

Perception-Driven Estimation of Terrain Motion Resistance for UGVs

Tom Bourbon, Stéphanie Aravecchia, Cédric Pradalier

Abstract—Accurate estimation of wheel–terrain interaction parameters is important for efficient navigation of Unmanned Ground Vehicles in unstructured outdoor environments. In this paper, we propose a hybrid data-driven and model-based method to estimate a priori *motion resistance*, a terrain-specific parameter representing the force opposing wheel motion, which is largely influenced by terrain class and geometry. The proposed method relies on learning motion resistance from proprioceptive feedback collected on reference terrains. This learned model is then transferred to new environments, where motion resistance is inferred from exteroceptive observation, including LiDAR and cameras, leveraging terrain geometry and class information. To capture uncertainty from terrain roughness and sensor noise, we evaluate two probabilistic models predicting motion resistance distributions: a Gaussian-MLP and a Gaussian Process Regressor. Their robustness to domain shifts is assessed by measuring performance degradation as the target diverges from the source domain. Extensive off-road field experiments validate the method’s effectiveness, demonstrating accurate prediction of motion resistance and its potential for deployment.

Index Terms—Field robotics, Wheeled robots, Terrain classification, Off-road navigation

I. INTRODUCTION

Understanding and perceiving terrain is an important capability for autonomous ground robots operating in real-world environments. While numerous studies have explored this challenge in structured or urban settings, uncontrolled outdoor environments present a wider range of terrain types and unpredictable conditions making reliable terrain understanding particularly challenging. A comprehensive survey by [1] laid a strong foundation by reviewing terrain traversability analysis methods.

Achieving robust behavior in unstructured environments can benefit from leveraging prior knowledge of the terrain. Especially, such a knowledge of terrain parameters, including both surface type (e.g., grass, gravel, sand) and terrain geometry, can be directly leveraged within control laws. For example, [2] uses these two information sources to adjust each wheel’s speed based on experimentally derived optimal slip models for different soils. Vehicles operating in outdoor environments increasingly rely on near-to-far sensing [3]. This method leverages exteroceptive data from sensors such as LiDAR and cameras to infer terrain interaction characteristics before physical contact, using models trained from prior proprioceptive experience. LiDAR provides detailed

This research was supported by the French Agence Nationale de la Recherche, under grant ANR-23-CE33-0015 (project MUSCAA). Authors are with GeorgiaTech Europe - IRL2958 GT-CNRS, Metz, France.



Fig. 1: Husky A200 acquiring proprioceptive data in various outdoor environments.

geometric information about terrain [4] and roughness [5], while vision-based methods enable semantic classification of terrain types [6]. These rich exteroceptive data sources facilitate estimation of terrain-specific properties used for robust control and planning. In this study, we focus on motion resistance, the force opposing wheel movement caused by energy losses from deformation of both the wheel and terrain surface [7]. Motion resistance depends not only on terrain class but also on the robot’s pose, which affects wheel load distribution. To robustly estimate this parameter despite noisy sensor data and varying robot attitudes, we choose to learn probabilistic models of motion resistance conditioned on terrain geometry, robot attitude (i.e pitch and roll), and terrain class.

The main contributions of this paper are:

- A methodology to predict ground motion resistance parameters from exteroceptive sensing, leveraging knowledge acquired via proprioceptive measurements. Two probabilistic predictive models are considered: a Gaussian-output multilayer perceptron (Gaussian-MLP) and a Gaussian Process Regressor (GPR).
- A quantitative robustness analysis under domain shifts,

evaluating how performance degrades when a model trained on a source domain is applied to a target domain. This provides insight into the generalization capabilities of the proposed approach across diverse operational environments and can serve as a quantitative indicator of the predictive reliability.

- Experimental validation through field tests, demonstrating reliable inference of motion resistance and its reliability using only exteroceptive inputs.

II. RELATED WORK

Terrain characterization is important for autonomous robots to navigate complex environments safely and efficiently. It enables the prediction of traversability by identifying safe paths [8] and supports control strategies that dynamically adapt robot’s behavior to changing terrain conditions [9].

Beyond terrain classification, several works aim to infer physical ground parameters that influence robot mobility. For instance, [10] proposes predicting energy consumption along a known trajectory using terrain information extracted from a Digital Elevation Map (DEM). While their method demonstrates promising results across various terrain types, it lacks a fine-grained classification mechanism for generalization across diverse environments. Lateral slippage, particularly relevant for large vehicles like tractors operating in agricultural settings, is addressed in [11], where variations in terrain geometry and surface conditions are shown to significantly impact slip behavior. More generally, the challenge of modeling terrain-dependent mobility lies in the deformable and often non-elastic nature of unstructured terrain. As discussed in [7], these effects contribute to motion resistance—a set of forces opposing the robot’s motion—which serves as the central ground parameter considered in this work.

In terrain analysis, the terrain is often characterized under the current robot position [12], with proprioceptive sensors. In this work, on the contrary, we characterize the terrain in a near-to-far approach. Near-to-far learning consists in predicting ground-robot interaction on a patch of terrain where it will be in the future, based on its current observations from exteroceptive sensors. Such an approach is proposed in [13], where a model of the robot-ground interaction is learned based on some local features, such as terrain bumpiness measured with the IMU, and some remote features, such as visual appearance of the terrain using a color-based criterion. This method learns a model of ground-robot interaction highly dependent on the robot itself, whereas our proposed method characterizes the terrain itself.

Estimating terrain-related properties requires not only modeling physical interactions but also recognizing surface characteristics that influence them. While sensors such as RGB cameras, LiDAR, and other proprioceptive devices provide valuable and complementary information, effectively fusing these modalities remains a challenge. As a result, a wide range of perception-driven approaches have been developed, each leveraging different combinations of sensor

data and features to characterize terrain more accurately. Appearance-based methods extract visual features such as texture, color, and shape from RGB cameras to classify terrain types. They perform well in structured, flat environments [14], [15], but outdoor terrains often show subtle differences, complicating classification [16], [17]. While [16] uses aerial imagery for segmentation, our method relies on onboard sensors, capturing terrain near the robot and fitting a near-to-far perception strategy. Additionally, [17] splits images into tiles, offering coarser detail than semantic segmentation. Moreover, camera-only methods do not capture geometry details.

Geometry-based methods use LiDAR to directly capture the surface shape of the terrain, providing robustness to changes in lighting and weather conditions. For example, [18] proposes multi-frame fusion of LiDAR data to improve terrain representation. However, these methods are limited in their ability to capture non-geometric features such as texture or vegetation type, which can also significantly affect terrain traversability.

In challenging, unstructured environments such as those encountered in agriculture, many approaches adopt hybrid methods that combine geometry-based and appearance-based techniques to achieve more robust terrain characterization. Using both appearance and geometry information, [10] and [19] generate cost maps for different objectives: estimating energy consumption and predicting velocity loss during navigation through dense vegetation, respectively. In our work, we favored a hybrid approach using a LiDAR and a camera, leveraging the strengths of both sensors to provide a more comprehensive terrain understanding.

To integrate semantic and geometric information within a unified data structure, several representations have been proposed, each offering different trade-offs between accuracy, efficiency, and expressiveness. Digital Elevation Maps (DEMs) offer a lightweight 2.5D representation of surface topology and are frequently used in real-time terrain analysis due to their compactness and fast processing capabilities [18]. In our case, we adopt the DEM representation, as our focus lies on wheel-terrain interactions at the surface level. This makes DEMs particularly suitable, since they offer a good balance between efficiency and relevant geometric detail, and we prioritize lightweight processing over accurate modeling of vertical structures. This structure also allows for the integration of a semantic layer, resulting in a Semantic DEM.

III. METHODOLOGY

This section addresses four key aspects: (i) the formal definition of motion resistance and its formulation for a skid-steer platform, (ii) the learning process that leverages proprioceptive measurements to enable inference from exteroceptive data, (iii) a robustness analysis that evaluates the model’s sensitivity to distributional shifts between training and deployment environments, and (iv) a methodology for inferring

motion resistance and its reliability using exteroceptive data only.

A. Definition

In this section, we directly apply [7] to calculate the motion resistance for the skid-steer model. Here, we report the equations necessary to understand our methodology. The symbols corresponding to the variables are reported in Fig 2 for convenience.

When a tire rolls over the ground, both the tire and the surface can deform. This deformation causes the vertical force on the tire to shift slightly forward. As a result, the rolling motion encounters resistance, requiring a driving torque T_R to keep the vehicle moving. Thus, the necessary torque T_R is:

$$T_R = m_R \cdot r \cdot F_z \quad (1)$$




Fig. 2: Left: definition of the variables. Right: schematic of the skid-steer platform.

In DC brushed motors, the mechanical torque is proportional to the electrical current. For a skid-steer vehicle equipped with two independent drive motors, the driving torque on the side $s \in \{Left, Right\}$ is:

$$T_{R,s} = \tau \cdot k_t \cdot I_s \quad (2)$$

The terrain's geometry leads to variations in the vertical load distribution between the two wheel pairs. Under the quasi-static assumption, i.e., neglecting inertial effects, the platform is solely subjected to weight-induced forces.

$$F_{z,Left} = \frac{W}{2} \cos \phi \cos \theta + W \cos \theta \sin \phi \cdot \frac{h}{B} \quad (3)$$

$$F_{z,Right} = \frac{W}{2} \cos \phi \cos \theta - W \cos \theta \sin \phi \cdot \frac{h}{B} \quad (4)$$

Leading to the equation of motion resistance :

$$m_{R,s} = \frac{\tau \cdot k_t \cdot I_s}{r \cdot F_{z,s}(\theta, \phi)} \quad (5)$$

However, both terrain geometry and terrain class significantly influence motion resistance as demonstrated by field experiments reported in the Preliminary Results section IV-B2. We denote the terrain class as $c_j \in \mathcal{C} = \{c_1, c_2, \dots, c_n\}$, where each class c_j differs in mechanical and physical properties such as deformability, adhesion, and slip, affecting the resistance experienced during motion. We thus define the motion resistance function $m_R^{(j)}$ for class c_j as:

$$m_{R,s}^{(j)}(\theta, \phi) = \frac{\tau \cdot k_t \cdot I_s^{(j)}(\theta, \phi)}{r \cdot F_{z,s}(\theta, \phi)} \quad (6)$$

B. Learning motion resistance distribution

We present two probabilistic models, a Gaussian-based multilayer perceptron (Gaussian-MLP) and a Gaussian Process Regressor (GPR), that both learn the motion resistance distribution conditioned on triplets (θ, ϕ, c_j) . These models are trained from proprioceptive measurements to enable inferring motion resistance using exteroceptive inputs only. This probabilistic formulation is introduced to account for the noise observed in motion resistance estimates, which arises from sensor inaccuracies and terrain irregularities (e.g., small bumps or uneven surfaces), leading to inherent variability in the proprioceptive signals.

1) *Dataset*: We define our dataset for a class c_j as

$$\mathcal{D}_j = \{(x_i, y_i)\}_{i=1}^N,$$

where N denotes the total number of samples. Each input vector

$$x_i = \begin{bmatrix} \theta_i \\ \phi_i \end{bmatrix} \in \mathbb{R}^2, \quad \text{with } \theta_i = \text{pitch}, \quad \phi_i = \text{roll},$$

represents the vehicle's orientation angles, and the corresponding output $y_i \in \mathbb{R}$ denotes the motion resistance, computed from proprioceptive measurements such as the current injected into the wheels, using the model-based formulation described in Equation (5).

Our modeling approach is based on the assumption that data collected at fixed pitch, roll, and terrain class follow a normal distribution (see the hypothesis validation in Preliminary results section IV-B1).

2) *MLP-Gaussian models*: We consider a probabilistic Multi-Layer Perceptron (Gaussian MLP) model that predicts both the mean $\hat{\mu}^{(j)}(\mathbf{x})$ and the standard deviation $\hat{\sigma}^{(j)}(\mathbf{x})$ of the target distribution, given an input $\mathbf{x} = (\theta, \phi)$ and a terrain class c_j .

The model is trained by minimizing the negative log-likelihood of the Gaussian distribution:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left[\frac{(y_i - \mu(\mathbf{x}_i))^2}{2\sigma(\mathbf{x}_i)^2} + \frac{1}{2} \log \sigma(\mathbf{x}_i)^2 \right] \quad (7)$$

where y_i denotes the observed motion resistance measurement corresponding to the input \mathbf{x}_i .

Gaussian MLP models learn a parametric function defined as:

$$\mathcal{M}_j^{\text{MLP}} : \mathbb{R}^2 \rightarrow \mathbb{R}^2, \quad \mathcal{M}_j^{\text{MLP}}(\theta, \phi) = \left[\hat{\mu}_{\text{MLP}}^{(j)}(\theta, \phi), \hat{\sigma}_{\text{MLP}}^{(j)}(\theta, \phi) \right] \quad (8)$$

3) *Gaussian Process Regressor*: As a second model, we use a Gaussian Process Regressor (GPR) to model the mapping from the input angles (θ, ϕ) to the output y . We use a standard kernel combining a Radial Basis Function (RBF) kernel with a White Noise kernel:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\ell^2}\right) + \sigma_n^2 \delta_{ij} \quad (9)$$

where ℓ is the length scale of the RBF kernel, σ_n^2 the noise level, and δ_{ij} is the Kronecker delta. The kernel is defined as:

$$k = \text{RBF}(\ell = 1.0) + \text{WhiteKernel}(\sigma_n^2 = 10^{-2}) \quad (10)$$

Gaussian Process models estimate a non-parametric function given by:

$$\mathcal{M}_j^{\text{GPR}} : \mathbb{R}^2 \rightarrow \mathbb{R}^2, \quad \mathcal{M}_j^{\text{GPR}}(\theta, \phi) = \left[\hat{\mu}_{\text{GPR}}^{(j)}(\theta, \phi), \hat{\sigma}_{\text{GPR}}^{(j)}(\theta, \phi) \right] \quad (11)$$

C. Robustness analysis

In this section, we focus on evaluating the robustness of the learned model. Specifically, we assess the reliability of the predictions when subject to distributional shifts between the training terrain and the deployment terrain.

1) *Data-distribution shift*: Given a set of sample points $\{\mathbf{x}_i\}_{i=1}^n$, the probability density function $p(\mathbf{x})$ at a point $\mathbf{x} \in \mathbb{R}^d$ can be estimated by:

$$\hat{p}(\mathbf{x}) = \frac{1}{nh^d} \sum_{i=1}^n K\left(\frac{\mathbf{x} - \mathbf{x}_i}{h}\right) \quad (12)$$

where:

- $K(\cdot)$ is the kernel function, typically a probability density function centered at zero,
- $h > 0$ is the bandwidth parameter controlling the kernel width,
- d is the dimension of the data.

For the isotropic multivariate Gaussian kernel, we have:

$$K(\mathbf{u}) = \frac{1}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2}\mathbf{u}^\top \mathbf{u}\right) \quad (13)$$

which leads to:

$$\hat{p}(\mathbf{x}) = \frac{1}{n(2\pi)^{d/2}h^d} \sum_{i=1}^n \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_i\|^2}{2h^2}\right) \quad (14)$$

The bandwidth h controls the smoothness of the estimate: a small h leads to a noisy estimate, while a large h results in an over-smoothed density.

To estimate the similarity between two datasets $A = \{\mathbf{x}_i^A\}_{i=1}^{n_A}$ and $B = \{\mathbf{x}_j^B\}_{j=1}^{n_B}$, we compute the average density

of points in B with respect to A using Kernel Density Estimation (KDE):

$$\text{KDE}(B, A) = \frac{1}{n_B} \sum_{j=1}^{n_B} \hat{p}_A(\mathbf{x}_j^B) \quad (15)$$

This metric reflects how well the samples in B are supported by the distribution represented by A . In our work, this formulation is used in two complementary ways: (i) to compare entire datasets (e.g., estimating the similarity between training and test sets), and (ii) to evaluate whether a test point lies in a high-density region of the training set. In the latter case, we simply evaluate the density $\hat{p}_A(\mathbf{x})$ of a test point \mathbf{x} under the distribution of A , without requiring a reference set B . A higher density indicates that the point is well-supported by the training data, and can be assimilated to greater confidence, as shown in the Results Section IV-D.

D. Inference from exteroceptive data

As in a *near-to-far* scenario, the objective is to predict the motion resistance of the terrain using solely exteroceptive data, and one of the trained models $\mathcal{M}_j^{\text{GPR}}$ or $\mathcal{M}_j^{\text{MLP}}$ described previously.

Terrain class and geometry are encoded in a semantic DEM. We assume the availability of such a semantic DEM, which can be reconstructed by first applying semantic segmentation to camera images, then projecting 3D points from the point cloud onto the segmented 2D images to assign semantic labels to each point. From this labeled point cloud, a gridded elevation map is generated by rasterizing elevation data onto a regular grid. Notably, this semantic DEM encodes: (i) Local terrain geometry, typically represented by surface normals derived from elevation gradients; (ii) Terrain class labels, such as vegetation, gravel, mud, etc.

Given such a representation, our goal is to estimate the motion resistance at arbitrary locations along the robot's future trajectory. To enable this behavior, we pre-compute the normal vector for each cell in the semantic DEM based on its neighbors. Given a query 2D-pose $P = [x, y, \psi]$ (robot position and yaw) along the trajectory, the semantic DEM returns $DEM(P) = (c_j, \phi, \theta)$, the terrain class and the corresponding pitch and roll of the robot. These values form the input features for the motion resistance prediction models.

E. Evaluation methodology

To demonstrate the quality of the inferences made by the Gaussian MLP and the Gaussian Process Regressor, we propose to use two baselines: i) a first baseline that does not account for the terrain geometry, and ii) a second, more informed baseline leveraging the distribution of neighboring data points to interpolate the predicted distribution.

1) *Baseline without geometry*: In order to highlight the value of using geometry, we consider as a baseline a model that predicts a constant mean μ for all pairs (θ, ϕ) of known class c_j . We refer to this baseline as baseline *mean*.

2) *Interpolation Baseline*: We implement a non-parametric baseline using bilinear interpolation over a discretized representation of the training set. First, the training data is binned along the pitch and roll axes using a fixed bin size. For each bin, we compute the sample mean and standard deviation of the motion resistance, producing a sparse grid of statistics over the feature space.

Let $X_{\text{train}} = [\theta_i, \phi_i]$ denote the training inputs, where θ and ϕ correspond to pitch and roll, respectively. We discretize the input space into a set of bins \mathcal{B} , and for each bin $b \in \mathcal{B}$, located at (θ_b, ϕ_b) , we estimate:

$$\mu_b = \mu(\theta_b, \phi_b) = \mathbb{E}[\text{MR}], \quad \sigma_b = \sigma(\theta_b, \phi_b) = \text{std}(\text{MR})$$

These statistics define the dataset for each class c_j :

$$\mathcal{S}_j = \{(\theta_b, \phi_b, \mu_b, \sigma_b) \mid b \in \mathcal{B}\} \quad (16)$$

At inference time, for each test point (θ, ϕ) and class c_j , we interpolate the mean and standard deviation values from \mathcal{S}_j using bilinear interpolation. This results in predicted values $\hat{\mu}^{(j)}(\theta, \phi)$ and $\hat{\sigma}^{(j)}(\theta, \phi)$. We refer to this baseline as *baseline interp*.

3) *Metrics*: To evaluate performance, we associate each test point to its corresponding bin, retrieve the empirical bin statistics from the training set, and compute the RMSE between predicted and empirical values:

$$\text{RMSE}_\mu = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{\mu}_i - \mu_i^{\text{bin}})^2}, \quad \text{RMSE}_\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_i - \sigma_i^{\text{bin}})^2}$$

IV. RESULTS

A. Experimental setup

Experiments were conducted using a Clearpath Husky A200 robotic platform equipped with a FLIR RGB camera, an Ouster OS1-16 LiDAR, an SBG Ellipse-D IMU, and a GPS RTK module for localization. All trials were performed at a constant speed of 1.0 m/s along a straight 5 m transect. The dataset includes time-synchronized measurements of platform attitude (pitch and roll) acquired at 20 Hz, and motor current intensity recorded at 10 Hz.

Assuming the terrain is locally planar over a $[1m \times 1m]$ area, the data is averaged using a 1-second sliding window to mitigate sensor noise and the natural roughness of outdoor surfaces.

In our experiments, we focus on four terrain classes, defining $\mathcal{C} = \{\text{asphalt, grass, tall grass, gravel}\}$, as illustrated in Figure 1. This selection is motivated by both practical and analytical considerations. These surfaces are commonly encountered in outdoor robotic applications, and they exhibit clearly distinct motion resistance profiles. Grass typically presents higher deformability and greater slip, resulting in higher resistance values, while asphalt is a rigid, non-deformable surface with low slip, leading to lower resistance. The strong contrast in the statistical distributions of motion

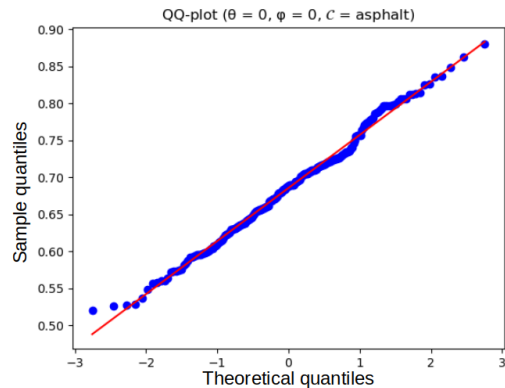


Fig. 3: Quantile-Quantile (QQ) plot comparing the empirical distribution of motion resistance within a bin to the theoretical Gaussian distribution.

resistance (Section IV-B2) between the classes makes them well-suited for validating the effectiveness of our model in capturing terrain-dependent behavior.

B. Preliminary Results

1) *Gaussian hypothesis*: The QQ-plot for an example bin $\theta = 0$, $\phi = 0$, and $c = \text{asphalt}$ is represented in Figure 3. It shows that the sample quantiles align with the theoretical normal distribution. The distribution corresponding to the bin can be approximated as a Gaussian distribution, enabling the application of Gaussian-based methods in the modeling and analysis conducted in this work.

2) *Influence of class and geometry on Motion Resistance*: To justify the motion resistance model introduced in Section III, Figure 4 qualitatively illustrates the influence of both terrain class and geometry on the predicted motion resistance across different terrain types. In Figure 4, we observe the impact of semantic class under constant geometric conditions, with the analysis conducted on flat terrain. Asphalt yields the lowest motion resistance values, centered around 0.6, while grass and gravel surfaces exhibit significantly higher values, approximately 0.8. Interestingly, the variance of motion resistance is much greater for grass and gravel than for asphalt. This higher spread can be attributed to two key factors: (i) more stochastic deformation behavior due to the compliant or loose nature of the surface material, and (ii) a finer-scale, irregular micro-geometry (e.g., bumps), which makes the local planarity assumption of the terrain surface less valid.

Figure 4 also illustrates the impact of terrain geometry on motion resistance. On the left, the terrain is nearly flat with pitch and roll equal to zero. On the right, the same semantic class is observed under a steeper geometry, with pitch equal to 9 degrees. A clear increase in predicted motion resistance is observed as the pitch increases for all classes. This demonstrates the sensitivity of motion resistance to local terrain inclination, even when the semantic label remains unchanged.

Complementary, Figure 5 illustrates the repeatability of motion resistance measurements under similar experimental conditions—specifically, when the pitch, roll, and terrain class (i.e., the semantic and geometric configuration) are the same.

These observations confirm the necessity of integrating both semantic and geometric features for accurate motion resistance prediction.

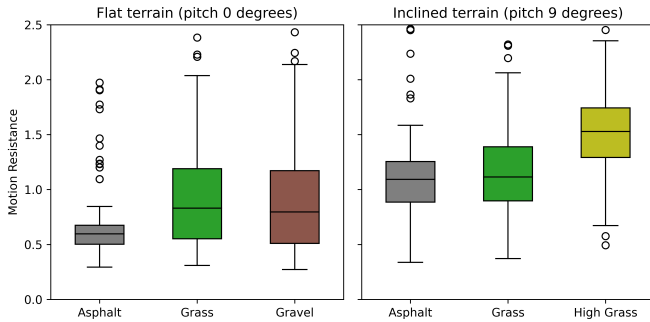


Fig. 4: Influence of class on motion resistance. Left: Distribution of the motion resistance on data collected on three classes and flat terrain. Right: Three different classes on inclined terrain.

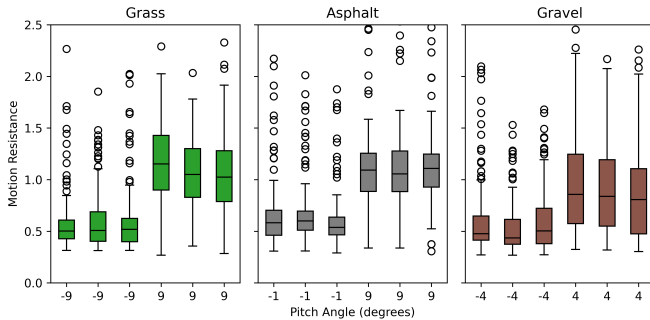


Fig. 5: Influence of geometry on motion resistance. Each graph represents a semantic class and contains boxplots of data collected at a constant pitch angle. Experiments were conducted with two different pitch angles, each repeated three times highlighting repeatability.

C. Learning motion resistance distribution

The following analysis quantitatively demonstrates the effectiveness of the proposed models in capturing motion resistance distribution by incorporating both semantic and geometric terrain information. The results presented in Table I show that the proposed models (MLP and GPR), as well as baseline *interp* are able to predict motion resistance distributions with a mean $RMSE_{\mu}$ around 0.1. These results were obtained from experiments conducted on inclined grass, where the average motion resistance is close to 1, corresponding to an approximate 10% error in the mean of predicted motion resistance distribution. A $RMSE_{\sigma}$ of 0.026, as obtained by GPR, corresponds approximately to a 25% error on the actual standard deviation. In contrast, the baseline *mean*,

which does not incorporate geometric information, achieves a significantly higher $RMSE_{\mu}$ of 0.166. This highlights the importance of leveraging terrain geometry, in addition to semantic class, for accurate motion resistance estimation.

The MLP used here has a single hidden layer with 32 neurons and trains in about one second. To account for seed-related variability, we train multiple models. In contrast, the Gaussian Process Regressor requires no weight training, making it well-suited for potential online adaptation.

While the baseline *interp* performs well, achieving an $RMSE_{\mu}$ comparable to the proposed models and a lower $RMSE_{\sigma}$, the following section highlights its lack of robustness under data distribution shifts.

	MLP	GPR	Baseline <i>interp</i>	Baseline <i>mean</i>
$RMSE_{\mu}$	0.094	0.104	0.098	0.166
$RMSE_{\sigma}$	0.032	0.026	0.010	0.13

TABLE I: Performance comparison of the models, evaluated with $RMSE_{\mu}$ and $RMSE_{\sigma}$ defined in Sec.III.

D. Data-distribution shift robustness

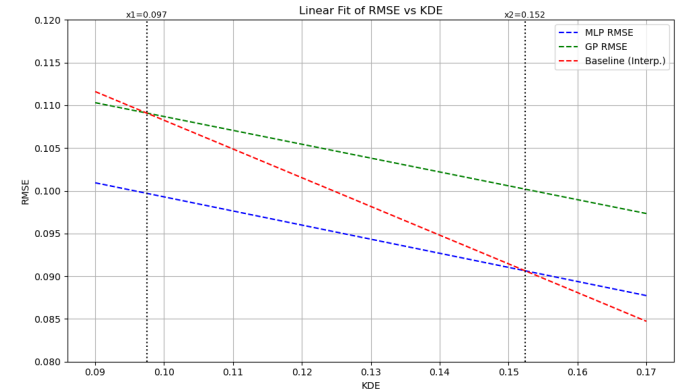


Fig. 6: Degradation of $RMSE_{\mu}$ as the Kernel Density Estimate (KDE) of the test set within the training distribution decreases.

As illustrated in Figure 6, baseline *interp* and the MLP perform comparably at higher KDE values. However, we observe that as the KDE decreases, i.e., as the support of set B becomes sparser within set A and requires more interpolation, the performance of the interpolation baseline deteriorates more sharply, with up to a 20% drop in accuracy.

Two notable KDE thresholds emerge from the analysis: below $x_2 = 0.152$, the MLP model outperforms the baseline interpolation, while below $x_1 = 0.097$, the Gaussian Process achieves better performance than the baseline interpolation. These thresholds indicate regions in the input space where the learned models provide significant improvements over baseline *interp*.

E. Inference from exteroceptive data validated with field test

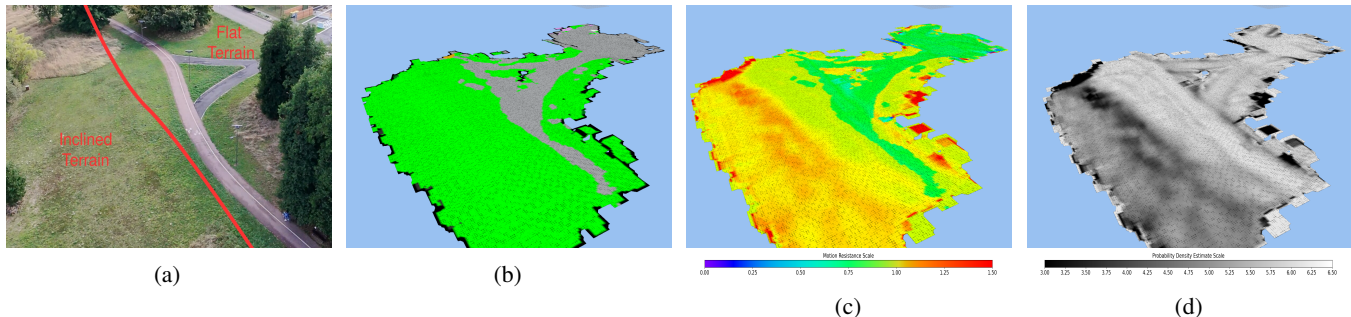


Fig. 7: Qualitative results for motion resistance mapping. From left to right: (a) Aerial view of the field experiment setup, (b) Semantic DEM, (c) Motion Resistance DEM (highest pitch direction), (d) Reliability DEM (highest pitch direction).

1) *Qualitative Evaluation of DEM layers:* This section presents qualitative results obtained under real outdoor conditions. We selected a test site consisting of two terrain classes—*asphalt* and *grass*—and varying geometry. An aerial view of the test site is shown in Figure 7a. In particular, the grassy area features an inclined slope with gradients ranging from 0° to 11° (on the left of the image).

Figure 7b shows the semantic DEM, where asphalt is represented in gray and grass in green. This semantic DEM is generated following the procedure described in Section III-D. From a qualitative standpoint, the segmentation appears generally accurate: the road and the adjacent bike lane are delineated, indicating recognition of asphalt surfaces. However, some areas appear to be poorly segmented due to inaccuracies from the semantic segmentation model when processing images captured by the embedded camera.

Figure 7c presents the predicted motion resistance derived from the semantic DEM. This prediction is obtained using a Gaussian MLP model, which leverages both semantic information and terrain geometry extracted from the semantic DEM. Since motion resistance depends on the local pitch and roll, it is inherently trajectory-dependent. To provide a trajectory-independent visualization, we represent here the motion resistance in the direction of steepest slope. On the flat region, a clear contrast in predicted resistance between asphalt and grass is observed, emphasizing the influence of terrain type on the model’s output. On the inclined grassy slope, the motion resistance varies consistently with the local slope gradient, demonstrating the model’s sensitivity to terrain geometry.

Finally, Figure 7d illustrates the reliability of the predicted motion resistance, based on the methodology described in Section IV-D. As with the motion resistance map, this reliability estimation is trajectory-dependent; hence, the visualization corresponds to the direction of steepest slope. For each cell, reliability is quantified by evaluating the probability density of its pitch-roll pair under the distribution observed in the training dataset, as defined in Equation 14. Regions shaded in lighter gray correspond to terrain geometries that were frequently encountered during training, implying a higher expected prediction accuracy. Conversely, areas approaching

black indicate rarely seen configurations, where the model is expected to be less confident in its predictions.

2) *Motion Resistance Inference Evaluation:* The objective of this validation field test is to compare the motion resistance perceived by the robot with the one predicted by our model. We selected a terrain composed of *grass* and *asphalt*, featuring varied geometry and located in a geographical area distinct from the one used during the model training. The robot performed 5-meters transects along predefined directions. During each traversal, the perceived motion resistance was continuously estimated from proprioceptive signals (motor currents, IMU) recorded onboard. Knowing the semantic DEM for the environment, the local pitch, roll, and terrain class encountered along the trajectory could be retrieved. Using this information, the predicted motion resistance was computed for each point along the path.

Figure 8 shows the time series comparison between the predicted and the perceived motion resistance along a test trajectory. The Root Mean Square Error (RMSE) between the predicted and perceived motion resistance is 0.13, which is satisfactory and of the same order of magnitude as the results presented in Section IV-C. Overall, the predictor successfully captures the general trend of the perceived resistance.

In addition, the predicted uncertainty (standard deviation) is consistent with the observed variations: only 5.60% of the measured values fall outside the 95% confidence interval, indicating that the uncertainty estimation is reliable.

Finally, Figure 9 highlights the relationship between the estimated reliability, obtained through Probability Density Estimate (PDE) (Equation 14) evaluation at the input point, and its distance to the predicted Gaussian distribution. The distance used here is defined as $D = \frac{|\mu - x|}{\sigma}$, where μ and σ are the mean and standard deviation of the predicted distribution, respectively.

We observe that most points remain within approximately two standard deviations of the predicted mean ($D < 2$) for PDE values above 2. Below this threshold ($PDE < 2$), the motion resistance experienced by the robot deviates significantly from the predicted distributions. This emphasizes the correlation between low PDE values and increased inaccuracy in the predictions.

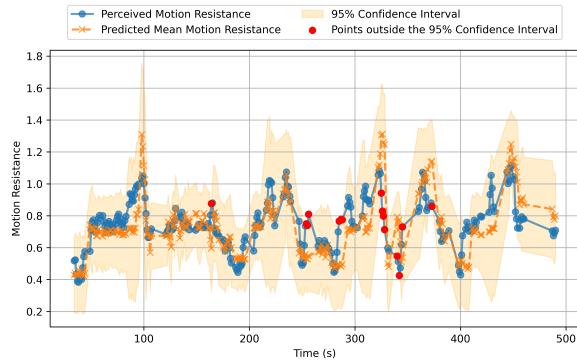


Fig. 8: Time series comparison between predicted and perceived motion resistance on a test trajectory.

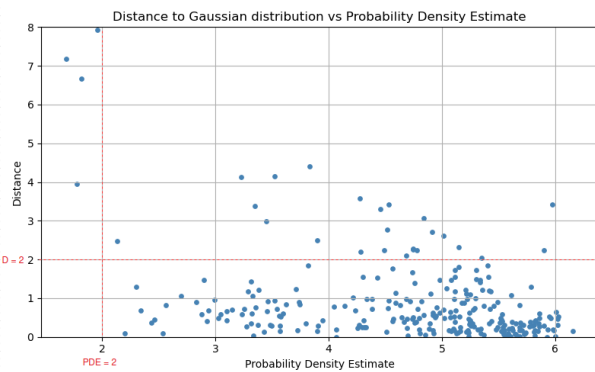


Fig. 9: Distance to the predicted motion resistance Gaussian distribution as a function of the PDE evaluation of the input within the training set.

V. CONCLUSION

In this work, we demonstrated the feasibility of predicting terrain motion resistance using exteroceptive sensing through a near-to-far learning method. We introduced two predictive approaches: a Gaussian Multi-Layer Perceptron and a Gaussian Process Regressor. Among the two, the Gaussian MLP achieved the best performance in our setting, with a RMSE below 10%, measured between the predicted and actual distribution means. We also introduced a methodology to evaluate the robustness of these models under geometric distribution shifts between the training and deployment terrains, leading to a reliability index during inference. Finally, we validated the predictive performance of both models through real-world field experiments, confirming the quality of motion resistance predicted distribution and their applicability in practical deployment scenarios. Future work will focus on enabling online model updates using sensory feedback, and extend the approach beyond longitudinal resistance for more flexible planning.

REFERENCES

[1] P. V. Borges, T. Peynot, S. Liang, B. Arain, M. Wildie, M. G. Minareci, S. Lichman, G. Samvedi, I. Sa, N. Hudson *et al.*, “A survey on terrain traversability analysis for autonomous ground vehicles:

Methods, sensors, and challenges,” *Field Robotics*, vol. 2, pp. 1567–1627, 2022.

[2] J. Kim and J. Lee, “Traction-energy balancing adaptive control with slip optimization for wheeled robots on rough terrain,” vol. 49, pp. 142–156. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1389041717302565>

[3] J. Chen, J. Frey, R. Zhou, T. Miki, G. Martius, and M. Hutter, “Identifying terrain physical parameters from vision – towards physical-parameter-aware locomotion and navigation.” [Online]. Available: <http://arxiv.org/abs/2408.16567>

[4] B. Yang, Q. Zhang, R. Geng, L. Wang, and M. Liu, “Real-time neural dense elevation mapping for urban terrain with uncertainty estimations,” vol. 8, no. 2. [Online]. Available: <http://arxiv.org/abs/2208.03467>

[5] K. Zhang, Y. Yang, M. Fu, and M. Wang, “Traversability assessment and trajectory planning of unmanned ground vehicles with suspension systems on rough terrain,” vol. 19, no. 20, p. 4372, publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1424-8220/19/20/4372>

[6] J. Zürn, W. Burgard, and A. Valada, “Self-supervised visual terrain classification from unsupervised acoustic feature learning.” [Online]. Available: <http://arxiv.org/abs/1912.03227>

[7] G. Reina and R. Galati, “Slip-based terrain estimation with a skid-steer vehicle,” *Vehicle System Dynamics*, vol. 54, no. 10, pp. 1384–1404, 2016.

[8] M. Elnoor, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, “Pronav: Proprioceptive traversability estimation for legged robot navigation in outdoor environments,” 2024. [Online]. Available: <https://arxiv.org/abs/2307.09754>

[9] J.-F. Tremblay, T. Manderson, A. Noca, G. Dudek, and D. Meger, “Multimodal dynamics modeling for off-road autonomous vehicles,” 2021. [Online]. Available: <https://arxiv.org/abs/2011.11751>

[10] M. Wei and V. Isler, “Predicting energy consumption of ground robots on uneven terrains,” *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 594–601, 2022.

[11] R. Lenain, B. Thuilot, C. Cariou, and P. Martinet, “High accuracy path tracking for vehicles in presence of sliding: Application to farm vehicle automatic guidance for agricultural tasks,” vol. 21, no. 1, pp. 79–97. [Online]. Available: <http://link.springer.com/10.1007/s10514-006-7806-4>

[12] D. LaRocque, W. Guimont-Martin, D.-A. Duclos, P. Giguère, and F. Pomerleau, “Proprioception is all you need,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2024, p. 11686–11693. [Online]. Available: <http://dx.doi.org/10.1109/IROS58592.2024.10801407>

[13] A. Krebs, C. Pradalier, and R. Siegwart, “Adaptive rover behavior based on online empirical evaluation: Rover–terrain interaction and near-to-far learning,” *Journal of Field Robotics*, vol. 27, no. 2, pp. 158–180, 2010.

[14] C. Guo, S. Wang, X. Zhao, and H. Zhang, “Tire-road peak adhesion coefficient estimation based on fusion of vision, lidar and dynamics.” [Online]. Available: <https://www.ssrn.com/abstract=4951581>

[15] S. Zhao, J. Zhang, Y. Jiang, C. He, and J. Han, “Tire-road friction coefficients adaptive estimation through image and vehicle dynamics integration,” vol. 224, p. 112039. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0888327024009373>

[16] M. Wei and V. Isler, “Building energy-cost maps from aerial images and ground robot measurements with semi-supervised deep learning,” vol. 5, no. 4, pp. 5136–5142. [Online]. Available: <https://ieeexplore.ieee.org/document/9131808/>

[17] A. Kurup, S. Kysar, J. Bos, P. Jayakumar, and W. Smith, “Supervised terrain classification with adaptive unsupervised terrain assessment,” vol. 3, no. 5, pp. 2337–2344. [Online]. Available: <https://www.sae.org/content/2021-01-0250>

[18] H. Xue, H. Fu, L. Xiao, Y. Fan, D. Zhao, and B. Dai, “Traversability analysis for autonomous driving in complex environment: A LiDAR-based terrain modeling approach,” vol. 40, no. 7, pp. 1779–1803. [Online]. Available: <http://arxiv.org/abs/2307.02060>

[19] F. Philippe, J. Laconte, P.-J. Lapray, M. Spisser, and J.-P. Lauffenburger, “Collision-aware traversability analysis for autonomous vehicles in the context of agricultural robotics,” 2025. [Online]. Available: <https://arxiv.org/abs/2410.03370>