test loss of 0.15. In contrast, our test loss (Red and Blue) continues to decrease and eventually converged to 0.09. This is aligned with our hypothesis in Section III-C, that our model would generalize better compared to AngleInput because the Fourier basis functions promise an upper-bound for the Lipschitz constant of the mapping  $\phi \to \gamma^i_{\phi}$ , whereas the AngleInput has to learn this mapping implicitly.

### C. Planning in Simulation

To test integration with a planner, we designed a challenging task where the vehicle needs to cross a randomly generated boulder field, with 20m radius and 1500 randomly placed obstacles, from 100 uniformly placed starting positions. Fig. 6 shows the overview of the environment. The obstacles are placed so densely that finding a clear path is challenging. Therefore, to complete the task at high speed, the planner must accurately distinguish between obstacles that can be safely



Fig. 6. (a): Overview of the Boulder Field (40m) test environment with goals (blue) and starting points (orange), overlaid with an elevation map (grayscale). (b): Overview of the environment from driver's POV.

Algorithm	$\bar{v}$	S.R. (%) ↑	Vel (m/s) ↑
Ours	4	91	<u>2.6</u>
AngleInput	4	<u>85</u>	2.5
AngleFree	4	73	3.2
	2	79	1.9
Elev	4	73	3.2
	2	75	2.1

TABLE I. Success rate (S.R.) of the models in the boulder field test. Our method achieves best success rate, and overall angle-dependent models perform better than those that do not consider angle. **Best**, and second best.

traversed and those that pose significant risks. In Section E we provide an open loop experiment with ablation study to further demonstrate our approach. Besides *AngleInput*, we consider 2 baselines: (i) a model trained to predict the distributions without modeling the angle dependency (*AngleFree*), and (ii) a heuristic baseline that chooses the point with lowest maximum elevation under vehicle's wheel footprint (*Elev*). Since the obstacles are placed on flat ground, *Elev* is equivalent to evaluating stepping difficulty as in [8], [10]. We use a MPPI [12] planner with Ackermann kinematics, and the planning objective includes the distance to the goal, a velocity cost for driving over speed limit  $\bar{v}$ , and a risk cost:

$$C_{\text{risk}}(t) = \begin{cases} \max(h_{\text{FL}}, h_{\text{FR}}, h_{\text{RL}}, h_{\text{RR}}) & \text{for } Elev\\ \text{CVaR}_{\alpha}(\gamma_{\phi_t}) & \text{otherwise} \end{cases}$$
(6)

$$Cost = \underbrace{w_{\text{goal}} \sum_{t=0}^{T} D(\mathbf{x}, \mathbf{x}_{\text{goal}})}_{\text{Goal Cost}} + \underbrace{w_{\text{v}} \sum_{t=1}^{T} \mathbb{1}(v_t > \bar{v})}_{\text{Velocity Cost}}$$
(7)  
$$+ w_{\text{risk}} \sum_{t=1}^{T} v_t C_{\text{risk}}(t).$$

$$\underbrace{t=0}_{\text{Risk Cost}}$$

where  $w_{\text{goal}}, w_{\text{v}}, w_{\text{risk}}$  are weights for associated cost term. Notice we scale the Risk Cost term by the current vehicle velocity  $v_t$  to penalize traversing risky terrain at high speed.

In Table I we report the success rate (S.R., reached goal without damage) and average vehicle velocity (Vel) for the methods tested. Our model achieves higher success rate compared to *AngleInput*, presumably because of its stronger generalization ability. Most importantly, the performances of models that captures angle-dependency (ours and *AngleInput*) are better than *AngleFree* and *Elev* because their Risk Cost term cannot distinguish between safe obstacles from risky



Fig. 7. (a): Overview of the environment. The left half of the environment is filled with risky obstacles (Red), and the right half is filled with safe ones (Blue). Arrow indicates expected angle of approach. (b): Background denotes CVaR predicted by our method at the expected angle of approach. The planner is able to identify risky obstacles and decelerate preemptively. Max Tire Force: 1426 N. (c): Background denotes CVaR predicted by the *AngleFree*. It can not identify risky obstacles and the vehicle drives at a constant speed. Max Tire Force: 2937 N. (d): Background denotes elevation. Elevation baseline treat them as equally risky and traverse with a constant speed. Max Tire Force: 4367 N.

ones. As a result, the vehicle attempts to traverse the boulder field at a roughly constant velocity, i.e. the Risk Cost in these cases degenerates to velocity cost. To test if the faster speed of *AngleFree* and *Elev* caused their low success rates, we tuned the weights and decreased  $\bar{v}$  from 4 to 2 to enforce a slower velocity. However, their success rates do not significantly increase even traversing at a lower velocity, which suggests that they are ineffective risk estimators in this environment.

To further help understand how *AngleFree* and *Elev* underestimate traversal risk, we crafted an environment, where the vehicle has to go through a narrow passage way filled with obstacles. We picked two obstacles with same elevation yet one appears to be safe and the other is visually risky, and use each of them fill half of the environment. Detail of the environment is shown in Fig. 7a, where red denotes the risky obstacles and blue denotes safe obstacles, when



Fig. 8. Environment overview with overlaid point cloud (blue, colored by elevation), physical obstacle (cyan box), and lethal obstacles (shaded). Our method (Green) consistently chooses the safe path, whereas *Elev* (Orange) can choose the risky path as the obstacles have same elevation change. All trajectories are shown in the white box and background color denotes CVaR of the expected angle of traversal (Left to Right).

approached from the expected orientation (from right to left). The CVaR predicted by our model (Fig. 7b) correctly identifies the risky obstacles from the safe ones, and is manifested by the vehicle's preemptive deceleration. During the traversal, our model induced the lowest force on the tire 1426 N. The CVaR predicted by the AngleFree model (Fig. 7c) however, is very similar for all obstacles because the obstacles are all safe to traverse from certain orientations, and as a result the vehicle drives at a constant (high) velocity through the obstacles. Similar to AngleFree, the obstacles have same height and thus Elev drives the vehicle at a constant speed (until it hits an obstacle very hard and thus slowed down). Both AngleFree and Elev results in significantly higher tire force (2937 N and 4367 N respectively), signifying higher probability of vehicle damage. In other words, in order for AngleFree and Elev to safely drive the vehicle through complex obstacles, their maximum speed must be set conservatively such that the vehicle can safely traverse the most risky obstacle possible. However, empirically it is difficult to find such a threshold that guarantees safety.

#### D. Planning on Hardware

In this experiment, we demonstrate our model on a custombuilt wheeled robot (specification can be found in Section F), and compare its performance qualitatively to *Elev*. We create an environment where a robot has to pick one out of 2 possible paths to reach its goal, where the left path is safe and the right path is risky. For the obstacle, we cut rigid foam into a shape that is smooth on one side yet sharp the opposite way. We place the obstacle on the left path as shown in the cyan box in Fig. 8. To avoid vehicle damage and allow the robot to smoothly traverse, we *virtually* place the same obstacle rotated 180 degrees on the right path as the risky obstacle. The point clouds (pre-collected with [26]) are overlaid in Fig. 8 for demonstration purposes. The planner uses unicycle dynamics and the planning objective is the same as Eq. (8) except we do not include the velocity terms. The planner samples 500 rollouts with a horizon of 5.0s, and runs at 50 Hz with our model (70 Hz with *Elev*). In the white box in Fig. 8 we show the trajectories of our method (Green) and *Elev* (Orange). Our model correctly identifies the safe path for all 5 trials, yet *Elev* chooses the risky path for 3 out of 5 trials. While our model drives the vehicle safely and smoothly through the safe obstacles, *Elev* would have likely caused vehicle damage by ramming into sharp edges of an obstacle (red circle).

## V. CONCLUSION

In this paper, we propose SPARTA, a general framework for estimating terrain traversal risk conditioned on vehicle angle of approach. Specifically, we identify and exploit the geometric structure in the problem setting, and propose to predict a smooth analytical function over the 1-Sphere to achieve better data efficiency and generalization performance. During planning, the function is queried to construct a categorical distribution for a risk variable of interest at any angle of approach, and the distribution's tail risk (CVaR) is considered to account for aleatoric uncertainty. Experiments in simulation and on hardware demonstrate the effectiveness of the proposed approach. Most importantly, off-road problems are naturally rich in geometric structure, and our formulation is general enough such that it finds application in numerous different problem settings. Future work may further explore the geometric structure in other off-road autonomy problems.

#### VI. LIMITATION

Our work has several limitations. For example, in this work we assume the environment is fully observable, i.e. a prebuilt point cloud map is accessible during planning. While the point cloud map we can build on the fly is sufficiently dense, the model's performance without a prebuilt map is still to be tested. A potential solution to this problem is to penalize terrain with limited observation during planning or incorporate depth inpainting as in [6], [7]. In addition, this work focused on estimating traversability given geometric information, whereas semantic information is also important to consider in many off-road settings. We anticipate that future work could incorporate semantic information by performing sensor fusion with point pillar feature in birds-eye-view [7]. Last but not least, tire deformation is a metric that's only available in simulation. However, we are not interested in predicting tire deformation in the real world, but rather using it as a surrogate metric for terrain property (i.e. how risky is the terrain). Future work could explore other variables to quantify vehicle damage.

#### VII. ACKNOWLEDGEMENTS

This research was sponsored by the DEVCOM Army Research Laboratory (ARL) under SARA CRA W911NF-24-2-0017. Distribution Statement A: Approved for public release; distribution is unlimited.

#### REFERENCES

- L. F. Oliveira, A. P. Moreira, and M. F. Silva, "Advances in agriculture robotics: A state-of-the-art review and challenges ahead," *Robotics*, vol. 10, no. 2, p. 52, 2021.
- [2] I. Lluvia, E. Lazkano, and A. Ansuategi, "Active mapping and robot exploration: A survey," *Sensors*, vol. 21, no. 7, p. 2445, 2021.
- [3] G. Wang, W. Wang, P. Ding, *et al.*, "Development of a search and rescue robot system for the underground building environment," *Journal of field robotics*, vol. 40, no. 3, pp. 655–683, 2023.
- [4] X. Cai, S. Ancha, L. Sharma, *et al.*, "Evora: Deep evidential traversability learning for risk-aware off-road autonomy," *IEEE Transactions on Robotics*, 2024.
- [5] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, "Fast traversability estimation for wild visual navigation," *arXiv preprint arXiv*:2305.08510, 2023.
- [6] B. Li, H. Lu, H. Wang, *et al.*, "Terrain-net: A highlyefficient, parameter-free, and easy-to-use deep neural network for ground filtering of uav lidar data in forested environments," *Remote Sensing*, vol. 14, no. 22, p. 5798, 2022.
- [7] J. Frey, M. Patel, D. Atha, et al., "Roadrunner-learning traversability estimation for autonomous off-road driving," *IEEE Transactions on Field Robotics*, 2024.
- [8] D. D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick, and A.-A. Agha-Mohammadi, "Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation," *arXiv preprint arXiv:2103.02828*, 2021.
- [9] A. Datar, C. Pan, and X. Xiao, "Learning to model and plan for wheeled mobility on vertically challenging terrain," *IEEE Robotics and Automation Letters*, 2024.
- [10] D. Lee, I. M. A. Nahrendra, M. Oh, B. Yu, and H. Myung, "Trg-planner: Traversal risk graph-based path planning in unstructured environments for safe and efficient navigation," *IEEE Robotics and Automation Letters*, 2025.
- [11] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, 2019, pp. 5745–5753.
- [12] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in 2016 IEEE international conference on robotics and automation (ICRA), IEEE, 2016, pp. 1433–1440.
- [13] G. Williams, N. Wagener, B. Goldfain, et al., "Information theoretic mpc for model-based reinforcement learning," in 2017 IEEE international conference on robotics and automation (ICRA), IEEE, 2017, pp. 1714–1721.
- [14] BeamNG GmbH, BeamNG.tech, version 0.25.0.0, Jun. 15, 2022. [Online]. Available: https://www. beamng.tech/.

- [15] X. Cai, M. Everett, J. Fink, and J. P. How, "Risk-aware off-road navigation via a learned speed distribution map," in 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2022, pp. 2931–2937.
- [16] M. Caron, H. Touvron, I. Misra, et al., "Emerging properties in self-supervised vision transformers," in *Proceedings of the IEEE/CVF international conference* on computer vision, 2021, pp. 9650–9660.
- [17] R. Yang, L. Cao, J. YANG, *et al.*, "Rethinking fourier transform from a basis functions perspective for longterm time series forecasting," *Advances in Neural Information Processing Systems*, vol. 37, pp. 8515–8540, 2024.
- [18] Z. Chen, Z. Li, L. Song, *et al.*, "Neurbf: A neural fields representation with adaptive radial basis functions," in *Proceedings of the IEEE/CVF International Conference* on Computer Vision, 2023, pp. 4182–4194.
- [19] B. Hu, X. Zhu, D. Wang, et al., "Orbitgrasp: SE(3)-equivariant grasp learning," arXiv preprint arXiv:2407.03531, 2024.
- [20] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Machine learning*, vol. 110, no. 3, pp. 457–506, 2021.
- [21] W. Rudin, *Fourier analysis on groups*. Courier Dover Publications, 2017.
- [22] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [23] G. Khromov and S. P. Singh, "Some fundamental aspects about lipschitz continuity of neural networks," arXiv preprint arXiv:2302.10886, 2023.
- [24] M. Belkin, D. Hsu, S. Ma, and S. Mandal, "Reconciling modern machine-learning practice and the classical bias-variance trade-off," *Proceedings of the National Academy of Sciences*, vol. 116, no. 32, pp. 15849– 15854, 2019.
- [25] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research: The 18th International Symposium ISRR*, Springer, 2020, pp. 75–84.
- [26] K. Chen, R. Nemiroff, and B. T. Lopez, "Direct lidarinertial odometry: Lightweight lio with continuoustime motion correction," in 2023 IEEE international conference on robotics and automation (ICRA), IEEE, 2023, pp. 3983–3989.
- [27] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 697–12 705.

## APPENDIX A DERIVATION OF LIPSCHITZ CONSTANT

With a slight abuse of notation F, Eq. (3) can be written in the following way:

$$G(\phi) = \frac{a_i^0}{2} + \sum_{k=1}^n [a_i^k \cos(k\phi) + b_i^k \sin(k\phi)]$$
(8)

$$F(\phi) = \sigma(G(\phi)) \tag{9}$$

where  $\sigma(\cdot)$  is the sigmoid function and  $a_i^0$ ,  $a_i^k$ , and  $b_i^k$  are the Fourier coefficients predicted by our network.

We will derive an upper bound on the Lipschitz constant of  $F(\phi)$ , via the chain rule:

$$|F'(\phi)| = |\sigma'(G(\phi))| \cdot |G'(\phi)|.$$
(10)

First we differentiate  $G(\phi)$  w.r.t.  $\phi$ :

$$G'(\phi) = \sum_{k=1}^{n} k a_i^k \cos(k\phi) - k b_i^k \sin(k\phi)$$
(11)

$$= \sum_{k=1}^{n} k[a_i^k \cos(k\phi) - b_i^k \sin(k\phi)].$$
 (12)

For every  $k \in [1, n]$ , let:

$$R_k = \sqrt{(a_i^k)^2 + (b_i^k)^2}$$
(13)

$$\Delta_k = \arctan(\frac{-b_i^k}{a_i^k}),\tag{14}$$

Then Eq. (12) can be rewritten as a phase shift leveraging the cosine addition formula:

$$G'(\phi) = \sum_{k=1}^{n} k [R_k \cos(k\phi) \frac{a_i^k}{R_k} - R_k \sin(k\phi) \frac{b_i^k}{R_k}]$$
(15)  
$$= \sum_{k=1}^{n} k [R_k \cos(k\phi) \cos(\Delta_k) + R_k \sin(k\phi) \sin(\Delta_k)]$$
(16)

$$=\sum_{k=1}^{n} kR_k \cos(k\phi - \Delta_k).$$
(17)

Notice that the maximum value of a  $\cos(\cdot)$  is 1, as a result, Eq. (17) is upper bounded by:

$$G'(\phi) \le \sum_{k=1}^{n} kR_k \tag{18}$$

$$=\sum_{k=1}^{n} k\sqrt{(a_i^k)^2 + (b_i^k)^2}.$$
(19)

Because the sigmoid function is Lipschitz with a Lipschitz constant of  $\frac{1}{4}$ , by the chain rule:

$$|F'(\phi)| = |\sigma'(G(\phi))| \cdot |G'(\phi)|$$
(20)

$$\leq \frac{1}{4} \sum_{k=1}^{n} k \sqrt{(a_i^k)^2 + (b_i^k)^2}.$$
 (21)

As a result the Lipschitz Constant of  $F(\phi)$  is bounded by:

$$L^{i} \leq \frac{1}{4} \sum_{k=1}^{n} k \sqrt{(a_{i}^{k})^{2} + (b_{i}^{k})^{2}}.$$

Note that the peak of the derivative of  $G(\phi)$  and the sigmoid function  $\sigma(\cdot)$  may not coincide, and as a result we expect the Lipschitz constant achieved in practice to be smaller.

## APPENDIX B MODEL ARCHITECTURE

All the learned models described in this paper (Ours, *AngleInput*, and *AngleFree*) predict categorical distributions composed of B = 8 bins, parameterized by the concentration parameters  $[\gamma_{\phi}^1, \ldots, \gamma_{\phi}^8]$ .

All models share the same Point Pillars encoder architecture [27], which extracts a feature representation from the input point cloud Q. Specifically, the point cloud Q is first uniformly scaled to fit within a cube defined by  $x \in [-0.5, 0.5], y \in$ [-0.5, 0.5], and  $z \in [0, 0.5]$ . This normalized point cloud is then discretized into uniformly spaced pillars in the x-y plane, each with a grid size of 0.1m. Points within each pillar are augmented with  $x_c, y_c, z_c, x_p$ , and  $y_p$  where the c subscript denotes distance to the arithmetic mean of all points in the pillar and the p subscript denotes the offset from the pillar's x, y center [27]. Notice we exclude the reflectance feature, yielding an 8-dimensional augmented input representation for each lidar point. For each pillar, the encoder samples up to 32 points and generates a corresponding 32-dimensional feature vector. The resulting output from the Point Pillars encoder is a pseudo-image tensor with dimensions  $10 \times 10 \times 32$ .

Next, we detail the remaining components of the model architecture.

#### A. Neural Network Used in SPARTA

The pseudo-image from the Point Pillars encoder [27] is subsequently processed by a two-layer convolutional backbone, specified as follows:

- Input Channels: [32, 64]
- Output Channels: [64, 256]
- Kernel Size: [3, 3]
- Stride: [2, 2]
- Padding: [1, 1]

Each convolutional layer in the backbone employs batch normalization and ReLU activation. The resulting output tensor, of dimensions  $3 \times 3 \times 256$ , is passed through a max pooling operation to yield a compact 256-dimensional point cloud feature representation.



Fig. 9. Runtime of *AngleInput* (Green) increases with model complexity, while the runtime of ours (Red) remains roughly constant. This is because during planning, we only need to perform a low dimensional dot product and a sigmoid pass during inference, whereas *AngleInput* requires repeated neural network inference. The x-axis denotes additional hidden layers in the decoder and additional frequencies in our model, i.e.,  $n \in [3, 13]$ .

To represent the smooth analytical mapping  $\phi \rightarrow \gamma_{\phi}^{i}$ , we employ Fourier basis functions up to a maximum frequency of n = 3, specifically the basis set  $\{1, \cos(\phi), \sin(\phi), \cos(2\phi), \sin(2\phi), \cos(3\phi), \sin(3\phi)\}$ . The choice of n is justified in Section E with an ablation study. Consequently, for each bin  $i \in [1, 8]$ , the network predicts 7 Fourier coefficients. Thus, the SPARTA decoder consists of a multilayer perceptron (MLP) that transforms the 256-dimensional point cloud features into a 56-dimensional output vector  $[\mathbf{F}^{1}, \ldots, \mathbf{F}^{8}]$ . This MLP decoder contains one hidden layer with 512 neurons and utilizes layer normalization and ReLU activation functions.

## B. AngleInput

The pseudo-image from the Point Pillars encoder [27] is subsequently processed by a three-layer convolutional backbone, specified as follows:

- Input Channels: [32, 64, 128]
- Output Channels: [64, 128, 256]
- Kernel Size: [3, 3, 3]
- Stride: [2, 2, 2]
- Padding: [1, 1, 0]

The encoder for angle of approach  $\phi$  is a MLP that maps 1-dimensional input  $\phi$  to a 32-dimensional embedding with one hidden layer with 16 neurons. The decoder merges the point cloud feature (256-dimensional) and the angle feature (32-dimensional), and outputs the 8 concentration parameters  $[\gamma_{\phi}^1, \ldots, \gamma_{\phi}^8]$ , with a hidden layer with 64 neurons. Both MLPs utilizes layer normalization and ReLU activation functions.

## C. AngleFree

The AngleFree baseline network uses the same convolutional backbone as AngleInput. A decoder head transforms the 256-dimensional point cloud feature to the 8 concentration parameters  $[\gamma_{\phi}^1, \ldots, \gamma_{\phi}^8]$ , with two hidden layers with [128, 64] neurons.

## APPENDIX C INFERENCE EFFICIENCY

To demonstrate the query efficiency of our model compared to AngleInput, we measure their respective inference runtimes. Specifically, for AngleInput, we measure the runtime of encoding the approach angle, decoding the concentration parameters, and computing the categorical distribution (Eq. (1)). For our model, we measure the runtime associated only with computing the concentration parameters (Eq. (3))and the categorical distribution (Eq. (1)). Note other components (e.g. extracting feature vector from point cloud) of both models can be precomputed or reused. Considering that more complex point clouds require higher-dimensional point cloud features and deeper network architectures, we perform a stress test by progressively increasing the number of layers in both models, as well as the maximum frequency n in our model. Both models process 25,000 queries with average runtimes reported in Fig. 9. Our model demonstrates an orderof-magnitude faster runtime compared to AngleInput. Most importantly, the runtime of our model remains constant despite increased complexity, as it only involves a single batched, low-dimensional dot product and sigmoid computation <sup>6</sup>. In contrast, the runtime of AngleInput increases with complexity.

## APPENDIX D DATA COLLECTION DETAIL

In this section we provide more details of our data collection process using the automotive simulator *BeamNG.tech* [14]. We intend to generate random obstacles for the vehicle to drive over, and collect the tire deformation during the traversal. We choose a patch size of  $1 \times 1m$  that covers the footprint of the vehicle wheel, and randomly place (non-overlapping) obstacles in the patch. Maximum elevation of the obstacle is set to 0.5m(the vehicle's minimum ground clearance). For every patch, we collect its point cloud and downsample to a maximum of 1024 points. To minimize the influence of vehicle dynamics (for the purpose of generalization to unseen hardware), we place the vehicle such that only one wheel contacts the obstacle. The deformation extent  $d_w = \frac{r_w - r_{\text{inner}}}{r_{\text{outer}} - r_{\text{inner}}}$  of all contacting tire node w are recorded during the traversal. We discard all samples below 0.2 (the tire deformation when the vehicle is static on flat ground) and divide the rest into a histogram with 8 bins, which forms the empirical distribution y. In total, 24,000 samples  $\{\mathbf{Q}_k, \mathbf{y}_k, \phi_k\}$  are collected. We augment the samples by applying 8 random rotations (to both the point cloud and the angle of approach), resulting in a total of 216,000 samples. We further add random perturbation to the point clouds for more robust performance. We use 90% of the data for training and 10% for testing.

#### Appendix E

## ABLATION STUDY

In this experiment, we test our model's ability to identify safe traversal points and angles of approach in a row of packed,

<sup>&</sup>lt;sup>6</sup>The dimension of the dot product depends only on the choice of maximum frequency n. In fact the dimension is exactly 2n + 1.



Fig. 10. (a): Overview of the test environment and the vehicle. (b)-(d): Example point cloud (colored by elevation) and the CVaR (Viridis arces around the point clouds) predicted by our model at the candidate angles of approach.

Algorithm	$\alpha$	n	Suc. ↑	Dmg.↓
Ours	0.9	3	95	5
	0.9	1	91	9
	0.9	5	94	6
	0	3	89	11
AngleInput	0.9	-	92	8
	0	-	89	11
Elev	-	-	73	27

TABLE II. Number of success (Suc.) and vehicle damage (dmg.) in simulation. Our model outperforms *AngleInput* and *Elev*. The performance of the models decrease when we ablate CVaR with mean. **Best**, <u>Second best</u>.

randomly placed obstacles. An overview of the environment, a few example point clouds (colored by elevation), and the CVaRs predicted by our model are shown in Fig. 10. This environment is simpler than the boulder field in Section IV-C and thus helpful in isolating the influence of the variables we are interested in. We consider two baselines, the AngleInput and a heuristic baseline that chooses the point with lowest maximum elevation under vehicle's wheel footprint (Elev). Since the obstacles are placed on flat ground, the Elev baseline is equivalent to finding the plan with lowest stepping difficulty as in [8], [10]. As shown in Table II, our model achieves higher success rate (S.R., the vehicle drives over the obstacle without damage) than the AngleInput, but both learning-based method significantly outperforms Elev. This is expected because the obstacles are spawned with similar sizes, and as a result Elev degenerates to random selection, leading to much lower success rate and thus high risk of vehicle damage.

a) Ablating CVaR: To demonstrate the effectiveness of using CVaR during planning to capture worst-case risk, we ablate it by setting  $\alpha = 0$ , which is equivalent to using the mean of the predicted distribution. As reported in Table II, the performance of both our model and AngleInput dropped. This is because accounting for worst-case (tail) risk is crucial in safety-critical tasks, and CVaR is a principled way for representing tail risk in robotics [25].

b) Ablating Maximum Frequency: As demonstrated in Section III-B and Section III-C, the choice of maximum frequency n is closely related to the smoothness and generalization of the mapping from angle of approach  $\phi$  to each of the concentration parameters  $\gamma_{\phi}^{i}$ . Therefore, we adjust the number



Fig. 11. Test  $\text{EMD}^2$  loss of our model with varying maximum frequency. They achieve similar test loss.

of Fourier bases n in Eq. (3) and analyze the results in this section. We consider the cases n = 1 and n = 5, and visualize the test  $EMD^2$  loss in Fig. 11. Overall, we observe similar test EMD<sup>2</sup> loss for all three models. In Table II we include their performances on the aforementioned experiment, on which they achieve success rate 91% and 94%, respectively. Notice the model with n = 1 achieves lower success rate compared to the other two models. We hypothesize this is because the basis functions can be too simple to capture the change of  $\gamma_{\phi}^{i}$ with  $\phi$ , i.e. it might not be expressive enough to parameterize the mapping we are interested in. Although the model with maximum frequency n = 5 achieves similar success rate as ours (n = 3), an overly high maximum frequency n introduces higher upper bound on the Lipschitz constant of the mapping and can introduce more abrupt changes in  $\gamma^i_\phi$  with small changes in  $\phi$ , as explained in Section III-C. Therefore, to balance expressivity and generalization, we choose n = 3in our model and our experiments. However, n is an easily changed hyperparameter that can be adjusted depending on the frequency composition of the mapping function we are trying to approximate.

### APPENDIX F

## HARDWARE SPECIFICATION

Our robot is built on top of the AgileX Scout Mini chassis, a skid-steering wheeled platform. It has a footprint of approximately  $0.6 \times 0.6m$  and can achieve a maximum speed of 3 m/s. The robot is equipped with an Ouster OS1 LiDAR with 128 channels. The LiDAR runs with a horizontal resolution of 1024 at 10 Hz, and we use its internal 6-axis IMU (100 Hz).



Fig. 12. (a): Obstacle and directions of the traversals. (b): The linear traction distribution of the forward (Blue) and backward (Orange) traversal. The backward distribution has significantly more probability mass in the lower bins. An algorithm not modeling the angle-dependency of linear traction (Green) would not be able to identify the obstacle as traversable from the forward orientation.

The point cloud and IMU measurements are processed using DLIO [26] for prebuilding the point cloud map and online localization. The onboard PC is an ASUS ROG NUC 970, with an Intel Core Ultra 9 CPU (16-Core), 32 GB RAM, and a Nvidia GeForce RTX 4070 GPU (8 GB).

#### Appendix G

#### GEOMETRIC STRUCTURE IN OFF-ROAD AUTONOMY

In this section, we demonstrate other problem settings that could benefit from modeling angle-dependency. We use the angle-dependency of linear traction (the robot's ability to execute commanded velocity on the given terrain [4]) as an example, and show how our approach can be applied to this problem setting. To construct an empirical dataset, we command the vehicle to drive at a constant linear velocity of 0.5 m/s over the obstacle, along the forward (Blue) and the backward (Orange) direction shown in Fig. 12a. We record the achieved linear velocity (estimated via DLIO [26]), and compute the linear traction as the ratio between the achieved linear velocity and the commanded linear velocity (0.5 m/s). The data points are discretized into histograms and then normalized into categorical distributions. As shown in Fig. 12b, the backward traversal (Orange) achieved significantly lower linear tractions compared to the forward traversal (Blue). This matches our intuition because the obstacle is smooth from the forward direction and sharp from the other. A system that models the angle-dependency of linear traction would thus be able to capture this difference. Notice SPARTA can be applied to this problem setting with minimal changes (e.g. adjusting input point cloud size) using an empirical dataset of linear traction distributions. However, systems that do not model angle-dependency (Green) [4] could be overly conservative and never command the robot to traverse the obstacle from the forward direction, which is traversable in practice.



Fig. 13. Example predictions of SPARTA. *Rainbow* denotes point elevation, and *Viridis* circles around the point clouds denote the CVaR at the corresponding angle of approach. The CVaR matches our intuition because it indicates the obstacles are more risky from orientations with sharp edges and less risky from orientations with smooth elevation change.

## APPENDIX H EXAMPLE CVAR PREDICTIONS OF SPARTA

In Fig. 13, we visualize some example obstacles in our test set and the CVaR predicted by our model with angle of approach  $\phi \in \{0, 1, \dots, 359\}$ . The *Rainbow* color map denotes the elevation of the point (Red denotes high elevation), and *Viridis* denotes the predicted CVaR (Yellow corresponds to high CVaR). A trend we notice from the examples is that our model predicts lower CVaR for angles where the obstacle has a smoother elevation change. This is aligned with our intuition, further supporting the validity of our approach for identifying risky obstacles and angles of approach.