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.

Single-Channel Lunar Penetrating Radar for Rover Localization via Revisit-Aware Amplitude-Phase Learning

Xiaohang Qiu
National University of Defense Technology
Yuwei Chen
National University of Defense Technology
Chunyu Ding
Shenzhen University
Pengyu Zhang
National University of Defense Technology
Xiaotao Huang
National University of Defense Technology
Tian Jin
tianjin@nudt.edu.cn

National University of Defense Technology

Article

Keywords: Chang'E-4, single-channel Lunar Penetrating Radar, rover localization, geophysical navigation, amplitude-phase disentanglement, autonomous planetary exploration

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11642499/v1

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

Additional Declarations: Competing interest reported. C.D. serves as an Editorial Board Member of npj Space Exploration and is a Guest Editor of the Collection "Geophysical Exploration Technologies for Space Missions". The manuscript was and will be assessed in line with the journal's standard editorial processes, including its policy on competing interests. C.D. has had no involvement in, and will have no involvement in, the editorial handling or decision-making for this manuscript.

Single-Channel Lunar Penetrating Radar for Rover Localization via Revisit-Aware Amplitude-Phase Learning

Xiaohang Qiu[1†], Yuwei Chen[1†], Chunyu Ding[2], Pengyu Zhang[1], Xiaotao Huang[1], Tian Jin[1*]

[1]College of Electronic Science and Technology, National University of Defense Technology, Changsha, China.
[2]Institute for Advanced Study, Shenzhen University, Shenzhen, China.

*Corresponding author(s). E-mail(s): tianjin@nudt.edu.cn;
†These authors contributed equally to this work.

Abstract

GNSS-denied lunar rover localization relies heavily on surface vision, which degrades under extreme solar angles, featureless regolith, and permanent shadows. Exploiting stable, visually invariant subsurface features via single-channel Lunar Penetrating Radar (LPR) offers a promising auxiliary means, yet repurposing in-situ geophysical probing remains unexplored, because revisit benchmarks are scarce and single-channel echoes are strongly attenuated. Here we establish the first revisit-aware localization benchmark from Chang'E-4 Yutu-2 500 MHz LPR observations, keeping non-revisited training/validation traverses strictly disjoint from the physically revisited scenes used for evaluation. By disentangling single-channel sweeps into energy envelopes and continuous unit-circle phase manifolds with structure-preserving aggregation, our early-fusion representation captures high-contrast scatterers and faint stratigraphic horizons. Across lunar and terrestrial radar datasets, our approach reaches 78.05% top-1 recall on lunar revisits, outperforming raw-echo (70.73%) and the strongest learned baseline (73.17%). Existing single-channel LPR payloads can thus support localization without extra hardware in visually degraded scenarios.

Keywords: Chang'E-4, single-channel Lunar Penetrating Radar, rover localization, geophysical navigation, amplitude-phase disentanglement, autonomous planetary exploration

1 Introduction

Precision localization is a cornerstone for long-duration rover traverses, targeted sampling, and hazard avoidance in deep-space exploration[1–3]. Operating without global navigation satellite systems (GNSS), lunar rovers rely heavily on dead reckoning and visual-inertial odometry; yet optical sensors are degraded by extreme solar incidence angles, high-contrast cast shadows, and sparse regolith textures[2, 4, 5]. These visual vulnerabilities are compounded within permanently shadowed regions (PSRs), where the absence of natural light limits both imaging and the range of usable visual landmarks[6, 7], and around the abundant micro cold traps at the lunar poles[8], as well as across multi-lunation operations, where changing illumination drastically alters surface appearance[2]. Existing lunar localization frameworks primarily leverage crater landmarks[2–4], rover-to-orbital image registration[5], and, as reviewed elsewhere[2, 4], elevation-map and horizon matching to bound odometry drift. While effective at macroscopic scales, these surface-dependent modalities lack fine-grained, illumination-invariant loop-closure capability, motivating a complementary navigation paradigm anchored in an invariant physical domain: the planetary subsurface.

Planetary subsurface radars, routinely deployed as scientific payloads to probe internal geological structures, offer a largely untapped modality for autonomous localization without additional hardware. Constrained by payload volume, mass and power budgets, rovers typically carry compact single-channel profiling radars rather than wide multi-antenna arrays. Unlike optical appearance, electromagnetic echoes from these single-channel payloads reflect intrinsic dielectric permittivity contrasts, regolith stratification, and buried rock distributions that remain invariant to solar illumination, shadow dynamics, and surface weathering. Radio sounding of the Moon began with the Apollo 17 lunar sounder, a three-frequency (5, 15 and 150 MHz) synthetic-aperture radar designed to detect subsurface geologic structures[9], and continued from orbit with the Kaguya Lunar Radar Sounder, which revealed buried regolith layers beneath several hundred meters of nearside mare basalt[10]. Rover-borne sounders then resolved the shallow stratigraphy in situ, first at Chang'E-3[11, 12] and subsequently at Chang'E-4[13]; the same class of instrument now operates on Mars aboard Tianwen-1 and Perseverance[14, 15]. In the Chang'E-4 mission, the Yutu-2 rover carries a dual-frequency Lunar Penetrating Radar (LPR), whose high-frequency channel (CH-2B, 500 MHz) resolves fine-scale regolith architecture within the top few tens of meters[16]. While Yutu-2 LPR observations have yielded landmark discoveries in farside ejecta and lava-flow layering[17, 18], buried crater architecture[19, 20], dielectric inversion and paleoregolith horizons[21, 22], and decimeter-scale regolith heterogeneity[23, 24], as summarized in a recent review[25], prior signal processing efforts have focused almost exclusively on echo quality control, clutter suppression, and geological interpretation[26–28]. Repurposing these in-situ geophysical profiles into an operational map/query localization loop, where a rover recognizes previously traversed subsurface structures at trajectory scenes to eliminate cumulative odometry drift, marks a critical leap toward integrated science-engineering autonomy for future lunar and Martian missions.

On Earth, localizing ground-penetrating radar (LGPR) has evolved from raw trace correlation[29, 30] and standardized multi-weather benchmarking[31] to deep

metric learning that incorporates spatiotemporal priors, sequence continuity, and weather-resilient feature alignment[32–37]. However, directly transferring terrestrial LGPR pipelines to extraterrestrial missions confronts two fundamental barriers. First, unlike structured terrestrial settings with dense engineered features, planned survey routes, and centimeter-level external ground truth, lunar rover traverses are largely exploratory with scarce self-intersections, uneven wheel-terrain coupling, and complex dead-reckoning uncertainties. Naively partitioning consecutive radar windows along a single pass for training and testing introduces severe spatial data leakage, allowing networks to memorize transient antenna-ground coupling artifacts, and fails to reflect genuine pass-to-pass revisit discrepancies. Second, whereas terrestrial localizing radars frequently employ wide multi-channel antenna arrays to capture three-dimensional (3D) cross-track spatial diversity[31], planetary rovers rely on single-channel radar profiling whose data are inherently characterized by narrow two-dimensional (2D) vertical cross-sections, absent cross-track azimuthal resolution, and sparse spatial features. Operating on such low-dimensional single-channel B-scans creates a severe information bottleneck: attenuation is depth-dependent, and at 500 MHz the wavelength in regolith ($\approx$0.32 m for a relative permittivity of 3.5) is comparable to the $\approx$0.20 m median disturbance range (twice the autocorrelation length) inferred for the Chang'E-4 regolith from stochastic-medium modeling[23]; heterogeneities of this size lie beyond the Rayleigh limit[38], so scattering falls in the resonance (Mie) regime. Feature enhancement must therefore simultaneously preserve localized high-contrast scatterers and faint stratigraphic horizons from a single transceiver trace.

From an electromagnetic wave-propagation perspective, radar amplitude and instantaneous phase encode complementary physical properties of the subsurface. The analytic signal envelope measures instantaneous echo energy governed by strong dielectric contrasts and boulder diffractions, whereas instantaneous phase captures the geometric continuity of weak, attenuated stratigraphic interfaces independent of amplitude decay. In geological interpretation of Chang'E-4 LPR profiles, cosine-phase attributes have proven instrumental in tracing faint subsurface reflectors[27], and amplitude-phase decompositions have aided terrestrial hyperbolic target detection[39] and subsurface object classification[40]. Because single-channel radar profiles possess limited per-trace data volume and narrow footprints that are sensitive to rover yaw discrepancies and lateral offsets during revisits, raw waveforms alone leave the relevant structure implicit. Decomposing each trace into an analytic envelope and an instantaneous phase makes these two attributes explicit to the network, which may ease learning from scarce extraterrestrial training samples. Yet, how these amplitude-phase representations interact with neural fusion stages and spatial aggregation heads under genuine lunar revisit conditions remains an open question.

Here we bridge this gap by establishing the first leakage-free extraterrestrial revisit benchmark from Yutu-2 CH-2B LPR observations and introducing a physics-informed amplitude-phase learning framework for rover localization. To break the representation bottleneck of single-channel radar sweeps without adding antenna hardware, we disentangle each scalar radar trace into an analytic energy envelope and continuous unit-circle phase channels $(\sin\phi, \cos\phi)$, expanding single-channel profiles into

a three-channel physics-enhanced tensor. We systematically evaluate early signal-level lifting against mid-level dual-branch fusion across two aggregation mechanisms, depth-preserving Fully Connected (FC) projection and translation-invariant Generalized Mean (GeM) pooling, to examine how subsurface geology relates to the preferred aggregation head. Our main contributions are threefold:

1. **Extraterrestrial Revisit Benchmark:** We establish the first revisit-aware map/query localization benchmark using Chang'E-4 Yutu-2 LPR observations, strictly isolating non-revisited traverses for training/validation from physically revisited trajectory scenes for testing, while demonstrating why conventional single-pass validation protocols fail under pass-to-pass coupling variations.
2. **Physics-Informed Amplitude-Phase Learning:** We decouple and fuse analytic energy envelopes and continuous unit-circle phase manifolds across early and mid-level network stages, achieving a 78.05% top-1 localization recall on real lunar revisits without additional onboard hardware overhead.
3. **Domain-Dependent Aggregation:** Across lunar (Chang'E-4) and terrestrial (CMU) single-channel datasets, we observe a domain-dependent preference: stratified lunar regolith favors depth-preserving FC projection, whereas discrete terrestrial anomalies favor GeM pooling for amplitude-phase inputs.

2 Results

2.1 A revisit-aware benchmark for lunar LPR localization

To evaluate subsurface localization under realistic extraterrestrial operating conditions, we construct the first revisit-aware benchmark from the Chang'E-4 Yutu-2 high-frequency LPR (CH-2B, 500 MHz) data publicly released by NAOC, spanning surface operations through March 2025; the 500 MHz channel and its editing workflow are described elsewhere[16, 26]. Each raw A-scan comprises 2,048 temporal samples acquired at a 0.3125 ns interval[18]. To isolate valid subsurface reflections, we crop each trace to sample indices 90-1664 (≈28-520 ns two-way travel time). Physically, this time gate suppresses early-time antenna direct coupling and surface bounce prior to 28 ns while discarding late-time thermal noise beyond the effective penetration limit, corresponding to an estimated probing depth of ≈2.2-41.6 m under typical lunar regolith permittivity. Along the rover traverse, radar profiles are partitioned into sliding windows of 24 A-scans with a stride of 5 A-scans.

Naive random splitting along a single traverse leaks overlapping windows between training and test. In addition, intermittent slews, station stops and coordinate updates introduce position discontinuities along the traverse. We therefore enforce a strict clean-baseline and revisit-isolation protocol across the 1,759 m fused trajectory (Fig. 1 and Table 1). First, we identify and excise trajectory discontinuity jumps and anomalous station intervals along with bilateral 48-trace ($2W$) spatial protection buffers, retaining 1,473 m of high-integrity clean traverse. Second, we partition all non-revisited segments into spatially disjoint training (5,605 windows) and validation (1,474 windows) sets. Finally, we isolate seven physically revisited trajectory scenes (S1-S7,

comprising 226 test windows in total) where Yutu-2 re-traversed the same subsurface footprint, mostly through out-and-back turnarounds or multi-pass shuttle scans within a single drive. Within each revisit scene, the initial pass populates the reference subsurface map database and the subsequent passes serve as localization queries, strictly prohibiting any window from straddling revisit boundaries (see full traverse architecture in Fig. 1a; individual scene geometries are detailed in Methods). Following spatial map-coverage screening within a 3 m localization radius, retrieval performance is evaluated across all 82 valid later-pass revisit queries against the first-pass reference maps (Table 1). As illustrated by the paired pass-to-pass B-scan profiles in Fig. 1b, although surface-coupling variations and depth attenuation alter raw waveform amplitudes between passes, the analytic energy envelope and continuous cosine phase reliably preserve complementary boulder diffractions and deep stratigraphic horizons.

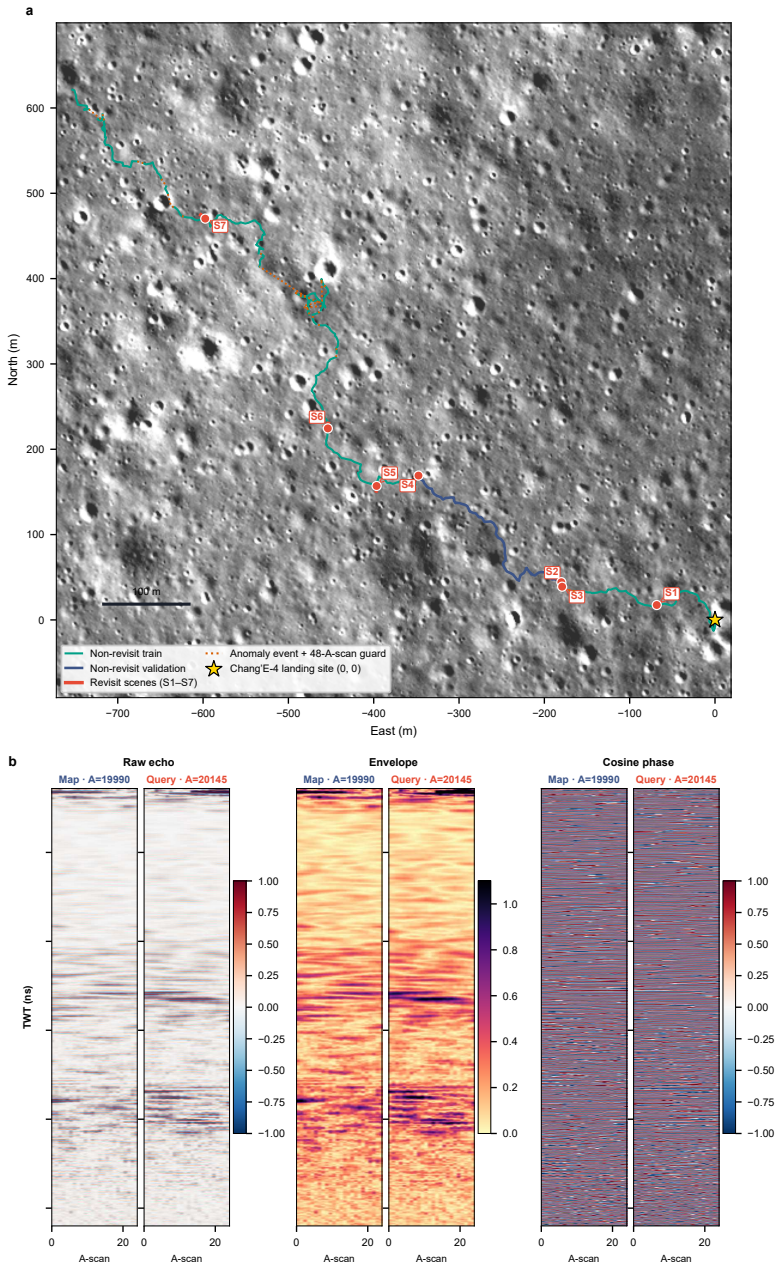


Fig. 1 Full-traverse architecture and pass-to-pass subsurface radar signatures of the Chang'E-4 Yutu-2 LPR revisit benchmark. a, Complete fused rover trajectory in the East-North coordinate frame overlaid on a high-pass filtered LROC NAC orthophoto of the Von Kármán crater floor. Excised kinematic jump anomalies and their bilateral 48-trace ($2W$) protection buffers are highlighted alongside retained clean non-revisited training/validation segments and the seven physically revisited scenes. **b**, Representative pass-to-pass CH-2B subsurface B-scan windows comparing the first-pass reference map and second-pass localization query across raw echoes (X_{raw}), analytic energy envelope ($A(t)$), and continuous cosine phase ($\cos\phi(t)$).

Table 1 Specifications and partitioning protocol of the Chang'E-4 LPR revisit localization benchmark.

Benchmark Parameter	Specification
Payload & channel	Yutu-2 LPR CH-2B (500 MHz center frequency)
Temporal coverage	Mission acquisition through March 2025
Raw A-scan sampling	2,048 samples at 0.3125 ns interval
Valid subsurface echo gate	Indices 90:1664 (≈28-520 ns two-way travel time)
Spatial windowing	24 A-scans per window, 5 A-scan stride
Total fused traverse length	1,759 m
Retained clean-fragment length	1,473 m
Dataset split (Train / Val / Test)	5,605 / 1,474 / 226 windows
Evaluable revisit queries	82 later-pass queries
Revisit test scenes	7 repeated-track scenes

2.2 Synergistic amplitude-phase learning enhances lunar revisit retrieval

To investigate how physics-informed feature enhancement can overcome the information bottleneck of single-channel lunar radar under pass-to-pass variations, we evaluate five input representations within a shared metric backbone: single-channel raw echoes (Raw), single-channel analytic energy envelope (Env), two-channel continuous unit-circle phase manifolds (Phase, encoded via $[\cos\phi, \sin\phi]$), three-channel signal-level early concatenation (Env+Phase Early), and feature-level dual-branch fusion (Env+Phase Dual). Each representation is paired with either depth-preserving Fully Connected (FC) projection or translation-invariant Generalized Mean (GeM) pooling (Table 2 and Fig. 2).

On the Chang'E-4 revisit test set, neither isolated attribute surpasses the raw waveform at top-1 retrieval when paired with FC aggregation: both Env+FC and Phase+FC yield a Test R@1 of 68.29%, trailing Raw+FC (70.73%). However, Phase+FC substantially boosts wider candidate retrieval, reaching 89.02% at R@5 compared to 82.93% for Raw+FC and 84.15% for Env+FC. The two attributes play different roles: the envelope emphasizes high-contrast scatterers, whereas phase preserves faint stratigraphic horizons over a wider neighborhood. Combining both attributes through signal-level early lifting (Env+Phase Early + FC) raises Test R@1 to **78.05%** and Test R@5 to **92.68%**.

By contrast, encoding amplitude and phase in separate branches before fusion (Env+Phase Dual + FC) lowers Test R@1 to 65.85%. Envelope peaks and phase transitions are expected to be tightly coupled at identical depth-time samples. We hypothesize that early pixel-aligned convolution captures this local coherence within the first receptive fields, whereas independent dual-branch downsampling disrupts depth-phase registration before fusion.

We further benchmark our optimal configuration (Env+Phase Early + FC) against representative terrestrial GPR localization and place-recognition architectures under the identical Chang'E-4 revisit protocol (Table 3 and Fig. 2b). Early

Table 2 Ablation study of electromagnetic input representations and spatial aggregation heads on the Chang'E-4 LPR revisit benchmark.

Input Representation	Aggregation Head	Val R@1 (%)	Test R@1 (%)	Test R@5 (%)	Test R@10 (%)
Raw echo (`raw`)	FC	95.79	70.73	82.93	98.78
Envelope (`env`)	FC	96.61	68.29	84.15	98.78
Continuous phase (`phase`)	FC	96.74	68.29	89.02	97.56
Early fusion (`env_phase`)	FC	95.93	**78.05**	**92.68**	**98.78**
Dual-branch (`ap_dual`)	FC	96.88	65.85	85.37	96.34
Raw echo (`raw`)	GeM	93.49	69.51	92.68	96.34
Envelope (`env`)	GeM	95.66	58.54	81.71	97.56
Continuous phase (`phase`)	GeM	97.29	57.32	78.05	92.68
Early fusion (`env_phase`)	GeM	97.01	46.34	76.83	92.68
Dual-branch (`ap_dual`)	GeM	96.20	58.54	84.15	97.56

Table 3 Comparison with GPR baseline methods on the Chang'E-4 LPR revisit benchmark.

Method	Test R@1 (%)	Test R@5 (%)	Test R@10 (%)	Training Setting
Ours (Env+Phase Early + FC)	**78.05**	**92.68**	**98.78**	Trained on CE-4 split
STU-Net[35]	73.17	91.46	98.78	Trained on CE-4 split
CMUNet[42]	71.95	90.24	95.12	Trained on CE-4 split
Baseline (Raw + FC)	70.73	82.93	98.78	Trained on CE-4 split
EDENet[41]	60.98	82.93	96.34	Trained on CE-4 split
SeqOT[43]	58.54	81.71	90.24	Trained on CE-4 split
EDENet-deep*	24.39	59.76	84.15	Trained on CE-4 split
DEC[44]	17.07	52.44	74.39	Non-parametric

*Five-stage variant of the EDENet encoder[41].

fusion gives the best R@1 and R@5 and matches the best R@10, ahead of STU-Net[35] (73.17%) and CMUNet (71.95%). EDENet[41], which extracts hyperbolic orientations with learnable Gabor filters, reaches only 60.98% R@1 here, suggesting that directional filtering is less effective than an explicit amplitude-phase representation for near-horizontal, strongly attenuated regolith. Meanwhile, the LiDAR-derived sequence descriptor SeqOT (58.54%) and the handcrafted energy-curve descriptor DEC (17.07%) degrade sharply under lunar pass-to-pass coupling noise. A five-stage EDENet variant (EDENet-deep) is included only as an encoder ablation; it reaches 24.39% R@1, indicating that deeper Gabor stacks do not compensate for the absence of phase input.

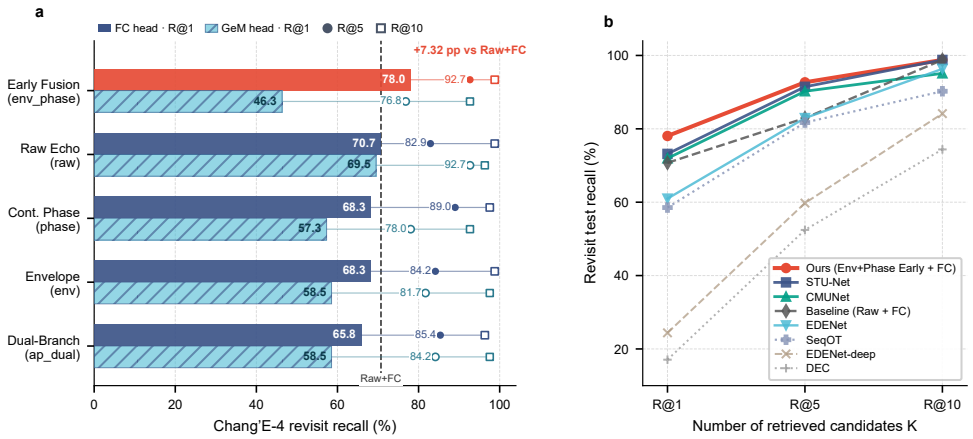


Fig. 2 Retrieval performance of amplitude-phase representations and baseline architectures on the Chang'E-4 revisit benchmark. a, Top-1, Top-5, and Top-10 recall across ten combinations of electromagnetic representations and spatial aggregation heads (bars, R@1, solid for FC and hatched for GeM; circles, R@5, filled for FC and open for GeM; open squares, R@10), highlighting the super-additive gain of early amplitude-phase fusion coupled with depth-preserving FC projection. **b**, Comparative evaluation against baseline methods at $K = 1, 5, 10$.

2.3 Domain-dependent aggregation and the validation-revisit gap

To verify whether our single-channel amplitude-phase enhancement paradigm generalizes across planetary bodies and to probe how subsurface geology interacts with neural aggregation, we evaluate the exact same representations and baseline architectures on the terrestrial CMU-GPR benchmark[42], which similarly acquires 2D single-channel B-scan profiles along survey trajectories, with the two CMU-GPR sites kept disjoint: n2g trains on NSH_b and tests on Gates_g, g2n reverses the roles, so no window from the test site is seen during training (Table 4 and Fig. 3a). On the terrestrial dataset, signal-level early fusion paired with GeM pooling (`Env+Phase Early + GeM`) performs best, with **91.70% mean R@1** and **93.60% mean R@5** (Table 4). EDENet[41] ranks second among all methods at 90.51% mean R@1, ahead of STU-Net[35] (83.60%) and CMUNet (78.56%), yet early fusion still surpasses it without specialized Gabor filter banks.

More importantly, juxtaposing the ablation matrices of lunar farside regolith (Chang'E-4, Table 2) and the terrestrial dataset (CMU, Table 4) shows that, for phase-bearing inputs, the preferred head reverses between domains: early fusion favors FC over GeM on lunar revisits (**78.05% vs. 46.34%**) but GeM over FC on the terrestrial set (**91.70% vs. 82.17%**). Raw echoes perform comparably under both heads on Chang'E-4 (70.73% vs. 69.51%) and favor FC on CMU (89.63% vs. 84.84%). Under FC, early fusion lowers CMU recall to 82.17% but raises Chang'E-4 recall to 78.05%.

This reversal is consistent with the contrasting subsurface morphologies of the two domains. Terrestrial subsurfaces (CMU) consist of indoor concrete floors in a basement and a parking garage, whose localized structural reflectors recur discretely along the track and are observed under small lateral offsets between runs[42]. GeM pooling acts as a spatial soft-maximum selector that isolates these localized hyperbolic signatures and is insensitive to along-track shifts of a reflector within the window; FC instead requires reflectors to fall at consistent grid positions. In contrast, the lunar farside regolith probed by Yutu-2 consists of quasi-continuous, horizontally stratified ejecta deposits and gradual dielectric transitions. Collapsing spatial dimensions via GeM discards the two-dimensional depth-versus-traverse coordinate grid, erasing the stratigraphic geometry essential for distinguishing lunar subsurface scenes. By preserving exact spatiotemporal coordinates, FC projection operates as a depth-aware structural template matcher that locks onto the vertical profile of regolith horizons.

Equally critical for mission engineering, our lunar benchmark exposes a severe generalization illusion when relying on conventional non-revisit validation splits (Fig. 3b). Across all ten configurations in Table 2, single-pass validation R@1 consistently exceeds 93.49%, offering virtually no predictive power for genuine revisit localization. The most dramatic collapse occurs in `Env+Phase Early + GeM`, which achieves a near-perfect **97.01% Validation R@1** on single-pass holdout segments, yet plummets by 50.67 percentage points to **46.34% Test R@1** on physically revisited scenes. Along contiguous single-pass validation segments, positive pairs are drawn from the same pass and partly share A-scans, so validation retrieval approaches near-duplicate matching; pitch/roll attitude and antenna-ground coupling also change little between neighboring windows. During a genuine revisit from a different heading or after a long gap, subtle variations in wheel-regolith contact and lateral trajectory offset alter local wave interference, causing GeM pooling to fail catastrophically. This indicates that single-pass validation can substantially overestimate revisit reliability, suggesting that revisit-aware evaluation should be the default protocol for planetary radar localization.

3 Discussion

3.1 Closing the science-engineering loop for planetary subsurface radars

Planetary subsurface radars such as Chang'E-4 LPR on the Moon[13, 16, 17] and Tianwen-1 RoPeR or Perseverance RIMFAX on Mars[14, 15] have mainly been used as scientific payloads: they acquire raw electromagnetic profiles along rover traverses, which are later inverted offline for geological interpretation and stratigraphy reconstruction[21–23, 45]. Under stringent spacecraft size, weight, and power constraints, dedicated navigation sensors cannot be arbitrarily added to deep-space rovers. Our results indicate that existing scientific radar payloads could serve as localization sensors without additional hardware. This differs from MarsLGPR[46], which regresses along-track displacement using a ground-based GPR at a Mars-analog site; matching against a prior map identifies a location rather than measuring travelled distance. Using a revisit-aware benchmark built from Yutu-2 high-frequency (500 MHz) observations, we find that shallow regolith echoes are repeatable across passes and

Table 4 Cross-domain ablation study and baseline comparison on the terrestrial CMU-GPR benchmark for the two leave-one-site-out settings. Avg R@K is the mean over the two settings.

Method / Representation	Head	n2g R@1 (%)	g2n R@1 (%)	Avg R@1 (%)	Avg R@5 (%)	Avg R@10 (%)
Amplitude-phase ablation within our metric backbone						
Raw echo (`raw`)	GeM	94.84	74.84	84.84	87.86	90.52
Envelope (`env`)	GeM	94.84	79.24	87.04	90.22	91.92
Continuous phase (`phase`)	GeM	92.53	86.74	89.63	93.46	94.79
Early fusion (`env_phase`)	**GeM**	**96.80**	**86.61**	**91.70**	**93.60**	94.53
Dual-branch (`ap_dual`)	GeM	95.55	84.67	90.11	91.94	94.29
Raw echo (`raw`)	FC	96.26	82.99	89.63	92.04	93.47
Envelope (`env`)	FC	92.53	76.20	84.36	86.41	88.50
Continuous phase (`phase`)	FC	93.06	70.31	81.69	85.07	86.55
Early fusion (`env_phase`)	FC	93.77	70.57	82.17	86.93	89.59
Dual-branch (`ap_dual`)	FC	94.84	71.99	83.42	87.30	89.88
Comparison with baseline methods						
EDENet[41]	—	93.25	**87.77**	90.51	93.28	**94.81**
STU-Net[35]	—	92.36	74.84	83.60	85.03	86.51
CMUNet[42]	—	83.84	73.29	78.56	84.65	87.28
SeqOT[43]	—	67.32	43.66	55.49	70.47	74.54
EDENet-deep*	—	63.77	43.53	53.65	67.17	75.91
DEC[44]	—	18.12	8.47	13.30	13.93	13.93

*Five-stage variant of the EDENet encoder[41].

can constrain cumulative odometry drift. This dual use links geophysical probing and spacecraft autonomy: the same in-situ measurement supports both navigation and scientific return.

3.2 Physical mechanisms of electromagnetic representation and spatial aggregation

Beyond empirical recall gains, our ablation and cross-planet evaluations reveal physical mechanisms that shape how neural networks represent subsurface radar waves. First, the super-additive gain from signal-level early amplitude-phase fusion is consistent with the local wave-propagation coherence of lunar regolith reflections. The analytic energy envelope captures dielectric contrast across boulders and layer interfaces, while the unit-circle phase manifold $(\cos\phi, \sin\phi)$ is insensitive to amplitude and therefore preserves structural continuity of deep horizons despite strong exponential attenuation. At subsurface reflectors, instantaneous envelope peaks coincide precisely in two-way travel time with diagnostic phase transitions. Concatenating envelope and continuous phase channels at the input stage lets first-layer convolutional filters jointly match complex-domain waveforms and exploit this depth-phase alignment; independent dual-branch downsampling, by contrast, blurs vertical registration before feature fusion. Expanding a single-channel sweep into envelope and phase channels is a parameter-free analytical transform: it makes amplitude and phase explicit to the network at the cost of one Hilbert transform per trace.

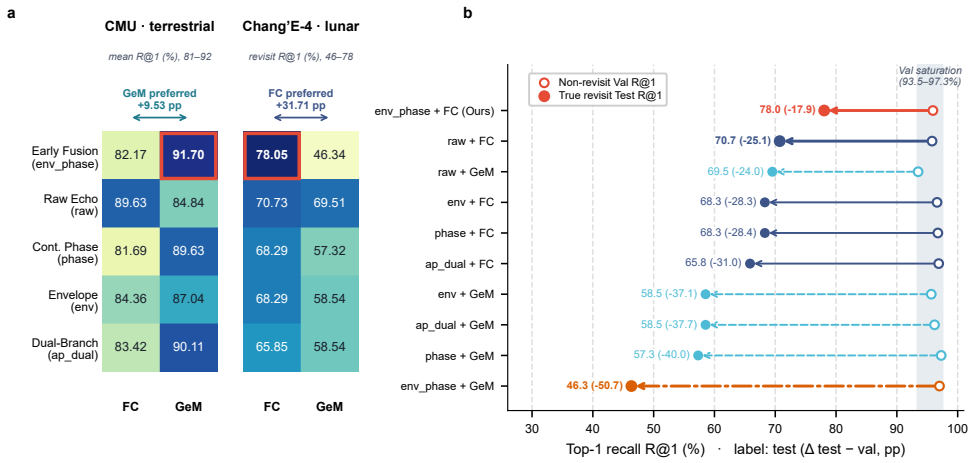


Fig. 3 Domain-dependent aggregation preferences and the single-pass validation illusion. a, Interaction heatmaps of electromagnetic input representations versus spatial aggregation heads on terrestrial radar (CMU, mean R@1 across n2g and g2n) and lunar farside regolith (Chang'E-4, revisit R@1). Each matrix is color-coded on its own scale (CMU 81-92%, Chang'E-4 46-78%) to expose the domain reversal: GeM pooling leads on terrestrial radar, whereas depth-preserving FC leads on stratified lunar regolith; red frames mark the best cell of each domain and the double-headed arrows give the head-to-head margin of the early-fusion row. **b**, Generalization gap between non-revisit validation R@1 (hollow markers) and true revisit test R@1 (filled markers) on Chang'E-4; each row is labeled with its test value and the change in percentage points, and the grey band marks the 93.5-97.3% validation saturation.

Second, the preferred aggregation head reverses between the terrestrial engineered dataset (CMU, GeM pooling) and lunar farside regolith (Chang'E-4, FC projection). This reversal is consistent with a link between subsurface morphology and aggregation choice. GeM pooling is translation-invariant and works well for isolated, high-contrast man-made targets, but it discards the two-dimensional depth-versus-traverse layout needed to characterize quasi-continuous, horizontally stratified lunar ejecta. For GeM, generalization drops by more than 50 percentage points between single-pass validation and genuine revisit testing. This gap points to a methodological pitfall in planetary radar learning: along contiguous single-pass segments, nearly stationary rover attitude and uniform antenna-ground coupling let GeM pooling heads memorize transient, pass-specific phase extrema. Under true pass-to-pass revisits, small heading differences and wheel-regolith contact variations change local wave interference, so single-pass validation metrics are a poor proxy for operational reliability.

3.3 Implications for future lunar polar and Martian missions

Subsurface radar echoes can also support rover localization in upcoming deep-space missions that operate in visually degraded or extreme environments. In planned lunar south pole expeditions and lava-tube explorations[47], rovers will routinely enter PSRs or low-angle illumination zones where cameras and visual-inertial odometry degrade in darkness or under long, dynamic shadows. High-frequency penetrating radar is already

planned as a primary payload in polar missions to prospect for subsurface water-ice volatiles; our framework allows the same radar sweeps to map volatile stratigraphy and provide illumination-independent loop closure. Similarly, Martian rovers equipped with subsurface radars (Tianwen-1 RoPeR and Perseverance RIMFAX)[14, 15] often traverse featureless aeolian dunes and dust-obscured plains where optical landmarks are scarce or temporary. Because shallow dielectric stratification and buried rock distributions are largely unaffected by solar illumination, dust deposition, and atmospheric opacity, a GPR localization module in a multi-sensor navigation stack could provide an additional constraint on dead-reckoning drift. Future multi-frequency architectures that combine high-frequency channels, such as Yutu-2 CH-2B at 500 MHz, with low-frequency channels such as CH-1 at 60 MHz or with full-polarimetric arrays could extend the depth and azimuthal sensitivity of subsurface place recognition.

3.4 Limitations and future research directions

Several methodological constraints of the current benchmark warrant consideration and point toward future research avenues. First, because extraterrestrial rovers operate without GNSS ground truth, the reference trajectory in our Chang'E-4 benchmark is derived from fused dead-reckoning propagation anchored by lander visual observations, orbital high-resolution imagery, and telemetry quality screening. To prevent kinematic discontinuities from corrupting spatial supervision, our clean-baseline protocol excises anomalous station stops and slip intervals, adopting a 3 m localization radius. In operational planetary navigation, this 3 m window-level retrieval serves as a global topological loop-closure trigger to arrest kilometer-scale dead-reckoning drift, after which fine sub-meter registration can be invoked. Extending the representation to accommodate unscreened velocity variations via spatial trace resampling and handling large intersection-angle discrepancies at high-obliquity trajectory crossings remain important next steps.

Second, while continuous unit-circle phase encoding eliminates $\pm\pi$ wrapping discontinuities, instantaneous phase becomes vulnerable to random thermal noise when the echo envelope decays below the receiver noise floor at extreme depths or within homogeneous low-loss strata. Future architectures could weight phase channels by local envelope energy, suppressing deep-time phase noise without sacrificing shallow stratigraphic sensitivity. Platform-induced phase distortion is a further source of pass-to-pass variability that could be estimated explicitly, as autofocus formulations do for near-field SAR[48]. Finally, embedding our lightweight amplitude-phase descriptor into an online, six-degree-of-freedom factor-graph simultaneous localization and mapping back-end, tightly coupled with inertial propagation and wheel odometry, would be a step toward real-time onboard planetary autonomy rather than offline revisit retrieval.

4 Methods

4.1 CH-2B LPR data acquisition and signal preprocessing

We utilize high-frequency Lunar Penetrating Radar (CH-2B, 500 MHz center frequency) observations acquired by the Chang'E-4 Yutu-2 rover and released through the

Lunar and Planetary Data Release System of the National Astronomical Observatories, Chinese Academy of Sciences (NAOC), encompassing surface operations through March 2025. Each raw A-scan trace contains 2,048 samples at $\Delta t = 0.3125$ ns[18]. Because raw planetary radar records contain repeated triggers and stationary acquisition intervals during rover pauses[26], we first remove duplicate traces recorded while the rover was stationary.

Each retained A-scan undergoes a standardized four-stage geophysical preprocessing pipeline: (1) direct-current (DC) bias subtraction to eliminate receiver baseline drift; (2) horizontal background removal to suppress time-invariant antenna ringing and direct air-ground coupling; (3) zero-phase band-pass filtering tailored to the 500 MHz operating bandwidth; and (4) Spreading and Exponential Compensation (SEC) time-varying gain to compensate for spherical wavefront spreading and dielectric attenuation in the lunar regolith. Following preprocessing, each trace is gated to sample indices 90-1664, corresponding to about 28-520 ns two-way travel time and $D = 1,575$ depth samples, to exclude early-time surface bounce and deep thermal noise.

To construct local subsurface sub-maps, gated B-scan profiles are segmented along the traverse using a sliding window of $W = 24$ consecutive A-scans with a stride of $S = 5$ A-scans. Each window $\mathbf{S} \in \mathbb{R}^{1575\times 24}$ is spatially resampled onto a uniform two-dimensional depth-traverse grid of size $H_0 \times W_0 = 200 \times 201$ and min-max normalized to $[0, 1]$. We omit dataset-wide Z-score standardization and per-window mean-variance normalization. This preserves the intrinsic physical amplitude contrast across subsurface strata, and the phase channels retain their circular geometry up to a fixed scaling and offset.

4.2 Dead-reckoning anchor fusion and coordinate construction

Because no lunar GNSS is available, rover trajectory coordinates are reconstructed by fusing wheel-inertial dead-reckoning odometry with periodic external anchors from lander visual tracking, orbital high-resolution imagery, and telemetry calibration. All positional states are projected into a local East-North-Vertical tangent frame. For any A-scan index i belonging to a continuous reference segment initiated at anchor index s, the fused horizontal position vector $\mathbf{p}_i = [E_i, N_i]^\top \in \mathbb{R}^2$ is computed via piecewise anchor-offset propagation:

$$\mathbf{p}_i = \mathbf{p}_s^{\text{ref}} + \mathbf{o}_i - \mathbf{o}_s, \tag{1}$$

where $\mathbf{p}_s^{\text{ref}}$ denotes the calibrated reference anchor coordinate at segment start s, and $\mathbf{o}_i$ represents the cumulative dead-reckoning odometry vector at trace i. Each radar window of W traces is assigned the coordinate of its central A-scan. The horizontal Euclidean distance between radar windows i and j is computed from their ENU coordinates.

Because the coordinates $\mathbf{p}_i$ are fused reference positions rather than centimeter-level RTK-GNSS ground truth, we cross-checked segment-level consistency against independent payload-derived position estimates before benchmark partitioning.

4.3 Clean-baseline protocol and leakage-free revisit split

Multi-lunation rover trajectories contain kinematic discontinuities caused by dead-reckoning drift between anchor updates or by abnormal wheel slip. To keep coordinate jumps from affecting spatial supervision, we apply a clean-baseline excision protocol. We detect all events with large anchor-to-odometry discrepancies and remove both the anomalous intervals and a safety buffer of $2W$ A-scans immediately before and after each event. This $2W$ margin ensures that no retained radar window overlaps or spans a kinematic discontinuity. Across the full 1,759 m fused traverse, removing these anomalous zones leaves 1,473 m of clean trajectory, summed over contiguous clean segments without bridging across gaps.

From the clean traverse, we identify seven revisit scenes (S1-S7; Table 5 and Fig. 4) in which Yutu-2 re-crossed the same subsurface footprint during out-and-back or multi-pass operations. In scenes S1-S5 and S7, the rover traversed a local segment and then reversed or turned around to re-scan the overlapping path; in S6 (Fig. 4f), it completed three consecutive overlapping passes.

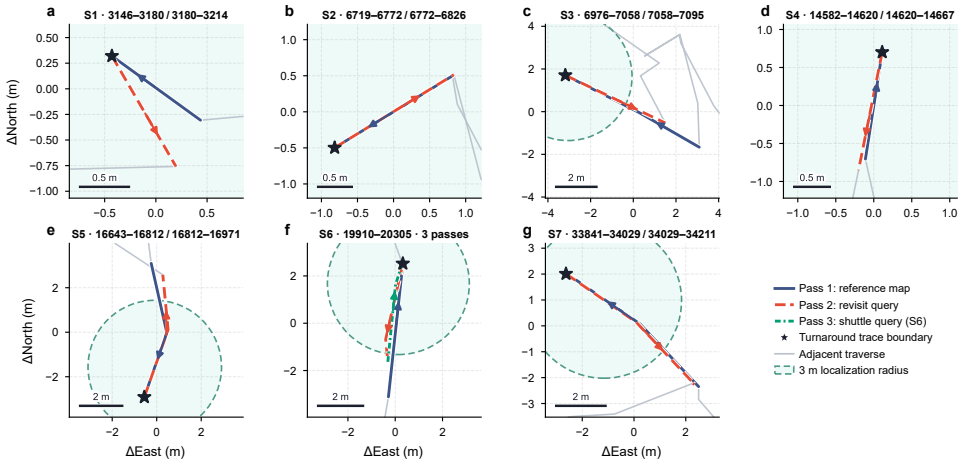


Fig. 4 Map/query partitioning across the seven Chang'E-4 revisit scenes. a-g, Local East-North reconstructions ($\Delta E, \Delta N$ relative to each scene centroid); panel titles give the A-scan ranges of individual passes. Deep blue solid lines show the first-pass reference map, crimson dashed lines the second-pass localization query, and the green dash-dotted line the third pass in **f**. Arrows indicate travel direction, black stars mark the turnaround trace boundary, and grey lines show the adjacent fused traverse. The dashed teal circle marks the 3 m localization radius around the query pass; in **a**, **b**, and **d** the scene lies entirely inside this radius.

Within each revisit scene, the initial pass is designated as the reference subsurface map database and subsequent passes serve as localization queries. To avoid boundary leakage at turnaround pivots, sliding windows are generated within individual pass intervals and do not cross turnaround boundaries. All remaining non-revisited clean

Table 5 Catalog of valid Chang'E-4 Yutu-2 revisit scenes, detailing A-scan index intervals across overlapping passes and maneuver geometries.

Scene ID	A-scan Index Ranges (Map; Query Passes)	Maneuver Type
S1	3146-3180 ; 3180-3214	Two-pass turnaround
S2	6719-6772 ; 6772-6826	Two-pass turnaround
S3	6976-7058 ; 7058-7095	Two-pass turnaround
S4	14582-14620 ; 14620-14667	Two-pass turnaround
S5	16643-16812 ; 16812-16971	Two-pass turnaround
S6	19910-20064 ; 20064-20186 ; 20186-20305	Three-pass shuttle
S7	33841-34029 ; 34029-34211	Two-pass turnaround

fragments are partitioned into spatially disjoint training ($5,605$ windows) and validation ($1,474$ windows) sets, whereas S1-S7 comprise the 226 test windows, yielding 82 evaluable second- and third-pass queries that possess valid first-pass map coverage within the 3 m localization radius.

4.4 Electromagnetic amplitude-phase encoding and network architecture

Yutu-2 CH-2B provides only single-channel 1D traces $s(t)$ along the rover track, without the cross-track resolution of terrestrial multi-channel radar arrays. To separate echo energy from phase continuity and increase the feature dimensionality of single-channel profiles, we transform each real-valued trace $s(t)$ into its complex analytic signal $z(t)$ using the one-dimensional Hilbert transform $\mathcal{H}[\cdot]$ along the depth-time axis:

$$z(t) = s(t) + j\,\mathcal{H}[s(t)] = A(t)\exp(j\phi(t))\,, \tag{2}$$

where the instantaneous energy envelope $A(t)$ and instantaneous phase $\phi(t) \in (-\pi, \pi]$ are given by:

$$A(t) = \sqrt{s^2(t) + \mathcal{H}[s(t)]^2}, \qquad \phi(t) = \operatorname{atan2}(\mathcal{H}[s(t)],\, s(t))\,. \tag{3}$$

Because raw phase angles $\phi(t)$ wrap abruptly at $\pm\pi$, producing high-frequency gradients during spatial convolution, we represent $\phi(t)$ on the unit circle via $[\cos\phi(t),\, \sin\phi(t)]$. Across the 2D single-channel window grid $H_0 \times W_0$, this analytical disentanglement expands scalar observations into multi-attribute tensors, yielding five input configurations X:

1. **Raw echo (`raw`):** single-channel normalized waveform $X_{\text{raw}} \in \mathbb{R}^{1\times H_0\times W_0}$;
2. **Envelope (`env`):** single-channel instantaneous amplitude $X_{\text{env}} = A \in \mathbb{R}^{1\times H_0\times W_0}$;
3. **Continuous phase (`phase`):** two-channel trigonometric phase tensor $X_{\text{phase}} = [\cos\phi, \sin\phi] \in \mathbb{R}^{2\times H_0\times W_0}$;
4. **Signal-level early fusion (`env_phase`):** sample-aligned three-channel concatenation $X_{\text{early}} = [A, \cos\phi, \sin\phi] \in \mathbb{R}^{3\times H_0\times W_0}$, processed jointly by the first convolutional layer;

5. **Feature-level dual-branch fusion (`ap_dual`):** parallel encoding of X_{env} and X_{phase} through independent convolutional stems $F_{\theta,\text{env}}$ and $F_{\theta,\text{phase}}$, followed by mid-level feature-map concatenation before spatial aggregation.

We evaluate these representations within a shared metric learning framework. Given an input tensor X, the convolutional backbone F_θ extracts a feature map $\mathbf{U} = F_\theta(X) \in \mathbb{R}^{C\times H'\times W'}$, which is aggregated by a head $H(\cdot)$ and ℓ_2-normalized to yield the global retrieval descriptor $\mathbf{f}$:

$$\mathbf{f} = \text{normalize}_2(H(F_\theta(X))) = \frac{H(\mathbf{U})}{\|H(\mathbf{U})\|_2}. \tag{4}$$

To examine how aggregation choice interacts with subsurface morphology, we compare two mechanisms for $H(\mathbf{U})$:

- **Fully connected (FC) projection:** flattens $\mathbf{U}$ into a vector of dimension $CH'W'$ and applies a learned linear projection W_{FC}. Because each feature retains a fixed position index along the depth (H') and traverse (W') axes, FC preserves the two-dimensional arrangement of quasi-horizontal lunar regolith layers.
- **Translation-invariant Generalized Mean (GeM) pooling:** collapses the spatial grid (H', W') into a channel-wise descriptor via learnable power-mean pooling over non-negative feature activations $u_{c,h,w}$:

$$\mathbf{g}_c = \left(\frac{1}{H'W'} \sum_{h=1}^{H'} \sum_{w=1}^{W'} u_{c,h,w}^{p_c} \right)^{1/p_c}, \tag{5}$$

  where $p_c > 0$ is a learnable pooling exponent for channel c. GeM is effective at isolating localized hyperbolic anomalies, but it is invariant to their spatial position along the depth and traverse axes.

4.5 Metric training and revisit evaluation protocol

Network parameters are optimized using triplet supervision tailored to the spatial scale of the Yutu-2 radar footprint. During training and validation on the non-revisited Chang'E-4 segments, window pairs (i, j) with horizontal separation $d_{xy}(i, j) \leq 3$ m are defined as positive matches, pairs in the transition buffer 3 m $< d_{xy}(i, j) <$ 5 m are ignored because their labels would be ambiguous, and pairs separated by $d_{xy}(i, j) \geq$ 5 m serve as negative samples. Because training and validation traverses are non-revisited, validation positive pairs necessarily consist of spatially proximate windows along the same pass; in contrast, test evaluation on S1-S7 queries later-pass traverses against the first-pass reference map database.

All comparative baselines, including STU-Net[35], CMUNet[42], SeqOT[43], our EDENet-deep, DEC[44] and the direction-aware multiscale EDENet[41], are evaluated under the identical Chang'E-4 revisit split and spatial thresholds. STU-Net is re-implemented on the same 200×201 windows as all other models, with its 11-antenna axis replicated from the single channel. For EDENet, we keep its original multiscale

learnable Gabor and direction-aware attention encoder, but adapt its input interface and metric supervision to our standardized 200×201 lunar radar windows. We also test a deeper five-stage variant, EDENet-deep. Across all models, optimal checkpoints are selected by maximum top-1 recall (R@1) on the non-revisit validation split; the revisit test scenes are held out from training and hyperparameter selection. At test time, a revisit query is deemed successfully localized at Recall@K (R@K, for $K \in \{1, 5, 10\}$) if at least one of its K nearest reference map neighbors in the ℓ_2 descriptor space resides within a 3 m physical radius of the query coordinate.

For cross-domain evaluations, we select the terrestrial CMU-GPR dataset[42] as it matches the 2D single-channel acquisition geometry of Yutu-2 LPR (in contrast to 3D multi-channel array benchmarks); training and test windows come from different sites, which rules out spatial overlap: for n2g the model is trained on NSH_b and tested on Gates_g, and for g2n the roles are reversed. Within the test site, the frames are pre-partitioned into a reference map set and a query set. We report Recall@1 for each setting and the mean Recall@K ($K \in \{1, 5, 10\}$) over the two.

Data Availability

Chang'E-4 Yutu-2 Lunar Penetrating Radar Level-2B products CE4_GRAS_LPR-2B are available from the Lunar and Planetary Data Release System of the National Astronomical Observatories, Chinese Academy of Sciences at https://moon.bao.ac.cn, and were used to build the lunar revisit benchmark, including the seven revisited test scenes S1-S7. The terrestrial CMU-GPR dataset is available at https://github.com/rpl-cmu/CMU-GPR-Dataset and was used for the cross-domain evaluation under the n2g and g2n settings.

Acknowledgements

We acknowledge the Lunar and Planetary Data Release System of the National Astronomical Observatories, Chinese Academy of Sciences for making the Chang'E-4 LPR Level-2B products publicly available, and the CMU AirLab for releasing the CMU-GPR dataset. This work was supported by the National Natural Science Foundation of China (Grant Nos. 62303475, 12473063 and 42634005), the Innovation Research Foundation of the National University of Defense Technology (Grant No. ZK25-29), and the Natural Science Foundation of Guangdong Province (Grant No. 2024A1515011275).

Author Contributions

X.Q. and Y.C. contributed equally to this work. X.Q. developed the method, implemented the code, performed the experiments, prepared all figures, and wrote the manuscript. Y.C. proposed the research idea, guided the experimental design and the analysis of results, and acquired funding. C.D. contributed to the interpretation of the lunar radar data, acquired funding and revised the manuscript. P.Z. contributed to the network design and the implementation of the baselines. X.H. co-supervised the project and revised the manuscript. T.J. supervised the project and revised the

manuscript. All authors discussed the results, reviewed the manuscript, and approved the submitted version.

Competing Interests

C.D. serves as an Editorial Board Member of npj Space Exploration and is a Guest Editor of the Collection "Geophysical Exploration Technologies for Space Missions". The manuscript was and will be assessed in line with the journal's standard editorial processes, including its policy on competing interests. C.D. has had no involvement in, and will have no involvement in, the editorial handling or decision-making for this manuscript.

References

[1] Farley, K. A. *et al.* Mars 2020 Mission Overview. *Space Science Reviews* **216**, 142 (2020). URL http://link.springer.com/10.1007/s11214-020-00762-y.

[2] Matthies, L. *et al. Lunar Rover Localization Using Craters as Landmarks. 2022 IEEE Aerospace Conference (AERO)*, 1–17 (IEEE, Big Sky, MT, USA, 2022). URL https://ieeexplore.ieee.org/document/9843714/.

[3] Ye, C. *et al.* Map-Guided Coarse-to-Fine Lunar Rover Localization via Multiview Crater Matching. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **18**, 19884–19896 (2025). URL https://ieeexplore.ieee.org/document/11103572/.

[4] Daftry, S. *et al. LunarNav: Crater-based Localization for Long-range Autonomous Lunar Rover Navigation. 2023 IEEE Aerospace Conference*, 1–15 (IEEE, Big Sky, MT, USA, 2023). URL https://ieeexplore.ieee.org/document/10115640/.

[5] Zhao, X., Cui, L., Wei, X., Liu, C. & Yin, J. Lunar Rover Cross-View Localization Through Integration of Rover and Orbital Images. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–14 (2024). URL https://ieeexplore.ieee.org/document/10681493/.

[6] Bickel, V. T., Moseley, B., Lopez-Francos, I. & Shirley, M. Peering into lunar permanently shadowed regions with deep learning. *Nature Communications* **12**, 5607 (2021). URL https://www.nature.com/articles/s41467-021-25882-z.

[7] Cauligi, A. *et al. ShadowNav: Crater-Based Localization for Nighttime and Permanently Shadowed Region Lunar Navigation. 2023 IEEE Aerospace Conference*, 1–12 (IEEE, Big Sky, MT, USA, 2023). URL https://ieeexplore.ieee.org/document/10115745/.

[8] Hayne, P. O., Aharonson, O. & Schörghofer, N. Micro cold traps on the Moon. *Nature Astronomy* **5**, 169–175 (2020). URL https://www.nature.com/articles/s41550-020-1198-9.

[9] Porcello, L. *et al.* The Apollo lunar sounder radar system. *Proceedings of the IEEE* **62**, 769–783 (1974). URL http://ieeexplore.ieee.org/document/1451447/.

[10] Ono, T. *et al.* Lunar Radar Sounder Observations of Subsurface Layers Under the Nearside Maria of the Moon. *Science* **323**, 909–912 (2009). URL https://www.science.org/doi/10.1126/science.1165988.

[11] Xiao, L. *et al.* A young multilayered terrane of the northern Mare Imbrium revealed by Chang'E-3 mission. *Science* **347**, 1226–1229 (2015). URL https://www.science.org/doi/10.1126/science.1259866.

[12] Fa, W., Zhu, M., Liu, T. & Plescia, J. B. Regolith stratigraphy at the Chang'E-3 landing site as seen by lunar penetrating radar. *Geophysical Research Letters* **42** (2015). URL https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2015GL066537.

[13] Li, C. *et al.* The Moon's farside shallow subsurface structure unveiled by Chang'E-4 Lunar Penetrating Radar. *SCIENCE ADVANCES* (2020).

[14] Li, C. *et al.* Layered subsurface in Utopia Basin of Mars revealed by Zhurong rover radar. *Nature* **610**, 308–312 (2022). URL https://www.nature.com/articles/s41586-022-05147-5.

[15] Hamran, S.-E. *et al.* Radar Imager for Mars' Subsurface Experiment—RIMFAX. *Space Science Reviews* **216**, 128 (2020). URL http://link.springer.com/10.1007/s11214-020-00740-4.

[16] Jia, Y. *et al.* The scientific objectives and payloads of Chang'E- 4 mission. *Planetary and Space Science* **162**, 207–215 (2018). URL https://www.sciencedirect.com/science/article/pii/S0032063317300211.

[17] Lai, J. *et al.* First look by the Yutu-2 rover at the deep subsurface structure at the lunar farside. *Nature Communications* **11**, 3426 (2020). URL https://www.nature.com/articles/s41467-020-17262-w.

[18] Cao, H. *et al.* From Schrödinger to Von Kármán: An Intriguing New Geological Structure Revealed by the Chang'e-4 Lunar Penetrating Radar. *Geophysical Research Letters* **50**, e2022GL101413 (2023). URL https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2022GL101413.

[19] Zhou, H. *et al.* Yutu-2 Radar Sounding Evidence of a Buried Crater at Chang'E-4 Landing Site. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–19 (2022). URL https://ieeexplore.ieee.org/document/9467667/.

[20] Zhang, L. *et al.* Subsurface structure of a buried Eratosthenian crater on the far-side of the Moon. *Earth and Planetary Science Letters* **668**, 119542 (2025). URL https://linkinghub.elsevier.com/retrieve/pii/S0012821X25003401.

[21] Chen, R. *et al.* Sub-surface stratification and dielectric permittivity distribution at the Chang'E-4 landing site revealed by the lunar penetrating radar. *Astronomy & Astrophysics* **664**, A35 (2022). URL https://www.aanda.org/10.1051/0004-6361/202142677.

[22] Qiu, X. *et al.* Quantitative Analysis of Subsurface Dielectric Properties by Chang'E-4 Lunar Penetrating Radar Over Lunar Days 24–31. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–13 (2025). URL https://ieeexplore.ieee.org/document/10909448/.

[23] Ding, C., Xiong, S., Li, J., Su, Y. & Huang, S. Yutu-2 radar observation of the lunar regolith heterogeneity at the Chang'E-4 landing site. *Astronomy & Astrophysics* **664**, A43 (2022). URL https://www.aanda.org/10.1051/0004-6361/202142803.

[24] Ding, C. *et al.* Rock Fragments in Shallow Lunar Regolith: Constraints by the Lunar Penetrating Radar Onboard the Chang'E-4 Mission. *Journal of Geophysical Research: Planets* **126**, e2021JE006917 (2021). URL https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2021JE006917.

[25] Lei, Z. & Ding, C. Yutu-2 Radar Observations at the Chang'E-4 Landing Site: The Shallow Geological Structure and Its Dielectric Properties. *Universe* **9**, 461 (2023). URL https://www.mdpi.com/2218-1997/9/11/461.

[26] Roncoroni, G. *et al.* High frequency Lunar Penetrating Radar quality control, editing and processing of Chang'E-4 lunar mission. *Scientific Data* **11**, 118 (2024). URL https://www.nature.com/articles/s41597-024-02963-4.

[27] Roncoroni, G. *et al.* Deep learning driven interpretation of Chang'E-4 Lunar Penetrating Radar. *Icarus* **421**, 116219 (2024). URL https://linkinghub.elsevier.com/retrieve/pii/S0019103524002793.

[28] Soldovieri, F., Gennarelli, G., Su, Y., Ding, C. & Du, W. Microwave tomography for Lunar Penetrating Radar data processing in Chang'e 4 mission. *Scientific Reports* **15**, 5219 (2025). URL https://www.nature.com/articles/s41598-025-89813-4.

[29] Cornick, M., Koechling, J., Stanley, B. & Zhang, B. Localizing Ground Penetrating RADAR: A Step Toward Robust Autonomous Ground Vehicle Localization. *Journal of Field Robotics* **33**, 82–102 (2016). URL https://onlinelibrary.wiley.com/doi/10.1002/rob.21605.

[30] Ort, T., Gilitschenski, I. & Rus, D. Autonomous Navigation in Inclement Weather Based on a Localizing Ground Penetrating Radar. *IEEE Robotics and Automation Letters* **5**, 3267–3274 (2020). URL https://ieeexplore.ieee.org/document/9013076/.

[31] Ort, T., Gilitschenski, I. & Rus, D. GROUNDED: A localizing ground penetrating radar evaluation dataset for learning to localize in inclement weather. *The International Journal of Robotics Research* **42**, 901–916 (2023). URL https://journals.sagepub.com/doi/10.1177/02783649231183460.

[32] Bi, B. *et al.* GPR Image Retrieval and Localization Based on Spatial Prior Constraints. *Journal of Physics: Conference Series* **2887**, 012016 (2024). URL https://iopscience.iop.org/article/10.1088/1742-6596/2887/1/012016.

[33] Zhang, P. *et al.* *Looking Beneath More: A Sequence-based Localizing Ground Penetrating Radar Framework. 2024 IEEE International Conference on Robotics and Automation (ICRA)*, 8515–8521 (IEEE, Yokohama, Japan, 2024). URL https://ieeexplore.ieee.org/document/10610174/.

[34] Zhang, P. *et al.* SeqLGPR: A Sequential Subsurface Feature-Based Framework for Vehicle Place Recognition. *IEEE Transactions on Radar Systems* **3**, 1159–1169 (2025). URL https://ieeexplore.ieee.org/document/11126435/.

[35] Chen, Y. *et al.* Spatial–Temporal U-Net for Localizing Ground-Penetrating Radar. *IEEE Geoscience and Remote Sensing Letters* **22**, 1–5 (2025). URL https://ieeexplore.ieee.org/document/11008670/.

[36] Bi, B. *et al.* Multiweather GPR Image Registration and Localization Based on Adaptive Hyperbolic Receptive Fields. *IEEE Transactions on Geoscience and Remote Sensing* **63**, 1–18 (2025). URL https://ieeexplore.ieee.org/document/11251001/.

[37] Chen, Y. *et al.* Weather-Resilient Localizing Ground-Penetrating Radar via Adaptive Spatio-Temporal Mask Alignment. *Remote Sensing* **17**, 2854 (2025). URL https://www.mdpi.com/2072-4292/17/16/2854.

[38] Grimm, R. E. *et al.* Absorption and scattering in ground-penetrating radar: Analysis of the Bishop Tuff. *Journal of Geophysical Research: Planets* **111**, 2005JE002619 (2006). URL https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2005JE002619.

[39] Zhang, P., Shen, L., Wen, T., Huang, X. & Xin, Q. Vector Phase Symmetry for Stable Hyperbola Detection in Ground-Penetrating Radar Images. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–12 (2022). URL https://ieeexplore.ieee.org/document/9628134/.

[40] Park, B., Kim, J., Lee, J., Kang, M.-S. & An, Y.-K. Underground Object Classification for Urban Roads Using Instantaneous Phase Analysis of Ground-Penetrating Radar (GPR) Data. *Remote Sensing* **10**, 1417 (2018). URL https://www.mdpi.com/2072-4292/10/9/1417.

[41] Zhang, P. *et al.* EDENet: Echo Direction Encoding Network for Place Recognition Based on Ground Penetrating Radar. *Proceedings of the AAAI Conference on Artificial Intelligence* **39**, 10067–10075 (2025). URL https://ojs.aaai.org/index.php/AAAI/article/view/33092.

[42] Baikovitz, A., Sodhi, P., Dille, M. & Kaess, M. *Ground Encoding: Learned Factor Graph-based Models for Localizing Ground Penetrating Radar. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 5476–5483 (IEEE, Prague, Czech Republic, 2021). URL https://ieeexplore.ieee.org/document/9636764/.

[43] Ma, J., Chen, X., Xu, J. & Xiong, G. SeqOT: A Spatial–Temporal Transformer Network for Place Recognition Using Sequential LiDAR Data. *IEEE Transactions on Industrial Electronics* **70**, 8225–8234 (2023). URL https://ieeexplore.ieee.org/document/9994714/.

[44] Li, H., Guo, J. & Song, D. *Subsurface Feature-based Ground Robot/Vehicle Localization Using a Ground Penetrating Radar. 2024 IEEE International Conference on Robotics and Automation (ICRA)*, 1716–1722 (IEEE, Yokohama, Japan, 2024). URL https://ieeexplore.ieee.org/document/10611579/.

[45] Ding, C., Chang, Y., Su, Y., Wang, J. & Xie, M. Rover-Mounted Radar Observation of Discrete Layers Within the Top 4 Meters of Regolith at the Chang'E-3 Landing Site, the Moon. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–16 (2024). URL https://ieeexplore.ieee.org/document/10433208/.

[46] Sheppard, A. & Skinner, K. A. MarsLGPR: Mars Rover Localization With Ground Penetrating Radar. *IEEE Transactions on Field Robotics* **2**, 906–919 (2025). URL https://ieeexplore.ieee.org/document/11225894/.

[47] Qiu, X. & Ding, C. Radar Observation of the Lava Tubes on the Moon and Mars. *Remote Sensing* **15**, 2850 (2023). URL https://www.mdpi.com/2072-4292/15/11/2850. Number: 11.

[48] Song, S., Sun, G., Dai, Y., Sun, S. & Jin, T. JSAF: Joint Structured Priors and Autofocus Framework for Near-Field Compressive Synthetic Aperture Radar Microwave Imaging. *IEEE Transactions on Microwave Theory and Techniques* **74**, 7488–7504 (2026). URL https://ieeexplore.ieee.org/document/11575246/.