

VAPOR: Legged Robot Navigation in Unstructured Outdoor Environments using Offline Reinforcement Learning

Kasun Weerakoon¹, Adarsh Jagan Sathyamoorthy¹, Mohamed Elnoor¹, and Dinesh Manocha²

Abstract—We present VAPOR, a novel method for autonomous legged robot navigation in unstructured, densely vegetated outdoor environments using offline Reinforcement Learning (RL). Our method trains a novel RL policy using an actor-critic network and arbitrary data collected in real outdoor vegetation. Our policy uses height and intensity-based cost maps derived from 3D LiDAR point clouds, a goal cost map, and processed proprioception data as state inputs, and learns the physical and geometric properties of the surrounding obstacles such as height, density, and solidity/stiffness. The fully-trained policy’s critic network is then used to evaluate the quality of dynamically feasible velocities generated from a novel context-aware planner. Our planner adapts the robot’s velocity space based on the presence of entrapment including vegetation, and narrow passages in dense environments. We demonstrate our method’s capabilities on a Spot robot in complex real-world outdoor scenes, including dense vegetation. We observe that VAPOR’s actions improve success rates by up to 40%, decrease the average current consumption by up to 2.9%, and decrease the normalized trajectory length by up to 11.2% compared to existing end-to-end offline RL and other outdoor navigation methods.

I. INTRODUCTION

Autonomous robot navigation in complex outdoor scenes is an essential capability for many applications, including precision agriculture [1], search and rescue operations in forested environments [2], reconnaissance[3], etc. There are two major challenges in navigating such scenarios. Firstly, the robot must perceive and differentiate non-solid/pliable obstacles (e.g. tall grass), from solid/non-pliable obstacles (e.g. trees) [4]. Pliable obstacles can be safely traversed through, whereas non-pliable obstacles must be avoided. Secondly, apart from avoiding collisions, the robot also faces challenges such as narrow passages, and scenarios where the vegetation could wrap/attach to the robot and entrap it. The robot’s navigation must be capable of handling such adverse situations.

To address the perceptual challenges in outdoors, methods based on image classification [4], semantic segmentation [5], and anomaly detection [6] using supervised learning have been employed. However, such works require extensive manual annotation and labeling to identify traversable terrain during training. Such models also may not align with the actual traversability capabilities of different robots due to varying dynamic constraints. This restricts the robot’s navigation and could lead to highly conservative, meandering trajectories [7], or freezing behaviors [8]. Imitation learning

Code, Video, and full technical report are available at <http://gamma.umd.edu/vapor/>.

This research was supported by Army Cooperative Agreement W911NF2120076, ARO grants W911NF2310046 and W911NF2310352. We acknowledge the support of the Maryland Robotics Center.

¹ Authors are with Dept. of Electrical and Computer Engineering, University of Maryland, College Park, MD, USA. kasunw@umd.edu, asathyam@umd.edu, melnoor@umd.edu

² Author is with Dept. of Computer Science, University of Maryland, College Park, MD, USA. dm@cs.umd.edu




Fig. 1: Legged robot trajectories generated when navigating through complex outdoor vegetation using VAPOR, end-to-end CQL-SAC [10], IQL [11], BCO [12], VERN [4], and DWA [13]. VAPOR identifies the entrapment in vines and uses holonomic actions to minimize the instability and current consumption instead of excessive angular actions taken by the other methods which result in further entrapment. Hence, in this scenario, VAPOR moves backward with minimal angular motion to reduce entanglement with vines.

techniques have also been proposed for outdoor navigation, but the resulting models may not generalize well [9].

Conversely, outdoor navigation methods based on online reinforcement learning (RL) [14], [15], [16] do not require human labeling since they are trained using a robot’s active interactions with a simulated environment. Nevertheless, such models exhibit severe performance degradation during real-world deployment due to sim-to-real transfer issues [17], [18]. Training complex models using RL requires high-fidelity simulations, which may not be available, especially for complex scenarios. To alleviate such shortcomings, offline-RL [19] methods have been proposed, where a model is trained using data collected in real-world environments, reducing the sim-to-real transfer issues and the need for manual labeling.

However, none of the existing methods for outdoor navigation have accounted for the constraints that complex real-world scenes impose on the robot’s velocities. For instance, while traversing through tall grass and bushes, angular motions could cause the vegetation to easily wrap around the robot restricting its motion. Furthermore, in cases with solid obstacles create narrow passages, the robot would have to rotate/reorient itself to maneuver the narrow free space. Therefore, the robot’s executable actions must be adapted based on the environment.

Main contributions: To address these challenges, we propose VAPOR, an offline RL-based trajectory evaluation model combined with a context-aware planner designed to generate dynamically feasible velocities to operate a legged robot in challenging outdoor scenes. VAPOR’s offline RL formulation allows it to be trained using data collected in the real world [20] that is automatically compiled for training,

alleviating sim-to-real transfer issues. The novel components of our work are:

- We propose a novel offline RL-based actor-critic network architecture to learn a Q-function to evaluate a legged robot’s candidate actions and velocities in terms of their ability to reach the goal, avoid solid, non-pliable vegetation, and other desirable behaviors. The network consists of spatial and channel attention layers to learn the spatial correlations in the input observation space. Our model is trained using real-world data collected in dense environments that are automatically compiled into states and actions between randomly chosen start and goal states. This alleviates the sim-to-real transfer issues prevalent in existing RL methods. This results in an improvement up to 40% in terms of success rate
- A novel observation space to sense dense vegetation consisting of robot-centric height and intensity cost maps obtained by processing lidar point clouds, a goal map indicating the distance and direction to the goal, and proprioceptive signals from the legged robot’s joints to indicate its stability. The height and intensity maps accurately represent the height and solidity or inversely, the pliability of the surrounding vegetation. The goal map and proprioception aid with spatially correlating the vegetation’s properties with the robot’s intended motion direction and stability during training.
- A novel context-aware motion planner that switches between (1). a holonomic velocity space to minimize the risk of entrapment in vegetation assessed from proprioceptive signals, and (2). a non-holonomic velocity space to navigate narrow passages between solid, non-pliable vegetation. Further, it generates dynamically feasible, smooth candidate actions/velocities to be assessed by VAPOR’s Q-function. VAPOR is evaluated on a real Boston Dynamics Spot robot in unstructured outdoor scenes.

II. RELATED WORK

In this section, we discuss the existing literature on vegetation perception in outdoor environments, and offline RL methods used for navigation. Finally, we discuss the existing holonomic planning methods.

A. Outdoor Vegetation Perception

Navigating robots in outdoor environments, particularly through vegetation, is a challenging task that has received increasing attention in recent years [4], [21], [22]. Existing approaches tackle this issue using various sensory modalities and learning techniques. For instance, [23] adopts a self-supervised approach to estimate the support surface in vegetation employing 3D point clouds and RGB images. Despite their promising results, the system requires manual labeling, which could be time-consuming and less scalable for real-world deployments. In [4], the authors use RGB images and 2D Lidar to create traversability cost maps in dense vegetation environments. Stone et al. [24] use an infrared (IR) sensor and an RGB camera for vegetation detection. While effective in certain conditions, these camera-based methods are often vulnerable to environmental factors such as changing lighting and motion blur, thereby limiting their robustness. Iberraken et al. [25] demonstrate the use of a 2D LiDAR to navigate through structured vineyard fields.

Some recent works have shifted their focus from external sensors to proprioceptive modalities to perceive vegetation [26], [27]. While proprioception offers reliable feedback about the robot’s internal state, it inherently lacks the capability for look-ahead predictions before traversing a given terrain, especially in the absence of exteroceptive sensors. Our work combines exteroception (3D point clouds) with proprioception for robust and efficient navigation through outdoor vegetation.

B. Offline RL based Robot Navigation

Reinforcement Learning (RL) has been fundamental to robot navigation [14], [28], [29], providing methods for autonomous decision-making based on interaction with the environment. However, traditional online RL often falls short in situations where real-time data collection is either impractical or lacks realistic simulation which increases the sim-to-real gap [30] e.g., navigating through dense vegetation or hazardous or complex terrains. On the other hand, offline RL has emerged as a promising alternative, designed to optimize policies based on pre-collected datasets. Among the foundational works in offline RL is the study by Levine et al. [19], which outlines the key methodologies and challenges. Methods based on Imitation Learning (IL) have also leveraged collected data [31], [32]. However, IL is often restricted by the limitations of the human operator (expert) who collected the data, meaning it cannot generally surpass the operator’s performance. In contrast, offline RL aims to optimize the target policy based on the dataset, offering the potential for more generalized and sometimes superior strategies [11], [33], [20]. Kostrikov et al. [11] introduce Implicit Q-Learning (IQL) which implicitly estimates the value function without querying the Q function of unseen actions. While IQL shows promise, it struggles in tasks with long planning horizons. Shah et al. [20] mitigate this limitation by combining IQL with topological graphs. Nevertheless, their method primarily relies on RGB images which are susceptible to lighting changes, motion blur, etc.

On the other hand, [10] developed Conservative Q-Learning (CQL) to enhance the robustness of the learned policy. This method lower-bounds the true value of its learned Q function. Following its superior performance with complex data distribution, we extend this method by employing data from a 3D LiDAR and a legged robot’s joint encoders.

C. Holonomic Planning

Traditional robotic planning often focuses on non-holonomic robotic planners [34], [35], [36], largely because many robots including wheeled robots, inherently possess non-holonomic constraints [37], [38], [39], [40]. Conversely, robots with higher degrees of freedom (e.g., legged or manipulator robots) can benefit from holonomic planning methods [41], [42]. However, such planners lack the ability to adapt the robot’s velocity space based on environmental constraints, especially in dense vegetation.

III. VAPOR: VEGETATION-AWARE PLANNING USING OFFLINE REINFORCEMENT LEARNING

A. Preliminaries

We mathematically formulate our navigation problem as a Markov Decision Process (MDP) with continuous states and actions. Our MDP can be defined as $\mathcal{M} := \{\mathcal{S}, \mathcal{A}, \mathbb{P}, r, \gamma\}$,

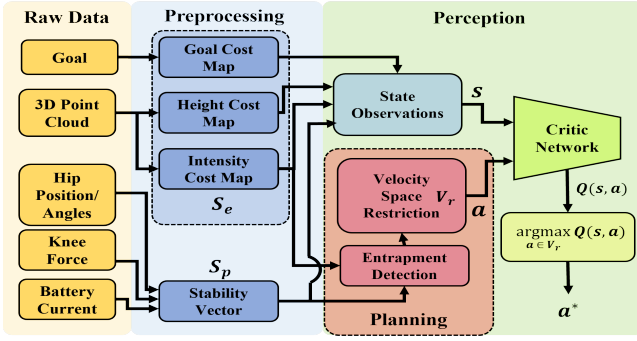


Fig. 2: Overall system architecture of VAPOR which uses a height and intensity cost map generated from 3D lidar, a goal cost map, and proprioception data from the robot as state inputs to train an Actor-Critic offline RL policy. Then, the fully trained critic network is used to evaluate the dynamically feasible actions generated by a planner. The planner uses instability detection using proprioception and intensity map to switch between a holonomic and non-holonomic action space to reduce the risk of entrapment.

where \mathcal{S}, \mathcal{A} denote state and action spaces, $\mathbb{P}(s'|s, a)$ represents the state transition dynamics between current state s , action a , and next state s' . $r(s, a)$ is the reward function, and $\gamma \in (0, 1)$ denotes the discount factor. The objective of RL is to learn a policy $\pi_\theta(a|s)$ parameterized by θ that maximizes the discounted cumulative reward return.

Offline RL particularly aims to learn policies from existing data sets instead of explicitly interacting with the environment. Hence, for a dataset $\mathcal{D} = \{(s_j, a_j, r_j, s'_j) | s_j, s'_j \in \mathcal{S}; a_j \in \mathcal{A}; j = 1, 2, \dots, N\}$, offline RL algorithms attempt to learn a policy $\pi_\theta(a|s)$ that maximizes the discounted reward return $R_t = \sum_{k=t}^T \gamma^{(k-t)} r_k(s_k, a_k)$ at time step t . However, leveraging the standard RL algorithms for offline RL leads to poor performance due to overfitting and distributional shifts [19]. In particular, the existing value-based off-policy RL methods typically overestimate the value function predictions for unseen outcomes, which results in erroneous and overly optimistic estimations [43]. To mitigate this issue, Conservative Q Learning (CQL) [10] regularizes the Q-values during training to learn conservative and lower-bound estimates of the value function. Hence, in this work, we incorporate CQL with Soft Actor-Critic (SAC) [44] as our base offline RL algorithm.

Hereafter, we use j, k as indices. Vectors are represented in bold, lowercase letters. All positions, velocities, and forces are represented w.r.t a rigid frame attached to the robot R (indicated in superscript) or relative to a cost map. The robot frame's x, y, z directions point forward, leftward, and upward respectively.

B. Dataset Generation

Our raw training data is collected by teleoperating a legged robot equipped with a 360° 3D LiDAR, and joint encoders for ~ 4 hours. We collect raw 3D point clouds, the robot's odometry, joint positions and velocities on the legs, and joint actuator current as the robot moves in random trajectories in vegetation including grass, bushes, and trees with varying density. Hence, the raw data set does not have any goal-conditioning or goal-reaching policy.

To create goal-conditioned data set \mathcal{D} with a series of $\{s_j, a_j, r_j, s'_j\}$, we consider random trajectory segments from the raw dataset, i.e., we select a random state as the initial position and a future sample in the same raw trajectory as the goal. This subsequent goal sample is selected such that it




Fig. 3: Robot-centric cost map state observations from exteroceptive sensing from Scenario 3 in Fig. 5: [Left] Point cloud-based intensity cost map C_i that indicates the density of the surrounding objects using lidar reflectance; [Center] Height cost map C_h that represents the maximum height of the objects derived from the point cloud; [Right] Goal cost map C_g that indicates the distance to the goal from the robot's neighborhood. Light colors indicate higher costs and dark areas represent lower costs.

is $\sim 8 - 20$ meters away from the robot's initial position, and our processed dataset $\mathcal{D} = \{(s_j, a_j, r_j, s'_j) | j = 1, 2, \dots, N\}$ is obtained. We explain the details of the state observations, actions, and reward formulation in the sub-sections below.

C. State Observations from Multi-sensor Data

Our state observations $s \in \mathcal{S}$ are obtained by preprocessing the raw sensory data collected from both the exteroceptive (point clouds) and proprioceptive (joint positions, forces) data from the robot. We denote the entire point cloud as \mathbf{P} , j^{th} reflected point's 3D location relative to the robot and intensity as $\mathbf{p}_j = \{x_j, y_j, z_j, i_j | x_j, y_j, z_j \in \mathbb{R}; i_j \in [0, i_m]\}$. Proprioceptive sensing is obtained from the robot's joint positions $h_1^{x/y}, h_2^{x/y}, h_3^{x/y}, h_4^{x/y}$, force feedback f_1, f_2, f_3, f_4 , and the battery's current consumption I_b .

We preprocess the aforementioned sensory data to generate two types of state observations: 1.) S_e : A set of robot-centric cost maps that reflect the solidity and height of the surrounding objects, and distance to the goal using exteroceptive sensors; 2.) S_p : A vector that quantifies the robot's stability using proprioception. Hence, our final state observations $s = [S_e, S_p] \in \mathcal{S}$.

1) *Layered Cost Maps from Exteroception*: Navigation in outdoor vegetation requires sensing the height and solidity of the vegetation in the robot's vicinity. Moreover, spatial information of the goal location is necessary to perform successful goal-reaching tasks. Hence, we propose three robot-centric 2D cost maps, intensity cost map (C_i), height cost map (C_h), and goal cost map (C_g), to represent the solidity and height of the surrounding objects/vegetation, and distance and direction to the goal respectively.

All three cost maps C_i, C_h and C_g are $n \times n$ matrices with the robot positioned at the center $(n/2, n/2)$ as depicted in Fig. 3. The elements in each costmap lie in $[0, 100]$. A grid (l, m) in a cost map is related to a grid $grid_{l,m}^R$ of physical (x, y) locations relative to the robot as,

$$(x, y) \in grid_{l,m}^R$$

$$grid_{l,m}^R = [[x_{l,m}, x_{l,m} + \beta], [y_{l,m}, y_{l,m} + \beta]] \quad (1)$$

$$x_{l,m} = \left[\left(l - \frac{n}{2} \right) \cdot \beta \right] \text{ and } y_{l,m} = \left[\left(m - \frac{n}{2} \right) \cdot \beta \right],$$

where β is the side length of a square-shaped grid $grid_{l,m}^R$ in meters.

Intensity Cost Map (C_i): We employ the point cloud intensity values [45], [46], to construct an intensity cost map C_i . The LiDAR's reflectance power (i.e., intensity) is directly proportional to the solidity of the corresponding objects. Hence, we observe that grass, bushes, and trees result in

distinct intensities (see Fig. 3a). We calculate elements of \mathcal{C}_i as,

$$\mathcal{C}_i(l, m) = \frac{\sum_{x_j} \sum_{y_j} i_j}{\beta^2} \quad \forall \mathbf{p}_j \in \mathbf{P} \text{ and } x_j, y_j \in \text{grid}_{l,m}^R. \quad (2)$$

Height Cost Map (\mathcal{C}_h): We generate \mathcal{C}_h to represent the maximum heights of the objects in each grid location $\text{grid}_{l,m}^R$. To this end, element (l, m) of \mathcal{C}_h is obtained by,

$$\mathcal{C}_h(l, m) = \max(z_j) \quad \forall \mathbf{p}_j \in \mathbf{P} \text{ and } x_j, y_j \in \text{grid}_{l,m}^R, \quad (3)$$

where higher values in $\mathcal{C}_h(l, m)$ indicate taller objects.

Goal Cost Map (\mathcal{C}_g): Each location (l, m) in the goal cost map represents $\text{grid}_{l,m}^R$'s distance to the goal (x_g^R, y_g^R) . Its value is calculated as,

$$\mathcal{C}_g(l, m) = \frac{\alpha_g \cdot \left(\sqrt{(x_g^R - x_{l,m})^2 + (y_g^R - y_{l,m})^2} \right)}{d_{tot}}, \quad (4)$$

where d_{tot} is the total distance to the goal from the robot's starting position and $\alpha_g \in \mathbb{R}$ is a tunable weight parameter.

Finally, we obtain our state observation from the exteroception S_e by concatenating the derived cost maps. Hence, $S_e = \{\mathcal{C}_i, \mathcal{C}_h, \mathcal{C}_g\}$ of shape $n \times n \times 3$.

2) *Stability Observation from Proprioception:* To estimate the robot's stability in vegetation, we incorporate data acquired from the robot's joint positions, forces, and battery current for proprioceptive sensing. To this end, we process the raw proprioceptive data $H_{prop} = [h_1^{x/y}, h_2^{x/y}, h_3^{x/y}, h_4^{x/y}, f_{1/2/3/4}, I_b]$, as performed in [26]. Principal Component Analysis (PCA) is then applied to the processed data to reduce its dimensions to two primary axes. Subsequently, we extract the variances (σ_{PC1}^2 and σ_{PC2}^2) of the dimension-reduced data along the principal components, and define our resulting proprioceptive state observation vector as, $S_p = [\sigma_{PC1}^2, \sigma_{PC2}^2, \sigma_{PC1}^2, \sigma_{PC2}^2] \in \mathbb{R}^+$. We observe that highly stable terrains such as asphalt lead to lower variances, and unstable terrain leads to higher values.

Lastly, we derive our final state observation as $s = [S_e, S_p]$ by combining both exteroceptive and proprioceptive state observations.

D. Offline Reinforcement Learning Using CQL-SAC

Our network architecture is based on CQL-SAC [10] that incorporates two critic networks and an actor network. The policy actor-network (i.e., $\pi_\theta(a|s)$) estimates the parameters θ of the policy distribution, which provides the conditional probability of taking action a given the state observation s . In our context, this policy distribution is Gaussian parameterized by the mean μ_θ and standard deviation σ_θ . Further, the two critic networks are Q networks (i.e., $Q_1(s, a; \pi_\theta), Q_2(s, a; \pi_\theta)$) that uses state-action pairs s, a as inputs to estimate the expectation of the value function. We design the actor and critic networks as follows.

1) *Actor and Critic Networks:* Our actor and critic network architecture with layer dimensions is presented in Fig. 4. In both networks, we use two separate network branches to process the exteroceptive S_e and proprioceptive S_p observations in our input state s . We highlight the use of spatial and channel attention networks in the exteroception branch. Spatial attention blocks encode spatial neighborhood properties in individual cost maps and channel attention helps learn the correlations between the features between the cost

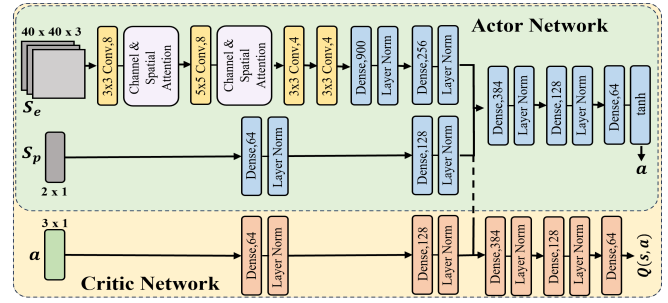


Fig. 4: Actor and Critic network architectures of our method. We incorporate two separate branches to process exteroception S_e , and proprioception S_p observations. We use spatial and channel attention to encode correlation among the layered cost maps in S_e .

maps. The outputs from the two branches are concatenated and processed using several linear layers to obtain the end-to-end action outputs.

Since the critic networks take both the action and state inputs, we use an additional branch to process the action by passing two linear layers through before concatenating with the state observation branches. All the hidden layers in the network are followed by *ReLU* activation.

2) *Reward Functions:* The reward function is formulated to obtain robot actions that lead to desired navigation behavior. In this work, we are primarily interested in three navigation behaviors: 1) Goal reaching; 2) Avoiding dense/solid objects while navigating through pliable vegetation; and 3) Minimizing the overall energy consumption. We introduce three reward terms r_{goal}, r_{veg} , and r_{energy} to achieve the aforementioned behaviors. Hence, the total reward r_{tot} obtained by the robot for a given sample is calculated as,

$$r_{tot} = \beta_1 r_{goal} + \beta_2 r_{veg} + \beta_3 r_{energy}, \quad (5)$$

where $\beta_1, \beta_2, \beta_3$ are tunable parameters to weigh the reward terms. We design r_{goal} based on the robot's current distance d_g to the goal to encourage moving towards the goal. Hence,

$$r_{goal} = \frac{\lambda_1 d_{tot}}{d_g} \mathbb{1}_{\{d_g > d_{th}\}} + \lambda_2 \mathbb{1}_{\{d_g \leq d_{th}\}}, \quad (6)$$

where $\lambda_1, \lambda_2 \in \mathbb{R}$ are adjustable parameters, $\mathbb{1}$ is an indicator function, and d_{th} is the goal reaching threshold.

The vegetation reward r_{veg} is a penalty for actions that navigate the robot in dense vegetation nearby (i.e., higher the density, lower the reward). To this end, we consider three circular neighborhoods with radii 0.5, 1.5, and 2.5 meters centered at the robot. Let, A_1, A_2 and A_3 denote the sets of grids corresponding to these neighborhoods in the intensity cost map \mathcal{C}_i . Then, r_{veg} is calculated as,

$$r_{veg} = - \sum_{k=1,2,3} \left(\frac{\eta_k}{|A_k|} \sum_{l,m \in A_k} \mathcal{C}_i(l, m) \right), \quad (7)$$

where the tunable parameters are set such that $\eta_1 > \eta_2 > \eta_3 \in \mathbb{R}$ to ensure higher penalties for the dense vegetation in the robot's nearby vicinity. $|A_k|$ denotes cardinality of the set A_k .

We incorporate r_{energy} to penalize actions consuming high amounts of energy (proportional to the current I_b) during navigation. We calculate r_{energy} as,

$$r_{energy} = -\epsilon I_b, \quad (8)$$

where $\epsilon \in \mathbb{R}^+$ is a weight parameter.

3) *Critic Networks for State-action Evaluation*: Even though we train an end-to-end navigation policy using CQL-SAC on our data set \mathcal{D} , we do not use the actions from the trained policy $\pi_\theta(a|s)$ in the actor-network for navigation. Instead, we leverage the Q-function $Q(s, a)$ learned by a critic network to evaluate the quality of the set of actions generated by a context-aware planner. Intuitively, $Q(s, a)$ indicates how well the action leads to desirable behaviors imposed by the reward function. Since CQL-SAC includes two critic networks and learned Q-functions ($Q_1(s, a; \pi_\theta)$ and $Q_2(s, a; \pi_\theta)$), we choose the critic network with the lowest training loss. We refer to its Q-function as $Q_{min}(s, a; \pi_\theta)$ from here on.

E. Context-Aware Planning

To generate dynamically feasible candidate actions to be evaluated using $Q_{min}(s, a; \pi_\theta)$, we formulate a novel context-aware planner. An action $a \in \mathcal{A}$ for our robot can be denoted as $a = (v_x, v_y, \omega_z)$. The planner uses a 3-dimensional velocity space ($V_s \subset \mathcal{A}$) defined as $V_s = \{(v_x, v_y, \omega_z) \mid -v_{max} \leq v_x, v_y \leq v_{max}, -\omega_{max} \leq \omega \leq \omega_{max}\}$. Here, v_x and v_y denote the linear velocities along the robot's x and y directions respectively, and ω_z represents the angular velocity about the vertical z-axis. v_{max} and ω_{max} are the maximum linear and angular velocity limits. Additionally, the planner uses the set of reachable/dynamically feasible velocities from the current velocities within an interval Δt based on acceleration limits as $V_r = \{[v_x - \dot{v}_{max}\Delta t, v_x + \dot{v}_{max}\Delta t], [v_y - \dot{v}_{max}\Delta t, v_y + \dot{v}_{max}\Delta t], [\omega_z - \dot{\omega}_{max}\Delta t, \omega_z + \dot{\omega}_{max}\Delta t]\}$. Here, \dot{v}_{max} , and $\dot{\omega}_{max}$ are the robot's maximum linear and angular acceleration limits.

The risk of entrapment in dense vegetation a robot faces is exacerbated when the robot performs angular motions because it aids the vegetation in helically twirling on to its legs (intuitively similar to rotating a fork on spaghetti). Therefore, in such scenarios, the robot's angular motion must be restricted. On the other hand, in scenarios with narrow passages, the rectangularly shaped robot must be capable of performing angular motions to traverse through. Such behaviors are also desirable when the robot is equipped with a sensor with a limited FOV that needs to be pointed in a specific direction. To accommodate both scenarios, we restrict V_s based on the following condition:

$$C : \sqrt{\sigma_1^2 + \sigma_2^2} > \Gamma, \mathcal{C}_i(l, m) \in [0.5i_m, 0.75i_m] \forall l, m \in A_2.$$

$$V_s = \begin{cases} \{(v_x, v_y, 0)\}, & \text{if } C \text{ is True,} \\ \{(v_x, 0, \omega_z)\}, & \text{Otherwise,} \end{cases} \quad (9)$$

where $v_x, v_y \in [-v_{max}, v_{max}]$, and $\omega \in [-\omega_{max}, \omega_{max}]$. The corresponding V_r is calculated from the restricted V_s by omitting either v_y or ω_z based on the environment. The best action a^* for the robot to execute given the current state s can then be found as,

$$a^* = \operatorname{argmax}_{a_k \in V_r} (Q_{min}(s, a_k)). \quad (10)$$

IV. RESULTS AND ANALYSIS

A. Implementation

Our CQL-SAC offline RL policy is implemented using PyTorch and our model is trained on a workstation with an

Metrics	Methods	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Success Rate (%)	DWA [13]	30	0	0	20
	VERN [4]	60	70	10	40
	BCO [12]	10	0	0	10
	IQL [11]	40	30	40	20
	CQL-SAC [10]	50	60	50	50
	VAPOR w/o Proprioception	50	40	50	30
	VAPOR w/o Attention	60	50	40	60
VAPOR (ours)	80	70	60	70	
Avg. Current Consumption (A)	DWA [13]	7.158	7.482	7.206	7.502
	VERN [4]	6.937	7.457	6.993	7.423
	BCO [12]	6.681	7.153	6.937	7.391
	IQL [11]	7.155	7.436	7.161	7.466
	CQL-SAC [10]	7.192	7.301	7.099	7.487
	VAPOR w/o Proprioception	7.013	7.464	6.792	7.408
	VAPOR w/o Attention	6.835	7.198	6.704	7.298
VAPOR (ours)	6.599	7.147	6.735	7.319	
Norm. Traj. Length	DWA [13]	1.327	1.655	0.428	1.421
	VERN [4]	1.105	1.327	1.517	1.365
	BCO [12]	0.425	1.398	0.422	0.643
	IQL [11]	0.735	0.686	1.761	0.892
	CQL-SAC [10]	0.897	1.245	1.453	1.277
	VAPOR w/o Proprioception	1.236	1.364	1.386	1.338
	VAPOR w/o Attention	1.125	1.223	1.374	1.294
VAPOR(ours)	1.065	1.238	1.289	1.256	

TABLE I: Navigation performance of VAPOR compared to other methods on various evaluation metrics in four test scenarios that are not included in the data set.

Intel Xeon 3.6 GHz processor and an Nvidia Titan GPU. For real-time deployment and inference, we use the Spot robot from Boston Dynamics equipped with a VLP16 Velodyne LiDAR, an onboard Intel NUC 11, which includes an Intel i7 CPU and an NVIDIA RTX 2060 GPU.

B. Comparison Methods and Evaluation Metrics

We compare our method's navigation performance with three recent offline RL algorithms: CQL-SAC (our end-to-end policy) [10], IQL [11], BCO [12], an autonomous imitation learning approach, VERN [4], an outdoor vegetation navigation algorithm, and the Dynamic Window Approach (DWA) [13] a classical model-based navigation approach that uses 2D LiDAR scans. We train all the aforementioned offline RL comparison methods on our data set \mathcal{D} using network architectures similar to ours for fair comparison. We further perform two ablation studies: VAPOR w/o Proprioception; and VAPOR w/o attention to highlight the benefits of our approach. Our metrics for evaluation are:

Success Rate - The number of times the robot reached its goal while avoiding collisions with *solid and dense vegetation* over the total number of attempts.

Avg. Current Consumption - The average battery current consumption during a navigation task (i.e., $\sum_{traj} I_b$) in Amperes (A).

Normalized Traj. Length - The robot's trajectory length normalized using the straight-line distance to the goal for both successful and unsuccessful trajectories.

C. Testing Scenarios

We compare our method's navigation performance in real-world outdoor test scenarios that are not included in the training data set. At least 10 trials are conducted in each scenario.

- **Scenario 1** - Contains narrow passages between shrubs, and trees in a mulch surface.
- **Scenario 2** - Dense bushes that lead to entrapment, sparse grass, and trees.
- **Scenario 3** - Thin grass, shrubs, and trees with narrow openings under low light conditions.
- **Scenario 4** - Dense grass, fallen branches, vines, and trees.

D. Analysis and Comparison

We evaluate our method's navigation performance qualitatively in Fig. 5 and quantitatively in Table I. Scenario 4 is




Fig. 5: Trajectories generated when navigating through complex outdoor vegetation using various comparison methods. The trajectories drawn from the robot’s rear indicate that it has moved backward. VAPOR is able to use holonomic action in dense vegetation and vines (Scenarios 2 and 4) to reduce the risk of entrapment while others use angular velocities to reach the goal which results in instability and navigation failures. In the presence of narrow spaces in Scenarios 1 and 3, VAPOR uses non-holonomic actions to navigate through.




Fig. 6: Sample velocities v_x generated by VAPOR and CQL-SAC (end-to-end) policy during a navigation task. The end-to-end policy actions demonstrate significantly high variance which indicates dynamically infeasible velocities to the robot leading to high vibrations and high motor current consumption. In contrast, VAPOR generates smooth, dynamically feasible actions.

Methods	Inference Time (ms)
VERN [4]	84.612
BCO [12]	3.622
IQL [11]	3.951
VAPOR w/o Attention	8.820
VAPOR (Ours)	8.934

TABLE II: Inference time comparison between ours and other methods when executing in robot’s onboard computer. VAPOR has a lightweight network that can execute in real time while providing accurate predictions.

presented in the Fig. 1. We observe that VAPOR demonstrate the highest success rate compared to other methods in all four scenarios that include diverse and unseen vegetation. Since the data set does not include expert demonstrations specifically collected with the behaviors imposed by reward functions, behavioral cloning with BCO [12] shows the lowest success rate due to its attempt to imitate the data set trajectories without the knowledge of the rewards. In contrast, offline RL methods such as IQL and CQL-SAC attempt to perform the navigation tasks at a reasonable success rate. Even though VERN demonstrates the second-best success rate in Scenarios 1 and 2, it performs poorly in low light conditions in Scenario 3 and trees covered with leaves in Scenario 4 due to the erroneous vegetation classification from its vision-based system. DWA freezes in tall and dense vegetation in Scenarios 2 and 3 identifying such regions as obstacles from the 2D LiDAR scan.

Benefits of Proprioception: We observe that VAPOR’s performance in terms of success rate and current consumption degrades in dense vegetation without the proprioception state observations. Further, our planner uses proprioception to restrict the angular velocities during entrapment in scenarios 2 and 4 which leads to a higher success rate and low current consumption than VAPOR without proprioception. VERN and DWA lead to entrapment in scenarios 2 and 4 due to lack of vegetation awareness from proprioception. Moreover, in stable conditions such as scenario 3, our planner uses angular velocities to move between the trees which creates a narrow

passage.

Benefits of Attention: Our method without attention demonstrates a relatively low success rate, high power consumption, and longer trajectory lengths particularly due to the lack of feature encoding capabilities between the cost map inputs than when spatial and channel attention are included. We observe that VAPOR without attention deviates from the goal in some trials due to lack of spatial aware encoding from the goal cost map.

End-to-end RL vs Ours: We observe that end-to-end RL policies generate dynamically infeasible actions for the robot’s motors though the actions reflect the behavior imposed by the rewards (See Fig. 6). This leads to jerky motion due to motor vibrations (see Fig. 5) and high avg. current consumption. In contrast, VAPOR’s planner ensures that the actions are dynamically feasible which results in lower current consumption than all end-to-end RL models.

Inference Time: VAPOR has a lightweight network that can execute in real-time ($\sim 112\text{Hz}$) on the robot’s onboard computer while providing accurate predictions as shown in Table II and I. Vision-based methods such as VERN [4] have a significantly lower inference time due to computationally heavy backbones. In contrast, VAPOR incorporate relatively lower dimensional state inputs that can represent 360° view of the robot’s vicinity and a lightweight network to obtain comparable or better navigation performance.

V. CONCLUSIONS, LIMITATIONS AND FUTURE WORK

We present VAPOR, an offline-RL based method for legged robot navigation in outdoor vegetation. Our method uses randomly collected real-world data to train a navigation policy that can reach local goals while avoiding dense and solid vegetation. Instead of end-to-end actions from the policy, its fully trained critic network is used to evaluate dynamically feasible actions generated by a planner. The planner is capable of adaptively switching between holonomic and non-holonomic action to minimize entrapment in unstructured vegetation. We deploy our method into a Boston Dynamics Spot robot and evaluate it in real outdoor vegetation to demonstrate benefits.

Our method has a few limitations. Our planner cannot provide any theoretical guarantees on the behavior since the state-action evaluations are obtained from a Q function trained on a data set. Even though our method generalizes well compared to vision-based and supervised learning methods, large data sets are required for training. Further, our method cannot detect thin poles or string fences due to the low resolution of the lidar and lack of scene awareness.

REFERENCES

- [1] N. S. Naik, V. V. Shete, and S. R. Danve, "Precision agriculture robot for seeding function," in *2016 international conference on inventive computation technologies (ICICT)*, vol. 2. IEEE, 2016, pp. 1–3.
- [2] S. Karma, E. Zorba, G. Pallis, G. Statheropoulos, I. Balta, K. Miki, J. Vamvakari, A. Pappa, M. Chalaris, G. Xanthopoulos, et al., "Use of unmanned vehicles in search and rescue operations in forest fires: Advantages and limitations observed in a field trial," *International journal of disaster risk reduction*, vol. 13, pp. 307–312, 2015.
- [3] S. Li, C. Feng, Y. Niu, L. Shi, Z. Wu, and H. Song, "A fire reconnaissance robot based on slam position, thermal imaging technologies, and ar display," *Sensors*, vol. 19, no. 22, p. 5036, 2019.
- [4] 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.
- [5] 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.
- [6] L. Wellhausen, R. Ranftl, and M. Hutter, "Safe robot navigation via multi-modal anomaly detection," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1326–1333, 2020.
- [7] A. J. Sathyamoorthy, K. Weerakoon, T. Guan, J. Liang, and D. Manocha, "Terrapin: Unstructured terrain navigation using online self-supervised learning," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 7197–7204.
- [8] A. J. Sathyamoorthy, U. Patel, T. Guan, and D. Manocha, "Frozone: Freezing-free, pedestrian-friendly navigation in human crowds," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4352–4359, 2020.
- [9] A. Hussein, E. Elyan, M. M. Gaber, and C. Jayne, "Deep imitation learning for 3d navigation tasks," *Neural computing and applications*, vol. 29, pp. 389–404, 2018.
- [10] A. Kumar, A. Zhou, G. Tucker, and S. Levine, "Conservative q-learning for offline reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1179–1191, 2020.
- [11] I. Kostrikov, A. Nair, and S. Levine, "Offline reinforcement learning with implicit q-learning," *arXiv preprint arXiv:2110.06169*, 2021.
- [12] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation," *arXiv preprint arXiv:1805.01954*, 2018.
- [13] 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, 1997.
- [14] 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.
- [15] K. Weerakoon, S. Chakraborty, N. Karapetyan, A. J. Sathyamoorthy, A. Bedi, and D. Manocha, "Htron: Efficient outdoor navigation with sparse rewards via heavy tailed adaptive reinforce algorithm," in *Conference on Robot Learning*. PMLR, 2023, pp. 1629–1639.
- [16] J. Liang, K. Weerakoon, T. Guan, N. Karapetyan, and D. Manocha, "Adaptiveveon: Adaptive outdoor local navigation method for stable and reliable actions," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 648–655, 2022.
- [17] J. Liang, U. Patel, A. J. Sathyamoorthy, and D. Manocha, "Crowd-steer: Realtime smooth and collision-free robot navigation in densely crowded scenarios trained using high-fidelity simulation," in *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, 2021, pp. 4221–4228.
- [18] S. Chakraborty, A. S. Bedi, K. Weerakoon, P. Poddar, A. Koppel, P. Tokkar, and D. Manocha, "Dealing with sparse rewards in continuous control robotics via heavy-tailed policy optimization," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 989–995.
- [19] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.
- [20] D. Shah, A. Bhorkar, H. Leen, I. Kostrikov, N. Rhinehart, and S. Levine, "Offline reinforcement learning for visual navigation," *arXiv preprint arXiv:2212.08244*, 2022.
- [21] J. Iqbal, R. Xu, S. Sun, and C. Li, "Simulation of an autonomous mobile robot for lidar-based in-field phenotyping and navigation," *Robotics*, vol. 9, no. 2, p. 46, 2020.
- [22] G. Kahn, P. Abbeel, and S. Levine, "Badgr: An autonomous self-supervised learning-based navigation system," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.
- [23] 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.
- [24] D. L. Stone, G. Shah, Y. Motai, and A. J. Aved, "Vegetation segmentation for sensor fusion of omnidirectional far-infrared and visual stream," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 2, pp. 614–626, 2019.
- [25] D. Ibrakken, F. Gaurier, J.-C. Roux, C. Chaballier, and R. Lenain, "Autonomous vineyard tracking using a four-wheel-steering mobile robot and a 2d lidar," *AgriEngineering*, vol. 4, no. 4, pp. 826–846, 2022.
- [26] M. Elnoor, A. J. Sathyamoorthy, K. Weerakoon, and D. Manocha, "Pronav: Proprioceptive traversability estimation for autonomous legged robot navigation in outdoor environments," *arXiv preprint arXiv:2307.09754*, 2023.
- [27] Z. Jian, Z. Liu, H. Shao, X. Wang, X. Chen, and B. Liang, "Path generation for wheeled robots autonomous navigation on vegetated terrain," *arXiv preprint arXiv:2306.08977*, 2023.
- [28] D. Dugas, J. Nieto, R. Siegwart, and J. J. Chung, "Navrep: Unsupervised representations for reinforcement learning of robot navigation in dynamic human environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7829–7835.
- [29] U. Patel, N. K. S. Kumar, A. J. Sathyamoorthy, and D. Manocha, "Dwa-rl: Dynamically feasible deep reinforcement learning policy for robot navigation among mobile obstacles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6057–6063.
- [30] T. Zhang and H. Mo, "Reinforcement learning for robot research: A comprehensive review and open issues," *International Journal of Advanced Robotic Systems*, vol. 18, no. 3, p. 17298814211007305, 2021.
- [31] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4693–4700.
- [32] D. Silver, J. A. Bagnell, and A. Stentz, "Applied imitation learning for autonomous navigation in complex natural terrain," in *Field and Service Robotics: Results of the 7th International Conference*. Springer, 2010, pp. 249–259.
- [33] J. Li, C. Tang, M. Tomizuka, and W. Zhan, "Hierarchical planning through goal-conditioned offline reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10216–10223, 2022.
- [34] S. LaValle, "Planning algorithms," *Cambridge University Press google schola*, vol. 2, pp. 3671–3678, 2006.
- [35] J. Canny, *The complexity of robot motion planning*. MIT press, 1988.
- [36] D. Manocha, *Algebraic and numeric techniques in modeling and robotics*. University of California, Berkeley, 1992.
- [37] K. Yang, S. Moon, S. Yoo, J. Kang, N. L. Doh, H. B. Kim, and S. Joo, "Spline-based rrt path planner for non-holonomic robots," *Journal of Intelligent & Robotic Systems*, vol. 73, no. 1–4, pp. 763–782, 2014.
- [38] A. Khan, I. Noreen, and Z. Habib, "On complete coverage path planning algorithms for non-holonomic mobile robots: Survey and challenges," *J. Inf. Sci. Eng.*, vol. 33, no. 1, pp. 101–121, 2017.
- [39] H.-W. Chae, J.-H. Choi, and J.-B. Song, "Robust and autonomous stereo visual-inertial navigation for non-holonomic mobile robots," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 9, pp. 9613–9623, 2020.
- [40] S. Eshthardian and S. Khodaygan, "A continuous rrt*-based path planning method for non-holonomic mobile robots using b-spline curves," *Journal of Ambient Intelligence and Humanized Computing*, vol. 14, no. 7, pp. 8693–8702, 2023.
- [41] R. Holmberg and O. Khatib, "Development and control of a holonomic mobile robot for mobile manipulation tasks," *The International Journal of Robotics Research*, vol. 19, no. 11, pp. 1066–1074, 2000.
- [42] M. Alireza, D. Vincent, and W. Tony, "Experimental study of path planning problem using emcoa for a holonomic mobile robot," *Journal of Systems Engineering and Electronics*, vol. 32, no. 6, pp. 1450–1462, 2021.
- [43] R. F. Prudencio, M. R. Maximo, and E. L. Colombini, "A survey on offline reinforcement learning: Taxonomy, review, and open problems," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [44] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [45] L. Di Giammarino, I. Aloise, C. Stachniss, and G. Grisetti, "Visual place recognition using lidar intensity information," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4382–4389.
- [46] T. D. Barfoot, C. McManus, S. Anderson, H. Dong, E. Beerepoot, C. H. Tong, P. Furgale, J. D. Gammell, and J. Enright, "Into darkness: Visual navigation based on a lidar-intensity-image pipeline," in *Robotics Research: The 16th International Symposium ISRR*. Springer, 2016, pp. 487–504.