

Safe Active Navigation and Exploration for Planetary Environments Using Proprioceptive Measurements

Matthew Jiang^{*†}, Shipeng Liu^{*†}, Feifei Qian^{*†}

^{*}University of Southern California, Los Angeles, California 90089, USA

[†]Equal contribution

[†]Corresponding Author

Email: {jiangmy, shipengl, feifeiqi}@usc.edu

Abstract—Legged robots can sense terrain through force interactions during locomotion, offering more reliable traversability estimates than remote sensing and serving as scouts for guiding wheeled rovers in challenging environments. However, even legged scouts face challenges when traversing highly deformable or unstable terrain. We present Safe Active Exploration for Granular Terrain (SAEGT), a navigation framework that enables legged robots to safely explore unknown granular environments using proprioceptive sensing, particularly where visual input fails to capture terrain deformability. SAEGT estimates the safe region and frontier region online from leg-terrain interactions using Gaussian Process regression for traversability assessment, with a reactive controller for real-time safe exploration and navigation. SAEGT demonstrated its ability to safely explore and navigate toward a specified goal using only proprioceptively estimated traversability in simulation.

I. INTRODUCTION

Planetary exploration often relies on remote or visual sensing to plan safe paths, as in the Mars 2020 rover’s use of HazCams and NavCams [23]. However, failures like the Spirit rover’s entrapment [8] reveal that vision systems cannot capture subsurface properties of deformable terrains. In contrast, legged robots use proprioceptive feedback [22, 11, 6, 7] during each step to directly assess terrain properties, offering sensing capabilities that are unaffected by the limitations of vision-based systems. This makes them well-suited for scouting roles in planetary exploration. However, despite their proven mobility in unstructured [14, 18] and soft granular terrains [28, 21, 20], legged robots still face locomotion challenges on the soft and cohesive surfaces [5] of planetary environments. This highlights the need for a safe and active navigation framework to support their role as scouts with only point-sensor measurements in planetary exploration.

To enable legged robots to navigate and actively scout based on proprioceptive measurements, this work introduces a safe navigation framework composed of two main components. We first employ a Gaussian Process using point-wise proprioceptive measurements to create a safe region map, where we know with a tunable level of confidence that the robot can walk safely, and a frontier map indicating where the robot can expand the safe region by exploring further. Second,

we combine these regions with a reactive, geometry-based local controller, SAEGT, enabling legged robots to actively map terrain mechanical properties, update safe and frontier regions in real time, and safely explore while supporting other navigation objectives.

II. RELATED WORK

Robot exploration and navigation frameworks generally can be divided into two categories: geometric-based [19] and learning-based approaches [35]. In this work, we are interested in the first one. These approaches often require the robot to construct a map from exteroceptive sensors (cameras, LiDAR) [31] and analyze the traversability. Gaussian Process-based methods have emerged as powerful tools to estimate a continuous map [19], and traversability is then often analyzed based on slope, roughness, and step height [36, 34]. However, on deformable terrains, the robot must rely on proprioceptive measurements to assess terrain properties, which can only be obtained upon contact. This paper presents a Gaussian Process-based method to generate a safe region and a frontier region, enabling the robot to actively map and explore using proprioceptive feedback.

After generating the traversability map, a planner is needed to compute the robot’s path. If the goal is to efficiently map the environment, the planner seeks to maximize information entropy [12, 1, 15]; if there is a global goal, it focuses on ensuring safe traversal to the goal location. Two main approaches are often used: (1) plan a trajectory using methods such as RRT [17], RRT* [16], Dijkstra [10], or A* [13], followed by a controller to track the trajectory; or (2) directly apply potential field-based methods to convert the occupancy map into a vector field to control the robot [2, 3, 32, 33]. Our method integrates the identified safe region and frontier region with potential-field-based methods, directly generating velocity control commands for the robot to either navigate to a global goal or explore the environment.

III. PROPRIOCEPTIVE TRAVERSABILITY MAPPING FOR GRANULAR TERRAINS

In this section, we present a granular-terrain traversability mapping framework composed of two tightly integrated components: (i) Gaussian Process-based traversability mapping from per-step proprioceptive sensing, which incrementally constructs a global traversability map \mathcal{T}_t ; and (ii) region detection and expansion using a Bayesian safety optimization algorithm, which iteratively expands the safe region S_t by leveraging both \mathcal{T}_t and the associated uncertainty estimates. This process simultaneously identifies the associated frontier region G_t for targeted exploration. The frontier points, $\mathbf{g}_t \in G_t$, are potential targeted exploration goals whose traversal can enlarge the robot's safe region, S_t .

A. Proprioceptive Traversability Mapping

The robot operates as a point at position $\mathbf{x}_t \in \mathbb{R}^2$ within workspace \mathcal{W} , discretized into grid D with resolution δ . The proprioceptive system infers terrain traversability y_t through leg-terrain interactions [22, 24], combined with terrain dynamics models [20] to predict entrapment risk [25].

The traversability function $f(\mathbf{x})$ represents a scalar score derived from the terrain dynamics model's assessment of mechanical properties. Higher values of $f(\mathbf{x})$ correspond to terrain with lower predicted risk of entrapment or mobility failure. We assume traversability measurements are available as noisy observations at each location \mathbf{x}_t :

$$y_t = f(\mathbf{x}_t) + \varepsilon_t \quad (1)$$

where $f(\mathbf{x}_t) \in \mathbb{R}$ represents the true terrain traversability at location \mathbf{x}_t , and $\varepsilon_t \sim \mathcal{N}(0, \sigma_{\text{noise}}^2)$ accounts for measurement noise and processing uncertainty in the proprioceptive sensing and terrain dynamics modeling pipeline.

The terrain traversability map \mathcal{T}_t is maintained as a Gaussian Process posterior over measurement history $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^t$, using the RBF kernel:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right) \quad (2)$$

The GP provides predictive mean $\mu_t(\mathbf{x})$ and variance $\sigma_t^2(\mathbf{x})$ at any location \mathbf{x} of the terrain traversability, yielding both terrain traversability predictions and uncertainty estimates that guide safe exploration.

B. Region Detection

1) *Traversability Confidence Interval Construction:* The Gaussian Process traversability mapping provides predictive mean $\mu_t(\mathbf{x})$ and uncertainty $\sigma_t(\mathbf{x})$. To determine safe and frontier regions, we construct confidence bounds on the predicted traversability values at each position \mathbf{x} .

A confidence parameter β controls the exploration-safety trade-off, with lower values corresponding to more conservative safety assessments (e.g., $\beta = 1$ corresponds to approximately 68% confidence [26]):

$$c_t(\mathbf{x}) = \left[\mu_{t-1}(\mathbf{x}) \pm \sqrt{\beta} \sigma_{t-1}(\mathbf{x}) \right] \quad (3)$$

The confidence bounds are iteratively intersected with prior confidence sets to ensure consistency:

$$C_t(\mathbf{x}) = C_{t-1}(\mathbf{x}) \cap c_t(\mathbf{x}) \quad (4)$$

From $C_t(\mathbf{x})$, the lower and upper bounds of traversability are defined as:

$$\ell_t(\mathbf{x}) = \min C_t(\mathbf{x}), \quad u_t(\mathbf{x}) = \max C_t(\mathbf{x}) \quad (5)$$

2) *Safe Region Detection:* Following previous work [30], the lower confidence bounds are used to iteratively expand regions that are safe with high probability for the robot to traverse.

The safe region expansion relies on the assumption that the true traversability $f(\mathbf{x})$ is Lipschitz continuous with constant L , meaning $|f(\mathbf{x}) - f(\mathbf{x}')| \leq L\|\mathbf{x} - \mathbf{x}'\|$ for all points \mathbf{x}, \mathbf{x}' in the workspace [29]. For granular terrains, this reflects the physical notion that soil mechanical properties change gradually over space due to similar formation processes, loading history, and environmental conditions affecting neighboring regions. The safe region S_t is constructed by expanding from previously safe regions:

$$S_t = \bigcup_{\mathbf{x} \in S_{t-1}} \{\mathbf{x}' \in D \mid \ell_t(\mathbf{x}) - L\|\mathbf{x} - \mathbf{x}'\| \geq h\} \quad (6)$$

where $\ell_t(\mathbf{x})$ is the lower confidence bound of traversability, h is the safety threshold, and L is the Lipschitz constant of the traversability function $f(\mathbf{x})$.

The robot begins with an initial small safe region S_0 based on its starting location and initial measurements. This region is then expanded iteratively as new measurements reduce uncertainty and allow the lower confidence bounds to exceed the safety threshold in previously unknown regions.

3) *Frontier Region Detection:* To identify regions that could potentially expand the safe region, we utilize the expansion objective/acquisition function [30]:

$$g_t(\mathbf{x}) = |\{\mathbf{x}' \in D \setminus S_t : u_t(\mathbf{x}) - L\|\mathbf{x} - \mathbf{x}'\| \geq h\}| \quad (7)$$

which evaluates how large the existing safe region could expand if point \mathbf{x} is sampled. Using the upper confidence bound $u_t(\mathbf{x})$, the potential safe region expansion is evaluated under the most favorable traversability scenario. To consider the expansion objective, let the frontier region G_t consist of safe points with positive expansion potential:

$$G_t = \{\mathbf{x} \in S_t \mid g_t(\mathbf{x}) > 0\} \quad (8)$$

These frontier points represent candidate sampling locations that could enlarge the verified safe region when visited.

IV. REACTIVE NAVIGATION FOR SAFE EXPLORATION

This section outlines the navigation strategy of SAEGT, which first selects a frontier subgoal that balances exploration and progress toward the goal. It then generates a potential field based on the selected subgoal and traversability prediction to guide the robot in safely exploring the environment.

A. Frontier Subgoal Selection

The robot selects the next sampling point by balancing two objectives:

- Expand the safe regions
- Make progress toward the mission goal

After identifying the safe region S_t and the frontier region $G_t \subseteq S_t$, the robot employs a hybrid strategy to guide exploration in a goal-directed yet uncertainty-aware manner.

1) *Subgoal Selection Procedure*: The subgoal selection procedure is as follows:

- 1) For each point $\mathbf{x} \in G_t$, compute its distance to the goal $\|\mathbf{x} - \mathbf{x}_g\|$.
- 2) Sort all points in G_t in ascending order of distance to the goal.
- 3) Consider only the top n nearest frontier points, where n is a configurable parameter.
- 4) Among these n frontier points, select the point with the largest confidence interval width $w_t(\mathbf{x}) := u_t(\mathbf{x}) - \ell_t(\mathbf{x})$.

This approach prioritizes frontier points that offer more information, while guiding the robot toward its goal using a distance heuristic. It mitigates the risk of wasting samples on irrelevant directions and prevents purely greedy selection based on proximity alone, which might lead into regions of high uncertainty without expanding the safe region.

Algorithm 1 Safe Active Exploration for Granular Terrain (SAEGT)

Require: Robot workspace D , GP prior (μ_0, k, σ_0) , Lipschitz constant L , initial safe region S_0 , safety threshold h , goal location \mathbf{x}_g , number of frontier points to consider n

- 1: $C_0(\mathbf{x}) \leftarrow [h, \infty), \forall \mathbf{x} \in S_0$
 - 2: $C_0(\mathbf{x}) \leftarrow \mathbb{R}, \forall \mathbf{x} \in D \setminus S_0$
 - 3: $c_0(\mathbf{x}) \leftarrow \mathbb{R}, \forall \mathbf{x} \in D$
 - 4: **for** each time step $t = 1, 2, \dots$ **do**
 - 5: $C_t(\mathbf{x}) \leftarrow C_{t-1}(\mathbf{x}) \cap c_{t-1}(\mathbf{x})$
 - 6: $S_t \leftarrow \bigcup_{\mathbf{x} \in S_{t-1}} \{\mathbf{x}' \in D \mid \ell_t(\mathbf{x}) - L\|\mathbf{x} - \mathbf{x}'\| \geq h\}$
 - 7: $G_t \leftarrow \{\mathbf{x} \in S_t \mid g_t(\mathbf{x}) > 0\}$
 - 8: Sort G_t by $\|\mathbf{x} - \mathbf{x}_g\|$ in ascending order
 - 9: Let $G_t^{(n)} \subseteq G_t$ be the top n nearest frontier points
 - 10: $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in G_t^{(n)}} w_t(\mathbf{x})$
 - 11: Navigate to \mathbf{x}_t using reactive controller
 - 12: Observe $y_t = f(\mathbf{x}_t) + \varepsilon_t$
 - 13: Update GP with (\mathbf{x}_t, y_t)
 - 14: Compute $c_t(\mathbf{x})$ for all $\mathbf{x} \in D$
 - 15: **end for**
-

This strategy combines uncertainty-driven sampling with spatial prioritization to ensure that each move both refines the terrain model and progresses toward the mission objective.

B. Reactive Navigation

While the planner identifies target locations, executing safe motion between sampling points requires a robust local navigation strategy. Classical global planners require full replanning after environmental updates, making them unsuitable for

dynamic terrain scenarios. Instead, we implement a reactive control framework adapted from Vasilopoulos et al. [33].

1) *Geometric Obstacle Construction*: To enable reactive navigation, the robot constructs geometric representations of safe and hazardous terrain from its evolving traversability map. This process transforms pointwise traversability predictions into polygonal approximations suitable for geometry-aware motion planning.




Fig. 1: Geometric processing pipeline for converting safe terrain points to navigable free space representations. (a) Visualization of a concave hull generated from point clusters. Image adapted from cubao [9]. Licensed under the MIT License. (b) Bounding box creation for workspace definition.

The conversion proceeds as follows:

- 1) **Cluster Formation**: The safe set S_t is partitioned into disjoint spatial clusters using KD-tree [4] nearest-neighbor queries. Each cluster represents a contiguous region of traversable terrain.
- 2) **Concave Hull Construction**: For each cluster, a concave hull is generated using the algorithm from [27], with parameters controlling shape resolution and computational complexity.
- 3) **Workspace Definition**: A bounding rectangle \mathcal{W}_e encloses all hull regions:

$$\mathcal{W}_e := \{\mathbf{x} \in \mathbb{R}^2 \mid \underline{H}_t \leq \mathbf{x} \leq \overline{H}_t\} \quad (9)$$

- 4) **Obstacle Extraction**: Regions not covered by safe hulls become pseudo-physical obstacles:

$$\mathcal{O}_t := \mathcal{W}_e \setminus \bigcup_i H_{ti} \quad (10)$$

2) *Diffeomorphic Reactive Control*: We adopt the reactive navigation framework from Vasilopoulos et al. [33], which constructs a diffeomorphic transformation between the physical workspace and a simplified model space. This enables obstacle avoidance using simple potential fields in the transformed domain. The method guarantees that smooth trajectories in model space correspond to safe, obstacle-free paths in physical space, provided the obstacle set is polygonal. This allows real-time control generation without full replanning as the known obstacle set evolves.

C. System Integration

The final system integrates global terrain expansion and local reactive control into a unified exploration loop: proprioceptive measurements update the GP-based terrain model, the

planner selects new targets from safe regions, and the reactive controller executes real-time movement.

This dual-loop architecture allows the robot to safely explore unknown, deformable environments while efficiently progressing toward its scientific objectives.

V. RESULTS

This section presents simulation-based validation of SAEGT across terrain scenarios that demonstrate its core capabilities in two scenarios: exploration with a certain goal/priority, and exploratory mapping without preassigned goals.

A. Simulation Setup

The robot is initialized within a verified safe region with randomly sampled proprioceptive measurements to bootstrap the Gaussian Process model. Test environments are represented as 2D grids where each cell contains terrain traversability values, with the traversability function $f(\mathbf{x})$ encoding relative terrain strength and safety threshold h defining traversability boundaries (e.g., $h = 1000$ for demonstration). $\beta = 3$ was used, meaning a 99.7% confidence interval was used.




Fig. 2: Ground truth traversability values for experiments.

B. Experimental Results

a) *Navigation Around Low-Traversability Zones:* The robot successfully navigates around low-traversability regions by iteratively expanding the safe region until reaching goals beyond obstacles. Two scenarios were completed in 210 and 461 iterations, respectively (Figure 4). The algorithm initially expands rapidly, then explores cautiously around low-traversability peripheries, and finally accelerates once past dangerous areas.

b) *Goal-Free Exploration:* Without predefined destinations, the robot selects points maximizing uncertainty reduction, rapidly expanding the known safe area (Figure 5). This demonstrates SAEGT’s capability for general environmental mapping, useful for initial site characterization before goal-directed missions.

The simulation experiments presented validate the intended capabilities of SAEGT under a variety of terrain conditions and traversability scenarios. The robot successfully performs

¹It is important to note that the gray and white areas are shown as ground truth to validate the algorithm, but are not information known to the robot during navigation, while the contents surrounded by the enclosing workspace (outlined by the black obstacles) are known and are created by SAEGT.




Fig. 3: A single frame showing the navigation regions and setup.¹ White = safe terrain ($f(\mathbf{x}) \geq h$); Gray = low-traversability terrain ($f(\mathbf{x}) < h$); Black = obstacles (\mathcal{O}_t); Green = local freespace; Light Blue = safe region (S_t); Red = frontier region (G_t); Purple = intermediate goal; Dark Blue = final goal.




Fig. 4: Navigation around low-traversability zone ($h = 1000$)

uncertainty-aware navigation around hazardous terrain, reuses previously acquired safe knowledge to complete tasks more efficiently, and performs autonomous environmental mapping using only local proprioceptive measurements.

VI. CONCLUSION

This paper presents SAEGT, a navigation framework for proprioceptive terrain-aware exploration in granular environments where visual sensing is unreliable. Designed to address

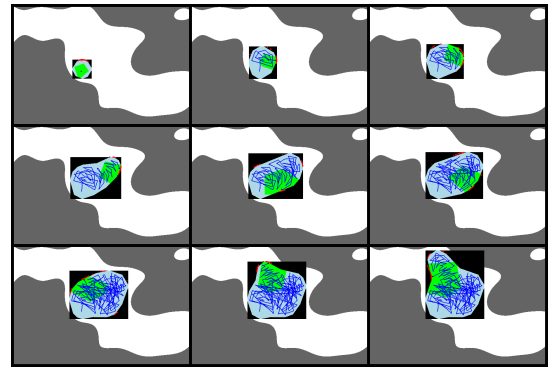


Fig. 5: Overall exploration progression over 9 intervals.

uncertainty, deformability, and incomplete knowledge in planetary terrains, SAEGT combines GP-based terrain modeling, confidence-guided safe region expansion, and reactive control to enable legged robots to explore safely using only local physical interaction.

REFERENCES

- [1] Mahmoud Ali, Hassan Jardali, Nicholas Roy, and Lantao Liu. Autonomous navigation, mapping and exploration with gaussian processes. *Robotics: Science and Systems (RSS)*, 2023.
- [2] Omur Arslan and Daniel Koditschek. Exact robot navigation using power diagrams. 05 2016. doi: 10.1109/ICRA.2016.7487090.
- [3] Omur Arslan and Daniel E Koditschek. Sensor-based reactive navigation in unknown convex sphere worlds. *The International Journal of Robotics Research*, 38(2-3): 196–223, 2019. doi: 10.1177/0278364918796267. URL <https://doi.org/10.1177/0278364918796267>.
- [4] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9): 509–517, September 1975. ISSN 0001-0782. doi: 10.1145/361002.361007. URL <https://doi.org/10.1145/361002.361007>.
- [5] Carol S. Breed, Maurice J. Grolier, and John F. McCauley. Morphology and distribution of common ‘sand’ dunes on mars: Comparison with the earth. *Journal of Geophysical Research: Solid Earth*, 84(B14):8183–8204, 1979. doi: <https://doi.org/10.1029/JB084iB14p08183>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/JB084iB14p08183>.
- [6] J. C. Bush, Y. Zhang, S. Liu, E. Fulcher, C. Franklin, D. J. Jerolmack, R. C. Ewing, K. R. Fisher, and F. Qian. Robotic Legs as Novel Planetary Instrumentation to Explore the Mechanical Properties of Regolith. In *54th Lunar and Planetary Science Conference*, volume 2806 of *LPI Contributions*, page 1733, March 2023.
- [7] J. C. Bush, Y. Zhang, J. D. Caporale, R. C. Ewing, K. R. Fisher, E. Fulcher, D. J. Jerolmack, N. A. Jones, D. E. Koditschek, S. Liu, B. McKeeby, M. Nachan, F. Rivera-Hernandez, J. Ruck, T. Shipley, S. Thompson, C. Wilson, and F. Qian. Relating Geotechnical Properties of Crusty Regolith to Morphology and Mineralogy Using a Robotic Leg Rheometer. In *55th Lunar and Planetary Science Conference*, volume 3040 of *LPI Contributions*, page 1798, March 2024.
- [8] John L. Callas. Mars exploration rover spirit end of mission report. Technical Report JPL-Publ-16-2, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, United States, December 2015. URL <https://ntrs.nasa.gov/api/citations/20160001767/downloads/20160001767.pdf>. Document ID: 20160001767.
- [9] cubao. Concave hull. https://github.com/cubao/concave_hull, 2024. Accessed April 30, 2025.
- [10] Edsger W Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: his life, work, and legacy*, pages 287–290. 2022.
- [11] E. Fulcher, J. D. Caporale, Y. Zhang, J. C. Bush, R. C. Ewing, K. R. Fisher, D. J. Jerolmack, S. Liu, M. Nachan, F. Rivera-Hernandez, J. Ruck, T. Shipley, C. Wilson, D. E. Koditschek, and F. Qian. Making Every Step an Experiment: Proprioceptive Sensing During Locomotion for Enhanced Mobility and Data Collection in Planetary-Analogue Explorations. In *55th Lunar and Planetary Science Conference*, volume 3040 of *LPI Contributions*, page 1806, March 2024.
- [12] Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake. Gaussian processes autonomous mapping and exploration for range-sensing mobile robots. *Autonomous Robots*, 42:273–290, 2018.
- [13] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [14] Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26):eaau5872, 2019.
- [15] Dohyun Jang, Jaehyun Yoo, Clark Youngdong Son, Dabin Kim, and H Jin Kim. Multi-robot active sensing and environmental model learning with distributed gaussian process. *IEEE Robotics and Automation Letters*, 5(4):5905–5912, 2020.
- [16] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- [17] Steven LaValle. Rapidly-exploring random trees: A new tool for path planning. *Research Report 9811*, 1998.
- [18] Joonho Lee, Marko Bjelonic, Alexander Reske, Lorenz Wellhausen, Takahiro Miki, and Marco Hutter. Learning robust autonomous navigation and locomotion for wheeled-legged robots. *Science Robotics*, 9(89): eadi9641, 2024.
- [19] Abe Leininger, Mahmoud Ali, Hassan Jardali, and Lantao Liu. Gaussian process-based traversability analysis for terrain mapless navigation, 2024. URL <https://arxiv.org/abs/2403.19010>.
- [20] Chen Li, Paul B Umbanhowar, Haldun Komsuoglu, Daniel E Koditschek, and Daniel I Goldman. Sensitive dependence of the motion of a legged robot on granular media. *Proceedings of the National Academy of Sciences*, 106(9):3029–3034, 2009.
- [21] Shipeng Liu, Boyuan Huang, and Feifei Qian. Adaptation of flipper-mud interactions enables effective terrestrial locomotion on muddy substrates. *IEEE Robotics and Automation Letters*, 2023.
- [22] Shipeng Liu, Jiaze Tang, Siyuan Meng, and Feifei Qian. Adaptive locomotion on mud through proprioceptive sensing of substrate properties. *arXiv preprint arXiv:2504.19607*, 2025.
- [23] Justin N. Maki, David Gruel, Constance McKinney, Michael A. Ravine, Michael Morales, Douglas Lee, Richard Willson, David Copley-Woods, Michael Valentin, Thomas Goodsall, James McGuire, R. Greg Sellar, Jake A. Schaffner, Ken D. Supulver, Jennifer Tu, Andres Wong, Paul Brugarolas, Samad A. Hayati,

- Michael Kasper, Subash Reddy, and Jeff Umland. The mars 2020 engineering cameras and microphone on the perseverance rover: A next-generation imaging system for mars exploration. *Space Science Reviews*, 216(8): 1–48, 2020. doi: 10.1007/s11214-020-00765-9.
- [24] E. Marteau, M. Golombek, P. Delage, C. Vrettos, K. Hurst, A. Gomez, N. Williams, P. Bailey, and P. Mishra. Initial results from the insight lander robotic arm soil mechanics experiments on mars. In *54th Lunar and Planetary Science Conference*, number 1597, The Woodlands, TX, USA, March 2023.
- [25] Yashodhan H Nevatia, Jeremi Gancet, Francois Bulens, Thomas Voegele, Roland U Sonsalla, Chakravarthini M Saaj, William A Lewinger, Marcus Matthews, Brian Yeomans, Yang Gao, et al. Improved traversal for planetary rovers through forward acquisition of terrain trafficability. In *IEEE International Conference on Robotics and Automation Planetary Rovers Workshop*, pages 6–10, 2013.
- [26] Jerzy Neyman. Outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 236(767): 333–380, 1937.
- [27] J.-S Park and S.-J Oh. A new concave hull algorithm and concaveness measure for n-dimensional datasets. *Journal of Information Science and Engineering*, 29:379–392, 03 2013.
- [28] Feifei Qian, Tingnan Zhang, Chen Li, Pierangelo Masarati, Aaron M Hoover, Paul Birkmeyer, Andrew Pullin, Ronald S Fearing, and Daniel I Goldman. Walking and running on yielding and fluidizing ground. *Robotics: Science and Systems*, page 345, 2013.
- [29] H.L. Royden and P. Fitzpatrick. *Real Analysis*. Prentice Hall, 2010. ISBN 9780131437470. URL <https://books.google.com/books?id=0Y5fAAAACAAJ>.
- [30] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 997–1005, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sui15.html>.
- [31] Monisha Mushtary Uttsha, Cedric Le Gentil, Lan Wu, and Teresa Vidal-Calleja. Gaussian process distance fields obstacle and ground constraints for safe navigation, 2024. URL <https://arxiv.org/abs/2410.17831>.
- [32] V. Vasilopoulos, G. Pavlakos, S. L. Bowman, J. D. Caporale, K. Daniilidis, G. J. Pappas, and D. E. Koditschek. Reactive Semantic Planning in Unexplored Semantic Environments Using Deep Perceptual Feedback. *IEEE Robotics and Automation Letters (RA-L)*, 5(3):4455–4462, 2020. doi: 10.1109/LRA.2020.3001496. URL <https://doi.org/10.1109/ICRA.2016.7487090>.
- [33] Vasileios Vasilopoulos, Georgios Pavlakos, Karl Schmeckpeper, Kostas Daniilidis, and Daniel E. Koditschek. Reactive navigation in partially familiar planar environments using semantic perceptual feedback, 2021. URL <https://arxiv.org/abs/2002.08946>.
- [34] Jingping Wang, Long Xu, Haoran Fu, Zehui Meng, Chao Xu, Yanjun Cao, Ximin Lyu, and Fei Gao. Towards efficient trajectory generation for ground robots beyond 2d environment. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7858–7864. IEEE, 2023.
- [35] Kasun Weerakoon, Adarsh Jagan Sathyamoorthy, Utsav Patel, and Dinesh Manocha. Terp: Reliable planning in uneven outdoor environments using deep reinforcement learning. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 9447–9453. IEEE, 2022.
- [36] Martin Wermelinger, Péter Fankhauser, Remo Diethelm, Philipp Krüsi, Roland Siegwart, and Marco Hutter. Navigation planning for legged robots in challenging terrain. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1184–1189. IEEE, 2016.

APPENDIX

This appendix provides detailed explanations for the key steps in the SAEGT algorithm presented in Algorithm 1.

A. Algorithm Step Annotations

Line 5: Update traversability confidence - This step intersects the current confidence interval with the previous one to maintain consistency across time steps. The intersection ensures that confidence bounds become progressively tighter as more measurements are collected.

Line 6: Expand safe region - The safe region expansion uses the Lipschitz continuity assumption to propagate safety guarantees from previously verified safe points to nearby locations. Points are added to the safe region if their lower confidence bound, adjusted for distance-based uncertainty, exceeds the safety threshold.

Line 7: Identify frontier region - Frontier points are identified as safe locations that have the potential to expand the safe region if visited. These points represent the boundary between known safe terrain and unexplored areas.

Line 8: Prioritize by goal distance - Frontier points are sorted by their Euclidean distance to the goal location to bias exploration toward the mission objective while maintaining safety.

Line 10: Select most informative - Among the nearest frontier candidates, the algorithm selects the point with the largest confidence interval width, maximizing information gain from the next measurement.

Line 12: Collect proprioceptive measurement - The robot obtains a noisy observation of terrain traversability through physical interaction with the terrain at the selected location.