ORIGINAL RESEARCH **Open access**

Efficient Autonomous UAV Exploration for Geospatial Mapping in Unknown Environments

Kim Jae-sung[1*], Park Soo-yeon[1], Lee Min-ho[2], Jung Hyun-woo[3]

Abstract

Autonomous exploration plays a vital role in mapping-related tasks such as information gathering, environmental surveillance, and search and rescue. Unmanned aerial vehicles (UAVs), valued for their low cost and superior maneuverability, have become central to these operations, particularly within complex or hostile settings. Nevertheless, current methods often suffer from drawbacks, such as repetitive exploration and inconsistent flight patterns. In this research, we introduce a hierarchical exploration methodology designed for UAVs with restricted fields of view in geospatial mapping scenarios. This methodology addresses these challenges through a hybrid viewpoint-generation process, an inventive frontier-exploration sequence, and a dual-stage global path-planning strategy. To strike a balance between exploration effectiveness and computational load, we use a hybrid scheme that blends collision-free spherical sampling with adaptive viewpoint construction based on stochastic differential equations. This scheme delivers high-grade candidate viewpoints while curbing computational demands. Moreover, we propose a novel heuristic evaluation function to order frontiers within compact areas, thereby supporting optimal path design. Drawing on this setup, the global coverage route is cast as a traveling salesman problem (TSP). The dual-stage global planning structure consists of a primary stage that executes a history-sensitive trajectory-refinement strategy with smoothing adjustments, followed by a secondary stage that applies a sliding-window TSP algorithm to formulate the global path. This arrangement alleviates motion variability triggered by recurrent heuristic modifications and bolsters flight consistency and path smoothness. To gauge the effectiveness of the suggested framework, we compare it with leading-edge techniques in both synthetic and real-world environments. Experimental results show that our method reduces travel distance and shortens exploration time, thereby improving overall exploration performance.

Keywords Mapping, Applications, Unmanned aerial vehicles (UAVs), Geospatial surveying applications, Path planning

*Correspondence:
Kim Jae-sung
jaesung.kim@snu.ac.kr
[1] Department of AI for Kimchi Cabbage and Radish Growth Prediction, College of Agriculture, Seoul National University, Seoul, South Korea
[2] Department of Smart Greenhouse Energy Optimization, College of Agriculture, Pusan National University, Busan, South Korea
[3] Department of Machine Learning for Yield Forecasting, College of Agriculture, Kyungpook National University, Daegu, South Korea

Introduction

Unmanned aerial vehicles (UAVs) have become crucial tools in modern geospatial surveying and mapping, owing to their versatility, economic efficiency, and ability to quickly and accurately capture spatial data. They perform a key function in topographic mapping [1], 3D reconstruction [1], infrastructure assessment [2], and environmental observation [3, 4], serving as an effective medium for acquiring geospatial data across wide-ranging, often hard-to-reach landscapes. Despite their rising prevalence, standard UAV mapping operations typically rely on predetermined flight trajectories or human piloting, which hampers their use in unfamiliar, GPS-unavailable, or structurally complex areas such as tunnels,

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

mines, and enclosed industrial facilities. This reliance restricts flexible data capture and can lead to patchy coverage, superfluous scanning, or partial spatial reconstruction, thereby weakening the geometric fidelity and productivity of mapping outputs. To address these drawbacks, self-governing UAV-based surveying has emerged as a primary area of inquiry in geomatics and remote sensing, aiming to enable intelligent scouting, flexible viewpoint scheduling, and optimal trajectory planning for precise 3D mapping in challenging environments. As a result, considerable investigative effort has been devoted by both civic and academic circles to autonomous exploration of mapping, resulting in a multitude of refined algorithms broadly categorized as frontier-driven, sampling-driven, and machine-learning-driven. Frontier-driven techniques endeavor to optimize exploration extent by pinpointing and probing the transitions between mapped and unmapped zones. While boundary grouping has demonstrated strong capability to improve exploration results, its high computational cost often lengthens planning durations, thereby reducing overall exploration speed and efficiency. Sampling-driven techniques, often deployed in greedy schemes, fixate on picking either the closest prospective goal or the one offering the greatest informational return as the next flight position. Though this tactic offers straightforward benefits in some contexts, it often misses viable targets in nearby areas, resulting in suboptimal global routing, wasteful backtracking, and tentative flight behavior—eventually retarding the entire exploration cycle. Learning-driven techniques, by contrast, show considerable promise in harnessing environmental context cues. Yet, their intense computational needs and inadequate real-time responsiveness hamper their practicality in field implementations. Among established works, FUEL [5]—a cutting-edge autonomous exploration framework for surveying and mapping—melds a stepwise frontier information structure (FIS) with a tiered planning mechanism to achieve swift, effective, and secure exploration in complex areas, yielding fluid, continuous flight paths. Despite these advancements, current methods still face several real-world obstacles. For example, the processing weight associated with ray-casting evaluation of candidate viewpoints limits gains in exploration speed mid-flight. Faulty boundary identification and chaotic exploration orders might also prompt the UAV to bypass minor zones, resulting in unnecessary or even counterproductive motion sequences. Furthermore, repeated re-optimization and heuristic recomputations habitually provoke tactical wavering, thereby undermining the system's overall alertness.

To address the fundamental inadequacies of current autonomous exploration frameworks, this paper proposes a hierarchical exploration architecture for limited-field-of-view (FOV) drones used in mapping and remote sensing. To reduce ray-casting overhead, the viewpoint creation component uses a hybrid sampling method that combines collision-safe spherical sampling with uniform sampling. Furthermore, in-flight energy expenditure is expected to yield more accurately calibrated candidate viewpoints. Upon producing these viewpoints, a frontier-heuristic appraisal function is suggested to polish the frontier exploration order by giving precedence to boundary segments and incorporating an original small-zone frontier detection scheme. To reduce hesitant stop-start flight behavior, the perceptual path history of each flight is recorded and tapped during replanning. Moreover, a flight-smoothness penalty is incorporated to advance trajectory steadiness and coherence. Against top-tier benchmarks, the proposed architecture slashes exploration time by 7–25% and route length by 9–27% over three dissimilar environments. Lastly, the real-world viability of the suggested methodology is corroborated through aerial trials performed in actual settings. The chief contributions of this work can be encapsulated as follows:

· A hybrid viewpoint-generation technique is proposed that combines collision-free spherical sampling with adaptive viewpoint construction based on stochastic differential equations. This technique effectively provides high-quality candidate viewpoints while reducing computational cost.

· A novel heuristic evaluation function is proposed to streamline frontier exploration ordering, prioritizing small-area frontiers to curb oscillatory movement and avoid unnecessary probing.

· A dual-stage global path planning structure is proposed to enhance flight stability and fluidity by coupling a history-informed trajectory enhancement technique with smoothing rectification and a sliding-window TSP scheme.

Related Works

© 2023 The Author(s). | Published by Agro Intelligence Research Press

The capacity of UAVs to explore and map environments autonomously hinges on onboard sensory systems that capture spatial information and guide navigation decisions. Over time, numerous techniques for autonomous exploration have emerged, generally classified into three paradigms: sampling-driven, frontier-driven, and learning-driven methodologies.

Sampling-driven techniques scatter random nodes throughout uncharted space and stochastically select potential observation points using the Rapidly-Exploring Random Tree (RRT) framework [6]. The next best view (NBV) emerges as the candidate delivering the maximal informational return, in keeping with principles intimately connected to the NBV formulation [7]. Seeking to optimize data acquisition, Bircher *et al.* [8] present an advanced, receding-horizon NBV planning tactic that probes unknown regions by weighing the information content along the entire path. Subsequent investigations have refined this concept by tailoring information yield measures to specific surveying objectives. To illustrate, Papachristos *et al.* [9] sought to reduce uncertainty in simultaneous localization and mapping across exploration runs, whereas Papachristos *et al.* [10] adapted the framework for visually impoverished conditions, such as underground excavations and dimly lit passages. Even so, these strategies commonly fall prey to confinement within local basins, culminating in exploration outcomes that fall short of optimal. The essential drawback of sampling-driven frameworks is their proneness to local extrema, repeatedly ensnaring the aerial vehicle within restricted regions and resulting in compromised exploration throughput [11]. To counteract this tendency, Witting *et al.* [11] drew upon historical flight paths to sidestep local entrapments and tuned the UAV's inclination toward uncharted territories to heighten sampling yield. In a parallel vein, Selin *et al.* [12] achieved faster appraisal of prospective information gains and preserved earlier computations by combining frontier search with NBVP, thereby alleviating the challenges posed by vast unmapped expanses. Further, Deng *et al.* [13] reconceived the information gain criterion as a differentiable expression, opening the door to jointly optimizing informational content alongside supplementary performance indicators. To further improve efficiency, Dharmadhikari *et al.* [14] harnessed motion primitives to generate high-quality viewpoints, ensuring fluid, uninterrupted traversal.

Notwithstanding these advances, a pivotal impediment remains unresolved: computational tractability. The voluminous viewpoint sampling and ray-casting procedures demanded by information gain estimation impose untenable processing burdens, particularly in expansive, geometrically intricate domains. In response, Xu *et al.* [15] combined progressive sampling with probabilistic roadmaps to improve computational throughput. Along similar lines, Zhu *et al.* [16] proposed a dynamically branching random tree coupled with a global map to enable swift backtracking and to retain previously constructed nodes for accelerated replanning. Respall *et al.* [17] proposed a blended scheme that combines sampling and frontier-based philosophies to reduce waste when canvassing extensive, unfamiliar territories. UFOExplore [18] maintained a stepwise roadmap dedicated to UAV guidance, facilitating nimbler reactions and steadier advance toward prime exploration targets. That said, these approaches inherently generate a profusion of unusable viewpoints during random generation, and the substantial computational cost of viewpoint authentication via ray casting remains a formidable obstacle to real-time operation on modestly powered UAV hardware.

The frontier-driven philosophy, widely acknowledged in the field, designates the interface between mapped and unmapped space as the frontier, singling out a segment of it as the next exploration destination. The exploration cycle concludes when no further frontiers populate the scene. First formulated by Yamauchi [19] and subsequently enhanced by Juliá *et al.* [20], this paradigm has undergone a range of modifications. As an example, Cieslewski *et al.* [21] engineered a scheme to select boundary zones within the FOV at each planning interval, damping velocity oscillations to maintain a uniform flight pace. Tang *et al.* [22] leveraged the wavefront algorithm's diffusion mechanism to localize frontier coordinates and labeled each frontier cell to sharpen search efficacy. Extending these ideas, Dai *et al.* [23] sidestepped the processing outlay tied to classical frontier grouping by implicitly clustering frontier voxel representations inside an octree structure and debuted a utility measure blending map entropy with transit duration to assess prospective observation spots.

Furthermore, Deng *et al.* [24] detailed an efficient frontier-driven exploration routine for computationally constrained platforms, demonstrating that slow frontier identification can lead to unnecessary retraversal. A fundamental fragility of orthodox frontier-driven formulations stems from their commitment to myopic policies fixated on maximizing instantaneous

© 2023 The Author(s). | Published by Agro Intelligence Research Press

informational payoff. While this outlook boosts computational speed, it systematically disregards the mission's long-horizon optimality, breeding wasteful to-and-fro displacements and inferior route design. Stated more concretely, such schemes exhibit scant spatial sensitivity to perimeter zones and fail to accord precedence to edge coverage, driving drones to repeatedly crisscross interior expanses before finally attending to boundaries. This behavioral signature appreciably bloats overall traveled distance and mission clock time. A dual-stage planning protocol was proposed by Meng *et al.* [25], in which frontier observation candidates are first sampled, and the globally shortest circuit that threads all such candidates is subsequently computed. In another incremental step, Batinovic *et al.* [26] leveraged Cartographer [27] to assemble OctoMap [28], enabling streamlined frontier aggregation and viewpoint selection across multiple resolutions. The aforementioned solutions nevertheless rely on a greedy calculus that prioritizes near-term reward maximization, which, while aiding computational responsiveness, leads to excessive backtracking and globally subpar exploration. Seeking to transcend this ceiling, FUEL was put forward [5]. That work weaves in incremental frontier identification, synthesizes global coverage paths after frontier aggregation, and re-expresses the global planning challenge as a variant of the traveling salesman problem (TSP) [25]. This design equips FUEL to emit not merely well-ordered exploration sequences but also brisk, flowing flight corridors within a compact temporal window. Subsequent contributions enriched the appraisal of the utility of frontier information [29, 30]. For instance, Zhang *et al.* [31] proposed an efficient Euclidean clustering algorithm for processing frontiers and isolating suitable viewpoints, paired with an approximate trajectory refinement procedure to rapidly compute the optimal tour that traverses all viewpoints. Yet, substantive gaps endure within contemporary frontier-driven frameworks. First, deficient flagging of confined, disjoint pockets: present-day schemes, FUEL included, habitually disregard frontiers lodged in constricted conduits and recessed corners, forcing expensive long-haul revisits that erode exploration economy. Second, path volatility arising from incessant replanning: prevailing architectures trigger wholesale global recomputation whenever new frontiers emerge, inducing hesitant flight behavior marked by jerky heading reversals and speed fluctuations. Although select works attempt to address fluidity via after-the-fact trajectory smoothing [31], they fall short of remedying the underlying unsteadiness bred by recurrent comprehensive replanning divorced from historical continuity.

Autonomous exploration is increasingly conceptualized as a problem of sequential decision-making. Propelled by breakthroughs in deep reinforcement learning, neural network-driven policies have begun matching human competence across an array of domains [32]. Kulkarni *et al.* [33] juxtaposed randomized roadmap pipelines with reinforcement learning formulations, underscoring the latter's edge in path-planning tasks. Ma *et al.* [34] presented a reinforcement learning algorithm with state-chain sequential feedback, streamlining self-directed exploration in static, cluttered, and uncharted environments. Tai and Liu [35] broke ground by applying reinforcement learning to UAVs for autonomous exploration via vision-based sensing. In another thrust, Zhu *et al.* [36] proposed a scheme that leverages graph neural networks coupled with reinforcement learning to address exploration decisions under localization uncertainty. Niroui *et al.* [37] fused frontier-based concepts with deep reinforcement learning, drawing on long short-term memory architectures to identify optimal frontier waypoints for exploration and dynamically modulating frontier selection gain coefficients to improve outcomes. Extending this line of work, Chen *et al.* [38] embedded the self-attention construct into autonomous exploration policies, extracting viable exploration itineraries via a self-attention model scaffold. Analogously, Li *et al.* [39] assembled an autonomous exploration pipeline comprising map-building, decision-making, and planning modules, enabling frictionless integration with broader mapping and navigation stacks in robotics. Adding further nuance, Rao *et al.* [40] incorporated an inverse dynamics model alongside multi-objective co-learning, equipping the visual reinforcement learning agent to infer associations linking navigation-relevant states and objectives while capitalizing on serendipitous positive exposures. Tao *et al.* [41] introduced an exploration architecture that marries deep learning for inferring plausible indoor layouts from accumulated observations with reinforcement learning for productive exploration.

Notwithstanding these achievements, learning-centric paradigms confront foundational hurdles in fielded operation. The foremost constraint lies in substantially compromised resilience and generalizability when encountering unfamiliar settings that deviate from the training distribution. This brittleness becomes especially pronounced within unstructured, intricate real-world locales, where the boundless spectrum of geometrical arrangements and environmental contexts dramatically outstrips the representativeness of any training corpus. Additionally, the processing toll of neural network forward passes, while gentler than training, nonetheless imposes constraints on resource-sparse UAV computation envelopes. These

© 2023 The Author(s). | Published by Agro Intelligence Research Press

factors collectively render purely learned methodologies less suited to the heterogeneous, unforeseeable mapping missions characteristic of operational geospatial surveying and infrastructure evaluation.

Materials and Methods

System overview

In the domain of autonomous UAV-based surveying and remote sensing, the three-dimensional region of interest is conceptualized as an entirely uncharted volume awaiting discovery by the aerial vehicle. The platform relies upon its embedded sensory suite to perceive and catalog observable voxel elements distributed throughout this volume. During exploration, the initially unmapped expanse is progressively classified into two mutually exclusive subsets: navigable void space and solid, obstructed space. The demarcation surface at the juncture between sensed and unsensed zones is termed the frontier, and the exploratory mission reaches its terminus once no such interfaces persist anywhere within the operational envelope. The autonomous exploration workflow is disaggregated into three core functional blocks: frontier maintenance, candidate observation site synthesis, and trajectory computation. The overarching system schematic is rendered in **Figure 1**. We advance a blended observation site synthesis technique that marries obstruction-free spherical sampling with stochastic differential equations, complemented by a frontier visitation ordering rule that accords precedence to perimeter regions and diminutive isolated sectors. This formulation reduces computational cost while simultaneously expanding the breadth of exploratory coverage. In addition, to attain more fluid exploratory flight corridors and suppress vacillatory flight patterns, we construct a dual-phase global planning scheme.

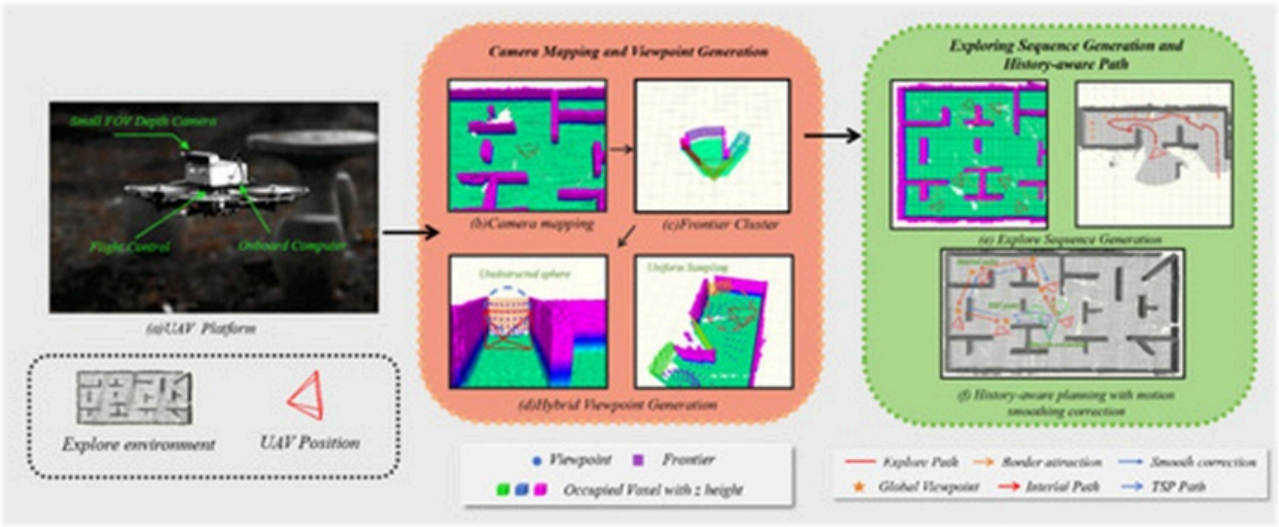

Figure 1. Overview of our system: (a) Our UAV hardware platform, (b) Utilize sensor-perceived environmental information to construct ESDF, (c) Generate frontiers using region-growing algorithm, (d) Generate candidate viewpoints by combining collision-free spheres and uniform sampling, (e) Prioritize exploration of boundaries and small region detection to generate better exploration sequences, and (f) Perform historical-perception planning based on smooth correction.

Hybrid viewpoint generation and decision-making

In most prevailing pipelines [5, 8, 16], observation sites are deployed to scrutinize unmapped voxels within the sensor's detection cone, thereby incrementally extending the exploration frontier. Under typical conditions, no intervening obstacles stand between the observation site and the frontier, facilitating more streamlined boundary elimination. Gauging the visibility of unseen voxels from prospective observation sites, however, typically entails recurrent, computationally burdensome ray-casting passes that substantially erode exploration speed while magnifying processing load. We propose

© 2023 The Author(s). | Published by Agro Intelligence Research Press

a hybrid observation site synthesis technique that extends the coverage of the known frontier within a constrained FOV regime while radically reducing computational demands, as shown in **Figure 2**.

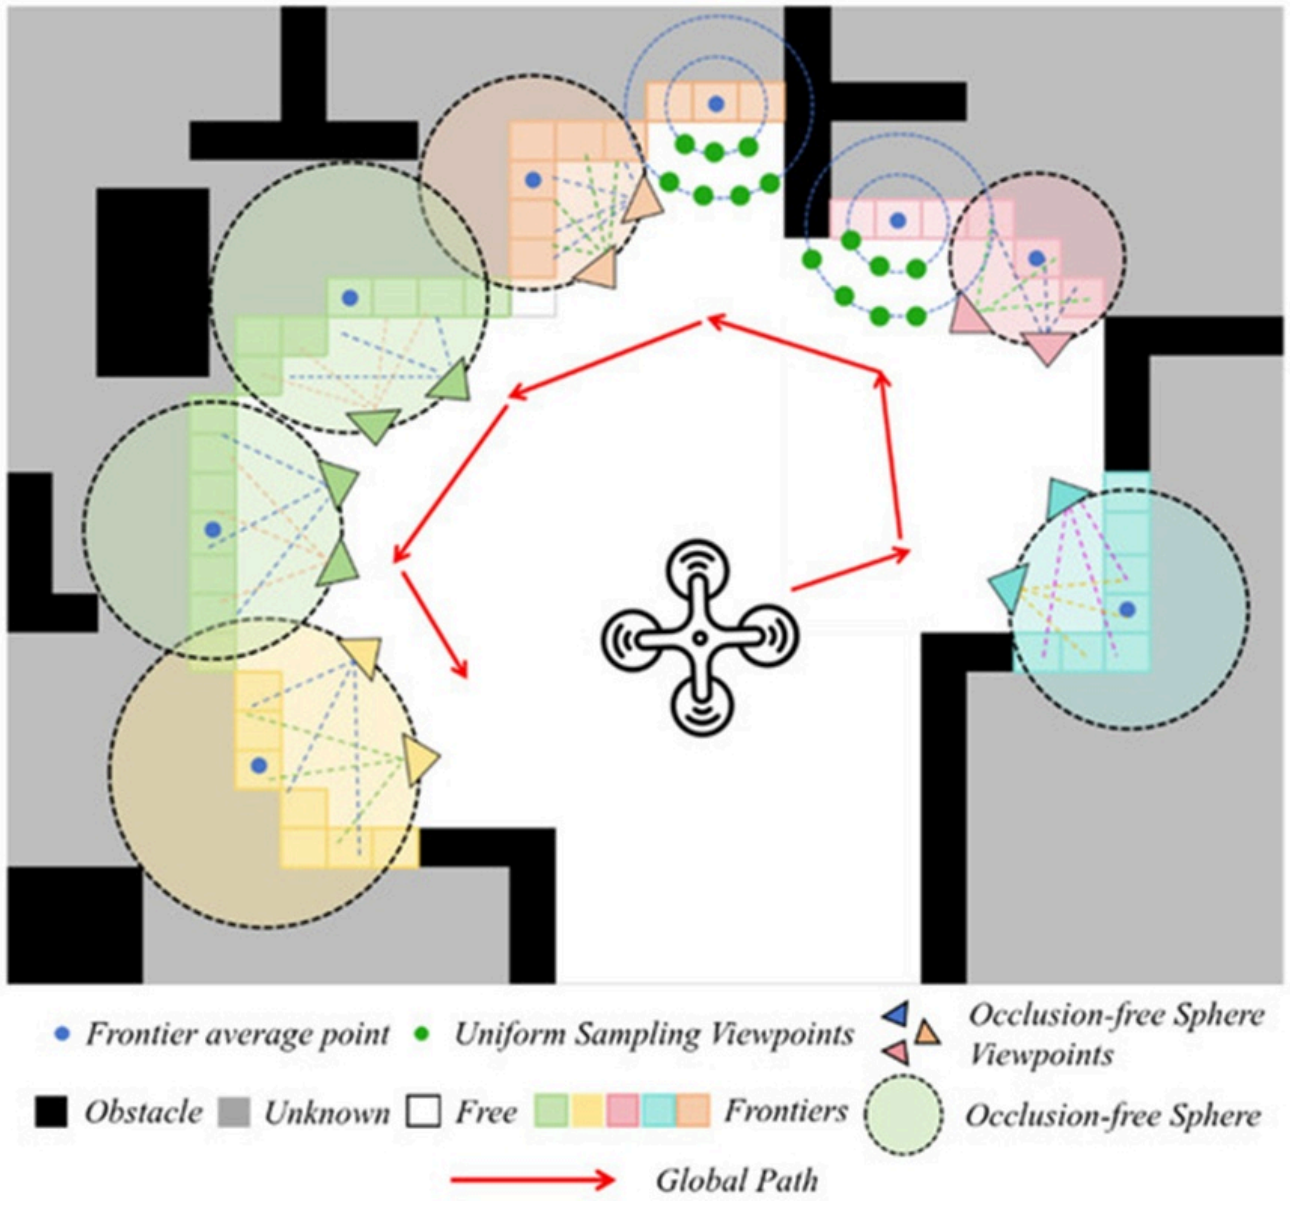

Figure 2. Illustration of hybrid view generation. View generation using a collision-free sphere and adaptive viewpoint generation based on stochastic differential equations.

Collision-free sphere

The obstruction-free sphere takes the geometric form of a spherical body whose centroid $ps \in \mathbb{R}3$ aligns with the center of mass of the designated frontier cluster, and whose radial extent r is expressed as:

$$r = \min\left(\|p_s - p_o\|_2\right) \quad (1)$$

Within which $p0$ designates the closest occupied voxel, and r captures the sphere's radial measure. Given that spheres intrinsically guarantee the absence of obstacles together with convexity [42], every straight path segment originating from

© 2023 The Author(s). | Published by Agro Intelligence Research Press

a position situated on or enclosed by the sphere and extending toward its frontier remains unimpeded. Premium observation sites can be synthesized with considerable efficiency by uniform sampling distributed across the spherical shell, obviating the need for computationally intensive ray casting for each constituent frontier voxel.

Adaptive viewpoint generation based on stochastic differential equations

The obstruction-free spherical tactic provides a computationally efficient alternative to expensive ray-casting methods. Its practical utility, however, is closely bound to the spatial proportions of the operating milieu. In broad, open terrain where the spherical radius is substantial, this technique provides sweeping coverage and proves extraordinarily capable.

In stark opposition, within geometrically elaborate and cramped environments—typified by slender passages and congested chambers—the spherical radius is often severely compressed to uphold a safety buffer from solid obstructions. This results in observation sites being placed extremely close to the frontier cluster boundary. The combined effect of such proximity and the narrow field of view (FOV) inherent in commonplace onboard sensory apparatus routinely produces fragmentary frontier inspection; upright structural features, for instance, may go entirely unobserved, causing the UAV to trace redundant orbits without registering any exploratory advancement, as portrayed in **Figure 3**. Beyond this, uniform sampling on a small-radius sphere yields a profusion of heavily overlapping, low-grade observation site proposals. To independently detect this problematic regime, we deploy a frontier surveillance cube, delineated by the spatial bounding envelope (the maximal and minimal x, y, and z coordinate values) that encompasses each frontier cluster. The cube detection pseudocode is detailed in **Algorithm 1**. Every observation site within this cube is classified as a near-frontier point, indicating the shortcomings of the rudimentary sphere sampler.

Require: Frontier cluster F, sphere center cs, sphere radius rs
Ensure: A Boolean value indicating whether cube behavior is detected
Calculate frontier centroid: $cf \leftarrow mean(F)$
Compute distance: $d \leftarrow \|cs - cf\|$
if $d > rs + \delta threshold$ then
return true
end if
$vsc \leftarrow (cf - cs)/\|cf - cs\|$
for each frontier point p in F do
$vsp \leftarrow (p - cs)/\|p - cs\|$
$\theta \leftarrow \arccos(vsc \cdot vsp)$
if $\theta > \theta max$ then
return true
end if
end for
return false

Algorithm 1. Cube Detection

© 2023 The Author(s). | Published by Agro Intelligence Research Press

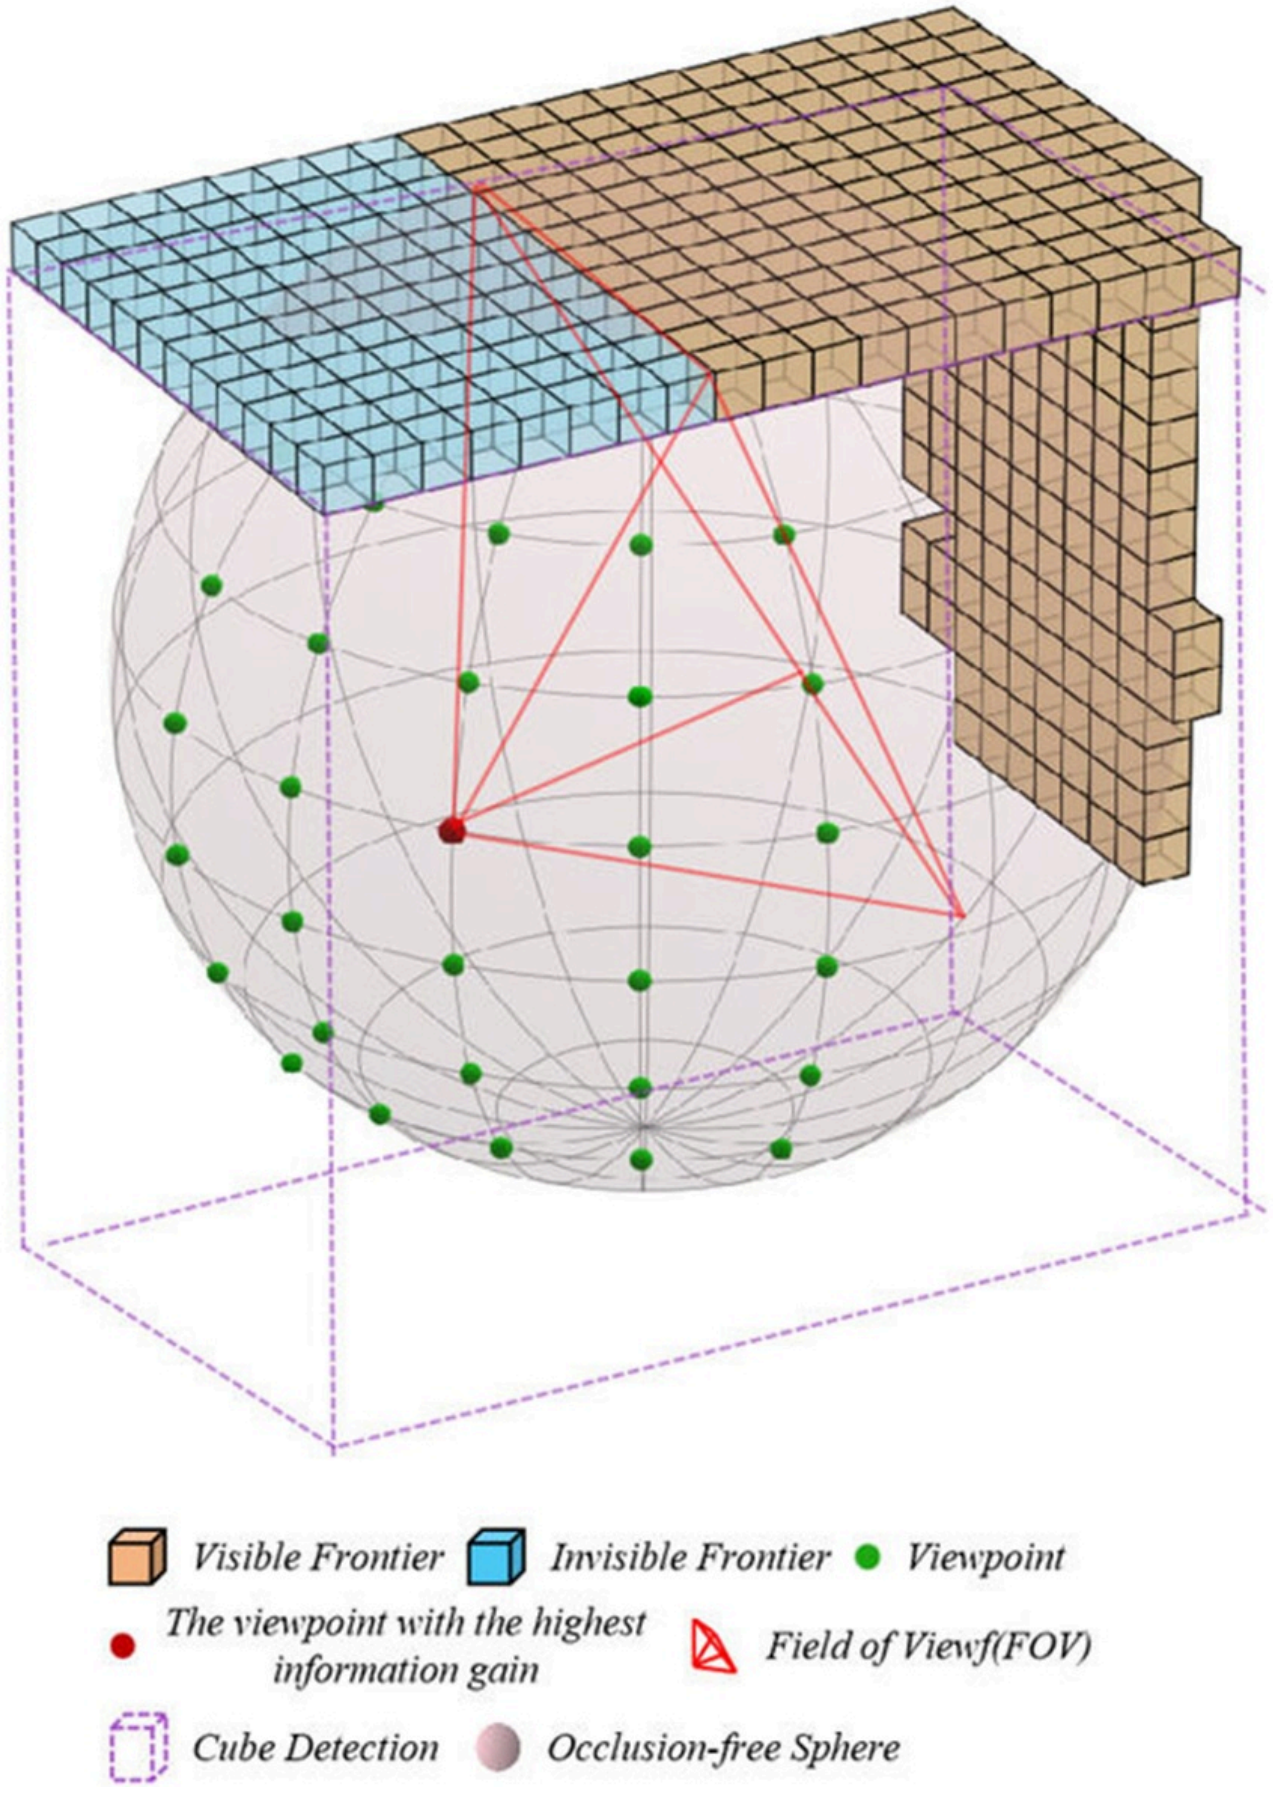

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Figure 3. Visualization of cube detection. A viewpoint on the surface of the sphere cannot fully cover the leading edge, resulting in circling behavior.

To address the complexities of these confined settings, an adaptive observation-site synthesis strategy grounded in stochastic differential equations (SDEs) is introduced. This strategy marks a conceptual pivot, reorienting the paradigm from passive, purely geometric sampling toward an active, information-steered mechanism capable of deftly handling intricate exploration dilemmas.

The heart of this strategy lies in recasting observation site synthesis as a stochastic dynamical process. Departing from uniform sampling conventions, the evolution of prospective observation sites in polar coordinates is governed by a purpose-built SDE. This SDE encompasses both deterministic advection and stochastic dispersion terms, crafting a refined equilibrium. The advection element shepherds observation sites toward sectors of substantial informational payoff by ascending the gradient of the local informational landscape, itself inferred from the spatial configuration of the frontier point set. This ensures that synthesized observation sites are not simply strewn about but are actively directed toward the most auspicious exploratory bearings, such as the entrances to tight corridors. Simultaneously, the dispersion element injects a calibrated amount of stochasticity proportionate to local perceptual ambiguity (quantified as information entropy). This randomized component fosters investigation of alternate pathways within zones of elevated uncertainty, shielding the system against premature convergence to a local maximum—a recurring pitfall of purely grasping heuristics. The SDE apparatus facilitates the generation of a lean yet varied portfolio of premium observation sites. It adaptively modulates sampling intensity, concentrating candidate locations within convoluted, high-return sectors while thinly distributing them across simpler, information-poor expanses.

To begin, the observation site synthesis procedure is formalized in polar coordinates as the interdependent progression of radial displacement $r(t)$ and azimuthal angle $\phi(t)$, orchestrated by the following stochastic differential equation:

$$\begin{cases} dr = \mu_r(r, \phi, \mathcal{I})dt + \sigma_r dW_t^{(1)} \\ d\phi = \mu_\phi(r, \phi, \mathcal{I})dt + \sigma_\phi dW_t^{(2)} \end{cases} \quad (2)$$

Where $\mathcal{I}$ embodies the local informational landscape, and $^{(1)}$, $^{(2)}$ signify uncorrelated Wiener processes. The advection components μr, $\mu\phi$ exploit informational gradients to channel observation sites into information-laden vicinities, whereas the dispersion components σr, $\sigma\phi$ strengthen the capability to delve into unmapped territory via manipulable stochasticity.

In particular, an advection term driven by the informational gradient is engineered to orient the trajectory of observation site evolution along the gradient of the local informational landscape $\nabla\mathcal{I}$. This informational gradient is assessed through the Fisher information array extracted from the frontier point collection.

$$F(x_c) = \sum_{x_i \subset \mathcal{F}} \exp\left(-\frac{\|x_i - x_c\|^2}{2\sigma^2}\right) \frac{(x_i - x_c)(x_i - x_c)^T}{\sigma^2 \|x_i - x_c\|^2} \quad (3)$$

Where xc pinpoints the centroid of the active frontier sector, and F signifies the frontier point collection ensemble. Eigendecomposition is performed on (xc), with the eigenvector corresponding to the dominant eigenvalue yielding the gradient $\nabla\mathcal{I}$, which reveals the bearing of the most rapid local informational accumulation. We structure the radial advection term as the sum of a gradient-pursuit constituent and a basal expansion constituent.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

As the observation site approaches an information-dense sector, such as a constricted channel, the forward projection of the image gradient heading $\nabla\mathcal{I}$, in conjunction with the radial basis vector er, induces a rapid radial compression, yielding concentrated observation site clustering around the passage ingress. On the flip side, throughout low-gradient open terrain, the basal expansion pace $v\exp and$ governs radial elongation, encouraging outward reconnaissance and warding off localized deadlock. Moreover, the angular advection term u adjusts the azimuthal coordinate through alignment with the gradient heading alongside exploitation of a golden-angle spiral scheme.

$$\mu_\phi = \kappa_\phi \nabla\mathcal{I} \cdot e_\phi + \omega_{\text{rotate}}\ , \quad e_\phi = (-\sin\phi, \cos\phi, 0) \tag{4}$$

Within which $wrotate=\theta g/\Delta t$ represents the golden-angle angular velocity, $\overset{\theta g}{\approx} 137.5^\circ$, the gradient alignment constituent channels the observation site toward information-abundant quadrants, and (b) the golden-angle scheme guarantees passability under conditions of negligible gradient signaling.

We formulate the dispersion term as random fluctuations reflecting environmental indeterminacy, expressed through:

$$\sigma_r = \eta_r H(\mathcal{I}), \quad \sigma_\phi = \eta_\phi H(\mathcal{I}) \tag{5}$$

Where $\overset{\mathscr{H}(\mathscr{I})}{= -\sum \mathrm{p}(\mathscr{I}) \log_{\mathrm{p}}(\mathscr{I})}$ captures the localized information entropy, embodying perceptual indeterminacy. Within high-entropy neighborhoods (such as uncharted frontier fringes), the dispersion term amplifies, leading to observation sites that branch away from the dominant gradient to probe complementary routes. By contrast, within low-entropy neighborhoods, the dispersion term naturally attenuates, suppressing non-constructive disturbances.

To conclude, physical constraints are imposed upon the radial and angular variables to satisfy the operational envelope of the sensor:

$$\sigma_r = \eta_r H(\mathcal{I}), \quad \sigma_\phi = \eta_\phi H(\mathcal{I}) \tag{6}$$

Where $r\max$ and $r\min$ align respectively with the furthest and nearest functional distances of the sensory payload, the Euler-Maruyama discretization scheme is applied to integrate the SDE numerically and produce candidate observation site coordinates:

$$x_{k+1} = x_c + r_{k+1}(\cos\phi_{k+1}, \sin\phi_{k+1}, 0) \tag{7}$$

Heuristic frontier cost calculation

After generating a candidate observation site set for each frontier grouping, the sites with the highest evaluation scores are selected as representative nodes that determine the sequential order of frontier visitation. A precise definition of the vocabulary used across this section and throughout the entire document is warranted:

· Frontier refers specifically to the interfacial surface where mapped navigable territory abuts uncharted space within the occupancy grid. These interfaces constitute the fundamental targets that exploration seeks to eliminate.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

· Boundary denotes the material spatial confines of the overall operating environment (for instance, the enclosing surfaces of an indoor space or the walls of a subterranean passage).

An intelligently ordered visitation sequence exerts a direct influence on the cumulative efficiency attained over the course of an exploratory mission. To accomplish this, we formulate a strategy that accounts for both the environmental constraints and the attributes of the frontier regions themselves. The strategy reduces the assessed cost of representative nodes near the environmental periphery. It increases the assessed cost of those located more centrally, thereby coaxing the aerial platform to channel its exploration effort outward toward the margins. In addition, we present an algorithmically efficient method for identifying compact, disconnected frontier segments. The method operates by computing vectors parallel to the frontier to determine whether a particular representative node lies within a detached segment, and assigning such segments the highest possible exploration precedence. Accordingly, throughout this section, the frontier visitation sequencing problem is formulated as the construction of a cost array, and the underlying mathematical treatment is expounded in full.

Environment boundary-aware cost

As demonstrated in **Figure 4b**, after traversing a lengthy flight path to reach the environmental periphery for the first time, the aerial platform deposits residual uncharted pockets in its wake. The practical outcome is that the platform ends up retracing its path toward the periphery at amplified expense after first canvassing interior sectors. To overcome this difficulty, we introduce a margin-focused exploration methodology coupled with a novel margin-cost evaluation procedure, aiming to reduce the considerable overhead associated with peripheral exploration. The procedure commences by computing the shortest distance linking the mean position pka of the FIS within the ensemble Fk to the environmental periphery:

$$d_{\min} = \min(d_{kx}, d_{ky}) \quad (8)$$

Where the quantities dkx, dky capture the shortest measured distances along the X and Y coordinate directions, respectively. Once the separations to the limits in x and y have been obtained, the margin cost is subsequently derived. The underlying principle is that proximity to the periphery yields lower cost, whereas distance from the periphery yields higher cost. The marginal cost is rendered as:

$$C_{bound} = (p_{kay} - (B_{\max y} + B_{\min y})/2)d_{\min}, d_{\min} = d_{kx} \quad (9)$$

$$C_{bound} = (p_{kax} - (B_{\max x} + B_{\min x})/2)d_{\min}, d_{\min} = d_{ky} \quad (10)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

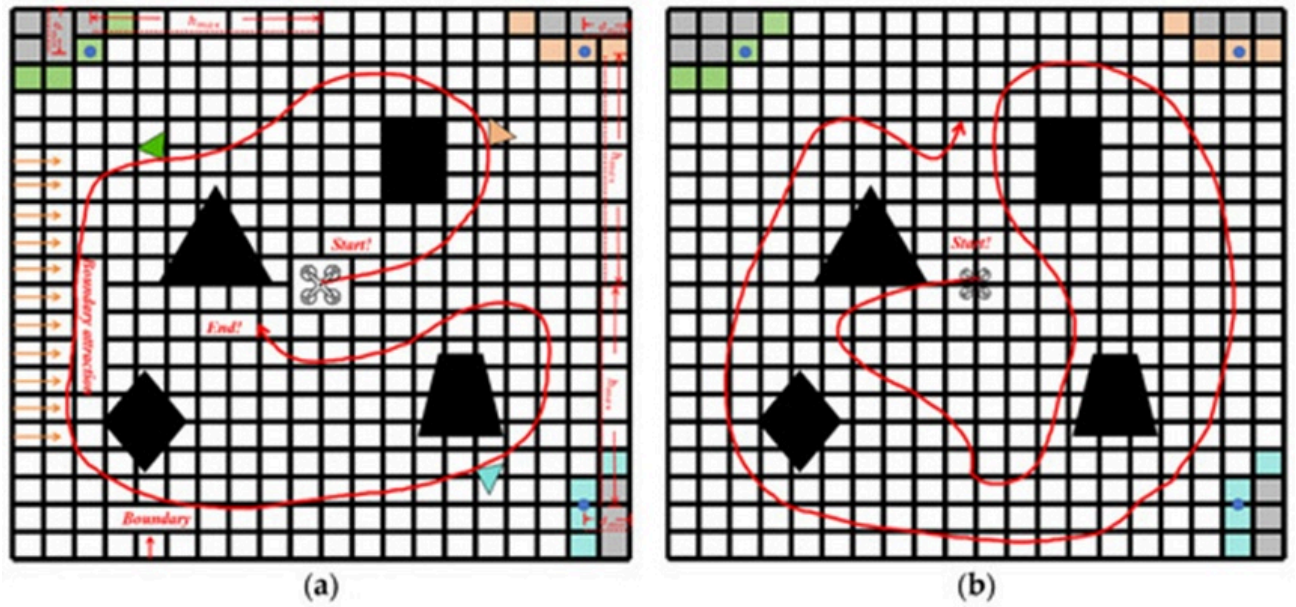

Figure 4. (a) Prioritize exploration of border areas; (b) The exploration sequence of frontiers is not optimized, leading to unnecessarily extended trajectories.

The symbols Bmax and Bmin correspond to the limiting values that demarcate the explorable domain. As shown in **Figure 4a**, the technique we advance enables the UAV to follow a markedly more logical visitation ordering when probing peripheral zones, based on margin costs shaped by synthetic attractive potentials, while sidestepping unproductive long-distance traversal.

Small area cluster detection

To operationalize the principle of small-area prioritization, we devised a detection routine that is uncomplicated, computationally economical, and resource-light, as depicted in **Figure 5**. Within our conceptual scheme, the phrase “edge points” is reserved strictly for the utmost extremity nodes of a frontier grouping when projected onto the coordinate axes. For every three-dimensional frontier grouping, we designate four unique extremity nodes: the node registering the smallest x-coordinate value, the node registering the largest x-coordinate value, the node registering the smallest y-coordinate value, and the node registering the largest y-coordinate value. The directed segments joining the x-extremity nodes (from minimum-x to maximum-x) and those joining the y-extremity nodes (from minimum-y to maximum-y) are referred to as the frontier-parallel vectors. The centroid of the frontier is then displaced by 1 m along the vector that is orthogonal to the frontier-parallel vector, with the resultant coordinate labeled the small-area detection site. Setting out from the frontier centroid and the small-area detection site alike, detection sweeps are performed at a step granularity of 0.1 m, within a radial span of 1 m and 1.5 m directed along the positive and negative orientations of the x and y axes, respectively. The procedure tests whether the detection site is registered to, or lies in proximity to, an occupied voxel. Should occupied or nearby voxels be detected along both the x-oriented and y-oriented sweep directions, the frontier grouping receives classification as a self-contained compact pocket. The penalty value assigned to a self-contained compact pocket is expressed through:

$$C_s = (h_{\max} - h_k) \cdot C_{small} \quad (11)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

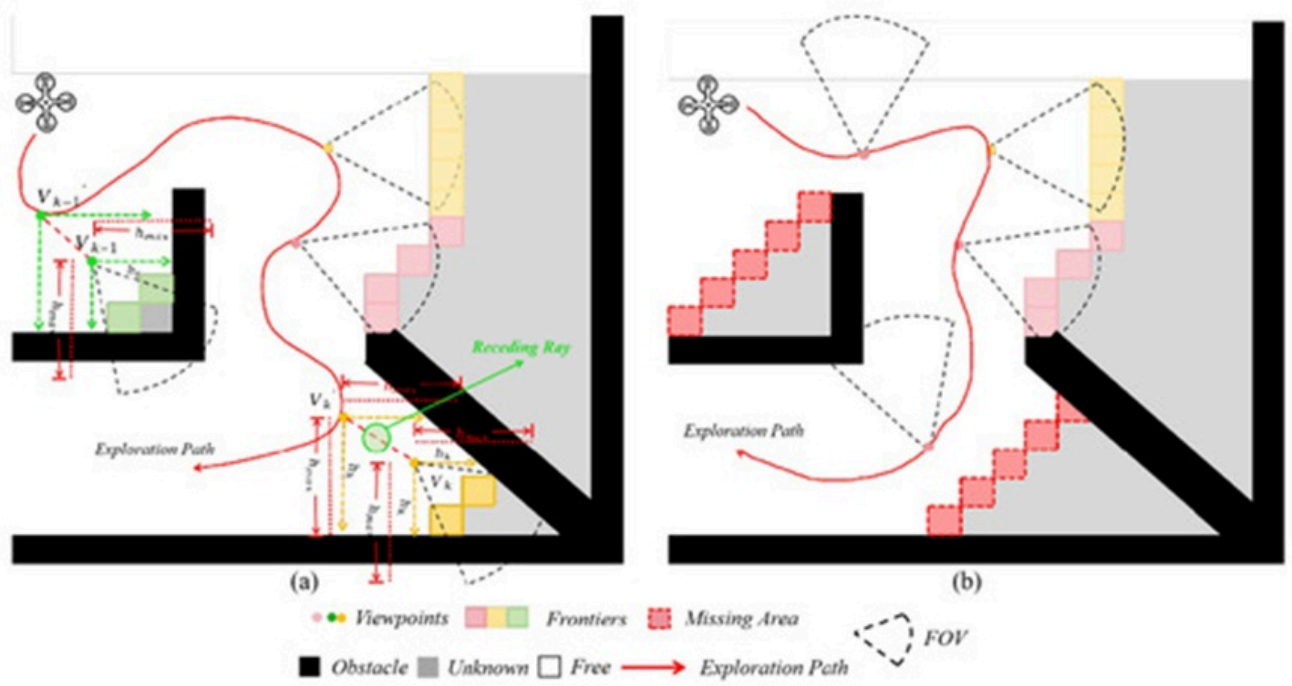

Figure 5. (a) The proposed novel small-area detection approach accurately identifies small areas, and (b) Lack of small area detection results in missing isolated areas.

Within this formulation, $h$max indicates the radial extent of the detection sweep, the metered separation to the occupied voxel is represented by h_k, and C_{small} denotes the prescribed constant penalty magnitude. The detailed implementation details of our small-area detection routine are provided in **Algorithm 2**.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Require: Set of frontier clusters $\{F1, F2, \ldots, Fn\}$, area threshold $Amin$
Ensure: Prioritized small-area clusters S
Initialize $S \leftarrow \phi$
for each cluster Fi do
  $P \leftarrow ConvexHull(Fi)$
  $Ai \leftarrow AreaOf(P)$
  if $Ai < Amin$ then
  $dnearest \leftarrow \min j \neq i \| center(Fi) - center(Fj) \|$
    if $dnearest > disolation$ then
      $S \leftarrow S \cup \{Fi\}$
      $priority(Fi) \leftarrow 1.0/Ai$
    end if
  end if
end for
$S \leftarrow SortByPriority(S)$
return $S=0$

Algorithm 2. Small-Area Detection

Cost matrix construction

As highlighted in the preceding discussion, accurately selecting the next exploratory observation site is critical to streamlining the broader exploration endeavor. The present study specifies a cost array, $Mtsp$, for solving the ATSP and establishing the visitation hierarchy, with the premier frontier groupings as prospective destinations. The transition cost extending from the momentary pose toward each prospective destination is structured in alignment with established conventions [5, 29], augmented by our novel boundary-sensitive penalty components, and takes the form:

$$M_{tsp}(0,k) = t_{lb}(V_0, V_k) + w_c \cdot c_{motion}(V_k) + w_{bound} \cdot c_{bound}(k) - w_s \cdot c_s(k) \quad (12)$$

(1) Path time cost: The present investigation adopts a distance–time penalty as the dominant criterion for ordering the principal frontier visitation. The path time constraint penalty tlb is assessed according to:

$$t_{lb}(V_i, V_j) = \max\left\{\frac{\text{length}(p_i, p_j)}{v_{\max}}, \frac{\min(|\theta_i - \theta_j|, 2\pi - |\theta_i - \theta_j|)}{\dot{\theta}_{\max}}\right\} \quad (13)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

In this notation, Vi designates the premier observation site contained within the frontier grouping, while pi records its spatial coordinates.

(2) Motion consistency cost: To preserve directional continuity throughout traversal, a motion consistency penalty term is incorporated that penalizes frontier groupings whose orientation markedly diverges from the current heading. The motion consistency penalty (Vk) is formally stated as:

$$c_{motion}(V_k) = \cos^{-1}\frac{(p_k - p_{uav}) \cdot v_{uav}}{\|p_k - p_{uav}\| \, \|v_{uav}\|} \quad (14)$$

Two-stage global path planning

To suppress indecisive flight patterns arising from repetitive heuristic boundary-cost refresh operations, we propose a dual-phase global route-construction strategy. During the initial phase, motivated by logged exploration histories [43], a history-conditioned route construction scheme is deployed. This scheme records past visitation orders and synthesizes a global traversal route that incorporates accumulated historical context, empowering the aerial platform to make more robust navigational decisions in geometrically complex settings. To further enhance route fluidity, a motion-smoothing penalty is introduced to suppress abrupt rotational maneuvers. Throughout the subsequent phase, at the juncture where the cost of the history-conditioned route exceeds that of the route derived by solving the Traveling Salesman Problem (TSP), a sliding-window TSP routine is invoked to rebuild the global traversal route. This tactic prevents the platform from becoming fixated on distant peripheral zones while disregarding nearby uncharted pockets, thereby increasing exploration throughput while reducing computational expense.

The framework we propose preserves each global traversal route and, at every replanning instance, constructs a fresh global route informed by the preceding visitation ordering. Given that computations are confined to the altered portions, unaffected segments remain untouched, such that the supplementary processing induces only a minimal growth in computation time. The refresh workflow is illustrated in **Figure 6**. To ensure the visitation ordering of freshly generated observation site groupings aligns with the historically conditioned trajectory, the historical groupings lying nearest to the new groupings are pinpointed, and the separations among them are evaluated. The newly formed groupings are subsequently arranged according to their vicinity to the nearest hub present in the historical trajectory. Across zones that have undergone only modest revision, the freshly formed groupings experience only slight positional adjustments. In general terms, hubs whose positional variation stays beneath d are treated as belonging to the identical frontier grouping as captured in the historically conditioned trajectory, and are termed the reference grouping. The platform's momentary position is likewise subsumed into this reference grouping. To visit every constituent within the freshly generated observation site groupings, they are linked to the reference grouping to yield an unbroken chain. Neighboring reference groupings serve as origin and termination nodes; groupings lying between them are resolved, producing a clearly defined visitation ordering for which the minimum-length traversal route is computed. By chaining these shortest sub-routes that bridge successive reference groupings, a refreshed historically conditioned trajectory is assembled. Finally, the motion consistency penalty, which bridges the platform's current pose to the initial node of the freshly built route, is evaluated to prevent abrupt directional discontinuities that could diminish exploration effectiveness. The historically conditioned trajectory is evaluated as:

$$C = C_{history} + w_{dtr} \cdot C_{dtr} \quad (15)$$

© 2023 The Author(s). | Published by Agro Intelligence Research Press

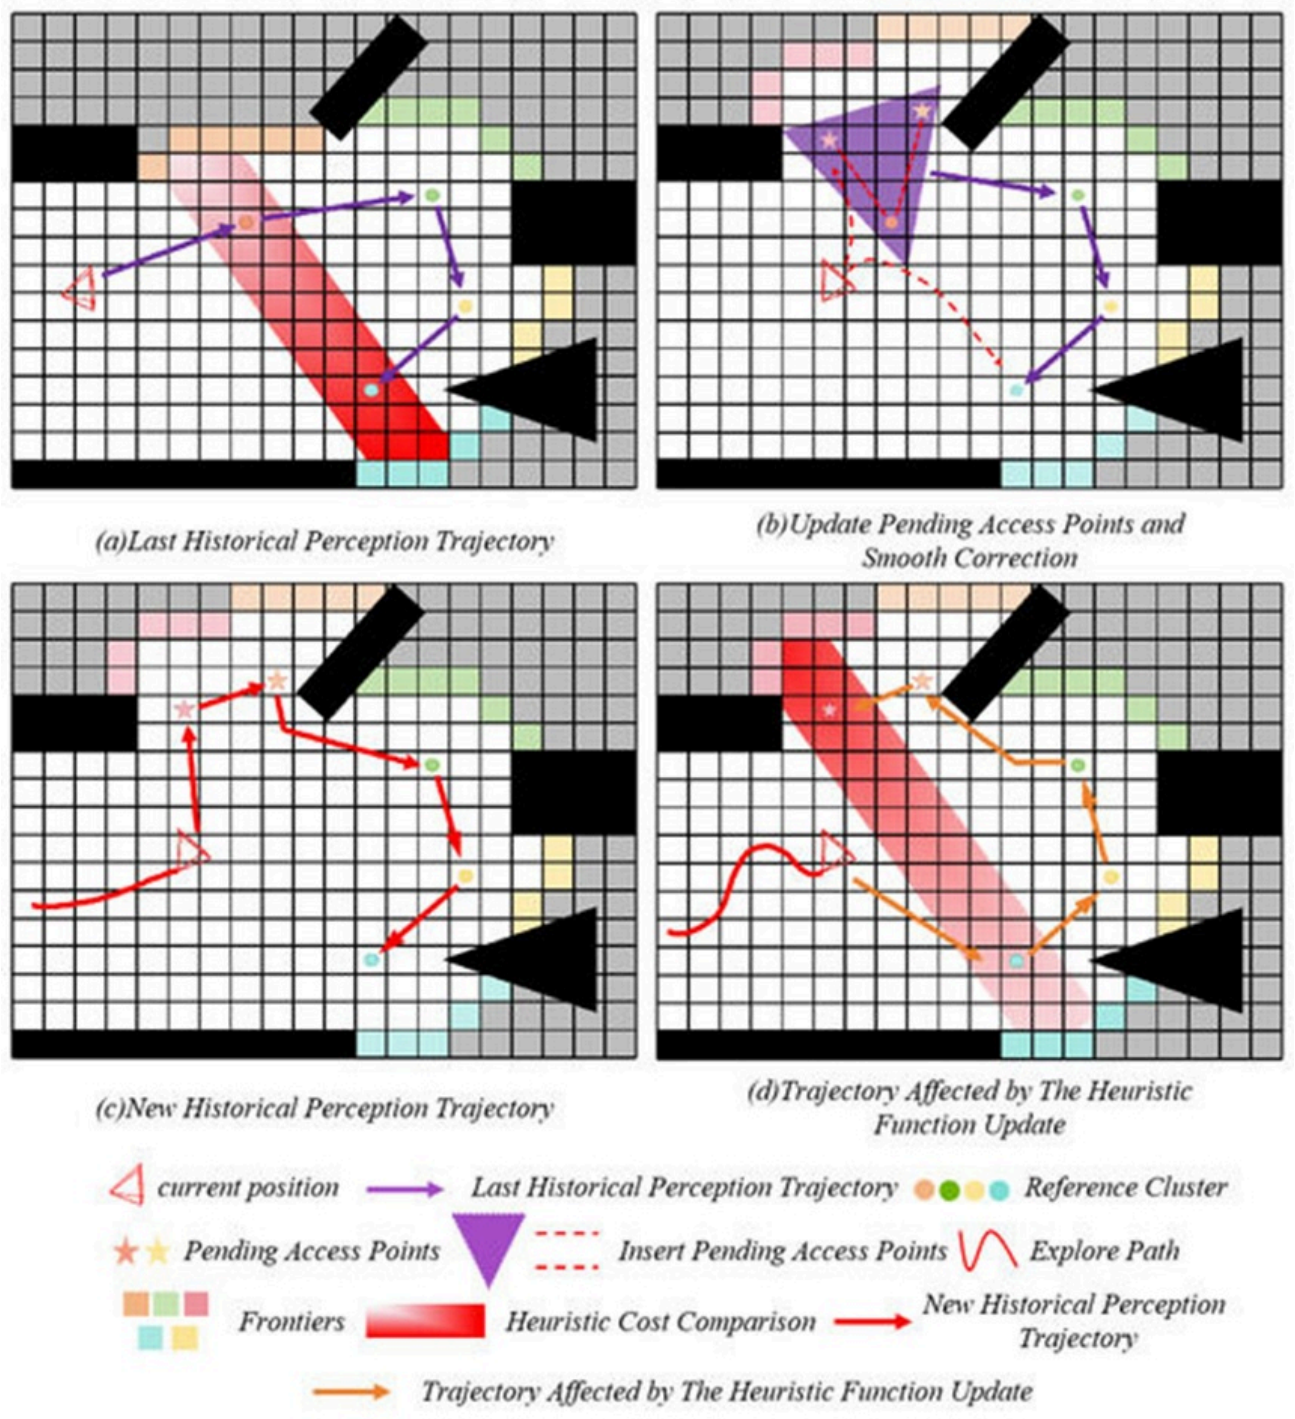

Figure 6. Visualization of historical-perception planning based on motion smoothing correction. Record historical trajectories and determine exploration paths based on smooth optimization costs. Without historical-perception paths, hesitant behavior will occur.

In this expression, $Chistory$ conveys the penalty ascribed to the historically conditioned trajectory, $CTSP$ captures the TSP-derived penalty for the historical ordering, and $Cdir$ represents the motion consistency penalty. The penalty value of the historically conditioned trajectory is weighed against that of the provisional shortest route $CTSP$: should the penalty difference prove negligible, the historically conditioned trajectory persists unaltered; should the difference prove considerable, a sliding-window TSP policy is invoked to reassemble the global traversal route.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

The sliding-window TSP policy is engineered to rapidly produce a near-optimal traversal ordering within a geographically constrained frame of reference. Concretely, it dynamically selects the k frontier groupings nearest to the UAV's momentary position, thereby forming a local optimization window W:

$$W = \{w_1, w_2, \ldots, w_k \mid \| w_i - p_{\text{uav}} \|_2 \leq \| w_j - p_{\text{uav}} \|_2, \forall i < j, w \in \mathcal{F}\} \quad (16)$$

Where $\mathcal{F}$ denotes the complete current ensemble of frontier groupings, and $puav$ marks the UAV's momentary location. The window capacity is fixed at k = 5, which bounds the combinatorial complexity of the resultant TSP instance to (k!). This design parameter guarantees real-time planning responsiveness by focusing computational resources on the most immediately pertinent frontiers, thereby avoiding the myopic choices characteristic of purely greedy heuristics while avoiding the combinatorial explosion of a full-scale global TSP solver. The window shifts forward in lockstep with the UAV's displacement, continuously refreshing the local planning backdrop and enabling logically coherent, extended-horizon exploration within bounded computational outlay per planning interval. The TSP instances generated by the sliding-window mechanism are fed into the Lin–Kernighan–Helsgaun (LKH) algorithm, a notably efficient heuristic tailored for symmetric TSP formulations. The LKH algorithm was selected for its proven ability to reliably deliver near-optimal configurations for moderately sized TSP instances, a characteristic well matched to our requirement for real-time route planning performance.

This policy uses the drone's instantaneous orientation as the coordinate origin and segments the local refinement window based on spatial adjacency. It ensures that the number of frontier groupings falling within the window respects the combinatorial tractability limits of the lightweight TSP engine. Inside this window, an efficient heuristic procedure resolves the optimal traversal ordering over frontier groupings, yielding a locally cost-minimizing traversal route. In the final step, frontier groupings outside the window are appended to form a complete, unbroken global exploration itinerary.

Results and Discussion

The soundness of the presented framework is subjected to rigorous scrutiny through a combination of physical field trials and simulated benchmark evaluations. To disentangle the specific contribution of each algorithmic building block, ablation studies compare the full-fledged framework with stripped-down baseline variants. Within the observation site appraisal procedure, we fix wen = 0.2, wv = 0.8. The heuristic function coefficients are set as wc = 0.5, we = 0.3, ws = 0.2, $Csmall$ = 5 throughout cost array assembly. For history-conditioned planning, $wdir$ = 0.5. With respect to SDE-governed observation site synthesis, the coefficient suite is arranged in the following manner: radial drift magnitude μr = 0.1, angular drift magnitude $\mu\theta$ = 0.05, radial diffusion magnitude σr = 0.2, angular diffusion magnitude $\sigma\theta$ = 0.1, and temporal discretization granularity Δt = 0.1. The sensory payload is configured to span an angular envelope of [80 × 60] deg, paired with a perceptual radius of 4.5 m. The NLopt library is tasked with solving the trajectory-refinement formulation, and the flight corridor is parameterized by a 3rd-order B-spline curve. To enable swift autonomous charting, the Fiesta framework is used to construct an occupancy network representation [44] incrementally. Alongside this, the ESDF map receives continual updates to furnish support for trajectory computation. Our simulation-based experimentation is performed on a laptop equipped with an Intel Core i5-12600KF processor, 32 GB of system memory, and the Ubuntu 20.04 operating system. Every software routine is written in C++.

Benchmark comparisons

We situated our presented framework in a comparative context alongside several reference methodologies, namely the classical frontier methodology [19], the rapid frontier methodology [21], ECHO [29], and FUEL [5]. A trio of archetypal simulated domains served as the evaluative arena: an office interior layout, a densely wooded outdoor scene, and a DARPA subterranean passage. The behavioral characteristics of these methodologies were scrutinized through

© 2023 The Author(s). | Published by Agro Intelligence Research Press

qualitative observation across the varied domains and appraised quantitatively via metrics capturing exploration clock time, traversed route length, and spatial volume mapped.

The qualitative head-to-head assessment against state-of-the-art methodologies provides an intuitive sense of the presented framework's practical merits. As shown in **Figure 7**, our framework achieves the fastest exploration wrap-up rate among the compared methodologies and the lowest peak duration across disparate unmapped settings. **Figure 8** further demonstrates that our framework navigates all three domains while tracing the most parsimonious, continuously flowing flight corridors. Particularly noteworthy, our framework correctly identifies and prioritizes detached compact pockets, ensuring their timely coverage while exhibiting virtually no backtracking behavior. In stark divergence, both FUEL and ECHO exhibit recurrent retraversal of previously charted zones. This tendency materializes because, when operating within unmapped environments, poorly calibrated heuristic cost appraisals habitually spawn myopic exploration conduct that neglects detached pockets offering only modest informational payoff and precipitate inefficient long-haul flight excursions. Set against these behavioral patterns, our presented boundary-cost expression and detached-pocket recognition mechanism—which affix heuristic cost penalties onto remote and isolated pockets—manifest substantial superiority, tangibly corroborating the operational value of our heuristic boundary-cost formulation.

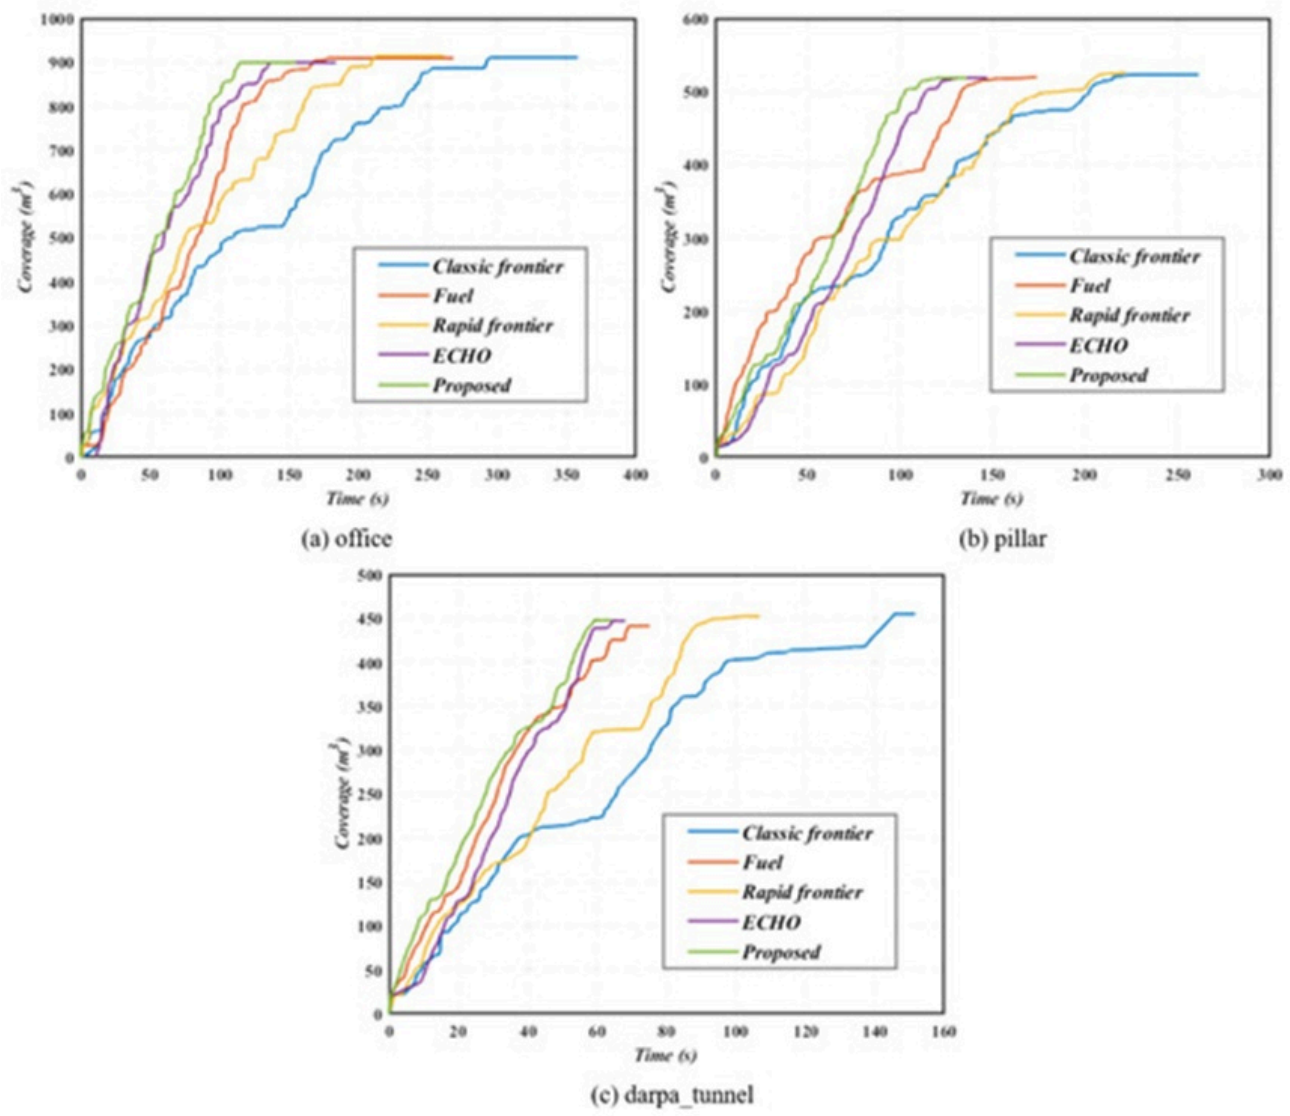

Figure 7. Comparison of exploration processes.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

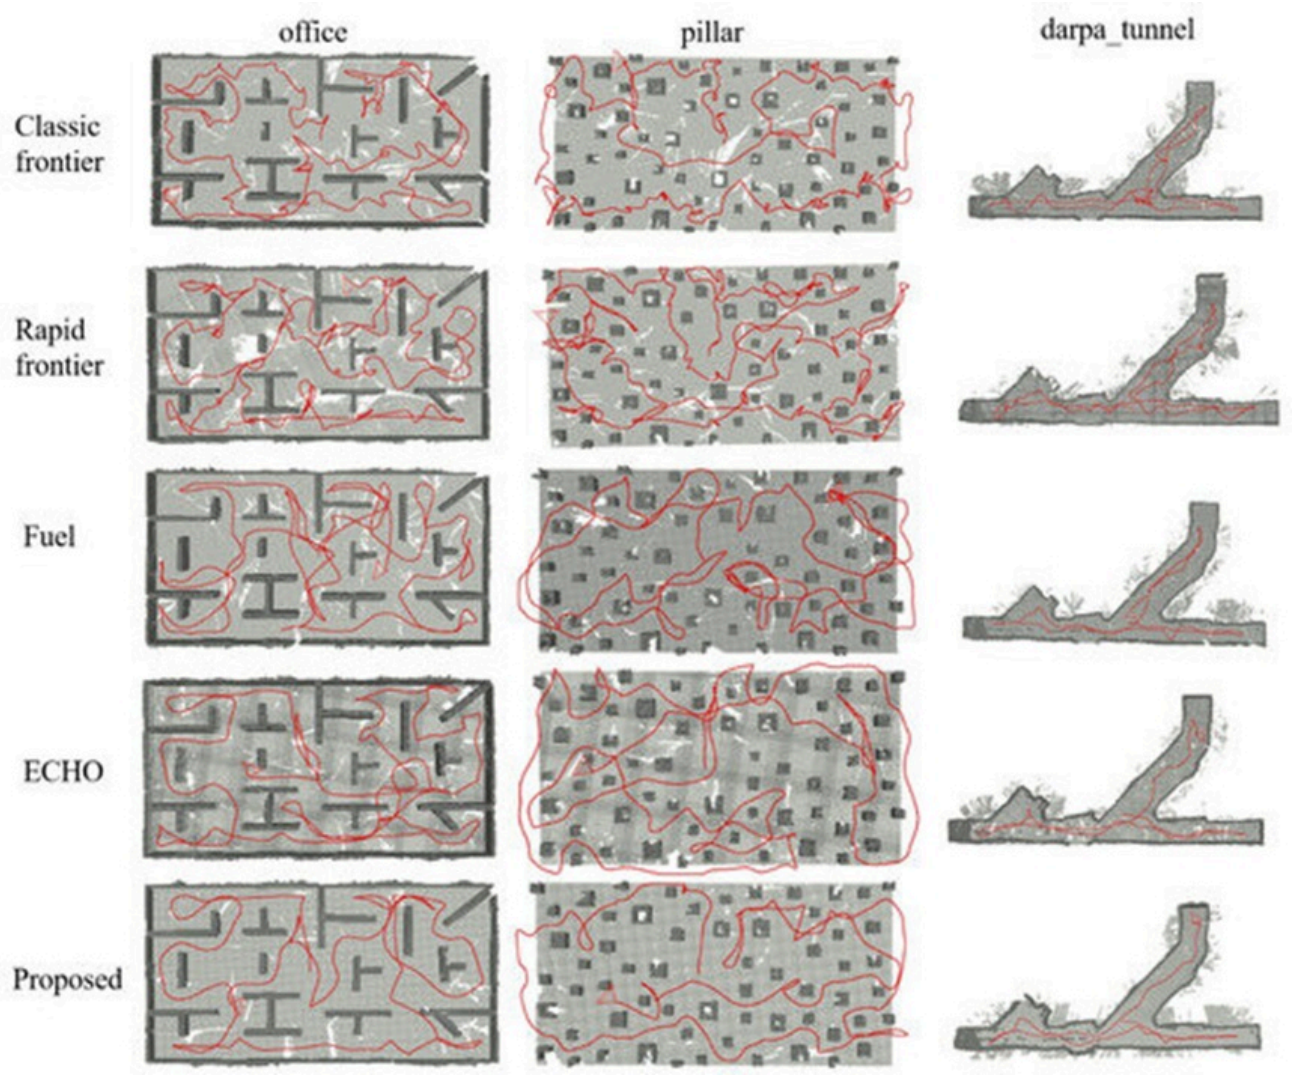

Figure 8. A comparative analysis of the proposed approach and other approaches in various simulated environments for path exploration.

Our quantitative experimentation compared the presented framework with multiple leading-edge methodologies with respect to exploration clock time, the length of the traversed route, and the spatial volume charted. The statistical measures of central tendency, dispersion, uppermost value, and lowermost value were marshaled as evaluative criteria to gauge the caliber of each exploration methodology. **Table 1** summarizes the outcomes from our framework and the baseline methodologies across the three contrasting simulated domains. It is pertinent to observe that the “Coverage (m^3)” indicator quantifies the cubic volume of the probed expanse, as estimated through the voxel-based occupancy representation. It reflects the tangible physical bulk of the environment that has been successfully charted, as validated against the ground-truth spatial boundary. Viewed holistically, our framework produced the shortest exploration clock time across all domains, indicating superior exploration throughput. This enhancement originates principally from the presented hybrid observation site synthesis strategy. In situations where the distance to the nearest obstruction is considerable, the collision-free spherical sampling technique sidesteps the need for burdensome ray casting to verify obstacle-free clearance between the aerial platform and candidate observation sites. Drawing observation sites along the spherical exterior shell provides all-encompassing perceptual coverage while preventing collisions. Even more crucially, we incorporate an adaptive observation site synthesis technique rooted in stochastic differential equations. This technique dynamically interlocks the progression of observation sites with the local informational topography, generating a sparse yet potent collection of high-information-yield candidate observation sites. As the observation site draws near information-

© 2023 The Author(s). | Published by Agro Intelligence Research Press

dense neighborhoods, abundant, tightly clustered, premium-grade candidate viewpoints are spawned; reciprocally, throughout information-impoverished neighborhoods, dense sampling is dampened to suppress superfluous redundancy. This machinery substantially reduces the computational cost of observation site synthesis while simultaneously providing premium-grade, information-rich observation sites. Beyond this, our framework consistently exhibits lower standard deviation, showing noticeably damped fluctuations compared with state-of-the-art methodologies. This improvement stems from the fact that every exploration run confronts idiosyncratic environmental configurations and decision-making dilemmas that can undermine exploration stability. By adopting a dual-phase global route construction strategy, our framework enables more decisive observation site selection, ensuring stable and resilient behavior across multiple trial repetitions.

Table 1. Exploration statistics in the three scenarios.

		Coverage (m3)				Path Length (m)				Exploration Time (s)			
Scene	**Approach**	**Min**	**Max**	**Std**	**Avg**	**Min**	**Max**	**Std**	**Avg**	**Min**	**Max**	**Std**	**Avg**
Office	Classic frontier	**903**	**916**	3.2	**909**	217	254	16.7	235	259	300	22.3	271
	Rapid frontier	900	911	**1.7**	908	279	312	11.3	298	205	230	7.0	217
	FUEL	898	911	2.1	905	208	225	5.6	217	144	174	2.0	159
	ECHO	895	910	3.1	905	189	207	3.1	198	120	148	2.8	126
	Proposed	899	912	2.3	907	**166**	**188**	**2.7**	**183**	**113**	**140**	**1.8**	**118**
Pillar	Classic frontier	**517**	**523**	**1.6**	517	194	224	5.3	200	248	261	12.5	250
	Rapid frontier	216	521	2.0	518	237	269	12.1	258	208	224	11.8	213
	FUEL	515	522	2.3	517	183	240	14.1	212	141	184	7.8	166
	ECHO	516	521	2.1	517	182	215	7.3	194	130	163	7.0	142
	Proposed	517	521	1.9	**519**	**158**	**185**	**4.8**	**175**	**119**	**150**	**5.2**	**133**
DARPA Tunnel	Classic frontier	**444**	**461**	3.1	**450**	120	130	4.2	124	135	149	9.7	140
	Rapid frontier	440	449	1.7	445	125	140	5.1	131	90	108	6.3	100
	FUEL	429	439	**0.9**	435	101	120	3.9	111	59	76	1.8	68
	ECHO	440	450	1.3	443	98	117	3.7	108	58	70	1.7	62
	Proposed	441	456	1.1	447	**93**	**107**	**3.3**	**101**	**56**	**63**	**1.5**	**58**

Ablation studies

© 2023 The Author(s). | Published by Agro Intelligence Research Press

The following segment details an ablation examination aimed at uncovering the relative importance of the three foundational pillars underpinning our presented framework: hybrid viewpoint generation and decision-making (HFGDC), heuristic frontier cost calculation (HFCC), and two-stage global path planning (TGPP).

To demonstrate the efficacy of the presented HFGDC technique, we compared the temporal cost of identifying candidate observation sites under conditions with and without HFGDC active within an unchanging exploration setting. The outcomes, visualized in **Figure 9**, make clear that HFGDC substantially outstrips uniform sampling in terms of temporal efficiency while simultaneously demonstrating compressed variability. During operation, HFGDC assembles a collision-free spherical volume and uniformly samples points from its exterior envelope. Furthermore, using cube and compact-sphere recognition, whenever frontier sectors or small obstructions are identified, an adaptive observation site synthesis strategy grounded in stochastic differential equations is invoked. This strategy dynamically couples the progression of observation sites to the local informational topography, yielding a lean set of candidate observation sites with elevated informational yield.

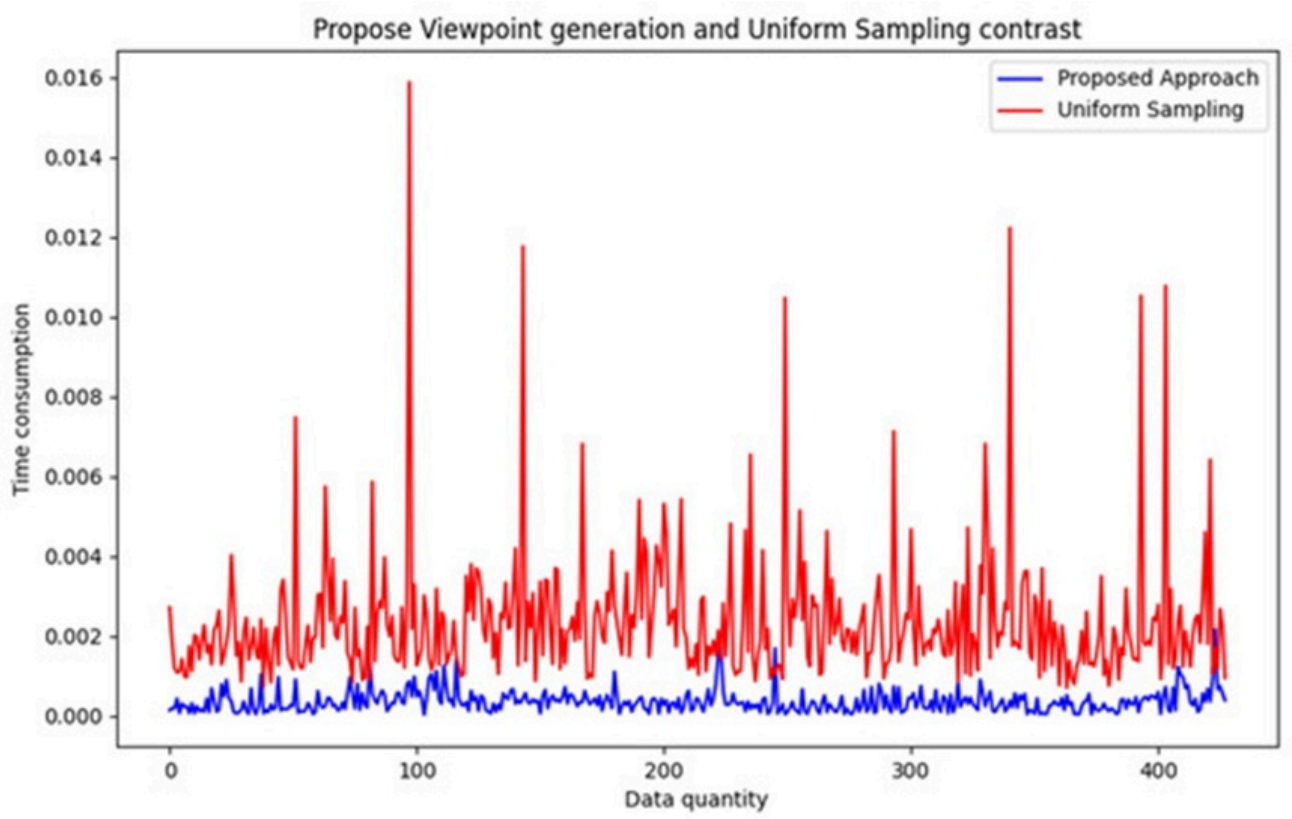

Figure 9. Comparison of the time consumption of our proposed viewpoint generation approach and the uniform sampling approach.

To more thoroughly scrutinize the resilience of our framework, we examined the impact of key parameters in the SDE-driven observation site synthesis module. **Table 2** collates the behavioral profile under varied parameter configurations within the Office domain. The outcomes bear out that our designated parameters (σr = 0.2, Δt = 0.1, μr = 0.1) realize a commendable middle ground. Depressed magnitudes of the radial diffusion coefficient σr yield inadequate exploratory stochasticity, stretching traversed route length as the planner becomes ensnared within local pockets. At the other end, elevated magnitudes degrade stability, as reflected in the inflated standard deviation of exploration clock time. The drift coefficient μr adeptly steers the compromise between exploration breadth and refinement depth, with our settled-upon magnitude producing the briefest exploration clock time.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Table 2. Analysis of SDE parameters in the office environment.

Configuration	Stability (Std. Time)	Coverage (m^3)	Path Length (m)	Exploration Time (s)
σr=0.2,Δt=0.1,μr=0.1 **(ours)**	2.3	907	190	128
Noise Level (σr)				
σr=0.1	4.2	902	205	142
σr=0.3	6.3	904	198	135
Step Size (Δt)				
Δt=0.05	3.1	905	196	136
Δt=0.2	4.5	903	201	138
Drift Coefficient (μr)				
μr=0.05	3.4	904	195	134
μr=0.15	3.8	906	193	132

In the next stage, we binarized the ESDF occupancy grid representation to furnish a more immediately graspable depiction of the advancement of uncharted space exploration. Beyond this, several zones within the office domain that are routinely susceptible to being bypassed were singled out for behavioral comparison between conditions with and without the HFCC module engaged. As illustrated in **Figure 10**, HFCC deploys a computationally lightweight yet exacting small-area recognition scheme, facilitating timely identification and probing upon the very first arrival within the target zone, thereby curtailing duplicate exploration effort. In marked contrast, without the HFCC module, several crucial zones went undetected, leading to repeated, entirely unnecessary exploration passes.

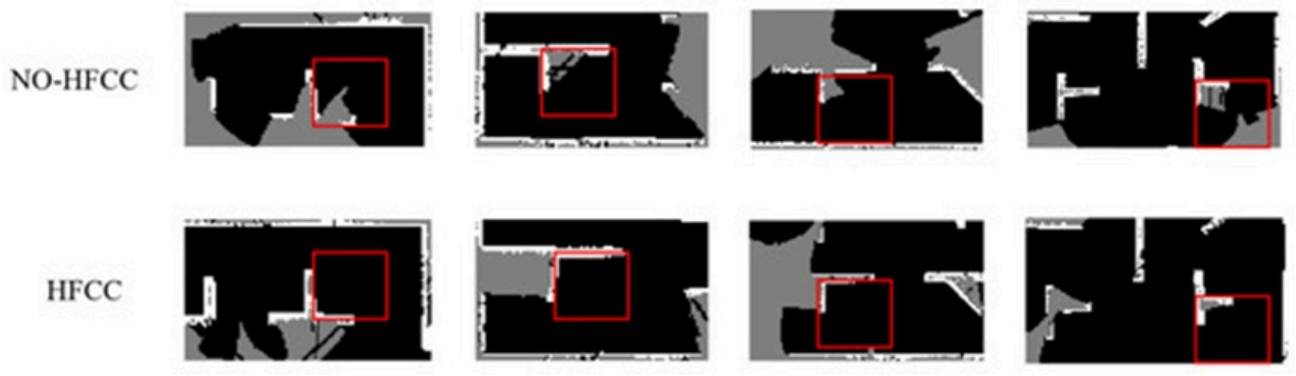

Figure 10. Comparison of the performance of our proposed approach in small area detection.

To quantify the computational throughput of our sliding-window TSP technique, we compare it with alternative planning schemes along two axes: theoretical complexity bounds and empirically measured wall time. **Table 3** presents a thorough comparative dissection conducted within the Office domain.

Table 3. Computational complexity and performance comparison of planning strategies.

Planning strategy	Search duration (s)	Computational complexity

© 2023 The Author(s). | Published by Agro Intelligence Research Press

One-shot TSP method	150	$O(n!)$
Stepwise incremental planner	156	$O(n^2)$
Proposed sliding-window TSP	140	$O(k!)$

The outcomes clearly show that our sliding-window TSP formulation delivers the most favorable computational profile while simultaneously safeguarding leading exploration throughput. The single-stage TSP variant, despite its theoretical guarantee of global optimality, becomes computationally unmanageable even with relatively modest frontier populations. The incremental planning alternative, while frugal in its computational appetite, succumbs to nearsighted decision tendencies that swell exploration clock time through wasteful retracing of previously covered ground. Our formulation converges to the optimal equilibrium point, yielding near-optimal localized itineraries with a computational footprint that remains invariant regardless of the number of frontiers in the global set.

An even-handed assessment of the presented framework's impact on exploration effectiveness was conducted by juxtaposing algorithmic behavior under uniform testing conditions, with each of the three cornerstone modules individually switched off. **Table 4** records the exploration clock time, route distance, and mapped volume for each arrangement. In the baseline condition where none of the modules are activated, the algorithm turns in a markedly weak showing, weaving the most disorderly and drawn-out flight corridors. This regression finds its primary roots in the steep computational levy imposed by substandard candidate observation sites, a jumbled visitation hierarchy that spawns incessant backtracking, and superfluous heuristic recomputation cycles that translate into wavering flight patterns. The HFGDC module, by virtue of its deployment of a blended observation site synthesis tactic that yields a measured quantity of premium-grade candidate observation sites, meaningfully compresses computational outlay while boosting both route economy and exploration pace. **Figure 11** captures the corresponding exploration corridors. Engaging the HFCC module provides a further boost to system behavior by contributing a refined frontier cost expression in tandem with a small-area identification apparatus, which together prioritize frontier peripheries and detached pockets, tangibly reducing wasted traversal and improving temporal efficiency. The configuration running TGPP alone delivers a visibly smoother flight corridor with barely perceptible hesitation, a trait directly linked to the presented dual-phase global route construction architecture. The source of this gain lies in the architecture's foundational tactic of threading historically accumulated trajectory data into the replanning loop. Through the preservation and strategic reuse of earlier exploration sequences via the history-sensitive refinement mechanism, the system attenuates the motion perturbations and abrupt orientation shifts that ordinarily accompany repeated heuristic boundary refresh cycles. This fosters more unified navigational choices and reinforced flight composure, even in the absence of the remaining presented modules.

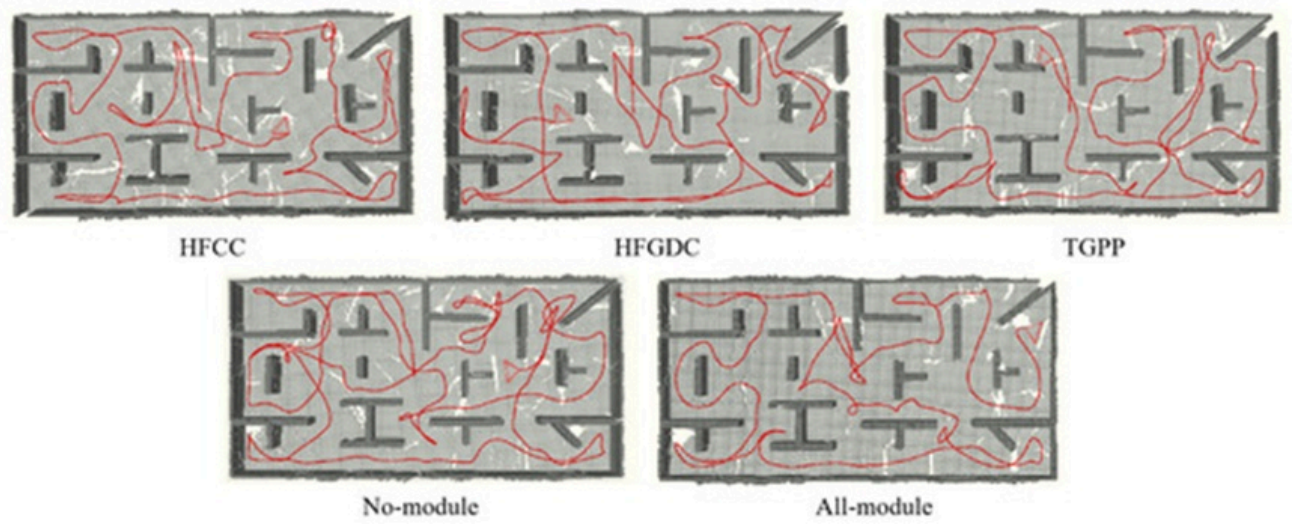

Figure 11. Exploration paths of ablation experiments, testing our approach without the HFCC module, HFGDC module, and HPP module.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

Table 4. Results of the ablation study.

Scene	Approach	Coverage (m^3)			Path length (m)			Exploration time (s)		
		Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
Office	Only HFGDC	903	911	907	184	196	190	123	144	128
	Only HFCC	900	911	908	175	194	185	125	146	134
	Only TGPP	897	910	906	191	210	200	130	150	140
	No-Module	898	911	905	198	237	217	144	174	157
	All-Module	899	912	907	166	188	**183**	113	143	**120**

In summary, the synergistic integration of all modules enables our algorithm to craft an effective, cohesive flight corridor, expediting the survey of unmapped expanses while minimizing wasteful return traversals.

Real-world experiments

To systematically benchmark the presented framework's capability across demanding real-world landscapes, a series of aerial sorties was conducted over three emblematic settings: a confined interior chamber, an exterior wooded grove, and a below-grade parking facility. These landscapes, marked by graduated levels of structural convolution, perceptual impediments, and piloting rigor, were deliberately curated to exercise the system's flexibility, tenacity, and in-flight poise from numerous complementary angles. The trial apparatus centered on a self-piloting multirotor airframe, as illustrated in **Figure 1**. The airframe was outfitted with an NVIDIA Jetson Orin NX edge processing unit functioning as the chief computational core, delivering formidable parallel throughput for high-bandwidth real-time sensing, spatial modeling, and live global route synthesis. Additionally, an Intel RealSense D435i RGB-D depth-sensing camera was tasked with acquiring temporally aligned depth streams, coupled with inertial measurement data. The globally synthesized itineraries were passed to the airframe's integrated Pixhawk autopilot module (running the PX4 firmware suite with its default MPC-based positional regulation loop), which oversaw real-time path-following and attitude stabilization.

To ensure the repeatability and neutrality of the findings, a minimum of 5 flight repetitions were performed in each test landscape. All sorties were executed under identical dynamic motion constraints, including control parameters such as peak translational velocity, rotational rate, and acceleration bounds, which were matched to those enforced throughout the simulation trials. The flight objectives tasked the system with autonomously canvassing unmapped volumes and delivering exhaustive environmental sweep and spatial reconstruction. From takeoff through mission termination, the entire apparatus operated without any external localization framework or human piloting, achieving completely autonomous, closed-loop functionality predicated solely on embedded sensing and computation. As depicted in **Figure 12**, the presented framework skillfully guided the aerial platform to achieve a brisk, productive area sweep across entirely uncharted and intricate terrain. The produced flight corridors showed balanced spatial dispersion, with negligible overlap or duplicated coverage, thereby substantiating the fielded effectiveness of the hybrid observation site synthesis and boundary-prioritization policies. The coverage accumulation profiles traced a reliably monotonic climb, indicating robust motion cohesion and planning reliability sustained throughout the full exploration window. Most revealingly, in situations involving occlusions or semi-enclosed alcoves, the exploration policy resolutely advanced into unsensed territory, thoroughly underscoring the presented framework's advantage in practical autonomous reconnaissance, detection, and cartographic reconstruction missions spanning interior, exterior, and subsurface domains.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

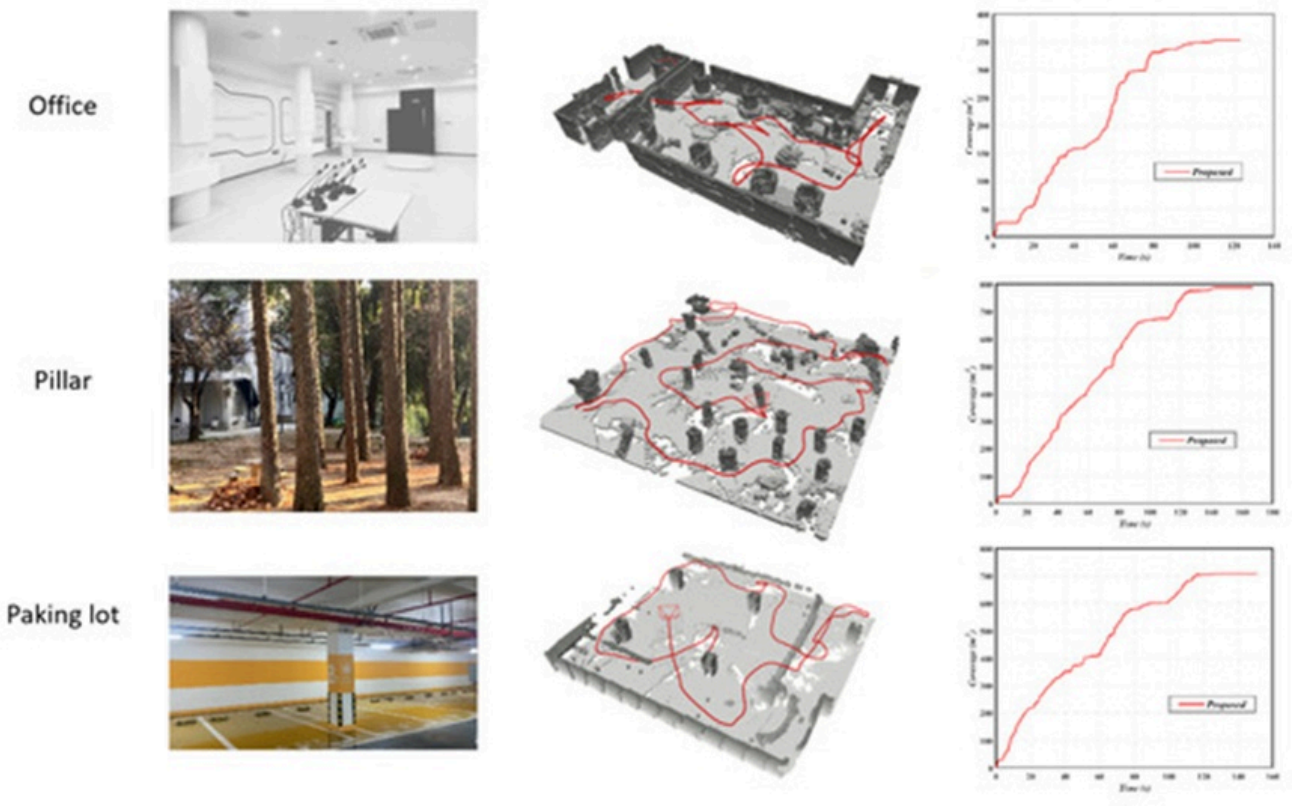

Figure 12. Exploration results in real-world experiments.

Conclusion

With the present body of work, we aspire to channel the broader research community's scrutiny toward the crafting of a layered exploration architecture expressly conceived for UAVs hampered by restricted fields of view. Constituting an original answer to the challenge of self-directed exploration within unmapped environments, the presented framework equips aerial vehicles to sense, investigate, and chart geometrically complex and unstructured settings with tangible proficiency. While established exploration techniques remain capable of autonomous detection and mapping, they routinely labor under the burdens of inefficiency and oscillatory instability. Throughout the present study, three defining design pillars underpin the encouraging outcomes secured. First and foremost, we espouse a blended tactic that fuses collision-free spherical sampling with adaptive observation site synthesis, springing from stochastic differential equations, thereby yielding premium-grade observation site proposals whilst keeping computational expenditure firmly in check. In the second instance, we advance a heuristic scoring function that prioritizes detached compact pockets, thereby paving the way for superior exploratory itineraries. Completing the set, we recast the global sweep itinerary as a Traveling Salesman Problem (TSP) and marshal a dual-phase global planning architecture to dampen motion perturbations kindled by unremitting heuristic updates, thereby ameliorating trajectory cohesion and in-flight stability. We hold the conviction that the presented architecture provides persuasive experimental substantiation and pragmatic engineering guidance for eventual transition into applied settings such as metropolitan cartography, post-disaster damage assessment, and autonomous surveillance rounds. While the existing investigation demonstrates stalwart behavior across a range of static settings characterized by elaborate geometry, including tunnel complexes and timbered expanses, supplementary inquiry remains imperative to address obstacles such as mobile impediments and fluctuating illumination regimes. Moreover, the fusion of a tightly coupled, resilient tracking regulator, purpose-optimized for directives issued by our planning module, represents a logical and pressing next step toward further augmenting holistic system efficacy. These facets, alongside the infusion of multi-UAV cooperative frameworks, dynamic obstacle negotiation routines, and semantic scene

© 2023 The Author(s). | Published by Agro Intelligence Research Press

comprehension modules, chart significant avenues for our forthcoming research aimed at enhancing the system's expandability and cognitive sophistication across increasingly demanding operational contexts.

Acknowledgements

None

Conflict of interest

None

Financial support

This research was funded by Hubei Provincial Technology Innovation Plan Project [grant number 2024BCB103] and Guangxi Science and Technology Major Project [grant number AA22068072].

Ethics statement

None

Received: 14 Jun 2022 Revised: 19 Aug 2022 Accepted: 05 Sep 2022
Published online: 25 November 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Abd Rahman MF. Technological innovations in UAVs for geospatial data collection: A review of current trends. 2025. https://www.researchgate.net/publication/388325527_Technological_Innovations_in_UAVs_for_Geospatial_Data_Collection_A_Review_of_Current_Trends (accessed on 3 December 2025).

2. He S, Tan X, Zhong Y, Huang M, Mei Z, Wan Y, et al. Evolutionary PSO-based emergency monitoring geospatial edge service chain in the emergency communication network. Int J Digit Earth. 2023;16:2797-817.

3. Martinez-Alpiste I, Golcarenarenji G, Wang Q, Alcaraz-Calero JM. Search and rescue operation using UAVs: A case study. Expert Syst Appl. 2021;178:114937.

4. Lyu M, Zhao Y, Huang C, Huang H. Unmanned aerial vehicles for search and rescue: A survey. Remote Sens. 2023;15:3266.

5. Zhou B, Zhang Y, Chen X, Shen S. FUEL: Fast UAV exploration using incremental frontier structure and hierarchical planning. IEEE Robot Autom Lett. 2021;6:779-86.

6. LaValle S. Rapidly-exploring random trees: A new tool for path planning. Annu Res Rep. 1998.

7. Connolly C. The determination of next best views. In: Proceedings of the 1985 IEEE International Conference on Robotics and Automation; St. Louis, MO, USA. 1985;2:432-5.

8. Bircher A, Kamel M, Alexis K, Oleynikova H, Siegwart R. Receding horizon next-best-view planner for 3D exploration. In: Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA); Stockholm, Sweden. 2016:1462-8.

9. Papachristos C, Khattak S, Alexis K. Uncertainty-aware receding horizon exploration and mapping using aerial robots. In: Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA); Singapore. 2017:4568-75.

10. Papachristos C, Khattak S, Mascarich F, Alexis K. Autonomous navigation and mapping in underground mines using aerial robots. In: Proceedings of the 2019 IEEE Aerospace Conference; Big Sky, MT, USA; 2019. p. 1-8.

11. Witting C, Fehr M, Behnemann R, Oleynikova H, Siegwart R. History-aware autonomous exploration in confined environments using MAVs. In: Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Madrid, Spain; 2018. p. 1-9.

12. Selin M, Tiger M, Duberg D, Heintz F, Jensfelt P. Efficient autonomous exploration planning of large-scale 3-D environments. IEEE Robot Autom Lett. 2019;4:1699-706.

13. Deng D, Duan R, Liu J, Sheng K, Shimada K. Robotic exploration of unknown 2D environment using a frontier-based automatic-differentiable information gain measure. In: Proceedings of the 2020 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM); Virtual. 2020. p. 1497-503.

14. Dharmadhikari M, Dang T, Solanka L, Loje J, Nguyen H, Khedekar N, et al. Motion primitives-based path planning for fast and agile exploration using aerial robots. In: Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA); Paris, France. 2020. p. 179-85.

15. Xu Z, Deng D, Shimada K. Autonomous UAV exploration of dynamic environments via incremental sampling and probabilistic roadmap. IEEE Robot Autom Lett. 2021;6:2729-36.

16. Zhu H, Cao C, Xia Y, Scherer S, Zhang J, Wang W. DSVP: dual-stage viewpoint planner for rapid exploration by dynamic expansion. In: Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Prague, Czech Republic. 2021. p. 7623-30.

17. Respall VM, Devitt D, Fedorenko R, Klimchik A. Fast sampling-based next-best-view exploration algorithm for a MAV. In: Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA); Xi'an, China. 2021. p. 89-95.

18. Duberg D, Jensfelt P. UFOExplorer: Fast and scalable sampling-based exploration with a graph-based planning structure. IEEE Robot Autom Lett. 2022;7:2487-94.

19. Yamauchi B. A frontier-based approach for autonomous exploration. In: Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97: "Towards New Computational Principles for Robotics and Automation"; Monterey, CA, USA; 1997. p. 146-51.

20. Juliá M, Gil A, Reinoso O. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. Auton Robot. 2012;33:427-44.

21. Cieslewski T, Kaufmann E, Scaramuzza D. Rapid exploration with multi-rotors: a frontier selection approach for high speed flight. In: Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Vancouver, BC, Canada; 2017. p. 2135-42.

22. Tang C, Liang Y, Yu S, Sun R, Zheng J. Autonomous 3D exploration of indoor environment based on wavefront algorithm. In: Proceedings of the 2020 IEEE International Conference on Networking, Sensing and Control (ICNSC); Nanjing, China; 2020. p.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

1-6.

23. Dai A, Papatheodorou S, Funk N, Tzoumanikas D, Leutenegger S. Fast frontier-based information-driven autonomous exploration with an MAV. In: Proceedings of the 2020 IEEE International Conference on Robotics and Automation (ICRA); Paris, France; 2020. p. 9570-6.

24. Deng D, Xu Z, Zhao W, Shimada K. Frontier-based automatic-differentiable information gain measure for robotic exploration of unknown 3D environments. arXiv. 2020;arXiv:2011.05288.

25. Meng Z, Qin H, Chen Z, Chen X, Sun H, Lin F, et al. A two-stage optimized next-view planning framework for 3-D unknown environment exploration, and structural reconstruction. IEEE Robot Autom Lett. 2017;2:1680-7.

26. Batinovic A, Petrovic T, Ivanovic A, Petric F, Bogdan S. A multi-resolution frontier-based planner for autonomous 3D exploration. IEEE Robot Autom Lett. 2021;6:4528-35.

27. Hess W, Kohler D, Rapp H, Andor D. Real-time loop closure in 2D LIDAR SLAM. In: Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA); Stockholm, Sweden. 2016. p. 1271-8.

28. Hornung A, Wurm KM, Bennewitz M, Stachniss C, Burgard W. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. Auton Robot. 2013;34:189-206.

29. Yu J, Shen H, Xu J, Zhang T. ECHO: An efficient heuristic viewpoint determination approach on frontier-based autonomous exploration for quadrotors. IEEE Robot Autom Lett. 2023;8:5047-54.

30. Zhao Y, Yan L, Xie H, Dai J, Wei P. Autonomous exploration approach for fast unknown environment mapping by using UAV equipped with limited FOV sensor. IEEE Trans Ind Electron. 2023;71:4933-43.

31. Zhang H, Wang S, Liu Y, Ji P, Yu R, Chao T. EFP: Efficient frontier-based autonomous UAV exploration strategy for unknown environments. IEEE Robot Autom Lett. 2024;9:2941-8.

32. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, et al. Human-level control through deep reinforcement learning. Nature. 2015;518:529-33.

33. Kulkarni P, Goswami D, Guha P, Dutta A. Path planning for a statically stable biped robot using PRM and reinforcement learning. J Intell Robot Syst. 2006;47:197-214.

34. Ma X, Xu Y, Sun G, Deng L, Li Y. State-chain sequential feedback reinforcement learning for path planning of autonomous mobile robots. J Zhejiang Univ Sci C. 2013;14:167-78.

35. Tai L, Liu M. Mobile robots exploration through CNN-based reinforcement learning. Robot Biomim. 2016;3:24.

36. Zhu D, Li T, Ho D, Wang C, Meng M.Q.-H. Deep reinforcement learning supervised autonomous exploration in office environments. In: 2018 IEEE International Conference on Robotics and Automation (ICRA); Brisbane, QLD, Australia. 2018. p. 7548-55.

37. Niroui F, Zhang K, Kashino Z, Nejat G. Deep reinforcement learning robot for search and rescue applications: exploration in unknown cluttered environments. IEEE Robot Autom Lett. 2019;4:610-7.

38. Chen F, Martin JD, Huang Y, Wang J, Englot B. Autonomous exploration under uncertainty via deep reinforcement learning on graphs. In: Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Las Vegas, NV, USA; 2020. p. 6140-7.

39. Li H, Zhang Q, Zhao D. Deep reinforcement learning-based automatic exploration for navigation in unknown environment. IEEE Trans Neural Netw Learn Syst. 2020;31:2064-76.

© 2023 The Author(s). | Published by Agro Intelligence Research Press

40. Rao Z, Wu Y, Yang Z, Zhang W, Lu S, Lu W, et al. Visual navigation with multiple goals based on deep reinforcement learning. IEEE Trans Neural Netw Learn Syst. 2021;32:5445-55.

41. Tao Y, Iceland E, Li B, Zwecher E, Heinemann U, Cohen A, et al. Learning to explore indoor environments using autonomous micro aerial vehicles. In: Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA); Yokohama, Japan; 2024. p. 15758-64.

42. Tang B, Ren Y, Zhu F, He R, Liang S, Kong F, et al. Bubble explorer: fast UAV exploration in large-scale and cluttered 3D-environments using occlusion-free spheres. In: Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Detroit, MI, USA; 2023. p. 1118-25.

43. Luo Y, Zhuang Z, Pan N, Feng C, Shen S, Gao F, et al. Star-searcher: a complete and efficient aerial system for autonomous target search in complex unknown environments. IEEE Robot Autom Lett. 2024;9:4329-36.

44. Han L, Gao F, Zhou B, Shen S. FIESTA: Fast incremental Euclidean distance fields for online motion planning of aerial robots. In: Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Macau, China; 2019.

© 2023 The Author(s). | Published by Agro Intelligence Research Press