

Slope-aware Adaptive Gaussian Process Sampling for Robotic Information Gathering on Rough Terrain

Minori Tazaki and Genya Ishigami

Abstract—Wheeled robots are increasingly utilized for environmental exploration and data collection in environments inaccessible to humans, such as the lunar surface. While adaptive sampling methods like the Upper Confidence Bound (UCB) can balance the trade-off between exploration and measurement data exploitation to model Regions of Interest (ROI), they do not explicitly incorporate robot traversability into the trade-off. Consequently, when faced with multiple scientifically promising locations, the UCB may select paths for the robot that are costly in terms of travel time, leading to time-inefficient surveys, particularly in sloped terrain. To address this limitation, we propose the Slope-aware Upper Confidence Bound (SaUCB), a novel acquisition function that integrates a traversability score based on terrain slope directly into the decision-making process. This allows the robot to explicitly balance exploration, exploitation, and traversal time. Through an extensive simulation study in realistic geological features of lunar terrain models, we demonstrate that the proposed approach generates significantly more time-efficient survey paths. Our method also demonstrates an improved trade-off between investigation time and modeling accuracy within the ROI compared to conventional approaches.

I. INTRODUCTION

In preparation for future sustained human activities on the lunar surface, the construction of lunar infrastructure has been proposed and scheduled [1]. Soil investigation on the lunar surface is an essential prerequisite for such infrastructure development for the following two fundamental reasons:


- Identifying routes for construction vehicles and suitable sites for infrastructure:

In general, firmer soil conditions prevent wheeled robots from sinking into sandy ground and reduce the risk of vehicle immobilization [2]. Furthermore, infrastructure construction sites require locations where the ground will not subside due to the weight of structures or collapse due to seismic activity [3]. Therefore, construction vehicle operation areas and infrastructure construction sites should be set in regions with high soil stiffness.


- Determination of lunar soil collection regions:
Transporting infrastructure materials from Earth would involve a substantial payload mass, leading to prohibitive launch costs. Therefore, “In-Situ Resource Utilization (ISRU)” has been proposed involving the collection and use of regolith as construction materials [4]. For regolith collection, regions with low bulk density and soft soil are preferable, as they are expected to reduce excavation

*This work was not supported by any organization

M. Tazaki and G. Ishigami are with the Space Robotics Group, Department of Mechanical Engineering, Keio University, Yokohama 223-8522, Japan minori.tazaki@keio.jp, ishigami@mech.keio.ac.jp



(a) Naive methods showing two extremes: one risks a hazardous path by treating all risks equally (UCB w/o slope constraints), while the other is too cautious and misses the ROI (UCB w/ slope constraints).



(b) Our method selectively finds a moderate risk path, with the flexibility for more aggressive or conservative paths by tuning its parameters.

Fig. 1: A conceptual comparison of (a) naive and (b) proposed adaptive sampling strategies.

resistance and improve energy efficiency. Therefore, the regolith collection area should be set in regions with low soil stiffness.

Consequently, these investigations motivate more detailed examinations of regions with either high or low soil stiffness compared to other areas. We refer to these regions as Regions of Interest (ROI). However, investigating these ROIs is not straightforward, as the magnitude of terrain slope significantly influences rovers traversing planetary surfaces. For this reason, the guidance, navigation, and control system of the robot has prioritized terrain with low slope angles, leveraging topographic information acquired through orbiting satellite data or real-time sensor measurements from onboard rover systems [5], [6]. This practice of selectively traversing only specific terrain classes with gentle slope angles poses a significant limitation, as it curtails the acquisition of highly

applicable knowledge. This is because lunar soil investigations should be driven by user requirements to explore arbitrary terrain classes, rather than restricted to easily accessible ones [7]. Therefore, lunar soil investigation for infrastructure construction requires a sampling strategy that balances the dual objectives of “exploration” for broad map coverage and “exploitation” for detailed ROI investigation.

One methodology to achieve this is adaptive sampling using Confidence Bounds (CBs), methods that include Upper Confidence Bound (UCB) for maximization problems and Lower Confidence Bound (LCB) for minimization problems [8], [9]. Both methods enable adaptive sampling by balancing exploration and exploitation. This balance is guided based on Gaussian Processes (GPs) [10], [11], which are used to estimate and model the spatial correlation structure of an environmental field. Adaptive sampling is frequently regarded as a non-parametric approach, offering the flexibility to dynamically select sampling points based on measurement data without predetermined parametric model assumptions. Indeed, knowledge about lunar soil stiffness is extremely limited, making methods relying on predetermined parametric models inapplicable. This necessitates a non-parametric approach, for which GPs are well-suited. GPs allow us to model the unknown soil stiffness by defining a prior over functions and inferring a posterior distribution directly from measurement data. Therefore, this research employs GP-UCB, which combines GPs with the UCB algorithm.

Robotic information gathering is a key capability across various domains, including environmental monitoring and modeling [12]–[14], search and rescue operations [15], and mapping applications [16], [17]. Gaussian Process Regression (GPR) is widely utilized to model spatial phenomena from discrete measurements collected by robots.

Bai *et al.* proposed a method using GP-UCB to enhance the value of information obtained by robots during autonomous exploration of unknown environments [18]. Furthermore, Tan *et al.* developed an approach to increase sampling density in ROI by incorporating the Cross-Entropy method with GP-UCB [19]. Marchant and Ramos proposed Distance-based Upper Confidence Bound (DUCB), which adds a travel distance term from the robot’s current position to the UCB formulation, thereby promoting energy-efficient sampling planning [20]. Agrawal *et al.* applied an on-the-way adaptive sampling strategy—OAS-GPUCB—for bathymetry mapping, achieving reduced mapping error while significantly cutting the vehicle’s travel distance by acquiring depth data at fixed intervals during transit [21].

While existing methods like UCB and DUCB can treat slope as a traversal constraint, they do not directly integrate it into the core exploration-exploitation trade-off. However, slope is not merely a constraint but a critical factor that dictates travel time, as rovers must reduce their speed on steep terrain to ensure stability [22]. Therefore, using standard UCB increases the required time for robotic investigation, which not only increases the robot’s energy consumption and risk of operational hazards (e.g., getting stuck), but also creates bottlenecks in the mission schedule by monopolizing scarce

robotic resources. Hence, in lunar soil investigation, a trade-off exists between survey accuracy and investigation time. To improve the Pareto front of this trade-off, it is necessary to explicitly consider exploration, exploitation, and slope angles simultaneously, thereby constructing methods that enhance ROI survey accuracy while reducing investigation time.

The main contributions of this research are as follows:

- We propose an algorithm that incorporates a slope-aware term into UCB, enabling a sampling plan that addresses the trade-off between investigation time and modeling error (Fig. 1).
- Our method outperforms baselines, including UCB, by achieving a greater Hypervolume on the evaluation metrics such as time-RMSE and time-F1 score fronts within the ROI.

II. PROBLEM FORMULATION

The robotic soil stiffness mapping problem is formulated as a bi-objective optimization problem of finding the set of locations X that minimizes both modeling error and investigation time:

$$\min_{X \subseteq \mathcal{C}} (E(X), T(X)). \quad (1)$$

In this formulation, $\mathcal{C} \subseteq \mathbb{R}^2$ is the continuous search space of all potential sampling coordinates. The outcome of the sampling process is a set of N visited locations, $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\} \subseteq \mathcal{C}$. At the i -th location $\mathbf{x}_i \in X$, the robot obtains a soil stiffness value ζ_i , yielding a final dataset $S = \{(\mathbf{x}_1, \zeta_1), \dots, (\mathbf{x}_N, \zeta_N)\}$. The two objective functions, $E(X)$ and $T(X)$, are determined by the chosen locations X . $E(X)$ evaluates the modeling error of a GP model trained on the resulting dataset S , while $T(X)$ evaluates the total time required to visit all locations in X , including both travel and measurement times.

In this multi-objective context, we use the concept of Pareto optimality. A solution X_1 is said to dominate another solution X_2 if it is no worse in all objectives and strictly better in at least one:

$$\begin{aligned} E(X_1) \leq E(X_2) \wedge T(X_1) \leq T(X_2), \\ \text{and } E(X_1) < E(X_2) \vee T(X_1) < T(X_2). \end{aligned} \quad (2)$$

Let the objective vector for any solution X be denoted by $f(X) = (E(X), T(X))$. The set of all non-dominated solutions constitutes the Pareto-optimal set, \mathcal{P} , defined as:

$$\mathcal{P} = \{X \subseteq \mathcal{C} \mid \nexists X' \subseteq \mathcal{C} \text{ s.t. } f(X') \prec f(X)\}. \quad (3)$$

Projecting this set into the objective space yields the Pareto front, F :

$$F = \{f(X) \mid X \in \mathcal{P}\}. \quad (4)$$

The Pareto front F represents the optimal trade-offs between modeling error and investigation time.

In practice, obtaining the true Pareto front F^* is infeasible. Therefore, this research aims to develop an adaptive sampling strategy that constructs a superior approximate Pareto front, generating a set of solutions that dominates those produced by naive methods, thereby shifting the approximate front closer to the origin and towards the true Pareto front.

III. METHODOLOGY

A. Overview

The adaptive sampling method proposed in this study is outlined in Algorithm 1. The procedure begins with an initialization step to prevent boundary over-exploration [23]—a phenomenon where the GP’s high uncertainty at the map edges, far from any measurements, draws the sampling path outwards and away from the primary area of interest.

To mitigate this, we first define a set of m virtual observation points, S_{virtual} , located along the map boundaries:

$$S_{\text{virtual}} = \{s_i = (\mathbf{x}_i, \zeta_0) \mid i = 1, 2, \dots, m\}, \quad (5)$$

where \mathbf{x}_i are the coordinates of the i -th boundary point and ζ_0 is an initial soil stiffness value. These virtual points are temporarily used to anchor the GP model (lines 3-4). While a larger m more effectively prevents boundary selection, excessively large values can suppress necessary exploration near boundaries and increase the GP’s computational cost. The values for m and ζ_0 are therefore determined empirically to ensure favorable performance. To ensure the final map is unbiased by these virtual points, they are excluded from the training set used for final model generation.

Following initialization, the main adaptive sampling loop begins and iterates until a termination condition, such as the maximum number of measurements, is met (lines 5-13). Let S_t be the set of t real measurements collected so far. Each iteration of the loop consists of two main steps:

- 1) **Gaussian Process Modeling:** At each step t , the GP model, \mathcal{GP} , is updated using the set of all measurements gathered so far, S_t , combined with the set of virtual boundary points, S_{virtual} (lines 8-9). This model then provides a full posterior distribution over the entire soil stiffness field, which is defined by its mean vector $\boldsymbol{\mu}$ and covariance matrix Σ . This process quantifies the predicted values and their uncertainties and is detailed in Section III-B.
- 2) **Bayesian Optimization for Point Selection:** Next, Bayesian Optimization (BO) is employed to determine the subsequent sampling location, \mathbf{x}_{t+1} (line 10). The BO framework uses an acquisition function that leverages the GP’s posterior distribution (defined by its mean $\boldsymbol{\mu}$ and covariance Σ) and a traversability score τ to select a point that optimally resolves the trade-offs between exploitation, exploration, and travel time. This selection mechanism is further explained in Section III-C, III-D.

B. Gaussian Processes

A Gaussian Process, denoted as $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, is a non-parametric model completely characterized by its mean function $\mu(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$. We model the environmental phenomenon under investigation as a latent function $g(\mathbf{x})$. Each observation ζ_i is assumed to be the true function value corrupted by Gaussian noise ϵ_i :

$$\zeta_i = g(\mathbf{x}_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_n^2). \quad (6)$$

Algorithm 1 Adaptive Sampling Framework

Require: Number of virtual boundary points m , Initial stiffness value ζ_0 , Robot’s initial position \mathbf{x}_1 , Maximum number of measurements N , Terrain map \mathcal{M} , Maximum Travel radius R , Weighting factor α

Ensure: Final adaptive sample set S_N , Final GP model $\mathcal{GP}_{\text{final}}$

```

// Initialization
1:  $S_0 \leftarrow \emptyset$   $\triangleright$  Initialize the set of real measurements
2:  $\mathbf{x}_{\text{current}} \leftarrow \mathbf{x}_1$ 
3:  $S_{\text{virtual}} \leftarrow \{(\mathbf{x}_i, \zeta_0) \mid i = 1, \dots, m\}$ 
4:  $\mathcal{GP} \leftarrow \text{TRAINGP}(S_{\text{virtual}})$   $\triangleright$  Anchor GP with virtual points
5: for  $t = 1, 2, \dots, N$  do
6:    $\zeta_t \leftarrow \text{MEASURESTIFFNESS}(\mathbf{x}_{\text{current}})$ 
7:    $S_t \leftarrow S_{t-1} \cup \{(\mathbf{x}_{\text{current}}, \zeta_t)\}$ 

// Gaussian Process Modeling
8:    $S_{\text{train}} \leftarrow S_t \cup S_{\text{virtual}}$   $\triangleright$  Combine real and virtual data
9:    $\mathcal{GP} \leftarrow \text{UPDATEGP}(S_{\text{train}})$   $\triangleright$  Retrain GP and re-optimize  $\theta$ 

// Bayesian Optimization
10:   $\mathbf{x}_{t+1} \leftarrow \text{SELECTNEXTLOCATION}(\mathcal{GP}, \mathbf{x}_{\text{current}}, \mathcal{M}, R, \alpha)$ 
11:  Move robot to  $\mathbf{x}_{t+1}$ 
12:   $\mathbf{x}_{\text{current}} \leftarrow \mathbf{x}_{t+1}$ 
13: end for
14:  $\mathcal{GP}_{\text{final}} \leftarrow \text{TRAINGP}(S_N)$   $\triangleright$  Train only on real measurements
15: return  $S_N$  and  $\mathcal{GP}_{\text{final}}$ 

```

This formulation models the environment as a GP, which defines a distribution over functions. Let $S_t = \{(\mathbf{x}_i, \zeta_i)\}_{i=1}^t$ be the set of t observations, with the corresponding inputs and outputs collected into vectors $X_t = [\mathbf{x}_1, \dots, \mathbf{x}_t]^\top$ and $\zeta_t = [\zeta_1, \dots, \zeta_t]^\top$. The learned GP model can be used to predict a Gaussian distribution over $g(\mathbf{x}^*)$ at any new sampling location \mathbf{x}^* . The joint distribution of the observed values ζ_t and the values ζ^* at any set of unobserved (candidate) locations X^* is given by:

$$\begin{pmatrix} \zeta_t \\ \zeta^* \end{pmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_X & K(X_t, X^*) \\ K(X^*, X_t) & K(X^*, X^*) \end{bmatrix}\right). \quad (7)$$

Here, $K(X_t, X^*)$ is the covariance matrix between the training and candidate locations, while $K(X^*, X^*)$ is the covariance matrix of the candidate locations. The term $K_X = K(X_t, X_t) + \sigma_n^2 I$ represents the covariance matrix of the noisy observations. Based on this formulation, the posterior predictive mean and covariance for the GP are expressed as:

$$\boldsymbol{\mu}^* = K(X^*, X_t) K_X^{-1} \zeta_t, \quad (8)$$

$$\Sigma^* = K(X^*, X^*) - K(X^*, X_t) K_X^{-1} K(X_t, X^*). \quad (9)$$

The choice of the covariance function, or kernel, is crucial as it directly influences the predictive performance of the GP. The kernel should encode prior beliefs about the properties of the underlying function, such as its smoothness and scale. In this work, we employ the widely-used Radial Basis Function (RBF) kernel, defined as:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right), \quad (10)$$

where the signal variance σ_f^2 is the prior variance of the function, controlling the magnitude of its fluctuations around

the mean. The length-scale ℓ determines the input distance at which the function's values become significantly uncorrelated, thereby governing the smoothness of the function. These parameters, along with the noise variance, form the hyperparameters $\theta = \{\sigma_f, \ell, \sigma_n\}$.

These hyperparameters are learned from the data by maximizing the log marginal likelihood (LML):

$$\theta^* = \arg \max_{\theta} \text{LML}(\zeta_t, X_t, \theta), \quad (11)$$

with,

$$\text{LML}(\zeta_t, X_t, \theta) = -\frac{1}{2} \zeta_t^\top K_X^{-1} \zeta_t - \frac{1}{2} \log |K_X| - \frac{t}{2} \log 2\pi. \quad (12)$$

As new observations are incorporated into the dataset, these hyperparameters are re-optimized. This process allows the model to adapt to the new data, thereby improving its modeling accuracy.

C. Bayesian Optimization

Bayesian Optimization (BO) is a sample-efficient framework for finding the global optimum of an unknown function. BO consists of two core components: a probabilistic surrogate model to approximate the unknown function, and an acquisition function to guide where to sample next. In this work, we use a GP as the surrogate model for the latent soil stiffness function g .

The acquisition function uses the GP's posterior predictions, specifically the mean $\mu(\mathbf{x})$ and standard deviation $\sigma(\mathbf{x})$, to quantify the utility of sampling at any given point. A popular choice for this is the Upper Confidence Bound (UCB) acquisition function [24], which explicitly balances the trade-off between exploring regions of high uncertainty (exploration) and sampling in regions of high predicted value (exploitation). The UCB is defined as:

$$a_{\text{UCB}}(\mathbf{x}) = \mu(\mathbf{x}) + \kappa_t \sigma(\mathbf{x}). \quad (13)$$

The next sampling point, \mathbf{x}_{t+1} , is chosen by maximizing this function, $\mathbf{x}_{t+1} = \arg \max_{\mathbf{x}} a_{\text{UCB}}(\mathbf{x})$. The parameter κ_t controls the exploration-exploitation trade-off. We adopt a dynamic schedule for κ_t based on the theoretical guarantees presented in [25], empirically scaling the value by a factor of two to tune it for our problem.

However, while UCB provides a principled framework for efficient modeling, its standard form does not account for the physical cost, such as travel time, required to acquire samples.

D. Slope-aware Upper Confidence Bound (SaUCB)

To incorporate the investigation time $T(X)$ from our problem formulation in (1) into the decision-making process, we propose a novel acquisition function that explicitly penalizes high traversal costs. Empirical studies show that traversal time for wheeled rovers is critically dependent on terrain slope [26]. We therefore define a traversal cost metric based on the absolute slope angle, $|\phi(\mathbf{x})|$, between the robot's

current position \mathbf{x}_t and the candidate point \mathbf{x} . The slope angle is computed as:

$$\phi(\mathbf{x}) = \arctan\left(\frac{z(\mathbf{x}) - z(\mathbf{x}_t)}{\|\mathbf{x} - \mathbf{x}_t\|}\right), \quad (14)$$

where $z(\mathbf{x})$ denotes the terrain elevation at \mathbf{x} .

Our proposed Slope-aware Upper Confidence Bound (SaUCB) is a weighted sum of the normalized UCB value and a normalized traversability score, expressed as:

$$a_{\text{SaUCB}}(\mathbf{x}) = (1 - \alpha) \tilde{a}_{\text{UCB}}(\mathbf{x}) + \alpha \tilde{\tau}(\mathbf{x}), \quad (15)$$

where $\alpha \in [0, 1]$ is the weighting factor expressing the slope awareness. The two components are normalized to the range $[0, 1]$ as follows:

$$\tilde{a}_{\text{UCB}}(\mathbf{x}) = \frac{a_{\text{UCB}}(\mathbf{x}) - a_{\text{UCB},\min}}{a_{\text{UCB},\max} - a_{\text{UCB},\min}}, \quad (16)$$

$$\tilde{\tau}(\mathbf{x}) = 1 - \frac{|\phi(\mathbf{x})| - |\phi|_{\min}}{|\phi|_{\max} - |\phi|_{\min}}. \quad (17)$$

Note that the traversability score $\tilde{\tau}(\mathbf{x})$ is defined such that a lower slope results in a higher score. The next sampling location \mathbf{x}_{t+1} is then chosen by maximizing the SaUCB acquisition function:

$$\mathbf{x}_{t+1} = \arg \max_{\mathbf{x} \in \mathbf{A}_t} a_{\text{SaUCB}}(\mathbf{x}). \quad (18)$$

Here, \mathbf{A}_t is the set of feasible candidate points for the next measurement location \mathbf{x}_{t+1} . A candidate point $\mathbf{x} \in \mathbf{A}_t$ satisfies the following three conditions:

- 1) $\|\mathbf{x} - \mathbf{x}_t\| < R$, (maximum travel radius)
- 2) $|\phi(\mathbf{x})| < \phi_{\text{lim}}$, (slope limit)
- 3) $\mathbf{x} \notin \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\}$. (not previously visited)

To ensure feasibility, the path to \mathbf{x} is discretized into segments based on the height map resolution; the slope limit in 2) applies to each segment, while $\phi(\mathbf{x})$ in (14) is the average of their slopes. In the above normalizations (16) and (17), the minimum and maximum values ($a_{\text{UCB},\min}$, $a_{\text{UCB},\max}$, $|\phi|_{\min}$, $|\phi|_{\max}$) are determined over all feasible candidate points $\mathbf{x} \in \mathbf{A}_t$.

IV. SIMULATION EXPERIMENT

In this section, we evaluate the performance of the proposed SaUCB algorithm through simulation experiments. The objective is to demonstrate its effectiveness in balancing exploration, exploitation, and traversal time on realistic lunar terrain models derived from actual mission data.

A. Experimental Setup

The simulation environment consists of a 250 m \times 250 m square map centered at the origin. The robotic survey commences at the center of the map, $\mathbf{x}_1 = (0 \text{ m}, 0 \text{ m})$, and iteratively performs $N = 50$ soil stiffness measurements. The time required for a single measurement is fixed at $T_{\text{measure}} = 1$ minute.

The robot's travel speed v is modeled as a linear function of the absolute slope angle $|\phi|$ between two points, given by:

$$v(\phi) = -\frac{v_{\max} - v_{\min}}{\phi_{\text{lim}}} |\phi| + v_{\max}. \quad (19)$$

Here, v_{\max} and v_{\min} represent the maximum and minimum speeds, and ϕ_{lim} is the maximum traversable slope angle. Based on specifications for future lunar rovers [27], these parameters are set to $v_{\max} = 8.4$ m/min, $v_{\min} = 0.84$ m/min, and $\phi_{\text{lim}} = 15$ degrees.

To mitigate boundary over-exploration, virtual boundary observations with a value of $\zeta_0 = 1.0$ are placed along the map edges. For the 250 m edges, we place 13 points at regular intervals on each side, resulting in a total of $m = 52$ virtual points. The initial hyperparameters for the GP are set to $(\sigma_f, \ell, \sigma_n) = (1.0, 1.0, 1.0)$.

1) *Experimental Maps*: The 79 experimental maps used in this study consist of a height map and a corresponding ground-truth SF map; their generation process is illustrated in Fig. 2. These maps are generated from Site06 (Nobile rim 1), a candidate landing site for the NASA Artemis program [28]. Specifically, we extract 250 m \times 250 m square map sections from a Digital Elevation Model (DEM) of the site, which provides elevation data at a 5 m resolution.

The ground-truth soil stiffness map is synthetically generated based on geomorphological principles of crater maturation. On the Moon, soil stiffness is strongly influenced by regolith depth and its bulk density. It is known that regolith is generated on the upper slopes of craters and subsequently migrates to and accumulates on the lower slopes [29]. Consequently, upper slopes, where the regolith layer is thin and influenced by bedrock, are expected to be stiff. Conversely, lower slopes, where fresh material accumulates, are expected to be soft.

In this study, we adopt the Scaling Factor (SF) as an index of soil stiffness [30]. The ground-truth SF map is generated such that values are continuously distributed to be high on upper slopes and low on lower slopes. The maximum and minimum SF values are set to 4.0 and 1.5, respectively, based on studies of terrestrial soil analogues [31]. For each map, the ROI is assumed to be the area containing the top 15% of SF values, corresponding to the stiffest ground suitable for infrastructure and roadway.

2) *Performance Evaluation*: To assess the trade-off between mission time-efficiency and scientific return, we evaluate the performance of each method using metrics for investigation time and modeling accuracy.

a) *Investigation Time*: The total time is the sum of the time spent traveling between measurement points and the time spent performing the measurements. It is calculated as:

$$\text{Time} = \sum_{t=1}^{N-1} \frac{\|\mathbf{x}_{t+1} - \mathbf{x}_t\|}{v(\phi_t)} + N \cdot T_{\text{measure}}, \quad (20)$$

where the first term is the total travel time, calculated from the distance between consecutive measurement points and the speed $v(\phi_t)$, defined in (19), corresponding to the slope ϕ_t of that path segment. The second term is the total measurement time.

b) *Modeling Accuracy*: We evaluate modeling accuracy from two perspectives: the fidelity of the entire map (global accuracy) and the specific accuracy within the ROI (ROI-specific accuracy).

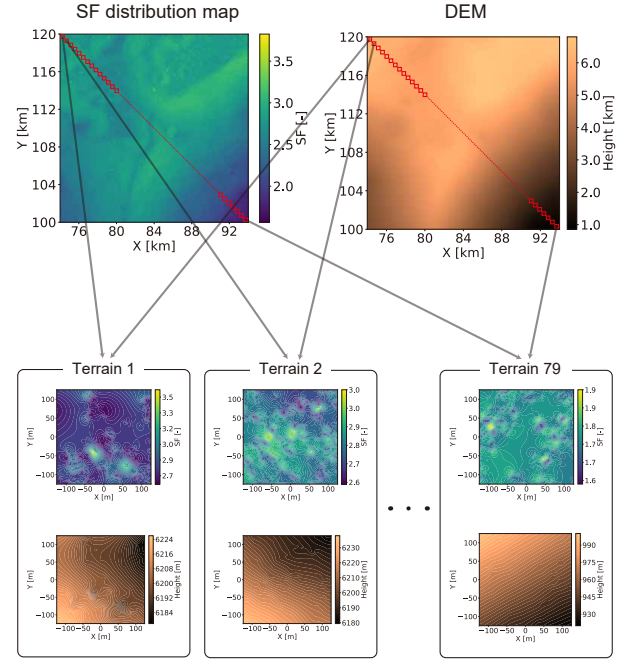


Fig. 2: Generation of the 79 experimental terrains from a diagonal transect of Site06. The large-scale maps (top of the figure) show the synthetic SF distribution and the source DEM. Each resulting terrain patch is composed of a local ground-truth SF map (upper panel) and a height map (lower panel).

Global Accuracy Metrics

- **Root Mean Squared Error (RMSE)** measures the overall fidelity of the predicted map. It is defined as:

$$\text{RMSE} = \sqrt{\frac{1}{M} \sum_{k=1}^M (\hat{\zeta}_k - \zeta_k)^2}, \quad (21)$$

where $\hat{\zeta}_k$ and ζ_k are the predicted and true SF values at grid point k , and M is the total number of grid points in the map.

- **Structural Similarity Index (SSIM)** assesses the perceptual similarity between the predicted and ground-truth maps. It is given by:

$$\text{SSIM}(p, gt) = \frac{(2\mu_p\mu_{gt} + C_1)(2\sigma_{p,gt} + C_2)}{(\mu_p^2 + \mu_{gt}^2 + C_1)(\sigma_p^2 + \sigma_{gt}^2 + C_2)}, \quad (22)$$

where $\mu_p, \mu_{gt}, \sigma_p^2, \sigma_{gt}^2$, and $\sigma_{p,gt}$ are the local means, variances, and cross-covariance for corresponding windows p and gt in the predicted and ground-truth maps, respectively. C_1 and C_2 are stabilization constants.

ROI-Specific Accuracy Metrics

- **ROI-RMSE (R-RMSE)** is a variant of RMSE that evaluates modeling error exclusively within the ROI, defined as:

$$\text{R-RMSE} = \sqrt{\frac{1}{M_{\text{ROI}}} \sum_{k \in \text{ROI}} (\hat{\zeta}_k - \zeta_k)^2}, \quad (23)$$

where M_{ROI} is the number of grid points within ROI.

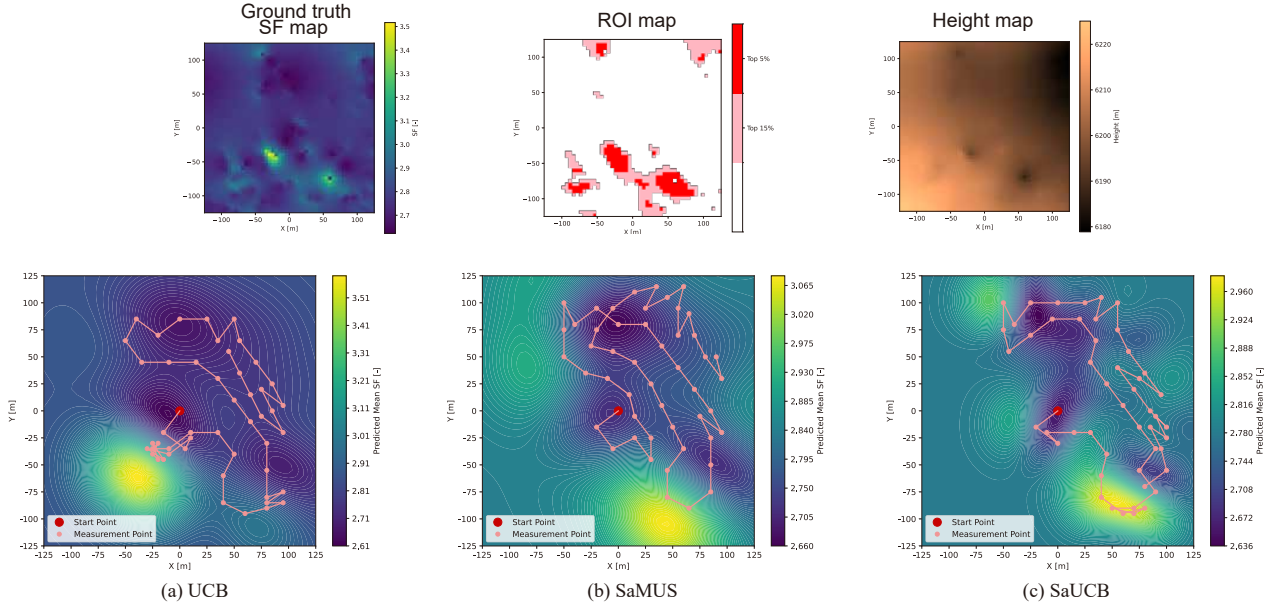


Fig. 3: SF regression results are shown as heatmaps for three different acquisition functions: (a) UCB, (b) SaMUS, and (c) SaUCB, with the sequence of measurement points overlaid. These results were obtained from the northwesternmost terrain of Site06. The ground-truth SF, ROI, and height maps used as input are shown at the top. In the ROI map, the light pink and red areas indicate the top 15% and 5% of SF values, respectively. All results were obtained using $R = 25$ m, with $\alpha = 0.2$ for SaMUS and SaUCB.

- **F1 Score** evaluates the method's ability to correctly classify the ROI. It is the harmonic mean of precision and recall:

$$F1 = \frac{2TP}{2TP + FP + FN}, \quad (24)$$

where TP, FP, and FN are counts of True Positive, False Positive, and False Negative pixels in ROI classification.

3) *Overall Performance Comparison:* Since modeling error and investigation time are conflicting objectives, a comprehensive evaluation requires analyzing the Pareto front. This study employs the Hypervolume indicator (Hv) [32] to quantitatively compare the quality of the approximate Pareto fronts generated by different methods. In a two-dimensional objective space, the Hv is defined as the area of the region dominated by a set of non-dominated solutions on the Pareto front F and bounded by a predefined reference point r :

$$Hv(F) = \Lambda \left(\bigcup_{f \in F} [f, r] \right), \quad (25)$$

where Λ denotes the Lebesgue measure. A larger Hv indicates a superior trade-off, meaning that the corresponding Pareto front is closer to the ideal point $(0, 0)$. Since the scales of the two objectives on each Pareto front (investigation time and the respective modeling accuracy metric) differ significantly, each objective is normalized by its corresponding coordinate of the reference point before the Hv calculation. This ensures that both objectives contribute equally to the metric. Additionally, the maximization metrics, SSIM and F1 Score, are converted to minimization objectives by subtracting each metric's value from 1.

TABLE I
HYPERVOLUME FOR DIFFERENT ACQUISITION FUNCTIONS

Acquisition function	Global		ROI	
	RMSE	SSIM	R-RMSE	F1
UCB	0.379	0.363	0.351	0.333
SaMUS	0.641	0.651	0.578	0.603
SaUCB	0.648	0.653	0.672	0.619

B. Result and Discussions

We compared the performance of our proposed method (SaUCB) against two baselines: the standard UCB and Slope-aware Maximum Uncertainty Sampling (SaMUS). SaMUS, a pure exploration strategy, uses an acquisition function identical to SaUCB's in (15) but lacks the exploitation term, $\mu(\mathbf{x})$. For each of the 79 terrain maps, we performed simulations across a range of parameters: 10 values for the maximum travel radius $R \in \{5, 10, \dots, 50\}$ m for all methods, and 6 values for the weighting parameter $\alpha \in \{0, 0.2, \dots, 1.0\}$ for SaMUS and SaUCB. We compared the performance of the methods by averaging the Hv across the 79 maps. The Hv was calculated for the Pareto fronts formed by investigation time and each modeling accuracy metric.

Table I presents the average Hv values from both global and ROI-specific perspectives. Both SaMUS and our proposed SaUCB outperformed the standard UCB across all metrics. This is because incorporating a slope-aware term into the sampling strategy significantly improves traversability, leading to a major reduction in investigation time that outweighs any minor increase in modeling error. While SaMUS and SaUCB performed similarly on global metrics, SaUCB

demonstrated a significant advantage in the ROI-specific metrics, particularly for R-RMSE. This improvement is attributed to the exploitation term in SaUCB, which effectively directs the robot to conduct more measurements within the ROI. As a result, SaUCB reduced the modeling error within the ROI without a significant penalty to the investigation time. Fig. 3 provides visual evidence of these differing behaviors. The UCB method focused on surveying the ROI but occasionally selected paths with low traversability. SaMUS, conversely, maintained high traversability by avoiding steep slopes but failed to focus on the ROI. In contrast, our proposed SaUCB found a more balanced path, successfully navigating the trade-off between focusing on the ROI and maintaining high traversability.

These results demonstrate that SaUCB effectively optimizes for both low investigation time and low modeling error within the ROI. This indicates that our proposed method is particularly well-suited for targeted environmental modeling tasks on rough terrain with wheeled robots. Furthermore, the proposed method is computationally feasible for onboard processing: the approach described in this paper was validated on an Intel Core i7-1260P CPU (2.10 GHz), where the essential computations per sampling iteration (lines 6-12 in Algorithm 1) required only about 100 ms, which would correspond to approximately one second on a space-hardened.

V. CONCLUSIONS

We proposed an adaptive sampling method, called SaUCB, which integrates a traversability score based on terrain slope into the UCB acquisition function. This approach guides a rover to efficiently map soil stiffness by balancing exploration, exploitation, and traversal time. Simulation experiments on realistic lunar terrain have validated its effectiveness, demonstrating that SaUCB achieves a superior trade-off between investigation time and modeling accuracy within the ROI compared to standard methods.

Future work will extend SaUCB in three directions. First, real-world validation is needed to verify the robustness of the method against practical uncertainties in localization and terrain slope estimation, as well as under realistic vehicle dynamics. Second, the framework should be expanded for multi-robot collaborative exploration to enhance mission efficiency. Finally, extending the acquisition function to co-optimize for other scientific metrics, such as visual data, would broaden its applicability for multi-modal robotic exploration.

REFERENCES

- [1] International Space Exploration Coordination Group (ISECG), Ed., *Global Exploration Roadmap*. ISECG, 2024.
- [2] J. Guo *et al.*, "In-situ evaluation of terrain mechanical parameters and wheel-terrain interactions using wheel-terrain contact mechanics for wheeled planetary rovers," *Mech. Mach. Theory*, vol. 145, p. 103696, 2020.
- [3] G. F. SOWERS, "Geotechnical issues in site selection and project development," *Bulletin of the Association of Engineering Geologists*, vol. 31, no. 2, pp. 209–216, 1994.
- [4] S. J. Dyke *et al.*, "Establishing standards for lunar in situ resource utilization structural materials," *AIAA J.*, vol. 62, pp. 2414–2423, 2024.
- [5] J. Richter *et al.*, "Multi-objective global path planning for lunar exploration with a quadruped robot," in *Proc. Int. Conf. Space Robot. IEEE*, 2024, pp. 48–55.
- [6] H. Inotsume *et al.*, "Robust path planning for slope traversing under uncertainty in slip prediction," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3390–3397, 2020.
- [7] O. Lamarre *et al.*, "Safe mission-level path planning for exploration of lunar shadowed regions by a solar-powered rover," in *Proc. IEEE Aerosp. Conf.*, 2024, pp. 1–14.
- [8] P. Auer, "Using confidence bounds for exploitation-exploration trade-offs," *J. Mach. Learn. Res.*, vol. 3, no. Nov, pp. 397–422, 2002.
- [9] V. Dani *et al.*, "Stochastic linear optimization under bandit feedback," in *Proc. Annu. Conf. Learn. Theory*, no. 101, 2008, pp. 355–366.
- [10] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [11] M. Ghaffari Jadidi *et al.*, "Gaussian processes autonomous mapping and exploration for range-sensing mobile robots," *Auton. Robots*, vol. 42, pp. 273–290, 2018.
- [12] J. Das *et al.*, "Data-driven robotic sampling for marine ecosystem monitoring," *Int. J. Robot. Res.*, vol. 34, no. 12, pp. 1435–1452, 2015.
- [13] R. Ouyang *et al.*, "Multi-robot active sensing of non-stationary gaussian process-based environmental phenomena," in *Proc. Int. Conf. Auton. Agents Multi-Agent Syst.*, 2014.
- [14] M. F. Mysorewala *et al.*, "Multi-scale adaptive sampling with mobile agents for mapping of forest fires," *J. Intell. Robot. Syst.*, vol. 54, pp. 535–565, 2009.
- [15] A. A. Meera *et al.*, "Obstacle-aware adaptive informative path planning for uav-based target search," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 718–724.
- [16] D. Mansfield *et al.*, "Autonomous cooperative mapping of gps-denied cluttered environments using gaussian process regression," in *Proc. IEEE Int. Conf. Control Autom.*, 2024, pp. 546–551.
- [17] A. Munir and R. Parasuraman, "Exploration–exploitation tradeoff in the adaptive information sampling of unknown spatial fields with mobile robots," *Sensors*, vol. 23, no. 23, p. 9600, 2023.
- [18] S. Bai *et al.*, "Information-theoretic exploration with bayesian optimization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1816–1822.
- [19] Y. T. Tan *et al.*, "Gaussian process adaptive sampling using the cross-entropy method for environmental sensing and monitoring," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 6220–6227.
- [20] R. Marchant and F. Ramos, "Bayesian optimisation for intelligent environmental monitoring," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 2242–2249.
- [21] R. Agrawal *et al.*, "Oas-gpubc: On-the-way adaptive sampling using gpubc for bathymetry mapping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024, pp. 1265–1270.
- [22] M. Sutoh *et al.*, "Traveling performance estimation for planetary rovers over slope," in *Proc. IEEE/SICE Int. Symp. Syst. Integr.*, 2011, pp. 884–889.
- [23] E. Siivola *et al.*, "Correcting boundary over-exploration deficiencies in bayesian optimization with virtual derivative sign observations," in *Proc. IEEE Int. Workshop Mach. Learn. Signal Process.*, 2018, pp. 1–6.
- [24] A. Cicirello and F. Giunta, "Machine learning based optimization for interval uncertainty propagation," *Mech. Syst. Signal Process.*, vol. 170, p. 108619, 2022.
- [25] N. Srinivas *et al.*, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [26] M. Sutoh *et al.*, "Traveling performance evaluation of planetary rovers on loose soil," *J. Field Robot.*, vol. 29, no. 4, pp. 648–662, 2012.
- [27] M. Robinson and J. Elliott, "Intrepid Planetary Mission Concept Study Report," Arizona State University and Jet Propulsion Laboratory, California Institute of Technology, Tech. Rep., 2020.
- [28] S. Creech *et al.*, "Artemis: an overview of nasa's activities to return humans to the moon," in *Proc. IEEE Aerosp. Conf.*, 2022, pp. 1–7.
- [29] A. Ikeda *et al.*, "Topographic degradation processes of lunar crater walls inferred from boulder falls," *J. Geophys. Res. Planets*, vol. 127, no. 10, p. e2021JE007176, 2022.
- [30] C. Li *et al.*, "A terradynamics of legged locomotion on granular media," *Science*, vol. 339, no. 6126, pp. 1408–1412, 2013.
- [31] Y. Morita and G. Ishigami, "Development and demonstration of robotic investigation tool for terrain mechanical property," in *Proc. 11th Asia-Pacific Reg. Conf. ISTVS*, 2022.
- [32] A. P. Guerreiro *et al.*, "The hypervolume indicator: Problems and algorithms," *arXiv preprint arXiv:2005.00515*, 2020.