

Dom, cars don't fly!—Or do they? In-Air Vehicle Maneuver for High-Speed Off-Road Navigation

Anuj Pokhrel, Aniket Datar, and Xuesu Xiao

Abstract—When pushing the speed limit for aggressive off-road navigation on uneven terrain, it is inevitable that vehicles may become airborne from time to time. During time-sensitive tasks, being able to fly over challenging terrain can also save time, instead of cautiously circumventing or slowly negotiating through. However, most off-road autonomy systems operate under the assumption that the vehicles are always on the ground and therefore limit operational speed. In this paper, we present a novel approach for in-air vehicle maneuver during high-speed off-road navigation. Based on a hybrid forward kinodynamic model using both physics principles and machine learning, our fixed-horizon, sampling-based motion planner ensures accurate vehicle landing poses and their derivatives within a short airborne time window using vehicle throttle and steering commands. We test our approach in extensive in-air experiments both indoors and outdoors, compare it against an error-driven control method, and demonstrate that precise and timely in-air vehicle maneuver is possible through existing ground vehicle controls.

I. INTRODUCTION

Off-road navigation presents various challenges that sharply contrast those encountered in on-road or indoor scenarios. In unstructured off-road environments, robots must detect and avoid obstacles, evaluate the traversability of varied terrain, and continuously adapt to complex vehicle-terrain interactions. Tackling all these challenges is essential to prevent terminal states that can jeopardize the mission and damage the robot, such as vehicle rollover and getting stuck.

One particular challenge of off-road navigation is addressing terrain unevenness. Current state-of-the-art approaches typically rely on perception-based [1] traversability estimation [2]–[5] to avoid uneven terrain or enforce slow speed [6]–[9] to prevent catastrophic failures. By circumventing or slowing down on uneven terrain, these approaches have yet to push the limits of off-road vehicles’ capabilities to quickly traverse through challenging off-road environments.

During time-sensitive off-road missions where achieving the physically feasible maximum speed is necessary, interacting with uneven terrain will cause robots to become airborne from time to time. Additionally, leveraging appropriate terrain structure to take off the vehicle can also efficiently circumvent difficult terrain (in the z direction instead of on the x-y plane) without compromising path length or traversal time. However, after air time, deviations from an appropriate landing pose, depending on the receiving terrain geometry (e.g., nose- or tail-down pitch and sideways roll on a flat terrain, or vice versa), can significantly impact landing safety.

All authors are with the Department of Computer Science, George Mason University {apokhre, adatar, xiao}@gmu.edu

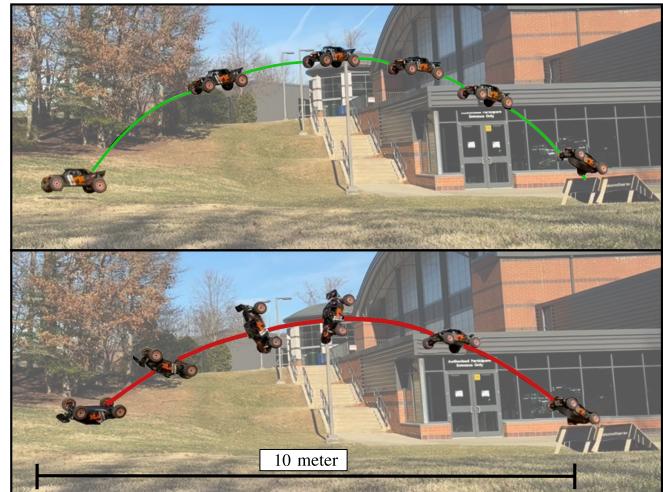


Fig. 1: In-air vehicle maneuvers are critical in ensuring safe vehicle landing during high-speed off-road navigation. Top: Precise and timely maneuvers prepare the robot to land with minimal impact. Bottom: Improper maneuvers cause the robot to land on its back, terminating mission execution and risking vehicle damage.

(see examples in Fig. 1). A critical yet often overlooked aspect of high-speed off-road navigation is the maneuver of a robot during these aerial phases. Despite limited work on planning before losing ground contact [10], to the best of our knowledge, no prior work has investigated ground vehicle in-air maneuver to facilitate safe landing.

To this end, we ask the question *how can we control the in-air attitude of a ground robot in a short amount of airborne time only with existing vehicle controls?* To address this question, we present a novel in-air vehicle attitude planning and control approach by re-purposing existing throttle and steering actions. Our novel forward kinodynamic model uses the inertial and gyroscopic effects of the spinning wheels to derive the robot’s angular accelerations. These accelerations are used in a Newtonian physics model to calculate the robot’s future states. Leveraging our model, we also develop a fixed-horizon, sampling-based motion planner specifically designed for quick in-air goal convergence to prepare the vehicle orientation for proper landing. Our experiment results on a gimbal platform demonstrate the ability of our method to perform precise and timely in-air maneuvers, showing a significant improvement compared to a traditional error-driven approach. Furthermore, our outdoor demonstration

shows our method's ability to generalize under real-world disturbances like air resistance and uneven weight distribution. Our contributions can be summarized as:

- A hybrid dynamics model combining physics principles and data-driven learning for in-air vehicle maneuver driven only by the vehicle throttle and steering actions;
- A fixed-horizon, sampling-based motion planner that converges to an appropriate goal state for safe landing;
- A set of real-world robot experiments on a gimbal platform and outdoors to demonstrate the effectiveness of our kinodynamic model along with our motion planner.

We review related work in modeling off-road robot dynamics and motion planning for off-road navigation.

A. Dynamics Modeling for Off-Road Robots

Accurate dynamics modeling is fundamental to effective robot control, particularly in complex and uncertain off-road environments [11]. Classical wheeled robot dynamics models [12], such as the Ackermann steering and the bicycle model, have demonstrated their utility for conventional ground vehicle navigation through $\mathbb{SE}(2)$ formulations that effectively describe planar motion [13].

Off-road navigation demands models to also address the challenges posed by vertically challenging terrain [6], [14], wheel slippage [15], and ditch crossing [4]. Consequently, data-driven approaches have gained prominence, leveraging sensor data to capture complex dynamics that is difficult to model analytically. Yet, purely data-driven methods can be limited by scarce training data and challenges in generalizing across diverse terrain. To address these issues, recent work has focused on physics-informed learning [16], [17], which integrates prior physics knowledge into the learning process to regularize model behavior. Our work also incorporates fundamental physics principles directly into the model structure, thereby combining the benefits of physical consistency with data-driven flexibility.

When pushing vehicle mobility limits for time-sensitive tasks, high-speed off-road navigation underscores the need for dynamics models operating in $\mathbb{SE}(3)$ as aggressive maneuvers introduce rapid state changes and complex, dynamic vehicle-terrain interactions [18], [19].

II. RELATED WORK

We review related work in modeling off-road robot dynamics and motion planning for off-road navigation.

A. Dynamics Modeling for Off-Road Robots

Accurate dynamics modeling is fundamental to effective robot control, particularly in complex and uncertain off-road environments [11]. Classical wheeled robot dynamics models [12], [20], [21], such as the Ackermann steering and the bicycle model, have demonstrated their utility for conventional ground vehicle navigation through $\mathbb{SE}(2)$ formulations that effectively describe planar motion [13], [22].

Off-road navigation demands models to also address the challenges posed by vertically challenging terrain [7], [8],

[14], [23], wheel slippage [15], and ditch crossing [4]. Consequently, data-driven approaches have gained prominence, leveraging sensor data to capture complex dynamics that is difficult to model analytically. Yet, purely data-driven methods can be limited by scarce training data and challenges in generalizing across diverse terrain. To address these issues, recent work has focused on physics-informed learning [16], [24], which integrates prior physics knowledge into the learning process to regularize model behavior. Our work also incorporates fundamental physics principles directly into the model structure, thereby combining the benefits of physical consistency with data-driven flexibility.

When pushing vehicle mobility limits for time-sensitive tasks, high-speed off-road navigation underscores the need for dynamics models operating in $\mathbb{SE}(3)$ as aggressive maneuvers introduce rapid state changes and complex, dynamic vehicle-terrain interactions [18], [19]. Recognizing that safety is essential for high-speed navigation, CAHSOR [9] restricts the robot's action to the maximum permissible limit for a given terrain using a $\mathbb{SE}(3)$ model to ensure continuous wheel-terrain contact. Furthermore, other models assume airborne dynamics dictated simply by gravity and therefore pre-plan the aerial trajectory before taking off [10], but they lack strategies for controlled landing with in-air control. Unlike any prior work, this paper investigates high-speed in-air vehicle maneuver with existing controls in order to safely land with appropriate vehicle configurations.

B. Motion Planning for Off-road Navigation

Off-road navigation presents unique challenges due to the complex and dynamic nature of the terrain. To address these challenges, a motion planner can take advantage of proprioceptive and exteroceptive sensors to generate traversability estimation [3], [5], [25], [26] and cost maps [1], [27] that can enable online adaptation of control strategies [28] and risk-aware trajectory planning [29], [30]. Despite their theoretical optimality, using search-based planners like A* to generate trajectories can be computationally expensive when addressing the non-convex optimization problem in off-road navigation. Error-driven strategies, like PID controllers, are often employed due to their simplicity and low computational overhead. However, when lacking an explicit robot dynamics model, error-driven methods are ill-equipped to tackle the complex and non-linear aggressive maneuvers, especially considering coupled state dimensions and actuation limits.

In contrast, model-based approaches such as the Model Predictive Path Integral (MPPI) planner [18] have been used extensively in off-road navigation [4], [7], [9], [10]. MPPI leverages sampling to address the non-convex solution space while incorporating robot dynamics and biases its sampling based on the previous best action. With its ability to explore non-convex regions and the use of GPU parallelization [31] to accelerate the evaluation of candidate trajectories, it is well-suited for off-road navigation. Nonetheless, those planners do not consider the limited time horizon for in-air maneuvers which is unique and paramount for our problem

as an airborne off-road ground robot has a strict and short time window to prepare its configuration for safe landing.

III. APPROACH

In this section, we first define the in-air vehicle maneuver problem with airborne kinodynamic modeling. Then, we discuss the physics principles of why and how in-air vehicle maneuver is possible through existing vehicle controls, i.e., throttle and steering, based on the inertial and gyroscopic effects of the spinning wheels. Motivated by the difficulty in accurately and analytically modeling changing quantities purely based on physics, we present PHysics and Learning based model for In-air vehicle maneuver (PHLI, pronounced as “fly”), a precise and efficient hybrid modeling approach. Finally, we introduce our Dom Planner (*named after Dominic Toretto, who makes cars fly. So does our planner.*) for vehicle trajectory planning to reorient the robot from its current configuration to the goal precisely at when a limited airborne time window expires.

A. Problem Formulation

We first formulate a discrete-time forward kinodynamic modeling problem based on a bicycle model, where the subsequent state, $\mathbf{s}_{t+1} \in \mathcal{S}$, is derived from the current state, $\mathbf{s}_t \in \mathcal{S}$, and current action, $\mathbf{a}_t \in \mathcal{A}$, with \mathcal{S} and \mathcal{A} as the state and action spaces respectively. Considering that during aerial phases a conventional ground vehicle does not have control over the three translational components in $\mathbb{SE}(3)$ (which is determined solely by gravity), we include in the vehicle state \mathbf{s}_t a tuple of the three angles in $\mathbb{SO}(3)$, i.e., roll, pitch, and yaw, together with their corresponding angular velocities. We also include rotation per minute of the wheels, rpm, and the front steering angle, ψ , in the state space, i.e., $\mathcal{S} \subset \mathbb{R}^8$. The state at time t is denoted as

$$\mathbf{s}_t = (\text{roll}_t, \dot{\text{roll}}_t, \text{pitch}_t, \dot{\text{pitch}}_t, \text{yaw}_t, \dot{\text{yaw}}_t, \text{rpm}_t, \psi_t) \in \mathcal{S}.$$

The vehicle action, \mathbf{a}_t , comprises the rate of change in wheel rpm (throttle) and in steering angle ψ , which are common controls available for ground robots. The action in the action space, $\mathcal{A} \subset \mathbb{R}^2$, is thus defined as a two-dimensional tuple:

$$\mathbf{a}_t = (\text{rpm}_t, \dot{\psi}_t) \in \mathcal{A}.$$

A forward kinodynamics model is defined as a function, $f_\theta : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, parametrized by θ , such that:

$$\mathbf{s}_{t+1} = f_\theta(\mathbf{s}_t, \mathbf{a}_t). \quad (1)$$

After an aerial phase, the vehicle should land on the ground with an appropriate configuration. Unlike traditional navigation planning problems, in which the goal is to achieve a desired state with certain notion of minimal cost (shortest path, lowest energy, etc.), for our in-air maneuver problem the goal is to achieve a goal configuration precisely when a small airborne time window expires, i.e., the landing time, T . Therefore, the goal is to achieve:

$$\mathbf{s}_T^g = (\text{roll}_T^g, \dot{\text{roll}}_T^g, \text{pitch}_T^g, \dot{\text{pitch}}_T^g, \text{yaw}_T^g, \dot{\text{yaw}}_T^g, \text{rpm}_T^g, \psi_T^g). \quad (2)$$

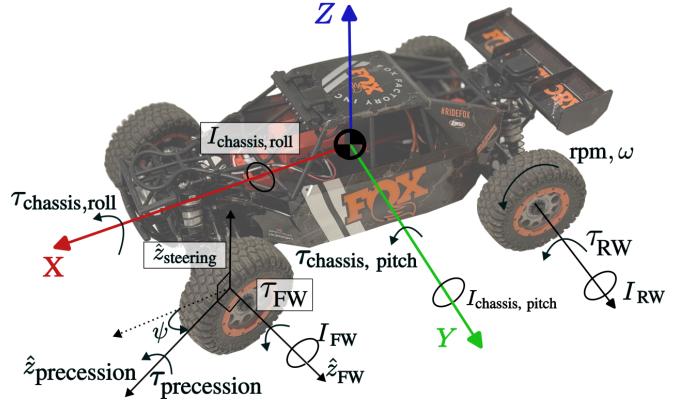


Fig. 2: Simplified Bicycle Model with Torques Acting on the Wheels and Chassis due to Wheel Acceleration and Steering.

Notice that \mathbf{s}_T^g must be achieved precisely at T , not later. If it is achieved at $t < T$, the vehicle needs to assure it maintains the same state at the landing time T . For example, when the landing region is horizontal, roll_T^g and pitch_T^g become zero, while $\dot{\text{roll}}_T^g$ and $\dot{\text{pitch}}_T^g$ should be as close to zero as possible to minimize impact. yaw_T^g and $\dot{\text{yaw}}_T^g$ are often less crucial since they do not cause significant impact during landing. rpm_T^g mostly needs be positive to avoid flipping over the vehicle upon ground contact with forward momentum, and ψ_T^g depends on what the immediate ground maneuver necessary for the vehicle to execute right after landing is, e.g., to quickly swerve to avoid an upcoming obstacle.

Therefore, the in-air vehicle maneuver problem for safe landing can be formulated as

$$\begin{aligned} \mathbf{a}_0^*, \dots, \mathbf{a}_{T-1}^* &= \arg \min_{\mathbf{a}_0, \dots, \mathbf{a}_{T-1}} \sum_{t=0}^T c(\mathbf{s}_t, \mathbf{s}_T^g), \\ \text{s. t. } \mathbf{s}_{t+1} &= f_\theta(\mathbf{s}_t, \mathbf{a}_t), \quad \mathbf{s}_0 \text{ is given, and } \mathbf{s}_T = \mathbf{s}_T^g. \end{aligned} \quad (3)$$

c is a state-wise cost function to be minimized, subject to the constraints by the forward kinodynamics, given initial state, and achieving the goal state at the end of the aerial phase T . Notice the difference of Problem (3) to traditional receding-horizon planning problems, where T is a receding horizon which gradually moves towards the goal. Here T is the end of the entire horizon of the problem, i.e., the landing point.

B. Physics Principles

We discuss the physics principles of why and how in-air vehicle maneuver of \mathbf{s}_{t+1} is possible with $\mathbf{a}_t = (\text{rpm}_t, \dot{\psi}_t)$.

1) *Inertial Effect*: During aerial phases, the only external force acting on the vehicle is gravity (air resistance is overlooked for simplicity). Therefore, with a constant wheel speed, the vehicle maintains its angular momentum by the principle of conservation of angular momentum. All four wheels are rotating at the same speed, because the differentials will not differentiate left and right wheel speed due to a lack of ground contact. Therefore they can be simplified as a front and a rear wheel (FW and RW, bicycle

model) with their angular momentum magnitude given by

$$L_{\text{FW/RW}} = I_{\text{FW/RW}} \cdot \omega,$$

and direction given by either wheel's spin axis, where $I_{\text{FW/RW}}$ is each wheel's moment of inertia along their spin axis and ω is their angular velocity. When the wheels are accelerated by applying a motor torque, equal and opposite reaction torques, τ_{FW} and τ_{RW} , are applied on the chassis by the wheels along their spin axes. These torques are dependent on the change in angular momentum of the respective wheel $dL_{\text{FW/RW}}$:

$$\tau_{\text{FW/RW}} = -\frac{dL_{\text{FW/RW}}}{dt} = -\frac{I_{\text{FW/RW}} \cdot d\omega}{dt} = -I_{\text{FW/RW}} \cdot \frac{2\pi}{60} \cdot \text{rpm}.$$

We can resolve these two torques in individual components, $\tau_{\text{FW/RW, roll}}$ and $\tau_{\text{FW/RW, pitch}}$, aligned with the vehicle's principle axes. The rear wheel's contribution to roll, $\tau_{\text{RW, roll}}$, will always be zero as the rear wheel's spin axis is always perpendicular to the vehicle's roll axis. Its contribution to pitch is $\tau_{\text{RW, pitch}} = \tau_{\text{RW}}$. For the front wheel turned at a steering angle ψ , $\tau_{\text{FW, roll}} = \tau_{\text{FW}} \cdot \sin(\psi)$ and $\tau_{\text{FW, pitch}} = \tau_{\text{FW}} \cdot \cos(\psi)$. $\tau_{\text{FW/RW, roll/pitch}}$ around the front or rear wheel has a resultant torque around the pitch and roll axis passing through the vehicle center of gravity (as the free-flying vehicle body will rotate around the axis with minimum moment of inertia):

$$\begin{aligned} \tau_{\text{chassis, roll}} &= \tau_{\text{FW}} \cdot \sin(\psi) = -I_{\text{FW}} \cdot \sin(\psi) \cdot \frac{2\pi}{60} \cdot \text{rpm}, \\ \tau_{\text{chassis, pitch}} &= \frac{\tau_{\text{FW}} \cdot \cos(\psi) + \tau_{\text{RW}}}{2} \\ &= -(I_{\text{FW}} \cdot \cos(\psi) + I_{\text{RW}}) \cdot \frac{\pi}{60} \cdot \text{rpm}, \end{aligned}$$

considering the roll axis of the front wheel and of the vehicle chassis align for a bicycle model and assuming the front and rear wheels are equal distance away from the vehicle center of gravity. These torque components will cause the vehicle to rotate around the roll and pitch axis based on the chassis moment of inertia around each of these axes, $I_{\text{chassis, roll/pitch}}$:

$$\begin{aligned} \ddot{\text{roll}} &= \frac{\tau_{\text{chassis, roll}}}{I_{\text{chassis, roll}}} = \frac{-I_{\text{FW}} \cdot \sin(\psi) \cdot 2\pi}{I_{\text{chassis, roll}} \cdot 60} \cdot \text{rpm}, \\ \ddot{\text{pitch}} &= \frac{\tau_{\text{chassis, pitch}}}{I_{\text{chassis, pitch}}} = \frac{-(I_{\text{FW}} \cdot \cos(\psi) + I_{\text{RW}}) \cdot \pi}{I_{\text{chassis, pitch}} \cdot 60} \cdot \text{rpm}. \end{aligned} \quad (4)$$

We deduce that roll and pitch are proportional to rpm and dependent on the \sin or \cos of the steering angle ψ due to the inertial effect. See Fig. 2 for a graphical illustration.

2) *Gyroscopic Effect*: Precession describes a phenomenon observed in spinning objects when an external torque is applied perpendicular to the spin axis. This perpendicular torque causes a change in the axis of rotation and therefore angular momentum to rotate towards the axis of the applied torque while keeping the magnitude of the angular momentum constant. This rotation of the spin axis is called precession. For our case of in-air vehicle maneuver, the spinning object is the front wheel, and the external torque perpendicular to the wheel spin axis is from the steering servo to steer the front wheel.

During in-air maneuvers, the front wheel spinning with angular velocity ω has angular momentum

$$L = I_{\text{FW}} \cdot \omega \cdot \hat{z}_{\text{FW}},$$

where \hat{z}_{FW} is the unit vector of the front wheel's spin axis. When steering the front wheel at a rate $\dot{\psi}$, the spin axis rotates towards the steering axis. This steering motion introduces a precession angular velocity

$$\Omega = \dot{\psi} \cdot \hat{z}_{\text{steering}},$$

where $\hat{z}_{\text{steering}}$ is the unit vector along the steering axis. The reaction torque on the vehicle chassis because of precession is determined by

$$\tau_{\text{precession}} = L \times \Omega = I_{\text{FW}} \cdot \omega \cdot \hat{z}_{\text{FW}} \times \dot{\psi} \cdot \hat{z}_{\text{steering}}.$$

Since \hat{z}_{FW} and $\hat{z}_{\text{steering}}$ are orthogonal, the magnitude of their cross product is 1 and the resultant direction is $\hat{z}_{\text{precession}}$ (see Fig. 2). $\tau_{\text{precession}}$ has a component on both roll and pitch axes:

$$\tau_{\text{precession, roll}} = \cos(\psi) \cdot \tau_{\text{precession}} = \cos(\psi) \cdot I_{\text{FW}} \cdot \omega \cdot \dot{\psi},$$

$$\tau_{\text{precession, pitch}} = \sin(\psi) \cdot \tau_{\text{precession}} = \sin(\psi) \cdot I_{\text{FW}} \cdot \omega \cdot \dot{\psi}.$$

The $\tau_{\text{precession, pitch}}$ component is around the front wheel (away from the pitch axis at vehicle center of gravity), hence the resultant torque $\tau_{\text{chassis, pitch}}$ is given by:

$$\tau_{\text{chassis, pitch}} = \frac{\tau_{\text{precession, pitch}}}{2},$$

as the rear wheel does not contribute to precession. Similar to the inertial effect, the direction of $\tau_{\text{precession, roll}}$ aligns with $\tau_{\text{chassis, roll}}$, thus $\tau_{\text{chassis, roll}} = \tau_{\text{precession, roll}}$. Therefore,

$$\begin{aligned} \ddot{\text{roll}} &= \frac{\tau_{\text{chassis, roll}}}{I_{\text{chassis, roll}}} = \frac{\tau_{\text{precession, roll}}}{I_{\text{chassis, roll}}} \\ &= \frac{\cos(\psi) \cdot I_{\text{FW}} \cdot \omega \cdot \dot{\psi}}{I_{\text{chassis, roll}}} = \frac{\cos(\psi) \cdot I_{\text{FW}} \cdot 2\pi}{I_{\text{chassis, roll}} \cdot 60} \cdot \text{rpm} \cdot \dot{\psi}, \\ \ddot{\text{pitch}} &= \frac{\tau_{\text{chassis, pitch}}}{I_{\text{chassis, pitch}}} = \frac{\tau_{\text{precession, pitch}}}{2 \cdot I_{\text{chassis, pitch}}} \\ &= \frac{\sin(\psi) \cdot I_{\text{FW}} \cdot \omega \cdot \dot{\psi}}{2 \cdot I_{\text{chassis, pitch}}} = \frac{\sin(\psi) \cdot I_{\text{FW}} \cdot \pi}{I_{\text{chassis, pitch}} \cdot 60} \cdot \text{rpm} \cdot \dot{\psi}. \end{aligned} \quad (5)$$

We deduce that roll and pitch are proportional to rpm and $\dot{\psi}$ and dependent on ψ due to the gyroscopic effect. The final roll and pitch are a combined effect of Eqns. (4) and (5).

3) *Lack of Yaw Control*: In theory, the vehicle action $\mathbf{a}_t = (\text{rpm}_t, \dot{\psi}_t)$ does not have direct control over the yaw acceleration, $\ddot{\text{yaw}}$. But in practice, $\ddot{\text{yaw}}$ can be affected by system noises such as air resistance and weight distribution. Furthermore, initial yaw position, yaw_0 , and yaw rate, $\dot{\text{yaw}}_0$, will still affect $\ddot{\text{yaw}}$ based on its inertia.

C. PHLI

Calculating roll and pitch requires accurate moment of inertia measurements for the vehicle chassis and for the wheels. These measurements are difficult to analytically derive or precisely measure due to changes in the center of gravity and weight distribution by steering and as a result of

suspension movement. Furthermore, expansion of the tires at high speeds due to centrifugal forces changes the moment of inertia of the wheels as well. Hence, we decompose the forward kinodynamics model $\mathbf{s}_{t+1} = f_\theta(\mathbf{s}_t, \mathbf{a}_t)$ into two parts, g_ϕ and h_ξ , i.e., $\mathbf{s}_{t+1} = h_\xi(g_\phi(\mathbf{s}_t, \mathbf{a}_t), \mathbf{s}_t, \mathbf{a}_t)$, and then employ a data-driven and a physics-based approach to derive g_ϕ and h_ξ respectively.

g_ϕ takes the state and action as input and derives the resultant acceleration of roll, pitch, and yaw

$$\ddot{\text{roll}}_t, \ddot{\text{pitch}}_t, \ddot{\text{yaw}}_t = g_\phi(\mathbf{s}_t, \mathbf{a}_t), \quad (6)$$

where model parameters ϕ are learned in a data-driven manner by minimizing a supervised loss. The ground truth acceleration values are derived from the recorded Inertial Measurement Unit (IMU) data via differentiation over time.

Once g_ϕ is trained, its output angular accelerations are fed into h_ξ , along with the state \mathbf{s}_t and action \mathbf{a}_t . For $\mathbf{s}_{t+1} = h_\xi((\text{roll}_t, \text{pitch}_t, \text{yaw}_t), \mathbf{s}_t, \mathbf{a}_t)$, we have

$$\begin{aligned} \dot{\text{roll}}_{t+1} &= \dot{\text{roll}}_t + \ddot{\text{roll}}_t \cdot dt, \\ \text{roll}_{t+1} &= \text{roll}_t + \dot{\text{roll}}_t \cdot dt + \frac{1}{2} \ddot{\text{roll}}_t \cdot dt^2, \\ \text{rpm}_{t+1} &= \text{rpm}_t + \dot{\text{rpm}}_t \cdot dt, \\ \psi_{t+1} &= \psi_t + \dot{\psi}_t \cdot dt, \end{aligned}$$

with analogous update equations for pitch and yaw. The only parameter for h_ξ is the integration interval dt . This completes PHLI's transition to next state \mathbf{s}_{t+1} given current state \mathbf{s}_t and action \mathbf{a}_t through g_ϕ and h_ξ .

D. Dom Planner

Given the short airtime, we adopt a fixed-horizon, instead of receding-horizon, planning approach, where the horizon is determined by the discretized time remaining until landing. The planner samples action sequences over this fixed horizon and uses the learned PHLI to rollout the corresponding state trajectories. Each trajectory is then evaluated using a cost function. The planner selects the trajectory with the minimal cost, executes its first action, and replans with the updated (and reduced) remaining time. The next planning cycle samples around the best action from the last cycle.

One critical consideration of Dom Planner is that, in addition to general physical actuation limits (e.g., limited motor current and torque causing that the actions can never exceed certain thresholds), the under-actuated vehicle controls of rpm and $\dot{\psi}$ are further constrained due to state-dependent actuation limits:

$$\begin{aligned} \text{rpm}_{\min} - \text{rpm}_t &\leq \text{rpm}_t^{\text{feasible}} \leq \text{rpm}_{\max} - \text{rpm}_t, \\ \psi_{\min} - \psi_t &\leq \dot{\psi}_t^{\text{feasible}} \leq \psi_{\max} - \psi_t, \end{aligned}$$

which means if rpm_t (ψ_t) is at the minimal value, $\text{rpm}_t^{\text{feasible}}$ ($\dot{\psi}_t^{\text{feasible}}$) cannot be negative, and if rpm_t (ψ_t) is at the maximal value, $\text{rpm}_t^{\text{feasible}}$ ($\dot{\psi}_t^{\text{feasible}}$) cannot be positive. This limited range of feasible rpm and $\dot{\psi}$ reduces possible in-air maneuver options and makes simple error-driven controllers, like PID controllers, inappropriate to enable complex in-air vehicle maneuvers.

To enable goal convergence within a short aerial phase, the state-wise cost function defined in the fixed-horizon, instead of receding-horizon, problem in Eqn. (3) is formulated as:

$$c(\mathbf{s}_t, \mathbf{s}_T^g) = \sum_{i=1}^K w_i(t) \cdot c_i(\mathbf{s}_t, \mathbf{s}_T^g), \quad (7)$$

which is a combination of the costs of K state dimensions (in our case, $K = 8$). A key difference from conventional state-wise cost function is that each state dimension is dynamically weighed by a weight term $w_i(t)$ as a function of time t . For example, for initial time steps, it is more important to quickly align the vehicle angles to prepare for a safe landing pose, while the focus will gradually shift to the angular velocities, wheel rpm, and steering components later on to ensure landing with minimal impact on the vehicle.

Additionally, based on the sampled trajectory with the lowest cost, Dom Planner checks its feasibility of reaching the goal state within the time remaining till landing by computing the difference between the final state \mathbf{s}_T and goal state \mathbf{s}_T^g . This feasibility check allows the system to either choose a viable trajectory or, if necessary, select an alternate goal to mitigate landing impact.

IV. IMPLEMENTATIONS

We present the implementation details of our PHLI and Dom Planner onboard a 1/5th-scale vehicle.

A. PHLI Implementations

The learnable function of PHLI, g_ϕ , is a 4-layer multi-layer perceptron that produces a three-dimensional output of predicted angular accelerations, $(\dot{\text{roll}}_{t+1}, \dot{\text{pitch}}_{t+1}, \dot{\text{yaw}}_{t+1})$. g_ϕ is trained to minimize the difference between the predicted angular accelerations and the derived ground truth from IMU. The outputs of g_ϕ is fed into h_ξ to calculate the next state of the vehicle \mathbf{s}_{t+1} .

B. Dom Planner Implementation

Dom Planner's implementation is shown in Algorithm 1. Line 1 first defines Dom Planner's parameters, including integration interval, sample count, actuation limits, sampling range, and PHLI. Actuation limits include rpm rate limit rpm_{\min} and rpm_{\max} (± 5000), rpm limit rpm_{\min} (0) and rpm_{\max} (1980), steering rate limit ψ_{\min} and ψ_{\max} (± 6.5 rad/s), and steering limit ψ_{\min} and ψ_{\max} (± 0.65 rad). Line 2 specifies the algorithm input. Given gravity, time till landing is determined by the initial take-off velocity and terrain geometry. Initial state is the robot state at take-off. Goal state can be defined by the geometry of the receiving terrain. Last best action is the action executed in the last planning cycle. The planning cycle is initiated as soon as the vehicle becomes airborne. The planner calculates the fixed landing horizon $H = \frac{T}{dt}$ by discretizing the time till landing with the interval dt in line 3. For each planning cycle, the algorithm uniformly samples 4000 candidate input pairs $(\text{rpm}_i, \dot{\psi}_i)$ within the ranges determined by σ_{rpm} and $\sigma_{\dot{\psi}}$ centered around the last best action $(\text{rpm}_{\text{best}}, \dot{\psi}_{\text{best}})$ in lines 4 and 5. For each sample (line 6), the planner rolls out a trajectory over the

Algorithm 1 Dom Planner

```

1: Parameters: integration interval  $dt$  (0.2), sample count
    $N$  (4000), actuation limits  $\lambda$ , sampling range  $\sigma_{\text{rpm}}$  (2000)
   and  $\sigma_{\dot{\psi}}$ (0.2), and PHLI  $f_\theta$ 
2: Input: Time till landing  $T$ , initial state  $\mathbf{s}_0$ , goal state  $\mathbf{s}_T^g$ ,
   and last best action  $(\text{rpm}_{\text{best}}, \dot{\psi}_{\text{best}})$ 
3: Compute fixed landing horizon:  $H = \frac{T}{dt}$ 
4:  $\text{rpm}_l, \text{rpm}_h = \text{rpm}_{\text{best}} \pm \sigma_{\text{rpm}}$   $\triangleright$  sampling range for rpm
5:  $\dot{\psi}_l, \dot{\psi}_h = \dot{\psi}_{\text{best}} \pm \sigma_{\dot{\psi}}$   $\triangleright$  sampling range for  $\dot{\psi}$ 
6: for  $(\text{rpm}_i, \dot{\psi}_i), i \in [1, N]$ , sampled from range  $[\text{rpm}_l, \text{rpm}_h]$  and  $[\dot{\psi}_l, \dot{\psi}_h]$  do
7:    $T_i = \{\mathbf{s}_0\}$ 
8:    $U_i = (\text{rpm}_i, \dot{\psi}_i)$ 
9:   for  $t \in [0, H - 1]$  do
10:     $\mathbf{s}_t = \text{CLAMP\_STATE}(\mathbf{s}_t, \lambda)$ 
11:     $\mathbf{a}_t = \text{CLAMP\_ACTION}(\text{rpm}_i, \dot{\psi}_i, \mathbf{s}_t, \lambda)$ 
12:     $\mathbf{s}_{t+1} = f_\theta(\mathbf{s}_t, \mathbf{a}_t)$ 
13:     $T_i.\text{add}(\mathbf{s}_{t+1})$ 
14:   end for
15:    $C_i = \text{CALCULATE\_COST}(T_i, \mathbf{s}_T^g)$ 
16: end for
17:  $T_{\text{best}} = T_{\arg \min_i(C_i)}$   $\triangleright$  minimal-cost trajectory
18:  $U_{\text{best}} = U_{\arg \min_i(C_i)}$   $\triangleright$  best action of this cycle
19: Return  $T_{\text{best}}, U_{\text{best}}$ 

```

horizon H (line 9) using our PHLI f_θ to compute the state transition dynamics (line 12). At every time step, the functions `CLAMP_STATE` and `CLAMP_ACTION` enforce the state and action limits (lines 10-11). After rolling out the full trajectory, `CALCULATE_COST` evaluates each trajectory's cost according to Eqn. (7) (line 15). For $w_i(t)$ in Eqn. (7), we simply use higher and lower weights for the angular position and velocity components respectively in the first half of the rollouts, and vice versa in the second half. We select the best trajectory corresponding to the minimum cost (line 17) and execute the corresponding action (line 18). The best trajectory and action is returned to initialize the last best action $(\text{rpm}_{\text{best}}, \dot{\psi}_{\text{best}})$ for the next planning cycle. We re-plan at 50 Hz and update the time horizon at every cycle based on the remaining time.

C. Robot, Gimbal, and Dataset

We implement PHLI and Dom Planner on a 1/5-scale Losi DBXL E2, 4WD Desert Buggy platform, with a top speed of 80+ km/h. The vehicle is equipped with a 9-DoF IMU, Intel RealSense D435i, NVIDIA Jetson Orin NX for perception and planning, two wheel encoders for front and rear wheels respectively, and an Arduino Mega micro-controller for all low-level actuators and the wheel encoders. For simplicity, we only allow the wheels to rotate forward, i.e., $\text{rpm} \in [0, \text{rpm}_{\text{max}}]$.

To collect a dataset while ensuring the safety of the robot, we construct a 2-axis 1.3 m \times 1 m \times 0.65 m aluminum gimbal platform capable of rotating in the roll and pitch axis. The purpose of the gimbal is to simulate weightlessness and in-

TABLE I: Quantitative Comparison of Dom Planner and PID.

Metric	Dom	PID
TT Error (rad) \downarrow	0.23 \pm 0.13	0.69 ± 0.82
TT Completion Time (sec) \downarrow	12.4 \pm 5.60	23.7 ± 14.1
TT Success Rate \uparrow	5/5	2/5
RSC Time (sec) \downarrow	23.3 \pm 11.2	51.2 ± 16.6
RSC Difference (rad) \downarrow	0.18 \pm 0.41	0.79 ± 0.81
RSC Success Rate \uparrow	5/5	4/5
TGR Time Difference (sec) \downarrow	0.36 ± 1.80	-
TGR State Difference (rad) \downarrow	0.13 ± 0.29	-
SS Correction Time (sec) \downarrow	0.74 \pm 0.74	1.20 ± 0.98
SS Reaction Latency (sec) \downarrow	0.18 \pm 0.13	0.65 ± 0.22
Outdoor Landing Roll (rad) \downarrow	0.08 ± 0.13	-
Outdoor Landing Pitch (rad) \downarrow	0.13 ± 0.21	-

air dynamics. When mounting the robot on the gimbal with robot's center of gravity aligned to the roll and pitch axis of the gimbal, the robot can freely rotate around the roll and pitch axis. Because existing vehicle controls do not have a direct effect on vehicle yaw acceleration and the gimbal introduces significant extra moment of inertial along the yaw axis, we do not include the yaw angle on the gimbal.

While the robot is mounted on the gimbal, the robot is commanded with a diverse range of rpm and $\dot{\psi}$ to capture all possible ways to change the roll and pitch of the vehicle, including behaviors at the extremes of rpm and steering $\dot{\psi}$. The robot configurations during the execution of these commands are recorded and processed to create our tuple of current state, current action, and next state ground truth of the angular accelerations to train g_ϕ . Our dataset includes one hour of gimbal data, split 85/15 for training and validation.

V. EXPERIMENTS

Our proposed method is validated through extensive experiments conducted in both indoor and outdoor environments. For practical reasons, the majority of our experiments are performed indoors on the custom-built gimbal, with additional real-world validation provided by outdoor experiments.

A. Indoor Experiments

We design four test scenarios on the gimbal to evaluate different aspects of our approach. Trajectory Tracking (TT) follows a predefined continuous trajectory. Rapid State Change (RSC) switches quickly to a random state and hold it for two seconds. Timed Goal Reaching (TGR) assesses the ability to reach a goal at a defined time. State Stability (SS) assesses the ability to maintain a goal state when facing external disturbances. We compare our method against a PID controller as the baseline and report the results in Table I. For each experiment, we manually set the initial launch angles, angular velocities, wheel rpm, and steering angle $\dot{\psi}$.

a) Trajectory Tracking Experiment: We define a continuous, closed-loop path as a sequence of goal states that involve continuous variations in both roll and pitch. Fig. 3 (left, top and middle) compares the desired trajectory (blue) with our method (green) and the PID baseline (red).

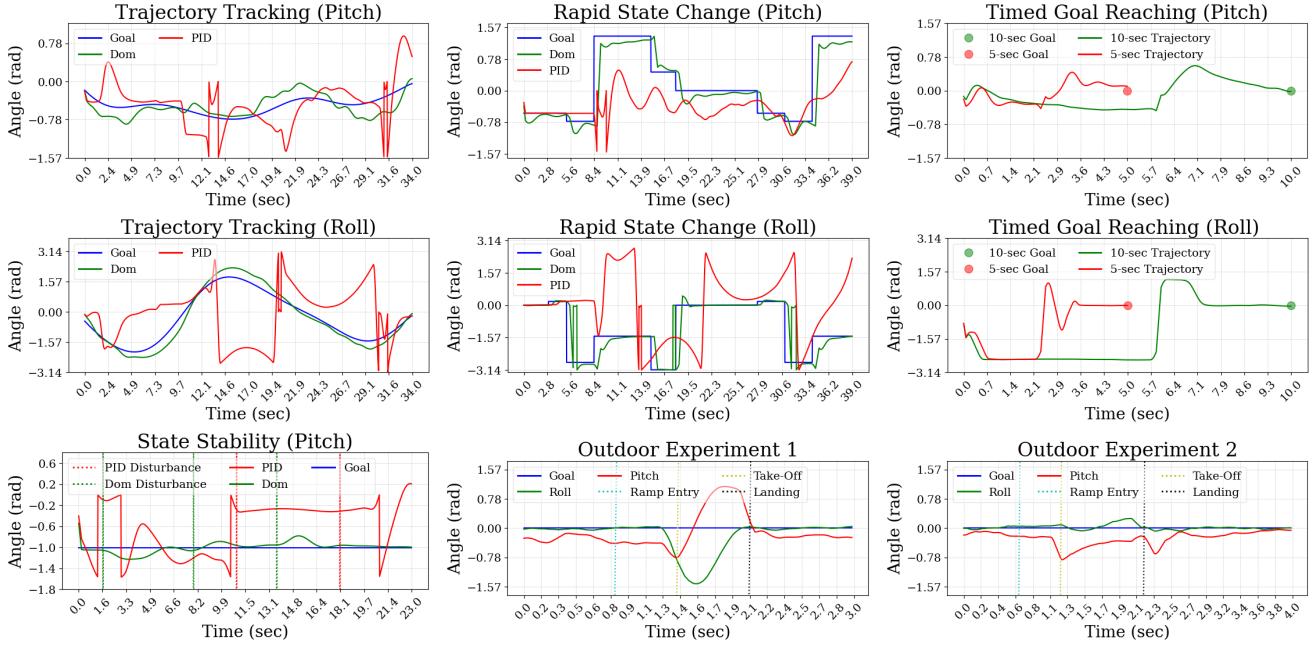


Fig. 3: Qualitative Comparison of Roll and Pitch Angle Achieved by Dom Planner vs PID w.r.t. Goal Angle for Trajectory Tracking, Rapid State Change, Timed Goal Reaching, State Stability, and Outdoor Experiments

Dom Planner tracks the desired trajectory closely, while the baseline PID controller deviates significantly and sometimes exceeds state limits, resulting in failure cases where the robot becomes stuck. We report the mean and standard deviation of the state error, the average time to complete the trajectory (for successful trials), and the overall success rate in Table I.

b) Rapid State Change Experiment: The planner randomly selects a new goal state from a predefined set of goal states. After reaching this state, the robot holds it for two seconds before selecting the next random goal. Fig. 3 (middle, top and middle) show Dom Planner achieves the new state in a more precise and timely manner than the PID baseline. We measure the mean and standard deviation of the time taken to reach each goal and the difference between the achieved state and the goal state along with the overall success rate as reported in Table I.

c) Timed Goal Reaching Experiment: This experiment tests the planner's ability to arrive at a specified goal state exactly at a predefined time. Fig. 3 (right, top and middle) qualitatively shows Dom Planner precisely reaches the goal at the right time (pitch) or reaches the goal earlier and stays there (roll). We report the mean and standard deviation of the difference between the desired and the actual arrival time and the difference between the achieved final state and the goal state in Table I.

d) State Stability Experiment: The planner first brings the robot to a goal state and holds it for ten seconds. During the holding phase, we manually disturb the vehicle by applying downward or upward force. Fig. 3 (left, bottom) demonstrates the Dom Planner reacts to three disturbances (vertical dashed lines) and quickly recovers from them, whereas the PID controller struggles to maintain the steady

state even without its two disturbances. We quantify recovery performance by measuring the mean and standard deviation of the correction time as well as the reaction latency. Table I shows, Dom Planner with PHLI consistently stabilizes the system at the goal state faster and with lower reaction latency compared to the PID baseline.

B. Outdoor Experiments

To validate our method in real-world scenarios, we accelerate the vehicle to about 14 m/s (50+ km/h) and launch it off a ramp set at 45° to the ground. The steering is manually controlled to align the vehicle with the ramp and control is passed to the planner as soon as the vehicle is airborne. We estimate time till landing using the equations of projectile motion and pass these inputs to the planner, which controls the vehicle mid-air for a safe landing at state $s_T^g = (0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 1000, 0.0)$, setting wheel rpm to 1000 to reduce impact at landing. The trials are conducted seven times with different launch poses and velocities. We record the entire trajectory until landing and report the average state upon landing in Table I. Fig. 3 (middle and right, bottom) presents two outdoor trials with their corresponding time step for ramp entry, take-off, and landing annotated as vertical dashed lines. In both cases, our Dom Planner is able to prepare the roll and pitch angles in-air to land the vehicle flat right at the landing time. To ensure the safety of the vehicle, PID is only used for indoor experiments.

C. Discussions

Based on our experimental observations, while the PID controller can reach the goal state under limited configurations, it often becomes stuck upon exceeding state limits—an issue that can lead to catastrophic failures during landing.

In contrast, our Dom Planner with PHLI, which leverages forward simulation until the final state with the help of an accurate dynamics model, achieves successful maneuvers by reliably reaching the goal state within the predefined time. Future work will focus on integrating the Dom Planner and PHLI into a full off-road navigation system.

VI. CONCLUSIONS

In this paper, we present the Dom Planner and PHLI to enable in-air vehicle maneuver for high-speed off-road navigation. Based on the precise in-air forward kinodynamics enabled by the hybrid PHLI model using physics-based and data-driven approaches, Dom Planner is able to accurately and timely maneuver the vehicle roll, pitch, yaw, and their velocities to desired states. Extensive experiments showcase that our method allows existing ground vehicle controls, i.e., throttle and steering, to prepare the vehicle in-air for safe landing during a short airborne period. *So, cars do fly.*

ACKNOWLEDGEMENT

This work has taken place in the RoboTiXX Laboratory at George Mason University. RoboTiXX research is supported by National Science Foundation (NSF, 2350352), Army Research Office (ARO, W911NF2320004, W911NF2420027, W911NF2520011), Air Force Research Laboratory (AFRL), US Air Forces Central (AFCENT), Google DeepMind (GDM), Clearpath Robotics, Raytheon Technologies (RTX), Tangenta, Mason Innovation Exchange (MIX), and Walmart.

REFERENCES

- [1] M. G. Castro, S. Triest, W. Wang, J. M. Gregory, F. Sanchez, J. G. Rogers, and S. Scherer, “How does it feel? self-supervised costmap learning for off-road vehicle traversability,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [2] J. Seo, S. Sim, and I. Shim, “Learning off-road terrain traversability with self-supervisions only,” *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4617–4624, 2023.
- [3] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, “Fast Traversability Estimation for Wild Visual Navigation,” in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [4] T. Han, A. Liu, A. Li, A. Spitzer, G. Shi, and B. Boots, “Model Predictive Control for Aggressive Driving Over Uneven Terrain,” *arXiv preprint arXiv:2311.12284*, 2024.
- [5] C. Pan, A. Datar, A. Pokhrel, M. Nazeri, and X. Xiao, “Traverse the non-traversable: Estimating traversability for wheeled mobility on vertically challenging terrain,” *arXiv preprint arXiv:2409.17479*, 2024.
- [6] A. Datar, C. Pan, and X. Xiao, “Learning to model and plan for wheeled mobility on vertically challenging terrain,” *IEEE Robotics and Automation Letters*, 2023.
- [7] A. Datar, C. Pan, M. Nazeri, A. Pokhrel, and X. Xiao, “Terrain-attentive learning for efficient 6-dof kinodynamic modeling on vertically challenging terrain,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024.
- [8] A. Datar, C. Pan, M. Nazeri, and X. Xiao, “Toward wheeled mobility on vertically challenging terrain: Platforms, datasets, and algorithms,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16 322–16 329.
- [9] A. Pokhrel, A. Datar, M. Nazeri, and X. Xiao, “CAHSOR: Competence-aware high-speed off-road ground navigation in SE (3),” *IEEE Robotics and Automation Letters*, 2024.
- [10] H. Lee, T. Kim, J. Mun, and W. Lee, “Learning Terrain-Aware Kinodynamic Model for Autonomous Off-Road Rally Driving With Model Predictive Path Integral Control,” *IEEE Robotics and Automation Letters*, 2023.
- [11] H. Taghavifar and A. Mardani, “Off-road vehicle dynamics,” *Studies in Systems, Decision and Control*, vol. 70, p. 37, 2017.
- [12] Y. Lu, D. Das, E. Plaku, and X. Xiao, “Multi-goal motion memory,” *arXiv preprint arXiv:2407.11399*, 2024.
- [13] S. Rabiee and J. Biswas, “A friction-based kinematic model for skid-steer wheeled mobile robots,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8563–8569.
- [14] M. Nazeri, A. Datar, A. Pokhrel, C. Pan, G. Warnell, and X. Xiao, “Vertiencoder: Self-supervised kinodynamic representation learning on vertically challenging terrain,” 2024.
- [15] J. Yi, H. Wang, J. Zhang, D. Song, S. Jayasuriya, and J. Liu, “Kinematic modeling and analysis of skid-steered mobile robots with applications to low-cost inertial-measurement-unit-based motion estimation,” *IEEE transactions on robotics*, 2009.
- [16] P. Maheshwari, W. Wang, S. Triest, M. Sivaprakasam, S. Aich, J. G. Rogers III, J. M. Gregory, and S. Scherer, “Piaug—physics informed augmentation for learning vehicle dynamics for off-road navigation,” *arXiv preprint arXiv:2311.00815*, 2023.
- [17] X. Cai, J. Queeney, T. Xu, A. Datar, C. Pan, M. Miller, A. Flather, P. R. Osteen, N. Roy, X. Xiao *et al.*, “Pietra: Physics-informed evidential learning for traversing out-of-distribution terrain,” *IEEE Robotics and Automation Letters*, 2025.
- [18] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, “Aggressive driving with model predictive path integral control,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016.
- [19] S. Talia, M. Schmittle, A. Lambert, A. Spitzer, C. Mavrogiannis, and S. S. Srinivasa, “Demonstrating hound: A low-cost research platform for high-speed off-road underactuated nonholonomic driving,” *Proc. Robot. Sci. Syst.*, 2024.
- [20] T. M. Howard and A. Kelly, “Optimal rough terrain trajectory generation for wheeled mobile robots,” *The International Journal of Robotics Research*, 2007.
- [21] Y. Lu and E. Plaku, “Leveraging single-goal predictions to improve the efficiency of multi-goal motion planning with dynamics,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023.
- [22] H. B. Pacejka and E. Bakker, “The magic formula tyre model,” *Vehicle system dynamics*, vol. 21, no. S1, pp. 1–18, 1992.
- [23] M. Nazeri, A. Pokhrel, A. Card, A. Datar, G. Warnell, and X. Xiao, “Vertiformer: A data-efficient multi-task transformer for off-road robot mobility,” *arXiv preprint arXiv:2502.00543*, 2025.
- [24] T. Kim, H. Lee, and W. Lee, “Physics embedded neural network vehicle model and applications in risk-aware autonomous driving using latent features,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4182–4189.
- [25] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, “Scate: A scalable framework for self-supervised traversability estimation in unstructured environments,” *IEEE Robotics and Automation Letters*, 2023.
- [26] S. Jung, J. Lee, X. Meng, B. Boots, and A. Lambert, “V-strong: Visual self-supervised traversability learning for off-road navigation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024.
- [27] X. Meng, N. Hatch, A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Chen, S. Deng, G. Okopal, D. Fox, B. Boots, and A. Shaban, “TerrainNet: Visual Modeling of Complex Terrain for High-speed, Off-road Navigation,” in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [28] S. Siva, M. Wigness, J. Rogers, and H. Zhang, “Enhancing consistent ground maneuverability by robot adaptation to complex off-road terrains,” in *Conference on Robot Learning*, 2021.
- [29] X. Cai, M. Everett, J. Fink, and J. P. How, “Risk-aware off-road navigation via a learned speed distribution map,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022.
- [30] S. Triest, M. G. Castro, P. Maheshwari, M. Sivaprakasam, W. Wang, and S. Scherer, “Learning risk-aware costmaps via inverse reinforcement learning for off-road navigation,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [31] G. Williams, A. Aldrich, and E. A. Theodorou, “Model predictive path integral control: From theory to parallel computation,” *Journal of Guidance, Control, and Dynamics*, 2017.