Learning Surface and Vertical Mobility for Enceladus Direct Ocean Access

Jack Naish
Independent
jrhn2@cantab.ac.uk

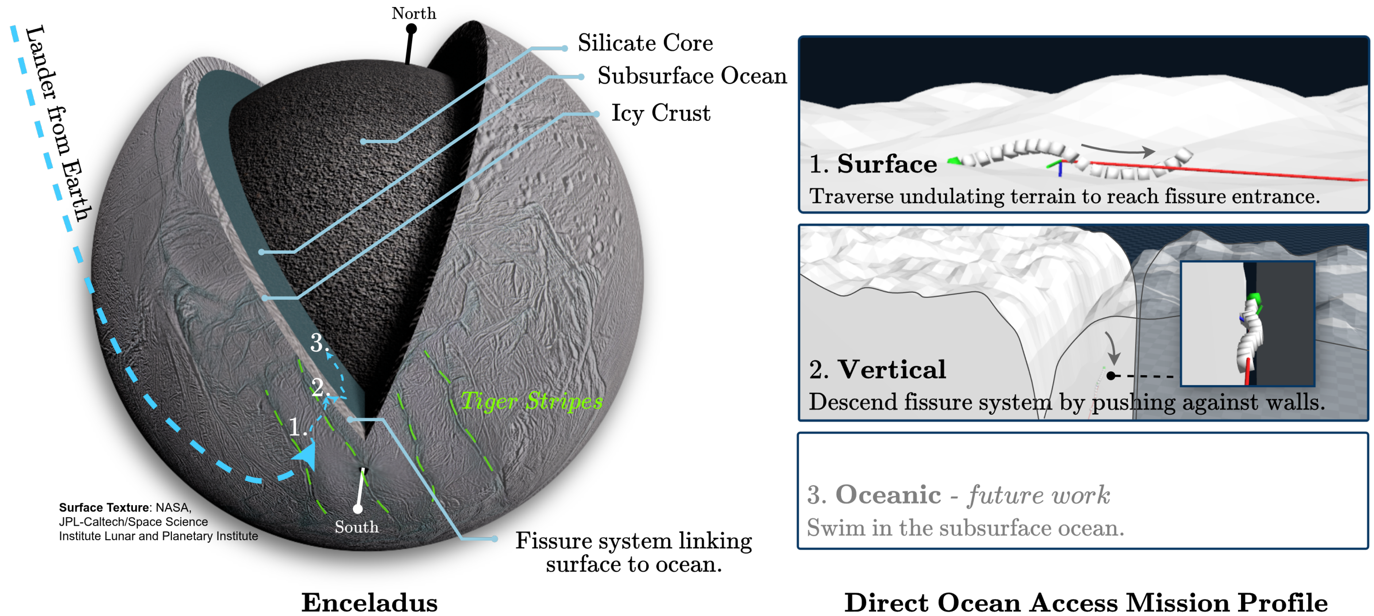

Fig. 1: Robotic exploration of Enceladus via direct ocean access, as originally proposed in [9] and [10].

Abstract*—Of the places in our solar system that may support life, Enceladus stands out due to the presence of complex organic molecules and possible hydrothermal vent activity in its subsurface ocean, with direct access theorized via fissure vents at its icy south pole. However, robotic exploration of the lunar interior requires first negotiating complex, undulatory surface terrain before transitioning into a vertical descent mode once inside the fissure system. Such a mission profile demands advances in multi-modal robotic locomotion which to-date are yet to be realized. In this work, we take a step towards enabling this vision by contributing a reinforcement learning controller capable of robust surface locomotion and vertical mobility with a snake-like morphology. Simulation experiments provide a proof-of-concept validation of our method. Our work takes a small step towards the broader idea of Robotic Exploration 3.0 [51]: *intelligent robots capable of adapting at mission-time, enabling one-shot exploration of celestial bodies.*

I. Introduction

Enceladus, a moon of Saturn, features a rocky silicate core in contact with a sub-surface liquid water ocean, encased in a cryogenic icy crust tens of kilometers thick [39]. A set of fissures, or *tiger stripes*, at the south pole vent gaseous plumes into space, suggesting a continuous path from the surface to ocean [30] (fig. 1). Combined with known complex organic molecules in plume samples [15] and theorized hydrothermal vent activity [31], Enceladus ranks highly on the list of places within reach that may support life beyond Earth [48] [8].

These factors motivate strong interest in robotic exploration [41] [71] [1] [42], particularly via direct ocean access [9]. Here, a robot must (1) traverse the undulating surface terrain from a delivery lander to fissure opening, (2) vertically descend the fissure while fighting turbulent, possibly supersonic [46], plume flow and (3) swim in the subsurface ocean.

Each phase requires fundamentally different locomotion modes (surface, vertical, and oceanic) making conventional morphologies — such as rovers, drones, or quadrupeds — unsuitable for the complete mission profile.

Several novel robotic mission concepts have been proposed in response [50] [67], most recently the Exo-Biology Extant Life Surveyor [10] [69]. EELS is a 4.4m, 100kg snake-like robot under development at the NASA Jet Propulsion Laboratory. With 46 articulated degrees of freedom, EELS is capable of potentially infinite gaits for traversing the surface, vertical, and oceanic biomes of Enceladus. Considerable progress has been made towards building specialized controllers for each mode, albeit in isolation. Surface locomotion was demonstrated through a variety of shape-based gaits [66] [47] [17], vertical mobility via screw propulsion [53], and swimming, while not yet demonstrated on EELS, has been studied on other snake-like platforms [28] [23].

However, these control methodologies are disjoint, creating discontinuities at the surface-vent and vent-ocean mode transitions where the robot must seamlessly adapt between different locomotion strategies. To date, entering vertical mobility requires 'placing' EELS in a crevasse with the correct pose, assisted by a wire harness. As highlighted during the 2023 EELS Athabasca field test, "*more work is needed to integrate surface and vertical components into a single robotic system capable of safely executing a full mission from lander to subglacial ocean*" [69, pg 8].

To address this challenge, and inspired by prior work [21], we propose a hierarchical learning-based control architecture. At a low level, reinforcement learning is used to train specialized locomotion policies for *both* surface and vertical mobility tasks by converting waypoints into commanded joint deflections based on proprioceptive observations. A high-level controller would then learn waypoint placement and which locomotion skill to employ[1] As such, this work offers the following contributions:

- An RL framework to produce robust serpentoid surface and vertical mobility over undulating terrain and through fissure-like gaps.
- Simulation experiments that demonstrate a surface-to-vertical mission profile from delivery lander, over terrain, down a fissure-vent opening, and into the lunar interior.

More broadly, this work builds towards the wider, paradigm-shifting idea of Robotic Exploration 3.0 (RE-3.0): *intelligent robots capable of adapting at mission-time, enabling one-shot exploration of celestial bodies without requiring many expeditionary campaigns* [51]. In our case, there are $n = 3$ things we desire our robot to do, but we don't want to send 3 separate robots [27]. Instead, a hyper redundant snake-like morphology, coupled with learnable autonomy, emerges as a viable solution to realizing RE-3.0 by learning surface and vertical locomotion with a single platform and control architecture.

II. Related Works

Learning robust surface and vertical mobility with a single platform has been a long-standing challenge for legged and limbless robots, especially in confined environments that limit actuation range and require precise contact control. To gauge progress in each dimension, fig. 2 informally compares locomotion distance records across prior surface and vertical robotics literature, colored by morphology. We believe such a comparison is useful because it shows perspective when comparing different control architectures relative to traversal distance requirements for direct ocean access: Enceladus's fissures are estimated to reach depths of 6km [49], while propulsive landing systems, we theorize, could provide targeting accuracy within tens of kilometers. Enabling autonomous robotic navigation to meet this spatial envelope remains largely unstudied.

The confined environment within Enceladus's vent system is theorized to constrict towards a narrow throat [32], making a snake-like morphology with small cross-sectional area an appealing choice.

Surface Mobility: At surface level over short spatial distances ($< 100m$), serpentoid locomotion typically employs "*shape-based*" gaits such as open-loop sidewinding [77] or lateral undulation [11]. Here, sinusoidal curves propagate down the backbone to produce forward motion. Common approaches rely on either expertly designed gait equations [52] or parametric approximations to fit the snake body against a reference shape or curve [20] [19] [76]. However, as predominantly open-loop approaches, adaptation is limited, constraining deployment to near-planar terrain.

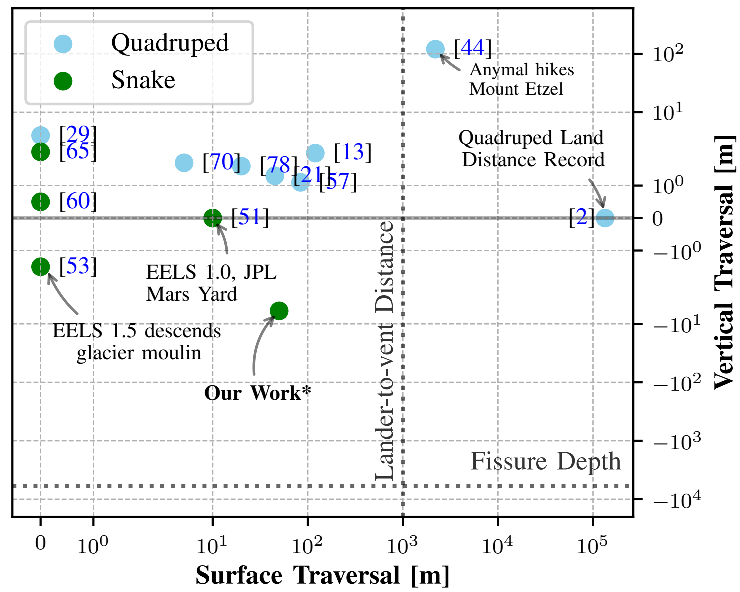

Fig. 2: Informal comparison of vertical and surface traversal distance records across prior robotics literature. The estimated depth of Enceladus's fissures is ≈ 6km [49], we theorize a propulsive lander would have an error margin ≈ 10km. **Asterisk denotes simulation only, no hardware validation.*

Snake robot hyper-redundancy also introduces many degrees of freedom (DOFs), creating high-dimensional action spaces. Prior works have addressed this through dimensionality-reduction techniques: backbone curves approximate the robot's complex body shape using a parameterized spline [20]; multi-agent RL was proposed to improve scalability by treating each mechanical module as an independent agent [75]; and central pattern generators (CPGs) encode gaits using compact frequency, phase shift, and amplitude parameter sets [26, 25, 72]. Augmentation with RL to automatically tune CPG parameters has also been explored [36, 54]. Despite such enhancements, rigidity to a geometric reference makes these techniques fragile when adapting to mobility-stressing terrains such as steep inclinations or undulating surfaces. In these settings, dynamic gait adjustments that conform with local surface curvature become critical.

One promising technique to bridge this gap are end-to-end learning-based approaches, which train a neural network to map pixels [6], tactile feedback [37], and or proprioceptive

[1] As this is ongoing work, the high-level controller remains unimplemented and is represented in by a human operator for now.

state [38] directly to joint positions or velocity commands. However, in all these works, as well as [5] [61], the environment has remained a flat plane. To the best of our knowledge, learning end-to-end shape-based gaits over challenging undulating surfaces has not yet been explored.

Vertical Mobility: Over short spatial distances ($< 10m$), snake robots exhibit remarkable vertical mobility for climbing trees [35] [33], pipes [65], gaps [59], ladders [64], steps [22], and steep inclinations [18]. Of these, internal pipe and gap crawling are most related to our work because they emphasize maintaining precise contact wrenches within narrow friction cones inside a confined environment while making and breaking contact.

Pipe climbing has been explored via LSTM-driven sine wave gaits [60] and arboreal concertina gaits [12], both incorporating geometric formalisms based on watching biological snakes in the wild. To climb walls, a lamprey-inspired gait used microspine anchoring to achieve near-vertical mobility [68]. However, relying on a single anchor is minimally redundant and, on Enceladus, potentially challenging to mount on a cryogenic ice wall.

To our knowledge, no prior work—simulated or real—demonstrates surface *and* vertical mobility capable of crossing complex terrain and descending un-bordered fissure-like spaces (ie, only two surfaces to push against). Direct ocean access requires traversing such spaces, in addition to translating in any lateral direction once inside the fissure.

Insights from Quadrupeds: Although the cross-sectional area of a quadruped is generally too large for Enceladus's fissures, recent innovations in RL for agile quadrupedal parkour have proved highly successful, offering valuable cross-domain insights. Many prior works train policies capable of surface mobility tasks for trotting, running, fall recovery [74], manipulation [43], often over challenging [34] or deformable [14] terrain, and even inside confined spaces [73] [45]; Vertical mobility is also prevalent, often climbing boxes [7], stairs [13], ladders [70], or inclined surfaces. To achieve such feats, a common theme is to first learn specialized policies for each low-level locomotion skill which are then either dynamically selected by a higher-level planner at deployment time [21], or, distilled into a single policy with additional fine-tuning during post-training [56].

Of these methods, we draw inspiration from the hierarchical approach, delegating low-level locomotion to specialized policies and relying on a high-level navigation module - an idea initially proposed in [21]. This approach forms the basis of our method, explained below.

III. Method

Our goal is to endow a snake robot with locomotion capabilities to traverse undulating surface terrain, reach and enter a fissure opening, initiate a controlled descent between two opposing vertical faces, and arrive at a spatial goal in minimal time. Throughout the descent, contact must be maintained with fissure walls to avoid falling, and for completeness, must also learn to climb. At surface level, Enceladus's surface terrain features deep undulations that can entrap the snake robot, thereby necessitating realtime gait adaptability to escape traps.

To accomplish this, we employ a hierarchical reinforcement learning (RL) strategy inspired by prior work [21]. Two low-level locomotion policies are trained to first master surface and vertical mobility independently, both capable of tracking towards any arbitrary goal location in their respective planes (fig. 3). Partitioning locomotion into task-specific policies like this considerably simplifies problem complexity compared to learning both simultaneously. At deployment time, a high-level controller (*for now represented by a human operator*) is responsible for goal placement and for triggering the surface-vertical mode transition upon reaching the fissure entrance[2].

Both low-level surface and vertical locomotion problems have shape-based gait solution policies which can be found through framing as a Markov Decision Processes [63]. During training, the agent starts from an initializing distribution ρ (Table V and VI). At each timestep, the agent receives state $\mathbf{s}_t \in \mathcal{S}$ and samples an action $\mathbf{a}_t \in \mathcal{A}$ from a policy, either $\pi_\theta^{\text{surface or vertical}}(\mathbf{a}|\mathbf{s}_t)$. Upon action execution, the agent transitions to a new state s_{t+1} and receives reward $r_{t+1} \in \mathcal{R}$. Off-the-shelf Proximal Policy Optimization from `RSL` [58] is used to find parameters θ that maximize the expected discounted sum of future reward. A multilayer perceptron with $[512, 256, 128]$ hidden units over three layers and elu activations is used to approximate the optimal π^*.

The same MDP observation and reward spaces are used for training surface and vertical mobility policies. However, we introduce the notion of principle axes $\mathcal{H}$ depending on the task: for surface mobility, $\mathcal{H} = \{x, y\}$ and for vertical, $\mathcal{H} = \{x, z\}$. All policies are trained in `Genesis-World` [4], a GPU-accelerated robotics simulator.

Observations: We define the observation vector $\mathbf{o} \in \mathbb{R}^{63}$ of a low-level policy as:

$$\mathbf{o}_t = \left[\mathbf{d}_{\text{goal}},\ \mathbf{g}_{\text{proj}},\ \mathbf{v}_{\text{com}},\ \mathbf{h},\ h_{\text{dot}},\ d_{\text{goal}},\ \mathbf{q},\ \dot{\mathbf{q}},\ \mathbf{a_{t-1}}\right]^\top \quad (1)$$

TABLE I: Notation

Symbol	Description
$\mathbf{p}_{goal}$	World goal position specified by navigation module.
$\mathbf{p}_{com}$	World snake center of mass
$\mathbf{d}_{\text{goal}}$	Vector from robot center-of-mass to goal (normed, in $\mathcal{H}$).
$\mathbf{g}_{\text{proj}}$	Gravity vector in robot's tail-link frame.
$\mathbf{v}_{\text{com}}$	Linear velocity of center of mass in the world frame.
$\mathbf{h}$	Heading vector of robot, normed, in principle plane $\mathcal{H}$.
h_{dot}	Cosine similarity between heading $\mathbf{h}$ and $\mathbf{d}_{goal}$
d_{goal}	Distance between center of mass $\mathbf{p}_{com}$ and goal $\mathbf{p}_{goal}$.
$\mathbf{q}, \dot{\mathbf{q}}, \tau$	Joint positions, velocities, torques.
$\mathbf{a}_{t-1}$	Previous action taken by the policy.

Importantly, observations exclusively utilize proprioceptive features, nothing is explicitly known about the geometry of surrounding terrain. Surprisingly, we find it's still possible for each policy to adapt its gait to the local environmental purely by knowing joint positions and the previous action. For example, at deployment time, reactive behavior can be observed

[2]In future work, the high-level human 'controller' could be replaced by a specialized task and mission planner similar to [24].

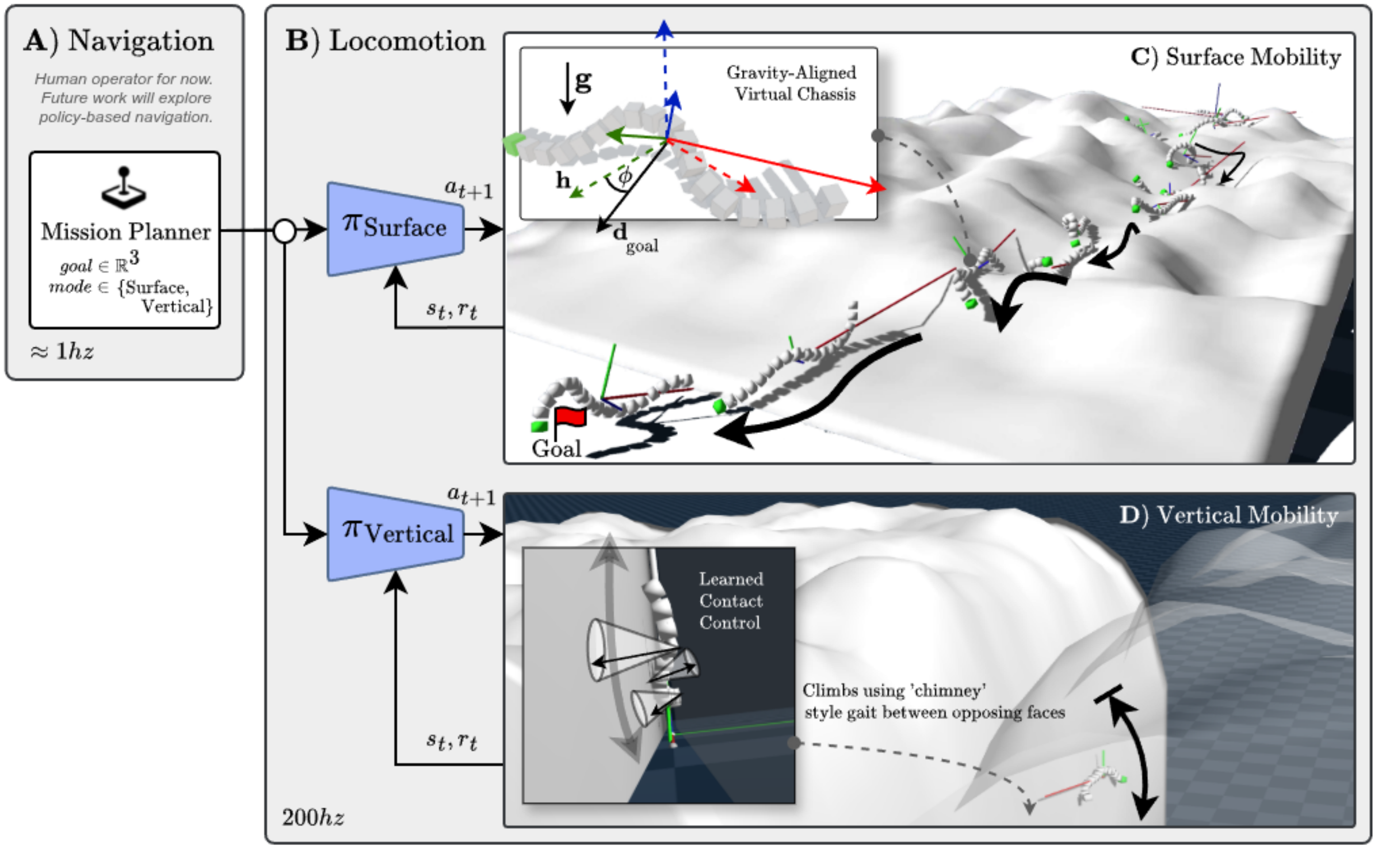

Fig. 3: **Methodology Overview**: Learning surface and vertical snake-robot locomotion. (**A**) Navigation module (human operator) provides high-level waypoint and mode switching commands - *future work aims to make this autonomous*. (**B**) Low-level locomotion module selects task based on commanded mode and converts desired goal into joint actuation. (**C**) Surface task trained to locomote over a flat plane before deployment on undulating terrain and obstacles. (**D**) Vertical task learns to traverse a simplified vent fissure formed between two opposing faces via a chimney gait and learned contact control.

where, upon getting stuck in a deep surface undulation, the agent is able make subtle gait adjustments in realtime to escape the hole and continue on its way.

Reward: Both surface and vertical tasks utilize the following reward function eq. (5). A progress term [62] incentivizes moving the agent's center of mass $\mathbf{p}_{com}$ to the goal $\mathbf{p}_{goal}$ in minimal time. Energy rewards minimize the magnitude of actions $\mathbf{a}_t$, ensuring smoothness and efficiency. Components are linearly scaled by $\alpha_1 = 1, \alpha_2 = 0.001$, obtained by manual tuning, and Δt the integration timestep.

$$v_t = \frac{\|\mathbf{p}_{\text{goal}} - \mathbf{p}_{\text{com}}\|_2}{\Delta t} \tag{2}$$

$$r_{\text{progress}} = v_{t-1} - v_t \tag{3}$$

$$r_{\text{energy}} = -\|\mathbf{a}_t\|_2^2 \tag{4}$$

$$r_t = \alpha_1 r_{\text{progress}} + \alpha_2 r_{\text{energy}} \tag{5}$$

Note this simplified design over the reward function in our previous work [47], dispensing with safety and heading components and featuring two terms rather than four. In hindsight, penalizing heading likely complicated learning by prohibiting agents from freely moving in a given direction with any orientation, an otherwise defining characteristic of robust shape-based gaits.

Virtual Chassis: Similar to our prior work [47], we build upon the established *virtual chassis* [55] for snake robots as our primary reference frame. However, a subtle but significant modification is made by ensuring the z up-axis is gravity aligned, finding this increases robustness in the presence of rolling gaits. This change also ensures the y-axis remains parallel to the ground, which forms an intuitive reference for the snake's orientation heading $\mathbf{h}$ in eq. (1).

Terminal Conditions: Agents are terminated after n timesteps. To avoid artificially correlated states during training, the timestep index of each vectorized environment is randomized before training. If an agent reaches a goal before episode termination, a new goal is uniformly sampled. See appendix parameters for randomization details.

Terrain Generator: To facilitate training and deployment in simulation, a parameterized terrain and fissure model was developed. The geometry is based upon [30], modelling tiger stripes as tidally-flexed vertical slots that puncture the ice shell. Due to Enceladus's 33-hour eccentric orbit around Saturn, tidal forces are imparted on the crust, causing compression and expansion at the vent interface. The width of our fissure is therefore variable. To further increase difficulty, layered perlin noise was imposed, creating surface undulations [16]. However, for vertical mobility, we found it necessary to

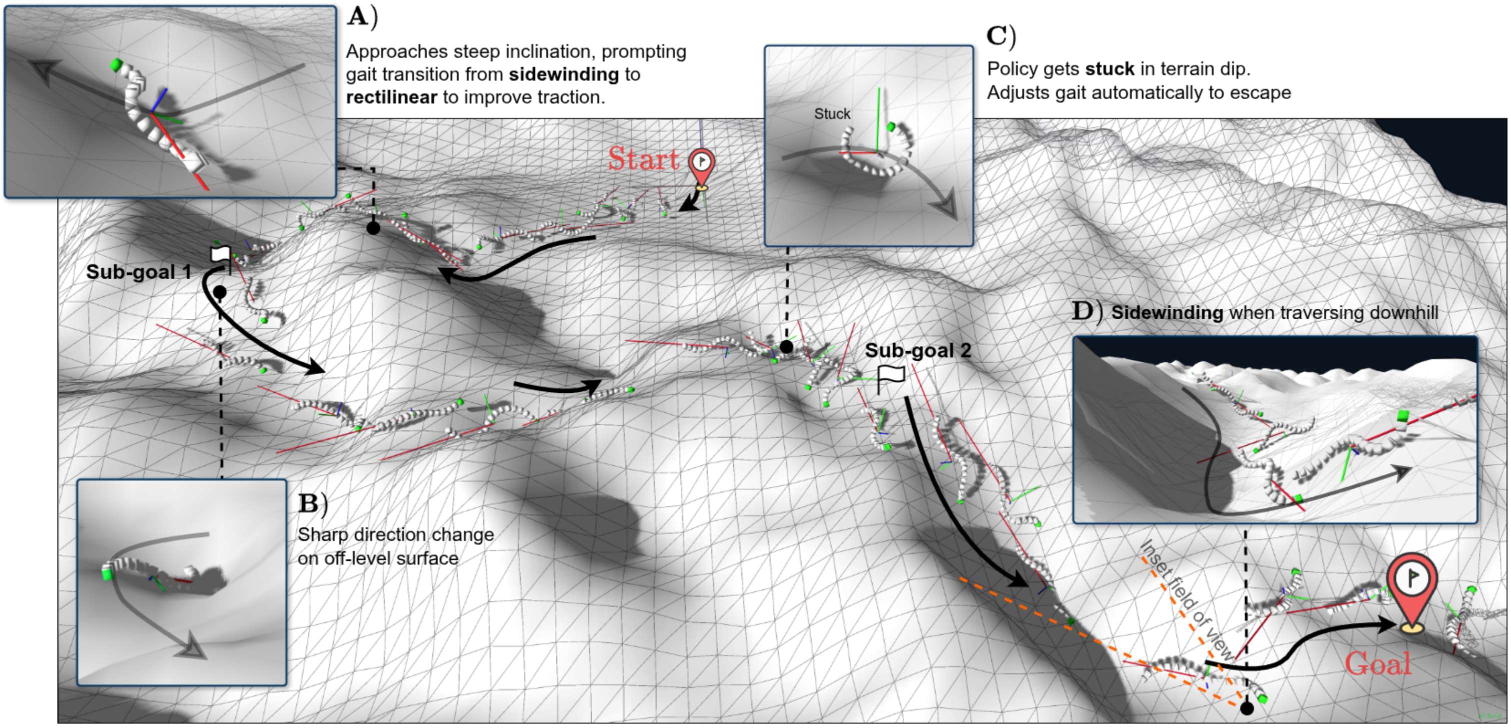

Fig. 4: **Surface Mobility**: Evaluation of surface mobility policy over randomly-generated perlin noise terrain. Starting from green flag, the agent (**A**) encounters a steep slope and adjusts gait from rectilinear to sidewinding, improving contract traction. (**B**) Performs a direction change on an off-camber slope after reaching first sub-goal. (**C**) Gets stuck in a pit. After a period of re-adjustment, policy switches gate to escape. (**D**) Enters valley, traverses off-planar surface and climbs out to reach the final goal. See Video S1 in table II for the full rollout.

increase the static and coupling terrain friction coefficient to 5 and 2 respectively. These values are far greater than the anticipated low-friction icy fissure walls of Enceladus and constitute a key limiting assumption.

IV. Results & Discussion

Our method is evaluated in simulation experiments that evaluate robustness of our learned surface and vertical mobility policies.

Surface Mobility: Figure 4 demonstrates mobility over randomized perlin terrain featuring steep undulations. Evidently, the policy can robustly translate towards specified sub-goals (white flags) while adapting its gait relative to proprioceptive observations. For example, at point A, a steep inclination is reached, prompting an apparent gait transition from sidewinding to rectilinear. We speculate this increases contact patch area and improves traction, a behavior also observed in biological snakes [40]. Similar adaptability was shown at point C, where upon briefly getting stuck in a pit, the policy appears to reconfigure itself to climb out and continue on its way. Video S1 in Table II shows full the rollout.

Note this learned surface mobility policy is blind, with no mapping capabilities to provide knowledge of surrounding terrain. The chosen gait must therefore be reasoned using proprioceptive feedback alone (eq. (1)). Despite this, robust surface mobility is still possible, reflecting a broader trend amongst other blind locomotion literature [34]. This highlights the promising potential of a hyper-redundant snake morphology coupled with learned autonomy for space exploration.

Surface-to-Vertical: Figure 5 demonstrates deployment on a simplified surface-to-vertical mobility task. After surface deployment (presumably from a delivery lander) the agent traverses undulating terrain, featuring inclinations of up to ≈ 20 degrees, while tracking towards the vent interface. After reaching the vent, the agent passively slides into the fissure, catches itself, and adopts a '*chimney*' climbing gait, enabling a controlled descent towards the fissure floor 4m below. By passive we mean the joint effort is zero.

Note, the passive slide into the fissure is not learned, but rather triggered by the navigation module and made easier by the smooth chamfered edges of the surface-vent opening. Further, the vent-opening itself features off-distribution terrain. During testing, we found the policy can occasionally 'fall' into a novel state and stall. Successfully entering the fissure also depends on the human operator's skill.

After clearing the vent opening and initiating a descent, the vertical policy systematically makes and breaks contact by pushing against the fissure walls in a rolling motion to maintain control. In prior works, this kind of coordinated behavior has, generally, only been possible through expertly designed shape-based gaits. Here, climbing behavior is emergent - we did not explicitly incentivize attributes such as

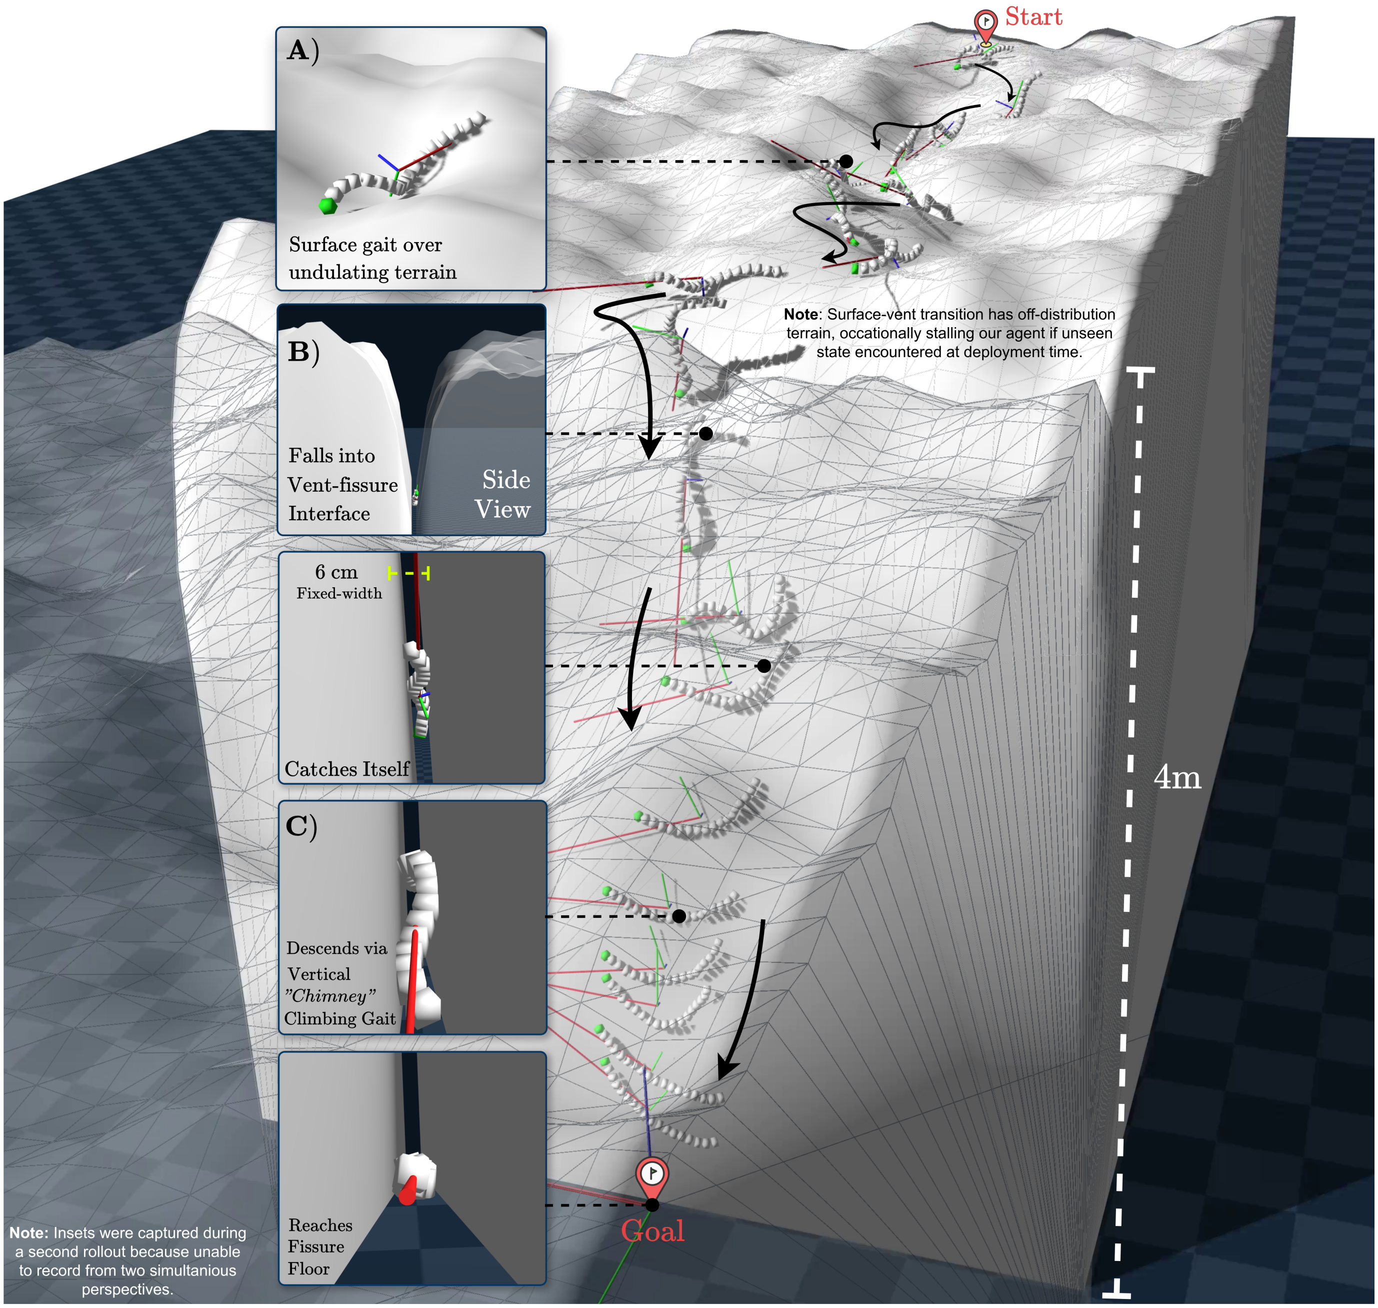

Fig. 5: **Surface-to-Vertical Mobility**: Deployment of our method on a surface and vertical mobility task. (**A**) The agent traverses over undulating surface terrain, featuring inclinations of up to ≈ 20 degrees, while tracking towards the vent interface. (**B**) Agent passively slides into the vent (joint effort zero), catches itself and (**C**) adopts a '*chimney*' climbing gait, enabling a controlled descent towards the fissure floor 4m below. Note, the surface-vent transition features off-distribution terrain which occasionally stalls our agent if an unseen state is encountered at deployment time.

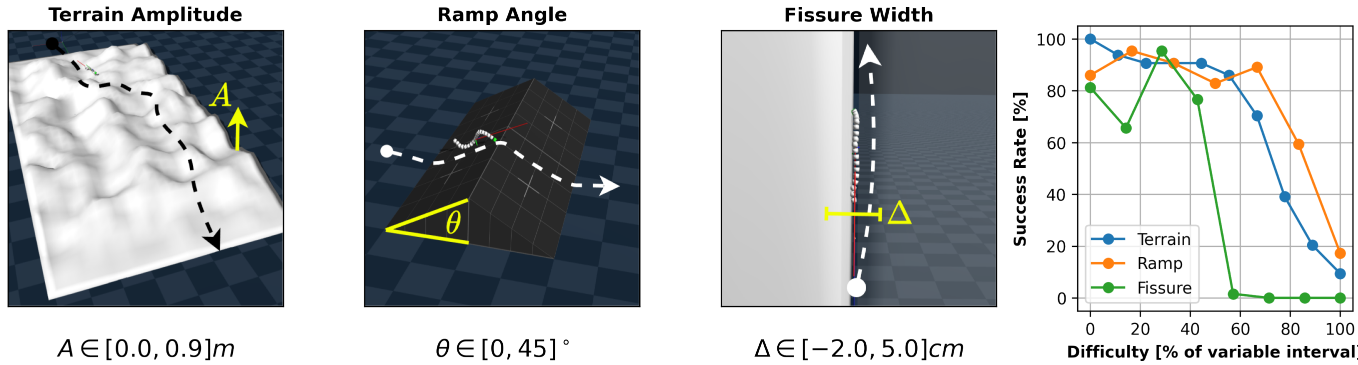

Fig. 6: **Policy Robustness**: Success rate of locomotion policy subject to increasing environment difficulty across three key features: (**A**) terrain perlin noise amplitude, (**B**) ramp angle, and (**C**) fissure gap width. Each data-point contains 64 rollouts.

sinusoidal motion or contact control (eq. (5)). Moreover, while this figure only shows a descent, our method is also capable of translating towards any point within the vertical fissure plane, including climbing and lateral translations (See Video S4 in Table II).

Figure 7 compares the distribution of joint angles during these surface and vertical mobility rollouts shown in fig. 5. Interestingly, the surface joint distribution (left) makes full use of the actuation range up to and including the revolute axis hard-stops at $\pm 45^\circ$. The distribution of joint angles is also symmetric, reflecting the symmetry, on average, of dominantly sinusoidal surface mobility gaits such as sidewinding and rectilinear motion. In contrast, the vertical mobility joint distribution (right) is asymmetric, reflecting the banana-like chimney climbing gait observed in fig. 5.

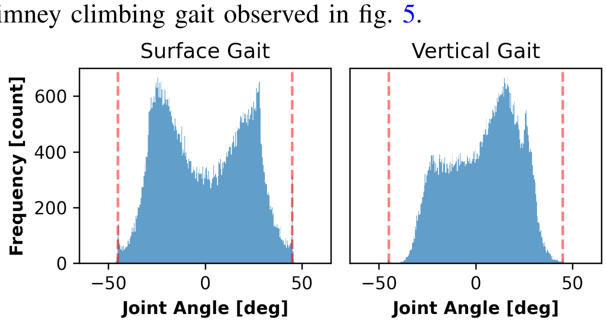

Fig. 7: (**Joint Distribution**): Average distribution of joint positions during surface (left) and vertical (right) mobility, for the rollout shown in fig. 5. Actuator hardstops marked by dashed lines.

Robustness: Policy robustness was evaluated by varying environment properties and measuring the resultant success percentage when traversing from a start to goal point, repeated over 64 rollouts. Terrain undulation difficulty was controlled by increasing perlin noise amplitude from 0m to 0.9m, inclined surfaces by increasing the gradient of a linear ramp obstacle up to 45°, and vertical mobility by varying the width of the fissure gap by -2cm and $+5$cm either side of that used during training (fig. 6).

In simulation, the surface policy is robust to undulations up to 0.6m in amplitude, before success rate falls of significantly. Similarly, inclinations up to $\approx 32^\circ$ can be traversed before performance degrades. Both characteristics are a marked improvement over our prior work [47], which could only manage a flat plane. However, the vertical policy appears to be very sensitive to fissure gap width. The policy is unable to complete even a single successful rollout for width variations beyond ± 1cm, suggestive of over-fitting (fig. 6 green).

Gait Analysis: Gait trajectories produced by our surface mobility policy exhibit the same turning mechanisms used by sidewinder rattlesnakes, previously observed in [3]. In biological snakes, *differential turns* are long, shallow course corrections that can continue over many gait cycles while *reversal turns* are sudden, sharp in nature for rapid direction changes [3].

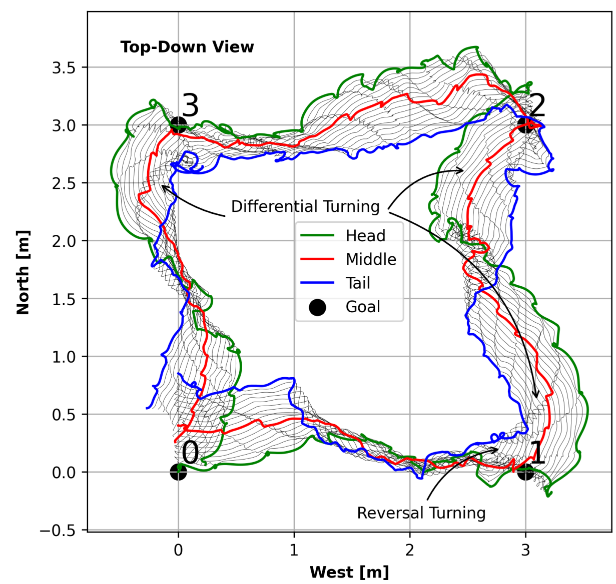

Fig. 8: **Gait Analysis**: Shape-based gait tracking waypoints in 0-1-2-3-0 sequence, reminiscent of sidewinder rattlesnake locomotion through differential and reversal turns.

Our policy exhibits both differential turning after passing waypoint 3, and a sharp reversal turn while approaching and passing point 1 (fig. 8). In our experiments, the internal policy mechanism that triggers which turning technique to be used appears subject to the snakes heading $\mathbf{h}$ in eq. (1) and the relative position of the goal.

V. Conclusions & Future Work

Motivated by Enceladus direct ocean access [9], we developed an RL controller capable of surface locomotion over undulating terrain (fig. 4) and vertical mobility in fissure-like environments (fig. 5). Robustness analysis (fig. 6) suggests our controller can handle variations in simulated terrain geometry and surface inclinations, but is sensitive to small ($\pm$1cm) changes in fissure width beyond those seen in training. We also observed an interesting parallel between the turning mechanism of real sidewinder desert snakes and our surface mobility policy (fig. 8).

However, to achieve the vertical mobility evident in fig. 5, several simplifications must be noted. First, the static and coupling friction coefficients of the fissure surfaces were increased to 5 and 2 respectively. This is unrealistically beyond the low-friction icy surfaces expected on Enceladus. Second, the navigation module, responsible for placing sub-goals at the vent entrance and fissure floor, and for triggering the low-level mode transition, was managed by a skilled human operator. Lastly, the entrance section of the fissure features off-distribution terrain not seen during training, occasionally stalling the vertical mobility policy if an unseen state is encountered at deployment time.

Future work could investigate further increasing vertical mobility robustness subject to increased fissure widths and reduced friction coefficients. Additional work is needed to replace our human navigation module with an autonomous equivalent, perhaps using the task and motion planner proposed in [69]. Lastly, quantifying the sim-to-real gap and deploying of our method on real hardware is an intriguing next step.

More broadly, we believe developing learning-based control systems for space robots will help realize the vision of Robotic Exploration 3.0 [51]. As humanity turns its gaze to the outer-solar system, ever-increasing communication delays demand ever-increasing levels of autonomy. This shifts mission-critical decision making from human operators on Earth to on-board the robotic explorers themselves. Learnable autonomy for robotic exploration opens the door to celestial destinations previously outside of our reach, expanding the frontier of human knowledge to new horizons in our solar system and beyond.

VI. Appendix

A. Supplementary Videos

The following supplementary videos are provided to support our results:

Behavior	Source
S1: Surface mobility (as in fig. 4)	YouTube
S2: Surface-to-vertical (as in fig. 5)	YouTube
S3: Vertical Mobility (timelapse)	YouTube
S4: Vertical Mobility (closeup)	YouTube

TABLE II: Supplementary Videos

B. Hyperparameters

Parameter	Value
Number of modules	18
Module Mass (kg)	10
Module Dimensions (m)	0.0325 x 0.04 x 0.04
Module I_{xx}	0.0133
Module I_{yy}	0.0175
Module I_{zz}	0.0175
Revolute Axes	$[0, 0, 1]$ if even, else [0,1,0]
Joint Force Range (N)	$[-100, 100]$
Stiffness K_p	300
Damping K_v	10

TABLE III: Snake robot configuration

Parameter	Value
Simulator Name	`genesis-world`
Release Version	`v0.2.1`
Integration Step (s)	0.005
Integration Number of substeps	2
Gravity (ms-2)	[0,0,-9.81]
Constraint Solver	Newton

TABLE IV: Simulator configuration

Parameter	Value
Roll	$\mathcal{U}(-\pi, \pi)$
Pitch	$\mathcal{U}(0, 0)$
Yaw	$\mathcal{U}(-\pi, \pi)$
Initial Joint Positions (rad)	$\mathcal{U}(0, 0)$
Initial Joint Velocities (rads-1)	$\mathcal{U}(0, 0)$
Goal $\mathbf{g} \in \mathbb{R}^2$	$g_x \sim \mathcal{U}(-5, 5), g_y \sim \mathcal{U}(-5, 5)$
OnGoalReached(...)	Wait for episode termination
Training Scene	Flat Plane

TABLE V: Initial State Distribution (surface mobility)

Parameter	Value
Roll	$\mathcal{U}(-\pi, \pi)$
Pitch	$\mathcal{U}(0, 0)$
Yaw	$\mathcal{U}(0, 0)$
Initial Joint Positions (rad)	$\mathcal{U}(0, 0)$
Initial Joint Velocities (rads-1)	$\mathcal{U}(0, 0)$
Goal $\mathbf{g} \in \mathbb{R}^2$	$g_x \sim \mathcal{U}(-2, 2), g_z \sim \mathcal{U}(1, 4)$
OnGoalReached(...)	Sample new goal
Training Scene	Vertically oriented, opposing planes 7cm apart

TABLE VI: Initial State Distribution (vertical mobility)

References

[1] Lecture: Exobiology Extant Life Surveyor (EELS) Robotic Architecture, January 2025. URL https://www.kiss.caltech.edu/lectures/2019_EELS.html.

[2] Farthest distance covered by a quadruped robot, June 2025. URL https://www.guinnessworldrecords.com/world-records/90801-farthest-distance-covered-by-a-quadruped-robot.

TABLE VII: Learning algorithm hyperparameters, adapted from `Legged Robotics RSL RL`

Parameter	Value
Algorithm	PPO
Clip parameter	0.2
Desired KL-Divergence	0.01
Entropy coefficient	0.01
Discount Factor	0.99
Lambda	0.95
Learning rate	0.001
Max gradient norm	1.0
Num learning epochs	5
Num mini batches	4
Learning Rate Schedule	adaptive
Clipped value loss	true
Value loss coefficient	1.0
Activations	elu
Actor hidden dims	[512, 256, 128]
Critic hidden dims	[512, 256, 128]
Initial noise std	1.0

[3] Henry C. Astley, Chaohui Gong, Jin Dai, Matthew Travers, Miguel M. Serrano, Patricio A. Vela, Howie Choset, Joseph R. Mendelson, David L. Hu, and Daniel I. Goldman. Modulation of orthogonal body waves enables high maneuverability in sidewinding locomotion. *Proc. Natl. Acad. Sci. U.S.A.*, 112(19):6200–6205, May 2015. doi: 10.1073/pnas.1418965112.

[4] Genesis Authors. Genesis: A universal and generative physics engine for robotics and beyond, December 2024. URL https://github.com/Genesis-Embodied-AI/Genesis.

[5] Zhenshan Bing, Christian Lemke, Zhuangyi Jiang, Kai Huang, and A. Knoll. Energy-efficient slithering gait exploration for a snake-like robot based on reinforcement learning, 2019.

[6] Zhenshan Bing, Christian Lemke, Fabric O. Morin, Zhuangyi Jiang, Long Cheng, Kai Huang, and Alois Knoll. Perception-Action Coupling Target Tracking Control for a Snake Robot via Reinforcement Learning. *Front. Neurorob.*, 14:591128, October 2020. ISSN 1662-5218. doi: 10.3389/fnbot.2020.591128.

[7] Ken Caluwaerts, Atil Iscen, J. Chase Kew, Wenhao Yu, Tingnan Zhang, Daniel Freeman, Kuang-Huei Lee, Lisa Lee, Stefano Saliceti, Vincent Zhuang, Nathan Batchelor, Steven Bohez, Federico Casarini, Jose Enrique Chen, Omar Cortes, Erwin Coumans, Adil Dostmohamed, Gabriel Dulac-Arnold, Alejandro Escontrela, Erik Frey, Roland Hafner, Deepali Jain, Bauyrjan Jyenis, Yuheng Kuang, Edward Lee, Linda Luu, Ofir Nachum, Ken Oslund, Jason Powell, Diego Reyes, Francesco Romano, Feresteh Sadeghi, Ron Sloat, Baruch Tabanpour, Daniel Zheng, Michael Neunert, Raia Hadsell, Nicolas Heess, Francesco Nori, Jeff Seto, Carolina Parada, Vikas Sindhwani, Vincent Vanhoucke, and Jie Tan. Barkour: Benchmarking Animal-level Agility with Quadruped Robots. *arXiv*, May 2023. doi: 10.48550/arXiv.2305.14654.

[8] Kalind Carpenter, Morgan L. Cable, Mathieu Choukroun, Saverio Iacoponi, Scott Moreland, Matt Heverly, Hari Nayar, Brett Kennedy, Joseph Bowkett, and William Reid. The Science Case for Surface Mobility on Icy Worlds. *Bulletin of the American Astronomical Society*, 53(4):186, May 2021. doi: 10.3847/25c2cfeb.40bebb9a.

[9] Kalind Carpenter, Morgan L. Cable, Mathieu N. Choukroun, Hiro Ono, Rohan A. Thakker, Michel D. Ingham, Patrick McGarey, Ansel Barchowsky, Saverio Iacoponi, Joseph J. Bowkett, William Reid, and Eloise Marteau. Venture Deep, the Path of Least Resistance: Crevasse-Based Ocean Access Without the Need to Dig or Drill. *Bulletin of the AAS*, 53(4), March 2021. doi: 10.3847/25c2cfeb.0183ab72.

[10] Kalind Carpenter, Andrew Thoesen, Darwin Mick, Justin Martia, Morgan Cable, Karl Mitchell, Sarah Hovsepian, Jay Jasper, Nikola Georgiev, Rohan Thakker, Ara Kourchians, Brian Wilcox, Michael Yip, and Hamid Marvi. Exobiology Extant Life Surveyor (EELS). *Proceedings*, pages 328–338, April 2021. doi: 10.1061/9780784483374.033.

[11] A. H. Chang, M. M. Serrano, and P. A. Vela. Shape-centric modeling of lateral undulation and sidewinding gaits for Snake Robots. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 12–14. IEEE. doi: 10.1109/CDC.2016.7799297.

[12] Shuoqi Chen and Aaron Roth. Gait Design of a Novel Arboreal Concertina Locomotion for Snake-like Robots. *arXiv*, September 2023. doi: 10.48550/arXiv.2309.06000.

[13] Xuxin Cheng, Kexin Shi, Ananye Agarwal, and Deepak Pathak. Extreme Parkour with Legged Robots. *arXiv*, September 2023. doi: 10.48550/arXiv.2309.14341.

[14] Suyoung Choi, Gwanghyeon Ji, Jeongsoo Park, Hyeongjun Kim, Juhyeok Mun, Jeong Hyun Lee, and Jemin Hwangbo. Learning quadrupedal locomotion on deformable terrain. *Sci. Rob.*, 8(74), January 2023. ISSN 2470-9476. doi: 10.1126/scirobotics.ade2256.

[15] David Deamer and Bruce Damer. Can Life Begin on Enceladus? A Perspective from Hydrothermal Chemistry. *Astrobiology*, 17(9):834, September 2017. doi: 10.1089/ast.2016.1610.

[16] Casey Duncan. Project title. https://github.com/caseman/noise, 2018.

[17] Nikola Georgiev, Torkom Pailevanian, Eric Ambrose, Avak Archanian, Hovhannes Melikyan, and Daniel Loret de Mola Lemu. EELS: A Modular Snake-like Robot Featuring Active Skin Propulsion, Designed for Extreme Icy Terrains. In *2024 IEEE Aerospace Conference*, pages 02–09. IEEE. doi: 10.1109/AERO58975.2024.10521082.

[18] Chaohui Gong, Matthew Tesch, David Rollinson, and Howie Choset. Snakes on an inclined plane: Learning an adaptive sidewinding motion for changing slopes. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 14–18. IEEE. doi: 10.1109/IROS.2014.6942697.

[19] Chaohui Gong, R. Hatton, and H. Choset. Conical sidewinding. *IEEE International Conference on Robotics and Automation*,

2012. URL https://www.semanticscholar.org/paper/Conical-sidewinding-Gong-Hatton/ae3f141cc7d4f324fb325e5e091c9c6963818ad6.

[20] Ross L. Hatton and Howie Choset. Generating gaits for snake robots by annealed chain fitting and Keyframe wave extraction. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 10–15. IEEE. doi: 10.1109/IROS.2009.5354257.

[21] David Hoeller, Nikita Rudin, Dhionis Sako, and Marco Hutter. ANYmal Parkour: Learning Agile Navigation for Quadrupedal Robots. *arXiv*, June 2023. doi: 10.48550/arXiv.2306.14874.

[22] Yuta Iguchi, Mizuki Nakajima, Ryo Ariizumi, and Motoyasu Tanaka. Step Climbing Control of Snake Robot with Prismatic Joints. *Sensors*, 22(13):4920, June 2022. ISSN 1424-8220. doi: 10.3390/s22134920.

[23] A. J. Ijspeert. AmphiBot II: An Amphibious Snake Robot that Crawls and Swims using a Central Pattern Generator. *Proceedings of the 9th International Conference on Climbing and Walking Robots (CLAWAR 2006)*, pages 19–27, 2006. URL https://infoscience.epfl.ch/entities/publication/4b556957-1f98-448e-8377-ad9b17ee20d8.

[24] Ashkan Jasour, Guglielmo Daddi, Masafumi Endo, Tiago S. Vaquero, Michael Paton, Marlin P. Strub, Sabrina Corpino, Michel Ingham, Masahiro Ono, and Rohan Thakker. Risk-aware Integrated Task and Motion Planning for Versatile Snake Robots under Localization Failures. *arXiv*, February 2025. doi: 10.48550/arXiv.2502.19690.

[25] Shuo Jiang, Adarsh Salagame, Alireza Ramezani, and Lawson Wong. Snake Robot with Tactile Perception Navigates on Large-scale Challenging Terrain. *arXiv*, December 2023. doi: 10.48550/arXiv.2312.03225.

[26] Shuo Jiang, Adarsh Salagame, Alireza Ramezani, and Lawson Wong. Hierarchical RL-Guided Large-scale Navigation of a Snake Robot. *arXiv*, December 2023. doi: 10.48550/arXiv.2312.03223.

[27] Jplraw. EELS: Robotic Exploration 3.0 - 2024 Conference Talk, April 2024. URL https://www.youtube.com/watch?v=rxzsJX3jNR8&t=1390s.

[28] E. Kelasidi, K. Y. Pettersen, J. T. Gravdahl, S. Str⊘ms⊘yen, and A. J. S⊘rensen. Modeling and propulsion methods of underwater snake robots. In *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 27–30. IEEE. doi: 10.1109/CCTA.2017.8062561.

[29] Brett Kennedy, Avi Okon, Hrand Aghazarian, Mircea Badescu, Xiaoqi Bao, Yoseph Bar-Cohen, Zensheu Chang, Borna E. Dabiri, Mike Garrett, Lee Magnone, and Stewart Sherrit. Lemur IIb: a Robotic System for Steep Terrain Access. In *Climbing and Walking Robots*, pages 1077–1084. Springer, Berlin, Germany, 2005. ISBN 978-3-540-26415-6. doi: 10.1007/3-540-26415-9_129.

[30] Edwin S. Kite and Allan M. Rubin. Sustained eruptions on Enceladus explained by turbulent dissipation in tiger stripes. *arXiv*, May 2016. doi: 10.1073/pnas.1520507113.

[31] NASA Jet Propulsion Laboratory. Exobiology Extant Life Surveyor (EELS) Robotic Architecture , 2019. URL https://youtu.be/7bdS_xpYz7A?t=930.

[32] NASA Jet Propulsion Laboratory. Exobiology Extant Life Surveyor (EELS) Concept of Operations on Enceladus , June 2022. URL https://youtu.be/e0D9IVo-E9M?t=90.

[33] Tin Lun Lam and Yangsheng Xu. A flexible tree climbing robot: Treebot - design and implementation. In *2011 IEEE International Conference on Robotics and Automation*, pages 09–13. IEEE, 2011. doi: 10.1109/ICRA.2011.5979833.

[34] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Sci. Rob.*, 5(47), October 2020. ISSN 2470-9476. doi: 10.1126/scirobotics.abc5986.

[35] Zhen Li, Guizhi Yang, Yuhuai Liu, Shengnian Lin, and Qiaoting Xu. Design of a Climbing Snake-Like Robot with Multi-Sensor Fusion. In *2024 11th International Forum on Electrical Engineering and Automation (IFEEA)*, pages 22–24. IEEE, 2024. doi: 10.1109/IFEEA64237.2024.10878733.

[36] Xuan Liu, Renato Gasoto, Ziben Jiang, C. Onal, and Jie Fu. Learning to locomote with artificial neural-network and cpg-based control in a soft snake robot, 2020.

[37] Xuan Liu, Cagdas D. Onal, and Jie Fu. Integrating Contact-aware Feedback CPG System for Learning-based Soft Snake Robot Locomotion Controllers. *arXiv*, September 2023. doi: 10.48550/arXiv.2309.02781.

[38] Yilang Liu and A. Farimani. An energy-saving snake locomotion gait policy using deep reinforcement learning, 2021.

[39] Jonathan I. Lunine. Ocean worlds exploration. *Acta Astronaut.*, 131:123–130, February 2017. ISSN 0094-5765. doi: 10.1016/j.actaastro.2016.11.017.

[40] Hamidreza Marvi, Chaohui Gong, Nick Gravish, Henry Astley, Matthew Travers, Ross L. Hatton, Joseph R. Mendelson, Howie Choset, David L. Hu, and Daniel I. Goldman. Sidewinding with minimal slip: Snake and robot ascent of sandy slopes. *Science*, 346(6206):224–229, October 2014. ISSN 0036-8075. doi: 10.1126/science.1255718.

[41] Jaret B. Matthews and Issa A. Nesnas. On the design of the Axel and DuAxel rovers for extreme terrain exploration. In *2012 IEEE Aerospace Conference*, pages 03–10. IEEE. doi: 10.1109/AERO.2012.6187039.

[42] Gareth Meirion-Griffith, Michael Garrett, Joseph Bowkett, Brendan Chamberlain-Simon, Blair Emanuel, Sisir Karumanchi, and William Reid. Actively Articulated Wheel-on-Limb Mobility for Traversing Europa Analogue Terrain. *Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2019*, August 2019. URL https://ntrs.nasa.gov/citations/20210016564.

[43] Russell Mendonca, Emmanuel Panov, Bernadette Bucher, Jiuguang Wang, and Deepak Pathak. Continuously Improving Mobile Manipulation with Autonomous Real-World RL. *arXiv*, September 2024. doi: 10.48550/arXiv.2409.20568.

[44] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *arXiv*, January 2022. doi: 10.1126/scirobotics.abk2822.

[45] Takahiro Miki, Joonho Lee, Lorenz Wellhausen, and Marco Hutter. Learning to walk in confined spaces using 3D representation. *arXiv*, February 2024. doi: 10.48550/arXiv.2403.00187.

[46] K. L. Mitchell, M. Ono, C. Parcheta, and S. Iacoponi. Dynamic pressure at enceladus' vents and implications for vent and conduit in-situ studies. In *Lunar and Planetary Science Conference*, volume XLVIII of *Lunar and Planetary Science*, page 2801, Houston, 2017. Lunar and Planetary Institute. LPI Contribution No. 2801.

[47] Jack Naish, Jacob Rodriguez, Jenny Zhang, Bryson Jones, Guglielmo Daddi, Andrew Orekhov, Rob Royce, Michael Paton, Howie Choset, Masahiro Ono, and Rohan Thakker. Reinforcement learning-driven parametric curve fitting for snake robot gait design. In Alessandro Abate, Mark Cannon, Kostas Margellos, and Antonis Papachristodoulou, editors, *Proceedings of the 6th Annual Learning for Dynamics; Control Conference*, volume 242 of *Proceedings of Machine Learning Research*, pages 1715–1727. PMLR, Jul 2024. URL https://proceedings.mlr.press/v242/naish24a.html.

[48] Engineering National Academies of Sciences. *Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032*. April 2022. ISBN 978-0-309-47578-5. doi: 10.17226/26522.

[49] Francis Nimmo. Solving the puzzle of Enceladus's active south pole. *Proc. Natl. Acad. Sci. U.S.A.*, 117(28):16107–16108, July 2020. doi: 10.1073/pnas.2011055117.

[50] Hiro Ono and Karl Kitchel. Enceladus Vent Explorer Concept - 2016 NIAC Phase I Study Journey to the Center of Icy Moons. *NTRS NASA*, January 2016. doi: 10.48550/arXiv.2403.03848.

[51] Masahiro Ono, Rohan Thakker, Nikola Georgiev, Peter Gavrilov, Avak Archanian, Tomas Drevinskas, Guglielmo Daddi, Michael Paton, Hovhannes Melikyan, Torkom Pailevanian, Christopher Lopez, Eric Ambrose, Bryson K. Jones, Phillipe Tosse, Matthew Gildner, Benjamin Hockman, Daniel Loret de Mola Lemus, Daniel Pastor Moreno, Tristan Hasseler, Yashwanth Kumar Nakka, Eloise Marteau, Benjamin Nuernberger, Martin Peticco, Morgan Cable, Pedro Proenca, Mike Malaska, Joseph Bowkett, Ashkan Jasour, Michel Ingham, Jeremy Nash, Dan Balentine, Ansel Barchowsky, Fredrik Bevreng, Kyle Botteon, Matthew Caballero, Kalind Carpenter, Mark Chodas, Adriana Daca, Jason Feldman, Alex Gardner, Austen Goddu, Abhinandan Jain, Curtis Jin, Maisha Khanum, Richard Kornfeld, Gary Mark, Benjamin Morell, Jack Naish, William Reid, and Rachel Etheredge. *To Boldly Go Where No Robots Have Gone Before – Part 1: EELS Robot to Spearhead a New One-Shot Exploration Paradigm with in-situ Adaptation*. doi: 10.2514/6.2024-1746. URL https://arc.aiaa.org/doi/abs/10.2514/6.2024-1746.

[52] Tony Owen. Biologically Inspired Robots: Snake-Like Locomotors and Manipulators by Shigeo Hirose Oxford University Press, Oxford, 1993, 220pages, incl. index (£40). *Robotica*, 12(3):282, May 1994. ISSN 1469-8668. doi: 10.1017/S0263574700017264.

[53] Michael Paton, Richard Rieber, Sarah Cruz, Matt Gildner, Chantelle Abma, and Kevin Abma. 2023 EELS Field Tests at Athabasca Glacier as an Icy Moon Analogue Environment. In *2024 IEEE Aerospace Conference*, pages 02–09. IEEE. doi: 10.1109/AERO58975.2024.10521174.

[54] K. Qiu, Hang Zhang, Yikai Lv, Yunkai Wang, Chunlin Zhou, and R. Xiong. Reinforcement learning of serpentine locomotion for a snake robot, 2021.

[55] David Rollinson and Howie Choset. Virtual chassis for snake robots. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 25–30. IEEE. doi: 10.1109/IROS.2011.6094645.

[56] N. Rudin, Junzhe He, Joshua Aurand, and Marco Hutter. Parkour in the wild: Learning a general and extensible agile locomotion policy using multi-expert distillation and rl fine-tuning, 2025.

[57] Nikita Rudin, David Hoeller, Marko Bjelonic, and Marco Hutter. Advanced Skills by Learning Locomotion and Local Navigation End-to-End. *arXiv*, September 2022. doi: 10.48550/arXiv.2209.12827.

[58] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 91–100. PMLR, 2022. URL https://proceedings.mlr.press/v164/rudin22a.html.

[59] Amir Shapiro, Aaron Greenfield, and Howie Choset. Frictional Compliance Model Development and Experiments for Snake Robot Climbing. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 574–579, April 2007. ISSN 1050-4729. doi: 10.1109/ROBOT.2007.363048.

[60] Ouyang Sheng and Wu Wei. Flexible adaptive control of snake-like robot based on lstm and gait, 2020.

[61] Junyao Shi, Tony Dear, and Scott David Kelly. Deep Reinforcement Learning for Snake Robot Locomotion. *IFAC-PapersOnLine*, 53(2):9688–9695, January 2020. ISSN 2405-8963. doi: 10.1016/j.ifacol.2020.12.2619.

[62] Yunlong Song, Mats Steinweg, Elia Kaufmann, and Davide Scaramuzza. Autonomous Drone Racing with Deep Reinforcement Learning. *arXiv*, March 2021. doi: 10.48550/arXiv.2103.08624.

[63] Richard S. Sutton and Andrew G. Barto. *Reinforce-*

ment Learning: An Introduction. The MIT Press, second edition, 2018. URL http://incompleteideas.net/book/the-book-2nd.html.

[64] Tatsuya Takemori, Motoyasu Tanaka, and Fumitoshi Matsuno. Ladder Climbing with a Snake Robot. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 01–05. IEEE. doi: 10.1109/IROS.2018.8594411.

[65] Tatsuya Takemori, Motoyasu Tanaka, and Fumitoshi Matsuno. Gait Design for a Snake Robot by Connecting Curve Segments and Experimental Demonstration. *IEEE Trans. Rob.*, 34(5):1384–1391, May 2018. doi: 10.1109/TRO.2018.2830346.

[66] Rohan Thakker, Michael Paton, Marlin P. Strub, Michael Swan, Guglielmo Daddi, and Rob Royce. EELS: Towards Autonomous Mobility in Extreme Terrain with a Versatile Snake Robot with Resilience to Exteroception Failures. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 01–05. IEEE. doi: 10.1109/IROS55552.2023.10341448.

[67] Peter Tsou, Donald E. Brownlee, Christopher P. McKay, Ariel D. Anbar, Hajime Yano, Kathrin Altwegg, Luther W. Beegle, Richard Dissly, Nathan J. Strange, and Isik Kanik. LIFE: Life Investigation For Enceladus A Sample Return Mission Concept in Search for Evidence of Life. *Astrobiology*, 12(8):730–742, August 2012. ISSN 1557-8070. doi: 10.1089/ast.2011.0813.

[68] Brian Van Stratum, Kourosh Shoele, and Jonathan E. Clark. Pacific Lamprey Inspired Climbing. *arXiv*, December 2022. doi: 10.1088/1748-3190/acd671.

[69] T. S. Vaquero, G. Daddi, R. Thakker, M. Paton, A. Jasour, M. P. Strub, R. M. Swan, R. Royce, M. Gildner, P. Tosi, M. Veismann, P. Gavrilov, E. Marteau, J. Bowkett, D. Loret de Mola Lemus, Y. Nakka, B. Hockman, A. Orekhov, T. D. Hasseler, C. Leake, B. Nuernberger, P. Proença, W. Reid, W. Talbot, N. Georgiev, T. Pailevanian, A. Archanian, E. Ambrose, J. Jasper, R. Etheredge, C. Roman, D. Levine, K. Otsu, S. Yearicks, H. Melikyan, R. R. Rieber, K. Carpenter, J. Nash, A. Jain, L. Shiraishi, M. Robinson, M. Travers, H. Choset, J. Burdick, A. Gardner, M. Cable, M. Ingham, and M. Ono. Eels: Autonomous snake-like robot with task and motion planning capabilities for ice world exploration. *Science Robotics*, 9(88):eadh8332, 2024. doi: 10.1126/scirobotics.adh8332. URL https://www.science.org/doi/abs/10.1126/scirobotics.adh8332.

[70] Dylan Vogel, Robert Baines, Joseph Church, Julian Lotzer, Karl Werner, and Marco Hutter. Robust Ladder Climbing with a Quadrupedal Robot. *arXiv*, September 2024. doi: 10.48550/arXiv.2409.17731.

[71] Rich Volpe. JPL Robotics: Ice Worm, May 2025. URL https://www-robotics.jpl.nasa.gov/what-we-do/research-tasks/ice-worm.

[72] Xiaodong Wu and Shugen Ma. Cpg-based control of serpentine locomotion of a snake-like robot*, 2009.

[73] Zifan Xu, Amir Hossain Raj, Xuesu Xiao, and Peter Stone. Dexterous Legged Locomotion in Confined 3D Spaces with Reinforcement Learning. *arXiv*, March 2024. doi: 10.48550/arXiv.2403.03848.

[74] Chuanyu Yang, Can Pu, Guiyang Xin, Jie Zhang, and Zhibin Li. Learning complex motor skills for legged robot fall recovery, 2023.

[75] Yuyou Zhang, Yaru Niu, Xingyu Liu, and Ding Zhao. COMPOSER: Scalable and Robust Modular Policies for Snake Robots. *arXiv*, October 2023. doi: 10.48550/arXiv.2310.00871.

[76] Weikun Zhen, Chaohui Gong, and H. Choset. Modeling rolling gaits of a snake robot. *IEEE International Conference on Robotics and Automation*, 2015. URL https://www.semanticscholar.org/paper/Modeling-rolling-gaits-of-a-snake-robot-Zhen-Gong/3bef0714a6737c1eec0a5eca4442133223d1687a.

[77] B. Zhong, Tianyu Wang, Jennifer M. Rieser, Abdul Kaba, H. Choset, and D. Goldman. Frequency modulation of body waves to improve performance of limbless robots, 2020.

[78] Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher Atkeson, Soeren Schwertfeger, Chelsea Finn, and Hang Zhao. Robot Parkour Learning. *arXiv*, September 2023. doi: 10.48550/arXiv.2309.05665.