2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

DRIVE: Data-driven Robot Input Vector Exploration

Dominic Baril[1], Simon-Pierre Deschênes, Luc Coupal, Cyril Goffin,
Julien Lépine, Philippe Giguère, François Pomerleau[1]

***Abstract*— An accurate motion model is a fundamental component of most autonomous navigation systems. While much work has been done on improving model formulation, no standard protocol exists for gathering empirical data required to train models. In this work, we address this issue by proposing Data-driven Robot Input Vector Exploration (DRIVE), a protocol that enables characterizing uncrewed ground vehicles (UGVs) input limits and gathering empirical model training data. We also propose a novel learned slip approach outperforming similar acceleration learning approaches. Our contributions are validated through an extensive experimental evaluation, cumulating over 7 km and 1.8 h of driving data over three distinct UGVs and four terrain types. We show that our protocol offers increased predictive performance over common human-driven data-gathering protocols. Furthermore, our protocol converges with 46 s of training data, almost four times less than the shortest human dataset gathering protocol. We show that the operational limit for our model is reached in extreme slip conditions encountered on surfaced ice. DRIVE is an efficient way of characterizing UGV motion in its operational conditions. Our code and dataset are both available online at this link: https://github.com/norlab-ulaval/DRIVE.**

I. Introduction

The ability to model the motion of uncrewed ground vehicles (UGVs) is fundamental to enabling localization [1], path planning [2] and path following [3]. Poor vehicle-terrain characterization will lead to significant modeling errors, potentially causing system failure [4]. With limited available information and sensory measurements on vehicle and ground properties, generating a reliable UGV motion model remains challenging. For most models, training on empirical data is required to reduce modeling error [5]. This task requires deploying a UGV in its operational environment and manually drive it for an extended period [6]. Since energy consumption and deployment time are critical for various UGV applications, facilitating this task is of high importance. Additionally, standardizing this process could help engineers to ensure that their systems comply with the ISO 34502:2022(E) standard on autonomous navigation.[2]

Most work on UGV motion modeling relies on manual driving to gather a training dataset, with little to no details on the driving protocol. Thus, we propose the *Data-driven Robot Input Vector Exploration (DRIVE)*, a protocol aiming to facilitate and standardize vehicle characterization with respect to the terrain, as illustrated in Figure 1. We start by identifying the true vehicle's input space, differing from the manufacturer's specifications. We then automatically send commands to the UGV to cover the entire true input space. This differs from the common manual driving approach, which tends to cover only forward driving, as shown by the red dots representing our previous work [7]. We show that broad input-space coverage offers significant modeling error reduction compared to narrow coverage. With this dataset, we train a learned vehicle slip model that maps UGV commands to resulting body velocities. The resulting trained parameters vary significantly depending on terrain, as highlighted by the green and blue diamond areas in Figure 1, representing navigation on gravel and snow respectively.

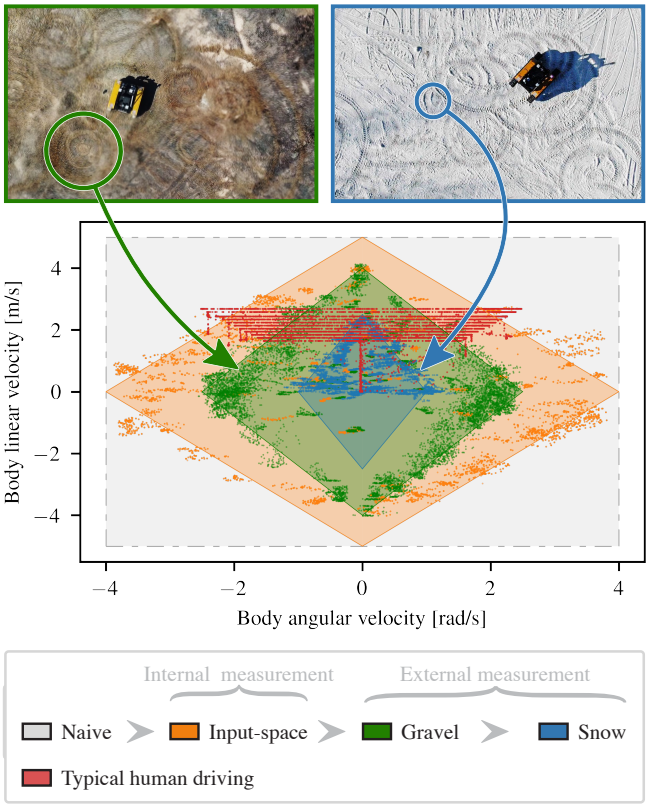

Fig. 1. Vehicle and terrain characterization done through DRIVE. The manufacturer-defined Naive input-space region is drawn in gray. The vehicle's true input-space, characterized through internal measurements, is shown in orange. Typical human driving is shown in red. The resulting body velocities are represented in green for gravel and blue for snow.

*This research was supported by the Fonds de recherche du Québec – Nature et technologies (FRQNT) and by the Natural Sciences and Engineering Research Council of Canada (NSERC) through grant CRDPJ 527642-18 SNOW (Self-driving Navigation Optimized for Winter).

[1]Northern Robotics Laboratory, Université Laval, Quebec City, Quebec, Canada { dominic.baril, francois.pomerleau }@norlab.ulaval.ca

[2]"ISO 34502:2022(E): Road vehicles — Test scenarios for automated driving systems — Scenario-based safety evaluation framework", 2022

The specific contributions of this paper are (i) DRIVE, a standardized UGV characterization and motion data generation protocol allowing motion models to be trained on the entire vehicle input space; (ii) a novel slip-based UGV motion prediction model, leveraging the accuracy of model-based approaches and the minimal system characterization requirement of learning-based approaches. We validate our contributions with an extensive experimental evaluation featuring three distinct UGVs, with weights ranging from $75\,\mathrm{kg}$ to $470\,\mathrm{kg}$, two types of ground interaction (i.e., wheels and tracks) and four different terrain types. Our observations rely on driving data totaling $7\,\mathrm{km}$ and $1.8\,\mathrm{h}$.

II. Related Work

Most vehicle motion modeling approaches can be divided into two distinct categories: model-based and learning-based. Both categories share the requirement of using empirical driving data to train their parameters and reduce modeling errors. For both categories, no standardized protocol exists for training dataset generation.

Model-based approaches can be split into two distinct categories: *kinematics* and *dynamics*. Kinematic models remain the most popular for UGVs due to their low computational complexity and number of parameters to train. For skid-steering mobile robots (SSMRs), Mandow *et al.* [8] reduced the model prediction error by $15\,\%$ compared to the manufacturer's model using a kinematic model empirically identifying vehicle slip and skid. Seegmiller *et al.* [9] proposed a similar additive slip approach, computing slip based on kinematic quantities, yielding a prediction error reduction between $70\,\%$ and $90\,\%$ depending on terrain type, again compared to the manufacturer's model. Bussmann *et al.* [10] extended the experimental validation for additive slip approaches and showed a similar performance for a $900\,\mathrm{m}$ experiment on off-road terrain. On the other hand, dynamic models account for various forces acting on the vehicle's body. Seegmiller *et al.* [4] proposed a multi-body full dynamic motion model with a generic formulation based on vehicle geometry and properties. This work has been extended by Yang *et al.* [11], showing simulation errors of less than $3.4\,\%$ for vehicle slip ratio. While being more accurate than kinematic models, dynamic models require extensive vehicle characterization effort and expertise. For all of the work mentioned above, empirical training data is acquired through a human driving the UGV with little to no guidelines, which motivates our standardized protocol.

Alternatively, **learning-based approaches** have been explored in the literature, yielding more accurate models for extreme UGV motion. In these approaches, part of the prediction is conducted through a nominal model, often represented by a unicycle model, with a module allowing system dynamics to be learned. Gaussian processes (GPs) have become a popular approach to learn system dynamics, both for vehicle slip in off-road driving [12] and tire forces in high-speed road racing [13]. McKinnon *et al.* [14] have proposed a similar approach, replacing GP learning with Bayesian linear regression (BLR). The lower computational complexity of BLR, makes it a more suitable approach for real-time UGV motion prediction. Alternatively, Djeumou *et al.* [15] have proposed a tire-force learning framework allowing autonomous drifting to be performed with $3\,\mathrm{min}$ of driving data. Deep learning has also been explored for motion prediction in off-road terrain. Williams *et al.* [6] have shown the ability to perform aggressive driving when relying on a $30\,\mathrm{min}$ training dataset. For increased resilience to sensor failure, Tremblay *et al.* [16] have proposed a multi-modal learned-dynamics model that leverages the various sensor measurements available for UGVs. Due to the importance of prediction uncertainty in enabling robust control for UGVs [3], this work focuses on BLR, which provides prediction uncertainty estimations [14]. Our novel slip-based BLR model allows us to leverage the minimal requirements of learning-based approaches in terms of system characterization, [14] as well as the improved accuracy of model-based approaches [9]. In this work, the approach of McKinnon *et al.* [14] is used as a comparison point, as it is the closest to our model formulation.

Although both model-based and learning-based approaches require empirical training data, only a few **dataset-gathering protocols** have been published. Voser *et al.* [17] have proposed to maintain a steady forward velocity while slowly increasing angular velocity, enabling generation of a quasi-steady-state empirical dataset. Wang *et al.* [18] have proposed a similar approach with a varying commanded curvature radius to empirically identify the relation between angular velocity and SSMR skid. These approaches only cover a small subset of the vehicle's input space. One can also find large, multimodal datasets allowing to train and evaluate models for off-road and extreme driving [19]. However, such datasets overrepresent forward motion, are limited to a specific UGV and would require new training data for any new vehicle configuration. Manual training data gathering guidelines have been proposed by Williams *et al.* [6], asking the driver to vary his driving style. However, these remain time-consuming and subject to input space coverage bias. We demonstrate that training a motion model with the DRIVE protocol allows increased motion prediction performance and fast training dataset gathering.

III. Methodology and Theory

In this section, we provide details on DRIVE, our automated vehicle characterization and training dataset-gathering protocol. We then describe our proposed slip-based BLR motion model. Due to the limited number of UGVs accessible to us, we focus on SSMRs.

The involved model variables are depicted in Figure 2. We limit the states of the vehicle to planar motion, such that the robot's state ${}^{\mathcal{G}}\boldsymbol{q} = [x, y, \theta]^T$ represents the pose of the vehicle in the global coordinate frame $\mathcal{G}$. The robot's body frame $\mathcal{R}$, has its x and y axis aligned with the vehicle's longitudinal and lateral directions respectively. For most SSMRs, the input vector is defined as $\boldsymbol{u} = [\omega_l, \omega_r]^T$, representing the left and right wheel angular velocities. State propagation, predicting the next state $\boldsymbol{q}_{t+dt}$ based on the

current state $\boldsymbol{q}_t$ and input $\boldsymbol{u}_t$ is computed as follows:

$$^{\mathcal{G}}\boldsymbol{q}_{t+dt} = {}^{\mathcal{G}}\boldsymbol{q}_t + {}^{\mathcal{G}}_{\mathcal{R}}\boldsymbol{T}\left({}^{\mathcal{G}}\theta_t\right){}^{\mathcal{R}}\boldsymbol{v}_t\ dt, \tag{1}$$

$$^{\mathcal{R}}\boldsymbol{v}_t = {}^{\mathcal{R}}\boldsymbol{f}_t(\boldsymbol{u}_t) - {}^{\mathcal{R}}\boldsymbol{g}_t, \tag{2}$$

where $^{\mathcal{R}}\boldsymbol{v}_t$ is the vehicle's translational and rotational body velocity, oriented in the vehicle's inertial frame $\mathcal{R}$, and $^{\mathcal{G}}_{\mathcal{R}}\boldsymbol{T}\left(^{\mathcal{G}}\theta_t\right)$ is a transformation matrix producing a rotation of the robot's angle in the world frame $^{\mathcal{G}}\theta$. The vehicle's body velocity $^{\mathcal{R}}\boldsymbol{v}_t \in \mathbb{R}^3$ is modeled as the commanded velocity $^{\mathcal{R}}\boldsymbol{f}_t(\boldsymbol{u}_t) \in \mathbb{R}^3$, from which we subtract the slip velocity $^{\mathcal{R}}\boldsymbol{g}_t \in \mathbb{R}^3$. The delay between UGV commands is represented by dt. The diamonds in the top-right inset of Figure 2 represent conceptually the set of possible commanded UGV body velocities $\mathcal{J}$ in orange and the set of actually possible body velocities $\mathcal{B}$. Thus, to characterize UGV motion, we require two things: a protocol to gather empirical data and a model to learn vehicle slip.

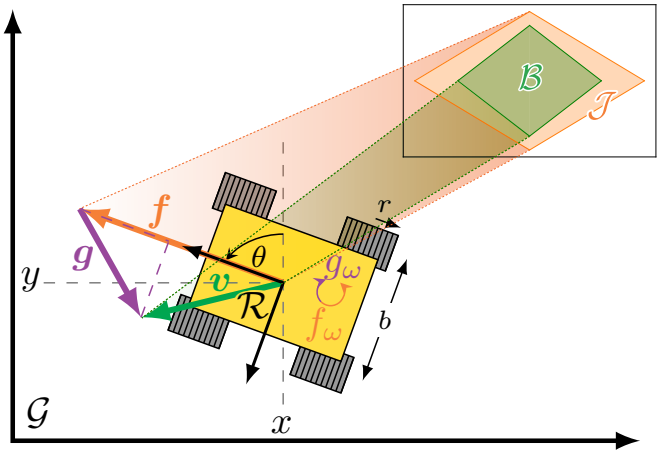

Fig. 2. Top view drawing of a SSMR. In orange is the commanded body velocity $^{\mathcal{R}}\boldsymbol{f}$ and the resulting body velocity $^{\mathcal{R}}\boldsymbol{v}$ is shown in green. The input-space $\mathcal{J}$ is shown in orange and the body velocity space $\mathcal{B}$ is shown in green. The difference between commanded and resulting body velocity is represented as the slip velocity $^{\mathcal{R}}\boldsymbol{g}$ in purple. All represented velocities have an angular component $(\cdot)_\omega$. Robot parameters are the wheel radius r and vehicle width b.

Our DRIVE protocol, described in Section III-A, automates the task of gathering a complete training dataset $\mathcal{D} = \{\boldsymbol{u}, \tilde{\boldsymbol{u}}, \boldsymbol{X}_x, \boldsymbol{g}_x, \boldsymbol{X}_y, \boldsymbol{g}_y, \boldsymbol{X}_\omega, \boldsymbol{g}_\omega\}$, where we concatenate n observations of slip input $\boldsymbol{X} \in \mathbb{R}^{n\times k}$ and slip velocities $\boldsymbol{g} \in \mathbb{R}^n$ for each vehicle dimension. The number of slip inputs per dimension is defined as k. This data is also used to train a powertrain model for each side of the SSMR by minimizing the error between commanded $\tilde{\boldsymbol{u}}$ and measured $\boldsymbol{u}$ vehicle input through the powertrain parameters described in Section III-B.1. We then use BLR to train our slip-based learning model, described in Section III-B.2.

A. Data-driven Robot Input Vector Exploration (DRIVE)

To learn any UGV slip model, we require a training dataset $\mathcal{D}$. Details on the composition of this dataset are given throughout Section III-B. It has been shown that numerous factors impact UGV dynamics, such as vehicle orientation [4], tire saturation [15] and ground properties [7]. However, coverage of the entire spectrum of dynamic features would require a large-scale training dataset, and extracting insights from this dataset would be plagued by the curse of dimensionality [20]. We simplify our problem by enabling stimulation of the entire input space $\mathcal{J}$, represented as the orange diamond in Figure 2. We rely on random, uniform sampling of UGV inputs to ensure broad input space coverage and to trigger dynamic, transitory behavior. Additionally, it was previously shown that a random search on multidimensional space is faster and offers performance similar to that of a grid search [21]. For models accounting for more features, an investigation of more targeted sampling approaches should be conducted. The system requirements to perform our protocol are as follows: (i) a sub-servo system mapping body-level commands to wheel commands; (ii) vehicle acceleration limits to reduce strain on vehicle components; (iii) an accurate localization system, estimating the robot position and velocity in a global frame; (iv) a safety operator to prevent the vehicle from leaving a predefined safe perimeter during the protocol. Lastly, we release our protocol as an open-source package to facilitate replicability.

The first step of DRIVE is to send high body longitudinal velocity commands to the platform in both directions to determine the wheel velocity limits ω_{min} and ω_{max}. This way, we define the UGV's true input-space $\mathcal{J}$ as combinations of left and right commanded wheel velocities, such that $\omega_{min} \leq \omega_l, \omega_r \leq \omega_{max}$ for SSMRs. This input-space $\mathcal{J}$ is shown in terms of body velocities as the orange diamond in Figure 2. Once the input-space limits are defined, we sample input vectors $\boldsymbol{u}_s$ from our calibration distribution, defined as a two-dimensional uniform distribution parametrized by vehicle input limits $\boldsymbol{u}_s \sim \mathcal{U}_2\left(\boldsymbol{\omega}_{min}, \boldsymbol{\omega}_{max}\right)$. Our goal is then to maintain each sampled input for a duration sufficient to gather steady-state motion data. Thus, we define training windows of 2 s, the same duration used in the seminal UGV path following work of Williams *et al.* [6]. We assume the UGV requires one training window to reach any desired wheel velocity. Since the majority of UGV motion is quasi-steady state [17], we define a training interval consisting of one transient training window and two steady training windows, lasting a total of 6 s. An example of two training intervals is shown in Figure 3. We use transient-state windows to train our powertrain model, described in Section III-B.1. We keep both steady-state and transient-state windows for our model training and evaluation to enable both steady-state and transient model training. This procedure is repeated until satisfactory model accuracy is reached. Section IV-C presents an analysis of model accuracy with respect to training dataset driving time, indicating model convergence with 46 sec of driving data. If the user's goal is simply to gather a large dataset, the protocol can run until UGV battery depletion.

B. Dynamics-aware slip-based model

1) Powertrain model: We start by defining a powertrain model to reduce the predicted wheel velocity $\hat{\omega}$ error, the result of which can be seen as both dashed lines in Figure 3. We use the same first-order plus dead time transient response

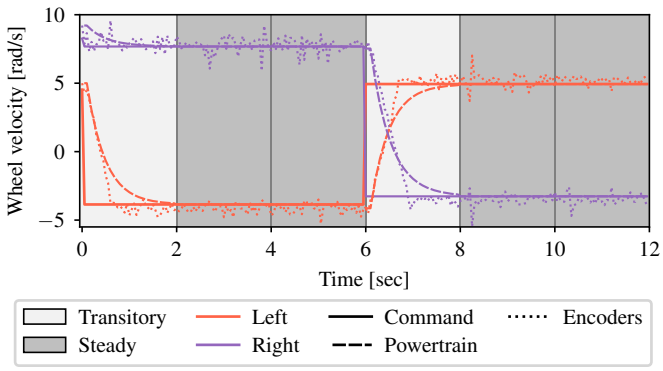

Fig. 3. Commanded, encoder-measured and modeled wheel velocities for both sides of a SSMR during two DRIVE training intervals. The powertrain model is described in Section III-B.1. Each training step consists of one transient-state window (in light gray) and two steady-state windows (in dark gray). Commands and measurements on the x-axis are acquired at a rate of 20 Hz.

model as used by Seegmiller *et al.* [5]:

$$\begin{aligned} \hat{\omega}_t &= \left(e^{\beta}\right)\omega_{t_0} + \left(1 - e^{\beta}\right)\tilde{\omega}_{t-\tau_d}\ , \\ \beta &= \frac{(t-\tau_d)}{\tau_c}\ , \end{aligned} \tag{3}$$

where $\hat{\omega}$, $\tilde{\omega}$ and ω are the predicted, commanded and measured wheel velocities, respectively. We also define the initial time t_0 and prediction horizon at time t. Here, the parameters that require characterization are the time constant τ_c and the time delay τ_d. One should note that these parameters are not considered symmetrical in our protocol and are trained independently for both sides of SSMRs. Thus, our protocol can identify vehicle powertrain asymmetries.

2) Body slip model: Next, we define a model enabling the computation of both the commanded body velocity ${}^{\mathcal{R}}\boldsymbol{f}_t$ and resulting slip velocity ${}^{\mathcal{R}}\boldsymbol{g}_t$ with respect to predicted input $\tilde{\boldsymbol{u}}_t$. For SSMRs, the commanded body velocity ${}^{\mathcal{R}}\boldsymbol{f}_t(\tilde{\boldsymbol{u}}_t)$ can be modeled through the ideal differential-drive model [8] as

$${}^{\mathcal{R}}\boldsymbol{f}_t(\boldsymbol{u}_t) = \begin{bmatrix} f_x \\ f_y \\ f_\omega \end{bmatrix} = r \begin{bmatrix} \frac{1}{2}, \frac{1}{2} \\ 0, 0 \\ -\frac{1}{b}, \frac{1}{b} \end{bmatrix} \begin{bmatrix} \hat{\omega}_{l_t} \\ \hat{\omega}_{r_t} \end{bmatrix}, \tag{4}$$

where r and b are the SSMR's wheel or track sprocket radius and vehicle width, respectively, as shown in Figure 2. We use the estimated wheel velocities through Equation 3 as the input vector $\tilde{\boldsymbol{u}}_t$. We consider slip in each dimension of the vehicle separately ${}^{\mathcal{R}}\boldsymbol{g}_t = [g_x, g_y, g_\omega]^T$, with the form

$$g_t = \boldsymbol{\gamma}^T\, {}^{\mathcal{R}}\boldsymbol{x}_t + \eta, \tag{5}$$

where $\boldsymbol{\gamma} \in \mathbb{R}^k$ are the weights associated with each slip input and $\eta \sim \mathcal{N}(0, \sigma^2)$. We draw inspiration from off-road vehicle dynamics work in the literature to define dynamics-aware basis functions for vehicle slip [9]. As shown by Seegmiller *et al.* [4], the following set of basis functions to estimate vehicle slip shows similar performance as fully dynamic models in off-road terrain. Firstly, for longitudinal slip ${}^{\mathcal{R}}\boldsymbol{g}_x$, we use the vehicle's rolling resistance, proportional to commanded body longitudinal velocity ${}^{\mathcal{R}}\boldsymbol{x}_x = f_x$. Secondly, for lateral slip ${}^{\mathcal{R}}\boldsymbol{g}_y$, we use centrifugal force ${}^{\mathcal{R}}\boldsymbol{x}_y = \psi = (f_x f_\omega)$, proportional to commanded longitudinal and angular velocities. Thirdly, for angular slip ${}^{\mathcal{R}}\boldsymbol{g}_\omega$, we use three distinct slip learning inputs ${}^{\mathcal{R}}\boldsymbol{x}_\omega = [\psi, f_x, f_\omega]$. The first angular slip input is the vehicle's centrifugal force ψ. We then add UGV asymmetry, which can be caused by manufacturing imperfections and mechanical wear, causing an angular velocity error proportional to the commanded longitudinal velocity f_x. Finally, we account for the vehicle's skid, leading to an error between the commanded angular velocity and the actual angular velocity f_ω. It should be noted that the vehicle gravity-dependent parameters, used by Seegmiller *et al.* [9], are missing in this work. The reason is that we simplify our calibration protocol to be executed on planar terrain. The remainder of this section describes how we learn slip for a single dimension, but the process is the same for all dimensions of slip.

We use Bayesian linear regression (BLR) to estimate the values for $\boldsymbol{\gamma}$ and σ^2. For a more in-depth explanation of BLR, refer to the book written by Murphy [22]. It can be shown that the posterior for learned parameters $p(\boldsymbol{\gamma}, \sigma^2|\mathcal{D}_\text{d})$ is distributed according to a Normal Inverse Gamma distribution $\text{NIG}(\boldsymbol{\gamma}, \sigma^2|\boldsymbol{\gamma}, \boldsymbol{K}, a, b)$, where

$$\begin{aligned} \boldsymbol{\gamma} &= \boldsymbol{K}\left(\boldsymbol{K}_0^{-1}\boldsymbol{\gamma}_0 + \boldsymbol{X}^T\boldsymbol{g}\right), \\ \boldsymbol{K} &= (\boldsymbol{K}_0^{-1} + \boldsymbol{X}^T\boldsymbol{X})^{-1}, \\ a &= a_0 + \frac{n}{2}, \\ b &= b_0 + \frac{1}{2}\left(\boldsymbol{\gamma}_0^T\boldsymbol{K}_0^{-1}\boldsymbol{\gamma}_0 + \boldsymbol{g}^T\boldsymbol{g} - \boldsymbol{\gamma}^T\boldsymbol{K}^{-1}\boldsymbol{\gamma}\right), \end{aligned} \tag{6}$$

where the estimated covariance of the distribution is represented by $\boldsymbol{K} \in \mathbb{R}^{k\times k}$. Priors for all parameters are defined by the $(\cdot)_0$ subscript. We define $\mathcal{D}_\text{d} = \{\boldsymbol{X}, \boldsymbol{g}\}$ as a training dataset consisting of vectors of n concatenated observed values for slip inputs $\boldsymbol{X}$ and observed slip velocities $\boldsymbol{g}$ for a specific dimension. The posterior equations can be used to train the BLR slip model for each dimension based on a training dataset $\mathcal{D}_\text{d}$. Once the model is trained, we can predict vehicle slip based on m test inputs $\tilde{\boldsymbol{X}} \in \mathbb{R}^{m\times k}$:

$$p\left(\hat{\boldsymbol{g}}|\tilde{\boldsymbol{X}}, \mathcal{D}_\text{d}\right) = \mathcal{T}\left(\boldsymbol{g}|\boldsymbol{X}\boldsymbol{\gamma}, \frac{b}{a}\left(\boldsymbol{I}_m + \boldsymbol{X}\boldsymbol{K}\boldsymbol{X}^T\right), 2a\right), \tag{7}$$

where $\mathcal{T}$ is a Student's t-distribution and $\hat{\boldsymbol{g}}$ represents a vector of m concatenated predicted slip velocities for a specific direction. We use an uninformative prior to ensure that our protocol requires as little expertise as possible to deploy. This consists of setting $a_0 = b_0 = 0$, $\boldsymbol{\gamma}_0 = \boldsymbol{0}$ and $\boldsymbol{K}_0 = \phi(\boldsymbol{X}^T\boldsymbol{X})^{-1}$ for any positive value ϕ. This allows our slip-based BLR model to be initialized with little knowledge of the UGV except for wheel radius r and vehicle width b.

IV. Results

In this section, we evaluate the improvement of motion prediction accuracy when training models with the DRIVE protocol. We also demonstrate that for off-road navigation of SSMRs, learning vehicle slip based on dynamics-aware basis functions is more accurate than learning on vehicle acceleration. Finally, we analyze the number of training data required to reach convergence with our model.

A. Experimental Setup

We have conducted experiments using three distinct UGV platforms, as shown in Figure 4. First, we evaluated our protocol and model on a *Clearpath Robotics* Warthog on wheels, weighing 470 kg, on gravel-covered terrain and an ice rink. The ice rink was leveled and recently resurfaced, leading to extreme vehicle slip. Next, we tested our protocol and model on smaller platforms, namely a wheeled *Clearpath Robotics* Husky, weighing 75 kg, and a tracked *Superdroid* HD2, weighing 80 kg, both on indoor tile and snow-covered terrain. The Warthog has a top speed of 5 m/s, which is around five times that of the HD2 at 1.2 m/s and of the Husky at 1 m/s. These platforms and terrains were selected to maximize the difference in properties between experiments. Localization ground truth is estimated through point-cloud registration with the iterative closest point (ICP) algorithm to use a common, centimeter-accurate ground truth [23] across all indoor and outdoor experiments. The localization system for the Husky and HD2 robots is described in [24] and for the Warthog in [25]. For every experiment, the recorded data was split into two halves, the training dataset and the evaluation dataset, to enable extensive model evaluation. Our experimental dataset totals over 7 km and 1.8 h of driving data across all platforms and terrain types.

Fig. 4. Three different commercial platforms that were used for the experimental work: a *Superdroid* HD2 (1), a *Clearpath Robotics* Husky (2), and a *Clearpath Robotics* Warthog mounted on wheels (3). The platforms weigh 80 kg, 75 kg and 470 kg, respectively.

B. Protocol performance analysis

First, we define a function to evaluate model prediction performance. While learned models are trained on single-step vehicle slip or acceleration, our goal is to use them to predict vehicle motion over a specific horizon. We train our model on single-step slip velocities to simplify the learning problem. Williams *et al.* [6] showed that this simplification allows sufficient prediction performances for high-speed UGV path following. Thus, we use the multi-step root mean squared error (MRMSE) ϵ to evaluate prediction errors [14], with our localization as ground truth:

$$\epsilon = \frac{1}{h}\sum_{j=1}^{h}\sqrt{({}^{\mathcal{G}}\boldsymbol{q}_j - {}^{\mathcal{G}}\hat{\boldsymbol{q}}_j)^T\boldsymbol{\Sigma}({}^{\mathcal{G}}\boldsymbol{q}_j - {}^{\mathcal{G}}\hat{\boldsymbol{q}}_j)}\ , \tag{8}$$

where h is the prediction window size. We define the measured state as ${}^{\mathcal{G}}\boldsymbol{q}$ and the model-predicted state as ${}^{\mathcal{G}}\hat{\boldsymbol{q}}$. All robots are commanded at 20 Hz and the prediction window length is set at 2 s reflecting the established path following work of Williams *et al.* [6]. We divide this error into translational MRMSE ϵ_T, for which $\boldsymbol{\Sigma} = \text{diag}(1, 1, 0)$ and rotational MRMSE ϵ_R, for which $\boldsymbol{\Sigma} = \text{diag}(0, 0, 1)$.

We compute MRMSE for three distinct training datasets gathering approaches. Table I shows the median prediction improvement for all robots and terrains. Translational and rotational prediction improvements are computed by comparing our DRIVE approach with the angular-focused [18] and linear-focused [6] approaches, respectively. The Warthog on ice experiment stands out due to the high prediction error, which is discussed in Section IV-C. Without considering this experiment, DRIVE offers a mean translation improvement of 31.8 % over the angular-focused approach and a mean rotation improvement of 43.6 % over the linear-focused approach. Figure 5 focuses on the HD2 on snow experiment, showing all three gathering approaches and the resulting prediction error. It can be seen that the DRIVE protocol significantly outperforms other training data-gathering methods, both for translation and rotation.

TABLE I

MEDIAN PREDICTION IMPROVEMENT FOR ALL ROBOTS AND TERRAINS.

Prediction improvement	Husky		HD2		Warthog	
	Tile	Snow	Tile	Snow	Gravel	Ice
Translation (%)	35	10	33	31	50	11
Rotation (%)	18	61	27	51	61	7

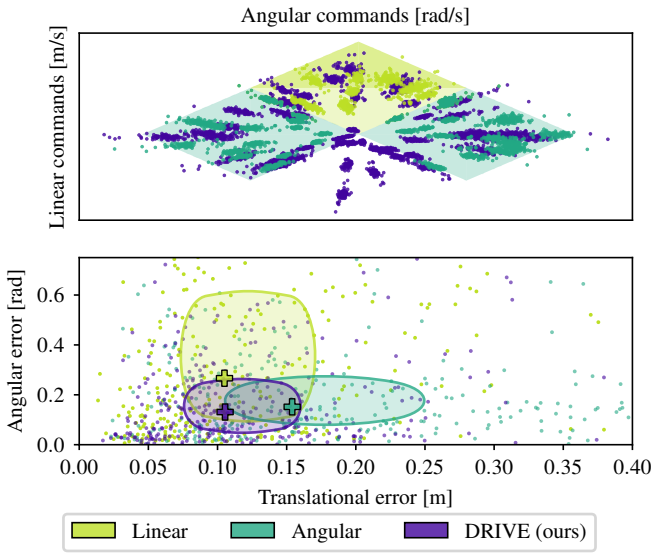

Fig. 5. Data-gathering protocol performance for the HD2 on snow experiment. The top subplot illustrates the three data-gathering methods compared in this work. In yellow, we have the linear-focused method. In teal, we have the angular-focused method. In blue-violet is our DRIVE approach. The crosses and regions on the bottom subplot show the medians and interquartile ranges for translational and angular prediction errors.

C. Slip-based learning predictive performance

In Section III-B, we propose a novel slip-based BLR model to predict UGV motion. Figure 6 shows this model's performance for both translational and rotational prediction, compared to the model proposed by McKinnon *et al.* [14],

which performs BLR on UGV actuator dynamics. We present three distinct datasets, namely HD2 on tile along with the wheeled Warthog on gravel and ice. The rightmost results combine the prediction errors for all experiments conducted in this work. We also show the performance of the model provided by manufacturers (i.e., Naive) and the improvement achieved through powertrain modeling.

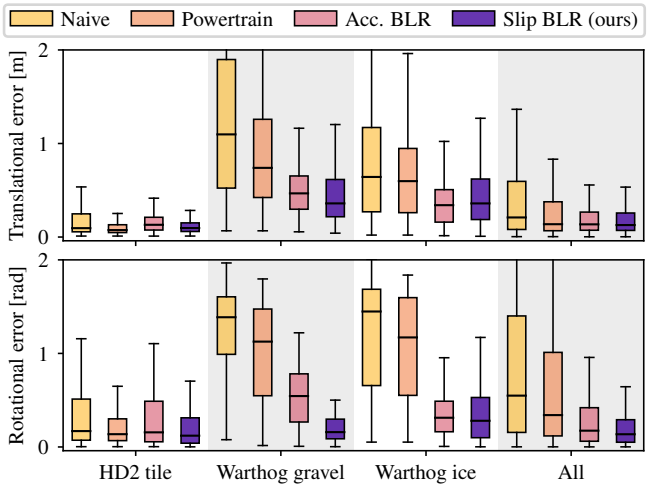

Fig. 6. Translational and rotational prediction errors for all models studied in this work. In yellow is the manufactured-defined naive model, in orange is the powertrain-aware model described in Section III-B.1, in red is the acceleration-based BLR model and in purple is our slip-based BLR model.

When accounting for all datasets, we observe a 34 % decrease in translation prediction error median and a 38 % decrease in rotation prediction error median when comparing the naive model with the powertrain-aware model. Also, our slip-based BLR approach leads to a 22 % decrease in rotation prediction error median and a 6 % decrease in translation prediction error median when compared to acceleration-based BLR. The Warthog in gravel shows the largest improvement between our slip BLR and acceleration BLR, with 71 % in rotation error median and 23 % in translation error median. In contrast, the HD2 on tile experiment shows a performance decrease for acceleration BLR and similar performance for slip BLR when compared to the powertrain model. Indeed, the indoor tile ground already had a low prediction error for the powertrain-aware model. Lastly, the ice rink experiment shows a similar performance between slip and acceleration BLR. This experiment corresponds to extreme slip, similar to a UGV driving over black ice for an extended duration. This result shows the limit of our slip-based BLR model which still performs similarly or better than other models. In this case, dynamic modeling could improve performance. Overall, we conclude that slip-based BLR offers improved performances for rotation prediction and similar performance in translation prediction over acceleration-based BLR, especially for driving at higher velocities on off-road terrains. For SSMRs in particular, rotation motion is the highest source of error due to the complexity of wheel-terrain skidding interactions [7], justifying the significance of our model.

Moreover, generating the training data is time and energy-consuming, which leads us to look for a trade-off between calibration duration and model prediction accuracy. Thus, we evaluated the relationship between training driving time and prediction accuracy. The results are shown in Figure 7. Three distinct experiments are presented, notably Husky on snow, HD2 on tile and Warthog on gravel. No other experiment is shown, to avoid cluttering, but similar results were observed. As specified in Section III-B.2, an uninformative prior is used for every platform, explaining the initially high errors.

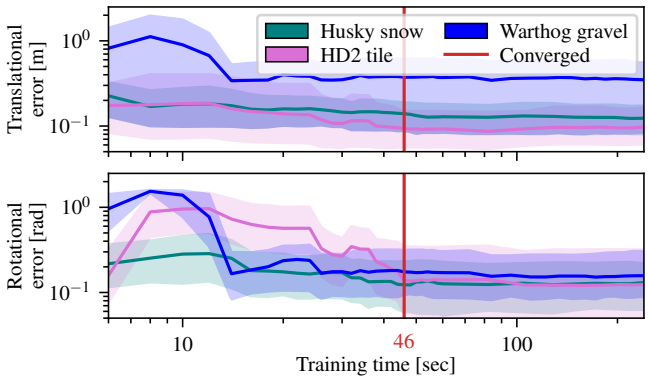

Fig. 7. The relation between training driving time and our slip-based BLR model prediction performance, for translation and rotation. Three datasets are shown, namely the Husky on snow in teal, the HD2 on tile in pink and the wheeled Warthog on gravel, in blue. We highlight at 46 s the converged value with the red line, for which our model converges for all UGVs tested. For all subplots, both axis are in log scale.

As shown by the red vertical line in Figure 7, the prediction accuracy stabilizes after a maximum of 46 s of driving time for all experiments. To compute this time, we evaluated the maximum driving time for which the gradient of translational and rotational prediction error with respect to training driving time was under $0.01\,\mathrm{m/s}$ and $0.01\,\mathrm{rad/s}$ for all shown experiments, indicating all models have converged. Thus, users of the DRIVE protocol SSMRs can expect that the slip-based BLR motion model has converged after 46 s of training data. This is almost four times shorter than 180 s, which is shortest training time reported in the literature [15].

V. Conclusion

In this paper, we propose *Data-driven Robot Input Vector Exploration (DRIVE)*, an automated vehicle characterization and training data generation protocol. We also propose a novel UGV prediction model called slip-based BLR. We show that training our model with our protocol offers improved prediction performances when comparing common training approaches and similar learning-based models. We also show that with our protocol, model convergence is reached with four times less driving time than the shortest similar protocol. Future work would include generalizing our protocol to any vehicle geometry (e.g., Ackermann steering) and adapting our model formulation for complete dynamic models for extreme slip situations such as driving on surfaced ice. Various input sampling strategies should be investigated in conjunction with dynamic models to further minimize training dataset driving time.

References

[1] F. Dümbgen, C. Holmes, and T. D. Barfoot, "Safe and Smooth: Certified Continuous-Time Range-Only Localization," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 1117–1124, Feb. 2023.

[2] R. Takemura and G. Ishigami, "Traversability-based Trajectory Planning with Quasi-Dynamic Vehicle Model in Loose Soil," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2021, pp. 8411–8417.

[3] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 411–444, 2022.

[4] N. Seegmiller and A. Kelly, "High-Fidelity Yet Fast Dynamic Models of Wheeled Mobile Robots," *IEEE Transactions on Robotics (T-RO)*, vol. 32, no. 3, pp. 614–625, Jun. 2016.

[5] N. Seegmiller, F. Rogers-Marcovitz, G. Miller, and A. Kelly, "Vehicle model identification by integrated prediction error minimization," *The International Journal of Robotics Research (IJRR)*, vol. 32, no. 8, pp. 912–931, Jul. 2013.

[6] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-Theoretic Model Predictive Control: Theory and Applications to Autonomous Driving," *IEEE Transactions on Robotics (T-RO)*, vol. 34, no. 6, pp. 1603–1622, Dec. 2018.

[7] D. Baril, V. Grondin, S.-P. Deschenes, J. Laconte, M. Vaidis, V. Kubelka, A. Gallant, P. Giguere, and F. Pomerleau, "Evaluation of Skid-Steering Kinematic Models for Subarctic Environments," in *17th Conference on Computer and Robot Vision (CRV)*, IEEE, May 2020, pp. 198–205.

[8] A. Mandow, J. L. Martinez, J. Morales, J. L. Blanco, A. Garcia-Cerezo, and J. Gonzalez, "Experimental kinematics for wheeled skid-steer mobile robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2007, pp. 1222–1227.

[9] N. Seegmiller and A. Kelly, "Enhanced 3D Kinematic Modeling of Wheeled Mobile Robots," in *Robotics: Science and Systems X (RSS)*, Robotics: Science and Systems Foundation, Jul. 2014.

[10] K. Bussmann, L. Meyer, F. Steidle, and A. Wedler, "Slip Modeling and Estimation for a Planetary Exploration Rover: Experimental Results from Mt. Etna," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2018, pp. 2449–2456.

[11] H. Yang, L. Ding, H. Gao, Z. Wang, Q. Lan, G. Liu, Z. Liu, W. Li, and Z. Deng, "High-Fidelity Dynamic Modeling and Simulation of Planetary Rovers Using Single-Input-Multi-Output Joints With Terrain Property Mapping," *IEEE Transactions on Robotics (T-RO)*, vol. 38, no. 5, pp. 3238–3258, Oct. 2022.

[12] J. Wang, M. T. H. Fader, and J. A. Marshall, "Learning-based model predictive control for improved mobile robot path following using Gaussian processes and feedback linearization," *Journal of Field Robotics (JFR)*, vol. 40, no. 5, pp. 1014–1033, 2023.

[13] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious Model Predictive Control Using Gaussian Process Regression," *IEEE Transactions on Control Systems Technology (T-CST)*, vol. 28, no. 6, pp. 2736–2743, Nov. 2020.

[14] C. D. McKinnon and A. P. Schoellig, "Learn Fast, Forget Slow: Safe Predictive Learning Control for Systems With Unknown and Changing Dynamics Performing Repetitive Tasks," *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 2, pp. 2180–2187, Apr. 2019.

[15] F. Djeumou, J. Y. Goh, U. Topcu, and A. Balachandran, "Autonomous Drifting with 3 Minutes of Data via Learned Tire Models," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2023, pp. 968–974.

[16] J.-F. Tremblay, T. Manderson, A. Noca, G. Dudek, and D. Meger, "Multimodal dynamics modeling for off-road autonomous vehicles," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 1796–1802.

[17] C. Voser, R. Y. Hindiyeh, and J. C. Gerdes, "Analysis and control of high sideslip manoeuvres," *Vehicle System Dynamics*, vol. 48, no. sup1, pp. 317–336, Dec. 2010.

[18] T. Wang, Y. Wu, J. Liang, C. Han, J. Chen, and Q. Zhao, "Analysis and Experimental Kinematics of a Skid-Steering Wheeled Robot Based on a Laser Scanner Sensor," *Sensors*, vol. 15, no. 5, pp. 9681–9702, Apr. 2015.

[19] S. Triest, M. Sivaprakasam, S. J. Wang, W. Wang, A. M. Johnson, and S. Scherer, "TartanDrive: A Large-Scale Dataset for Learning Off-Road Dynamics Models," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 2546–2552.

[20] R. Bellman, "Dynamic Programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.

[21] J. Bergstra and Y. Bengio, "Random Search for Hyper-Parameter Optimization," *Journal of Machine Learning Research (JMLR)*, vol. 13, no. 10, pp. 281–305, 2012.

[22] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.

[23] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing ICP variants on real-world data sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, Apr. 2013.

[24] K. Ebadi, L. Bernreiter, H. Biggie, G. Catt, Y. Chang, A. Chatterjee, C. E. Denniston, S.-P. Deschênes, K. Harlow, S. Khattak, L. Nogueira, M. Palieri, P. Petráček, M. Petrlík, A. Reinke, V. Krátký, S. Zhao, A.-a. Agha-mohammadi, K. Alexis, C. Heckman, K. Khosoussi, N. Kottege, B. Morrell, M. Hutter, F. Pauling, F. Pomerleau, M. Saska, S. Scherer, R. Siegwart, J. L. Williams, and L. Carlone, "Present and Future of SLAM in Extreme Environments: The DARPA SubT Challenge," *IEEE Transactions on Robotics (T-RO)*, vol. 40, pp. 936–959, 2024.

[25] D. Baril, S.-P. Deschênes, O. Gamache, M. Vaidis, D. LaRocque, J. Laconte, V. Kubelka, P. Giguère, and F. Pomerleau, "Kilometer-scale autonomous navigation in subarctic forests: challenges and lessons learned," *Field Robotics*, vol. 2, no. 1, pp. 1628–1660, Mar. 2022.