Title

Geometry-Aware Moving Viewpoint Particle Swarm Optimisation for Efficient UAV Facade Coverage

---

Abstract

Inspection of high-rise building façades is critical for structural assessment, emergency response, and risk mitigation in dense urban environments. While Unmanned Aerial Vehicles (UAVs) provide a safe and flexible platform for data acquisition, the effectiveness of UAV-based inspection is fundamentally constrained by coverage path planning. Conventional optimisation approaches typically operate over static, precomputed viewpoint sets or abstract search spaces, limiting adaptability in complex, concave, and partially occluded geometries. This work proposes a geometry-aware moving-viewpoint Particle Swarm Optimisation (VPSO) framework for UAV façade coverage planning. In the proposed formulation, viewpoints are treated as movable agents constrained by local surface geometry, and the search space is reduced from three spatial degrees of freedom to a single angular parameter. By embedding geometric feasibility directly into the optimisation process, the method integrates coverage optimisation with continuous surface-aware motion. The approach is evaluated against a conventional PSO baseline across three building geometries of increasing complexity. Results demonstrate improved optimisation stability, higher coverage efficiency, and enhanced robustness in non-convex and occluded environments. The proposed framework provides a foundation for unified coverage and trajectory optimisation in UAV-based structural inspection.

Keywords

UAV coverage planning; particle swarm optimisation; moving viewpoints; building inspection; geometry-aware optimisation

1. Introduction

Rapid population growth and urbanisation have led to an increase in the development of vertical infrastructure. Global trend analyses indicate a growing emphasis on vertical development in residential and commercial infrastructure (Qu et al., 2021). Although these buildings improve land-use efficiency, they also present novel challenges during emergencies, such as fires and earthquakes. Fire incidents in high-rise buildings are of particular concern because of vertical evacuation constraints and delayed awareness, which significantly affect the evacuation response (Ronchi and Nilsson, 2013). Importantly, such incidents are preventable and containable if detected early and accurately.

Façade inspection and rapid response to dynamic situations are crucial in both emergency response and routine inspection. However, traditional practices rely heavily on human access and judgment, which is dangerous and prone to errors. Sensors and automation for this specific purpose are sparse, costly and sometimes infeasible or incorrect in hazardous conditions. These limitations have motivated increasing interest in Unmanned Aerial Vehicles (UAVs) as a means

of rapidly acquiring visual and thermal information from building exteriors without exposing human operators to risk (Idris et al., 2023).

Despite advances in hardware, the effectiveness of UAV-based inspection is often limited by path planning constraints and coverage efficiency. In emergency scenarios, incomplete coverage can severely compromise decision-making. Consequently, the development of robust, efficient and adaptive path planning strategies constitutes an urgent problem for structural monitoring emergency response applications.

UAV path planning for surveying and inspection is usually regarded as a coverage path planning (CPP) problem, where the objective is to devise the most efficient path for complete surveillance of an area. CPP has been extensively studied for large-scale and relatively open-area applications, including agricultural monitoring, wildfire detection, and terrain mapping (Tan et al., 2021). In such applications, the search domain is predominantly planar, sparsely occluded and tolerant to small coverage gaps. Consequently, many CPP approaches reduce the three-dimensional path planning problem to two dimensions by limiting drone manoeuvres to horizontal planes at fixed altitudes, thereby sacrificing the benefits of vertical navigation for obstacle avoidance (Gugan and Haque, 2023). Furthermore, these methods lack sufficient spatial resolution to operate effectively in complex, occluded environments and are not robust enough to handle highly concave or convex geometries.

In contrast, path planning for buildings and urban environments introduces a fundamentally different set of challenges. Building facades present vertical surfaces with complex geometries, irregular occlusion and intricate architectural features. In emergency scenarios, some small unobserved regions may be critical. As a result, conventional methods for open-area applications cannot be directly transferred for use in urban environments.

A review of modern literature reveals that early CPP research relied primarily on deterministic and graph-based formulations, while recent approaches increasingly employ metaheuristic and learning-based methods such as Genetic Algorithms (GA), Particle Swarm Optimisation (PSO), Ant Colony Optimisation (ACO), and reinforcement learning (Tan et al., 2021). These methods are particularly prominent due to their flexibility and speed in handling multi-objective optimisation. However, they typically operate in abstract search spaces that may not directly incorporate the structural information of complex physical environments.

In structural inspection specifically, GA has been employed to select the optimal set of viewpoints from a precomputed viewpoint space and find the shortest path (Zhao et al., 2023). Further, a greedy pose refinement is performed to ensure optimality. ACO has been used for cooperative path planning with multiple UAVs to reduce redundancy between multiple agents and significantly decrease path length (Bui et al., 2024). Similar to GA-based methods, ACO typically operates on a pre-defined viewpoint space set and focuses on selecting the optimal set from the viewpoint space and visitation sequence. PSO and its variants have also been explored for path planning, including hybrid PSO with potential field methods and spherical-vector-based PSO (Phung and Ha, 2021; Yang et al., 2015). (Nguyen et al., 2022) Use spherical-vector-based PSO with game theory to coordinate multiple UAVs for construction site monitoring. While effective for maintaining formation during area coverage, this method lacks explicit view-planning and surface-aware mechanisms required for detailed inspection of complex building façades.

Despite the diversity of optimisation algorithms proposed for CPP planning, existing methods share several limitations that restrict their applicability for emergency scenarios. Most approaches tailored for complete coverage operate on a static, precomputed viewpoint space, which requires non-trivial preprocessing and inherently restricts exploration to sampled regions. As a result, important viewpoints may be missing due to noise, complex geometry, or sparse data.

In many conventional PSO-based formulations, viewpoint selection and path planning are typically solved separately and sequentially. This separation increases the effective search-space complexity, thus increasing the likelihood of convergence to suboptimal solutions, potentially excluding better solutions.

Furthermore, while metaheuristic methods are effective at exploring the complex search space, most modern algorithms lack explicit geometry-aware constraints. Viewpoint update calculations are performed in an abstract search space without incorporating local geometric information, which limits their overall effectiveness, especially in narrow, concave areas and areas with occlusion. In addition, the stochastic nature of these methods necessitates scenario-specific hyperparameter tuning to achieve stable performance. These characteristics limit the reliability of existing approaches in time-critical coverage tasks with uncertainty in both the agent and the environment.

This paper addresses the above gap by proposing a geometry-aware moving-viewpoint Particle Swarm Optimisation (VPSO) framework for UAV façade coverage planning. In the proposed formulation, viewpoints are treated as movable agents constrained by local surface geometry, and the search space is reduced from three spatial degrees of freedom to a single angular parameter. The method is evaluated against a conventional PSO baseline across three building geometries of increasing complexity to assess stability, coverage efficiency, and geometric robustness.

2. Automation in Building Management and Assessment

The inspection and assessment of buildings are central in structural engineering, particularly in the context of high-rise buildings and dense urban environments. Façade conditions influence not only serviceability and durability but also life safety during emergency events such as fires, earthquakes, and extreme weather. In post-event scenarios, timely identification of damaged regions, compromised structural elements, and hazardous façade components is essential for decisions related to evacuation, access restriction, and repair prioritisation.

Traditionally, building inspection has been primarily performed through manual visual assessments conducted using scaffolding, rope-access techniques, suspended platforms, or lift systems (Bortolini and Forcada, 2018). These methods are labour-intensive, time-consuming, and expose personnel to significant safety risks, particularly in high-rise and post-disaster environments. Inspection quality also depends heavily on inspector experience, accessibility constraints, and subjective judgment. In emergency scenarios, where rapid assessment is required, these constraints may delay critical decision-making and limit the comprehensiveness of evaluation. Furthermore, physical access limitations often result in partial inspections, especially in complex geometries, damaged or unsafe building elements or occluded structural regions. Thus, automation offers not only improved safety but also the potential for systematic,

repeatable, and comprehensive inspection coverage that is difficult to achieve through manual methods alone.

UAV-based inspection has been integrated into structural engineering workflows for acquiring visual, thermal, and geometric data in locations that are unsafe or inaccessible to human inspectors (Kim et al., 2024). It is being increasingly used for crack detection, damage classification, condition assessment, and digital documentation of structures (Iqbal et al., 2024; Villarino et al., 2025). To support these applications, path planning is often integrated with domain-specific representations such as Building Information Models (BIM), which provide a detailed three-dimensional digital twin for use in path planning algorithms (Huang et al., 2023). BIM-based frameworks support the incorporation of structural semantics and prioritisation of important regions, resulting in more robust geometry-based path planning. However, such approaches rely heavily on model accuracy, require substantial preprocessing and precomputation of data and assume static environments. Consequently, their applicability in emergency-response scenarios remains limited. Moreover, reliance on precomputed viewpoints introduces vulnerability when viewpoints are missing due to noise, occlusions, or sparse data.

From a structural engineering perspective, UAV path planning extends beyond navigation to issues of data quality and risk mitigation. Coverage gaps may lead to unassessed structural regions, while redundant viewpoints reduce operational efficiency during time-sensitive missions. Moreover, building façades often exhibit geometric complexity associated with structural behaviour, including discontinuities at openings, connection zones, and load-transfer regions. Particle Swarm Optimisation (PSO) is a population-based metaheuristic algorithm in which candidate solutions, referred to as particles, iteratively update their positions in a search space based on their own best-known positions and the performance of neighbouring particles (Kennedy and Eberhart, 1995). In UAV inspection planning, PSO is commonly applied to select or sequence viewpoints to maximise coverage while penalising redundancy (Huang and Fei, 2018). However, conventional PSO formulations typically operate over static, precomputed viewpoint sets or abstract search representations. Such formulations may require substantial preprocessing and do not explicitly incorporate local surface geometry. As a result, geometric constraints must be accounted for indirectly during optimisation, potentially reducing stability and limiting adaptability in complex building environments.

Path-planning strategies that adapt to façade geometry and prioritise informative viewpoints can significantly improve the reliability and interpretability of collected data. In this context, geometry-aware optimisation frameworks such as the one proposed in this work provide a direct link between path planning techniques and practical structural engineering applications.

3. Problem Formulation

3.1 Facade Representation

The building façade is represented as a three-dimensional point cloud

$$\boldsymbol{\mathcal{B}} = \{\, \mathbf{p}_{\boldsymbol{i}} \in \mathbb{R}^{\mathbf{3}} \mid \boldsymbol{i} = \mathbf{1}, \ldots, \boldsymbol{N} \,\}, \tag{1}$$

where each point corresponds to a sampled location on the exterior surface of the structure.

Point cloud representations are widely used in UAV path planning papers for building inspection due to their compatibility with Building Information Models (BIM), efficiency in computation, and storage requirements (Huang et al., 2023). Unlike mesh-based or voxel-based

models, point clouds avoid explicit surface connectivity assumptions, which may be unreliable in emergency or rapidly deployed scenarios.

In practice, the point cloud may exhibit non-uniform density, measurement noise, and missing regions due to occlusion. These factors make global surface reconstruction unreliable, motivating the use of local geometric analysis to guide the optimisation process.

For any point $\mathbf{p_i}$, a local neighbourhood $\boldsymbol{\mathcal{N}_i} \subset \boldsymbol{\mathcal{B}}$ is defined using a k-nearest-neighbour criterion. Local geometric properties such as surface orientation and curvature are inferred from these neighbourhoods and used to guide viewpoint motion.

3.2 Viewpoint and Visibility Model

A UAV viewpoint is defined as a tuple

$$\mathbf{v} = (\mathbf{x}, \mathbf{d}), \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^3$ denotes the UAV position and $\mathbf{d} \in \mathbb{R}^3$ is the viewing direction.

The initial set of viewpoints is constructed by offsetting the building points along the outer surface normals by a fixed distance d_0, referred to as the optimum viewing distance. This set of candidate viewpoints is referred to as the viewpoint space.

Given a viewpoint $\mathbf{v} = (\mathbf{x}, \mathbf{d})$, then a building point $\mathbf{p}_i \in \mathcal{B}$ is considered visible if it satisfies all the following conditions:

1. **Distance constraint**:

$$\| \mathbf{p}_i - \mathbf{x} \| \leq d_{\max}, \tag{3}$$

where $d_{\max}$ is the maximum sensing distance.

2. **Field-of-view constraint**:

$$\angle(\mathbf{d}, \mathbf{p}_i - \mathbf{x}) \leq \phi \tag{4}$$

where ϕ is half of the camera's field-of-view angle.

3. **Occlusion constraint**: The line segment connecting $\mathbf{x}$ and $\mathbf{p}_i$ does not intersect any other part of the building geometry.

The coverage function for a viewpoint $\mathbf{v}$ is thus defined as

$$\mathcal{C}(\mathbf{v}) = \{ \mathbf{p}_i \in \mathcal{B} \mid \mathbf{p}_i \text{ satisfies all the above visibility constraints} \}. \tag{5}$$

For practical feasibility, viewing angles relative to the horizontal plane are restricted to less than 45° for realistic UAV orientations.

3.3 Optimisation Objective

Let:

$$\mathcal{B} = \{ \mathbf{p}_i \in \mathbb{R}^3 \mid i = 1, \dots, N \} \tag{6}$$

denote the set of building points.

$$\mathcal{V} = \{ \mathbf{v}_j \mid j = 1, \dots, M \} \tag{7}$$

denote the set of viewpoints.

The total coverage set is defined as

$$\mathcal{C}(\mathcal{V}) = \bigcup_{\mathbf{v}_j \in \mathcal{V}} \mathcal{C}(\mathbf{v_j}). \tag{8}$$

where $\mathcal{C}(\mathbf{v}_j)$ is defined in Section 3.2.

The number of uniquely covered façade points is therefore $| C(\mathcal{V}) |$.

To penalise redundancy, let the number of points visible from each viewpoint be

$$n_j = | \{C(v_j)\} |. \tag{9}$$

Then, the duplicate coverage count is

$$D(\mathcal{V}) = \sum_{j=1}^{M} n_j \; - \; | \mathcal{C}(\mathcal{V}) |. \tag{10}$$

The optimisation objective (fitness function) used in PSO is then defined as

$$\max_{\mathcal{V}} \quad F(\mathcal{V}) = | \mathcal{C}(\mathcal{V}) | \; - \; \alpha D(\mathcal{V}) \tag{11}$$

Where $\alpha > 0$ is a hyperparameter controlling the trade-off between maximising coverage and reducing redundancy.

4. Geometric Surface Analysis

This section describes how local surface geometry is estimated from the building point cloud and how it is used to constrain viewpoint motion during optimisation. For a given viewpoint $\mathbf{v}$, the set of visible façade points $\mathcal{C}(\mathbf{v})$ is first determined as per section 3.2. This set constitutes the neighbourhood points used for geometry estimation.

$$\mathcal{N} = \mathcal{C}(\mathbf{v}) \tag{11}$$

4.1 Covariance-Based Local Geometry

Let

$$\bar{p} = \frac{1}{| \mathcal{N} |} \sum_{p_i \in \mathcal{N}} p_i \tag{12}$$

denote the mean of the neighbourhood points.

When the surface is approximately planar, local geometry is estimated using the covariance matrix of the neighbourhood points:

$$\mathbf{\Sigma} = \frac{1}{| \mathcal{N} | -1} \sum_{\mathbf{p}_i \in \mathcal{N}} (\mathbf{p}_i - \bar{p})(\mathbf{p}_i - \bar{p})^{\top}, \tag{13}$$

where $\bar{p}$ is the mean of the neighbourhood points.

Eigenvalue decomposition of $\mathbf{\Sigma}$ yields three eigenvectors and corresponding eigenvalues:

$$\Sigma = Q\Lambda Q^T. \tag{14}$$

Under the planar surface assumption, the eigenvector associated with the smallest eigenvalue corresponds to the local surface normal, while the remaining two eigenvectors define the local tangent plane. The tangent plane is used to constrain viewpoint motion around the building and prevent drift away from the building.

4.2 Curvature Handling via Local Sphere Fitting

In regions with high curvature or near edges, covariance-based normal estimation becomes unreliable due to significant variance along the normal direction. To address these cases, a local sphere is fitted to the neighbourhood points using a least-squares formulation.

The fitted sphere provides an estimated centre c and radius r. If the fitted radius is large or numerically unstable, the surface is treated as locally planar, and the covariance-based normal is retained. Otherwise, the normal direction of the surface is defined as the radial vector:

$$\bar{n} = \frac{\bar{p} - c}{\|\bar{p} - c\|} \tag{15}$$

The local tangent plane is then constructed using two orthogonal vectors perpendicular to $\bar{n}$. The scaling factors associated with the local coordinate directions are adjusted according to the fitted sphere geometry. The updated values corresponding to the tangent plane are

$$\sigma_1 = \sigma_2 = r, \tag{15}$$

The value corresponding to the normal direction is modified according to the sphere radius and the optimum viewing distance.

$$\sigma_3 = r \times \left(1 - e^{-r/d}\right) \tag{16}$$

Where *d* is the optimum viewing distance.

5. Proposed Moving-Viewpoint PSO Method

5.1 Viewpoint Movement Parametrisation

In conventional PSO-based path planning, particles typically represent discrete selections of viewpoints over a predefined static viewpoint space. Such formulations require the algorithm to discover geometric constraints imposed by the building façade implicitly, often without explicit knowledge of the surface geometry (Huang and Fei, 2018).

In contrast, the proposed method treats each viewpoint as a movable agent constrained to the local surface geometry. To achieve this, a single angular variable is used instead of three independent spatial coordinates. This parameterisation reduces the dimensionality of the search space from three translational degrees of freedom to a single angular degree of freedom, while preserving geometric feasibility. Consequently, the optimisation process is focused on coverage improvement without requiring the algorithm to learn building geometry implicitly.

5.2 Swarm Initialisation

Each particle in the swarm comprises a fixed number of moving viewpoints initialised at random feasible locations around the building. For controlled comparison with a conventional PSO baseline, the initial locations were randomly sampled from the same precomputed viewpoint space described in Section 3.2. For each viewpoint v, an angular parameter $\theta \in [0,2\pi)$ is assigned, which defines the direction of motion on the local tangent plane of the façade surface.

At initialisation, local surface geometry at each viewpoint is estimated using the methods described in Section 4. The resulting tangent plane and surface normal define a local coordinate frame. One tangent vector is selected and stored as a reference direction for that viewpoint. The angular parameter θ is then measured relative to this reference direction.

During optimisation, the PSO algorithm operates exclusively on the angular parameter θ. Feasibility of movement is maintained through constraints derived from the local surface geometry.

5.3 Geometry-Constrained Viewpoint Motion

Given the angular parameter θ and a stored reference vector $\mathbf{r}$ associated with each viewpoint, movement is executed by projecting the reference vector onto the local tangent plane. Let $\mathbf{n}$ denote the unit surface normal vector at the current viewpoint location. The projection of the reference vector onto the tangent plane is given by:

$$\mathbf{t_1} = \mathbf{r} - (\mathbf{r} \cdot \mathbf{n})\mathbf{n} \tag{17}$$

$$\mathbf{t_2} = \mathbf{n} \times \mathbf{t_1} \tag{18}$$

The vectors $\mathbf{t_1}$ and $\mathbf{t_2}$ are normalised and form an orthonormal basis of the local tangent plane. The direction of motion is then obtained as

$$d = \cos\theta\, \hat{t}_1 + \sin\theta\, \hat{t}_2\,. \tag{19}$$

The viewpoint location is updated by displacing it along d by a fixed step size. The step size is constant throughout the optimisation and can be selected based on problem scale and sensing resolution.

5.4 Local Hill Climbing

Following the geometry-constrained deterministic movement, a local stochastic refinement step is performed to reduce overlap and clustering.

At the updated viewpoint location, local surface geometry is re-estimated using the methods described in Section 4. A multivariate Gaussian distribution centred at the current viewpoint is constructed using the locally computed covariance matrix, defining a three-dimensional probability distribution in the vicinity of the viewpoint.

A fixed number of candidate viewpoints (16 in this work) are sampled from this Gaussian distribution. Since global fitness is not available at this step, refinement is performed using a local repulsion-based fitness function. Let d_k denote the Euclidean distance between the candidate viewpoint and the k^{th} neighbouring viewpoint. A repulsion radius R is defined, and all neighbouring viewpoints within this radius are penalised as:

$$\mathcal{R}(\mathbf{v}) = \sum_{k:d_k<R} \beta\left(1 - \frac{d_k}{R}\right), \tag{20}$$

The local refinement fitness function is defined as:

$$\mathcal{F}(\mathbf{v}) = \mathcal{C}(\mathbf{v}) - \mathcal{R}(\mathbf{v}), \tag{21}$$

where $\mathcal{C}(\mathbf{v})$ is the coverage function defined in section 3.2. To regulate the redundancy repulsion during optimisation, the redundancy coefficient β is adjusted over iterations:

$$\beta = \beta_0 e^{-iter/\tau} \tag{22}$$

where $iter$ denotes the current iteration and τ is a decay constant. Among the sampled candidates, the viewpoint yielding the highest fitness is selected as the final viewpoint position for the current iteration.

This local refinement enables exploitation around promising regions identified by PSO while preventing clustering and maintaining geometric feasibility.

5.5 PSO Update Rule

The angular parameter θ associated with each viewpoint is optimised using a standard PSO formulation with a local neighbourhood (lbest) topology of size 8 (Kennedy and Eberhart, 1995) (Schor et al., 2010). Each viewpoint maintains its personal best angular position, while neighbourhood best information is exchanged between nearby particles rather than relying on a single global best.

To ensure meaningful neighbourhood interactions, all particles at initialisation are reindexed so that spatially proximate viewpoints have the same index across particles. This matching is performed through the Hungarian algorithm (Kuhn, 1955), where the cost matrix consists of the pairwise Euclidean distances between viewpoints of each particle and those of a reference particle.

At each iteration, the PSO update follows the standard velocity–position update rule applied to the angular parameter:

$$v^{(t+1)} = \omega v^{(t)} + c_1 r_1\left(\theta_{pbest} - \theta^{(t)}\right) + c_2 r_2\left(\theta_{nbest} - \theta^{(t)}\right), \tag{23}$$

$$\theta^{(t+1)} = \theta^{(t)} + v^{(t+1)} \tag{24}$$

where ω is the inertia coefficient, c_1 and c_2 are cognitive and social coefficients, and $r_1, r_2 \sim \mathcal{U}(0,1)$. The angular parameter θ is wrapped around $[0,2\pi)$.

5.6 Algorithm Summary

The complete update cycle for each iteration consists of:

1. Local geometry estimation: Surface normal and tangent plane are computed at the current viewpoint location.
2. Geometry-constrained motion: The viewpoint is displaced along the tangent plane according to the current angular parameter θ with a fixed step size.
3. Local hill climbing: A Gaussian neighbourhood is sampled and refined using the fitness function as described in section 5.4.

4. Geometry re-estimation: Surface geometry is recomputed at the refined viewpoint location.
5. PSO update: The angular velocity and position are updated using neighbourhood best information.

By projecting all reference vectors and update directions onto the locally valid tangent plane at every stage, the optimisation operates exclusively in angular space while maintaining geometric feasibility. This separation prevents divergence from the façade surface and enables stable optimisation across convex, concave, and irregular building geometries.

Algorithm 1 Geometry-Aware PSO Implementation

1. **Initialisation**:
    a. Generate a swarm of S particles
    b. **For** each particle
        i. sample M feasible viewpoints.
        ii. Assign angular parameter $\theta \in [0,2\pi)$ to each viewpoint.
    c. Reindex viewpoints across particles using Hungarian matching.
2. **For** iteration $t = 1$ to T:
    For each viewpoint in each particle:
        a. Estimate local surface geometry
        b. Perform geometry constrained motion
        c. Perform local hill-climbing refinement
        d. Re-estimate local surface geometry
    Evaluate global fitness
    Update personal best and neighbourhood best solutions
    Update PSO velocity and θ
3. **Return** Particle with highest fitness

6. Simulation and Testing

Following the formulation in Section 5, this section evaluates the proposed algorithm through a series of simulations. In addition to absolute coverage, these experiments are designed to test optimisation stability, ability to handle non-convex geometry and most importantly, the redundancy of each viewpoint.

For this purpose, the proposed algorithm is compared with a conventional PSO across a set of buildings with increasing structural complexity. These studies are intended to expose characteristic behaviours, failure modes and trade-offs associated with each optimisation strategy.

6.1 Test Environments

To evaluate these metrics under increasingly non-convex building geometry, three representative buildings were studied. These case studies reflect the distinct classes of building geometry rather than serving as exhaustive benchmarks.

1. **Synthetic Cylindrical Building:** A smooth, convex cylindrical façade was used as a baseline environment. The building has a height of 100 m and a radius of 25 m, and is represented as a stochastic point cloud comprising approximately 10,000 points to

emulate sensing noise. This geometry contains no occlusions or concavities and serves as a sanity check for optimisation behaviour under ideal conditions.

2. **Rectangular High-Rise Building:** A real-world building model was used with predominantly planar structure, sharp edges and mild concavities. This geometry introduces local curvature discontinuities and partial occlusions commonly observed in real urban buildings.

3. **Twin-Tower Building Configuration:** A strongly non-convex structure consisting of two adjacent towers separated by narrow gaps. It includes occluded regions and deep concavities between towers, serving as a stress test for viewpoint selection and exposing failure modes associated with static viewpoint sampling. The dimensions of the bounding box of the building were 68x58x153 m, and it was sampled with 20,000 points.

All façades were represented as 3D point clouds, and the same viewing distance, field-of-view constraints, and visibility criteria were used across all experiments to ensure fairness.

6.2 Compared Methods

The proposed geometry-aware moving-viewpoint PSO is evaluated against the conventional PSO-based path planning algorithm, which serves as a representative baseline. Particle Swarm Optimisation has been widely applied to robotics path planning and other optimisation problems because it is easy to implement, has fewer parameters, and requires less computation time in high-dimensional optimisation tasks (Gad, 2022).

In the context of coverage path planning, PSO is commonly used to optimise objectives by selecting the optimal subset and visitation sequence over a predefined set of candidate viewpoints. In the baseline formulation, a static precomputed viewpoint set is required before optimisation. Each particle encodes a discrete selection over this viewpoint set, and the fitness function combines façade coverage with redundancy penalties to discourage overlapping views. This formulation reflects standard PSO-based inspection approaches reported in the literature (Rosas-Carrillo et al., 2025; Yang and Mao, 2024).

In contrast, the proposed method treats viewpoints as movable entities, only imposing constraints based on local façade geometry. No exhaustive precomputation of the viewpoint space is required; only a sufficiently general knowledge of the region of interest is required. The viewpoints adapt to the structure as viewpoint motion is explicitly guided by surface orientation and curvature, as described in Section 5.

6.3 Comparison Criteria

Each method was evaluated under conditions that are optimal and appropriate for its formulation, with care taken to ensure a fair comparison.

1. Both approaches were initialised with the same number of viewpoints and evaluated using identical sensing constraints, including field-of-view, visibility criteria, and maximum viewing distance.
2. Due to the fundamentally different search spaces, fitness functions were formulated independently for each method to achieve optimal performance, while PSO hyperparameters were kept identical across methods.

3. Each experiment was repeated multiple times with different random initialisations to assess robustness.
4. No geometry-specific parameter tuning was performed for the proposed VPSO across different building geometries.
5. To ensure a fair and meaningful comparison, both methods were evaluated under a comparable viewpoint budget. This reflects realistic constraints in emergency-response scenarios, where excessive redundancy and unbounded growth in the number of viewpoints are not acceptable. While conventional PSO can achieve higher absolute coverage by substantially increasing the number of viewpoints if redundancy penalties are reduced in the fitness function, such solutions incur excessive overlap. They are not considered meaningful for time-critical inspection tasks.

6.4 Evaluation Metrics

The proposed method is evaluated using a combination of quantitative and qualitative metrics that are selected to reflect the practical objectives and potential errors in UAV-based building inspection. In emergency-response scenarios, high coverage alone is insufficient; convergence characteristics, area and density of blind spots, efficiency, and path length are equally critical for measuring and evaluating performance. For this purpose, the following metrics are adopted.

Coverage

Coverage is defined as the fraction of façade surface elements that are visible from at least one selected viewpoint. This is the primary objective of coverage path planning and is widely used in coverage path planning and UAV-based inspection literature to quantify inspection completeness (Choset, 2001; Tan et al., 2021). The procedure for visibility and coverage calculations is elaborated in Section 3.

$$\text{Coverage} = \frac{\text{Number of unique visible building points}}{\text{Total number of building points}}$$

Coverage Efficiency

Coverage efficiency is defined as the ratio of coverage to the number of viewpoints. It is effectively a measure of redundancy and signifies the effectiveness of each viewpoint in increasing coverage. This is especially relevant in emergency scenarios as higher efficiency implies a lower total flight time due to the reduced number of viewpoints. This is not only critical for drones with limited flight times but also reduces time-to-detection. Reinforcement learning approaches often define reward structures that explicitly encourage covering new areas while penalising repeated coverage, a notion similar to coverage efficiency (Huang et al., 2025).

$$\text{Coverage Efficiency} = \frac{\text{Coverage}}{\text{Number of viewpoints}}$$

Coverage Stability

Optimisation stability is evaluated by analysing the change in coverage over iterations. The PSO algorithm is known to exhibit instability in convergence behaviour based on the algorithm parameters and solution space characteristics (Tarekegn et al., 2023). In particular, it is

qualitatively determined through the overall drop in coverage over iterations and is used to identify catastrophic degradation caused by unstable viewpoint distribution and repulsion effects due to redundancy penalties. Quantitatively, this metric is measured by the maximum drop in coverage between any two iterations. This metric reflects whether the number of iterations reliably improves coverage and, on a deeper level, whether the algorithm is truly optimising or improving coverage only due to stochastic variation.

Qualitative Geometric Robustness

Quantitative metrics alone are often insufficient to capture failure modes arising from complex façade geometry. Therefore, qualitative assessment is used to evaluate geometric robustness, focusing on behaviours such as:

- clustering or collapse of viewpoints,
- Area and distribution of blind spots or unobserved regions,
- failure to traverse concave or occluded areas

These behaviours are especially critical in emergency-response scenarios, where small unobserved regions may correspond to structurally significant damage. Qualitative analysis is therefore used to complement numerical metrics and provide insight into the practical reliability of each method.

6.5 Results

Case Study 1: Synthetic Cylindrical Building

Results on the cylindrical geometry are summarised in Table 1. Figure 1 presents the simulation outcome and viewpoint distribution of a representative run. The initial coverage for randomly initialised viewpoints was 72-73%.

Table 1: Results of Case Study 1

Method	Viewpoints	Coverage (%)	Coverage Efficiency (% per viewpoint)	Max Coverage Drop
Novel VPSO	300	80 - 81	0.30	<2%
PSO	260 - 320	72 - 77	0.21- 0.25	>25%

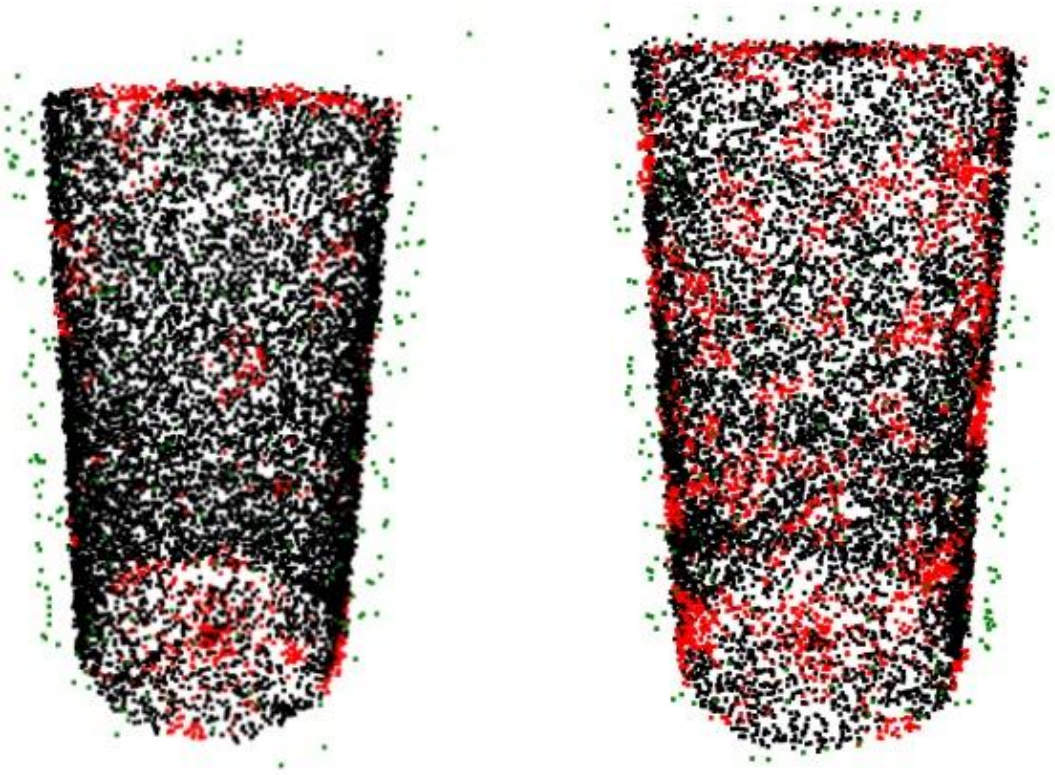

Figure 1: Result of Study 1: Conventional PSO (left), VPSO (right)

On this convex geometry, both methods produced relatively uniform viewpoint distributions with minimal clustering. Conventional PSO exhibited non-monotonic behaviour, with frequent drops due to redundancy penalties. VPSO demonstrated stable coverage improvement with less than 2% overall drop. Blind spots in PSO were randomly distributed, whereas in VPSO the remaining uncovered regions were concentrated near boundary areas, such as the roof and ground transitions

Case Study 2: Rectangular High-Rise Building

Results on the rectangular building geometry are summarised in Table 2. Figure 2 presents the simulation outcome and viewpoint distribution of a representative run. The initial coverage for randomly initialised viewpoints was 79-81% for 300 viewpoints and 80-82% for 350 viewpoints.

Table 2: Results of Case Study 2

Method	**Viewpoints**	**Coverage (%)**	**Coverage Efficiency (% per viewpoint)**	**Max Coverage Drop**
Novel VPSO	300	90-91	0.30	<3%
Novel VPSO	350	92-93	0.26	<3%
PSO	330 - 380	82 - 88	0.23- 0.25	>25%

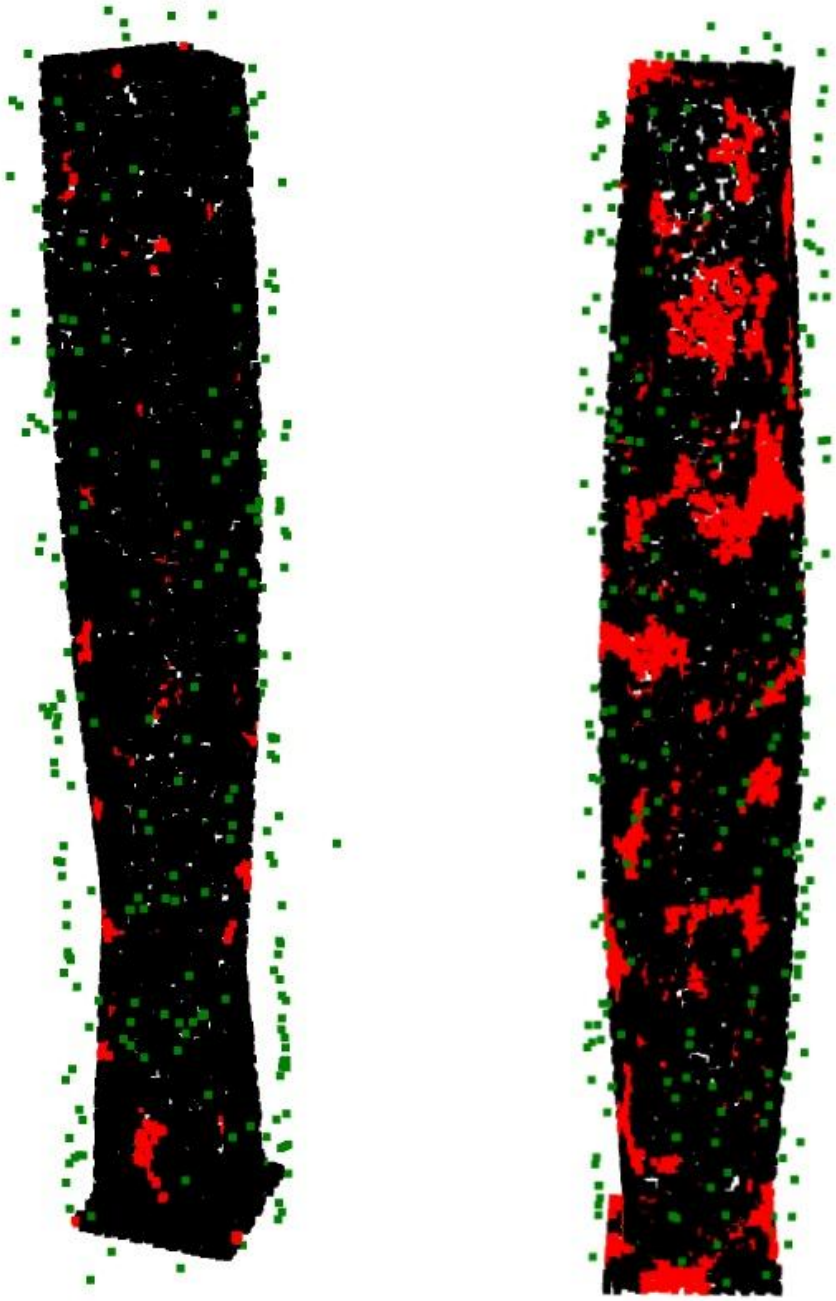

Figure 2: Result of Study 2: Conventional PSO (left), VPSO (right)

On the rectangular building with edges and mild concavities, differences between strategies became more apparent. Conventional PSO required careful calibration to avoid degradation and unstable convergence, and frequently exhibited coverage drops due to redundancy penalties. In contrast, VPSO maintained stable coverage evolution with minimal degradation.

Blind spots in conventional PSO were frequently located near edges and discontinuities. VPSO, by constraining viewpoint motion along the façade surface, redistributed viewpoints more coherently and exhibited fewer clustered blind spots. VPSO was able to traverse mildly non-convex regions through continuous motion along the façade surface, even in areas where initial sampling was sparse.

Case Study 3: Twin-Tower Geometry

Results on the twin tower building geometry are summarised in Table 3. Figure 3 presents the simulation outcome and viewpoint distribution of a representative run. The initial coverage for randomly initialised viewpoints was 79-80% for 800 viewpoints.

Table 3:Results of Case Study 3

Method	Viewpoints	Coverage (%)	Coverage Efficiency (% per viewpoint)	Max Coverage Drop
Novel VPSO	800	86-87	0.11	<5%
PSO	760-860	77 - 83	0.09-0.10	>15%

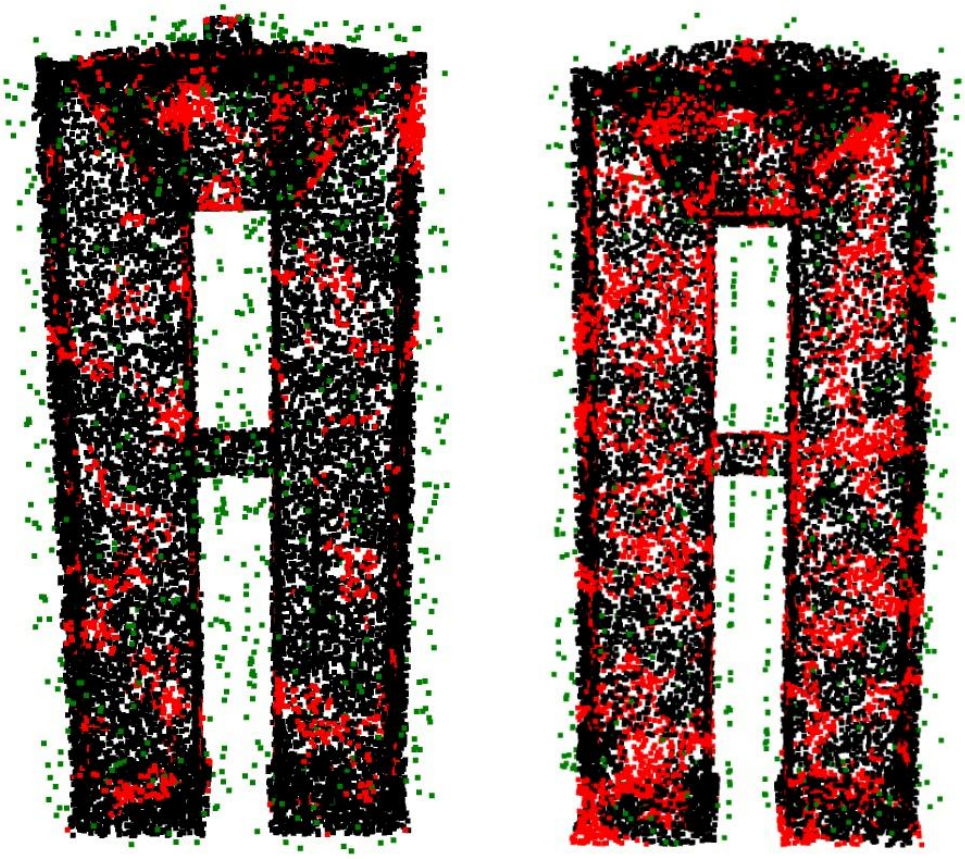

Figure 3: Result of Study 3: Conventional PSO (left), VPSO (right)

The twin-tower configuration exposed the limitations of static-viewpoint optimisation methods. In this geometry, conventional PSO struggled to improve coverage beyond the randomly initialised configuration significantly. VPSO was able to progressively migrate viewpoints into concave and partially occluded regions through continuous geometry-constrained motion. Although some blind spots persisted, coverage improvements were consistently observed across runs. The advantage of moving-viewpoint optimisation was most evident in this geometry, where static sampling imposed structural limitations on exploration.

6.6 Summary of Experimental Findings

Across increasing geometric complexity, several consistent trends were observed. First, VPSO demonstrated stable coverage evolution with minimal degradation across iterations, whereas conventional PSO exhibited frequent non-monotonic behaviour and sensitivity to redundancy penalties. Second, VPSO achieved comparable or higher coverage, using fewer or similarly bounded viewpoints, resulting in improved coverage efficiency across most runs. Third, the advantage of geometry-aware viewpoint motion became more pronounced in non-convex and occluded environments, where static-viewpoint PSO struggled to meaningfully improve initial coverage.

These observations suggest that the proposed formulation modifies not only optimisation performance but also the structural characteristics of the search space, particularly in complex façade geometries.

7. Results and Discussion

This section discusses the algorithm behaviour observed in Section 6 and interprets the results. Rather than focusing solely on absolute coverage values, the discussion emphasises robustness, efficiency, and failure modes of the algorithm which are central to building inspection and emergency-response applications.

7.1 Coverage Behaviour and Optimisation Stability

Across all tested geometries, the most immediate distinction between the conventional PSO formulation and the proposed VPSO was the stability of coverage during optimisation.

In the conventional PSO, coverage evolution was frequently non-monotonic and highly dependent on careful tuning of hyperparameters. Although high coverage values were often reached early in optimisation, these were coupled with a significant increase in viewpoints. Often, this was followed by substantial drops caused by redundancy penalties or unstable viewpoint redistribution. Increasing the number of iterations does not improve final coverage; instead, the algorithm tends to oscillate around suboptimal configurations or drastically reduce the number of viewpoints. These observations indicate that, for the conventional PSO, iteration count alone is not a reliable indicator of convergence quality.

In contrast, VPSO exhibited stable and repeatable coverage behaviour. Coverage increased progressively and remained bounded during optimisation. Minor fluctuations were caused by redundancy regulation rather than instability. By constraining the motion of viewpoints along the façade surface, VPSO prevents large discontinuous changes in visibility and implicitly regularises the optimisation process.

7.2 Coverage Efficiency and Redundancy Reduction

A significant advantage of VPSO observed across experiments is its ability to achieve comparable or higher coverage with substantially fewer viewpoints. While conventional PSO often relied on increased viewpoint count to compensate for redundancy and blind spots, VPSO was able to position viewpoints more effectively through continuous motion along the façade surface.

This improvement in coverage efficiency is not achieved through explicit redundancy penalties alone, but rather emerges from the structure of the algorithm. By restricting viewpoint movement updates to geometry-consistent directions and reducing search space dimensionality

to a single parameter, VPSO promotes coherent distribution of viewpoints rather than relying on stochastic relocation. As a result, additional viewpoints contribute meaningfully to coverage instead of overlapping with existing fields of view.

From a practical UAV deployment perspective, improved coverage efficiency directly translates to reduced sensing time and flight trajectories, which is critical in time-sensitive scenarios such as fire detection.

7.3 Access to Non-Precomputed Viewpoints

A central motivation for developing VPSO was addressing a key limitation of static-viewpoint methods, including conventional PSO, which constrain optimisation to a predefined set of viewpoints. This limitation was particularly evident in concave and occluded regions, such as along the edges and in highly curved geometries, where effective viewpoints may be absent from initial sampling.

Experimental results on non-convex geometries, particularly in the twin-tower configuration with narrow concavities, demonstrated that VPSO progressively moved viewpoints into regions that were not explicitly represented in the initial viewpoint set. By treating viewpoints as movable agents constrained to local surface geometry, the algorithm continuously adapts their positions based on surface orientation and curvature rather than relying solely on discrete sampling.

This difference highlights a change in problem formulation rather than parameter tuning. By allowing viewpoints to move, the optimisation is no longer reliant on exhaustive precomputation of viewpoints and enables adaptive coverage for complex geometries.

7.4 Robustness to Complex Geometry and Local Optima

Another notable outcome was the consistency of VPSO results across repeated runs, even on complex geometries. While the exact final viewpoint locations varied, the achieved coverage remained similar across runs. In contrast, the conventional PSO baseline exhibited higher run-to-run variance, with final coverage strongly dependent on initial particle distribution and stochastic updates.

This empirical consistency suggests increased robustness to local optima and poor initial configurations. By constraining movement to a continuous surface manifold and reducing the dimensionality of the search space, VPSO appears to reduce the likelihood of becoming trapped in poor configurations that are common in high-dimensional, discontinuous fitness landscapes.

7.5 Computational Considerations

VPSO incurs a higher computational cost per iteration due to local geometry estimation and constrained motion updates. However, this increased cost is offset by the effectiveness of each iteration, which produces stable, geometry-consistent refinement rather than stochastic repositioning.

In contrast, while conventional PSO iterations are computationally cheaper, increased iteration counts do not consistently improve solution quality. From an optimisation perspective, VPSO trades raw iteration speed for more structured and reliable progress towards coverage objectives.

For UAV inspection scenarios, where total mission time and solution reliability are more critical than per-iteration speed, this trade-off can be advantageous. The improved stability and coverage efficiency reduce the need for re-optimisation and manual intervention.

7.6 Implications for Unified Coverage and Path Planning

A significant conceptual advantage of VPSO is its potential to unify coverage optimisation and path planning within a single optimisation framework. Because viewpoints are movable and maintain geometric continuity, additional objectives such as total flight distance, travel time or energy consumption can be incorporated directly into the optimisation process without requiring a post-processing stage.

In many conventional PSO-based inspection approaches, coverage optimisation and path planning are treated as sequential phases: viewpoints are first selected from a static set, after which a routing problem (e.g., TSP) is solved independently. This separation may introduce additional local optima and limit the ability to balance coverage quality against trajectory efficiency during optimisation.

While this work focuses primarily on coverage behaviour, the experimental results suggest that geometry-aware moving viewpoints provide a natural foundation for single-phase multi-objective optimisation. This is particularly attractive for autonomous UAV inspection and emergency-response missions, where complete coverage and path efficiency are both critical.

7.7 Limitations and Scope of Interpretation

The present experimental evaluation prioritises conceptual validation over exhaustive benchmarking across large-scale real-world datasets. While the tested buildings capture varying levels of geometric complexity, broader evaluation across models remains an area of future work.

In addition, other objectives such as path length and energy consumption were not fully incorporated into the current formulation. These factors are important and require further investigation.

Nevertheless, the observed trends are consistent across multiple geometries and runs, supporting the central claims regarding stability, efficiency, and geometric robustness.

8. Conclusions and Future Work

This work presented a geometry-aware moving-viewpoint Particle Swarm Optimisation (VPSO) framework for UAV-based façade inspection in complex urban environments. The proposed formulation integrates geometric constraints directly into the optimisation process, enabling stable and efficient inspection in complex building environments.

The main contributions of this work are summarised as follows:

- Moving-viewpoint formulation: Viewpoints are treated as movable agents rather than static selections from a precomputed set.
- Geometry-aware motion: Viewpoint updates are constrained by locally estimated surface geometry.
- Dimensionality reduction: The search space is reduced from three spatial parameters to a single angular parameter.

- Improved optimisation behaviour: Experimental results demonstrate enhanced stability and coverage efficiency across geometries.

Future work will focus on incorporating trajectory-related objectives such as flight time and path length into the unified optimisation framework, exploring learning-based strategies for enhanced local refinement, and validating the approach through deployment on real-world structural inspection scenarios.

References

Bortolini, R., Forcada, N., 2018. Building Inspection System for Evaluating the Technical Performance of Existing Buildings. J. Perform. Constr. Facil. 32, 04018073. https://doi.org/10.1061/(ASCE)CF.1943-5509.0001220

Bui, D.N., Duong, T.N., Phung, M.D., 2024. Ant Colony Optimization for Cooperative Inspection Path Planning Using Multiple Unmanned Aerial Vehicles, in: 2024 IEEE/SICE International Symposium on System Integration (SII). pp. 675–680. https://doi.org/10.1109/SII58957.2024.10417512

Choset, H., 2001. Coverage for robotics – A survey of recent results. Ann. Math. Artif. Intell. 31, 113–126. https://doi.org/10.1023/A:1016639210559

Gad, A.G., 2022. Particle Swarm Optimization Algorithm and Its Applications: A Systematic Review. Arch. Comput. Methods Eng. 29, 2531–2561. https://doi.org/10.1007/s11831-021-09694-4

Gugan, G., Haque, A., 2023. Path Planning for Autonomous Drones: Challenges and Future Directions. Drones 7. https://doi.org/10.3390/drones7030169

Huang, C., Fei, J., 2018. UAV Path Planning Based on Particle Swarm Optimization with Global Best Path Competition. Int. J. Pattern Recognit. Artif. Intell. 32, 1859008. https://doi.org/10.1142/S0218001418590085

Huang, J., Li, H., Chen, C., Liu, Y., Zhang, X., 2025. An Improved Deep Reinforcement Learning-Based UAV Area Coverage Algorithm for an Unknown Dynamic Environment. Appl. Sci. 15, 8942. https://doi.org/10.3390/app15168942

Huang, X., Liu, Y., Huang, L., Stikbakke, S., Onstein, E., 2023. BIM-supported drone path planning for building exterior surface inspection. Comput. Ind. 153, 104019. https://doi.org/10.1016/j.compind.2023.104019

Idris, N., Kaamin, M., Sarif, A., Nawawi, M., Noh, H., Mokhtar, M., Kadir, A., 2023. High-Rise Building Inspection by Using Unmanned Aerial Vehicles. Constr. Technol. Archit. 4, 195–209. https://doi.org/10.4028/p-dcbmba

Iqbal, M., Bachan, N., Khan, M., Subhani, M., 2024. Applications of UAVs and AI in Civil Engineering: A Systematic Review of Current Uses and Future Directions in Road Construction. https://doi.org/10.31224/4046

Kennedy, J., Eberhart, R., 1995. Particle swarm optimization, in: Proceedings of ICNN'95 - International Conference on Neural Networks. Presented at the ICNN'95 - International Conference on Neural Networks, pp. 1942–1948 vol.4. https://doi.org/10.1109/ICNN.1995.488968

Kim, S.Y., Yun Kwon, D., Jang, A., Ju, Y.K., Lee, J.-S., Hong, S., 2024. A review of UAV integration in forensic civil engineering: From sensor technologies to geotechnical, structural and water infrastructure applications. Measurement 224, 113886. https://doi.org/10.1016/j.measurement.2023.113886

Kuhn, H.W., 1955. The Hungarian method for the assignment problem. Nav. Res. Logist. Q. 2, 83–97. https://doi.org/10.1002/nav.3800020109

Nguyen, L.V., Herrera, I.T., Le, T.H., Phung, M.D., Aguilera, R.P., Ha, Q.P., 2022. Stag hunt game-based approach for cooperative UAVs. https://doi.org/10.22260/ISARC2022/0051

Phung, M.D., Ha, Q.P., 2021. Safety-enhanced UAV Path Planning with Spherical Vector-based Particle Swarm Optimization. Appl. Soft Comput. 107, 107376. https://doi.org/10.1016/j.asoc.2021.107376

Qu, J., Wang, Z., Du, P., 2021. Comparative Study on the Development Trends of High-rise Buildings Above 200 Meters in China, the USA and the UAE. Int. J. High-Rise Build. 10, 63–71. https://doi.org/10.21022/IJHRB.2021.10.1.63

Ronchi, E., Nilsson, D., 2013. Fire evacuation in high-rise buildings: a review of human behaviour and modelling research. Fire Sci. Rev. 2, 7. https://doi.org/10.1186/2193-0414-2-7

Rosas-Carrillo, A.S., Solís-Santomé, A., Silva-Sánchez, C., Camacho-Nieto, O., 2025. UAV Path Planning Using an Adaptive Strategy for the Particle Swarm Optimization Algorithm. Drones 9, 170. https://doi.org/10.3390/drones9030170

Schor, D., Kinsner, W., Anderson, J., 2010. A study of optimal topologies in swarm intelligence, in: CCECE 2010. Presented at the CCECE 2010, pp. 1–8. https://doi.org/10.1109/CCECE.2010.5575132

Tan, C., Mohd-Mokhtar, R., Arshad, M.R., 2021. A Comprehensive Review of Coverage Path Planning in Robotics Using Classical and Heuristic Algorithms. IEEE Access 9, 119310–119342. https://doi.org/10.1109/ACCESS.2021.3108177

Tarekegn, D., Tilahun, S., Gemechu, T., 2023. A Review on Convergence Analysis of Particle Swarm Optimization. Int. J. Swarm Intell. Res. 14. https://doi.org/10.4018/IJSIR.328092

Villarino, A., Valenzuela, H., Antón, N., Domínguez, M., Méndez Cubillos, X.C., 2025. UAV Applications for Monitoring and Management of Civil Infrastructures. Infrastructures 10, 106. https://doi.org/10.3390/infrastructures10050106

Yang, L., Qi, J., Yong, X., 2015. A literature review of UAV 3D path planning. Proc. World Congr. Intell. Control Autom. WCICA 2015, 2376–2381. https://doi.org/10.1109/WCICA.2014.7053093

Yang, M., Mao, P., 2024. UAV 3-D Path Planning Based on Particle Swarm Optimization Algorithm in Complex Terrain, in: 2024 6th International Conference on Internet of Things, Automation and Artificial Intelligence (IoTAAI). Presented at the 2024 6th International Conference on Internet of Things, Automation and Artificial Intelligence (IoTAAI), pp. 221–226. https://doi.org/10.1109/IoTAAI62601.2024.10692576

Zhao, Y., Lu, B., Alipour, M., 2023. UAS-based Automated Structural Inspection Path Planning via Visual Data Analytics and Optimization. https://doi.org/10.48550/arXiv.2312.15109