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.

Adaptive spatial partitioning under the GeoSOT-3D global grid for efficient low-altitude environment modelling

Qiurong Wang
Xi'an Physical Education University
Zhiyong Liu
Xianyang Polytechnic Institute
Ningshe Zhao
zhaoningshe@gmail.com

Xi'an University

Article

Keywords: GeoSOT-3D, discrete global grid system, adaptive mesh refinement, low-altitude airspace, semantic-driven discretisation, path planning

Posted Date: May 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9453228/v1

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

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Scientific Reports on August 20th, 2026. See the published version at https://doi.org/10.1038/s41598-026-67401-4.

Abstract

Allocating finite representational capacity across spatial domains where geometric complexity and semantic importance vary independently is a recurring problem in spatial modelling. We address this problem in the context of low-altitude airspace, where representational fidelity must trade off against computational efficiency. Building on the GeoSOT-3D global grid encoding system, we propose TSDP (Terrain-Semantic Dual-driven Partitioning), an adaptive partitioning framework that fuses terrain roughness and semantic importance into a multi-criteria weighted decision function. The current implementation operates as a horizontal-plane quadtree subdivision with elevation as a voxel attribute, complemented by the TS-AdaptiveSubdivide algorithm that formalises adaptive mesh generation as spatial-complexity minimisation under fidelity constraints. Experiments on a synthetic scene and a real-world Xi'an district show that TSDP achieves a favourable accuracy–efficiency trade-off: 60% voxel compression with semantic accuracy (F1 = 0.918) comparable to a uniform grid (F1 = 0.933) and competitive with single-driver baselines that excel on either accuracy or efficiency alone but not both. Path-planning time decreases by approximately 40% while mean safety distance to obstacles increases roughly threefold. More broadly, the results indicate that spatial discretisation efficiency in such environments is governed jointly by geometric and semantic factors rather than by either alone, supporting the design of digital-twin systems for low-altitude airspace and urban air mobility management.

Introduction

With the rapid growth of the low-altitude economy[1], emerging applications such as Urban Air Mobility (UAM), unmanned-aerial-vehicle logistics, and emergency response have placed increasingly urgent demands on the digitisation of low-altitude airspace (generally below 1,000 m)[2,3]. China's revised Civil Aviation Law (effective 2026) formally recognises low-altitude transportation, and low-altitude airspace has evolved from a mere aviation transition zone into a critical carrier of three-dimensional intelligent transport networks, requiring both high-precision geometric modelling and fast dynamic computation. At a more fundamental level, this challenge reflects a recurring problem in spatial modelling: how to allocate finite representational capacity across regions whose geometric complexity and semantic importance are not spatially aligned.

The low-altitude environment exhibits pronounced spatial complexity: terrain undulations form continuous elevation gradient fields; building clusters constitute discrete obstacle sets; and meteorological disturbances, dynamic no-fly zones, and air corridors introduce spatio-temporal heterogeneity[3,4]. This multi-scale, multi-modal interplay challenges conventional 3D geographic information methods.

Discrete Global Grid Systems (DGGS) provide a principled mathematical framework for representing geospatial data on a globally consistent reference frame. The concept has been formalised in the Open Geospatial Consortium DGGS abstract specification[8], which defines hierarchical spatial partitioning

schemes with globally unique cell indexing and multi-resolution support[11,12]. The hierarchical coordinate system of Dutton[12] and the geodesic DGGS framework of Sahr et al.[13] established the formal basis for recursive subdivision of the spheroid into uniquely indexed cells, while hexagonal DGGS[14] further improved isotropy and neighbour uniformity. Production-scale open-source implementations of these principles—Uber's H3 hexagonal grid system[9] and Google's S2 geometry library[10]—have demonstrated the practical scalability of DGGS in large-scale geospatial computing, even though their detailed designs are documented primarily as software rather than as peer-reviewed publications. Among these systems, GeoSOT and its 3D extension GeoSOT-3D[4–7] retain the regular quadtree structure favoured by raster GIS pipelines while inheriting the global integer encoding common to all DGGS approaches; voxels (treated as columnar pseudo-3D cells in a 2.5D quadtree representation, where horizontal subdivision is adaptive and elevation is stored as an attribute) obtain globally unique, multi-scale codes that simplify cross-platform spatial-data exchange (concept illustrated in Supplementary Fig. S1). Compared with hexagonal DGGS such as H3, GeoSOT-3D trades isotropy for integer encoding and raster compatibility, an attractive property for the regulatory and audit requirements of airspace management. However, the standard GeoSOT-3D subdivision is static and isotropic[5], creating a contradiction common to all uniform DGGS schemes: excessive resolution in flat areas wastes resources, while fixed resolution in complex areas produces geometric distortion and semantic ambiguity.

Adaptive Mesh Refinement (AMR), introduced by Berger and Oliger[15] and later generalised by Berger and Colella[16], provides a systematic framework for dynamically allocating computational resolution based on local solution features. Subsequent developments, including adaptively refined Cartesian solvers[17], tree-based fluid solvers[18] and octree-based fluid simulation[19], have demonstrated the effectiveness of AMR in handling multi-scale phenomena while maintaining computational efficiency. These approaches typically rely on local error estimators to guide refinement decisions, allowing resolution to concentrate in regions of high gradient or complexity. Geometry-adaptive refinement on irregular domains[20] and multi-resolution terrain models[24] have transferred these ideas to graphics and surface modelling. Yet AMR in the geosciences has remained largely tied to continuous fields and rarely incorporates discrete semantic priors. This integration gap is largely orthogonal to the comprehensive treatment of multidimensional indexing structures by Samet[21], the R-tree dynamic spatial index of Guttman[22], and the foundational quadtree formulation of Finkel and Bentley[23], which address data-structure efficiency rather than the fusion of geometric and semantic refinement criteria pursued here.

Despite substantial progress in both DGGS and adaptive mesh refinement, a unified framework for jointly allocating spatial resolution based on both geometric variation and semantic importance remains limited. Three specific gaps motivate this work: (1) the integration of a globally unified encoding frame with locally adaptive representation has received relatively little attention; (2) existing adaptive methods focus predominantly on continuous geometric fields, with discrete semantic objects under-represented as refinement criteria[15–20]; and (3) lightweight formulations that jointly target accuracy and efficiency are scarce. Existing low-altitude modelling approaches—oblique-photography meshes, point-cloud voxelisation, and standardised semantic 3D city models such as CityGML[25] with associated levels of detail[26,27]—each address subsets of these concerns, but to our knowledge a model that combines DGGS

encoding with semantically-driven AMR has not been reported (a comparative summary is provided as Supplementary Table S1).

This work addresses the general problem of adaptive spatial discretisation when geometric variation and semantic importance are heterogeneously distributed and only weakly correlated. We instantiate this problem on the GeoSOT-3D global grid for low-altitude airspace and contribute: (1) a dual-driven metric system fusing terrain geometry with semantic features in a constrained-optimisation framework, building on AMR error-estimator design[15,18] and DGGS multiscale indexing[8,13]; (2) the TS-AdaptiveSubdivide algorithm for efficient recursive subdivision and conditional merging (Fig. 1); (3) systematic experiments under synthetic and real-world conditions with full ablation and significance testing.

Scope. This paper focuses on horizontal-plane quadtree partitioning under the GeoSOT-3D encoding system; elevation is incorporated as a voxel attribute through the terrain-roughness term in the decision function. A full octree extension to volumetric three-dimensional partitioning, comparable to the octree-based fluid simulations of Losasso et al.[19], is deferred to future work.

Methods

Problem formalisation

Let the study area Ω be a closed rectangular region in (longitude, latitude) coordinates aligned with the GeoSOT-3D global grid. An initial uniform subdivision at level L_0 yields N base voxels $G_0 = \{V_i\}$. A *dynamic subdivision grid* G* is the multi-scale voxel set obtained by finite recursive subdivisions or merges, satisfying three conditions: (i) every terminal voxel $V \in G^*$ has a level $l \in [L_0, L_{max}]$; (ii) the closures of any two terminal voxels intersect at most on a shared face or edge (interior-disjointness); and (iii) the union of the closed terminal voxels exactly covers Ω. Boundary alignment with Ω is exact because Ω is chosen along GeoSOT-3D cell boundaries; no clipped or partial voxels arise.

Hanging nodes — vertices of a fine voxel that lie on the edge of a coarser neighbour — are permitted in G* without imposing a 2:1 level-balance constraint, since the path-planning layer (defined later) operates on voxel centroids via a KD-Tree adjacency graph rather than on shared faces; this means that level differences across neighbours do not require ghost cells or interpolation operators of the kind used in finite-volume AMR. The cost of permitting unbalanced refinement is purely combinatorial (more terminal voxels than a balanced tree would need); the benefit is a strictly smaller |G*| in scenes where features have sharp scale separation, such as a tall building adjacent to flat ground.

Constrained optimisation

We aim to minimise |G*| under fidelity constraints — a geometric bound $|R(V) - \hat{R}(V^*)| \le \varepsilon_T$ (the per-voxel terrain reconstruction error must not exceed ε_T) and a semantic bound stating that no terminal voxel

should fragment a single semantic object below a coverage tolerance ε_S. Because both ε_T and ε_S are problem-specific abstractions whose direct enforcement requires global error estimation at every step, we follow the local error-estimator pattern established in classical AMR[15] and replace the constrained minimisation with a per-voxel greedy decision function that acts as a tunable surrogate:

$$\Phi(V) = \alpha \cdot R(V) / R_{max} + \beta \cdot S(V) / S_{max} + \gamma \cdot (1 - l / L_{max}) \quad (1)$$

where $\alpha + \beta + \gamma = 1$, and R_{max}, S_{max} are 95th-percentile values across G_0 used to normalise the heterogeneous metrics onto a common scale. The depth-penalty term $\gamma \cdot (1 - l/L_{max})$ is an empirical regulariser, not derived from a fidelity bound: it is included to bias the recursion toward earlier termination at deeper levels, providing a bounded-depth guarantee (no voxel can recurse beyond L_{max} regardless of R(V) and S(V)) and reducing the risk of pathological over-refinement when the local terrain or semantic signal contains noise. The setting $\gamma = 0.1$ means depth contributes at most 0.1 to $\Phi(V)$, making it a soft secondary criterion behind R(V) and S(V); ablating γ to 0 produces qualitatively similar grids in our synthetic experiments but increases voxel count in the 5–10% range. Subdivision occurs when $\Phi(V) > \tau$; the threshold τ jointly controls the effective ε_T and ε_S in (1), and its sensitivity is characterised in the Parameter sensitivity subsection.

Approximation regime. We do not claim that the per-voxel greedy minimisation in (1) achieves a provable approximation ratio relative to the constrained optimum $|G^*|^*$. The substitution is justified pragmatically: (i) the constraints ε_T, ε_S are themselves problem-specific abstractions whose direct enforcement would require global error estimation at every refinement step, eliminating the locality that makes adaptive schemes scalable; (ii) the per-voxel form of (1) is consistent in spirit with the local-error-estimator pattern of classical AMR[15],[16] and tree-based fluid solvers[18],[19], where similar locality assumptions are accepted in practice despite the absence of closed-form approximation guarantees on irregular domains; and (iii) the resulting empirical behaviour is characterised by the parameter sensitivity scan (Fig. 3) and the multi-baseline ablation (Tables 2–3), which together quantify the operational envelope of the approximation rather than its worst-case bound. We discuss the implications of this design choice openly in the Engineering boundaries paragraph of the Discussion.

Feature metrics

Terrain roughness:

$$R(V) = w_1 \cdot \Delta H/D + w_2 \cdot ||\nabla H|| + w_3 \cdot |\kappa| \quad (2)$$

with $w_1 = 0.7$, $w_2 = 0.2$, $w_3 = 0.1$, combining the normalised elevation range over voxel diagonal D, the least-squares gradient magnitude, and the absolute mean curvature obtained from a quadric-surface fit.

Semantic importance:

$$S(V) = \max_k (O_k \cdot W_k) + \lambda \cdot \mathrm{sqrt}[\, (1/M) \cdot \Sigma_k (O_k \cdot (W_k - \bar{W}))^2 \,] \quad (3)$$

where the operator $\max_k$ denotes the maximum over the M semantic classes (i.e., the dominant occupied class contributes its full weight), sqrt[·] denotes the positive square root, $O_k \in \{0, 1\}$ is the occupancy indicator for class k in voxel V, W_k is the domain weight of class k (Table 1), M is the total number of semantic classes, $\bar{W} = (1/M)\ \Sigma\ W_k$ is the unconditional mean weight (computed once over all classes, not per voxel), and $\lambda = 0.3$ controls how strongly variance among occupied classes contributes to the importance score.

Table 1
Domain importance weights for semantic feature classes. Where a range is given, the value within the range is selected during preprocessing based on local conditions (e.g., the specific class of weather risk or facility); a single representative value is used per voxel during scoring.

Feature class	Weight	Rationale
Permanent no-fly zone	1.0	Absolute prohibition
Restricted flight zone	0.9	Dynamic control
Tall buildings (> 50 m)	0.8	Major obstacles
Medium buildings (20–50 m)	0.6	Moderate obstacles
Weather-risk / critical facility	0.5–0.6	Flight stability
Low-rise buildings (< 20 m)	0.3	Minor impact
Open ground / water	0.10–0.15	Lowest risk

TS-AdaptiveSubdivide algorithm

The full pseudo-code is shown in Fig. 1, using the following notation: Q is a working queue holding voxels awaiting evaluation, C is the candidate terminal-voxel set accumulated during top-down processing, and GeoSOT_Subdivide(V) denotes the standard GeoSOT-3D split operation that produces 4 children at level l + 1 by halving the longitude and latitude extents of V. The algorithm proceeds in four steps: (i) initialise Q with all G_0 voxels and compute R_{max}, S_{max}; (ii) top-down — for each dequeued V, compute Φ(V); if Φ(V) > τ and l < L_{max}, call GeoSOT_Subdivide(V) and enqueue the 4 children, otherwise mark V as a terminal-state element and add it to C; (iii) bottom-up merging — process C from the deepest level upwards: if all 4 sibling voxels $\{V_{sib}\}$ satisfy $\Phi(V_{sib}) < \tau_{merge}$ and share consistent semantic labels, remove them from C and label their parent V_{parent} as terminal, adding it to C; (iv) output G* = C.

Adjacency between terminal voxels uses a KD-Tree on terminal centroids (search radius 1.6× cell diagonal) with symmetrisation: if voxel A is found in B's neighbour list, B is added to A's list as well, ensuring the resulting graph is undirected. This centroid-based adjacency is computationally efficient

and well suited for irregular multi-resolution grids, although it is an approximation that does not strictly enforce face-adjacency topology as in classical structured meshes; in practice the 1.6× search radius proved sufficient to capture all geometrically meaningful neighbours across the level distributions reported in our experiments. The A* path cost between adjacent voxels V_i and V□ is:

$$\mathrm{cost}(V_i, V_j) = d \cdot (1 + \mu \cdot \mathrm{risk}(V_j)) \quad (4)$$

where d is the Euclidean centroid distance and $\mathrm{risk}(V_j)$ is the maximum semantic weight of V_j; the risk multiplier μ = 5 was chosen so that traversing a high-risk voxel (W = 1.0) costs six times its geometric distance, ensuring semantic-driven detours dominate when alternatives exist. No-fly voxels are treated as non-traversable nodes by assigning them a sentinel cost (10^9) that effectively excludes them from the search graph while keeping all arithmetic finite. The number of terminal voxels |G*| is bounded above by $4^{L_{max}-L_0}$, and A* visits each voxel at most once with O(log |G*|) heap operations, giving an overall complexity of O(|G*| log |G*|); in practice the greedy criterion prunes most subtrees early, so |G*| stays well below the worst-case bound (see Tables 2–3).

Experimental setup

Synthetic scene

5 × 5 km, 0–180 m altitude; fractal-noise terrain (175 m relief), 15 random buildings (20–80 m), one no-fly zone, one restricted zone, and one river (illustrated in Supplementary Fig. S2). For each repetition, building positions and start–goal pairs were re-sampled uniformly within the scene; the terrain field itself was held fixed across repetitions to isolate the variability of the placement and planning components.

Real-world scene

Xi'an Bell Tower–Giant Wild Goose Pagoda district (108.85°–109.00° E, 34.18°–34.30° N; ~13.8 × 13.3 km). Data sources are summarised in Supplementary Table S2 and include ASTER GDEM V3 (30 m)[28], OpenStreetMap buildings (~ 4,200)[29], eAIP/DJI FlySafe airspace, and OpenStreetMap waterways. The DEM of the study area is shown in Supplementary Fig. S3. The DEM was held fixed; randomisation across the 30 repetitions was over the start–goal pair locations only.

Baselines: (1) Uniform grid (~ 50 m); (2) Terrain-AMR (α = 0.9, β = 0, γ = 0.1); (3) Semantic-Only (α = 0, β = 0.9, γ = 0.1). **TSDP configuration**: L5–L9, α = 0.6, β = 0.3, γ = 0.1, τ = 0.35, τ_merge = 0.20.

Platform

Intel i7-12700H, 32 GB RAM, single-threaded Python 3.13 (no GPU acceleration); reported wall-clock times exclude I/O.

Metrics and statistics

(i) Elevation RMSE — the root-mean-square deviation between the voxel-reconstructed elevation field (each terminal voxel assigned the mean elevation over its DEM coverage) and the original GDEM V3 input; this measures how faithfully the discretised grid represents the source terrain, and is therefore lower than the absolute vertical accuracy of GDEM V3 itself (≈ 8.68 m, see Supplementary Table S2). (ii) Maximum elevation error — the worst-case absolute deviation across all terminal voxels, capturing peaks/cliffs missed by coarse cells. (iii) Building F1 — see protocol below. (iv) NFZ boundary error — the mean perpendicular distance between the voxel-discretised no-fly-zone boundary and the ground-truth polygon boundary, sampled at 1 m intervals. (v) A* runtime, expanded nodes, and total voxel count — measured per Python's $time.perf_counter()$ and standard A* counters, single-threaded, with peak resident set size obtained from $tracemalloc$. Each experimental setting was repeated 30 times. Distribution normality was tested per metric using the Shapiro–Wilk test; in all cases Shapiro–Wilk rejected normality ($W < 0.95$, $p < 0.05$), motivating non-parametric pairwise comparisons via the Wilcoxon signed-rank test. All ± values denote standard deviation (s.d.); exact p values are reported throughout.

Building F1 protocol

A building is counted as "detected" when its OpenStreetMap footprint intersects at least one terminal voxel labelled *building* in the horizontal plane. Precision is the fraction of building-labelled voxels whose footprint actually contains a building, and recall is the fraction of buildings detected. Because all tested grid resolutions are sufficient to cover every building footprint, recall is 1.000 across methods, and F1 differences arise from precision—specifically, from the fraction of empty voxels incorrectly assigned to building proximity.

Results

Synthetic environment

TSDP refines to L8–L9 in undulation zones, densifies around buildings, and retains L5–L6 in flat areas (Fig. 2). The total terminal voxel count is 36,782 (64.1% compression), with level distribution L5: 12.3%, L6: 31.2%, L7: 28.7%, L8: 19.5%, L9: 8.3%.

Multi-baseline comparison results are reported in Table 2. TSDP RMSE (2.41 m) is statistically indistinguishable from Terrain-AMR (2.38 m; Wilcoxon $p = 0.23$). Building F1 (0.85) significantly exceeds Terrain-AMR (0.74; $p = 0.0003$) and modestly improves on Semantic-Only (0.83; $p = 0.038$). Compression (64.1%) outperforms both single-driver baselines (Terrain-AMR 58.8%, Semantic-Only 55.2%; both $p < 0.001$).

Table 2
Multi-baseline comparison on the synthetic environment. Values after ± denote s.d. across 30 repetitions.

Dimension	Metric	Uniform	Terrain-AMR	Semantic-Only	TSDP
Geometric	RMSE (m)	2.82 ± 0.15	2.38 ± 0.11	2.79 ± 0.14	2.41 ± 0.12
	Max. error (m)	8.7	5.9	8.5	6.3
Semantic	Building F1	0.73	0.74	0.83	0.85
	NFZ boundary error (m)	12.5	12.1	8.2	7.8
Efficiency	Total voxels	102,400	42,156	45,893	36,782
	Compression	—	58.8%	55.2%	64.1%

Parameter sensitivity

The optimal threshold range is $\tau \in [0.30, 0.40]$; balanced weights ($\alpha = 0.6$, $\beta = 0.3$) yield the best trade-off (60.2% compression). Perturbation of W_k by ± 20% over 100 trials yields $F1 \in [0.82, 0.87]$ (s.d. = 0.013). The full sensitivity surface is shown in Fig. 3, including a (α, β) heatmap that confirms the chosen operating point lies within a stable plateau.

System performance

Empirical scaling across scenes of 2, 5, 8, and 10 km side length (each repeated 5 times) shows TSDP follows an empirical sub-linear trend with side length (best-fit power-law exponent ≈ 0.68 against scene area), versus the near-linear trend of the uniform baseline (≈ 0.95); the corresponding 10 km path planning takes 3.85 s under TSDP, approximately 43% of the uniform baseline. Dynamic re-planning takes 0.86 ± 0.18 s (s.d.), a 45.3% reduction relative to uniform (Wilcoxon $p = 0.0007$), affecting approximately 28% of the domain on average. The Spearman rank correlation between $\Phi(V)$ and the final voxel level, computed across all 36,782 terminal voxels of the synthetic scene, is $\rho = 0.91$ ($p < 10^{-6}$); we note this correlation is high by construction—$\Phi(V)$ is the input to the subdivision rule that determines level—and it is reported here as an internal-consistency check rather than an independent validation. Peak resident set size is 35.8% of the uniform baseline, with disk I/O reduced by 55–65% across the tested scene sizes.

Real-world scene

In the Xi'an district, elevations range from 380 to 800 m with approximately 4,200 buildings and 6 airspace zones. The City Wall–Bell Tower core refines to L8–L9, the Pagoda periphery to L7–L8, and flat northern areas remain at L5–L6. TSDP level distribution is highly skewed toward fine resolutions: L9 contains 23,348 voxels (79.6%), L8 contains 4,664 (15.9%), L7 contains 1,085 (3.7%), L6 contains 205

(0.7%), and L5 contains 29 (0.1%), summing to the 29,331 total reported in Table 3. This concentration confirms that adaptive refinement focuses computational resources on geometrically and semantically complex regions.

A visual comparison of the four partitioning strategies is shown in Fig. 4, and quantitative results are reported in Table 3. A close-up local comparison between uniform sampling and TSDP is provided as Supplementary Fig. S4.

Table 3
Real-world experimental results on the Xi'an district. Values after ± denote s.d.

Dimension	Metric	Uniform	Terrain-AMR	Semantic-Only	TSDP
Geometric	RMSE (m)	3.35	5.15	3.38	3.97
	Max. error (m)	50.7	65.9	52.8	48.7
Semantic	Building F1	0.933	0.842	0.935	0.918
	Precision	0.875	0.727	0.877	0.849
	Recall	1.000	1.000	1.000	1.000
Efficiency	Total voxels	73,682	4,872	65,121	29,331
	Compression	—	93.4%	11.6%	60.2%
Path	A* time (s)	10.61 ± 0.30	1.03 ± 0.04	11.72 ± 0.23	6.26 ± 0.21
	Expanded nodes	93,538	3,789	86,612	29,946
	Path length (m)	17,541	16,414	17,581	17,088
	Mean safety (m)	23.6	194.9	27.5	80.8

TSDP achieves 60.2% compression with F1 = 0.918, slightly below the uniform baseline (0.933) and Semantic-Only (0.935), but with the lowest maximum elevation error (48.7 m) among all methods. Planning time is reduced by 41.1% relative to uniform (6.26 ± 0.21 s vs. 10.61 ± 0.30 s; Wilcoxon $p < 10^{-5}$), and mean safety distance to obstacles improves from 23.6 m to 80.8 m. The higher real-world RMSE (3.97 m) compared to the synthetic scene reflects the 30 m horizontal resolution limit of GDEM V3 rather than methodological degradation; we note however that this same data limitation applies to the uniform baseline (3.35 m), and the residual gap is discussed below.

Path-planning comparison

Path-planning performance on the synthetic environment (Supplementary Table S3) shows TSDP reduces planning time by 40.1% (1.42 ± 0.11 s vs. 2.37 ± 0.19 s for uniform, both s.d.; Wilcoxon p = 0.0002) and increases the minimum obstacle clearance from 4.2 m to 8.5 m, eliminating all clearance violations (defined as the path approaching within 5 m of an obstacle: 3 such violations under the uniform baseline, 0 under TSDP across the 30 repetitions). Path lengths differ by less than 2% across

methods, confirming that the safety improvements do not come at the cost of route efficiency. Note that the synthetic scene reports minimum clearance whereas the real-world experiment (Table 3) reports mean safety distance: minimum clearance is the worst-case approach distance along a planned path and is most relevant in tightly cluttered environments, while mean safety distance averages the approach distance over the path and better reflects overall route quality in the larger-scale Xi'an scene. A visual comparison of paths produced by each method on the Xi'an scene is shown in Fig. 5.

Discussion

The dual experiments validate TSDP's value as a **middle-ground operating point** between extreme single-driver methods. Terrain-AMR achieves the highest compression (93.4%) but at a substantial cost to building detection (F1 = 0.842), while Semantic-Only preserves semantic accuracy at the price of poor compression (11.6%). TSDP fills the gap, suitable for scenarios that require both semantic reliability and computational economy—onboard navigation (35.8% memory footprint) and large-scale path planning (41.1% speedup) being the most direct applications. More broadly, the ablation pattern across the three baselines suggests that spatial discretisation efficiency in low-altitude environments is governed not by geometric variation alone but by the joint distribution of geometric variation and semantic constraints, indicating a dual-factor regime in which neither factor can be neglected without measurable accuracy cost.

A nuance worth highlighting concerns the mean safety distance metric. Terrain-AMR's 194.9 m mean safety distance superficially appears superior to TSDP's 80.8 m, but this advantage is mechanically produced by the coarseness of its grid—coarse voxels physically displace the planner away from obstacles, at the cost of path-resolution (only 3,789 expanded nodes). TSDP's 80.8 m mean safety distance, by contrast, is achieved while preserving 29,946 expanded nodes, reflecting genuine semantic-driven avoidance rather than a granularity artefact. This distinction matters in practice: planners that operate on Terrain-AMR grids cannot represent fine-scale corridors, whereas TSDP retains the option to plan at high resolution where it matters.

Where TSDP does not dominate. In scenes where terrain undulation is the only meaningful feature and semantic objects are sparse, Terrain-AMR alone is sufficient; the additional semantic-driven refinement in TSDP yields little benefit and slightly increases voxel count. Conversely, in scenes dense with semantic objects but with negligible terrain relief, Semantic-Only achieves comparable accuracy at lower implementation cost. The advantage of TSDP is concentrated in mixed scenes where both factors matter, which describes most real urban airspace environments. Practitioners deploying TSDP should therefore characterise the dominant feature mode of their target environment before committing to dual-driven partitioning.

Relation to other DGGS approaches. TSDP's reliance on the GeoSOT-3D 2^n quadtree is a deliberate trade-off in the broader DGGS design space. Hexagonal systems such as Sahr's hexagonal DGGS[14] offer better isotropy and uniform neighbour distances, while production-scale geospatial computing platforms such

as Uber H3[9] and Google S2[10] have demonstrated the practical scalability of icosahedral and hexagonal indexing. The semantic-driven refinement mechanism we describe is in principle portable to those systems—the decision function $\Phi(V)$ makes no assumption about cell shape—and porting TSDP to other DGGS such as H3 or S2, where differences in cell topology do not affect the general applicability of the decision function, is a natural follow-up. Conversely, the integer-encoding advantages of GeoSOT-3D (lossless cross-platform exchange via fixed-length codes) are particularly suited to the regulatory and audit requirements of low-altitude airspace management, where unambiguous referencing of voxel identity matters more than geometric isotropy.

Engineering boundaries. Three constraints currently shape the practical reach of TSDP. First, the linear form of $\Phi(V)$ is a deliberate simplification of what could be a nonlinear interaction between $R(V)$ and $S(V)$; for instance, a tall building on steep terrain may warrant refinement beyond what the linear sum predicts. Second, manual calibration of W_k limits portability across regulatory jurisdictions, since "no-fly zone" itself is not a globally uniform category. Third, the per-voxel greedy criterion does not provide the formal a-posteriori error guarantees available in classical AMR[15,16]; $\Phi(V) > \tau$ is an empirical proxy for the unmet fidelity constraints rather than an upper bound on the resulting error. These boundaries do not preclude practical deployment but they should temper claims of optimality.

Application potential spans large-scale urban air traffic management (10 km scenes), emergency re-planning (sub-second response), onboard navigation under memory constraints, and multi-resolution collaborative decisions across heterogeneous platforms. The integer encoding inherited from GeoSOT-3D enables efficient cross-platform data exchange that resource-specific representations such as oblique-photography meshes cannot provide.

More generally, this work can be viewed as a unified framework for multi-objective spatial discretisation under heterogeneous geometric and semantic constraints; while we instantiate it on GeoSOT-3D for low-altitude airspace, the same dual-driven decision principle applies to any DGGS or AMR scheme where multiple, weakly correlated factors compete for finite representational capacity.

Several limitations frame the scope of future work. First, the linear form of $\Phi(V)$ cannot capture nonlinear interactions between terrain and semantics; machine-learning extensions—including reinforcement-learning auto-tuning and Bayesian optimisation of weights—are warranted. Second, W_k requires manual calibration per environment; a learned calibration would improve generalisability. Third, only one city has been tested in detail, and cross-city transfer studies are needed to confirm robustness. Fourth, the GDEM V3 30 m resolution sets an upper bound on terrain fidelity; airborne LiDAR data would lift this ceiling. Fifth, as noted in the Scope statement, the current implementation is a horizontal-plane quadtree; a full octree-based volumetric extension and 4D spatio-temporal modelling are the next major development steps. Sixth, ground truth for F1 is derived from OpenStreetMap; independent validation using high-resolution remote sensing would strengthen the semantic accuracy claims.

Future directions include reinforcement-learning and Bayesian auto-tuning of $(\alpha, \beta, \gamma, \tau)$; 4D spatio-temporal models incorporating dynamic airspace; GPU-parallel implementations for real-time large-scale

deployment; multi-city generalisation studies; and multi-source data fusion combining LiDAR, oblique photography, and live air-traffic feeds.

Declarations

Competing interests

The authors declare no competing interests.

Funding

This work was supported by the Xi'an Social Science Fund (Grant No. 25ZL71).

Author Contribution

N.Z. conceived the study and supervised experiments. Q.W. implemented the model and conducted experiments. Z.L. participated in data collection and analysis. All authors reviewed the manuscript.

Acknowledgement

We thank NASA for ASTER GDEM V3 data and the OpenStreetMap community for building and waterway vectors.

Data Availability

All datasets used in this study are publicly available: ASTER GDEM V3 from NASA Earthdata (https://search.earthdata.nasa.gov/), OpenStreetMap data via Overpass API (https://www.openstreetmap.org/), and airspace data from public eAIP and DJI FlySafe sources. Source code for the TS-AdaptiveSubdivide algorithm and all data processing scripts is available at https://github.com/dongmh/geosot3d-tsdp and archived at Zenodo (https://doi.org/10.5281/zenodo.20025515).

References

1. Bauranov, A. & Rakas, J. Designing airspace for urban air mobility: a review of concepts and approaches. *Prog Aerosp. Sci.* **125**, 100726. https://doi.org/10.1016/j.paerosci.2021.100726 (2021).
2. Doole, M., Ellerbroek, J. & Hoekstra, J. Estimation of traffic density from drone-based delivery in very low level urban airspace. *J. Air Transp. Manag*. **88**, 101862.

https://doi.org/10.1016/j.jairtraman.2020.101862 (2020).

3. Wan, L. et al. Airspace conflict detection based on GeoSOT-3D. *J. Command Control*. **11**, 738–745 (2025).
4. Han, B. & Qu, T. GeoSOT-3D grid modeling for spatial AI. *Int. Arch. Photogramm Remote Sens. Spat. Inf. Sci.* **XLVIII-4-2024**, 233–238. https://doi.org/10.5194/isprs-archives-XLVIII-4-2024-233-2024 (2024).
5. Song, S. et al. Global remote sensing data subdivision organization based on GeoSOT. *Acta Geod. Cartogr. Sin*. **43**, 869–876 (2014).
6. Liu, Z., Nan, Y. & Xie, R. Airspace conflict detection via subdivision grids. *Shipboard Electron. Eng.* **42**, 157–162 (2022).
7. Zhou, J. et al. Multi-mode 3D extension for equal-area global grids. *Acta Geod. Cartogr. Sin*. **53**, 173–188 (2024).
8. Open Geospatial Consortium. Topic 21: Discrete Global Grid Systems Abstract Specification, Version 2.0 (OGC 15-104r5. (2017). https://docs.ogc.org/as/15-104r5/15-104r5.html
9. Brodsky, I. *H3: A Hexagonal Hierarchical Geospatial Indexing System*Uber Engineering, Open-source implementation; software repository: (2018). https://github.com/uber/h3; documentation: https://h3geo.org
10. Google & Google Inc.. *S2 Geometry Library: Hierarchical Spatial Indexing on the Sphere* (, Open-source implementation; software repository: (2017). https://github.com/google/s2geometry; documentation: https://s2geometry.io
11. Goodchild, M. F. Discrete global grids: a web of geospatial knowledge. *Geomatica* **54**, 419–426 (2000).
12. Dutton, G. *A Hierarchical Coordinate System for Geoprocessing and Cartography* (Springer, 1999). https://doi.org/10.1007/978-3-642-60214-1
13. Sahr, K., White, D. & Kimerling, A. J. Geodesic discrete global grid systems. *Cartogr. Geogr. Inf. Sci.* **30**, 121–134. https://doi.org/10.1559/152304003100011090 (2003).
14. Sahr, K. Hexagonal discrete global grid systems for geospatial computing. *Arch. Photogramm Cartogr. Remote Sens.* **22**, 363–376 (2011).
15. Berger, M. J. & Oliger, J. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **53**, 484–512. https://doi.org/10.1016/0021-9991(84)90073-1 (1984).
16. Berger, M. J. & Colella, P. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.* **82**, 64–84. https://doi.org/10.1016/0021-9991(89)90035-1 (1989).
17. De Zeeuw, D. L. & Powell, K. G. An adaptively refined Cartesian mesh solver for the Euler equations. *J. Comput. Phys.* **104**, 56–68. https://doi.org/10.1006/jcph.1993.1005 (1993).
18. Popinet, S. Gerris: a tree-based adaptive solver for the incompressible Euler equations in complex geometries. *J. Comput. Phys.* **190**, 572–600. https://doi.org/10.1016/S0021-9991(03)00298-5 (2003).

19. Losasso, F., Gibou, F. & Fedkiw, R. Simulating water and smoke with an octree data structure. *ACM Trans. Graph*. **23**, 457–462. https://doi.org/10.1145/1015706.1015745 (2004).
20. Liu, Z., Tian, F. B. & Feng, X. Geometry-adaptive mesh refinement for immersed boundary lattice Boltzmann methods. *Comput. Methods Appl. Mech. Eng.* **392**, 114662. https://doi.org/10.1016/j.cma.2022.114662 (2022).
21. Samet, H. *Foundations of Multidimensional and Metric Data Structures* (Morgan Kaufmann, 2006).
22. Guttman, A. R-trees: a dynamic index structure for spatial searching. In *Proc. ACM SIGMOD*, 47–57 (ACM, 1984). https://doi.org/10.1145/602259.602266
23. Finkel, R. A. & Bentley, J. L. Quad trees: a data structure for retrieval on composite keys. *Acta Inf.* **4**, 1–9. https://doi.org/10.1007/BF00288933 (1974).
24. Pajarola, R. & Gobbetti, E. Survey of semi-regular multiresolution models for interactive terrain rendering. *Vis. Comput.* **23**, 583–605. https://doi.org/10.1007/s00371-007-0147-2 (2007).
25. Gröger, G. & Plümer, L. CityGML—interoperable semantic 3D city models. *ISPRS J. Photogramm Remote Sens.* **71**, 12–33. https://doi.org/10.1016/j.isprsjprs.2012.04.004 (2012).
26. Biljecki, F., Ledoux, H. & Stoter, J. An improved LOD specification for 3D building models. *Comput. Environ. Urban Syst.* **59**, 25–37. https://doi.org/10.1016/j.compenvurbsys.2016.04.005 (2016).
27. Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S. & Çöltekin, A. Applications of 3D city models: state of the art review. *ISPRS Int. J. Geo-Inf*. **4**, 2842–2889. https://doi.org/10.3390/ijgi4042842 (2015).
28. Tachikawa, T. et al. ASTER Global Digital Elevation Model V3 [Data set]. *NASA EOSDIS Land Processes DAAC* (2020). https://doi.org/10.5067/ASTER/ASTGTM.003
29. OpenStreetMap contributors. Planet dump. (2025). https://www.openstreetmap.org

Figures

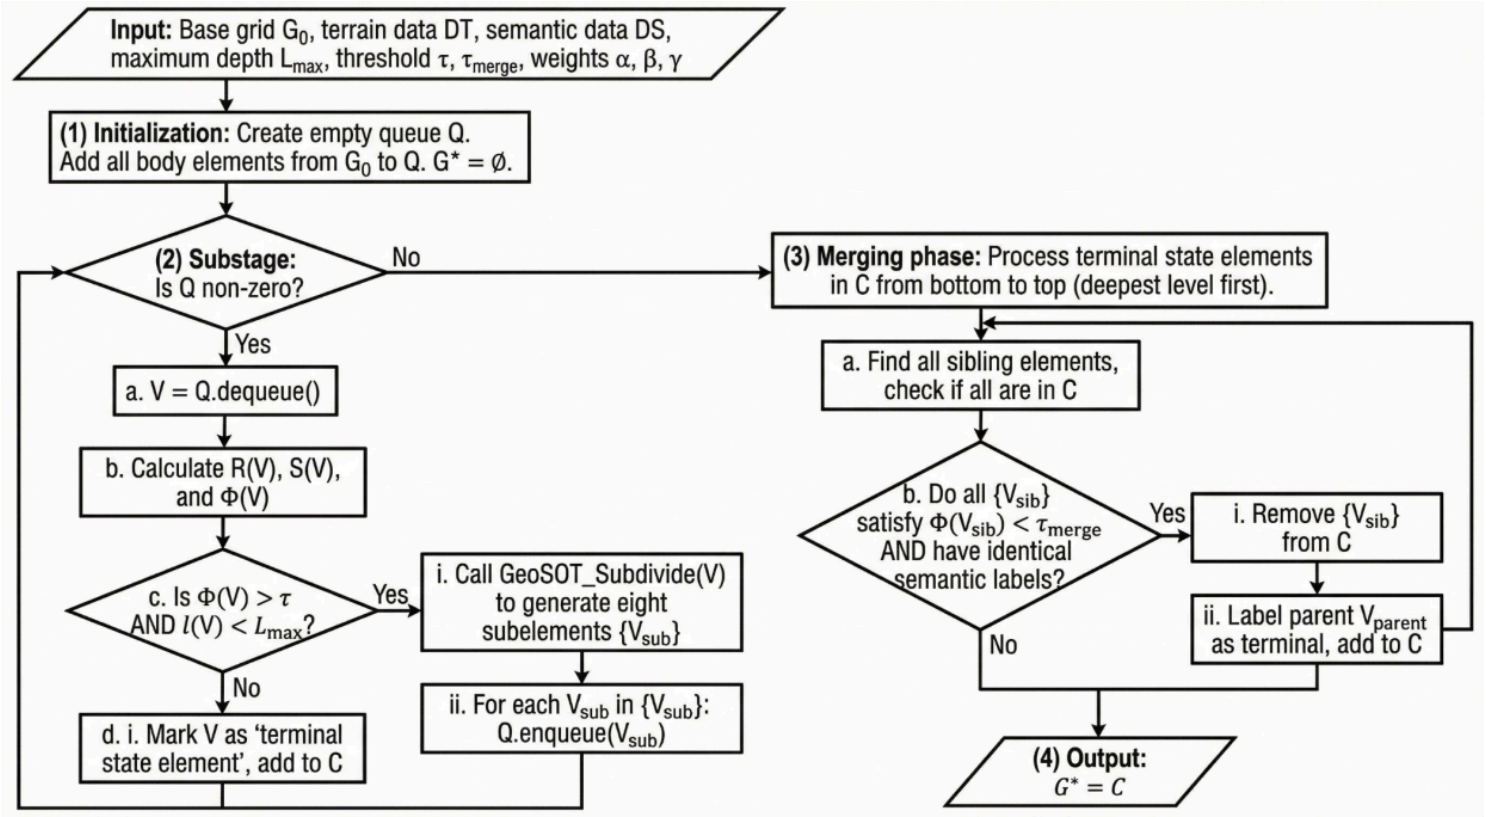


Figure 1

TS-AdaptiveSubdivide algorithm flowchart. Phase I performs top-down subdivision driven by Φ(V); Phase II performs bottom-up sibling merging when Φ < τ_merge and labels are consistent.

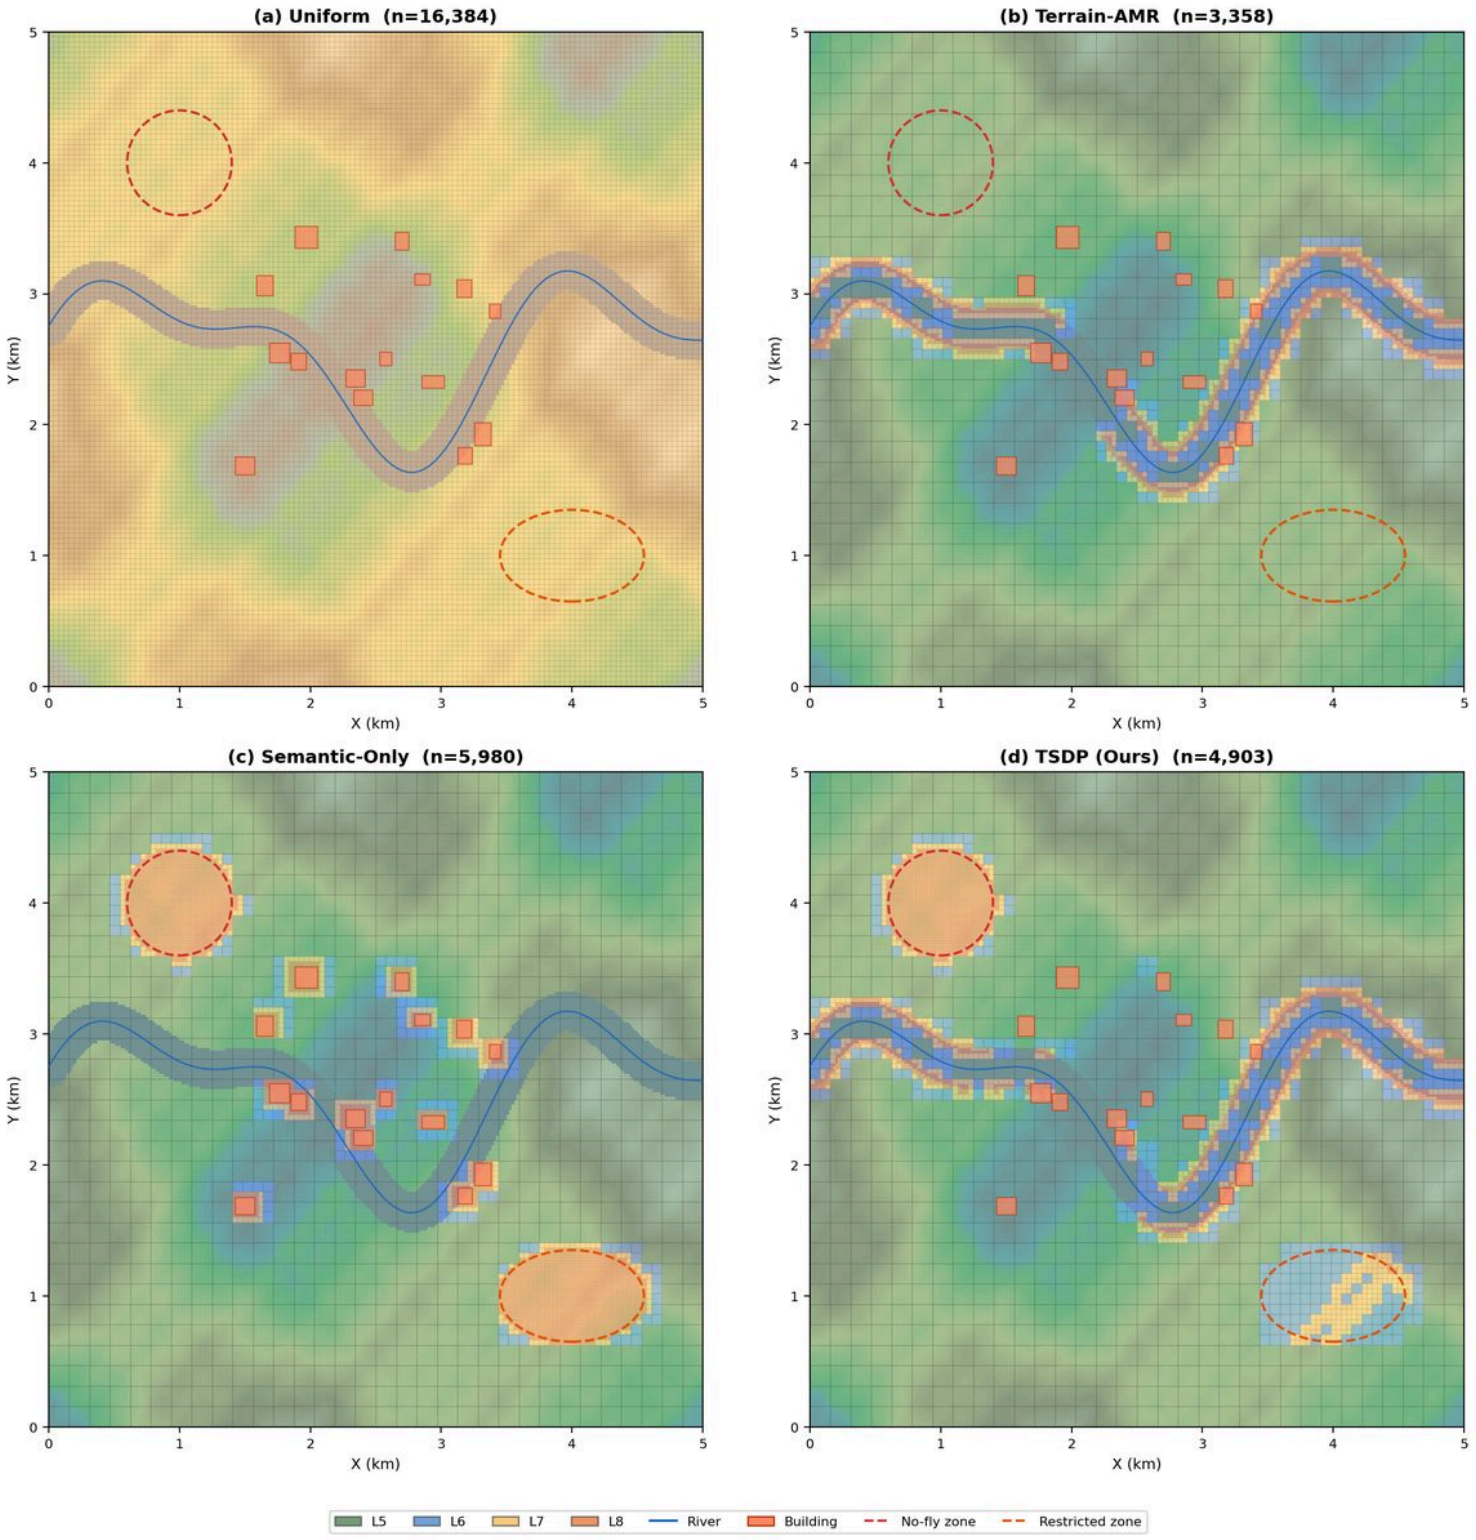


Figure 2

TSDP subdivision result on the synthetic scene, showing automatic multi-resolution refinement: dense voxels concentrate in undulation zones and around buildings, while flat regions retain coarse voxels.

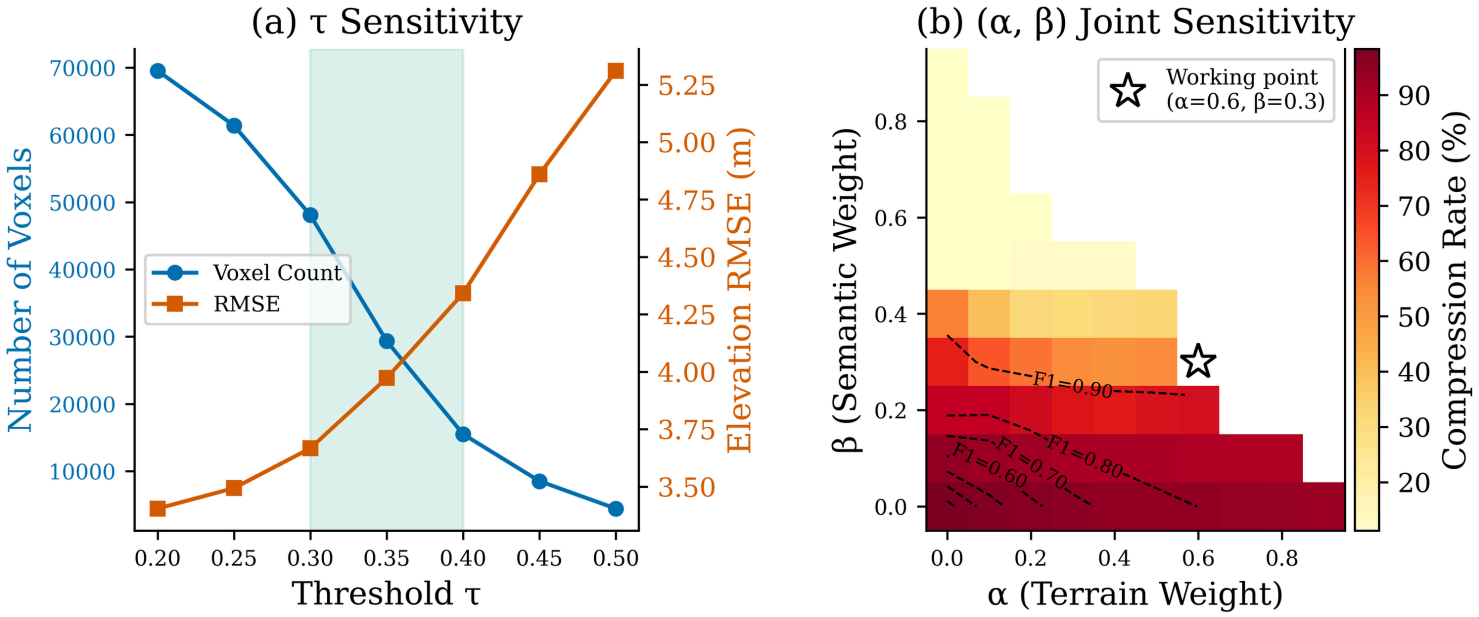


Figure 3

Parameter sensitivity analysis. (a) τ sensitivity scan showing voxel count (blue, left axis) and elevation RMSE (orange, right axis) as functions of the subdivision threshold τ; the shaded green region highlights the optimal range τ ∈ [0.30, 0.40]. (b) Joint (α, β) sensitivity heatmap. Colour encodes the resulting compression rate; dashed contours indicate Building F1. The chosen working point (α = 0.6, β = 0.3, marked with ☆) lies on a stable plateau where compression and F1 are jointly favourable.

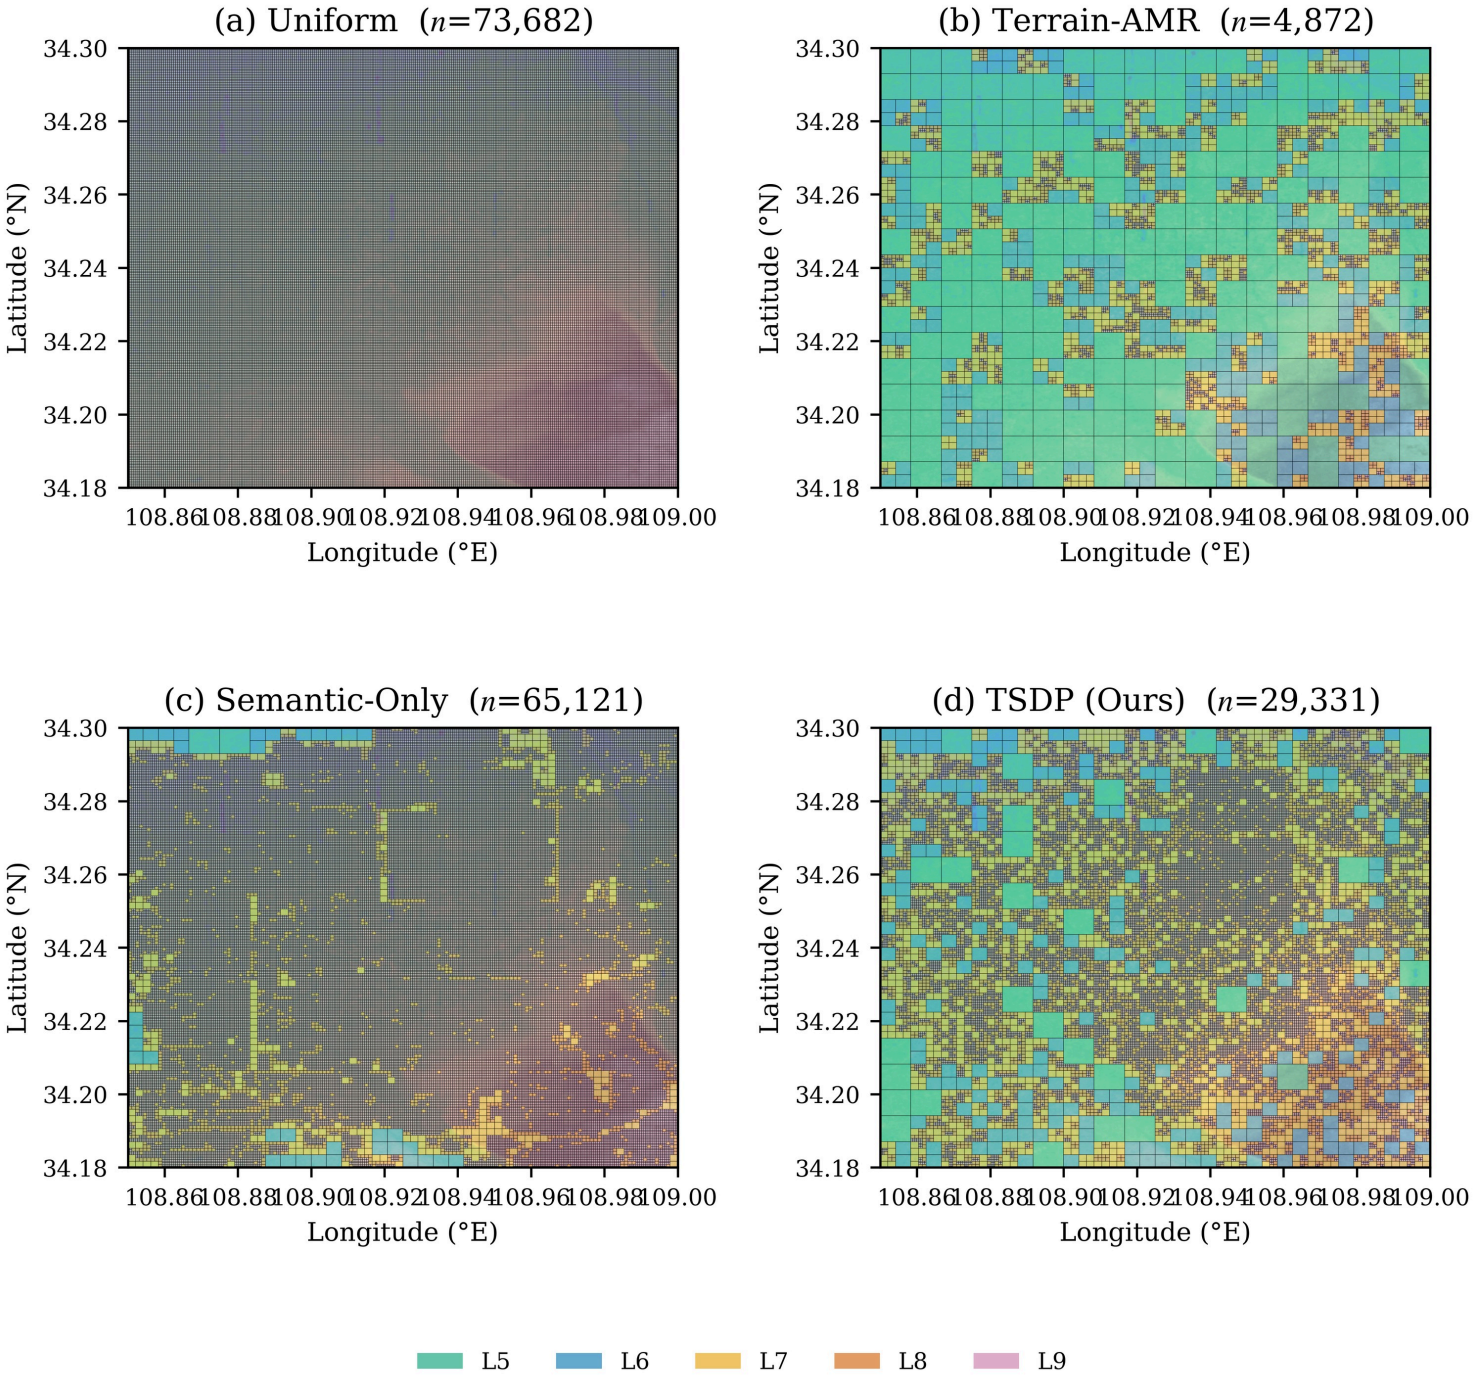


Figure 4

Real-world grid comparison on the Xi'an district. (a) Uniform grid, n = 73,682; (b) Terrain-AMR, n = 4,872; (c) Semantic-Only, n = 65,121; (d) TSDP, n = 29,331. Cell colour encodes the GeoSOT-3D subdivision level (legend below). TSDP concentrates fine voxels (L7–L9, yellow/orange/pink) along the historic core and major obstacles while keeping peripheral flat areas at coarse levels (L5–L6, green/blue).

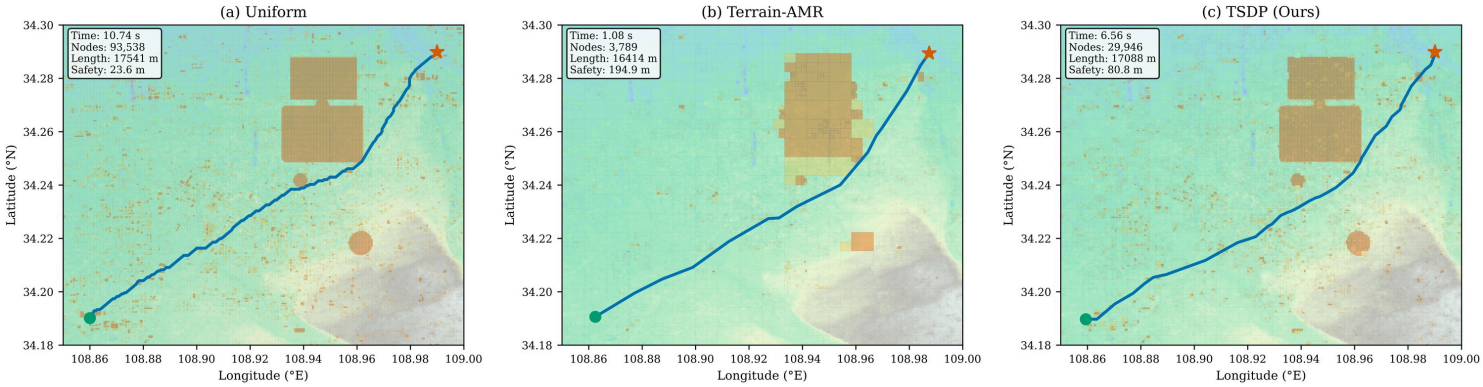


Figure 5

Path planning on the Xi'an scene. (a) Uniform grid; (b) Terrain-AMR; (c) TSDP. The blue line denotes the planned path; the green dot and orange star mark start and goal respectively. Inset boxes report planning time, expanded nodes, path length, and mean safety distance for the displayed run. TSDP follows a route comparable in length to the uniform path but with substantially larger margins around buildings and the no-fly zone.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- GeoSOT3DTSDPSupplementary.docx