

# MonoForce: Self-supervised Learning of Physics-informed Model for Predicting Robot-terrain Interaction

Ruslan Agishev<sup>1,2</sup>, Karel Zimmermann<sup>1</sup>, Vladimír Kubelka<sup>2</sup>, Martin Pecka<sup>1</sup> and Tomáš Svoboda<sup>1</sup>

**Abstract**—While autonomous navigation of mobile robots on rigid terrain is a well-explored problem, navigating on deformable terrain such as tall grass or bushes remains a challenge. To address it, we introduce an explainable, physics-aware and end-to-end differentiable model which predicts the outcome of robot-terrain interaction from camera images, both on rigid and non-rigid terrain. The proposed MonoForce model consists of a black-box module which predicts robot-terrain interaction forces from onboard cameras, followed by a white-box module, which transforms these forces and a control signals into predicted trajectories, using only the laws of classical mechanics. The differentiable white-box module allows backpropagating the predicted trajectory errors into the black-box module, serving as a self-supervised loss that measures consistency between the predicted forces and ground-truth trajectories of the robot. Experimental evaluation on a public dataset and our data has shown that while the prediction capabilities are comparable to state-of-the-art algorithms on rigid terrain, MonoForce shows superior accuracy on non-rigid terrain such as tall grass or bushes. To facilitate the reproducibility of our results, we release both the code and datasets.

## I. INTRODUCTION

The ability to predict the outcome of robot-terrain interaction from onboard sensors is essential for many fundamental functionalities of autonomous navigation, such as traversability estimation, planning and control. Over the last few decades, a wide variety of *white-box* [1], [2], [3] and *black-box* [4], [5], [6] motion models have been proposed, each with its strengths and weaknesses. We propose to get the best of both worlds by introducing a *grey-box*, explainable, physics-aware, and end-to-end differentiable model that enables self-supervised learning.

We focus on predicting the robot’s trajectory on both rigid and deformable terrain from a single image captured by an onboard camera. The proposed model is based on a physics-aware simulation conditioned by the terrain image, see Figure 1 for details. The model consists of a *MonoForce predictor*, which delivers robot-terrain reaction forces given the monocular image, state and control, and a *physics engine*, which integrates these forces and estimates the resulting robot’s trajectory. While the physics engine is mostly deterministic, the transformation from image to reaction forces

<sup>1</sup>The authors are with the VRAS group, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic (e-mail: agishrus@fel.cvut.cz; zimmerk@fel.cvut.cz; peckama2@fel.cvut.cz; svobodat@fel.cvut.cz) (Corresponding author: R. Agishev.)

<sup>2</sup>The authors are with the RNP Lab of the AASS Research Centre, Örebro University, Örebro, Sweden (e-mail: vladimir.kubelka@oru.se; agishrus@fel.cvut.cz)

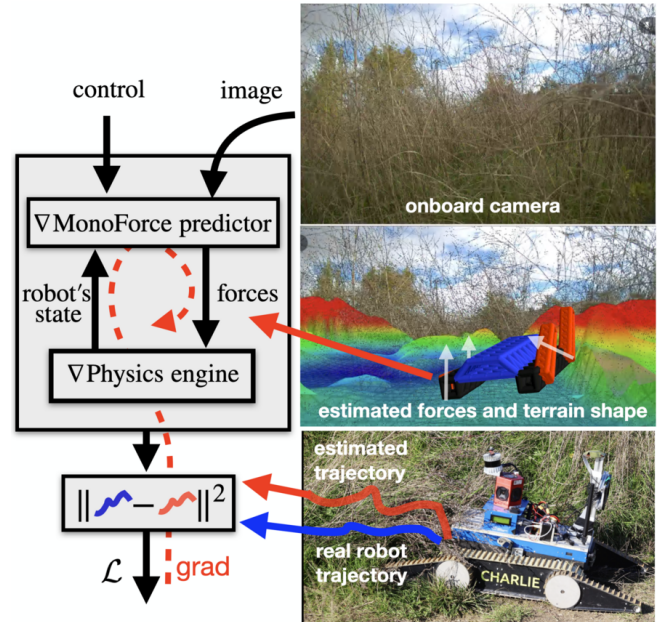


Fig. 1. **Architecture overview:**  $\nabla$ MonoForce predictor delivers robot-terrain reaction forces conditioned by an image from the onboard monocular camera, by the robot state, and by the control input.  $\nabla$ Physics engine integrates these forces to estimate the resulting robot trajectory. Self-supervised loss measures physical consistency between forces and the real robot trajectory through the differentiable physics engine.

is unknown. A natural question emerges: “How can we train the MonoForce predictor?”.

It has been shown that supervised learning of robot-terrain reaction forces is possible for legged robots [6] because the force appears only on its small-sized feet that are in contact with the terrain. Such force is inherently well-localized in the camera frame, and its magnitude and direction can be measured accurately through torque sensors. Therefore, the training data can be easily provided. However, direct measurement of contact forces for wheeled or tracked robots [7] requires highly specialized sensors and is typically hard to achieve in outdoor scenarios. Thus, the contact forces have to be inferred from robot trajectories. Since both the MonoForce predictor and the physics engine are fully differentiable in our implementation, we train the predictor by backpropagating the trajectory error through the physics engine into the predictor parameters. In such a setup, the physics engine plays the role of a self-supervised loss that measures the consistency between predicted forces and ground truth trajectories. The resulting setup then resembles the well-known MonoDepth model [8], which learns to

predict depth from the image by optimizing *color consistency* between the predicted depth and the image delivered by another calibrated camera onboard.

The most straightforward architecture of the motion model would be a black-box, deep-convolutional network that would transform images and control commands into a 3D force field over the state space. However, since there is an inherent training/testing distribution mismatch due to the natural absence of robot-endangering samples in the real training data, a good generalization of the force predictor is crucial. Consequently, we search for architectures that combine black-box deep convolution layers with the white-box laws of classical mechanics and camera geometry. The resulting model builds on the Lift-Splat-Shoot architecture [9] that predicts the terrain properties, such as the shape of the terrain or its rigidity, from a single image and delivers forces that are directly computed from these properties. We emphasize that the resulting model is also directly differentiable with respect to the robot model. Therefore, joint optimization of the robot model, such as its center of gravity or moment of inertia, is also possible. The model is also directly differentiable with respect to control; therefore, its plug-and-play application in state-of-the-art MPC controllers [10] is at hand. **Our main contributions** are as follows:

**Grey-box model:** A novel end-to-end differentiable grey-box model for predicting robot behavior on complex, non-rigid terrains from a monocular image.

**Self-supervised learning:** A physics-aware, self-supervised learning setup that delivers robot-terrain reaction forces and terrain model for tracked and wheeled robots with no contact force or torque sensors.

**Experimental evaluation on non-rigid terrains:** The proposed model outperforms other state-of-the-art methods on non-rigid terrains, such as grass or soft undergrowth that deforms when traversed by the robot.

See the project repository<sup>1</sup> for codes, dataset and supplementary materials such as videos and presentations.

## II. RELATED WORK

Despite being extremely important for predicting robot behavior on complex terrains, most of the so far introduced motion models fall into one of two categories: white-box or black-box models. We argue that this is one of the most substantial reasons for failure in outdoor robotics. This work is a step towards filling the grey-zone gap in order to combine the strengths of both types of models.

White-box models often provide good generalizations if prior assumptions, such as terrain rigidity, are satisfied. For example, Fabian et al. [1] introduce the concept of robot heightmap and a novel iterative geometric method that allows fast 6DOF pose estimation on terrain heightmaps. For skid-steered wheeled robots, Dogru and Marques [2] propose a kinematic model that takes into account both the geometry and weight distribution of the robot. They test their approach on solid surfaces with different slip properties.

However, the unaccounted and hard-to-model influence of non-rigid terrain, which is typically present in vegetation-rich environments, limits the applicability of these white-box models.

Black-box models can be trained on an arbitrary domain; however, the training data are restricted to robot-non-threatening samples. For example, Wellhausen et al. [4] project foothold positions of a walking quadruped robot into the images of its on-board camera. Then, this data is used to train a convolutional network to predict terrain properties relevant for the walking robot. Castro et al. [5] propose a method to learn a proprioceptive-based traversability cost from exteroceptive visual and geometrical data, in a self-supervised manner. They demonstrate the method on two different all-terrain robots and it leads to smoother navigation. The lack of extreme training examples (e.g. robot falling or rolling over) put together with poor generalization outside the training domain lead to the fact that black-box models usually do not perform well in corner-case situations.

This issue can be worked around by providing manual annotations for real data. Palazzo et al. combine manually annotated real data with a simulator using domain transfer [11]. However, manual annotations for traversability are extremely expensive and error-prone. A cheaper source of manual annotations are expert driver trajectories, that can be used to learn to predict costmap using Inverse Reinforcement Learning [12]. Although IRL can theoretically learn from robot-endangering samples (since the driver will avoid them), it is often ill-conditioned and the assumption that the driver provides optimal trajectories with respect to a consistent latent cost is often violated.

Explainable (grey-box) models also start appearing in the robotics literature. Zeng et al. [13] construct a DNN with partially explainable internal structure and guide the network to learn not only plans for driving an autonomous car, but also the intermediate products – detecting other cars, object bounding boxes and motion vectors. However, this method needs lots of manually annotated data and externally provided maps to train the explainable inner parts correctly. We argue that self-supervised methods are more suitable for outdoor driving where it is much more problematic to provide high-quality manual annotations or semantic maps.

Even though manually supervised approaches will always be superior to self-supervised methods (if sufficient amount and quality of manually annotated data is provided), we do not consider them to be a direct competitor to the proposed method because manual annotations do not scale well and cannot be easily transferred among different domains (environments and robots).

Finally, the proposed  $\nabla$ physics module is implemented as a differentiable ODE solver. The existing end-to-end differentiable physics solvers include Neural ODE framework [14] or differentiable simulators such as Google’s brax [15] or Nvidia’s WARP [16]. While the feedforward pass is usually based on Runge-Kutta integrator with a dynamic temporal step, backpropagation through physics is typically tackled by the implicit function theorem. We found out that for our task,

<sup>1</sup><https://github.com/ctu-vras/monoforce>

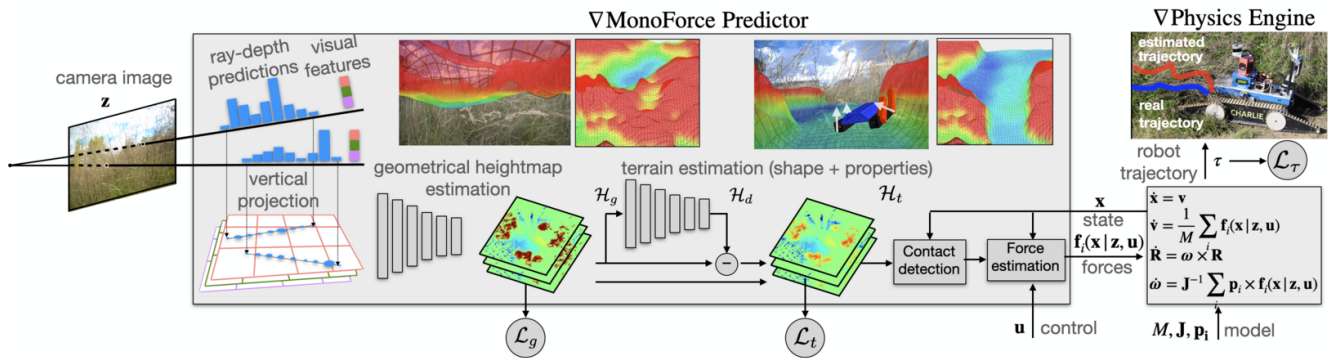


Fig. 2. **Detailed architecture overview of “ $\nabla$ MonoForce predictor” and “ $\nabla$ Physics engine”:** Neural network estimates depth predictions and rich visual features for each pixel ray. Depth-weighted visual features are vertically projected on 2.5D representation, and a geometrical heightmap is estimated. This heightmap is further refined through terrain predictor. It delivers terrain properties such as the heights of the rigid layer of terrain hidden under the vegetation or stiffness and dampening. Given state, control and terrain properties, forces at robot-terrain contacts are estimated. Finally,  $\nabla$ Physics engine integrates these forces to estimate the resulting robot trajectory. We employ two losses: Trajectory loss, which measures the distance between the predicted and real trajectory, and an auxiliary geometrical loss, which measures the distance between the predicted geometrical heightmap and lidar-estimated heightmap.

the custom-built Euler integrator with fixed temporal step delivers similar accuracy while being approximately  $10\times$  faster.

### III. THEORY

A detailed overview of the proposed architecture, which converts images and control commands into robot trajectories, is shown in Figure 2. Given an input image  $\mathbf{z}$  and robot control  $\mathbf{u}$ , the proposed architecture successively estimates geometric heightmap  $\mathcal{H}_g$ , terrain heightmap  $\mathcal{H}_t$ , robot-terrain forces  $\mathbf{f}_i$  and trajectories  $\tau$ . The geometrical map  $\mathcal{H}_g$  is the multichannel 2D array, which contains the heights of the environment observed in the camera in the first channel and visual features in the remaining channels (see Section III-A for details). Similarly, the terrain map  $\mathcal{H}_t$  is the multichannel 2D array, which contains the heights at which the terrain is assumed to start generating forces against the robot in the first channel and visual features in the rest. The intuition is that it corresponds to a not fully flexible layer of terrain (e.g. mud) that is hidden under the fully flexible layer (e.g. small vegetation). It is estimated by subtracting the estimated heightmap decrease  $\mathcal{H}_d$  from the geometrical heightmap. Both maps are augmented by visual features that are successively converted into terrain properties. The part of the architecture that predicts terrain properties is called the *terrain encoder* (see Section III-A) for details). Given the predicted terrain properties, state of the robot and control commands (e.g. track speed or flipper position), the forces  $\mathbf{f}_i$  acting on the robot are computed (see Sections III-C and III-D). Finally, the physics engine solves the robot motion equation and estimates the trajectory corresponding to the delivered forces (see Section III-B). The learning process of the proposed architecture is detailed in Section III-F.

#### A. Terrain encoder

The proposed architecture starts by converting pixels from a 2D image plane into a heightmap with visual features.

Since the camera is calibrated, there is a substantial geometrical prior that connects heightmap cells with the pixels. We incorporate the geometry through the Lift-Splat-Shoot architecture [9]. This architecture uses known camera intrinsic parameters in order to estimate rays corresponding to particular pixels – pixel-rays, Figure 2. For each pixel-ray, the convolutional network then predicts depth probabilities (blue histogram) and visual features (colored vector). Visual features are vertically projected on a virtual heightmap for all possible depths along the corresponding ray. The depth-weighted sum of visual features over each heightmap cell is computed, and resulting multi-channel array is further refined by deep convolutional network in order to estimate the terrain heightmap  $\mathcal{H}_t$ . We call this architecture *terrain encoder*. The terrain encoder can be described as the following function:  $\mathcal{H}_t = h_{\mathbf{w}}(\mathbf{z})$ , where  $\mathbf{w}$  is a vector of the model weights.

#### B. $\nabla$ Physics engine

The higher the complexity of the physical model, the more complicated physical interactions can be modeled, and the lower inductive bias is consequently introduced. However, too complex models require detailed input parameters, such as the complete topological structure of the vegetation in front of the robot, which cannot be reconstructed from a single camera image. Learning to predict unobservable properties is prone to overfitting. To find a good trade-off between the inductive bias and the overfitting, we introduce a collection of physics models with increasing complexity.

We model the robot as a rigid body represented by a set of mass points  $\mathcal{P} = \{(\mathbf{p}_i, m_i) \mid \mathbf{p}_i \in \mathbb{R}^3, m_i \in \mathbb{R}^+, i = 1 \dots N\}$ , where  $\mathbf{p}_i$  denotes coordinates of the  $i$ -th 3D point and  $m_i$  its mass. We employ common 6DOF dynamics of a rigid body [17] are as follows:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{v} & \dot{\mathbf{R}} &= \boldsymbol{\omega} \times \mathbf{R} \\ \dot{\mathbf{v}} &= \frac{1}{M} \sum_i \mathbf{f}_i & \dot{\boldsymbol{\omega}} &= \mathbf{J}^{-1} \sum_i \mathbf{p}_i \times \mathbf{f}_i \end{aligned} \quad (1)$$

where  $\mathbf{x}, \mathbf{v} \in \mathbb{R}^3$  are its position and velocity,  $\mathbf{R} \in \mathcal{SO}_3$

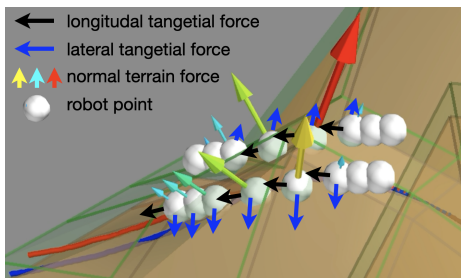


Fig. 3. **Three types of forces** acting on each robot point which is in contact with the terrain. Only track points are visualized for the brevity.

is its rotation represented by  $3 \times 3$  orthogonal matrix and  $\omega \in \mathbb{R}^3$  is its angular velocity. We denote  $\mathbf{f}_i \in \mathbb{R}^3$  all external forces acting on  $i$ -th mass point. Although the forces and the trajectory are assumed to be a function of time, we drop the time index for simplicity. Total mass  $M = \sum_i m_i \in \mathbb{R}^+$  and moment of inertia  $\mathbf{J} \in \mathbb{R}^{3 \times 3}$  are assumed to be known static parameters of the rigid body since they can be identified independently in laboratory conditions. Note that the proposed framework allows backpropagating the gradient with respect to these quantities, too, which makes them jointly learnable with the rest of the architecture. The trajectory of the rigid body is the solution of differential equations (1), that can be obtained by any ODE solver for given external forces and initial state (pose and velocities).

When the robot is moving over a terrain, two types of external forces are acting on the point cloud  $\mathcal{P}$  representing its model: (i) gravitational forces and (ii) robot-terrain interaction forces. The former is defined as  $\mathbf{f}_{gi} = [0, 0, -m_i g]^\top$  and acts on all the points of the robot at all times, while the latter is the result of complex physical interactions that are not easy to model explicitly and act only on the points of the robot that are in contact with the terrain. There are three types of robot-terrain interaction forces: (i) normal terrain force that prevents the penetration of the terrain by the robot points, (ii) longitudinal tangential force that generates forward acceleration when the tracks are moving, and (iii) the lateral tangential force that prevents side slippage of the robot; see Figure 3 for details.

### C. Normal reaction forces

One extreme option is to predict the 3D force vectors  $\mathbf{f}_i$  directly by the convolutional network, but we decided to enforce additional prior assumptions to reduce the risk of overfitting. These prior assumptions are based on common intuition from the contact dynamics of flexible objects. In particular, we assume that the magnitude of the force that the terrain exerts on the point  $\mathbf{p}_i \in \mathcal{P}$  increases proportionally to the deformation of the terrain. Consequently, the network does not directly predict the force, but rather predicts the height of the terrain  $h \in \mathbb{R}$  at which the force begins to act on the robot body and the stiffness of the terrain  $e \in \mathbb{R}^+$ . We understand the quantity  $e$  as an equivalent of the spring constant from Hooke's spring model. Given the stiffness of the terrain and the point of the robot that penetrated the

terrain by  $\Delta h$ , the reaction force is calculated  $e \cdot \Delta h$ .

Since such a force, without any additional damping, would lead to an eternal bumping of the robot on the terrain, we also introduce a robot-terrain damping coefficient  $d \in \mathbb{R}^+$ , which similarly reduces the force proportionally to the velocity of the point that is in contact with the terrain.

**Terrain model:** In order to apply these forces, we train the terrain encoder to predict a 4-channel heightmap that contains these quantities (height, stiffness, damping and friction coefficient). The heightmap is in the robot's coordinate frame from the time the input camera image has been captured. During the solution of differential equations (1), we apply forces  $\mathbf{f}_i$  computed from the terrain that is in contact with point  $\mathbf{p}_i = [\mathbf{p}_{xi}, \mathbf{p}_{yi}, \mathbf{p}_{zi}]^T$ . We will always denote the quantities corresponding to point  $\mathbf{p}_i$  as  $h_i, e_i, d_i$ . We considered two possible terrain models: (i) piece-wise constant model, in which the corresponding quantity is uniquely determined by rounding on the heightmap grid and (ii) bilinear model, in which the quantities are estimated by bilinear interpolation from neighboring bins. Since the former exhibited significantly lower stability, less realistic behavior, and diminishing gradients during backpropagation, we stayed with the latter; see Figure 5 for an example of the bilinear terrain model.

**Vertical terrain force model:** The most naive application of the principles outlined above assumes that the terrain induces only vertical forces. The vertical force  $\mathbf{f}_{zi}$  acting on the  $i$ -th point  $\mathbf{p}_i$  of the robot body is then computed as follows:

$$\mathbf{f}_{zi} = -m_i g + \begin{cases} e_i(h_i - \mathbf{p}_{zi}) - d_i \dot{\mathbf{p}}_{zi} & \text{if } \mathbf{p}_{zi} \leq h_i \\ 0 & \text{if } \mathbf{p}_{zi} > h_i \end{cases}$$

See Figure 4 for a visual sketch of terrain properties and forces on a simplified 2-point robot body. These vertical forces can influence only the robot's roll, pitch, and height; therefore, its initial velocity in  $(x, y, \text{heading})$ -space remains constant in the ODE solution.

**Normal terrain force model:** The more advanced model applies similar forces in the normal direction  $\mathbf{n}_i$  of the terrain surface, where the  $i$ -th point is in contact with the terrain.

$$\mathbf{f}_i = [0, 0, -m_i g] + \begin{cases} (e_i \Delta h_i - d_i (\dot{\mathbf{p}}_i^\top \mathbf{n}_i)) \mathbf{n}_i & \text{if } \mathbf{p}_{zi} \leq h_i \\ \mathbf{0} & \text{if } \mathbf{p}_{zi} > h_i \end{cases},$$

where terrain penetration  $\Delta h_i = (h_i - \mathbf{p}_{zi}) \mathbf{n}_{zi}$  is estimated by projecting the vertical distance on the normal direction.

### D. Tangential (driving) forces

Our tracked robots move by moving the main tracks and flippers (auxiliary tracks). The flipper motion is purely kinematic in our model. This means that in a given time instant, their pose is uniquely determined by a 4-dimensional vector of their rotations, and they are treated as a rigid part of the robot.

The motion of the main tracks is transformed into forces tangential to the terrain. In particular, if the point on the main track is in contact with the terrain, two tangential forces are

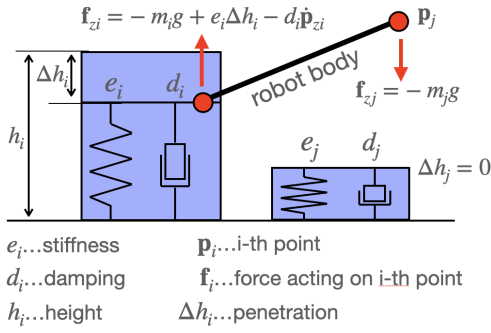


Fig. 4. **Vertical terrain force model:** Simplified 2D sketch demonstrating vertical forces acting on a robot body consisting of two points  $\mathbf{p}_i$  and  $\mathbf{p}_j$ .

generated: longitudinal force  $f_{long}$ , which is parallel with the current orientation of the robot, and lateral force  $f_{lat}$ , which is perpendicular to the current orientation of the robot.

**The longitudinal force** delivers forward acceleration of the robot when robot tracks (either on flippers or on main tracks) are moving. When robot point  $\mathbf{p}_i$ , which belongs to a track, is in contact with terrain with friction coefficient  $\mu_i$ , the magnitude of the resulting longitudinal force is computed as follows:  $f_{long,i} = \mu_i m_i g (\sigma(u - v_x) - 0.5)$ , where  $u$  is the velocity of the track, and  $v = [v_x, v_y]$  is the velocity of the point  $\mathbf{p}_i$  with respect to the terrain transformed into the robot coordinate frame. Consequently, the higher the difference  $u - v_x$ , the higher the longitudinal force. The maximum of this force is given by the gravitation force and the friction coefficient.

**The lateral force** avoids the side slippage of the robot on the terrain. When robot point  $\mathbf{p}_i$  is in contact with terrain, the magnitude of the resulting lateral force is computed as follows:  $f_{lat,i} = \mu_i m_i g (\sigma(-v_y) - 0.5)$ . Similarly, the higher the side velocity  $v_y$ , the higher the force that prevents this motion. This model can be understood as a simplified Pacejka’s tire-road model [18] that is popular for modeling tire-road interactions.

### E. Implementation

**Differentiable implementation of ODE solver:** The overall learning behavior (speed and convergence) is mainly determined by (i) the way the differentiable ODE solver is implemented and (ii) the horizon over which the difference between trajectories is optimized (i.e., the distance of control waypoints from the initial state). We denote current state of the robot as  $\mathbf{s} = [\mathbf{x}, \mathbf{v}, \mathbf{R}, \boldsymbol{\omega}]$  and waypoint  $\mathbf{u}$  representing the control signal. We implemented the forward model  $\dot{\mathbf{s}} = g(\mathbf{s}, \mathbf{u})$  corresponding to differential equations (1) in PyTorch. We experimented with two possible implementations of the ODE solver. The first one was based on the Neural ODE framework [14]. This implementation leverages the existing ODE solver available in Scipy. The backward pass is computed by calling the solver on the adjoint ODE constructed through PyTorch’s autodiff functionality. The second implementation directly implements the Euler integration method with a fixed step size in PyTorch that

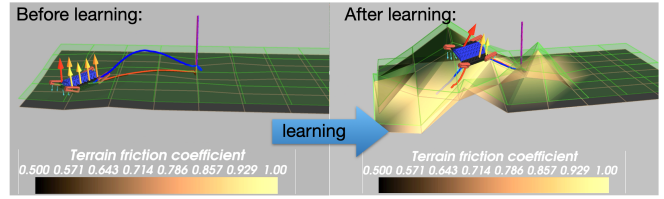


Fig. 5. **Terrain shape computed by backpropagating through  $\nabla$ Physics:** Shape of the terrain (border of the area where terrain forces start to act) outlined by mesh and interaction forces outlined by different colors. The optimized trajectory is in red, and the ground truth trajectory is in blue.

builds the full computational graph of the forward integration and then uses the autograd functionality to get the gradient. Since the former turned out to be significantly faster and achieved comparable accuracy, we employed it for most of our experiments.

### F. Learning and losses

**Trajectory loss:** Since the architecture (Figure 2) is end-to-end differentiable, we can directly learn to predict all these intermediate outputs just from trajectory error

$$\mathcal{L}_\tau = \|\tau - \tau^*\|^2, \quad (2)$$

where  $\tau^*$  is the ground truth trajectory obtained by a common SLAM procedure from onboard measurements only. In order to make the implementation more efficient and explainable, we employ several regularization losses.

**Terrain loss:** Since backpropagating through the physics engine is time-consuming, we first precompute terrain heightmaps  $\mathcal{H}_t^*$  by backpropagating from each ground truth trajectory toward terrain properties, see Figure 5 for an example. These terrain heightmaps serve as an auxiliary ground truth that prevents costly backpropagation through the physics engine in every iteration or if any learning hyperparameters are changed. Consequently, we introduce terrain loss

$$\mathcal{L}_t = \|\mathbf{W}_t \circ (\mathcal{H}_t - \mathcal{H}_t^*)\|^2 + \lambda_d \|\mathcal{H}_d\|^2, \quad (3)$$

where  $\mathbf{W}_t$  denotes the array, which masks out heightmap cells untouched by the ground truth trajectory, and  $\circ$  is the element-wise product. The term  $\|\mathcal{H}_d\|^2$  just provides additional regularization by providing a penalty for decreasing the geometrical heightmap.

**Geometrical loss:** In order to further regularize the learning procedure, we leverage existing lidar measurements in our dataset and construct ground truth geometrical heightmap  $\mathcal{H}_g^*$  from lidar point clouds using the same method as in [19]. Consequently, we introduce geometrical loss

$$\mathcal{L}_g = \|\mathbf{W}_g \circ (\mathcal{H}_g - \mathcal{H}_g^*)\|^2, \quad (4)$$

where  $\mathbf{W}_g$  denotes the array, which masks out heightmap cells measured by the lidar scan.

For the terrain encoder training, we do not require any tedious manual annotations. The geometric heightmap,  $\mathcal{H}_g$ , is directly estimated from a lidar scan, while only a recorded (with the help of SLAM or GPS) robot trajectory  $\tau$  is

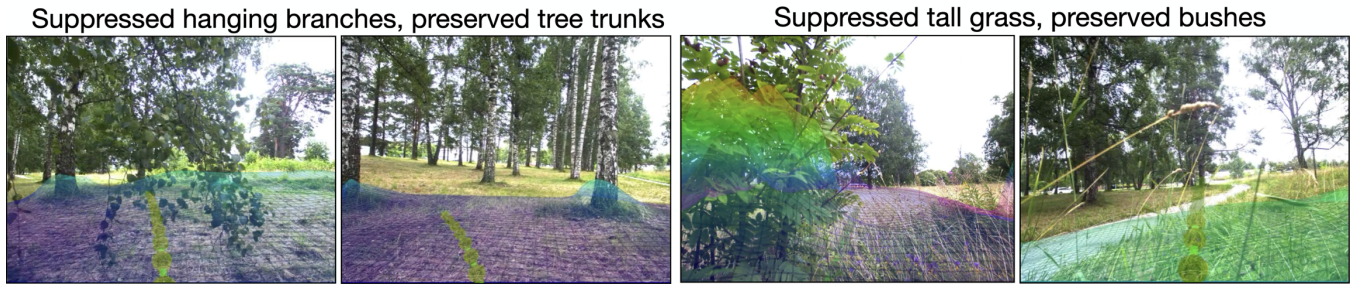


Fig. 6. **Qualitative results:** Predicted terrain heightmap and robot trajectory given given camera images. The terrain heightmap corresponds to the expected heights at which the terrain begins to generate counter-penetration forces. Maximum height limited to 1m for practical reasons.



(a) *Wheeled Husky robot*      (b) *Tracked MARV robot*

Fig. 7. Robot platforms used to collect RobinGas data sequences.

required to obtain the terrain properties  $\mathcal{H}_t$  term. Therefore, the terrain encoder training process is fully self-supervised.

**Prediction horizon:** The length of the horizon employed in learning should correspond to the horizon at which the method is intended to be used. Nevertheless, we observed that both of the previous implementations suffered from diminishing gradients when the horizon was too long. Consequently, we split the training trajectories into 1-second-long chunks. Once trained, this function can be directly usable as a node expansion model in any state-of-the-art planner, MDP controller or terrain traversability predictor (e.g. by thresholding the maximum roll and pitch of the robot in the resulting trajectory). Figure 6 shows qualitative examples of the predicted terrain and robot trajectories.

## IV. EXPERIMENTS

### A. Data

We pretrain the terrain encoder on the large-scale outdoor dataset RELLIS-3D [20]. It is a multimodal dataset collected in an off-road environment containing accurately localized (with RTK GPS) 13,556 lidar scans and corresponding time-synchronized RGB images. Despite the amount and annotation quality of the data provided in the RELLIS-3D sequences, it lacks examples of a robot moving over hills, obstacles, and traversing high grass. To fill this gap in autonomous off-road navigation, we release new data sequences, called RobinGas, containing sensory data from forest and field scenarios recorded with the two robot platforms: *tracked* (MARV robot), *wheeled* (Clearpath Husky robot), Figure 7. The data set contains point cloud scans from Ouster OS0-128, OS0-32 lidars and corresponding RGB images from Basler and IDS cameras installed on the

robots (front, rear, left, and right). For each lidar-images data sample, we additionally record robot localization data that was obtained using ICP SLAM [21]. It showed to be a better choice than GPS that is unreliable near or under forest canopy. With the help of the  $\nabla$ physics module, we obtain height values for each cell of the terrain traversed by the robot following a 10 s-long trajectory. The height values are being optimized in such a way that the trajectory predicted by the  $\nabla$ physics module is as close as possible to the SLAM-generated ground-truth, as described in Section III-A.

After training the MonoForce terrain encoder and the LSS [9] model on RELLIS-3D dataset, we fine-tune them on our RobinGas data sequences. For evaluation, we select the best-performing model on validation part of the RobinGas data in terms of criterion:  $\mathcal{L} = \mathcal{L}_t + \lambda_g \mathcal{L}_g$ , where  $\mathcal{L}_t$  and  $\mathcal{L}_g$  are defined by equations (3), (4),  $\lambda_g$  is a hyperparameter.

### B. Trajectory Estimation Accuracy

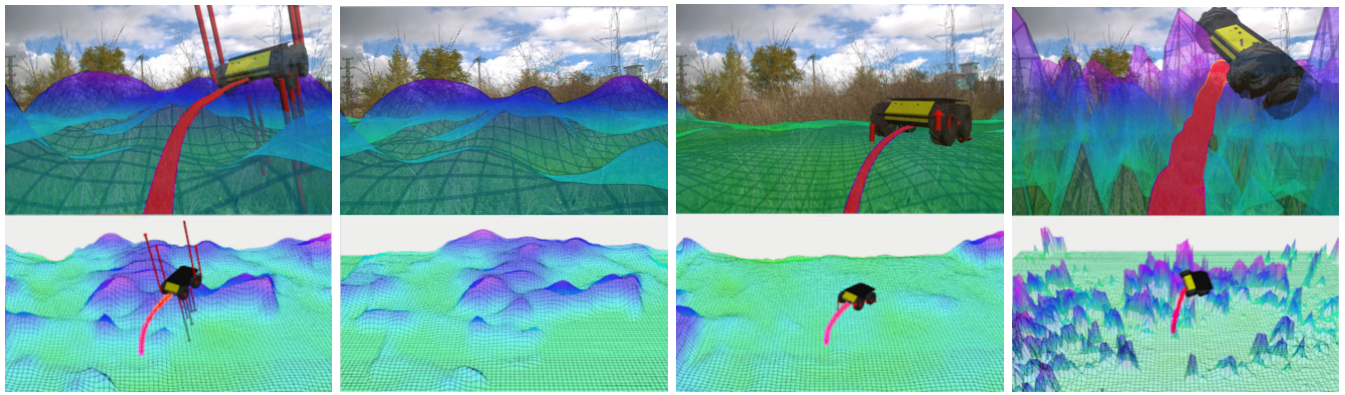
Experiments are conducted according to the pipeline represented in Figure 1. The following terrain encoding methods are compared:

- heightmap interpolation from a lidar scan [19],
- model [22] providing geometrical heightmap  $\mathcal{H}_g$  from a lidar scan,
- LSS model [9] providing  $\mathcal{H}_g$  from RGB images,
- MonoForce (**ours**) predictor providing terrain properties  $\mathcal{H}_t$  from RGB images.

Given the  $\mathcal{H}_t$ , robot model, and a sequence of control commands driving the robot, the  $\nabla$ physics module outputs the predicted trajectory  $\tau = \{(\mathbf{x}_i, \mathbf{R}_i)\}_i$ . This trajectory is then compared to  $\tau^* = \{(\mathbf{x}_i^*, \mathbf{R}_i^*)\}_i$ , the ground-truth trajectory computed by the SLAM pipeline [21]. We compute mean translation and orientation errors  $\Delta \mathbf{x}, \Delta \mathbf{R}$  for robot locations in the time moments corresponding to ground truth poses  $i \in [1 \dots T]$ . The metrics are computed for the trajectory parts that are inside the predicted heightmap region.

$$\Delta \mathbf{x} = \frac{1}{N} \sum_{i=1}^T \|\mathbf{x}_i - \mathbf{x}_i^*\| \quad (5)$$

$$\Delta \mathbf{R} = \frac{1}{N} \sum_{i=1}^T \arccos \frac{\text{tr}(\mathbf{R}_i^T \mathbf{R}_i^*) - 1}{2} \quad (6)$$



(a)  $\mathcal{H}_g$  by LSS [9]: Predicted geometrical heightmap, predicted trajectory, terrain interaction forces. (b)  $\mathcal{H}_d$ : Heightmap difference predicted by (MonoForce, **ours**). (c)  $\mathcal{H}_t = \mathcal{H}_g - \mathcal{H}_d$ : Predicted supporting terrain, predicted trajectory, heightmap, predicted trajectory on the terrain interaction forces (MonoForce, **ours**). (d)  $\mathcal{H}_g$  by [19]: Interpolated lidar supporting terrain, predicted trajectory, heightmap, predicted trajectory on the terrain interaction forces (MonoForce, **ours**).

Fig. 8. **Qualitative comparison of different heightmap estimation methods:**  $\mathcal{H}_g$  estimation with LSS model [9], our MonoForce terrain encoder, and  $\mathcal{H}_g$  interpolation from lidar cloud [19]. Four RGB images are given as input to LSS and MonoForce. *Top row:* front camera image with projected heightmaps and trajectories. *Bottom row:* estimated heightmaps in 3D, predicted trajectories for given control input, and normal component of robot-terrain interaction forces, red arrows at contact points. The length of the arrows denotes the force magnitude. In all scenarios, the Husky robot starts from a stationary state and is given constant linear and angular velocity inputs  $\mathbf{u} = \{(0.5 \text{ m/sec}, 0.13 \text{ rad/sec})\}_t, t \in [0..7 \text{ sec}]$ .

TABLE I  
TRAJECTORY ESTIMATION ACCURACY ON RELIS-3D DATA.

input	terrain encoder	$\tau$ pred.	$\Delta\mathbf{x}$ [m]	$\Delta\mathbf{R}$ [deg]
point cloud	$\mathcal{H}_g$ interp. [19]	$\nabla\text{physics}$	0.14	<b>4.06</b>
point cloud	$\mathcal{H}_g$ pred. [22]	$\nabla\text{physics}$	0.20	4.63
RGB images	$\mathcal{H}_g$ , LSS [9]	$\nabla\text{physics}$	<b>0.12</b>	4.59
RGB images	$\mathcal{H}_t$ , MonoForce	$\nabla\text{physics}$	0.13	4.53

TABLE II  
TRAJECTORY ESTIMATION ACCURACY ON ROBIN GAS DATA.

input	terrain encoder	$\tau$ pred.	$\Delta\mathbf{x}$ [m]	$\Delta\mathbf{R}$ [deg]
point cloud	$\mathcal{H}_g$ interp. [19]	$\nabla\text{physics}$	0.13	4.74
point cloud	$\mathcal{H}_g$ pred. [22]	$\nabla\text{physics}$	0.13	4.32
RGB images	$\mathcal{H}_g$ , LSS [9]	$\nabla\text{physics}$	0.11	5.80
RGB images	$\mathcal{H}_t$ , MonoForce	$\nabla\text{physics}$	<b>0.07</b>	<b>2.62</b>

Table I contains the results for the test part of RELIS-3D data. It could be noticed that the LSS [9] and MonoForce methods are on-par with the methods predicting geometrical heightmap from lidar scans in terms of position estimation metric,  $\Delta\mathbf{x}$ . It could be explained by the fact that the RELIS-3D environments do not have many examples of the robot driving through flexible terrain. So the actual robot footprint supporting terrain does not differ much from its geometrical shape. The robot’s orientation estimation w.r.t. supporting terrain remains challenging ( $\Delta\mathbf{R}$  column in the Table I) for the image-based model in comparison to the geometrical ones.

On the contrary, the results for the RobinGas data indicated in the Table II are more interesting. Both image-based approaches ([9] and MonoForce) outperform the geometrical methods ([19], [22]) for terrain shape estimation on RobinGas data. Qualitative results are shown in Figure 8. They provide comparison between our method output and the two geometrical heightmap estimation approaches: the

TABLE III  
IMPACT OF MONOFORCE COMPONENTS ON TRAJECTORY ACCURACY

method	$\Delta\mathbf{x}$ [m]	$\Delta\mathbf{R}$ [deg]
Geom. $\mathcal{H}$ est. + Terrain est. + $\nabla\text{physics}$	<b>0.07</b>	2.62
Geom. $\mathcal{H}$ est. + $\nabla\text{physics}$	0.11	5.80
Terrain est. + $\nabla\text{physics}$	0.09	<b>2.56</b>

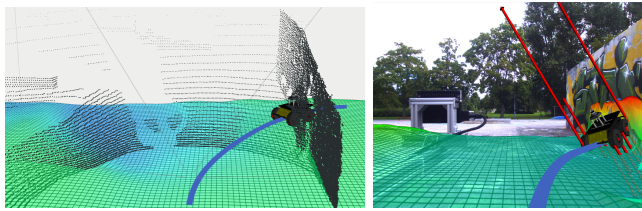
LSS [9] model taking as input RGB images and heightmap interpolation from lidar scan [19]. As seen in the figure, our (MonoForce) system is able to correctly predict supporting terrain for a robot moving through high grass. On the other hand, it is often not possible to predict feasible trajectories using the geometry-only methods, Figure 8(d).

### C. Ablation Study

We conduct the ablation studies to investigate an impact of individual components of the MonoForce pipeline (introduced in Figure 2) on the trajectory prediction accuracy. The results are presented in the Table III. The evaluation is done in the same way as described in the Sec. IV-B on the RobinGas dataset. The top row corresponds to the complete MonoForce architecture results. The second row denotes trajectory accuracy degrades when the terrain properties estimation step is switched-off (or only geometrical  $\mathcal{H}_g$  heightmap is used to predict trajectories). It is not obvious though from the third row results whether providing the model the knowledge about heightmap geometry helps to predict more accurate trajectories. However, a model trained without geometrical information struggles to generalize to out-of-distribution situations not present in the dataset, for example a robot crashing into rigid obstacles, Fig. 9.

## V. CONCLUSION AND FUTURE WORK

In this work, we have presented *MonoForce*, an explainable, physics-aware and end-to-end differentiable model that



(a) LSS model trained without geometrical loss (Table III, 3-rd row) (b) MonoForce result (Table III, 1-st row)

Fig. 9. Out-of-training distribution (robot-endangering) trajectory prediction. *Left*: model trained without geometrical loss is not able to represent a wall in the heightmap. *Right*: MonoForce model predicts the crash correctly.

predicts robot's trajectory from monocular camera images. As a *grey-box* model, it benefits from the end-to-end trainability in different domains while it still retains the determinism of its physics engine and explainability of its estimates of terrain properties. The training process is self-supervised; it only requires the monocular camera images, lidar scans for evaluating the geometrical loss and reference trajectories for the terrain loss. The model learns to recognize non-rigid obstacles similar to those that have been driven over in the reference trajectories, while keeping the understanding and prediction capabilities for obstacles it has not encountered yet. It treats them as rigid obstacles until future data prove otherwise.

In the experiments, we have shown that MonoForce generates accurate geometric and terrain heightmaps that in turn serve as a basis for robot-terrain interaction force and trajectory estimates. These estimates are valid both on rigid and non-rigid terrains. On the rigid terrain that is present in RELLIS-3D dataset, our results are comparable to the lidar-based methods and to the LSS model. The strength of MonoForce becomes apparent on non-rigid terrains, where we outperform the other methods with average 0.07 m translation and 2.62° rotation error in the predicted robot trajectories.

One of the directions for future work is seen in designing different trajectory patterns in the data collection phase that would better cover the estimated terrain heightmap and thus result in more informative terrain loss. We also plan to generalize the definition of supporting terrain to be defined by thresholding the predicted reaction force, which should allow representing a richer variety of terrain properties. It is also possible to define multiple layers of the terrain, each representing an independent physical trait and generating its own reaction force.

### Acknowledgments

This work was co-funded by the Grant Agency of the CTU in Prague under Project SGS24/096/OHK3/2T/13, the European Union under the project Robotics and advanced industrial production (reg. no. CZ.02.01.01/00/22\_008/0004590), the Czech Science Foundation under Project 24-12360S, and the European Union's Horizon Europe Framework Programme under the RaCOON project (ID: 101106906).

### REFERENCES

- [1] S. Fabian, S. Kohlbrecher, and O. Von Stryk, "Pose prediction for mobile ground robots in uneven terrain based on difference of heightmaps," in *2020 IEEE Int. Symp. on Safety, Security, and Rescue Robotics (SSRR)*, 2020, pp. 49–56.
- [2] S. Dogru and L. Marques, "An improved kinematic model for skid-steered wheeled platforms," *Autonomous Robots*, vol. 45, no. 2, pp. 229–243, 2021.
- [3] A. Manoharan, A. Sharma, H. Belsare, K. Pal, K. M. Krishna, and A. K. Singh, "Bi-level trajectory optimization on uneven terrains with differentiable wheel-terrain interaction model," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) under review*, 2024, pp. 1–8.
- [4] L. Wellhausen, A. Dosovitskiy, R. Ranftl, K. Walas, C. Cadena, and M. Hutter, "Where should I walk? Predicting terrain properties from images via self-supervised learning," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, 2019.
- [5] M. Guaman Castro, S. Triest, W. Wang, J. M. Gregory, F. Sanchez, J. G. Rogers III, and S. Scherer, "How does it feel? self-supervised costmap learning for off-road vehicle traversability," in *ICRA*, 2023.
- [6] A. Li, C. Yang, J. Frey, J. Lee, C. Cadena, and M. Hutter, "Seeing through the grass: Semantic pointcloud filter for support surface learning," *arXiv preprint arXiv:2305.07995*, 2023.
- [7] D. Inoue, M. Konyo, K. Ohno, and S. Tadokoro, "Contact points detection for tracked mobile robots using inclination of track chains," in *2008 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 2008, pp. 194–199.
- [8] C. Godard, O. Mac Aodha, M. Firman, and G. J. Brostow, "Digging into self-supervised monocular depth prediction," 2019.
- [9] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [10] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable mpc for end-to-end planning and control," in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [11] S. Palazzo, D. C. Guastella, L. Cantelli, P. Spadaro, F. Rundo, G. Muscato, D. Giordano, and C. Spampinato, "Domain adaptation for outdoor robot traversability estimation from rgb data with safety-preserving loss," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 10014–10021, 2020.
- [12] D. Silver, J. A. Bagnell, and A. Stentz, "Learning from demonstration for autonomous navigation in complex unstructured terrain," *The Int. J. of Robotics Research*, vol. 29, no. 12, pp. 1565–1592, 2010.
- [13] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *2019 IEEE/CVF Conf. on Comput. Vision and Pattern Recognit. (CVPR)*, 2019, pp. 8652–8661.
- [14] M. N. Ricky T. Q. Chen, Brandon Amos, "Learning neural event functions for ordinary differential equations," in *ICLR*, 2021.
- [15] C. D. Freeman, E. Frey, A. Raichuk, S. Girgin, I. Mordatch, and O. Bachem, "Brax - a differentiable physics engine for large scale rigid body simulation," 2021. [Online]. Available: <http://github.com/google/brax>
- [16] M. Macklin, "Warp: A high-performance python framework for gpu simulation and graphics," <https://github.com/nvidia/warp>, March 2022, nVIDIA GPU Technology Conference (GTC).
- [17] K. Bagi, "The contact dynamics method," in *Computational Modeling of Masonry Structures Using the Discrete Element Method*. IGI Global, 2016, pp. 103–122.
- [18] H. Pacejka, *Tire and Vehicle Dynamics*. Elsevier, 2012.
- [19] Y. Wang, H. Li, X. Ning, and Z. Shi, "A new interpolation method in mesh reconstruction from 3d point cloud," in *Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry*, 2011, pp. 235–242.
- [20] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, "Rellis-3d dataset: Data, benchmarks and analysis," 2020.
- [21] 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, 2013. [Online]. Available: <https://doi.org/10.1007/s10514-013-9327-2>
- [22] V. Šalanský, K. Zimmermann, T. Petříček, and T. Svoboda, "Pose consistency KKT-loss for weakly supervised learning of robot-terrain interaction model," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5477–5484, 2021.