Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Risk-Aware Lunar Rover Path Planning with Edge-Level Slope Risk and a Path-Level Cold-Exposure Budget

Muhammad Fazil

fazil@my.swjtu.edu.cn

University of Peshawar

Akhtar Hussain

University of Peshawar

Research Article

Keywords: lunar rover, path planning, chance constraints, permanently shadowed regions, elevation uncertainty, resource-constrained search

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11497970/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Risk-Aware Lunar Rover Path Planning with Edge-Level Slope Risk and a Path-Level Cold-Exposure Budget

Muhammad Fazil[1*] and Akhtar Hussain[2]

[1*]School of Information Science and Technology, Southwest Jiaotong University, Chengdu, Sichuan, China.
[2]Department of Physics, University of Peshawar, Peshawar, Khyber Pakhtunkhwa, Pakistan.

*Corresponding author(s). E-mail(s): fazil@my.swjtu.edu.cn;

Abstract

Rovers approaching lunar permanently shadowed regions must plan over uncertain elevation maps under thermal limits that apply to the whole traverse rather than to single steps. We study a grid planner, ARTP-Lunar, that converts elevation uncertainty into a closed-form per-edge probability of exceeding a slope limit and enforces a path-level cold-exposure budget through label-setting A* over (cell, cold-bin) states, retaining an admissible Euclidean heuristic. Optional dust and water-ice terms are assumed illumination-based proxies. Parameters were frozen on separate validation maps. On 40 previously unseen synthetic maps, with five comparisons fixed in advance, safe-success rose from 0.72 (fixed k-sigma) and 0.70 (terrain-aware) to 0.90 (Holm-adjusted $p < 0.001$ for all five), and cold-budget violations among successful routes fell from 0.32 to 0.02 relative to the same planner without the budget. Removing the uncertainty handling or the budget erased the gain, but replacing the closed-form probability with a k-sigma indicator changed safety by no measurable amount ($+0.03$, 95% CI $[-0.03, 0.08]$) and lowered route energy by about 325 J. In an exploratory sweep the advantage narrowed at 0.2 m elevation error and vanished at 0.4 m. The dust and ice terms had no practical benefit. All terrain is synthetic and all dust, thermal, and ice models are unvalidated proxies, so the results characterise planner behaviour rather than lunar performance.

Keywords: lunar rover, path planning, chance constraints, permanently shadowed regions, elevation uncertainty, resource-constrained search

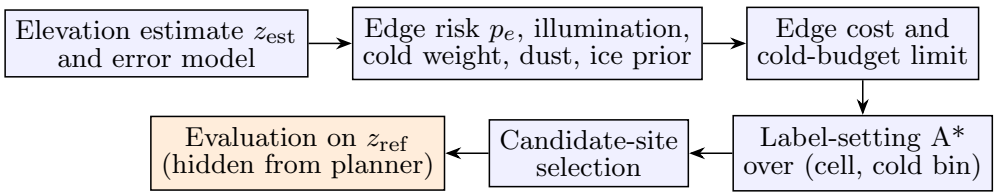


Fig. 1 Planning and evaluation pipeline. The planner uses only z_{est} and its assumed error model; safety is evaluated on the hidden reference terrain z_{ref}.

1 Introduction

Permanently shadowed regions (PSRs) near the lunar poles may preserve water ice and other volatiles, but rover access requires traverses near the sunlit–shadow boundary, where illumination is grazing, terrain relief is difficult to judge, and thermal conditions change over tens of metres [1–5]. Elevation products have finite accuracy, and treating them as exact can lead a planner across ground steeper than it believes [6]. In addition, resources are consumed over the *whole route*: solar-powered rovers in shadow are limited by accumulated energy [7, 8], and cold exposure is a similar cumulative limit, so many cold cells may be crossed provided their accumulated exposure remains within budget. Lunar dust and candidate ice-rich sites introduce further operational considerations [3, 9, 10].

Grid-based rover planning commonly combines terrain costs with A* and uses slope or roughness penalties, often with a fixed k-sigma margin on estimated slope [11–13]. A fixed margin is a deterministic test with one safety factor, so it does not weigh the probability of a violation, or its consequence, against the other costs of a route, and cumulative exposure cannot be enforced exactly by additive edge costs alone because feasibility depends on the route already traversed.

ARTP-Lunar addresses both issues: it converts the elevation-error model into a closed-form per-edge slope-violation probability, enforces a path-level cold-exposure budget by searching over (cell, cold-bin) states, and optionally includes dust and ice terms derived from illumination (Fig. 1). We evaluate it on procedurally generated terrain with parameters frozen on validation maps, followed by a confirmatory run on fresh maps with five pre-specified comparisons and ablations.

Our contributions are: (i) closed-form edge-level slope risk with a per-edge chance constraint and an admissible heuristic; (ii) a path-level cold-exposure budget enforced by label-setting search; (iii) a pre-specified evaluation protocol with frozen parameters, held-out and fresh maps, and map-level cluster statistics; and (iv) an ablation study identifying which components affect safety and route cost. Handling uncertainty and enforcing the budget account for the safety gain, whereas the closed-form probability adds no measurable safety over a k-sigma indicator and the dust and ice terms do not help. The advantage disappears at large elevation error. Because the terrain is synthetic and the lunar layers are unvalidated proxies, the findings describe planner behaviour under stated assumptions, not mission performance on the Moon.

Section 2 reviews related work, Section 3 presents the method, Section 4 describes the protocol, Section 5 reports results, and Section 6 discusses limitations.

2 Related Work

Planetary rover navigation commonly combines terrain assessment with search over a costed grid, using slope, roughness, obstacles, and sometimes energy; these methods generally treat the elevation map as exact [11–15]. Our distance-only and terrain-aware baselines represent this family. Planning under terrain uncertainty has used fixed margins and chance constraints; under a Gaussian belief, a margin is a binary test with one global factor, whereas a violation probability can be carried into the cost and compared across edges with different uncertainty [7, 16–19]. We use the probability in both ways and retain a fixed-margin baseline, without claiming that chance constraints are new to rover planning. A cold-exposure limit is a resource-constrained shortest-path constraint, which ordinary cell-based A* cannot enforce exactly because the best route into a cell depends on resource already consumed; state augmentation over cell and consumed resource is therefore a natural exact-in-belief approach [20, 21]. Lunar polar planning has considered illumination, energy, communication, and thermal constraints, often with time-varying illumination or power models [5]. Dust is a recognised lunar hazard, and ice prospecting uses remote-sensing evidence and terrain accessibility to select sites [3, 9, 10]. Our contribution is the combination of per-edge slope risk with an exact path-level cold budget, together with an evaluation that identifies which components matter; the closest prior work applies chance constraints to mission-level planning in lunar shadowed regions, there for random rover faults and solar energy rather than terrain uncertainty [7], and chance-constrained dynamic programming to risk-aware space exploration [17].

3 Problem Formulation and Method

3.1 Problem setting

The terrain is a square grid with cell size h. The planner searches an 8-connected graph over free cells, does not cut corners, and treats known rocks as obstacles; steep ground is not excluded a priori. Let z_{ref} be the reference elevation and $z_{\mathrm{est}} = z_{\mathrm{ref}} + \varepsilon$ the estimate available to the planner, where ε is zero-mean Gaussian with standard deviation σ and correlation $\rho(r) = \exp(-(r/\ell)^2)$ at separation r cells. The planner receives z_{est} and an assumed model (σ_a, ℓ_a), with $\sigma_a = \kappa\sigma$; z_{ref} is never visible to the planner and is used only for evaluation.

For edge $e = (a, b)$, let $r_e \in \{1, \sqrt{2}\}$ be its length in cells, $d_e = hr_e$ its horizontal length, Δz_e its elevation difference, and $\theta_e = \arctan(\Delta z_e / d_e)$ its slope. With slope limit L, violation is equivalent to $|\Delta z_e| > t_e$, where

$$t_e = d_e \tan L. \tag{1}$$

Violations are judged on z_{ref}. A route is safe if no edge violates the slope limit and, when applicable, its accumulated cold exposure remains within budget.

3.2 Per-edge slope risk

The planner models the elevation difference as Gaussian,

$$\Delta z_e \sim \mathcal{N}\big(\Delta\hat{z}_e, s_e^2\big)\,, \qquad s_e = \sigma_a\sqrt{2(1-\rho(r_e))}, \tag{2}$$

where $\Delta\hat{z}_e$ is the estimated elevation difference. Under this belief the probability that the edge violates the slope limit is

$$p_e = \Pr(|\Delta z_e| > t_e) = \mathrm{Q}\left(\frac{t_e - \Delta\hat{z}_e}{s_e}\right) + \mathrm{Q}\left(\frac{t_e + \Delta\hat{z}_e}{s_e}\right), \tag{3}$$

with Q the Gaussian tail function. If $s_e = 0$, p_e reduces to the indicator of $|\Delta\hat{z}_e| > t_e$. An edge is forbidden when $p_e > \alpha_{\max}$; setting $\alpha_{\max} = 1$ disables the constraint.

Proposition 1 *For a route π with n admitted edges,*

$$\Pr(\exists e \in \pi : |\Delta z_e| > t_e) \le \sum_{e\in\pi} p_e \le n\alpha_{\max}. \tag{4}$$

Proof The first inequality is the union bound; the second follows from $p_e \le \alpha_{\max}$ on every admitted edge. □

The bound is loose for long routes and depends on the correctness of the belief in Eq. (2), which uses a flat prior on elevation differences and ignores terrain smoothness.

3.3 Base edge cost

For an admitted edge entering cell b,

$$c_e = d_e + \lambda_E\frac{E_e}{mg} + C_v(1+\beta H(b))p_e + \lambda_R d_e\frac{R(b)}{R_{\mathrm{ref}}}, \tag{5}$$

in metre-equivalents. Here E_e is a simplified energy estimate,

$$\frac{E_e}{mg} = d_e\max(0, \mu\cos\theta_e + \sin\theta_e)\,, \tag{6}$$

with rolling-resistance coefficient μ, rover mass m, and gravitational acceleration g. The risk term weights p_e by consequence C_v, increased by the factor $(1+\beta H(b))$ where the surroundings are steep: $H(b) \in [0,1]$ is the mean estimated slope over the 3×3 neighbourhood of b divided by L. The term $R(b)$ is the root-mean-square residual of a plane fitted to the 3×3 window of z_{est} around b. The Euclidean heuristic is $h(n) = h\|n-g\|$. Fixed values are $h = 1\,\mathrm{m}$, $L = 20°$, $m = 100\,\mathrm{kg}$, $g = 1.62\,\mathrm{m/s^2}$, $\mu = 0.15$, $\lambda_E = 1$, $\lambda_R = 0.5$ and $R_{\mathrm{ref}} = 0.1\,\mathrm{m}$; C_v, β and $\alpha_{\max}$ are selected on validation maps.

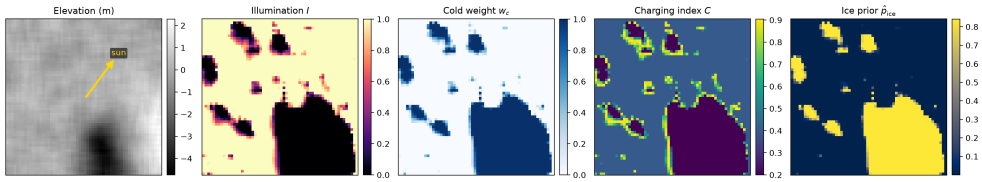


Fig. 2 Lunar layers on one synthetic map, computed from the reference elevation: elevation, illumination index I, cold weight w_c, charging index C, and ice prior. All layers are assumed illumination-based proxies, not validated models.

Proposition 2 *If every term added to d_e in Eq. (5) is non-negative, then $c_e \geq d_e$ and the Euclidean heuristic is admissible and consistent.*

Proof Each added term is non-negative, so $c_e \geq d_e$; by the triangle inequality, the Euclidean distance never exceeds the remaining route cost, giving admissibility and consistency. □

3.4 Lunar proxy layers

Illumination, temperature, dust, and ice are optional, non-negative extensions of Eq. (5). All are assumed proxies computed from a smoothed elevation map: a horizon search gives an illumination index $I(c) \in [0, 1]$; temperature is linear in I; and the cold weight $w_c(c)$ measures how far a cell lies below the operating temperature. The cold time on an edge entering b is

$$\tau_e = \frac{d_e}{v} w_c(b). \tag{7}$$

A qualitative charging index $C(b)$ peaks near the terminator. It enters a slip-and-sinkage proxy D_e, which also grows with slope, and a contamination probability $p_e^c = 1 - \exp(-r_c d_e C(b))$ with $r_c = 0.02/\text{m}$. The planner's ice prior $\hat{p}_{\text{ice}}$ decreases with temperature. The terms are added as

$$c_e \leftarrow c_e + \lambda_D d_e D_e + C_{\text{dust}} p_e^c + \lambda_T d_e w_c(b), \tag{8}$$

with $\lambda_D = 1$, $C_{\text{dust}} = 20$ and $\lambda_T = 1$, so Proposition 2 still holds. Key assumptions are: sun elevation 5°; temperatures 250, 70, and 200 K for sun, shadow, and operation; speed $v = 0.05\,\text{m/s}$; and cold budget $B = 600\,\text{s}$ with $N = 5$ bins. Every value is an unvalidated assumption.

3.5 Path-level cold-exposure budget

The accumulated cold time is $\Gamma(\pi) = \sum_{e \in \pi} \tau_e$, and the thermal limit is $\Gamma(\pi) \leq B$. Because this constraint depends on the whole route, we search over labels (n, g, γ) holding a cell, accumulated cost, and exact cold time. Labels are grouped into states

Algorithm 1 Label-setting A* with a cold-exposure budget

Require: edge costs c_e, cold times τ_e, budget B, bins N, start s, target $g^\star$, heuristic h
Ensure: route from s to $g^\star$ with $\Gamma \le B$, or "no route"

create label $(s, 0, 0)$; set best$[(s, 0)] \leftarrow 0$; push with priority $h(s)$
while queue not empty **do**
  pop label (u, g_u, γ_u) with smallest $g_u + h(u)$
  $k \leftarrow (u, \min(N, \lfloor N\gamma_u / B \rfloor))$
  if k closed **or** best$[k] < g_u$ **then continue**
  end if
  if $u = g^\star$ **then return** route from parent labels
  end if
  close k
  for all free neighbours v with $p_e \le \alpha_{\max}$ **do**
    $\gamma' \leftarrow \gamma_u + \tau_{(u,v)}$
    if $\gamma' > B$ **then continue**
    end if
    $g' \leftarrow g_u + c_{(u,v)}$; $k' \leftarrow (v, \min(N, \lfloor N\gamma' / B \rfloor))$
    if k' open **and** $g' <$ best$[k']$ **then**
      update best$[k']$; append and push label
    end if
  end for
end while
return no route

(n, b) using N bins, $b = \min(N, \lfloor N\gamma/B \rfloor)$; only the lowest-cost label per state is retained, and an edge is expanded only if $\gamma' \le B$ (Fig. 3 and Algorithm 1).

Proposition 3 *Every route returned by Algorithm 1 satisfies $\Gamma(\pi) \le B$ under the planner's belief, and the number of states is at most $(N + 1)$ times the number of free cells.*

Proof Each expansion requires $\gamma' \le B$, so every returned route satisfies the budget; each state pairs a cell with one of $N + 1$ bins. □

Because labels in the same bin may have slightly different exact cold times, optimality is approximate up to this discretisation. Moreover, shadow is judged from z_{est}, so a returned route can still violate the budget on z_{ref}; we report such violations.

3.6 Ice-site selection

Planners with the ice objective choose among candidate sites: the nominal target and one further free cell sampled in the same distance band from the start (geometry only, not filtered by reachability). Each candidate ς is searched with Algorithm 1, and the planner selects

$$\varsigma^\star = \arg\min_{\varsigma} \left[C(\pi_\varsigma) - V_{\text{ice}} \hat{p}_{\text{ice}}(\varsigma) \right], \tag{9}$$

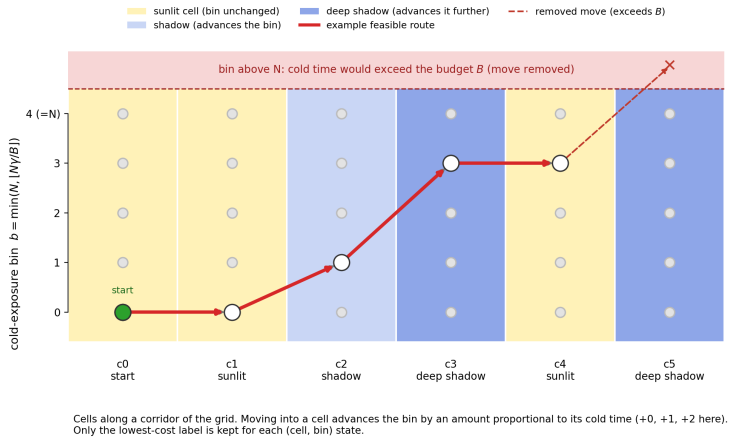


Fig. 3 State space of the budgeted search. Each cell is replicated per cold-exposure bin; shadowed movement advances the bin, moves exceeding B are removed, and only the lowest-cost label is retained per (cell, bin) state.

where $C(\pi_\varsigma)$ is the accumulated edge cost and $V_{\text{ice}} = 60$ converts the estimated ice probability into metre-equivalents. All other planners travel to the nominal target. The true ice probability is the prior evaluated on z_{ref} times an unobserved smooth random modulation, and a route reaches ice if the true probability at its site is at least 0.5. Choosing among targets can also let a planner avoid a difficult nominal target, so a gain from this step is not necessarily due to ice value.

4 Experimental Protocol

Terrain. Each map is a 60×60 grid with $h = 1\,\text{m}$. The reference terrain combines four Gaussian features (hills or depressions, height 2 m to 5 m, width 6 to 12 cells) and smoothed roughness (standard deviation 0.25 m); rocks occupy 2% of cells. In the base condition, elevation error is independent Gaussian with $\sigma = 0.1\,\text{m}$ and the planner's error model is correct ($\kappa = 1$). Start–target pairs are drawn at 50–80% of grid size and are not filtered for solvability, so unsolved pairs count as failures; plans are limited to 5×10^5 expansions and 60 s.

Planners. Distance-only A* uses horizontal distance and no slope limit. The terrain-aware planner uses cost $d_e + 4d_e(|\theta_e|/L)^2 + E_e/(mg)$ and forbids edges with $|\theta_e| > L$ on z_{est}. The fixed k-sigma planner inflates the slope, $\theta_e^{(k)} = |\theta_e| + k\, s_{\theta,e}$ with $s_{\theta,e} = \frac{180}{\pi}\frac{d_e s_e}{d_e^2 + \Delta z_e^2}$, uses $\theta_e^{(k)}$ in the same cost, and forbids edges with $\theta_e^{(k)} > L$; it uses the same s_e as ARTP-Lunar, so differences come from how the belief is used. Ablations of ARTP-Lunar are: *no uncertainty* ($s_e = 0$); *no thermal* (no cold cost and no budget); *no dust*; *no ice* (nominal target only); *k-sigma indicator* (probabilistic risk replaced by the hard test $\theta_e^{(k)} > L$, keeping the budget and all other terms); and three single-term ablations ($\beta = 0$, $\lambda_E = 0$, $\lambda_R = 0$).

Selection and stages. Parameters were selected on 12 validation maps in six conditions over the grid $C_v \in \{25, 100, 400\}$, $\alpha_{\max} \in \{0.1, 0.5, 1\}$, $\beta \in \{0, 2\}$ and

$k \in \{1.0, 1.645, 2.326\}$. A candidate was feasible if its success rate was at least 0.85 and at most 10% of its successful routes were unsafe (slope or cold violation). No candidate was feasible, so the fallback rule (lowest unsafe rate among candidates meeting the success floor) gave $C_v = 100$, $\alpha_{\max} = 0.5$, $\beta = 2$ and $k = 1.0$, which were then frozen. An exploratory test stage used 15 held-out maps and 31 one-factor-at-a-time conditions (elevation error, correlation length, assumed error ratio, roughness, rock fraction, slope limit, pair distance, grid size, $\alpha_{\max}$, sun elevation, cold budget, contamination rate). The confirmatory run used 40 previously unseen maps in the base condition only, with two start–target pairs and three noise draws per map (240 scenarios per planner).

Metrics. Success means the planner returns a route that reaches its target within the limits; failures count against the planner. Safe-success is success with no edge where $|\theta_e| > L$ and $\Gamma \leq B$, both evaluated on z_{ref}. Slope violation and cold violation are the fractions of successful routes that violate the respective limit, dust exposure is $\sum_e d_e D_e$, and ice reached is success at a site with true ice probability at least 0.5 (over all scenarios).

Statistics. The map is the unit of analysis. We report map-level means, 95% percentile cluster-bootstrap confidence intervals from 2,000 resamples, and two-sided paired sign-flip permutation tests with 10,000 permutations; for metrics defined on successful routes, a scenario enters a paired comparison only if both planners succeed. Holm's correction is applied only to the five pre-specified confirmatory comparisons:

(1) ARTP-Lunar versus fixed k-sigma on safe-success (> 0);
(2) ARTP-Lunar versus terrain-aware on safe-success (> 0);
(3) ARTP-Lunar versus no-uncertainty on safe-success (> 0);
(4) ARTP-Lunar versus no-uncertainty on slope violation (< 0);
(5) ARTP-Lunar versus no-thermal on cold violation (< 0).

All other comparisons are exploratory, descriptive, and uncorrected. Calibration compares predicted p_e with empirical violation frequency on 300 randomly sampled edges per map and on edges of chosen routes.

5 Results

5.1 Confirmatory comparison

On the 40 confirmatory maps, ARTP-Lunar achieved safe-success 0.90, versus 0.72 for fixed k-sigma, 0.70 for terrain-aware planning, and 0.33 for distance-only A* (Table 1). Success rates were close to one, so differences arose mainly from unsafe routes. All five pre-specified comparisons moved in the predicted direction and remained significant after Holm correction (Table 2). Removing the uncertainty handling raised the slope-violation rate from 0.07 to 0.27 and lowered safe-success from 0.90 to 0.72, and removing the cold budget lowered safe-success to 0.65. In the exploratory stage, replacing the probabilistic risk with a k-sigma indicator lowered safe-success (0.77 against 0.94), but this did not replicate on the confirmatory maps: the difference in safe-success was $+0.03$ (95% CI $[-0.03, 0.08]$), in slope violation -0.03 ($[-0.08, 0.02]$) and in cold violation $+0.01$ ($[0.00, 0.03]$), and the k-sigma variant used about 325 J less energy per route (95% CI $[282, 375]$). At this noise level the

Table 1 Confirmatory base condition: 40 fresh maps, 240 scenarios per planner. Violations are conditional on success; other rates are over all scenarios.

Planner	Succ.	Safe succ.	Slope viol.	Cold viol.	Dust (m)	Ice reached
A* (distance)	1.00	0.33	0.60	0.14	13.4	0.33
Terrain-aware	1.00	0.70	0.13	0.18	10.7	0.32
Fixed k-sigma	0.97	0.72	0.08	0.20	11.1	0.32
ARTP-Lunar	0.98	**0.90**	0.07	0.02	13.4	0.31
No uncertainty	0.99	0.72	0.27	0.01	11.5	0.40
No thermal	1.00	0.65	0.04	0.32	11.8	0.38
No dust	0.98	0.90	0.06	0.02	13.6	0.32
No ice	0.97	0.84	0.12	0.03	16.0	0.31

Table 2 Pre-specified confirmatory comparisons (ARTP-Lunar minus comparison), 40 maps. Differences include 95% cluster-bootstrap intervals; Holm correction is over the five tests.

Comparison	Metric	Difference [95% CI]	Raw p	Holm p
vs fixed k-sigma	Safe-success	$+0.179$ $[+0.108, +0.258]$	0.0002	0.0006
vs terrain-aware	Safe-success	$+0.196$ $[+0.133, +0.267]$	0.0001	0.0005
vs no uncertainty	Safe-success	$+0.179$ $[+0.092, +0.271]$	0.0004	0.0006
vs no uncertainty	Slope violation	-0.193 $[-0.277, -0.110]$	0.0002	0.0006
vs no thermal	Cold violation	-0.307 $[-0.408, -0.208]$	0.0001	0.0005

closed-form probability therefore adds no measurable safety over a k-sigma indicator and costs energy; the gain comes from handling uncertainty at all and from the cold budget. The adaptive-scale, energy, and roughness terms produced differences within noise and are not claimed as contributions.

5.2 Cold budget, dust, ice, and cost

With the budget, 2% of successful ARTP-Lunar routes exceeded the cold limit on the reference terrain, versus 32% for the same planner without it; mean cold exposure was 248 s with the budget and 433 s without it. The residual violations reflect evaluation on true shadow while the planner estimates shadow from z_{est}. The assumed dust terms produced a statistically clear but practically negligible exposure reduction of 0.28 m on routes with about 13 m exposure, with no change in safe-success. Ice-site selection did not increase the ice-reached rate (difference 0.00, 95% CI $[-0.09, 0.10]$); its small safe-success gain over the no-ice ablation (+0.06, exploratory) is most likely due to access to a second candidate target rather than to ice value. These null or negligible effects hold under the stated proxy assumptions, not for real lunar dust or ice. ARTP-Lunar is also much more expensive: it expanded about 11,100 nodes per plan against about 280 for distance-only A* and 1,600 for fixed k-sigma, and needed about 1.2 s to

Table 3 Effect of true elevation error σ (exploratory stage, 15 maps). Safe-success and slope-violation rate among successes. Success rates for $\sigma = 0$, 0.05, 0.1, 0.2, 0.4 m: ARTP-Lunar 1.00, 1.00, 1.00, 0.99, 0.02; fixed k-sigma 1.00, 1.00, 0.99, 0.13, 0.00; terrain-aware 1.00, 1.00, 1.00, 0.99, 0.91. Distance-only A* ignores the estimate and has safe-success 0.40 at every σ.

	Terrain-aware		Fixed k-sigma		ARTP-Lunar	
σ (m)	Safe	Viol.	Safe	Viol.	Safe	Viol.
0.00	0.80	0.00	0.80	0.00	1.00	0.00
0.05	0.78	0.06	0.79	0.03	1.00	0.00
0.10	0.69	0.20	0.71	0.07	0.94	0.04
0.20	0.48	0.40	0.10	0.07	0.63	0.34
0.40	0.29	0.63	0.00	–	0.01	0.50

2.5 s per plan at 60×60 cells (parallel workers, so wall-clock times are indicative) and about 6 s at 90×90 cells in the exploratory run.

5.3 Elevation error and calibration

Table 3 varies the true elevation error while the assumed model remains correct. ARTP-Lunar was strongest at small error, narrowed its advantage at $\sigma = 0.2$ m, and found a route in only 2% of scenarios at $\sigma = 0.4$ m. This failure is consistent with the fixed per-edge chance constraint: with independent errors, a flat orthogonal 1 m step has $s_e \approx 0.57$ m and threshold $t_e \approx 0.36$ m, giving predicted risk about 0.52, above $\alpha_{\max} = 0.5$; only diagonal steps remain admissible, which cannot reach all grid cells. Calibration showed that the edge-risk model was conservative: on random edges, empirical violation frequency lay below predicted probability throughout, and on chosen-route edges it was at most about 0.11 even at predicted probabilities near 0.4. A plausible, untested explanation is that the flat prior overstates elevation-difference uncertainty by ignoring terrain smoothness.

6 Discussion and Limitations

The confirmatory evidence supports two findings: handling elevation uncertainty and enforcing the cold budget together account for the safe-success gain, since removing either erases it, and the budget removes most cold-budget violations. The closed-form probability itself is not shown to add anything. The k-sigma-indicator variant, which shares the budget and all other terms, matches ARTP-Lunar on safety in the confirmatory run and uses less energy. The gain over the stand-alone fixed-margin planner therefore comes mainly from the budget and from uncertainty handling, and partly from the second candidate target (safe-success 0.84 without ice selection); a fixed-margin planner with only the budget was not evaluated. The dust and ice results are negative or negligible under the implemented proxies, and should not be read as evidence that dust or ice are unimportant in reality.

The principal limitations are as follows. The terrain is synthetic, and the dust, thermal, charging, and ice models are unvalidated illumination-based proxies with static illumination. The cold budget is enforced only under the planner's belief and can be violated on the reference terrain. The fixed per-edge chance constraint with a flat prior becomes overly conservative at large elevation error: ARTP-Lunar fails at 0.4 m, where no planner is acceptable. The baselines are our own implementations of standard rules, not published planners, and parameter selection relied on a fallback rule because no candidate met the feasibility limits. The elevation-error and sweep results come from 15 exploratory maps and were not re-confirmed, and run time was not benchmarked on a dedicated machine. Future work should use real lunar elevation data with product-derived error models, replace the proxy layers with empirically calibrated thermal, dust, and ice models, and test a smoothness-aware or learned belief with an adaptive chance limit and time-dependent illumination.

7 Conclusion

We presented a lunar-rover planner that converts elevation uncertainty into closed-form per-edge slope risk and enforces a path-level cold-exposure budget through label-setting A* over (cell, cold-bin) states. On 40 previously unseen synthetic maps, with parameters frozen beforehand and five comparisons fixed in advance, it raised safe-success from 0.72 and 0.70 to 0.90 and reduced cold-budget violations from 0.32 to 0.02 relative to the same planner without the budget. Removing either the uncertainty handling or the budget erased the gain, whereas replacing the closed-form probability with a k-sigma indicator changed safety by no measurable amount, and the assumed dust and ice terms had no practical benefit. All conclusions apply to synthetic terrain and assumed proxy models, and therefore characterise planner behaviour rather than lunar mission performance.

Acknowledgements.

Declarations

- **Funding:** The authors received no funding for this work.
- **Conflict of interest/Competing interests:** The authors declare that they have no competing interests.
- **Ethics approval and consent to participate:** Not applicable.
- **Consent for publication:** Not applicable.
- **Data availability:** All terrain maps are synthetic and generated procedurally by the code described in Section 4; no external datasets were used.
- **Materials availability:** Not applicable.
- **Code availability:** The code is available from the corresponding author on reasonable request.
- **Author contribution:** M.F. developed the method, implemented the planner, ran the experiments and wrote the manuscript. A.H. contributed to the review and editing of the manuscript.

References

[1] Watson, K., Murray, B.C., Brown, H.: The behavior of volatiles on the lunar surface. Journal of Geophysical Research **66**(9), 3033–3045 (1961)

[2] Paige, D.A., Siegler, M.A., Zhang, J.A., Hayne, P.O., Foote, E.J., *et al.*: Diviner lunar radiometer observations of cold traps in the moon's south polar region. Science **330**(6003), 479–482 (2010) https://doi.org/10.1126/science.1187726

[3] Colaprete, A., Schultz, P., Heldmann, J., *et al.*: Detection of water in the LCROSS ejecta plume. Science **330**(6003), 463–468 (2010) https://doi.org/10.1126/science.1186986

[4] Li, S., Lucey, P.G., Milliken, R.E., *et al.*: Direct evidence of surface exposed water ice in the lunar polar regions. Proceedings of the National Academy of Sciences **115**(36), 8907–8912 (2018)

[5] Mazarico, E., Neumann, G.A., Smith, D.E., Zuber, M.T., Torrence, M.H.: Illumination conditions of the lunar polar regions using LOLA topography. Icarus **211**(2), 1066–1081 (2011)

[6] Smith, D.E., Zuber, M.T., Neumann, G.A., *et al.*: Initial observations from the lunar orbiter laser altimeter (LOLA). Geophysical Research Letters **37**(18), 18204 (2010)

[7] Lamarre, O., Malhotra, S., Kelly, J.: Safe Mission-Level Path Planning for Exploration of Lunar Shadowed Regions by a Solar-Powered Rover. arXiv:2401.08558 (2024)

[8] Lamarre, O., Malhotra, S., Kelly, J.: Recovery Policies for Safe Exploration of Lunar Permanently Shadowed Regions by a Solar-Powered Rover. arXiv:2307.16786 (2023)

[9] Colwell, J.E., Batiste, S., Horányi, M., Robertson, S., Sture, S.: Lunar surface: Dust dynamics and regolith mechanics. Reviews of Geophysics **45**(2), 2006 (2007) https://doi.org/10.1029/2005RG000184

[10] Colaprete, A., *et al.*: The volatiles investigating polar exploration rover (VIPER) mission update. In: 53rd Lunar and Planetary Science Conference (2022). Abstract 2675

[11] Hart, P.E., Nilsson, N.J., Raphael, B.: A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics **4**(2), 100–107 (1968)

[12] Maimone, M., Cheng, Y., Matthies, L.: Two years of visual odometry on the Mars exploration rovers. Journal of Field Robotics **24**(3), 169–186 (2007)

[13] Papadakis, P.: Terrain traversability analysis methods for unmanned ground vehicles: A survey. Engineering Applications of Artificial Intelligence **26**(4), 1373–1385 (2013)

[14] Plonski, P.A., Tokekar, P., Isler, V.: Energy-efficient path planning for solar-powered mobile robots. Journal of Field Robotics **30**(4), 583–601 (2013)

[15] Lamarre, O., Limoyo, O., Marić, F., Kelly, J.: The canadian planetary emulation terrain energy-aware rover navigation dataset. The International Journal of Robotics Research **39**(6) (2020)

[16] Blackmore, L., Ono, M., Williams, B.C.: Chance-constrained optimal path planning with obstacles. IEEE Transactions on Robotics **27**(6), 1080–1094 (2011)

[17] Ono, M., Pavone, M., Kuwata, Y., Balaram, J.: Chance-constrained dynamic programming with application to risk-aware robotic space exploration. Autonomous Robots **39**(4), 555–571 (2015)

[18] Fan, D.D., Otsu, K., Kubo, Y., Dixit, A., Burdick, J., Agha-mohammadi, A.-a.: STEP: Stochastic traversability evaluation and planning for risk-aware off-road navigation. In: Robotics: Science and Systems (RSS) (2021)

[19] Fankhauser, P., Bloesch, M., Hutter, M.: Probabilistic terrain mapping for mobile robots with uncertain localization. IEEE Robotics and Automation Letters **3**(4), 3019–3026 (2018)

[20] Irnich, S., Desaulniers, G.: Shortest path problems with resource constraints. In: Desaulniers, G., Desrosiers, J., Solomon, M.M. (eds.) Column Generation, pp. 33–65. Springer, New York (2005). https://doi.org/10.1007/0-387-25486-2_2

[21] Pugliese, L.D.P., Guerriero, F.: A survey of resource constrained shortest path problems: Exact solution approaches. Networks **62**(3), 183–200 (2013)