RISK AWARE REACTIVE NAVIGATION FOR GRANULAR TERRAIN EXPLORATION

by

Matthew Jiang

Submitted to the UNIVERSITY OF SOUTHERN CALIFORNIA In Partial Fulfillment of Graduation Requirements for HONORS DEGREE

Viterbi School of Engineering

May 2025

Research Advisor: Feifei Qian
Signature: ______________________

Program Director: Sandeep Gupta
Signature: ______________________

UNIVERSITY OF SOUTHERN CALIFORNIA

DEPARTMENTAL APPROVAL

of an honors undergraduate thesis submitted by

Matthew Jiang

This thesis has been reviewed by the research advisor, engineering honors research advisor, and honors program director, and it has been found to be satisfactory.

Prof. Feifei Qian, Research Advisor

Prof. Sandeep Gupta, Program Director

Acknowledgments

I would like to express my deepest gratitude to my thesis advisor, Professor Feifei Qian, and my Ph.D. student mentor, Shipeng Liu, for their exceptional mentorship, intellectual guidance, and unwavering support throughout the course of this research.

This work would not have been possible without the support of the RoboLAND Lab at the University of Southern California, whose resources and collaborative environment were instrumental to the development and validation of this work.

I also wish to acknowledge the LASSIE project team for their vision in shaping the broader context of this research, and for providing an ambitious platform that inspired the pursuit of safe, proprioceptive navigation for planetary environments.

Lastly, I extend heartfelt appreciation to my family and friends for their camaraderie, encouragement, and grounding presence. Your support helped sustain the energy and purpose that brought this thesis to completion.

Table of Contents

List of Figures

Abstract

This thesis presents Safe Reactive Planning for Granular Terrain (SRPGT), a navigation framework for autonomous exploration in granular environments where proprioceptive sensing is the main mode of terrain mapping. The method is motivated by challenges faced in planetary exploration, particularly in scenarios where deformable terrain properties cannot be reliably inferred from visual input. Developed within the context of NASA's LASSIE (Legged Autonomous Surface Science In Analogue Environments) project, SRPGT enables a legged robot to traverse unfamiliar regolith terrain using only proprioceptive feedback from its limbs.

The core contribution is a dual-layer architecture combining global terrain expansion with real-time reactive control. The global layer models terrain risk using a Gaussian Process updated through in-place mechanical tests at each step. It incrementally expands a certified safe set of traversable regions using a confidence-aware exploration strategy inspired by Safe Bayesian Optimization. The local control layer uses a diffeomorphism-based reactive controller adapted from Voronoi and power diagram methods, allowing the robot to navigate through complex, concave environments without requiring global replanning.

Simulations demonstrate that SRPGT successfully balances exploration and exploitation, avoids unsafe regions, and generalizes across navigation and pure exploration tasks. The algorithm achieves real-time performance and robustness without vision, supporting future missions in low-light or dust-obscured extraterrestrial environments. This work establishes a foundation for proprioception-driven planning, expanding autonomous mobility capabilities in safety-critical planetary operations.

Chapter 1

Introduction

The future of space exploration increasingly relies on our ability to establish a sustainable human presence beyond Earth. As missions become more ambitious, with visions of long-term habitation on the Moon, Mars, and other celestial bodies, the importance of robust robotic systems capable of autonomous exploration becomes ever more critical. Before building the infrastructure that future astronauts will depend on, we must first navigate and understand unfamiliar and often hazardous terrains, frequently without direct human supervision.

For decades, wheeled rovers have served as humanity's primary tools for planetary exploration. Their successes, such as Spirit and Opportunity's extensive traverses across the Martian surface and Curiosity and Perseverance's ongoing missions, have significantly expanded our knowledge of other worlds. However, as we move toward the goal of sustained human presence, the limitations of traditional rovers have become increasingly apparent. These systems depend heavily on human oversight and primarily on camera-based sensors, interpreting their environment largely through visual data. While vision provides rich spatial information, it offers only an indirect understanding of the terrain.

This reliance on vision presents significant risks when operating in granular, deformable environments. Materials such as regolith, volcanic ash, and fine particulate dust, which are common on planetary surfaces, can be deceptive. Two patches of ground may appear identical but respond very differently under load. One may be stable and supportive, while the other collapses under minimal pressure. Cameras, regardless of resolution or sophistication, are inherently limited in their ability to infer critical mechanical properties such as cohesion, bearing strength, or subsurface instability.

A poignant reminder of these challenges comes from the story of the Spirit rover, which in 2009 became irretrievably stuck in Martian sand despite careful navigation. The episode captured public attention and was memorialized in popular culture, such as in XKCD's comic tribute [1]:

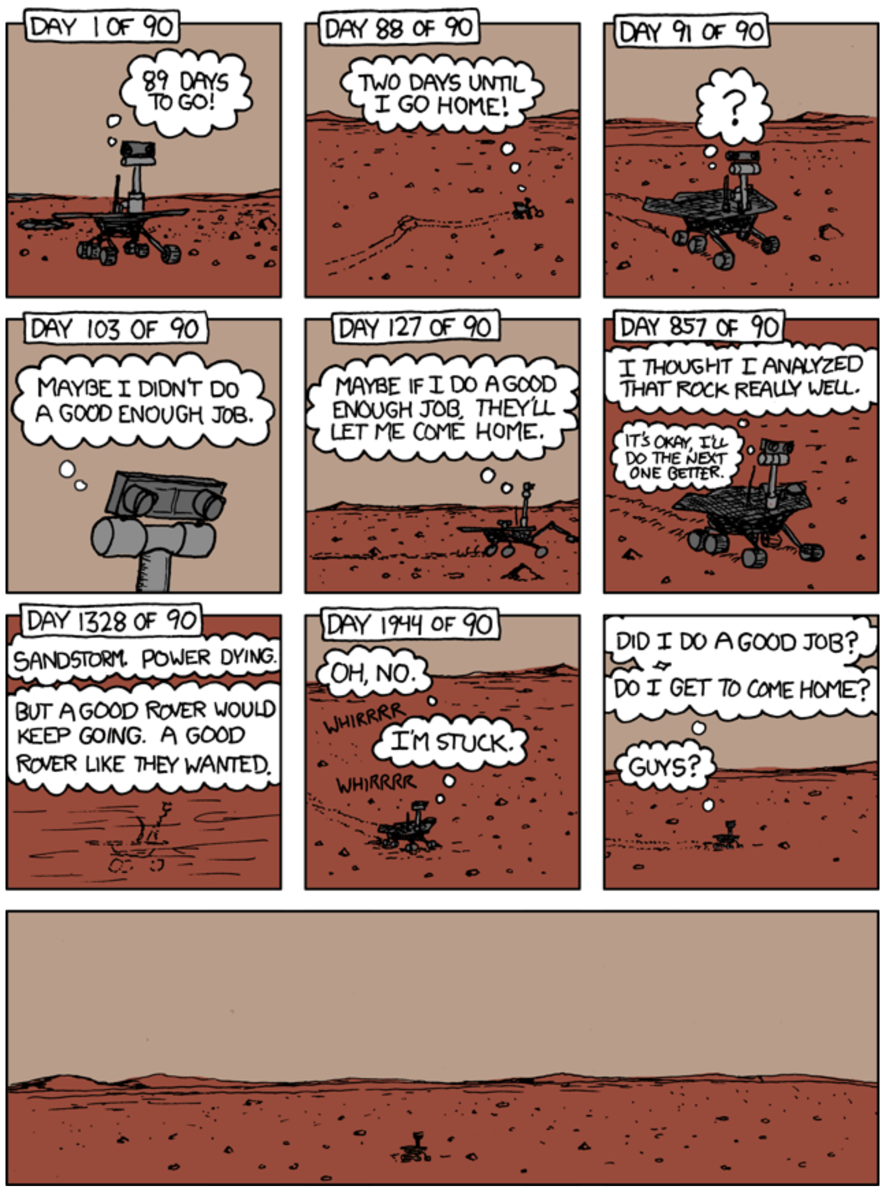

Figure 1.1: XKCD comic #695: A tribute to Spirit, highlighting the emotional weight and technical challenge of navigating granular terrain [1].

Though humorous in style, the comic conveys a sense of quiet tragedy. It conveys the undeniable truth: our robotic explorers are vulnerable to many hazards when faced with environments they cannot fully perceive or understand.

Moreover, granular terrains are dynamic and can change significantly across distances. Small disturbances, such as the pressure from a robot's wheel or foot, can cause localized deformation, slippage, or collapse. These changes often occur without any visual indication. Dust accumulation and variable lighting conditions, including low sun angles and deep shadows, further complicate visual perception, making reliance on vision alone increasingly dangerous.

Recognizing these challenges, a new paradigm for planetary exploration is necessary, one that moves beyond passive observation to active physical interrogation of the environment. NASA's Legged Autonomous Surface Science In Analogue Environments (LASSIE) project [2] embodies this approach. Instead of simply observing, legged robots interact with the terrain, using their limbs not only for locomotion but also as scientific instruments capable of performing mechanical tests.

Legged robots offer a unique advantage for such interactions. Every step becomes an opportunity to measure the mechanical response of the terrain. By pressing into the surface, applying controlled forces, and recording the resulting displacements, these robots can infer properties such as stiffness, cohesion, and internal friction. Unlike visual sensing, which depends on interpreting surface features, proprioceptive sensing provides direct measurements of terrain characteristics critical to safe navigation and environmental understanding.

The LASSIE project's quadrupedal platforms, such as the Ghost Robotics Spirit shown in Figure 1.2, are equipped with actuators capable of precise force control. These robots systematically perform mechanical measurements during locomotion, continuously gathering data about the terrain they traverse.

As shown in Figure 1.3, each step provides new information about terrain stiffness and stability. These measurements not only ensure the safety of the robot's immediate move-

Figure 1.2: Ghost Robotics Spirit, the quadrupedal robot used in LASSIE experiments on Mount Hood.

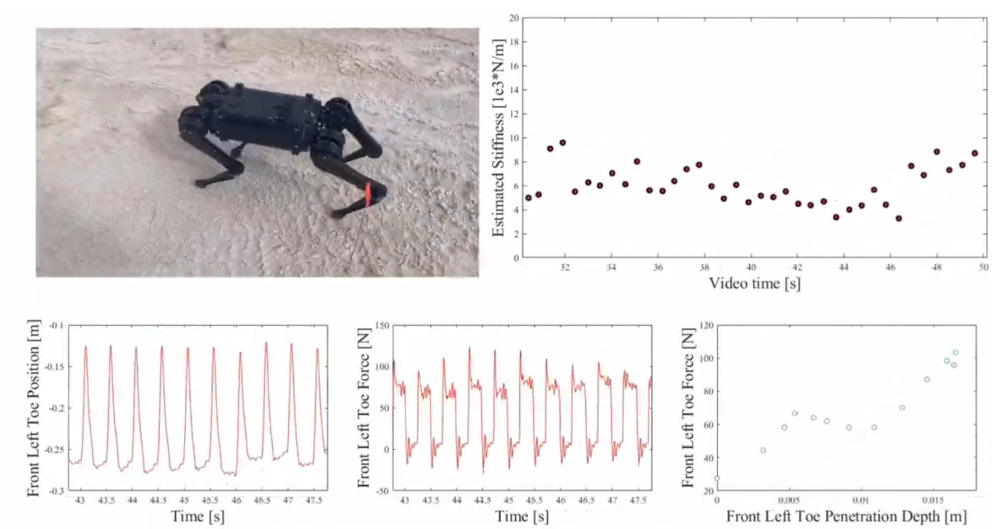

Figure 1.3: A LASSIE quadrupedal robot recording terrain stiffness during each step.

ments but also build a continuously updated map of the environment. Over time, the robot constructs a detailed understanding of the traversability and mechanical behavior of the landscape.

Several challenges remain in leveraging proprioceptive sensing for autonomous navigation. Most existing navigation algorithms are designed with the assumption that the robot has access to either global terrain maps or rich local sensing from sources such as LIDAR or cameras. These sensors provide spatial context beyond the robot's current location, allowing planners to construct paths based on look-ahead information. In contrast, proprioceptive sensing is inherently localized. The robot receives information only at the exact point of contact with the terrain, with no visibility into the safety of neighboring or distant regions.

This limitation introduces several technical challenges:

1. **Sparse and Delayed Observability:** Because the robot can only sense terrain after physically interacting with it, the environment must be discovered incrementally and cautiously.

2. **Uncertainty in Terrain Modeling:** Local measurements are subject to noise and variability, making it difficult to construct a reliable global map of terrain safety.

3. **Balancing Exploration and Exploitation:** The robot must decide whether to explore unknown regions to reduce uncertainty or exploit known safe areas to reach the goal, while maintaining a margin of safety.

4. **Real-Time Decision Making:** Planning must occur under strict time constraints and be continuously updated in response to new measurements.

5. **Safety Guarantees with Minimal Lookahead:** Without visibility into terrain ahead, the robot must ensure safety using only current and past measurements.

This thesis builds upon the LASSIE exploration loop by introducing a navigation strategy called *Safe Reactive Planning for Granular Terrain* (SRPGT) that places proprioceptive sensing at the core of the robot's decision-making process. SRPGT enables the robot to use each proprioceptive measurement to inform both its immediate path planning and its broader exploration objectives. Specifically, SRPGT generates safe, dynamically updated paths based on local terrain measurements, while also selecting optimal locations for conducting further mechanical tests. Through this active exploration loop, the robot incrementally expands its map of traversable terrain while minimizing risk.

Through the integration of proprioceptive measurements and adaptive navigation, this research aims to transform the process of planetary exploration from a passive visual survey to an active, resilient, and intelligent engagement with the environment. Such capabilities are not technological conveniences but essential requirements for building a sustainable human presence beyond Earth.

The remainder of this thesis is organized as follows. Chapter 2 provides a literature review of related work in autonomous navigation. Chapter 3 details the proposed methodology, including the design of the SRPGT framework and its integration into the LASSIE exploration loop. Chapter 4 presents the simulation setup and interprets simulation results. Finally, Chapter 5 analyzes limitations of SRPGT, and outlines opportunities for future research.

Chapter 2

Literature Review

Autonomous navigation in unknown, unstructured environments poses a complex challenge, especially when sensing is limited to proprioceptive input and the terrain is composed of granular material that can deform unexpectedly underfoot. In such settings, traditional visual mapping approaches often fail, and the robot must rely entirely on local interaction to evaluate risk and plan its movement. Addressing this problem requires methods that can reason under uncertainty, ensure safety without full knowledge of the environment, and adapt motion in real time based on new evidence.

This chapter reviews two key areas of related work that inform the development of SRPGT. The first involves Bayesian optimization and Gaussian Process-based modeling to handle uncertainty and guide safe exploration. The second includes reactive control strategies, particularly those based on Voronoi and power diagram constructions, which enable real-time, geometry-aware navigation in partially known or dynamic spaces. These approaches together form the conceptual backbone of SRPGT, though the system departs from prior work by operating entirely on proprioceptive sensing in visually degraded environments.

2.1 Bayesian Optimization and Gaussian Processes in Navigation

Bayesian optimization techniques have become powerful tools for exploration under uncertainty. They use statistical models—typically Gaussian Processes (GPs)—to approximate unknown functions and select informative or optimal sampling points. In robotic navigation, this framework allows the robot to learn about the environment while avoiding unsafe terrain.

Muenprasitivej et al. [3] employ GPs to model terrain elevation and uncertainty for bipedal robots. Their method integrates footstep planning with information gain objectives, enabling safe exploration even in unstable regions. This work highlights how GP-based models can unify terrain understanding with locomotion constraints, but it assumes exteroceptive input such as elevation maps.

Uttsha et al. [4] generalize this idea by creating GP-based distance fields from 3D point clouds. Their system constructs smooth elevation and obstacle maps that support planning for legged and wheeled robots. The result is a flexible, continuous representation of traversability that allows robust trajectory optimization in uneven terrain. This framework, however, requires point cloud sensing, which may not be feasible in degraded visual conditions.

Leininger et al. [5] extend these ideas using Sparse GPs to build terrain cost maps and steer an RRT* planner. Their approach selects subgoals along the frontier to guide the robot while minimizing exposure to uncertain regions. However, it relies on global terrain observations and sampling-based planning, which can be computationally expensive and poorly suited to real-time reactivity.

All of these methods demonstrate the power of GPs for terrain-aware planning, but most require exteroceptive sensors and rely on offline or batch planning. In contrast, SRPGT uses proprioceptive data only, incrementally building a risk model in real time and making decisions at the resolution of physical contact.

2.1.1 Gaussian Process Upper Confidence Bound (GP-UCB)

The GP-UCB algorithm provides a principled approach for choosing sampling points when both uncertainty and performance must be considered. It selects the next query location x by maximizing:

$$a_{\mathrm{UCB}}(x) = \mu_{n-1}(x) + \beta\, \sigma_{n-1}(x)$$

where μ and σ represent the GP's posterior mean and uncertainty, and β scales the confidence margin. This algorithm efficiently balances exploration (uncertain points) and exploitation (high-performing predictions).

In SRPGT, a similar confidence-aware sampling strategy is used, but applied spatially to terrain traversal. Instead of seeking optima, the goal is to incrementally grow a certified safe set based on proprioceptive risk estimates.

2.1.2 Safe Bayesian Optimization (SafeOpt)

SafeOpt builds on GP-UCB by enforcing that each function evaluation satisfy a minimum safety constraint:

$$f(x) \geq h$$

with high probability. It maintains and expands a set of safe decisions while seeking high performance, avoiding evaluations in dangerous regions.

SRPGT adapts this paradigm to navigation by only sampling locations where terrain risk is predicted to fall below a safety threshold. While not a direct implementation of SafeOpt, the core idea of conservative, uncertainty-aware expansion is central to its planning mechanism.

2.2 Reactive Navigation Using Voronoi-Based Methods

Whereas GP-based methods focus on estimating terrain and selecting strategic waypoints, reactive navigation methods address a complementary need: executing safe, collision-avoiding motion in real time as new obstacles are discovered.

2.2.1 Navigation in Convex Worlds

Arslan and Koditschek [6] propose a method that uses power diagrams to define convex, obstacle-free regions around the robot. Within these regions, a local optimization strategy guides the robot toward its goal. The method guarantees convergence from almost all starting positions and is highly efficient due to its geometric construction.

Arslan and Koditschek [7] extend this framework to unknown environments by using local sensing to infer separating hyperplanes and construct reactive control fields. Their system allows real-time motion in cluttered spaces using only local information.

These methods inspire the reactive control layer in SRPGT. By constructing convex approximations of safe space based on proprioceptive estimates rather than direct obstacle detection, SRPGT achieves comparable responsiveness using internal sensing alone.

2.2.2 Navigation in Concave and Partially Known Worlds

Vasilopoulos et al. [8] combine semantic SLAM with reactive control to navigate through cluttered, dynamic environments. Their robot adjusts its plan in real time based on object detection and scene understanding, enabling smooth motion even in rapidly changing spaces.

This work is extended in Vasilopoulos et al. [9], where robots use semantic labels to decide whether to trust a prior map or act reactively. This hybrid strategy improves efficiency by leveraging prior structure while maintaining responsiveness to new hazards.

In contrast, SRPGT constructs hazard boundaries from terrain risk estimates, not visual or semantic input. The use of proprioceptive sensing makes it viable in visually degraded environments such as planetary surfaces, where dust, shadowing, or lack of light renders cameras ineffective.

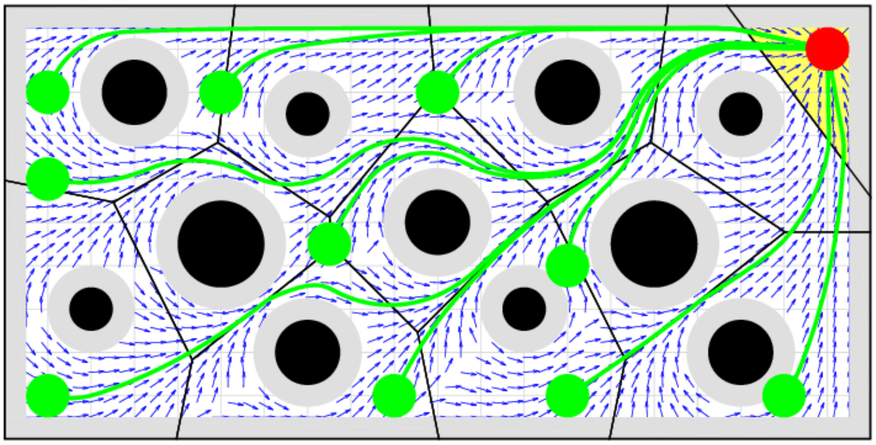

Figure 2.1: Exact robot navigation using power diagrams, generated by disks representing obstacles (black) and the robot (red at the goal). The power cell (yellow) associated with the robot defines its obstacle free convex local neighborhood, and the continuous feedback motion towards the metric projection of a given desired goal (red) onto this convex set asymptotically steers almost all robot configurations (green) to the goal without collisions along the way. The grey regions represent the augmented workspace boundary and obstacles, and the arrows show the direction of the resulting vector field. [6]

2.2.3 Other Voronoi-Based Techniques

Garrido and Moreno [10] use Voronoi diagrams with the Fast Marching Method to compute paths with maximal clearance in real time. Their technique supports continuous updates as new obstacles are sensed, reinforcing the need for responsiveness in dynamic settings.

SRPGT shares this emphasis on responsiveness, but relies on inferred terrain risk to shape its navigation domain, rather than explicit obstacle maps. The underlying philosophy—a robot should move cautiously through space while dynamically updating its understanding of nearby hazards—is closely aligned.

2.3 Summary and Motivation for This Work

Prior work in terrain-aware navigation has demonstrated the power of both probabilistic modeling and reactive control. Gaussian Processes enable safe and data-efficient exploration by modeling uncertainty, while Voronoi-based planners enable fast, obstacle-avoiding motion with minimal computation. However, these approaches have typically been applied in isolation, or require vision or global maps to function effectively.

SRPGT integrates the most compelling aspects of these two paradigms: it uses confidence-guided expansion of safe terrain based on proprioceptive input, and it executes motion in real time using a reactive controller adapted from power diagram-based methods. This combination enables the robot to safely explore deformable granular terrain with no visual information, no global map, and no need for complete re-planning in response to environmental change.

In doing so, this work contributes a unified navigation architecture that is both safe and adaptive, capable of expanding its operational zone and navigating within it using only what the robot feels through its limbs.

Chapter 3

Methods

Exploration on granular planetary terrain is a task marked by uncertainty, variability, and risk. Unlike rigid or structured environments where surface characteristics are relatively predictable, granular media such as regolith can shift suddenly under pressure. These shifts often occur without any visual cues. A successful navigation algorithm must not only plan paths toward a goal but must also adapt its behavior in real time based on what the robot learns through physical interaction.

This chapter presents the complete method developed to address these challenges. The system combines safe Bayesian global planning with geometry-based local control. The fundamental idea is that a legged robot, through proprioceptive interaction, can actively map the terrain's mechanical properties, update its internal risk estimates, and select both where to go and how to get there, while maintaining operational safety and supporting scientific objectives.

3.1 Terrain-Aware Navigation as a Sequential Decision Problem

Legged robots are uniquely suited to granular terrain because their limbs can serve two functions: enabling mobility and acting as embedded sensors. Each step provides an opportunity for measurement and becomes an input to an evolving internal model of terrain risk.

Navigation in this context can be treated as a sequential decision-making process under uncertainty. The robot must choose its next move based on local terrain observations, estimated risks, and proximity to the goal, while ensuring that all movement adheres to strict safety constraints. This challenge is met by combining two core components:

1. A probabilistic planner that incrementally expands a certified safe set of traversable terrain using terrain observations and confidence-based reasoning.

2. A reactive controller based on diffeomorphic transformation that enables real-time, geometry-aware path execution within certified safe regions.

3.2 Problem Statement

The robot operates within a two-dimensional discretized domain $D \subset \mathbb{R}^2$, where each point has an unknown terrain property, such as shear strength or cohesion, that influences traversability. The objective is to move from an initial location x_0 to a goal location x_g, ensuring that all intermediate positions lie within regions classified as safely traversable.

We define $f(x)$ as a scalar-valued function representing the traversability of terrain at location x, where higher values correspond to more easily traversable regions. This function is treated abstractly and may encode various underlying physical or empirical terrain characteristics depending on the application context. A fixed threshold h is used to determine safety: terrain is considered safe if $f(x) \geq h$, and unsafe if $f(x) < h$. The values of $f(x)$ and h are not assumed to carry physical units or universal semantics; rather, they serve as internally consistent constructs for distinguishing safe and unsafe terrain within the chosen representation framework.

Key assumptions and constraints are as follows:

- Terrain risk is a continuous but unknown function over space.

- Proprioceptive terrain measurements are available only at the robot's current location.
- Navigation must maintain high-confidence safety guarantees at all times.

3.3 Confidence-Guided Terrain Expansion

The planner maintains a probabilistic model of terrain risk using a Gaussian Process (GP). At each step, it updates terrain predictions and refines a subset of the environment certified as safe.

This method draws inspiration from prior work on safe exploration algorithms such as SafeOPT and SafeUCB [11, 12], but differs in structure and execution. In particular, the focus of this work is not on maximizing a reward function, but on expanding reachable terrain while ensuring probabilistic safety at every decision point.

3.3.1 Gaussian Process Terrain Modeling

The robot uses proprioceptive data to train a Gaussian Process (GP) model over D, producing a mean prediction $\mu_t(x)$ and standard deviation $\sigma_t(x)$ at each location $x \in D$. The GP uses a radial basis function (RBF) kernel:

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

where σ^2 is the kernel variance and ℓ is the lengthscale, both selected manually based on the expected variability of the proprioceptive signal. These hyperparameters are fixed at deployment time.

The model is updated recursively using all past measurements for computational efficiency. At each step, the robot collects a new proprioceptive terrain measurement y_t at position x_t, then incorporates this into the GP model to update posterior predictions across the domain.

3.3.2 Safe Set Expansion

A fixed exploration parameter β_t is used to control the width of confidence intervals:

$$Q_t(x) = \left[\mu_{t-1}(x) \pm \sqrt{\beta_t}\,\sigma_{t-1}(x)\right] \tag{3.1}$$

This formulation is adapted from the SafeOPT algorithm introduced by Sui et al. [11], where confidence intervals are used to balance safety with informative sampling. The value of β_t should be chosen via empirical tuning to balance conservative exploration with steady progress.

The resulting confidence bounds are intersected with prior confidence sets to form updated constraints:

$$C_t(x) = C_{t-1}(x) \cap Q_{t-1}(x) \tag{3.2}$$

From these sets, the lower and upper bounds are defined as:

$$\ell_t(x) = \min C_t(x), \quad u_t(x) = \max C_t(x) \tag{3.3}$$

and the interval width is computed as:

$$w_t(x) = u_t(x) - \ell_t(x) \tag{3.4}$$

Algorithm 1 Safe Optimization for 2D Navigation

Require: Sample set D, GP prior (μ_0, k, σ_0), Lipschitz constant L, seed set S_0, safety threshold h, goal location x_g, number of expanders to consider n

$C_0(x) \leftarrow [h, \infty)$, $\forall x \in S_0$
$C_0(x) \leftarrow \mathbb{R}$, $\forall x \in D \setminus S_0$
$Q_0(x) \leftarrow \mathbb{R}$, $\forall x \in D$
for each time step $t = 1, 2, \ldots$ **do**
    $C_t(x) \leftarrow C_{t-1}(x) \cap Q_{t-1}(x)$
    $S_t \leftarrow \bigcup_{x \in S_{t-1}} \left\{x' \in D \,\middle|\, \ell_t(x) - L\|x - x'\| \geq h\right\}$
    $G_t \leftarrow \left\{x \in S_t \,\middle|\, g_t(x) > 0\right\}$
    Sort G_t by $\|x - x_g\|$ in ascending order
    Let $G_t^{(n)} \subseteq G_t$ be the top n nearest expanders
    $x_t \leftarrow \arg\max_{x \in G_t^{(n)}} w_t(x)$
    Observe $y_t = f(x_t) + n_t$
    Compute $Q_t(x)$ for all $x \in S_t$
end for

To identify G_t, the function

$$g_t(x) = \left|\left\{x' \in D \setminus S_t \,\middle|\, u_t(x) - L\|x - x'\| \geq h\right\}\right| \tag{3.5}$$

is defined, evaluating the expansion potential of each point.

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.

The robot selects the next point to sample by balancing two objectives:

- Reducing model uncertainty (exploration)
- Making progress toward the mission goal (exploitation)

After constructing the safe set S_t and identifying the candidate expanders $G_t \subseteq S_t$, the robot must select the next point at which to sample terrain. Rather than choosing purely based on uncertainty or purely based on proximity to the goal, SRPGT employs a hybrid strategy to guide exploration in a goal-directed yet uncertainty-aware manner.

3.3.2.1 Next Parameter Selection

The new selection procedure is as follows:

1. For each point $x \in G_t$, compute its distance to the goal $\|x - x_g\|$.

2. Sort all points in G_t in ascending order of distance to the goal.

3. Consider only the top n nearest expanders, where n is a configurable parameter.

4. Among these n expanders, select the point with the largest confidence interval width $w_t(x)$.

This approach prioritizes expanders that are spatially aligned with the mission objective while still favoring those that offer informative terrain data. It mitigates the risk of wasting samples on irrelevant or backward directions and prevents purely greedy selection based on proximity alone, which might lead into regions of high uncertainty without expanding the safe set.

By focusing on a small subset of candidate expanders that are both near the goal and informative, the planner ensures steady, efficient progress toward the goal while incrementally enlarging the verified safe zone.

3.4 Reactive Voronoi-Based Navigation

While the planner identifies the next point to move toward, executing safe and efficient motion between steps requires a robust local navigation strategy. Classical global planners, such as A* and RRT*, require full replanning after significant environmental updates, making them unsuitable for highly dynamic, partially known terrains.

Instead, a reactive control framework inspired by Vasilopoulos et al. [9] is implemented.

3.4.1 Synthesizing Pseudo-Physical Obstacles

To support reactive navigation, the robot must construct geometric representations of both safe and hazardous terrain regions from its evolving internal risk map. This process transforms pointwise risk predictions into polygonal approximations of safe space and its complement, enabling geometry-aware motion planning.

The environment is discretized into a two-dimensional grid. In simulation, this grid is aligned with the resolution of the ground truth map for consistency. In real-world deployments, however, the resolution becomes a tunable parameter. Higher resolutions enable more detailed boundary tracking but increase computational overhead, while lower resolutions reduce map fidelity in favor of speed and memory efficiency.

The conversion from safe points to obstacle geometry proceeds as follows:

1. **Cluster Formation:** The current safe set S_t is partitioned into disjoint spatial clusters $\{S_{ti}\}$ using a combination of KD-tree nearest-neighbor queries and union-find data structures. Each cluster represents a contiguous region of confidently traversable terrain.

2. **Concave Hull Construction:** For each cluster S_{ti}, a concave hull H_{ti} is generated using the algorithm described by Park and Oh [13], which takes as parameters a relative concavity level $\mathcal{C}$ and a local threshold L_{th}. These settings determine the hull's shape resolution and computational complexity. In our implementation, we select high concavity values to tightly wrap each cluster. While these parameters influence geometric fidelity, they do not affect overall safety guarantees and may be tuned based on application requirements.

3. **Workspace Definition:** A bounding rectangle $\mathcal{W}_e$ is formed to enclose all H_{ti} regions:

$$\mathcal{W}_e := \left\{ x \in \mathbb{R}^2 \,\middle|\, \min_i H_{ti}^{(1)} \leq x_1 \leq \max_i H_{ti}^{(1)},\ \min_i H_{ti}^{(2)} \leq x_2 \leq \max_i H_{ti}^{(2)} \right\} \tag{3.6}$$

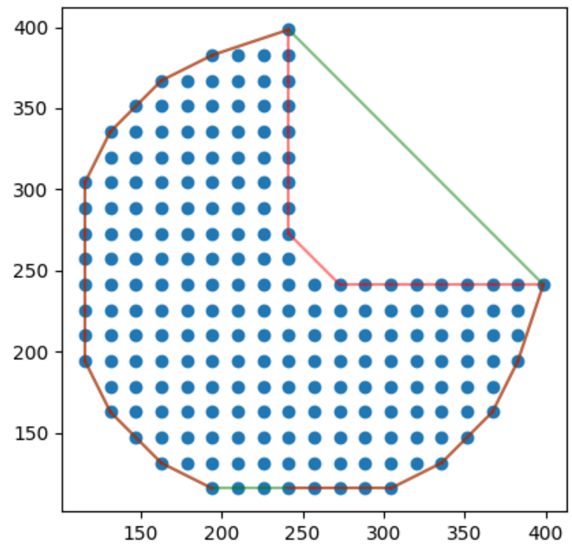

Figure 3.1: Visualization of a concave hull generated from point clusters. Image adapted from the `cubao/concave_hull` GitHub repository [14]. Licensed under the MIT License.

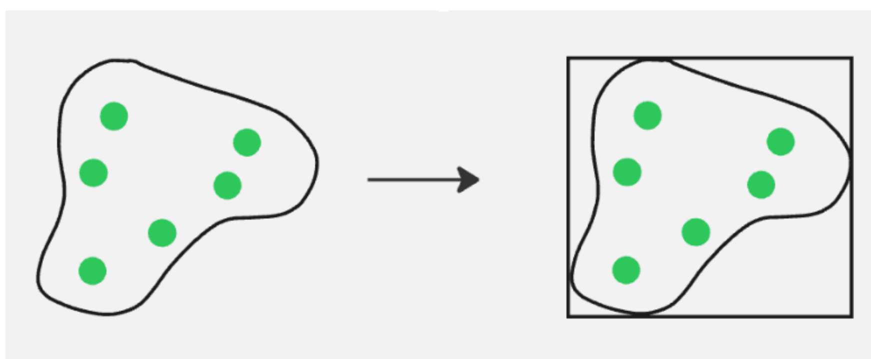

Figure 3.2: Visualization of bounding box creation. Green dots are fictional locations of safe samples, inner shape is depiction of safe area.

4. **Obstacle Extraction:** The regions not covered by the union of safe hulls are treated as pseudo-physical obstacles:

$$\mathcal{O}_t := \mathcal{W}_e \setminus \bigcup_i H_{ti} \tag{3.7}$$

5. **Polygon Simplification:** The obstacle set $\mathcal{O}_t$ is polygonal by construction but may contain excessive vertices. Each obstacle polygon is simplified using the Douglas-Peucker algorithm to reduce computational cost while preserving topology for subse-

quent triangulation and transformation.

The resulting polygon set $\mathcal{O}_t$ is used by the reactive controller to generate safe, smooth navigation commands around terrain hazards that have been inferred from contact-based measurements.

3.4.2 Diffeomorphic Mapping and Reactive Control

To enable real-time obstacle-aware motion without full replanning, we adopt a reactive navigation framework based on the approach proposed by Vasilopoulos et al. [9]. Their method provides a systematic means to guarantee obstacle avoidance and goal convergence in planar environments with polygonal obstacles, by leveraging a diffeomorphic transformation to map the physical space into a geometrically simplified model space.

We implement the core structure of their controller and adapt it to operate on pseudo-physical obstacles derived from proprioceptive terrain assessments. The application of this method in our context is novel, particularly in how the safe terrain regions are estimated and transformed dynamically based on risk-driven sampling.

Each polygon in the obstacle set $\mathcal{O}_t$ is first triangulated using the Ear Clipping method [15]. The resulting set of triangles is organized into a tree structure, where leaf triangles are recursively mapped onto their parents, and the root triangle is ultimately mapped onto a disk. This hierarchical transformation process defines a global diffeomorphism $\mathbf{h}(x)$ from the original workspace to a simplified model space free of obstacles.

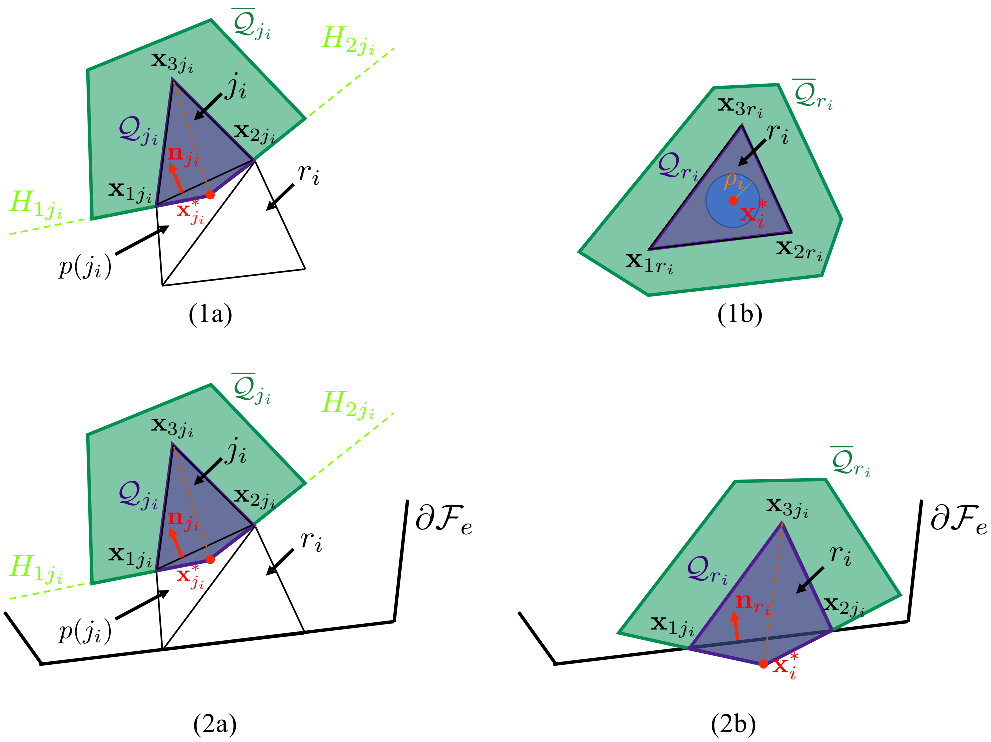

Figure 3.3: Top row: (1a) A leaf triangle is mapped onto its parent triangle; (1b) A root triangle is mapped onto a disk centered at $\mathbf{x}_i$ with radius ρ_i. Bottom row: (2a) A leaf triangle mapped onto its parent; (2b) A root triangle mapped onto the freespace border. Figure reproduced from Vasilopoulos et al. [9].

In this transformed space, navigation reduces to a gradient-based controller that drives the robot toward the goal projected into model space. The control law is defined by:

$$\mathbf{v}(\mathbf{y}) = -\left(\mathbf{y} - \prod_{\mathcal{LF}(\mathbf{y})} (\mathbf{y}_d)\right) \tag{3.8}$$

where $\mathbf{y} = \mathbf{h}(\mathbf{x})$ is the transformed robot position, $\mathbf{y}_d = \mathbf{h}(\mathbf{x}_d)$ is the transformed goal, and $\Pi_{\mathcal{LF}(\mathbf{y})}$ denotes the projection operator onto the local Voronoi cell around $\mathbf{y}$, denoted $\mathcal{LF}(\mathbf{y})$.

The corresponding control command in the original workspace is obtained by pulling back the model space vector field through the inverse Jacobian of the diffeomorphism:

$$\mathbf{u}(\mathbf{x}) = k \left[D_x \mathbf{h}\right]^{-1} \mathbf{v}(\mathbf{h}(\mathbf{x})) \tag{3.9}$$

where k is a user-defined gain parameter. The gain is kept constant and empirically tuned for the simulation, with the expectation that it would be adjusted in hardware based on robot actuation limits.

To enforce feasibility within physical constraints, a fixed maximum velocity cap is applied to $\mathbf{u}(\mathbf{x})$. This ensures that commanded motions remain within achievable limits, despite the abstraction of full actuation in the control model. This assumption is reasonable for legged robots such as quadrupeds, which can approximate omnidirectional planar motion using whole-body control.

While the mathematical structure of the diffeomorphic mapping and control law is adopted directly from Vasilopoulos et al. [9], its integration with a terrain risk-aware planning system and adaptation to proprioceptively inferred pseudo-obstacles represents a novel application in this work.

3.5 System Integration

The final system integrates global terrain expansion and local reactive control into a unified exploration loop:

- Proprioceptive measurements are gathered during motion.
- The GP-based model of terrain risk is updated after each measurement.
- The planner selects a new target location from safe or candidate expander sets.
- The reactive controller executes real-time movement toward the selected target.

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

Chapter 4

Results

This chapter presents simulation-based validation of SRPGT, designed to assess its performance in a variety of terrain scenarios. The experiments evaluate SRPGT in a variety of terrain scenarios that simulate challenges encountered during planetary exploration. These scenarios are designed to demonstrate the algorithm's core capabilities: its ability to expand safe regions under uncertainty, navigate around high-risk areas, reuse learned terrain knowledge, and perform exploratory mapping without preassigned goals.

4.1 Evaluating the Navigation Framework

4.1.1 Validation Objectives and Contribution Scope

SRPGT fills a gap not currently addressed by existing navigation algorithms: it enables **proprioceptive, confidence-aware, reactive navigation** in unknown, unstructured, and deformable environments where visual sensing is unreliable or unavailable.

While prior work has addressed uncertainty in terrain modeling, safe exploration, or reactive control in isolation, this framework integrates all three in a practical, simulation-validated approach. Specifically, the method combines:

- Uncertainty-guided subgoal selection based on terrain risk predictions;
- Active expansion of a certified safe set;

- Real-time navigation using diffeomorphic control in non-visual, proprioceptive settings.

This section presents simulation experiments designed to demonstrate these capabilities and compare their effectiveness to relevant prior work. The most closely related work is that of Leininger et al. [5], which combines Sparse Gaussian Processes with RRT*-based planning, with which the differences with SRPGT are discussed in subsubsection 4.1.4.2

4.1.2 Simulated Planetary Terrain and Initialization

To execute SRPGT, it is assumed that a set of initial known safe points is available to form the starting safe zone. In a real-world application, the robot would be physically placed within a verified safe region and would take several initial steps around this area to collect proprioceptive data, thereby building the initial terrain model. For simulation purposes, the robot is initialized in a designated safe region, and a set of randomly sampled points from a known safe region is used to initialize the Gaussian Process model, emulating initial contact-based measurements.

The test environments are represented as two-dimensional grids, where each cell contains a value corresponding to terrain properties such as risk or traversability. For simplicity, the ground truth environment is discretized in simulation, as defining a discrete grid is considerably more straightforward than specifying a continuous ground truth function. In practical real-world applications, the underlying terrain properties are continuous; however, they are discretized for computational purposes during processing and analysis.

For ease of implementation in simulation, the parameter set is defined as the collection of all possible coordinate points within the discrete ground truth map. This representation assumes that the discrete map provides the maximum available resolution of environmental information.

Because the traversability function $f(x)$ is only defined abstractly in the methods section, a specific interpretation must be assigned when instantiating it for simulation. In these experiments, $f(x)$ is constructed as a scalar field encoding relative terrain strength according to an internal definition. As a result, the safety threshold h used to distinguish between safe and unsafe regions must also be selected empirically. Its role is not to enforce an externally calibrated notion of safety, but rather to define relative risk boundaries under a particular terrain encoding. Thresholds such as $h = 1000$ are chosen for demonstration purposes and vary across scenarios to test different aspects of algorithm behavior.

With this simulation environment in place, the following scenarios were designed to evaluate distinct capabilities of SRPGT. Each experiment highlights a different navigation objective or environmental challenge relevant to planetary exploration.

4.1.3 Simulation Results

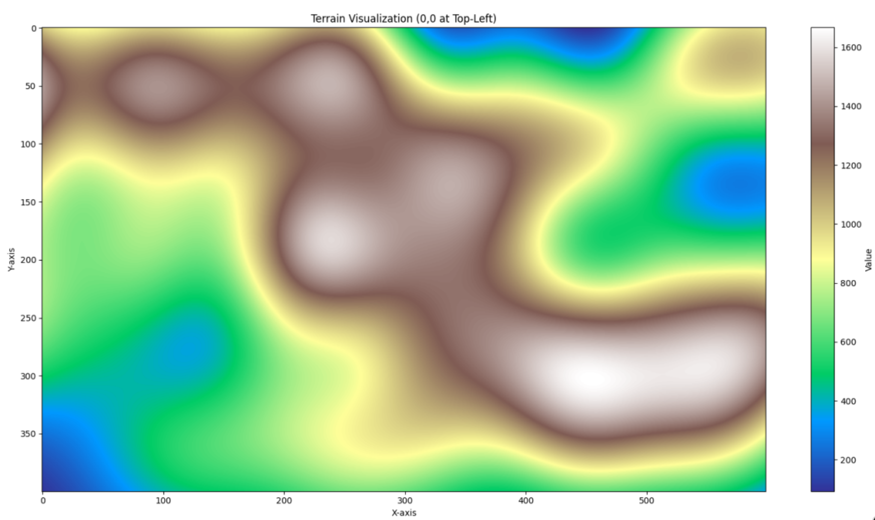

Figure 4.1: Visualization of ground truth risk values used in the following experiments, starting at $(230, 220)$, with a goal at $(104, 82)$.

4.1.3.1 Visualization Legend

To aid in interpreting the following figures, we define the color scheme used to represent key regions and elements within the navigation environment:

- **White:** Ground truth safe terrain ($f(x) \geq h$)

- **Gray:** Ground truth unsafe terrain ($f(x) < h$)
- **Black:** Pseudo-physical obstacles ($\mathcal{O}_t$)
- **Green:** Local freespace ($\mathcal{LF}(\mathbf{x})$ eroded by robot radius)
- **Light Blue:** Safe set (S_t)
- **Red:** Expanders ($G_t \subseteq S_t$)
- **Pink:** Intermediate goal (selected $x_t \in G_t$)
- **Dark Blue:** Final goal (x_g)

This legend applies uniformly across all following figures in this chapter unless otherwise noted.

4.1.3.2 General Navigation

This scenario shows the robot navigating in a scenario where the straight line path to the goal does not contain any risk zones. The robot quickly expands the safe zone as shown in Figure 4.2, reaching the point in 54 iterations.

4.1.3.3 Navigation Around a Risk Zone

In these scenarios, the robot must reach a designated goal point located beyond a high-risk region. While the direct path to the goal intersects with terrain classified as unsafe according to the risk estimation model, the robot successfully navigates around this obstacle.

Starting with a limited known safe region, the robot employs SRPGT to iteratively expand its accessible area until reaching the goal point. Figure 4.3 and Figure 4.4 illustrate this process through 9 frames captured at equal intervals throughout two navigation tasks.

The complete navigation path was generated after 210 and 461 planning iterations, where each iteration includes one terrain measurement and safe set update. Initially, the robot

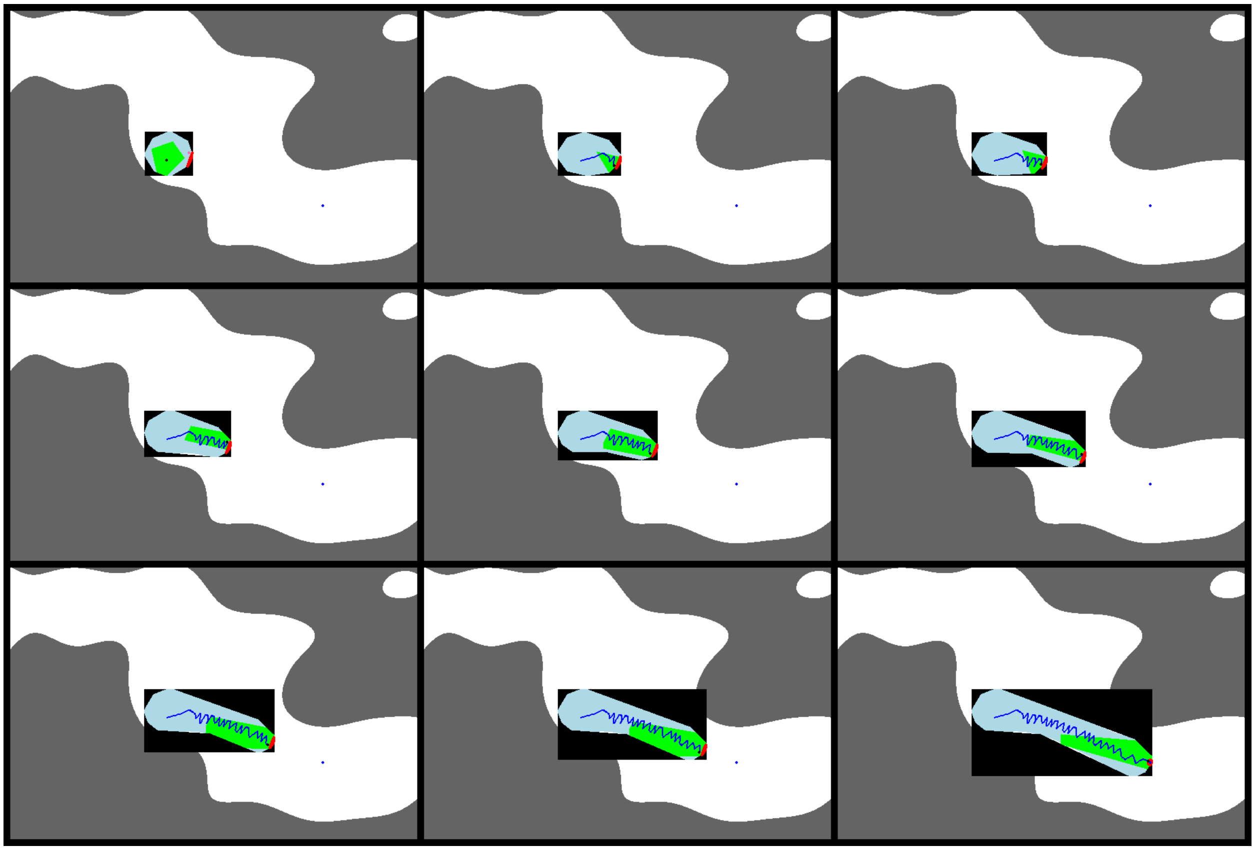

Figure 4.2: Nine frames showing the progression of the simple navigation task at evenly spaced intervals. Gray zones represent areas where traversability f(x) is less than 1000.

rapidly expands the explored safe zone until encountering regions approaching the risk threshold. As the robot gets closer to the risk zone, the algorithm selects exploration points around the periphery of the high-risk area, ensuring the robot maintains a safe distance from dangerous terrain while progressing toward the goal. Then, when the robot passes the risk area, each safe set expansion becomes much larger and less constrained, reaching the goal quickly.

A learned safe path can be reused in a subsequent navigation task, as can be seen in the following test scenario.

4.1.3.4 Traversing Across Known Safe Area to Reach Opposite Side

This experiment demonstrates how the robot leverages previously acquired knowledge about safe regions. After exploring for some time and creating a larger known safe area, the robot

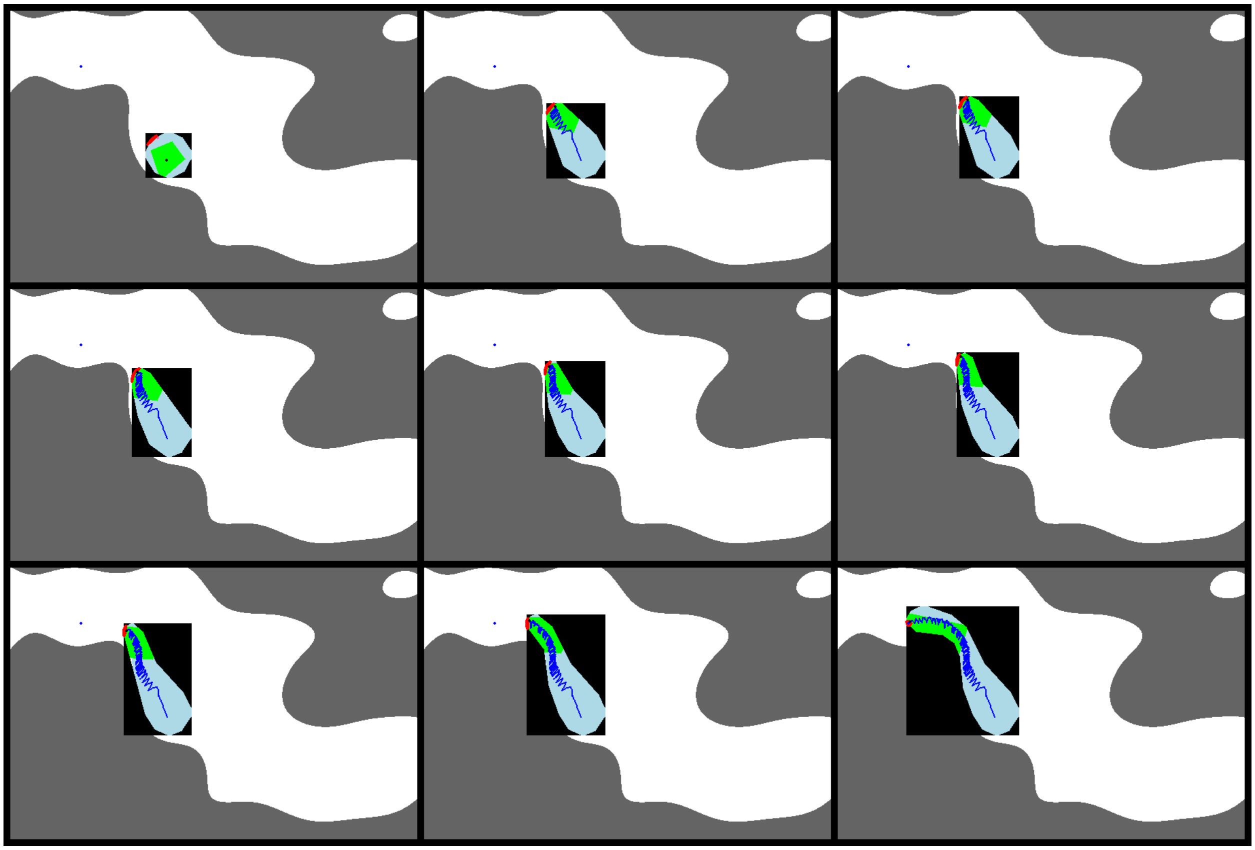

Figure 4.3: Nine frames showing the progression of the navigation task around a risk zone at evenly spaced intervals. Gray zones represent areas where traversability f(x) is less than 1000.

was positioned at a new starting point with the goal of reaching the opposite side of the environment.

As illustrated in Figure 4.5, the robot efficiently navigates across a previously learned safe corridor around the high-risk area. This traversal completed in only 10 iterations, as the robot was able to utilize the large safe zone to travel without stopping, only beginning to take new measurements after reaching the first sub-goal.

The robot maintains its safety-conscious behavior throughout the traversal, staying within the previously established safe zones and avoiding unnecessary re-exploration of the environment. This capability is particularly valuable in applications where repeated navigation through partially known environments is required, such as periodic inspection tasks or multi-objective missions.

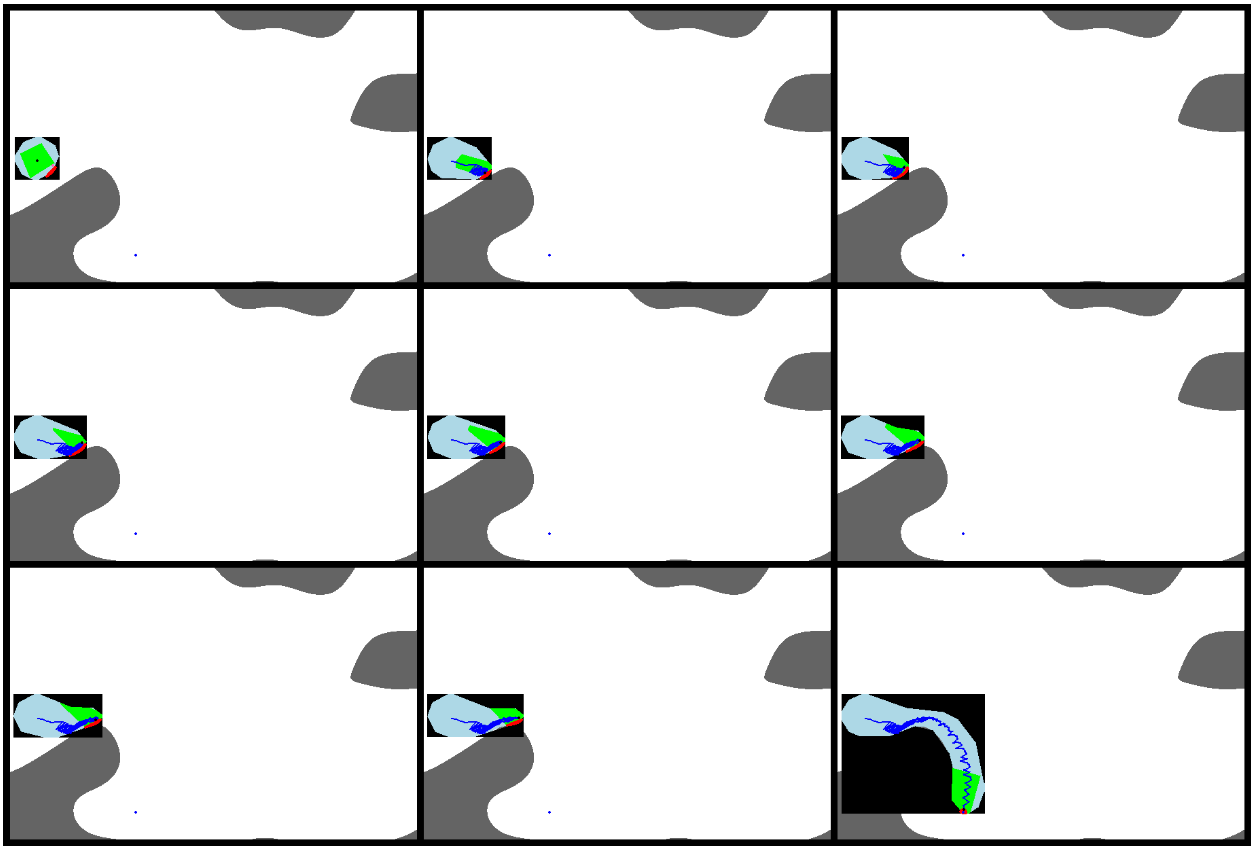

Figure 4.4: Nine frames showing the progression of the navigation task around a risk zone at evenly spaced intervals. Gray zones represent areas where traversability f(x) is less than 500.

4.1.3.5 Exploring Without a Goal in Mind

When given the task to simply explore without a goal in mind, the robot selects sub-goals that maximize terrain uncertainty reduction, independent of any goal location. As can be seen in Figure 4.6, the robot selects points where the most information can be revealed about the map. It expands the safe zone as quickly as possible, resulting in a larger known area, shown in Figure 4.7, faster than in a goal-oriented navigation task.

This exploration behavior demonstrates how SRPGT can be applied to general environmental mapping tasks without requiring predefined destinations. The algorithm prioritizes points at the frontiers of the known safe region, systematically expanding the mapped area while maintaining safety constraints. In just 9 iterations, the robot has successfully mapped

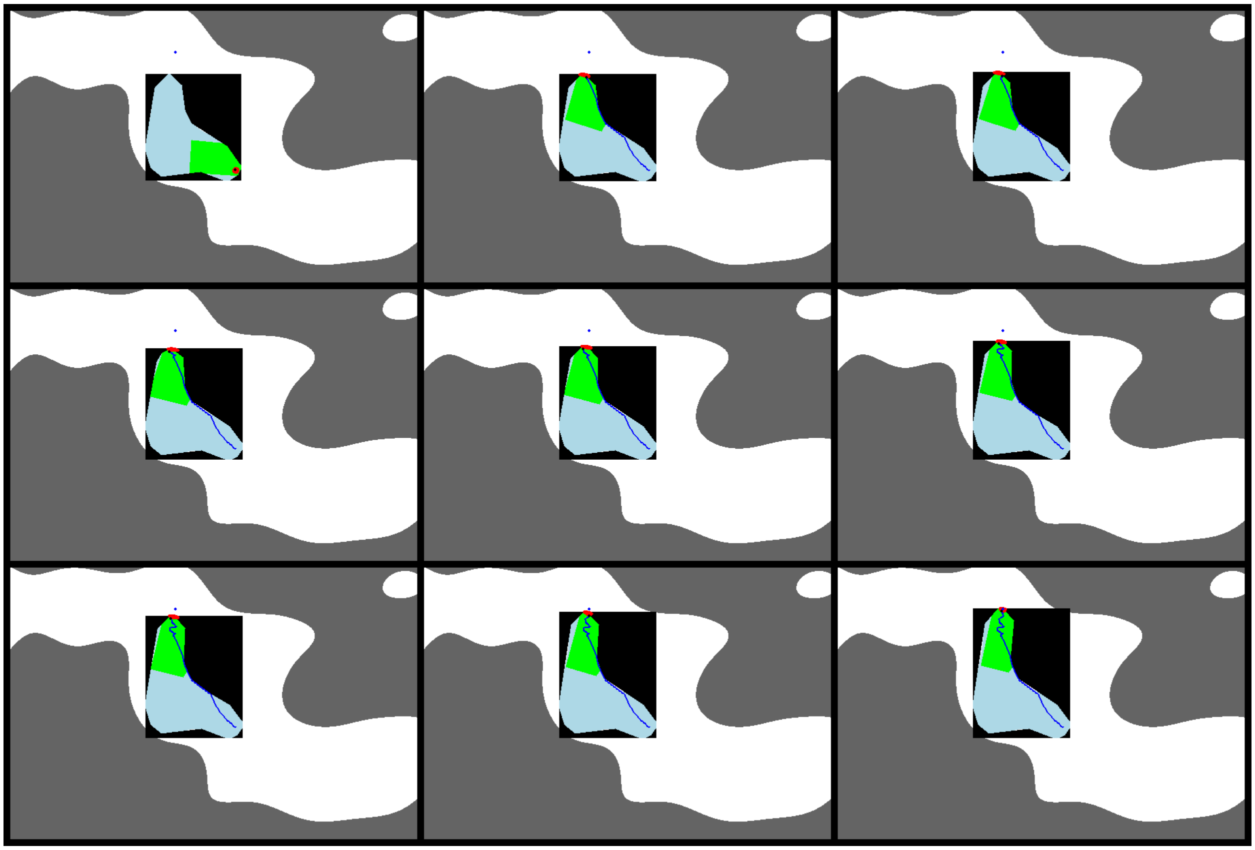

Figure 4.5: Nine frames showing the progression of the navigation task at evenly spaced intervals. Gray zones represent areas where traversability f(x) is less than 1000.

a significant portion of the accessible environment, creating a comprehensive safety map.

This scenario proves particularly useful for robots that have just deployed in a small safe area without prior inspection of the surroundings. Running an exploration task before assigning an explicit navigation goal allows the system to build a safety model of the environment, which can subsequently enhance the efficiency of goal-directed tasks by reducing the need for cautious exploration during navigation.

4.1.4 Comparisons to Baseline Methods

4.1.4.1 Comparison to Reactive Baseline Navigation

To further contextualize the performance of SRPGT, we compare it to a standard reactive navigation baseline shown in Figure 4.8 that lacks probabilistic terrain modeling and safe

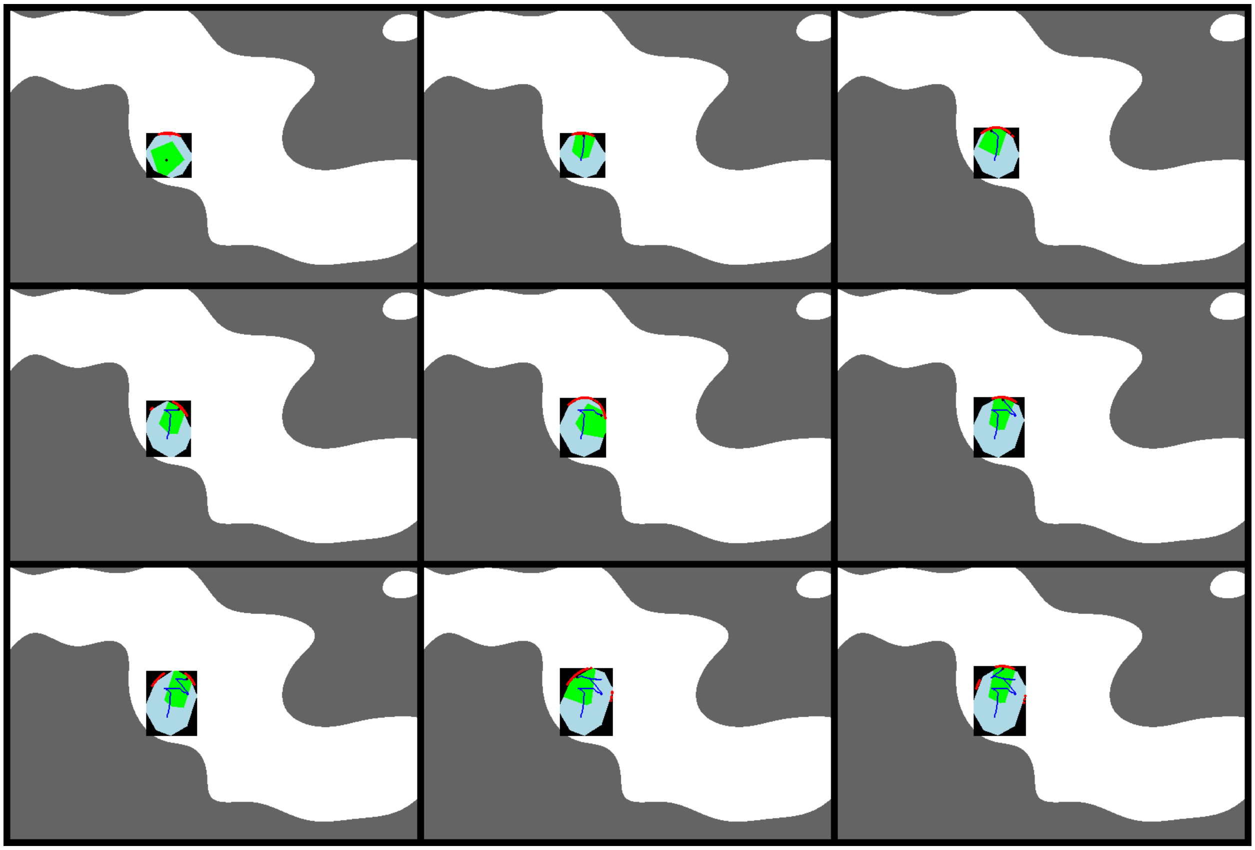

Figure 4.6: The first nine frames of the exploration task. Gray zones represent areas where traversability f(x) is less than 1000.

set reasoning. The baseline method operates under the assumption that the robot can safely navigate directly toward the goal along the shortest Euclidean path, adjusting only when immediate obstacles are encountered.

This naive strategy fails to account for terrain risk unless the risk manifests as an explicit obstacle in a local sensing field. As a result, when the shortest path intersects a high-risk region not detectable by general sensing methods, the robot proceeds directly through it, leading to unsafe behavior and frequent failures. In contrast, SRPGT actively models the risk of terrain using a Gaussian Process and expands a certified safe set around the robot. It selects subgoals based not solely on distance but on estimated safety and uncertainty.

The introduction of the safe zone representation and confidence-based subgoal selection allows SRPGT to anticipate and avoid unsafe terrain even before it is physically encoun-

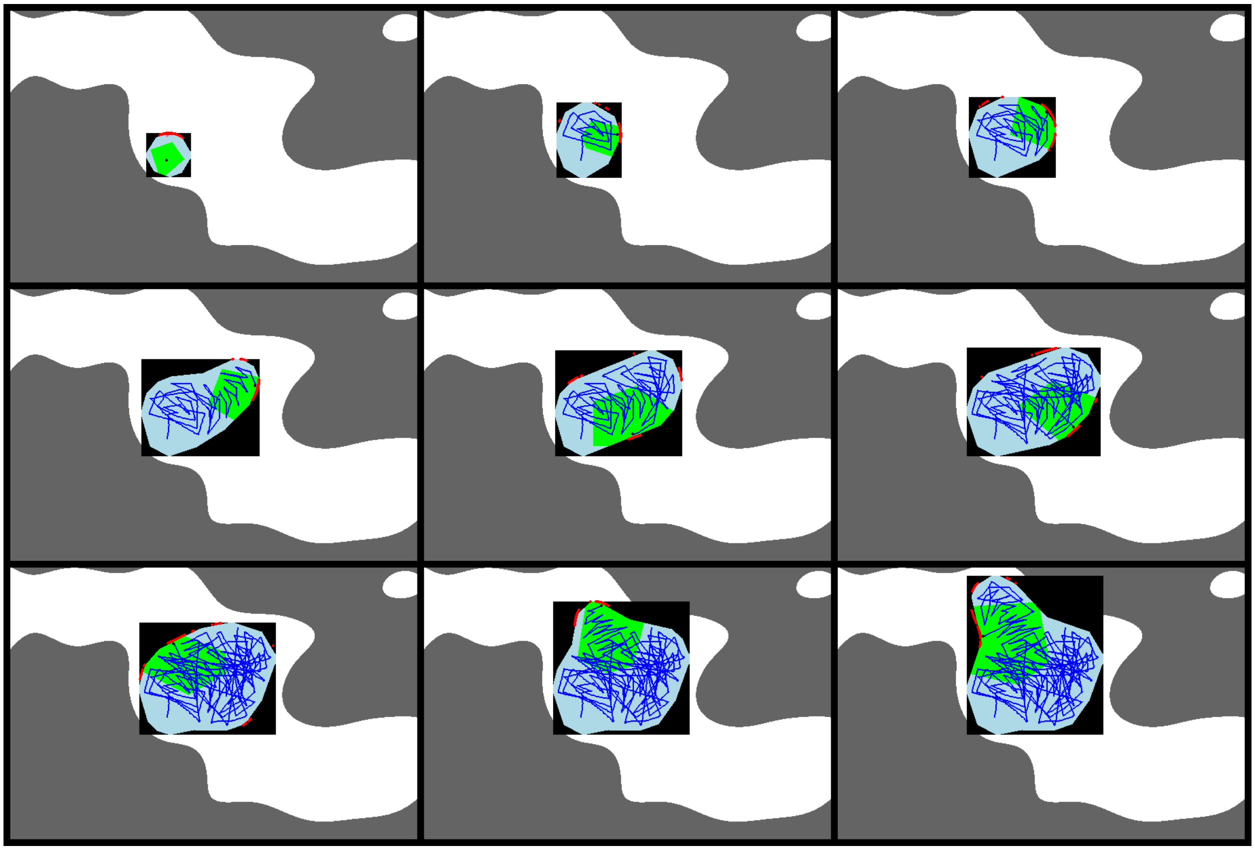

Figure 4.7: Nine frames showing the overall progression of the exploration task. Gray zones represent areas where traversability f(x) is less than 1000.

tered. Across all scenarios presented in this chapter—including those with complex terrain topographies and non-convex risk zones—SRPGT achieved a 0% failure rate, provided that the hyperparameters (e.g., Lipschitz coefficient L, exploration parameter β) were appropriately tuned to accurately reflect the underlying terrain characteristics.

This performance contrast highlights the fundamental advantage of combining proprioceptive sensing, probabilistic terrain modeling, and confidence-aware planning. Without these components, as shown in the baseline strategy, the robot is prone to entering unsafe terrain. With SRPGT, the robot instead exhibits foresight, adaptability, and safety-conscious behavior—capabilities essential for autonomous planetary exploration in uncertain and deformable environments.

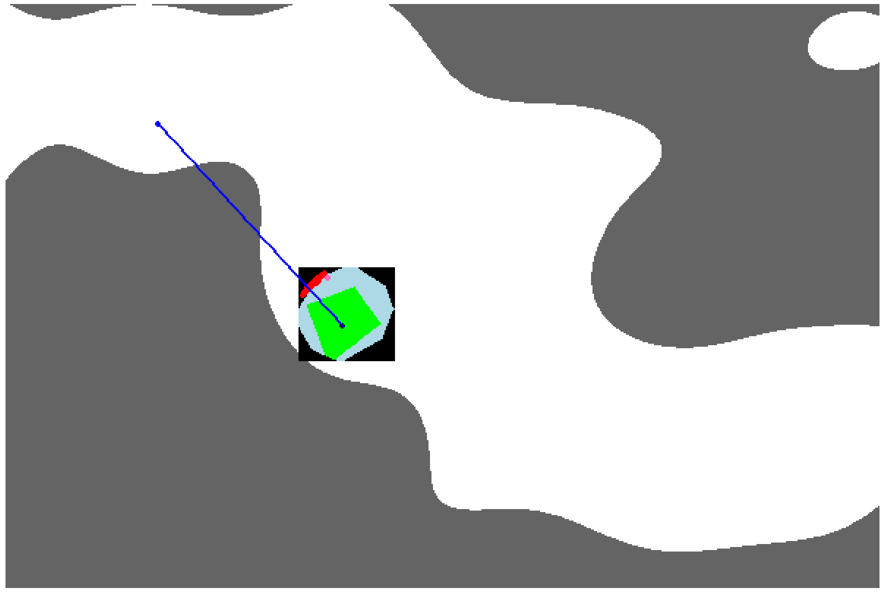

Figure 4.8: The baseline navigation result, as shown for the scenario in Figure 4.3.

Table 4.1: Performance Comparison Between SRPGT and Reactive Baseline

Metric	**Reactive Baseline**	**SRPGT**
Success Rate	Low; guaranteed failure when passing risky terrain	High; 0% failure rate with accurately tuned hyperparameters
Path Length	Shortest-path by default, but often unsafe or infeasible	Slightly longer paths that avoid risk zones; consistently feasible
Computational Efficiency	Low; minimal computation, but unsafe	Moderate; additional overhead for GP and confidence checks, but manageable in real time

4.1.4.2 Comparison to Prior Gaussian Process-Based Planning

While the reactive baseline illustrates the importance of terrain modeling and safety-aware planning, it is also valuable to compare SRPGT against more structured, probabilistically informed methods. One such method is the framework proposed by Leininger et al. [5], which combines Sparse Gaussian Processes (SGP) with RRT*-based motion planning. This approach offers global terrain reasoning and path planning capabilities using uncertainty-aware models. However, SRPGT distinguishes itself through a number of architectural and operational design choices that make it more suitable for granular terrain exploration using local feedback.

4.1.4.2.1 Planning Strategy. Leininger et al. [5]'s approach constructs global paths to the goal using RRT* over a probabilistic terrain cost map. This requires maintaining and updating a global terrain estimate and is most effective in structured or semi-known environments. In contrast, SRPGT takes a reactive planning approach that forgoes global path computation in favor of incremental, confidence-aware subgoal selection from a locally expanding safe set. This supports greater flexibility in dynamically evolving or unknown environments.

4.1.4.2.2 Uncertainty Handling. While both methods utilize Gaussian Process models to estimate terrain risk and uncertainty, SRPGT directly incorporates the GP confidence bounds into its decision-making loop. The confidence intervals define both the safe and expander sets that determine feasible and informative next steps. In the SGP-RRT* formulation, uncertainty is treated as a scalar cost modulation rather than as a constraint mechanism to enforce safety under uncertainty.

4.1.4.2.3 Sensing Assumptions. A critical difference lies in sensing modality. The SGP-RRT* planner assumes access to exteroceptive observations, such as visual or remote terrain features, to update the GP model. In contrast, SRPGT is built for scenarios in which such sensing may be unavailable or unreliable. It relies solely on proprioceptive feedback gathered through direct terrain interaction—making it particularly well-suited for planetary missions in dust-obscured or visually degraded environments.

Taken together, these distinctions highlight how SRPGT extends the capabilities of uncertainty-aware terrain modeling into domains where traditional planning pipelines fall short. The integration of proprioceptive sensing, real-time safe set expansion, and local control makes SRPGT more resilient in scenarios where structural assumptions break down.

4.2 Conclusions from Simulation Trials

The simulation experiments presented in this chapter validate the intended capabilities of SRPGT under a variety of terrain conditions and risk scenarios. The robot successfully performs 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.

Chapter 5

Conclusion and Future Work

5.1 Conclusion

This thesis presented a novel navigation framework for proprioceptive terrain-aware exploration in granular environments, where conventional visual sensing is unreliable or unavailable. The system was designed to address the core challenges of uncertainty, deformability, and incomplete knowledge that arise in planetary terrain scenarios. By integrating Gaussian Process-based terrain modeling, confidence-guided safe set expansion, and diffeomorphic reactive control, SRPGT enables legged robots to safely and adaptively explore unknown environments using only local physical interaction.

Through simulation-based validation, SRPGT demonstrated three key capabilities: safely navigating around high-risk terrain, efficiently reusing previously learned safe paths, and rapidly expanding known terrain through exploratory behavior. Unlike prior approaches that rely on global visual maps or precomputed paths, SRPGT achieves safety-aware navigation using real-time proprioceptive measurements and confidence-aware planning, filling a critical gap in current planetary robotics literature.

This work contributes a unified framework for autonomous, local-information-driven exploration under risk constraints, opening new possibilities for robotic operation in visually degraded and mechanically unstable environments.

5.2 Future Work

While the proposed method successfully addresses many core challenges, several directions remain open for further enhancement and investigation.

5.2.1 Adapting the Kernel Model for Terrain-Specific Exploration

One limitation of the current system lies in the use of a radial basis function (RBF) kernel with a fixed length scale. Although this kernel captures smooth terrain correlations effectively, it assumes spatial homogeneity across the entire domain. In practice, terrain conditions often vary dramatically—some areas may require finer resolution due to localized instability, while others may permit broader generalization.

To address this, future work could explore the use of attentive or adaptive kernels [16] that dynamically adjust the kernel length scale based on local terrain variation or uncertainty. This would allow SRPGT to quickly sweep through regions of low complexity while cautiously probing areas with steep gradients or sparse data. Such local adaptivity may reduce sampling requirements and improve both efficiency and safety in real-world deployments.

5.2.2 Bayesian Planning to Avoid Frontier Oscillation

Another improvement involves extending the planning strategy beyond the current greedy sampling approach. At present, SRPGT selects the next expander point based on the highest local uncertainty, which can lead to inefficient "zig-zagging" behavior when traversing across a discrete frontier of uncertain terrain. This behavior, illustrated in Figure 5.1, arises because uncertainty is reduced only in the immediate vicinity of each expander, causing its neighbors momentarily become the new most uncertain.

To overcome this limitation, future versions of SRPGT could incorporate probabilistic or belief-space planning frameworks, such as partially observable Markov decision processes (POMDPs). These models consider both the robot's current knowledge and its uncertainty

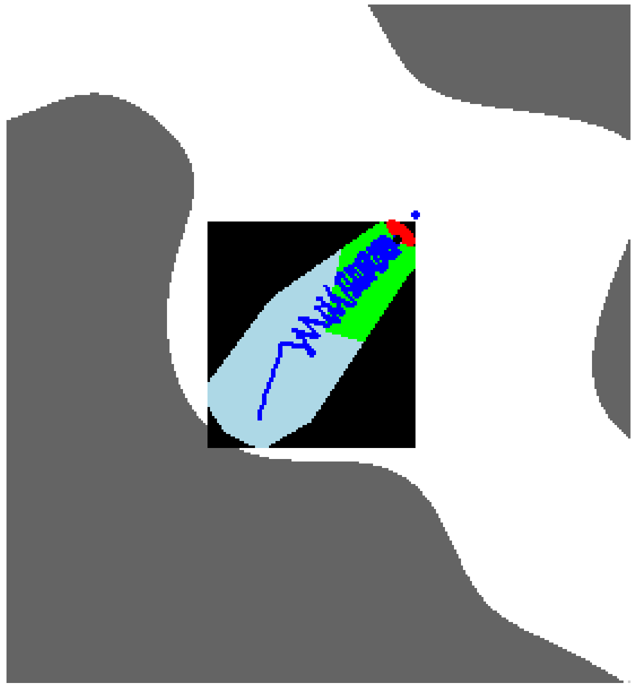

Figure 5.1: "Zig-Zag" behavior of the robot when searching in a straight line.

over future terrain, allowing actions to be chosen based on expected long-term information gain or goal reachability. Rather than selecting the immediate best expander point, the robot could simulate sequences of actions that maximize the likelihood of expanding the frontier coherently or reaching a distant goal safely.

Such methods would reduce redundant sampling and promote smoother, more globally directed navigation. Ultimately, this integration could yield an exploration strategy that balances near-term caution with long-term progress, improving performance in both open and constrained terrains.

5.2.3 Integration with Exteroceptive Sensing to Include Planning for Physical Obstacles

While SRPGT excels at evaluating terrain risk through proprioceptive interaction, it currently assumes that all safe areas are also free of physical obstructions. In practical planetary scenarios, however, safe terrain may still contain non-traversable obstacles such as rocks, outcrops, or structural debris. These features are often static, geometrically well-defined, and perceptible through standard sensing modalities such as stereo vision, LiDAR, or depth cameras—even in partially degraded visual conditions.

Future work could enhance the current framework by integrating conventional exteroceptive sensors to build geometric representations of physical obstacles within the safe set. By fusing proprioceptive terrain safety maps with geometric occupancy grids or semantic segmentation outputs, the robot could distinguish between terrain that is merely risky (e.g., soft or unstable) and terrain that is physically blocked. This would allow for more nuanced path planning decisions—for instance, avoiding areas that are both risky and cluttered, while still exploring marginally risky but obstacle-free terrain.

Furthermore, this integration would enable multi-layered constraint handling within the local control policy. The diffeomorphic mapping currently used to avoid terrain-based obstacles can be used to encode hard geometric constraints derived from point cloud clustering or shape modeling.

Incorporating geometric obstacle information opens the door to hybrid planning techniques, where a higher-level geometric planner filters feasible corridors while SRPGT locally selects among those corridors based on proprioceptive risk. This two-tiered approach could significantly improve global efficiency, especially in cluttered environments where terrain risk alone does not fully constrain the robot's motion.

Together, these enhancements would elevate SRPGT from a proprioceptive-only planner to a more complete terrain-aware motion planning framework, capable of leveraging multiple sensing modalities to ensure both geometric and mechanical safety.

Appendix A

Configuration File Parameters

The following configuration parameters are available in the code demonstration and can be modified through the provided configuration file. These settings control the environment, robot behavior, optimization strategy, and display preferences.

Environment Parameters

- **environment.FILENAME** - Specifies the terrain data file (e.g., `"terrain.csv"`).
- **environment.THRESHOLD** - Sets the safety threshold for classifying terrain as safe or unsafe.
- **environment.SIMPLIFICATION_CONSTANT** — Adjusts the simplification level of obstacle polygons derived from unsafe regions.

Robot Parameters

- **robot.ROBOT_RADIUS**: Defines the robot's physical radius, used for collision checking and power diagram erosion.
- **robot.MODE**: Specifies the robot's operation mode (e.g., `"navigate"` for goal-directed movement).

Optimization Parameters

- **optimization.NUM_EXPANDERS**: Determines the number of candidate points evaluated for safe set expansion.
- **optimization.KERNEL_VARIANCE**: Sets the variance parameter of the GP kernel for terrain risk modeling.
- **optimization.KERNEL_LENGTHSCALE**: Sets the length scale of the GP kernel, controlling spatial smoothness.
- **optimization.BETA**: Configures the confidence interval scaling factor used in the upper confidence bound (UCB).
- **optimization.LIPSCHITZ**: Defines the Lipschitz constant assumed during safe set growth calculations.

Display Parameters

- **display.BUFFER_SIZE**: Sets the buffer size for storing and visualizing simulation frames.

All parameters are editable prior to runtime and are loaded automatically at initialization. This modular design facilitates reproducibility and allows users to test different configurations without modifying source code.

Appendix B

Simulation Configuration

This appendix documents the configuration settings used to generate the simulation results presented in Chapter 4. The parameters were specified using a structured configuration file, which enabled reproducible control over the robot's behavior and the optimization environment.

Configuration File Structure

The configuration file is divided into thematic sections: `environment`, `robot`, `optimization`, and `display`. The following settings were used for Scenarios 1 and 2 (”Navigation Around a Risk Zone” and ”Traversing Across Known Safe Area”). For Scenario 3 (”Exploring Without a Goal in Mind”), the only change was setting `environment.MODE` to `"explore"`.

```
[environment]
FILENAME = "terrain.csv"
THRESHOLD = 1000
SIMPLIFICATION_CONSTANT = 4

[robot]
ROBOT_RADIUS = 2
MODE = "navigate"

[optimization]
```

```
NUM_EXPANDERS = 40
KERNEL_VARIANCE = 2
KERNEL_LENGTHSCALE = 30
BETA = 3
LIPSCHITZ = 0.003

[display]
BUFFER_SIZE = 1
```

References

[1] Randall Munroe. *Spirit*. Accessed: 2025-04-29. 2009. URL: https://xkcd.com/695/.

[2] K. R. Fisher et al. "LASSIE: Legged Autonomous Surface Science In Analogue Environments". In: *54th Lunar and Planetary Science Conference*. 2023000243. Contract Grant: 80NSSC22K1313. Johnson Space Center. The Woodlands, TX: Universities Space Research Association, Lunar and Planetary Institute, 2023.

[3] Kasidit Muenprasitivej et al. *Bipedal Safe Navigation over Uncertain Rough Terrain: Unifying Terrain Mapping and Locomotion Stability*. 2024. arXiv: 2403.16356 [cs.RO]. URL: https://arxiv.org/abs/2403.16356.

[4] Monisha Mushtary Uttsha et al. *Gaussian Process Distance Fields Obstacle and Ground Constraints for Safe Navigation*. 2024. arXiv: 2410.17831 [cs.RO]. URL: https://arxiv.org/abs/2410.17831.

[5] Abe Leininger et al. *Gaussian Process-based Traversability Analysis for Terrain Mapless Navigation*. 2024. arXiv: 2403.19010 [cs.RO]. URL: https://arxiv.org/abs/2403.19010.

[6] Omur Arslan and Daniel Koditschek. "Exact Robot Navigation Using Power Diagrams". In: May 2016. DOI: 10.1109/ICRA.2016.7487090.

[7] Omur Arslan and Daniel E Koditschek. "Sensor-based reactive navigation in unknown convex sphere worlds". In: *The International Journal of Robotics Research* 38.2-3 (2019), pp. 196–223. DOI: 10.1177/0278364918796267. eprint: https://doi.org/10.1177/0278364918796267. URL: https://doi.org/10.1177/0278364918796267.

[8] V. Vasilopoulos et al. "Reactive Semantic Planning in Unexplored Semantic Environments Using Deep Perceptual Feedback". In: *IEEE Robotics and Automation Letters (RA-L)* 5.3 (2020), pp. 4455–4462. DOI: 10.1109/LRA.2020.3001496.

[9] Vasileios Vasilopoulos et al. *Reactive Navigation in Partially Familiar Planar Environments Using Semantic Perceptual Feedback*. 2021. arXiv: 2002.08946 [cs.RO]. URL: https://arxiv.org/abs/2002.08946.

[10] Santiago Garrido and Luis Moreno. "Mobile Robot Path Planning using Voronoi Diagram and Fast Marching". In: May 2015. ISBN: 9781466686939. DOI: 10.4018/978-1-4666-8693-9.

[11] Yanan Sui et al. "Safe Exploration for Optimization with Gaussian Processes". In: *Proceedings of the 32nd International Conference on Machine Learning*. Ed. by Francis

Bach and David Blei. Vol. 37. Proceedings of Machine Learning Research. Lille, France: PMLR, 2015, pp. 997–1005. URL: https://proceedings.mlr.press/v37/sui15.html.

[12] Arpan Losalka and Jonathan Scarlett. "Benefits of Monotonicity in Safe Exploration with Gaussian Processes". In: *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence (UAI)*. 2023, pp. 1304–1314. URL: https://arxiv.org/abs/2211.01561.

[13] J.-S Park and S.-J Oh. "A New Concave Hull Algorithm and Concaveness Measure for n-dimensional Datasets". In: *Journal of Information Science and Engineering* 29 (Mar. 2013), pp. 379–392.

[14] cubao. *Concave Hull*. https://github.com/cubao/concave_hull. Accessed April 30, 2025. 2024.

[15] Hossam ElGindy, Hazel Everett, and Godfried Toussaint. "Slicing an ear using prune-and-search". In: *Pattern Recognition Letters* 14.9 (1993). Computational Geometry, pp. 719–722. ISSN: 0167-8655. DOI: https://doi.org/10.1016/0167-8655(93)90141-Y. URL: https://www.sciencedirect.com/science/article/pii/016786559390141Y.

[16] Weizhe Chen, Roni Khardon, and Lantao Liu. *AK: Attentive Kernel for Information Gathering*. 2022. arXiv: 2205.06426 [cs.RO]. URL: https://arxiv.org/abs/2205.06426.