

AMCO: Adaptive Multimodal Coupling of Vision and Proprioception for Quadruped Robot Navigation in Outdoor Environments

Mohamed Elnoor, Kasun Weerakoon, Adarsh Jagan Sathyamoorthy, Tianrui Guan, Vignesh Rajagopal, and Dinesh Manocha

Abstract—We present AMCO, a novel navigation method for quadruped robots that adaptively combines vision-based and proprioception-based perception capabilities. Our approach uses three cost maps: general knowledge map; traversability history map; and current proprioception map; which are derived from a robot’s vision and proprioception data, and couples them to obtain a coupled traversability cost map for navigation. The general knowledge map encodes terrains semantically segmented from visual sensing, and represents a terrain’s typically expected traversability. The traversability history map encodes the robot’s recent proprioceptive measurements on a terrain and its semantic segmentation as a cost map. Further, the robot’s present proprioceptive measurement is encoded as a cost map in the current proprioception map. As the general knowledge map and traversability history map rely on semantic segmentation, we evaluate the reliability of the visual sensory data by estimating the brightness and motion blur of input RGB images and accordingly combine the three cost maps to obtain the coupled traversability cost map used for navigation. Leveraging this adaptive coupling, the robot can depend on the most reliable input modality available. Finally, we present a novel planner that selects appropriate gaits and velocities for traversing challenging outdoor environments using the coupled traversability cost map. We demonstrate AMCO’s navigation performance in different real-world outdoor environments and observe 10.8%-34.9% reduction w.r.t. two stability metrics, and up to 50% improvement in terms of success rate compared to current navigation methods.

I. INTRODUCTION

Quadruped robots have been used for autonomous navigation due to their superior dynamics, which enables them to traverse diverse and challenging terrains in outdoor settings [1], [2], [3]. They can be used in applications including agriculture [4], planetary exploration [5], surveillance [6], and disaster response missions [7]. However, guiding these robots across such complex environments requires a comprehensive understanding of the scene, which encompasses how the terrain visually appears using vision-based sensors and how it physically *feels* under the robot’s feet using proprioception.

Traditionally, terrain traversability has been assessed using exteroceptive/visual sensors such as cameras [8], [9] and LiDARs [10], [11], which provide valuable information about the environment. Yet, these approaches often fall short in complex terrains where visual features do not provide sufficient information about the terrain, or when the visual characterization is incorrect. For instance, a terrain that appears to be solid and stable (like soil) but it could deform (is actually mud) under the robot’s weight. This leaves the

Technical report and video can be found at <http://gamma.umd.edu/amco/>
This work was supported in part by ARO Grant W911NF2310352 and Army Cooperative Agreement W911NF2120076.

Authors are with University of Maryland, College Park, MD, USA. {melnoor, kasunw, asathyam, rayguan, vigneshr, dmanocha}@umd.edu




Fig. 1: The navigation trajectories generated by our method, AMCO (trot: brown, crawl: red, amble: pink), ProNav [16] (trot: blue, crawl: purple, amble: light blue), GaNav [9] (orange), Spot’s Inbuilt (yellow), and RFC [20] (trot: green, crawl: light green, amble: light beige) in scenario 4, which contains different types of deformable terrains: (A) water puddle, (B) mud, (C) granular. AMCO, ProNav, and RFC can switch between three gaits (trot, crawl, and amble) during navigation as presented in three different colors. Our method AMCO updates its trajectory as traversing the terrain and uses the coupled vision and proprioception to avoid extremely deformable regions such as A and B. It balances between the trajectory length and maintaining stability to reduce potential robot failures on terrains with varying deformability. Conversely, GaNav [9] takes a longer path by avoiding mud and grass altogether. The other methods take a straight path towards the goal which leads to sinkage and failures.

robot vulnerable to sinkages, slips, and entrapments [5], [12], [13], [14].

On the other hand, proprioception, which can be captured from a legged robot’s internal joint encoders [15], [16], odometry errors [1], robot states [17], and contact vibrations [18], effectively measures the terrain traversability characteristics (softness, bumpiness, etc) at the robot’s present location [19]. Proprioception has previously been used to learn traversability costs for visual navigation [1], [17]. However, it can not be used to predict the terrain traversability ahead of the robot [16] without exteroceptive sensing.

To this end, many approaches have combined proprioceptive and visual inputs using self-supervised learning [1], [15], [17], [21], [22]. Those methods estimate traversability cost maps which only represent terrains’ *typical* traversability. However, they do not account for the potential changes in terrains’ physical properties, such as deformability [23] and trafficability [24], due to weather or other factors. This limitation restricts previous methods from adapting to traversability changes, potentially leading to failures. To

create a generalizable traversability estimation approach for stable navigation, the robot must also consider both the recent proprioceptive history, and the current instant's proprioceptive feedback from the terrain along with its *typical* traversability obtained from visual inputs.

Main Contributions: To this end, we present AMCO (Adaptive Multimodal COupling of vision and proprioception for stable robot navigation), a novel approach that adaptively integrates semantic segmentation and proprioceptive feedback of terrain into a traversability cost map. This multimodal perception representation is integrated with a local planner to select a stable gait and velocity combination for robot navigation. Our approach is inspired by how humans navigate through deformable, and occluded terrains like sand, muddy grass, and bushes, using: 1. our general knowledge of the terrain (e.g. *soil is usually stable to walk on*), 2. recent traversability history (e.g. *soil was wet and felt deformable a few time steps ago in a nearby location*), 3. current experience (e.g. *soil feels more deformable and non-traversable*). Here, both visual features and proprioception information inform our understanding of the terrain and guide our movement. The novel components of our work are:

- We introduce novel formulations of three traversability cost maps using vision and proprioception perception capabilities. A general knowledge map that uses semantic segmentation from an RGB camera to represent the terrain's traversability ahead of the robot. Next, we use the recent history of proprioception signals. We captured the duration of this prospective-informed traversability and its consistency with its semantic traversability class to build a history map. Also, we present a proprioceptive-based traversability cost map, which extrapolates the current robot proprioception to future time steps to robustly predict the traversability along the robot's trajectory.
- An adaptive multi-modal coupling of the three cost maps according to the vision-based sensor reliability. We incorporate a reliability estimation method that assesses the quality of RGB images and weighs the vision-based representations on the coupled maptraversability cost map. This ensures that the navigation decisions are based on the most reliable input modality available. We demonstrate that our adaptive coupling leads to improvements in navigating different types of deformable terrains.
- A navigation algorithm that adapts gait selection and dynamically calculates feasible velocities in real-time based on multi-modal coupled map to ensure stable navigation. We demonstrate our method on Boston Dynamic Spot and evaluate its performance in four unstructured outdoor environments. We observe 10.8%-34.9% reduction in IMU Energy density [25] and up to 50% improvement in navigation success rate.

II. RELATED WORKS

In this section, we discuss the traversability estimation methods for both exteroception and proprioception sensors. After that, we briefly discuss the existing methods for outdoor navigation.

A. Traversability Estimation

Traversability estimation tends to incorporate vision-based approaches that use external sensors and proprioception-based approaches. We review both perception methods in the following sections.

1) *Exteroception-based perception:* Exteroception has been widely used to estimate traversability in outdoor settings [11], [26], [27]. Many methods leverage semantic segmentation to classify different types of terrains in outdoor settings [9], [28], [29], [30]. For instance, ViTAL [31] leverages vision-based techniques for terrain-aware locomotion, which optimizes path planning across varied terrains with a focus on quadruped robots. Another noteworthy approach by Agarwal et al.[32], utilizes an egocentric camera to facilitate legged locomotion on rugged terrains through an end-to-end locomotion system. These methodologies predominantly utilize vision-based sensors, whose performance can degrade under various conditions such as light changes, motion blur, and occlusions [33].

2) *Proprioception-based Perception:* Proprioception, the internal sense of movement and position, plays a crucial role in quadruped robot navigation, as described in [1], [15], [34], [35]. It is defined as the robot's internal sense of movement and position [19]. While effective in terrain interaction, its limitations in predicting unencountered terrain necessitate external sensor integration for improved navigation. Wisth et al.'s development of a method combining proprioceptive feedback with visual, inertial, lidar, and odometry data exemplifies advancements in multi-modal sensor fusion [21]. This approach, further refined by cross-modal learning techniques by Loquercio et al. [36], enhances environmental understanding and adaptability. Moreover, studies by García et al. [37] and Karnan et al. [34], [38] illustrate the benefits of using proprioceptive data for gait adjustments and incorporating human preferences into navigation systems, showcasing the potential of learning-based methods for better terrain adaptation. This collective research shows the significance of integrating proprioceptive feedback with external sensors for advanced robotic navigation and terrain adaptability.

Other works include fused proprioception with visual data in particular to improve navigation. As highlighted by Fu et al. [15], integrating these sensory inputs improves spatial awareness and movement strategies. This approach is crucial for navigating adverse conditions, with Teng et al. [39] and Jin et al. [40] showcasing the importance of proprioception in maintaining balance and enhancing resilience. Assessing terrain supportiveness through proprioceptive data is essential for safe and effective navigation. Homberger et al. [41] discuss the role of proprioception in evaluating surface stability. Recently, Dey et al. [12] introduced a predictive method that utilizes proprioceptive signals to foresee and counteract potential navigation failures, emphasizing the need for anticipation in challenging terrains. While these methods integrate proprioception with exteroception, many of them assume the consistency between the two sensory modalities which is reasonable for some terrains. In our case, we consider terrains that could have different traversability under different weather conditions such as sand, mud, grass, and covered grounds.

B. Outdoor Navigation

Recent advancements in outdoor navigation for autonomous systems have showcased the integration of several methods and hardware designs [5], [31], [42], [43]. For instance, Overbye et al. [44] works on path optimization for ground vehicles in off-road terrains uses vehicle kinematic constraints and advanced algorithms for improved efficiency. On the other hand, Wellhausen et al. [45] developed robust navigation systems for legged robots, employing reachability-based methods and neural network-based learned motion costs to navigate unpredictable terrains in term of traversability. Miki et al. [46] advanced sensory integration techniques, employing attention-based recurrent encoders and lidar intensity maps for enhanced environmental interaction. Additionally, Valsecchi et al. [5] investigated the potential of quadruped robots for haptic inspection of planetary soils. Also, Kolvenbach et al. in [47] have developed a terrain-aware locomotion controller. This controller combines proprioceptive and exteroceptive inputs using a recurrent encoder, which enables the legged robots to navigate challenging terrains. Our navigation approach is different from these methods. We compute the best gait and velocity pairs based on the semantic traversability class and proprioception. This helps the robot to maintain stability as well as heading toward the goal.

III. BACKGROUND

In this section, we explain symbols, notations, and the preliminary concepts used in our work.

A. Symbols and Notations

We highlight the symbols of images and cost maps we use in Table (I).

Symbols	Definitions
I_{RGB}^t	Input RGB image.
I_{SD}^t	Image after applying semantic segmentation.
C_g^t	A general knowledge map based on vision . IV-A
C_h^t	A history map that's based on vision and proprioception. IV-B
C_p^t	A proprioception map that's based on proprioception. IV-D
$C_{coupled}^t$	A coupled map that's computed after the adaptive coupling. IV-E

TABLE I: List of symbols used in our approach. I_{RGB}^t and I_S^t are three-channel images of size $H \times W$. All the cost maps are grayscale of size $H \times W$ where H, W are the height and width of the images and the cost maps.

B. Setup and Conventions

The quadruped robot in our formulation has a coordinate frame attached to its center of mass with X, Y, and Z directions facing forward, leftward, and upward, respectively. The robot is equipped with sensors to measure the positions, velocities and forces exerted at each joint in the i^{th} leg along each axis. Additionally, the robot measures the current \mathcal{I}^t consumed from the battery. At every time instant t , these quantities are used as raw proprioceptive signals to *feel* the terrain and estimate its traversability (section III-C). The robot is also equipped with an RGB camera that provides an image I_{RGB}^t to view the terrains ahead, and a 3D lidar to detect obstacles. The robot can execute three kinds of gaits: **1. trot**, where two of the robot's feet make contact

with the ground while walking, **2. crawl**, where three feet touch the ground for enhanced stability but with a restricted maximum velocity, and **3. amble**: a gait similar to crawl but with unrestricted maximum speed.

We categorize the various terrains that the robot navigates into the following types: 1. stable, 2. granular (e.g. sand, snow), 3. poor foothold (e.g. rocks), 4. high resistance (e.g. dense vegetation). We use i, j to indicate grid locations in costmaps.

C. Proprioception-based Terrain Analysis

We utilize the approach proposed in [16] to estimate a terrain τ 's traversability by processing the raw proprioceptive signals (defined in Section III-B) from a legged robot and transforming them into two dimensional point \mathbf{p}_τ^t using Principal Component Analysis [48]. [16] shows that the L-2 norm ($\|\mathbf{p}_\tau^t\|_2$) indicates a terrain's traversability. After estimating a terrain's traversability, [16] empirically demonstrates that using a certain gait g to traverse a terrain type τ (III-B) results in a bounded distribution of points in the PCA space, that can be modeled as a 2D Gaussian ellipse $\mathcal{E}(\tau, g)$. [16] also shows that the area of a gait-terrain ellipse is directly proportional to the level of stability the robot experiences while traversing the terrain using the gait (see Fig. 2).




Fig. 2: The gaussian ellipses for the PCA components of four different terrains [16]. For each terrain, three types of gait data are shown using different ellipse boundaries. The solid line denotes trot, the dashed line denotes amble, and the dash-dot line is for crawl.

D. Vision-based Traversability Types

As a preliminary for our vision module, we use a terrain semantic segmentation method to classify the terrain type. Particularly, we incorporate [9], a semantic segmentation method that uses group-wise attention to differentiate the traversability of the terrains and provides a segmentation image. We train this segmentation method using RUGD dataset [49] using 4 terrain groups as presented in Table II.

Semantic Categories	Terrain Type
Concrete, Asphalt	Stable
Sand, Dirt, Gravel, Mulch	Granular
Rocks, Rockbed	Poor Foothold/Rocky
Grass, Bush	High Resistance/Vegetation

TABLE II: Semantic categories and their corresponding terrain type τ .

IV. OUR APPROACH

In this section, we discuss how we construct our adaptive multimodal cost map using semantic segmentation and proprioception. We highlight the symbols and notation of our approach in table I.

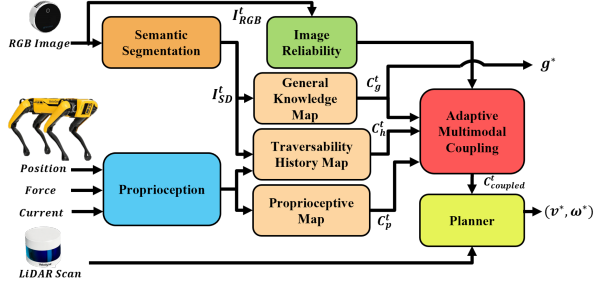


Fig. 3: The overall system architecture AMCO that couples vision and proprioception perception modules. We use the robot’s joint position, forces, and battery current to compute our traversability measure. We also use a semantic segmentation module to classify the terrains. After that, three traversability cost maps are created: general knowledge map (vision-based), history map (vision and proprioception), and proprioception map. All of them are coupled based on the reliability of the input RGB images to create a coupled cost map for the local planner.

A. General Knowledge Map

Given an RGB image I_{RGB}^t , we process it through a semantic segmentation model (III-D) to output a segmented image I_S^t . Each pixel in I_S^t categorizes the terrain τ into the four terrain classes defined in Section III-B. Next, we create a general knowledge map that represents the typical traversability of the terrains ahead of the robot as a costmap. It is constructed by assigning a traversability cost to each segmented terrain class τ based on the smallest-area ellipse [16] corresponding to it. This ellipse represents the distribution of points obtained when using the most stable gait for τ (see Fig. 2).

To create general knowledge map, we first discretize I_S^t into grids of $n \times n$ pixels as shown in Fig. 5.c to obtain I_{SD}^t . Each grid in the general knowledge map is assigned a cost value based on the terrain classes that the pixels in the corresponding grid in I_{SD}^t belong to. Consequently, (C_g^t) is constructed as,

$$C_g^t(i, j) = \frac{\gamma}{n^2} \sum_{(a,b) \in I_{SD}^t(i,j)} \min_{g \in \mathcal{G}} (\mathbf{Ar}(\mathcal{E}^t(\tau_{a,b}, g))), \quad (1)$$

where, $C_g^t(i, j)$ denotes the traversability cost for the $(i, j)^{th}$ grid in the general knowledge map at time t . γ is a scaling factor. \mathcal{G} represents the set of all gaits, and $\mathbf{Ar}(\mathcal{E}^t(\tau_{a,b}, g))$ refers to the area of the ellipse associated with using gait g in terrain type $\tau_{a,b}$ present at pixel (a, b) .

Additionally, by detecting the predominant terrain τ_{pred} in a 1×1 square meter area in front of the robot in I_{SD}^t , we can choose the most stable gait to traverse it. This is done based on the gait g that leads to the smallest area for the ellipse $\mathcal{E}^t(\tau_{pred}, g)$. That is,

$$g^* = \operatorname{argmin}_{g \in \mathcal{G}} \mathbf{Ar}(\mathcal{E}^t(\tau_{pred}, g)). \quad (2)$$

B. Recent Traversability History Map

We define traversability history of a terrain type τ as a list of its traversability ($\|\mathbf{p}_\tau^t\|_2$) accumulated over a duration T_τ by the robot. The intuition behind T_τ is that the longer the robot has traversed a terrain τ , the more reliable its traversability cost estimate becomes. This history is then

coupled with the general knowledge map C_g^t . This is because, intuitively, if soil is particularly muddy and unstable on a given day—indicating low stability—it is justified to reflect this observation in the traversability cost for all the soil regions ahead of the robot, even if typically soil is perceived to be more stable by the costs in C_g^t . Building on this intuition, we formulate the recent traversability history as a costmap C_h^t for a terrain τ as,

$$C_h^t(i, j) = \begin{cases} \alpha \cdot \sum_{l=0}^{T_{\tau_{a,b}}} ((\|\mathbf{p}_\tau^l\|_2 - C_g^l(i, j))), & \text{if } (a, b) \in I_{SD}^t(i, j) \text{ and } \tau_{a,b} \in \Gamma, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Here, Γ is the set of terrains that the robot has traversed in the past, and $T_{\tau_{a,b}}$ is the time the robot spent traversing terrain $\tau_{a,b}$. $C_h^t(i, j)$ represents the adjusted traversability cost at grid (i, j) and time t , and similarly for $C_g^t(i, j)$. α is an adjustable factor. Equation 3 ensures that a traversability history cost is applied only to the regions with previously traversed terrains in front of the robot.

C. Reliability Estimation of RGB images

Assessing the reliability of RGB images is essential for vision-based sensing. The image’s brightness and its blurriness can affect the semantic segmentation (see figure 4) and consequently the reliability of our general knowledge map (IV-A) and history map (IV-B). Here, we discuss how we evaluate those two visual phenomena:

1) *Brightness Percentage*: We use two methods to calculate the brightness of an image. The first computes the mean brightness (r_{mean}) of an image as the average of the RGB values for each pixel, normalized to a percentage: $r_{mean} = \frac{1}{N} \sum_{i=1}^N \left(\frac{R_i + G_i + B_i}{3} \right)$, Where N denotes the number of pixels in the image. R_i , G_i , and B_i represent the red, green, and blue color values of the i^{th} pixel, respectively. The arithmetic mean is widely used to estimate the brightness, however, it can be misleading in outdoor environments as it treats all the RGB pixels uniformly. To improve the brightness estimation, we additionally use [50] to compute the luminance (luma) of an image as: $r_{luma} = \frac{1}{N} \sum_{i=1}^N (\rho_1 \cdot R_i + \rho_2 \cdot G_i + \rho_3 \cdot B_i)$. Similar to r_{mean} but with scalar weights: ρ_1, ρ_2, ρ_3 .




Fig. 4: Reliability Estimation of RGB images: we estimate the reliability of RGB images (top row) by measuring two factors, the image’s brightness, and blurriness as both affect semantic segmentation (bottom row). (a) and (b) show two brightness levels for an RGB image. (c) shows a sharp image (with almost no blur). (d) shows how rocky terrains could lead to blurriness in the image.




Fig. 5: Traversability cost maps of our method. (a) The input RGB image, (b) The segmented image with the color labels as (stable: green, vegetation: red, granular: yellow and poor foothold: orange), (c) The general knowledge map (vision-based), (d) The history map, (e) The proprioception map. (f) The coupled map. Dark blue reflects low traversability cost in (c-d) maps and conversely red indicates high traversability cost. As shown in the figure, concrete shows low cost in the general knowledge map and particularly evident in the proprioception-based cost map as the robots captured the terrain stability using proprioception.

2) *Blurriness*: An image that is perceived as sharp (not blurry) is typically associated with high fidelity and clear detail representation, which affects semantic segmentation (see figure 4). To quantify the sharpness/blurriness of an image, we use two methods. The first method employs the Haar wavelet transform [51] to derive an energy map E_{map} which can be used to estimate a blurriness score (r_{hwt}) of an image. Deep-learning methods were used previously to detect motion blur [52], [53]. However, the datasets for motion blur in outdoor settings are limited. Therefore, we create a dataset with images from unstructured outdoor environments with different degrees of motion blur corruption, and the % of corruption as their labels. Then, we train a ResNet50 [54] model on our dataset. We represent the output of the model as, $O = \Phi(I_{RGB}^t)$. We use both methods together to improve blurriness detection.

Overall, we formulate the reliability score ξ as follows:

$$\xi = \lambda_1 \cdot r_{mean} + \lambda_2 \cdot r_{luma} + \lambda_3 \cdot r_{hwt} + \lambda_4 \cdot O \quad (4)$$

Here, λ_1 , λ_2 , λ_3 , and λ_4 are coefficients that balance the contributions of brightness and blur to the overall reliability score.

D. Current Proprioception Map

Apart from the general knowledge map and history map, we also use the robot's current proprioceptive sensing $|\mathbf{p}_\tau^t|_2$ to estimate terrain traversability due to its reliability in indicating how the terrain *feels* regardless of the vision-based sensing. We construct a current proprioception map C_p by extrapolating $|\mathbf{p}_\tau^t|_2$ based on the distance $d_{i,j}$ from the robot's location without using the terrain segmentation (unlike C_g^t and C_h^t) as,

$$C_p^t(i, j) = (\mathcal{U} - \delta \cdot |\mathbf{p}_\tau^t|_2) \cdot (1 - d_{i,j}). \quad (5)$$

Here, \mathcal{U} is a positive value corresponding to the cost of traversing a moderately traversable terrain. We assume a moderately traversable terrain as we do not use the terrain information ahead from the terrain segmentation. μ is a scale factor for the traversability's impact on cost. $d_{i,j}$ is the normalized distance to the $(i, j)^{th}$ grid from the bottom center grid in front of the robot (see Fig. 5.e), with values ranging from 0 (directly in front) to 1 (furthest from the robot).

E. Adaptive Multimodal Coupling

The Adaptive Multimodal Coupling mechanism combines the general knowledge map C_g^t , history map C_h^t , and the proprioception map C_p^t into a coupled traversability cost map $C_{coupled}^t$. This integration is achieved through the reliability metric $\xi \in [0 - 1]$ (Section IV-C), which assesses the quality of RGB images to adaptively adjust the contributions of the vision-based components. The construction of $C_{coupled}^t$ is formulated as,

$$C_{coupled}^t(i, j) = \xi \cdot (C_g^t(i, j) + C_h^t(i, j)) + C_p^t(i, j). \quad (6)$$

Here, each $(i, j)^{th}$ grid in the input maps results in a corresponding $(i, j)^{th}$ grid in $C_{coupled}^t$. The intuition behind this formulation lies in the different reliability of the modalities. The components reliant on vision (C_g^t and C_h^t) are susceptible to variations in environmental conditions, such as motion blur, which can diminish their reliability. Conversely, the proprioception-based component C_p^t provides an accurate measure of the terrain's physical properties, unaffected by visual distortions. Thus, ξ serves as an adaptive scaling factor, reducing the influence of vision-based assessments when their reliability is compromised, and ensuring that coupled map reflects the most accurate traversability information possible.

F. Navigation using the Coupled Traversability Cost Map

We incorporate a modified version of the Model Predictive Controller (MPC) in [43] to generate actions (i.e., (v, ω) linear and angular velocity pairs) for the robot to prefer relatively stable terrains while reaching user-defined local goals. MPC in [43] calculates a restricted velocity search space \mathcal{V}_r by choosing only the obstacle-free and dynamically feasible actions from an action space $\mathcal{V}_s = \{(v, \omega) | v \in [0, v_{max}], \omega \in [-\omega_{max}, \omega_{max}]\}$, where v_{max} and ω_{max} are the robot's linear and angular velocity limits.

Then, we define an objective function $J(v, \omega)$ as follows to calculate the optimal action from the restricted velocity space \mathcal{V}_r .

$$J(v, \omega) = \beta_1 \cdot goal(v, \omega) + \beta_2 \cdot obs(v, \omega) + \beta_3 \cdot sur(v, \omega), \quad (7)$$

where $goal(v, \omega)$ and $obs(v, \omega)$ are cost terms that indicate the distance to the goal and proximity to an obstacle for a given action $(v, \omega) \in \mathcal{V}_r$. $\beta_i \forall i = 1, 2, 3$ are tunable weights. We introduce a traversability cost term $sur(v, \omega)$ to leverage the terrain perception information encoded in the coupled traversability cost map $C_{coupled}$. To calculate the $sur(v, \omega)$ for a given action (v, ω) , we first extrapolate the action for a ΔT time and obtain a trajectory $traj^R(v, \omega)$ w.r.t. the robot. Then, $traj^R(v, \omega)$ is projected to the cost map using perspective projection [55]. Let's denote the projected trajectory on the cost map as $traj^C(v, \omega) \in C_{coupled}$. Then, the terrain surface traversability cost $sur(v, \omega)$ is calculated as,

$$sur(v, \omega) = \frac{1}{|traj^C(v, \omega)|} \sum_{(i,j) \in traj^C(v, \omega)} C_{coupled}^t(i, j), \quad (8)$$

where $C_{coupled}^t(i, j)$ is the traversability cost of the $(i, j)^{th}$ location on the $C_{coupled}^t$ at time t, and $|traj^C(v, \omega)|$ is the




Fig. 6: Navigation trajectories of AMCO (Ours) and the other methods in three challenging outdoor scenarios. AMCO, ProNav, and RFC can switch between three gaits (trot, crawl, and amble) during navigation as presented in three different colors. AMCO, ProNav, and RFC can switch between three gaits (trot, crawl, and amble) during navigation as presented in three different colors. The trajectory colors are represented as follow: AMCO (trot: brown, crawl:red, amble: pink), ProNav [16] (trot: blue, crawl:purple, amble: light blue), GaNav [9] (orange), Spot's Inbuilt (yellow), and RFC [20] (trot: green, crawl:light green, amble: light beige). Vision-only methods such as GA-Nav fail under low light conditions while proprioception-only methods such as ProNav lead to robot failures in extremely unstable terrains since proprioception cannot estimate traversability beforehand. In contrast, AMCO minimizes the robot's instability during navigation by adaptively coupling both vision and proprioception.

cardinality of the $traj^C(v, \omega)$. Finally, the optimal navigation action (v^*, ω^*) is obtained as,

$$(v^*, \omega^*) = \underset{(v, \omega) \in \mathcal{V}_r}{\operatorname{argmin}} J(v, \omega). \quad (9)$$

After that, we send (v^*, ω^*) along with g^* (which is computed from equation 2) to the robot's controller.

V. RESULTS AND ANALYSIS

A. Implementation & Robot Setup

We implement AMCO on Boston Dynamics Spot, a quadruped robot with 12 degrees of freedom (DOF) which provides joint feedback (position, velocities, and forces) from the robot legs and the current consumed by the battery. Additionally, the robot is equipped with a VLP16 Velodyne lidar and a Zed 2i camera. AMCO runs on an Intel NUC edge computer equipped with an Intel i7 CPU, and an Nvidia RTX2060 GPU. We use the following other parameters for our implementation: $n = 40, \gamma = 1, \alpha = 4.5\rho_1 = 0.299, \rho_2 = 0.587, \rho_3 = 0.114, \delta = 31.875, \mathcal{U} = 127, \lambda_1 = 0.0008, \lambda_2 = 0.001, \lambda_3 = 0.0025, \lambda_4 = 0.003, v_{max} = 0.6 \text{ m/s}, \omega_{max} = 0.2 \text{ rad/s}, \beta_1 = 2.4, \beta_2 = 3.2, \beta_3 = 0.05$.

B. Comparison Methods

- **Spot's in-built planner:** A planner that uses RGB-D cameras to detect and avoid obstacles.
- **GA-Nav [9]:** A vision-based method that uses semantic segmentation of RGB images to compute traversability costs for different terrains. It is combined with a local planner [56] to compute trajectories on low-cost terrains and avoid obstacles.
- **ProNav [16]:** A proprioceptive-based method that uses joint feedback and battery's current to estimate traversability and changes to an appropriate gait while computing goal-conditioned velocities using [2].
- **Random Forest Classifier (RFC) [20]:** It uses proprioceptive signals to classify the terrain type. We used the dataset in [16] to train the classifier to identify between four different terrain types (granular, poor foothold, high resistance and stable). RFC is combined with [2] to generate goal-directed velocities (v^*, ω^*) .

C. Evaluation Metrics

- **Success Rate:** The ratio of successful navigation trials where the robot was able to reach its goal without freezing or colliding with obstacles.
- **Normalized Trajectory Length:** The ratio between the robot's trajectory length and the straight-line distance to the goal averaged across both successful and unsuccessful trials.
- **IMU Energy Density:** The mean of the aggregated squared acceleration values measured by the IMU sensors across the x, y, and z axes, calculated over the successful trials [25]: $E_i = \sum_{n=1}^N a_i^2$, and $E_{\text{Total}} = E_{ax} + E_{ay} + E_{az}$, where a_i represents one of the three acceleration signals (x, y, and z axes), and N is the IMU readings along the trajectory.
- **Vibration Cost:** The cumulative sum of deviations, measured from a predefined stable reference terrain, in the actuator positions at the robot's hip joints along its trajectory [16]:

$$\delta_j(t) = \begin{cases} |p^j(t) - \min(p_{\text{ref}}^j)| & \text{if } p^j(t) < \min(p_{\text{ref}}^j), \\ |p^j(t) - \max(p_{\text{ref}}^j)| & \text{if } p^j(t) > \max(p_{\text{ref}}^j). \end{cases}$$

Where, $p^j(t)$ is the position of joint j at time t , and p_{ref}^j represents reference positions from a stable terrain (concrete). The total Vibration Cost at time t is then computed as $\text{Vibration Cost}(t) = \sum_j \delta_j(t)$.

D. Test Scenarios

- **Scenario 1:** contains granular terrain (sand) and grass.
- **Scenario 2:** contains dense vegetation.
- **Scenario 3:** contains bushes, mulch, and concrete.
- **Scenario 4:** contains mud, grass and concrete.

E. Analysis and Discussion

We evaluated our method's navigation performance qualitatively in Figs. 1 and 6, and quantitatively in Table III. Across the four scenarios, our method AMCO consistently outperforms the other methods in terms of successful goal-reaching by effectively selecting appropriate gaits and velocities to reach the designated goal. Further, AMCO results in lower vibration experienced by the robot in all scenarios compared to the other methods.

In Scenario 1, GA-Nav generates significantly longer trajectories preferring grass over the granular sand terrain, and

Metrics	Method	Scen. 1	Scen. 2	Scen. 3	Scen. 4
Success Rate (%) ↑	In-built system	10	10	-	-
	GA-Nav[9]	80	30	20	60
	RFC[20]	40	50	60	30
	ProNav[16]	60	60	70	40
	AMCO w/o reliability	70	60	50	80
	AMCO w/o history	70	70	80	60
	AMCO (ours)	100	90	100	90
Normalized Trajectory Length	In-built system	1.215	1.420	-	-
	GA-Nav[9]	1.324	1.394	2.105	1.447
	RFC[20]	1.089	0.916	1.129	0.838
	ProNav[16]	1.063	0.948	1.036	0.861
	AMCO w/o reliability	1.192	1.131	1.235	1.217
	AMCO w/o history	1.230	1.178	1.189	1.253
	AMCO (ours)	1.171	1.064	1.093	1.184
IMU Energy Density ↓	In-built system	37506	51764	-	-
	GA-Nav[9]	18392	29354	21349	26096
	RFC[20]	23412	30233	26043	37412
	ProNav[16]	21342	26490	24392	35291
	AMCO w/o reliability	14323	25374	25034	24523
	AMCO w/o history	16271	24361	22375	25828
	AMCO (ours)	13928	19832	20126	23120
Vibration Cost ↓	In-built system	29.3	48.7	-	-
	GA-Nav[9]	14.3	28.5	23.9	19.4
	RFC[20]	15.9	29.2	37.4	36.7
	ProNav[16]	16.4	31.3	34.9	34.1
	AMCO w/o reliability	12.3	24.1	26.6	19.1
	AMCO w/o history	11.5	22.8	21.4	23.4
	AMCO (ours)	9.3	20.3	19.3	17.3

TABLE III: The table shows the quantitative results of AMCO and the other methods over 10 trials. We use four navigation metrics to evaluate our method’s performance. All metrics are averaged over both the successful and unsuccessful trails (reaching the goal).

demonstrates oscillatory behavior near the boundary between the two terrains. ProNav and RFC use relatively straight trajectories and switch to crawl gait after encountering heavy instability in the granular sand region. However, in certain instances, their inconsistent gait selection leads to sinkage in the sand. In contrast, AMCO switches to the crawl gait in advance using the general knowledge map to minimize the potential instabilities in the forthcoming terrain (i.e., sand).

Scenario 2 depicts a navigation task on concrete followed by a densely vegetated region. ProNav leads to leg entanglement in vegetation due to its delayed gait switching. GANav’s perception of vegetation is affected by image motion blur which leads to inconsistent segmentation. However, AMCO generates trajectories to maximize the stay in the concrete region while switching to the amble gait before entering the vegetated region to reduce the leg entanglement risks.

In Scenario 3, our method AMCO relies heavily on proprioception map for navigation due to low and varying lighting conditions which results in less reliable RGB images. Both Spot’s in-built system and GANav struggles to reach the goal due to erroneous perception from the camera images under challenging lighting. The complexity of the muddy terrain in **Scenario 4** poses significant challenges to all the methods. Particularly, RFC and ProNav methods experience considerable leg sinkage due to delayed gait switching. Although GA-Nav attempts to take longer trajectories through grass and concrete, the method struggles in muddy regions due to its fixed gait. In contrast, AMCO initially adjusts its gait using the general knowledge map and later changes the trajectory to avoid the heavily deformable muddy regions to reach the goal.

Benefits of Image Reliability: We observe that our method’s navigation performance is significantly affected without the image reliability estimation, especially under low-light conditions (scenario 3) and motion blur (scenarios 2 and 4). This is primarily due to the inconsistent segmentation impacts

on the general knowledge map, which eventually results in errors in the coupled traversability cost estimation. However, AMCO’s adaptive coupling which accounts for image reliability ensures prioritizing the proprioception under unreliable image inputs.

Benefits of Traversability History: We observe a notable decrease in success rate when traversability history is omitted from AMCO’s formulation. Particularly in scenarios 1 and 4, where the visually similar-looking granular terrain’s deformability varies from one region to another. However, the use of traversability history allows AMCO to avoid severely deformable regions based on recent experiences on the same terrain. Further, the recent history enables AMCO to update the forthcoming terrain conditions better, compared to relying only on fixed costs provided by image segmentation.

VI. CONCLUSION, LIMITATIONS & FUTURE WORK

We present AMCO, a novel approach for quadruped robot navigation through challenging outdoor environments by adaptively combining vision-based and proprioceptive-based perception capabilities. Our method effectively utilizes semantic segmentation from RGB images and proprioceptive inputs, to generate a coupled perception representation. This enables the selection of appropriate gaits and velocities, which significantly improve navigation success rates and stability across various terrains. Our method has some limitations. For instance, AMCO’s reliance on proprioceptive data under dark conditions limits its capabilities. It’s a good direction to explore integrating other sensor modalities in evaluating traversability such as (LiDAR, thermal or hyperspectral cameras). Furthermore, incorporating Vision Language Models (VLMs) could offer more insights into complex environments, due to their common reasoning capabilities and ability to classify more terrains.

REFERENCES

- [1] J. Frey, M. Mattamala, N. Chebrolu, C. Cadena, M. Fallon, and M. Hutter, “Fast traversability estimation for wild visual navigation,” *arXiv preprint arXiv:2305.08510*, 2023.
- [2] A. J. Sathyamoorthy, K. Weerakoon, T. Guan, M. Russell, D. Conover, J. Pusey, and D. Manocha, “Vern: Vegetation-aware robot navigation in dense unstructured outdoor environments,” *arXiv preprint arXiv:2303.14502*, 2023.
- [3] X. Cai, S. Ancha, L. Sharma, P. R. Osteen, B. Bucher, S. Phillips, J. Wang, M. Everett, N. Roy, and J. P. How, “Evora: Deep evidential traversability learning for risk-aware off-road autonomy,” *IEEE Transactions on Robotics*, 2024.
- [4] S. Sotnik and V. Lyashenko, “Agricultural robotic platforms,” 2022.
- [5] G. Valsecchi, C. Weibel, H. Kolvenbach, and M. Hutter, “Towards legged locomotion on steep planetary terrain,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 786–792.
- [6] L. Bruzzone and P. Fanghella, “Functional redesign of mantis 2.0, a hybrid leg-wheel robot for surveillance and inspection,” *Journal of Intelligent & Robotic Systems*, vol. 81, pp. 215–230, 2016.
- [7] T. Yoshiike, M. Kuroda, R. Ujino, H. Kaneko, H. Higuchi, S. Iwasaki, Y. Kanemoto, M. Asatani, and T. Koshiishi, “Development of experimental legged robot for inspection and disaster response in plants,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 4869–4876.
- [8] J. Sun, Y. Meng, J. Tan, C. Sun, J. Zhang, N. Ding, H. Qian, and A. Zhang, “A vision-based perception framework for outdoor navigation tasks applicable to legged robots,” in *2017 Chinese Automation Congress (CAC)*. IEEE, 2017, pp. 2894–2899.
- [9] T. Guan, D. Kothandaraman, R. Chandra, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, “Ga-nav: Efficient terrain segmentation for robot navigation in unstructured outdoor environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8138–8145, 2022.

- [10] 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.
- [11] A. J. Sathyamoorthy, K. Weerakoon, M. Elnoor, and D. Manocha, "Using lidar intensity for robot navigation," *arXiv preprint arXiv:2309.07014*, 2023.
- [12] S. Dey, D. Fan, R. Schmid, A. Dixit, K. Otsu, T. Touma, A. F. Schilling, and A.-A. Agha-Mohammadi, "Prepare: Predictive proprioception for agile failure event detection in robotic exploration of extreme terrains," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4338–4343.
- [13] K. Weerakoon, A. J. Sathyamoorthy, M. Elnoor, and D. Manocha, "Vapor: Legged robot navigation in unstructured outdoor environments using offline reinforcement learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 10 344–10 350.
- [14] P. Ewen, H. Chen, Y. Chen, A. Li, A. Bagali, G. Gunjal, and R. Vasudevan, "You've got to feel it to believe it: Multi-modal bayesian inference for semantic and property prediction," *arXiv preprint arXiv:2402.05872*, 2024.
- [15] Z. Fu, A. Kumar, A. Agarwal, H. Qi, J. Malik, and D. Pathak, "Coupling vision and proprioception for navigation of legged robots," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 273–17 283.
- [16] M. Elnoor, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, "Pronav: Proprioceptive traversability estimation for legged robot navigation in outdoor environments," *IEEE Robotics and Automation Letters*, 2024.
- [17] X. Yao, J. Zhang, and J. Oh, "Rca: Ride comfort-aware visual navigation via self-supervised learning," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 7847–7852.
- [18] C. A. Brooks and K. D. Iagnemma, "Self-supervised classification for planetary rover terrain sensing," in *2007 IEEE aerospace conference*. IEEE, 2007, pp. 1–9.
- [19] A. H. Al-dabbagh and R. Ronsse, "A review of terrain detection systems for applications in locomotion assistance," *Robotics and Autonomous Systems*, vol. 133, p. 103628, 2020.
- [20] C. Kertész, "Rigidity-based surface recognition for a domestic legged robot," *IEEE Robotics and automation letters*, vol. 1, no. 1, pp. 309–315, 2016.
- [21] D. Wisth, M. Camurri, and M. Fallon, "Vilens: Visual, inertial, lidar, and leg odometry for all-terrain legged robots," *IEEE Transactions on Robotics*, 2022.
- [22] S. Jung, J. Lee, X. Meng, B. Boots, and A. Lambert, "V-strong: Visual self-supervised traversability learning for off-road navigation," *arXiv preprint arXiv:2312.16016*, 2023.
- [23] J. Dallás, M. P. Cole, P. Jayakumar, and T. Ersal, "Terrain adaptive trajectory planning and tracking on deformable terrains," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 11, pp. 11 255–11 268, 2021.
- [24] F. J. Comin, W. A. Lewinger, C. M. Saaj, and M. C. Matthews, "Trafficability assessment of deformable terrain through hybrid wheel-leg sinkage detection," *Journal of Field Robotics*, vol. 34, no. 3, pp. 451–476, 2017.
- [25] P. Try and M. Gebhard, "A vibration sensing device using a six-axis imu and an optimized beam structure for activity monitoring," *Sensors*, vol. 23, no. 19, p. 8045, 2023.
- [26] E. I. Al Khatib, M. A. K. Jaradat, and M. F. Abdel-Hafez, "Low-cost reduced navigation system for mobile robot in indoor/outdoor environments," *IEEE Access*, vol. 8, pp. 25 014–25 026, 2020.
- [27] Y. D. Yasuda, L. E. G. Martins, and F. A. Cappabianco, "Autonomous visual navigation for mobile robots: A systematic literature review," *ACM Computing Surveys (CSUR)*, vol. 53, no. 1, pp. 1–34, 2020.
- [28] R. Miyamoto, M. Adachi, H. Ishida, T. Watanabe, K. Matsutani, H. Komatsuzaki, S. Sakata, R. Yokota, and S. Kobayashi, "Visual navigation based on semantic segmentation using only a monocular camera as an external sensor," *Journal of Robotics and Mechatronics*, vol. 32, no. 6, pp. 1137–1153, 2020.
- [29] W. Zhan, C. Xiao, Y. Wen, C. Zhou, H. Yuan, S. Xiu, X. Zou, C. Xie, and Q. Li, "Adaptive semantic segmentation for unmanned surface vehicle navigation," *Electronics*, vol. 9, no. 2, p. 213, 2020.
- [30] J. Frey, S. Khattak, M. Patel, D. Atha, J. Nubert, C. Padgett, M. Hutter, and P. Spieler, "Roadrunner—learning traversability estimation for autonomous off-road driving," *arXiv preprint arXiv:2402.19341*, 2024.
- [31] S. Fahmi, V. Barasuo, D. Esteban, O. Villarreal, and C. Semini, "c," *IEEE Transactions on Robotics*, 2022.
- [32] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Conference on Robot Learning*, 2022, pp. 403–415.
- [33] M. Aladem, S. Baek, S. A. Rawashdeh, *et al.*, "Evaluation of image enhancement techniques for vision-based navigation under low illumination," *Journal of Robotics*, vol. 2019, 2019.
- [34] K. Haresh, Y. Elvin, W. Garrett, B. Joydeep, and S. Peter, "Learning to extrapolate human preferences for preference aligned path planning," *arXiv preprint arXiv:2309.09912*, 2023.
- [35] J. Ren, Y. Liu, Y. Dai, and G. Wang, "Top-nav: Legged navigation integrating terrain, obstacle and proprioception estimation," *arXiv preprint arXiv:2404.15256*, 2024.
- [36] A. Loquercio, A. Kumar, and J. Malik, "Learning visual locomotion with cross-modal supervision," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7295–7302.
- [37] R. García, M. Estrada, M. Ebrahimi, F. Zuppichini, L. Gambardella, A. Giusti, and A. Ijspeert, "Gait-dependent traversability estimation on the k-rock2 robot," in *2022 26th International Conference on Pattern Recognition (ICPR)*. IEEE, 2022, pp. 4204–4210.
- [38] H. Karnan, E. Yang, D. Farkash, G. Warnell, J. Biswas, and P. Stone, "Sterling: Self-supervised terrain representation learning from unconstrained robot experience," in *Conference on Robot Learning (CoRL 2023)*, 2023.
- [39] S. Teng, M. W. Mueller, and K. Sreenath, "Legged robot state estimation in slippery environments using invariant extended kalman filter with velocity update," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3104–3110.
- [40] J. Jin, C. Zhang, J. Frey, N. Rudin, M. Mattamala, C. Cadena, and M. Hutter, "Resilient legged local navigation: Learning to traverse with compromised perception end-to-end," *arXiv preprint arXiv:2310.03581*, 2023.
- [41] T. Homberger, L. Wellhausen, P. Fankhauser, and M. Hutter, "Support surface estimation for legged robots," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8470–8476.
- [42] K. Weerakoon, A. J. Sathyamoorthy, M. Elnoor, and D. Manocha, "Adventr: Autonomous robot navigation in complex outdoor environments," in *International Symposium on Experimental Robotics*. Springer, 2023, pp. 219–228.
- [43] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics Automation Magazine*, vol. 4, no. 1, pp. 23–33, March 1997.
- [44] T. Overbye and S. Saripalli, "Path optimization for ground vehicles in off-road terrain," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7708–7714.
- [45] L. Wellhausen and M. Hutter, "Artplanner: Robust legged robot navigation in the field," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [46] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [47] H. Kolvenbach, C. Bärtschi, L. Wellhausen, R. Grandia, and M. Hutter, "Haptic inspection of planetary soils with legged robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1626–1632, 2019.
- [48] A. Maćkiewicz and W. Ratajczak, "Principal components analysis (pca)," *Computers & Geosciences*, vol. 19, no. 3, pp. 303–342, 1993.
- [49] M. Wigness, S. Eum, J. G. Rogers, D. Han, and H. Kwon, "A rugd dataset for autonomous navigation and visual perception in unstructured outdoor environments," in *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [50] S. Bezryadin, P. Bourov, and D. Ilinih, "Brightness calculation in digital image processing," in *International Symposium on Technologies for Digital Photo Fulfillment 2007*, vol. 1, 2007, pp. 10–15.
- [51] H. Tong, M. Li, H. Zhang, and C. Zhang, "Blur detection for digital images using wavelet transform," in *2004 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2004.
- [52] K. Purohit, A. B. Shah, and A. N. Rajagopalan, "Learning based single image blur detection and segmentation," in *IPCV Lab, Department of Electrical Engineering, Indian Institute of Technology Madras, India*, 2023.
- [53] S. Nayak and B. Ravindran, "Reinforcement learning for improving object detection," in *Indian Institute of Technology Madras, Chennai TN, India*, 2023.
- [54] S. R. J. S. Kaiming He, Xiangyu Zhang, "Deep residual learning for image recognition," *arXiv preprint arXiv:1512.03385*, 2015.
- [55] K. Weerakoon, A. J. Sathyamoorthy, J. Liang, T. Guan, U. Patel, and D. Manocha, "Graspe: Graph based multimodal fusion for robot navigation in unstructured outdoor environments," *arXiv preprint arXiv:2209.05722*, 2022.
- [56] K. Weerakoon, A. J. Sathyamoorthy, U. Patel, and D. Manocha, "Terp: Reliable planning in uneven outdoor environments using deep reinforcement learning," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 9447–9453.