

Multi-UAV Search Strategies for Anomaly Monitoring

İlham Mirzayev^{1*} , İsmayıl İsmayilov² , Emil Sadıgov¹ , Fərid Quliyev¹ 

Abstract. Teams of unmanned aerial vehicles with onboard vision are used to locate events of interest inside a bounded area, such as disaster victims, wildfire ignition points, and industrial leaks. The choice of search strategy is usually made without a formal basis. This study asks when a detection-driven adaptive strategy outperforms systematic geometric coverage. The spatial clustering of events is treated as a controlled variable generated by a Thomas cluster process. Lawnmower, spiral, and Hilbert-curve sweeps are compared against an adaptive strategy that adds bounded detection-triggered detours to a Hilbert sweep. In a 400 m arena the adaptive strategy found 23 percent more events than the best sweep under tight clustering. The advantage crossed to a deficit near a cluster scale of about 47 m. A frame-level analysis shows that this gain is not simply an artifact of the camera footprint. Co-detectability is nonetheless collinear with the clustering variable, so a second experiment separates the two effects through a minimum event separation of 90 m that exceeds the footprint diagonal. The advantage then moves from event count to detection timing. Clustered events were found 34 s earlier on average, and at the tightest clustering level half of the events were reached after 275 s rather than 483 s. Under uniform placement the detours became a bounded cost and the sweeps performed better.

Keywords: multi-UAV search, coverage path planning, adaptive search, spatial clustering, event detection, co-detection

¹National Aviation Academy, PhD student, Baku, Azerbaijan

²National Aviation Academy, Doctor of Technical Sciences, Baku, Azerbaijan

*Corresponding author. E-mail: ilham.mirzayev60230@naa.edu.az

Received: 21 May 2026; Accepted: 28 September 2026; Published online: 22 October 2026

© The Author(s) 2026. This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0).

Anomaliyaların monitorinqi üçün çoxsaylı PUA-larla axtarış strategiyaları

İlham Mirzəyev^{1*} , İsmayıl İsmayilov² , Emil Sadıqov¹ , Fərid Quliyev¹ 

Xülasə. Görüntü emalı sistemləri ilə təchiz olunmuş pilotsuz uçuş aparatlarından ibarət qruplardan bir çox hadisələrin, məsələn, fəlakət qurbanlarının, meşə yanğınlarının başlanğıc nöqtələrinin və sənaye sızmalarının aşkarlanması üçün istifadə olunur. Bu əməliyyatlarda axtarış strategiyasının seçimi adətən formal əsaslandırma olmadan həyata keçirilir. Bu tədqiqatda aşkarlamaya əsaslanan adaptiv strategiyanın hansı hallarda sistematik metodlardan daha yaxşı nəticə verdiyi araşdırılmışdır. Hadisələrin məkan üzrə qruplaşması Thomas klaster prosesi vasitəsilə idarə olunan dəyişən kimi modelləşdirilmişdir. Lawnmower, Spiral və Hilbert əyrisi üzrə axtarış strategiyaları Hilbert marşrutunun genişləndirilmiş adaptiv variantı ilə müqayisə edilmişdir. 400 m-lik sınaq ərazisində sıx qruplaşma şəraitində adaptiv strategiya ən yaxşı sistematik axtarış strategiyasından 23 faiz daha çox hadisə aşkar etmişdir. Klaster miqyası təxminən 47 m olduqda bu üstünlük əksinə çevrilərək çatışmazlığa keçmişdir.

Tək bir kamera görüntüsündən aşkarlanmış hadisələr ilə bağlı analizlər göstərir ki, bu üstünlük sadəcə kameranın əhatə sahəsinin nəticəsi deyil. Bununla belə, birgə aşkarlanma qruplaşma dəyişəni ilə xətti əlaqədə olduğundan, ikinci eksperimentdə hadisələr arasında kameranın əhatə sahəsinin diaqonalını aşan 90 m-lik minimum məsafə tətbiq edilərək bu iki təsir bir-birindən ayrılmışdır. Bu halda üstünlük aşkar edilmiş hadisələrin sayından onların aşkarlanma vaxtına keçmişdir. Sıx qruplaşmış hadisələr orta hesabla 34 saniyə daha tez aşkar edilmiş və ən sıx qruplaşma səviyyəsində hadisələrin yarısına 483 saniyə əvəzinə 275 saniyədə çatılmışdır. Hadisələrin ərazi üzrə bərabər paylandığı halda isə adaptiv yanaşmadakı əlavə dövrlər xərcə çevrilmiş və sistemə axtarış strategiyaları daha yaxşı nəticə göstərmişdir.

Açar sözlər: çox-PUA-lı axtarış, əhatə etmə marşrutunun planlaşdırılması, adaptiv axtarış, anomaliya klasterləşməsi, hadisələrin aşkar olunması, birgə aşkar olunma

¹Milli Aviasiya Akademiyası, doktorant, Bakı, Azərbaycan

²Milli Aviasiya Akademiyası, texnika elmləri doktoru, Bakı, Azərbaycan

*Məsul müəllif. E-poçt: ilham.mirzayev60230@naa.edu.az

Daxil oldu: 21 May 2026; Qəbul edildi: 28 Sentyabr 2026; Onlayn dərc edildi: 22 Oktyabr 2026

© Müəllif(lər) 2026. Bu açıq girişli məqalə Creative Commons Attribution–NonCommercial 4.0 Beynəlxalq Lisenziyasının (CC BY-NC 4.0) şərtlərinə uyğun olaraq yayımlanmışdır.

Introduction

Large and complex areas often require frequent monitoring to detect anomalies before they lead to serious consequences. Examples include disaster victims in search-and-rescue operations, wildfire ignition points, industrial leaks, and other events of interest that may be difficult to locate quickly over a wide area. Conventional monitoring methods can be costly or limited by coverage area, weather, resolution, or revisit time. Cooperative multi-UAV teams equipped with real-time computer vision offer a flexible alternative that can be deployed on demand to search large areas and locate anomalies more rapidly.

A central question in cooperative aerial monitoring is which search strategy to use. Systematic geometric sweeps, such as lawnmower patterns and space-filling curves, guarantee complete and uniform coverage of the survey area. Detection-driven adaptive search takes a different approach and uses real-time detections to redirect UAVs toward neighboring regions. Adaptive search relies on the assumption that shared causal factors make anomalies spatially clustered. This assumption is often reasonable in industrial settings, but it is not always satisfied.

The physical basis of clustering differs across anomaly types and monitoring scenarios. In disaster response, for example, people and damaged infrastructure may occur in groups because they are affected by the same event. During wildfire monitoring, ignition points and fire-related anomalies can be spatially concentrated within areas where environmental conditions favor ignition. In industrial monitoring, leaks, equipment failures, and other faults may cluster around particular facilities, pipeline sections, or production areas. Environmental monitoring can also produce clustered observations when multiple anomalies arise from the same local source or environmental condition. In contrast, some anomalies may occur independently across the monitored area. Isolated manufacturing defects, sporadic infrastructure damage, or individual security events are examples where detecting one anomaly provides little information about the location of another. The degree of spatial clustering therefore varies across anomaly types and operating conditions, making it an important factor in deciding whether detection-driven adaptive search can provide a useful advantage over systematic coverage.

This work makes five contributions. First, it establishes a controlled evaluation framework that parameterizes anomaly clustering with a Thomas process and drives a cooperative multi-UAV search with vision-in-the-loop detections. Second, it uses a paired-seed design in a first set of experiments to compare three standard coverage geometries against a minimal and interpretable adaptive strategy. Third, it separates theoretical geometric co-detectability from realized co-detection through a frame-level analysis, which shows that the realized artifact is small, lower for the adaptive method than for the Hilbert baseline, and collinear with the clustering treatment. Fourth, it resolves this collinearity with a separation-controlled redesign in a second set of experiments, where a hard one-footprint minimum target separation moves the adaptive advantage from target count to detection latency. Fifth, it isolates the underlying search mechanism through ablation in both experiments. These contributions apply directly to offshore spill surveillance, pipeline monitoring, and facility inspection.

1.1. UAV-based monitoring

Aerial monitoring under a hard time limit is established across several application domains. A survey of civil applications lists search and rescue, remote sensing, and civil infrastructure inspection among the mature uses of unmanned aircraft (Shakhathreh et al., 2019). Search and rescue has received its own review, covering platform types, sensor payloads from thermal and RGB imaging to LiDAR, target detection methods, and the assignment of tasks across a team (Lyu et al., 2023). The breadth of that treatment shows how much of the field concerns perception and platform design rather than the order in which ground is examined. Wildfire monitoring places the same problem under a moving deadline, and one system plans fixed-wing trajectories with a variable neighborhood search metaheuristic so that a fleet can track a spreading front (Bailon-Ruiz et al., 2022). Its planner couples models of terrain, wind, and fire propagation in order to predict where the fire will be, which is the opposite regime from the static targets studied here. The oil and gas value chain presents the same structure, and UAVs are now applied across it. A recent survey of UAV-based remote sensing for the petroleum industry identifies six application areas: offshore oil spill detection, oil leakage detection, pipeline monitoring, gas-emission sensing, remote facility inspection, and environmental monitoring (Asadzadeh et al., 2022). The same survey reviews the platforms and sensor payloads commonly used in each area. Operational experience with the technology predates these surveys, and early use of unmanned aircraft for oil spill surveillance established the concept in practice (Allen & Walsh, 2008). These studies fix the operational scenario adopted here. A team of UAVs performs a time-limited patrol over a large area in which the locations of anomalies are unknown before the mission. Depending on the application, the surveyed area may include a disaster zone, a forest affected by wildfire, an industrial facility, an infrastructure corridor, or another monitored region. Anomaly detection is performed during the mission so that confirmed detections can immediately influence the search trajectory rather than being processed only after the survey is completed. This setting also motivates a modular detector interface. Any perception system that can provide confirmed anomaly positions in real time can be substituted without changing the search layer evaluated in this paper.

1.2. Coverage path planning and systematic search geometries

Coverage path planning computes a path that passes a sensor over every point of a region. An early survey of the field groups the published methods by how each one decomposes the region before a path is planned (Choset, 2001). A later survey extends that grouping and adds the methods introduced in the intervening decade (Galceran & Carreras, 2013). The version written for unmanned aircraft catalogs the simple flight patterns, including back-and-forth and spiral paths, and reports energy consumption as the main optimization criterion under the limited endurance of the platform (Cabreira et al., 2019). None of the three treats the spatial distribution of the targets as an input to the choice of pattern.

Two constructions from this literature serve as baselines here. The first is the boustrophedon decomposition, in which each cell of an exact decomposition is covered by back-and-forth motions

(Choset & Pignon, 1998). The same decomposition underpins the standard treatment of complete coverage in known spaces, where covering every cell by back-and-forth motions is reported to be provably complete (Choset, 2000). A multi-robot extension divides the cells among a team so that repeated coverage is minimized (Rekleitis et al., 2008). That partition-then-cover structure is adopted directly in this study, with one sub-region assigned to each aircraft.

The second construction is the space-filling curve. A traversal of such a curve maps two-dimensional coverage onto a locality-preserving one-dimensional ordering with minimal backtracking. The reported method minimizes the robot-to-robot communication needed to organize the movements, and it guarantees coverage in a finite time that decreases as the reciprocal of the number of robots sharing the task (Spires & Goldsmith, 1998). Later work adapts the same curve to regions that contain obstacles (Joshi et al., 2019). Locality preservation is not only an efficiency property. A Hilbert traversal visits spatially adjacent cells consecutively, so a vehicle can leave the curve for a local excursion and resume it at the saved index without stranding distant unswept pockets. The adaptive strategy proposed here depends on exactly this property, and the Hilbert curve therefore serves both as a systematic baseline and as the exploration layer of the adaptive method.

Energy is the usual objective when a spiral is preferred. An energy-aware spiral planner for UAV photogrammetry achieved energy savings over back-and-forth coverage in reported experiments (Cabreira et al., 2018). Operational search and rescue doctrine takes a different route and standardizes a small set of patterns, including the parallel sweep, creeping line, expanding square, and sector search, together with the effective sweep-width concept (International Civil Aviation Organization & International Maritime Organization, 2006). Sweep width is the quantity that converts a sensor specification into a lane spacing, and the same reasoning fixes the lane spacing of all four methods compared here.

Across this literature, pattern selection is generally made by convention or by single-objective optimization of energy, time, or path length. There is no principled guidance for choosing a pattern from the expected spatial distribution of the targets, and there is no reporting convention for target spacing relative to the sensor footprint.

1.3. Adaptive and detection-driven search

Adaptive search updates the plan from observations instead of fixing it in advance. In the Bayesian formulation of coordinated search for a lost target, each platform maintains a probability density over the target location and acts on it (Bourgault et al., 2003). The density is updated by both detections and non-detections, so a sweep that finds nothing still changes where the team will look next. Search has also been treated as a sequence of decisions, with the policy chosen to minimize the expected time until the searcher locates the target or declares it absent (Chung & Burdick, 2012).

A third approach defines coverage with respect to an expected information density. The reported objective encodes both exploration and exploitation, and it trades between them in a single objective function whether the information is diffuse or localized (Miller et al., 2016). What these three formulations share is continuous replanning against a global utility, and none of them bounds how much systematic coverage is surrendered for a single exploitation decision. That bound is the design question addressed in this study.

Industrial deployments show the same unbounded convergence. A team driven by distributed consensus converges on a detected gas source (Facinelli et al., 2019). The proposed strategy differs in that it limits how many aircraft can leave the main search area to investigate a detected target. This allows the rest of the group to continue searching the area while some of the aircraft investigate the discovery. Reinforcement learning has been applied to the same cooperative search problem within a

digital-twin training framework (Shen et al., 2023). A multi-agent formulation of swarm target search follows the same direction (Hou et al., 2024).

The closest prior study evaluates a context-aware deep Q-network on discrete uniform and clustered target layouts (Mohanty et al., 2020). The study found that missions finished faster when targets were clustered together. It explained this by saying that several targets within the same cluster could be dealt with at the same time. However, that study only compared two levels of clustering. In this study, clustering is varied gradually across many levels. More importantly, the effect of dealing with several targets at the same time is directly controlled in the experiment instead of simply being used as an explanation for the results.

1.4. Search theory and the role of target spatial distribution

The prediction that motivates the experiments is based on classical search theory. The main idea is that search effort should be allocated according to where the target is most likely to be found (Koopman, 1957). This idea was later developed into a general theory of optimal search (Stone, 1975). This theory leads to two important cases. When targets are uniformly distributed, uniform coverage is optimal because every location has the same probability of containing a target. It leaves little for an adaptive strategy to exploit. In contrast, when targets are spatially clustered, some regions are more likely to contain targets than others. Concentrating search effort in these regions can be more efficient than maintaining uniform coverage. However, this difference matters mainly when search resources are limited. If a UAV has unlimited endurance and enough time to complete a full sweep, both strategies can eventually find all static targets. Therefore, the expected advantage of adaptive search is not that it can find targets that uniform search cannot reach, but that it can use limited flight time or energy more efficiently when targets are clustered.

Foraging theory also shows that the effectiveness of a search strategy depends on how targets are distributed. When targets are sparse and locations can be revisited, an inverse-square power-law distribution of flight lengths has been reported to provide an efficient search strategy (Viswanathan et al., 1999). As the conditions of the search change, a transition between different search regimes has been observed. In particular, the efficient movement pattern can shift from superdiffusive movement, which includes occasional long movements, to more direct ballistic movement (Santos et al., 2004). However, a wide-ranging search strategy that is effective for sparse targets is not necessarily optimal for all types of blind search (Palyulin et al., 2014). Together, these findings suggest that the spatial distribution of targets can determine which search strategy is most effective. This leads to the prediction that a similar transition may occur in geometric multi-UAV search as the level of target clustering changes. Identifying this transition is the main objective of this study.

1.5. Modeling clustered spatial distributions

Spatial point-process theory provides a systematic way to vary target clustering. Spatially independent targets are modeled using a homogeneous Poisson point process, which provides a baseline with no spatial clustering. Clustered target fields are modeled using Neyman–Scott Poisson cluster processes, with the Thomas process used as the specific model (Thomas, 1949). The properties of this process family are described in the standard reference on stochastic geometry (Chiu et al., 2013). In the Thomas process, offspring points are distributed around randomly placed parent points according to a Gaussian kernel. The spread of these offspring points is controlled by the parameter σ , allowing the level of clustering to be varied continuously. This parameter also provides a measure of the characteristic spatial scale of the target clusters and can therefore be related to the physical correlation lengths of the anomalies. To account for differences between the specified and resulting target distributions, the realized clustering statistic is reported together with the nominal σ . This is

important since the removal of overlapping targets and the treatment of targets near the search-area boundaries can change the final cluster size.

1.6. Position of this study

This study brings together several research directions that are usually considered separately. Systematic coverage path planning, adaptive search based on detection feedback, and controlled spatial clustering are evaluated within the same mission framework and under a common detection model. The study examines how the relative performance of these search strategies changes as the level of target clustering increases. This change is evaluated for specific geometric search patterns under a fixed time budget, while the sensor's ability to detect multiple nearby targets is explicitly controlled. The effect of simultaneous target detection is first measured, bounded, and then removed by construction. As a result, the observed transition between search strategies can be attributed primarily to the spatial distribution of targets rather than to differences in sensor geometry. The clustering-dependent crossover is therefore treated as the main object of measurement in the study.

2. Methodology

2.1. Problem statement

A bounded planar survey region (A) represents a monitored industrial area containing (N) anomalies at positions $x_1, \dots, x_N$. The anomalies represent events such as disaster victims, wildfire ignition points, industrial leaks, infrastructure damage, and equipment failures, and their spatial distribution can range from approximately independent to strongly clustered. A team of ($K=3$) identical UAVs operates within (A) at a fixed altitude. Each UAV flies at a constant cruise speed of approximately $v \approx 5$ m/s and carries a sensor with a $71.1 \text{ m} \times 53.3 \text{ m}$ footprint, giving a diagonal of approximately 89 m. A target is considered detected when it is first covered by the sensor footprint. Each trial is limited by a fixed mission time budget (T), which is shorter than the time T_{sweep} required to complete a full systematic sweep of the survey region. The ratio T/T_{sweep} is reported to express the available search time relative to the time required for complete coverage, allowing the results to be compared across different UAV platforms and operating conditions.

The cumulative number of detected targets at time t is denoted by $D(t)$. Two timing metrics are used to describe the speed of detection: T_{50} , the time required to detect 50% of the targets, and T_{90} , the time required to detect 90% of the targets. When a detection threshold is not reached before the end of the mission, the corresponding time is treated as censored at T .

Two sets of experiments are conducted. The first uses the total number of detected targets at the end of the mission, $D(T)$, as the primary performance metric. The second focuses on detection speed using T_{50} and T_{90} . These experiments are used to compare systematic coverage and adaptive search across different levels of spatial target clustering. The main objective is to identify the clustering level at which the relative performance of the two strategies changes. This performance crossover can then provide a practical rule for selecting a search strategy: systematic coverage is expected to be more suitable for dispersed anomalies, while adaptive search is expected to become more effective as anomalies become increasingly clustered.

2.2. Proximity assumption

An adaptive search strategy is useful only when detecting one anomaly provides information about the likely locations of other anomalies. In particular, the detection of one anomaly should increase the probability that additional targets are located nearby. This spatial proximity is represented by the

cluster scale σ , which controls the strength and spatial extent of target clustering. By varying σ continuously, different levels of spatial dependence can be represented within the same model. Small values of σ produce tightly clustered targets, so detecting one target provides strong information about nearby targets. As σ increases, the targets become more widely dispersed and the spatial relationship between them becomes weaker. At sufficiently large values of σ , the target distribution approaches spatial independence, where detecting one target provides little information about the locations of others (Figure 1).




Figure 1
*Example target layouts generated by the Thomas cluster process for
 Experiment 1 and Experiment 2 (left to right)*

2.3. Modeling anomaly spatial distribution

Parent placement. Target fields are generated using a Thomas cluster process (Thomas, 1949). First, $N_p = 4$ parent points are placed inside the square $[-h, h]^2$, where the half-span is $h = 0.4L$ and L is the side length of the arena. The first parent is placed at a uniformly random location within this square. To place each of the remaining parents, 128 candidate locations are generated uniformly, and the candidate that is farthest from the nearest already-placed parent is selected. This farthest-point sampling procedure keeps the four parent points separated and prevents the cluster centers from being placed too close to one another.

Offspring assignment. A total of $N = 12$ target points are generated, with exactly three targets assigned to each of the four parents. The assignment is made using $i \bmod N_p$, which means that the first target is assigned to the first parent, the second to the second parent, and so on. After all four parents have received one target, the assignment starts again from the first parent. This process continues until all 12 targets have been assigned. The assignment is therefore deterministic and balanced, ensuring that every parent produces exactly three offspring.

Offspring scatter. After each target has been assigned to a parent, its position is generated by adding a random two-dimensional offset to the parent location c . The offset is drawn from an isotropic Gaussian distribution with a standard deviation of σ along both coordinate axes. Thus, targets tend to remain close to their assigned parent, while the value of σ controls how widely they are spread around

it. A small σ produces tight clusters, whereas a larger σ produces more dispersed clusters. The resulting probability density of an offspring around its parent is given in Equation 1.

$$f(x | c) = \frac{1}{2\pi\sigma^2} e^{\left(-\frac{\|x-c\|^2}{2\sigma^2}\right)} \quad (1)$$

Boundary constraint. If an offspring is generated outside the box, it is rejected and a new position is drawn. The point is not moved, or clamped, to the nearest boundary. Clamping would cause offspring generated with large σ values to accumulate along the boundary and would therefore distort the intended spatial distribution. Rejection instead produces a Gaussian distribution conditioned on the offspring remaining inside the box. As σ increases, this conditioned distribution approaches a uniform distribution within the box.

Minimum separation constraint. An optional minimum pairwise distance $s_{\min}$ can be imposed between all targets. If a Gaussian draw produces a point that violates this minimum distance, the point is rejected and another draw is attempted. If the minimum separation cannot be satisfied after 400 attempts, a ring-search fallback is used. In this fallback, the point is placed on a circle centered on its parent, and the radius of the circle is gradually increased until a position is found that is both inside the box and sufficiently separated from all existing targets. A final check is then performed to verify that the required minimum separation has been satisfied. If the constraint is still violated, an error is raised and the invalid layout is not accepted.

The ring fallback is important when interpreting the effect of σ . When σ is small and $s_{\min}$ is large, the minimum-separation constraint can prevent targets from occupying their original Gaussian positions. Instead, the fallback moves them onto concentric rings around the parent. Thus, a hard minimum separation places a lower limit on the range of $s_{\min}$ values that can represent the intended Gaussian clustering without being strongly affected by the separation constraint. The two experiments use different minimum-separation values. Experiment 1 uses $s_{\min} = 12$ m, which still allows sensor-footprint overlap. Experiment 2 uses $s_{\min} = 90$ m, which is greater than the 89 m sensor-footprint diagonal and therefore removes single-frame co-detection.

Pairwise-distance law. The Gaussian scatter also provides a direct relationship between the clustering parameter σ and the distance between targets within the same cluster. Two offspring from the same parent are generated using independent Gaussian offsets. Therefore, the difference between their positions is also Gaussian, with a variance of $2\sigma^2$ along each coordinate axis. The distance d between the two offspring consequently follows a Rayleigh distribution with scale $\sigma\sqrt{2}$. The probability that two offspring from the same parent are located within a distance r of each other is therefore given by Equation 2.

$$P(d \leq r) = 1 - e^{\left(-\frac{r^2}{4\sigma^2}\right)} \quad (2)$$

Equation 2 quantifies the proximity between targets within the same cluster. For a given distance r , a smaller σ produces a higher probability that two targets are close to each other, while a larger σ produces a lower probability. This provides a quantitative representation of the proximity assumption used by the adaptive search strategy: detecting one target is more informative about the location of other targets when σ is small. The same relationship also shows how the clustering scale interacts with the hard minimum-separation constraint, since sufficiently small values of σ become difficult to realize when targets are required to remain at least $s_{\min}$ apart.

2.4. Search strategies

All strategies operate within the same multi-UAV framework. The survey region is divided into K sub-regions, with one sub-region assigned to each UAV. This decomposes the multi-UAV coverage problem into independent single-UAV coverage tasks (Rekleitis et al., 2008). Each UAV then follows the coverage pattern assigned to its sub-region independently. The sensing, target confirmation, flight control, and data logging modules are kept identical across all strategies. The only component that differs is the sequence of waypoints used to determine each UAV's trajectory. This ensures that differences in performance are caused by the search strategy rather than by differences in the underlying UAV or sensing systems.

2.4.1. Systematic geometric baselines

Three systematic coverage patterns are used as baselines. The lawnmower, or boustrophedon, pattern consists of parallel back-and-forth sweeps with lane spacing determined by the sensor swath, providing systematic and uniform coverage of the assigned sub-region (Figure 2). It serves as the primary baseline because of its widespread use in coverage path planning. The spiral pattern follows an inward or outward spiral trajectory within each sub-region and provides a systematic alternative for environments where targets may be concentrated toward the center. The Hilbert space-filling curve provides another systematic traversal pattern that preserves spatial locality while covering the assigned area. All three baselines use the same lane spacing, which is determined by the sensor footprint and the reliable-detection fraction. Therefore, the baselines have the same coverage density and differ only in the order in which the area is traversed.




Figure 2

Systematic coverage algorithms: lawnmower, spiral, and Hilbert traversal (left to right)

2.4.2. Detection-driven adaptive strategy

The proposed strategy uses the systematic Hilbert path as its main coverage plan and adds a local search response when a target is detected. Each UAV follows its assigned Hilbert trajectory in the same way as in the baseline, and all confirmed detections are stored in a shared registry. When a target is confirmed, the nearest available UAV temporarily leaves its planned trajectory and searches the area around the detected target. This local search is bounded so that the UAV does not move indefinitely away from its assigned path. After completing the local search, the UAV returns to the point where it left the Hilbert trajectory and continues along the original path. The Hilbert coverage path is therefore not replanned or replaced during the mission. Instead, the adaptive strategy adds short, detection-triggered detours to the existing systematic path. If additional targets are located near the confirmed target, the detour can result in additional detections. If no additional targets are present,

the UAV returns after only a limited additional flight distance, keeping its behavior close to that of the original Hilbert strategy. The following paragraphs describe the detection, detour, and return mechanisms in detail.

Exploration layer. A single Hilbert space-filling curve is first generated over the entire survey region, with the lane spacing determined by the sensor footprint. The resulting sequence of waypoints is then divided into K contiguous segments, with one segment assigned to each UAV. During normal exploration, each UAV follows the waypoints in its assigned segment in sequence. A waypoint is considered reached when the UAV comes within 5 m of its location. This exploration procedure is identical to that used by the systematic Hilbert baseline and is not changed by the adaptive strategy. Therefore, when no exploitation detours are active, the proposed strategy follows exactly the same waypoint sequence as the Hilbert baseline.

Exploitation layer. A confirmed target can trigger at most one local detour. Let c denote the position of the confirmed target. When a detour is triggered, the UAV closest to c is selected from the UAVs that are not already performing another detour. The selected UAV therefore does not have to be the UAV that detected the target. It temporarily leaves its assigned Hilbert path and performs one complete lap around the detected target. The lap consists of K_h waypoints placed on a circle with radius R_{halo} centered at c . If the target is close to the boundary of the survey region, the ring waypoints are clamped to remain inside the region so that the resulting path remains flyable. After completing the lap, the UAV returns to its original exploration path and continues from the exact waypoint index at which the detour began.

Fold rule. When a new target is confirmed, its position is compared with the centers of detours that have already been registered. If the new confirmation is within the specified fold radius of an existing detour center, no new detour is created. Instead, the new target is treated as another target belonging to the same local cluster, since the existing detour is already intended to search that area. A detour center is registered immediately when its confirmation passes the fold check, even if the corresponding detour has not yet started and is still waiting in the queue. As a result, later targets detected near a queued or active detour are folded into that existing detour rather than creating additional rings. The fold rule is applied only when new confirmations arrive; it does not remove or cancel a detour that has already been queued.

Effort bounds. The adaptive strategy also ensures that at least one UAV that is still performing its assigned sweep remains on the exploration path at all times. A UAV that has already completed its assigned segment is not counted toward this requirement. If a confirmed target cannot be serviced immediately because no suitable UAV is available, the corresponding detour is placed in a queue. The queued detour is dispatched when a UAV becomes available, unless a later confirmation near the same detour center is folded into the existing detour. These rules limit the amount of flight time that can be redirected from systematic exploration to local exploitation.

Halo radius. The radius of the local detour is chosen so that the UAV searches the region where additional targets from the same cluster are most likely to be found. The choice of R_{halo} differs between the two experiments. In Experiment 1, the halo radius is linked to the assumed cluster scale and is set to $R_{\text{halo}} = 30$ m. This radius is smaller than the effective sensor footprint radius, so one complete lap around the detected target provides coverage of the entire disc around the detection, including its center. In Experiment 2, the halo radius is determined by the geometry of the detection and separation constraints rather than by the cluster scale. The hard minimum separation is $s_{\text{min}} = 90$ m, which means that no other target can be located within 90 m of a confirmed target. A target closer than this distance would have been detected in the same sensor frame, and such target configurations are excluded by the generator. Therefore, any undetected target from the same cluster must be located at a distance $r \geq s_{\text{min}}$ from the confirmed target. The sensor reliably detects targets

within a band of approximately $w = 27$ m on either side of the UAV's flight path. To search for targets beginning at the 90 m exclusion boundary, the inner edge of this detection band is placed on that boundary. This geometric relationship gives Equation 3.

$$R_{\text{halo}} = s_{\text{min}} + w = 90 + 27 = 117\text{m} \quad (3)$$

- One complete lap therefore searches the first region outside the 90 m exclusion zone where another undetected target could be located. No flight time is spent searching the inner region because it is guaranteed to contain no additional target.
- **Detour cost.** The flight cost of each detour is bounded and can be calculated before the mission begins. Each detour consists of one complete lap around the confirmed target. The circular halo is represented by a regular K_h -gon, where K_h equally spaced waypoints are placed on the halo circle and connected by straight flight segments. The lap therefore follows the perimeter of the inscribed K_h -gon rather than the circumference of the circle. Its length is given by Equation 4.

$$L_{\text{ring}} = 2K_h R_{\text{halo}} \sin \left(\frac{\pi}{K_h} \right) \quad (4)$$

The distance required for the UAV to leave its original sweep and reach the nearest waypoint of the K_h -gon is added to the polygon perimeter, together with the distance required to return to the sweep after completing the lap. Including these entry and exit legs, the total detour distance is approximately 176 m per lap in Experiment 1 and approximately 716 m per lap in Experiment 2. At a constant speed of $v = 5$ m/s, these distances correspond to approximately 35 s and 143 s, respectively. The cost of each detour is therefore fixed by the predefined K_h -gon geometry and does not depend on how many additional targets are found during the lap. This bounded cost allows the adaptive strategy to investigate detected targets without allowing a single detection to generate unlimited additional search effort. When targets are clustered, one K_h -gon lap can detect nearby targets before the systematic sweep reaches the same region again. When a detected target is isolated, the detour still incurs only its predefined cost, keeping the adaptive strategy relatively close to the systematic Hilbert baseline in dispersed environments.

2.5. Detection model

Search strategy and target perception are treated as separate components of the experimental framework. The purpose of the study is to evaluate how different search strategies respond to the spatial distribution of targets, rather than to compare the performance of specific deep-learning detectors. Therefore, a lightweight detector based on hue–saturation–value (HSV) color segmentation is used. In the simulation, ground anomalies are represented by flat-colored markers with predefined and distinct HSV characteristics. A detection is registered when the number of pixels matching the target color within the camera footprint exceeds a predefined threshold. This produces detection events that depend on the UAV's position and sensor overflight while avoiding the need to train a deep-learning model. The detected image pixels are then projected into world coordinates using attitude-corrected ray casting, and the resulting detections are matched to the corresponding simulation timestamps. This deterministic perception pipeline provides a controlled proxy for the target detections that would be produced by a real-time computer-vision system in an operational deployment.

2.6. Metrics and statistical design

All performance metrics are computed from the same shared pose and detection log streams for all four strategies, ensuring that every strategy is evaluated using the same data and scoring procedure. Temporal performance is evaluated using three measures.

- **Per-target latency (L).** The time from mission start until each target is first detected is recorded as its latency. When a target is not detected before the mission ends, its latency is censored at the mission time limit. The resulting per-target latencies are reported together with their trial mean.
- **Milestone attainment (T_{50}, T_{90}).** These measures represent the time required to detect 50% and 90% of all targets, respectively. If a milestone is not reached before the mission time limit, its attainment time is censored at the budget limit. Each milestone is reported both as a censored mean time and as an attainment rate, defined as the fraction of trials in which the corresponding detection threshold is reached.
- **Cumulative detection curves ($D(t)$).** These curves show how the number of detected targets changes throughout the mission. They capture the complete detection progression and therefore distinguish strategies that achieve the same final number of detections but reach those detections at different rates.

Total fleet distance traveled is reported alongside the performance metrics to determine whether any advantage of the adaptive strategy is associated with increased travel. The crossover analysis relates strategy performance to the realized clustering statistics, including per-layout dispersion and nearest-neighbor distances. This allows the change in relative strategy performance to be examined as the actual spatial distribution of targets changes.

2.7. Experimental Setup

Three UAVs operate at an altitude of 30 m and a cruise speed of approximately 5 m/s. The onboard camera has a horizontal field of view of 1.74 rad, producing a ground footprint of 71.1 m $\times$ 53.3 m with an 89 m diagonal. Computer vision processing is restricted to the central 75% of the image to improve detection reliability, giving a reliable cross-track half-width of approximately $w \approx 27$ m. All four search strategies use the same sweep lane spacing, derived using a 0.95 overlap factor. Each target field contains 12 anomalies generated by a Thomas cluster process with four parent locations and three offspring per parent.

2.7.1. Experiment 1 (400 m arena)

Performance was evaluated using a full-factorial stochastic simulation study that varied six clustering levels, $\sigma \in \{8, 20, 40, 80, 160, 400\}$ m, four search strategies, and ten pseudo-random seeds. The static arena measured 400 m $\times$ 400 m, and each trial was limited to a 120 s flight budget, corresponding to approximately $0.55 T_{\text{sweep}}$. A minimum target separation of 12 m was imposed, which still allowed multiple targets to fall within the sensor footprint and be detected in the same frame. The primary performance metric was the cumulative number of detected targets $D(T)$ (Table 1).

Three structural limitations motivated the scale-up to a second experiment. First, the 89 m sensor-footprint diagonal at 30 m altitude was large relative to the spatial scale of the tighter clusters. As a result, clustering levels with $\sigma \leq 40$ m frequently produced frame-level co-detections. This introduced a sensor-dependent effect that could not be cleanly separated from the effect of spatial clustering itself. Increasing the minimum target separation to 90 m would have removed this co-detection effect, but this constraint could not be imposed reliably within the 400 m arena without substantially restricting the range of achievable clustered layouts. The resulting layouts would become more dispersed and would weaken the intended clustering variation. Second, the 120 s flight budget allowed the UAVs to complete only approximately 1.5 legs of the spiral pattern, making this baseline poorly represented under the chosen time budget. Third, ten-seed replication became unnecessarily redundant once the statistical analysis was based on pooled regime-level comparisons.

2.7.2. Experiment 2 (1000 m separation-controlled arena)

Experiment 2 expanded the arena to $1000 \text{ m} \times 1000 \text{ m}$ to address the geometric limitations identified in Experiment 1, while keeping the sensor, platform, and fleet parameters unchanged. The larger arena allowed a hard minimum target separation of $s_{\min} = 90 \text{ m}$ to be imposed. Because this separation exceeds the 89 m sensor-footprint diagonal, frame-level co-detection was eliminated. The clustering range was adjusted to $\sigma \in \{60, 120, 250, 1000\} \text{ m}$. Here, $\sigma = 60 \text{ m}$ represents the tightest clustering level that can be generated under the 90 m separation constraint, while $\sigma = 1000 \text{ m}$ provides the approximately uniform control condition. The flight budget was increased to 800 s, corresponding to approximately $0.54 T_{\text{sweep}}$ for the larger arena. This preserved a similar level of resource scarcity to Experiment 1 while allowing the spiral baseline to complete approximately four full rings, avoiding the structural limitation observed in the smaller arena. The adaptive detour geometry was also rescaled to account for the increased target separation. Replication was reduced to five paired seeds per experimental cell, and the primary metric focus was shifted from cumulative detection count to paired detection latency (Table 1). With frame-level co-detection removed, fixed-budget detection counts provided fewer distinct outcome levels, reducing their statistical resolution. Paired latency differences provided a more sensitive measure of the relative performance of the strategies under the same experimental conditions.

Table 1
Experiment settings

Parameter	Experiment 1	Experiment 2
Area	$400 \text{ m} \times 400 \text{ m}$	$1000 \text{ m} \times 1000 \text{ m}$
Min target separation	12 m	90 m
Mission budget	120 s	800 s
σ levels	8 / 20 / 40 / 80 / 160 / 400 m	60 / 120 / 250 / 1000 m
Seeds per cell	10	5
Trials	240	80
Primary metric	fraction found within budget	time-to-detection (latency, T50/T90, curves)
Adaptive detour geometry	$R_{\text{halo}} = 30\text{m}, K_h = 5$, fold 45 m	$R_{\text{halo}} = 117\text{m}, K_h = 8$, fold 60 m

3. Results

3.1. Scope of the evaluation

Four search strategies were evaluated. Three are fixed systematic coverage strategies: lawnmower, spiral, and a Hilbert space-filling curve. The fourth is the proposed adaptive strategy, which follows the Hilbert sweep and adds detection-triggered halo-ring detours. Target locations were generated using a Thomas cluster process, with the cluster scale σ serving as the single controlled variable. Small values of σ produce tightly clustered targets, whereas large values produce increasingly dispersed layouts that approach a uniform distribution. The same target layout was used for all four strategies within each seed, so the methods were evaluated on paired and directly comparable conditions.

The two experiments were conducted at different spatial scales. Experiment 1 used a $400 \text{ m} \times 400 \text{ m}$ arena with three UAVs, a 120 s mission budget, and a 12 m minimum target separation. It evaluated six clustering levels with ten paired seeds per level and strategy, resulting in 240 trials. Experiment 2 used a $1000 \text{ m} \times 1000 \text{ m}$ arena with an 800 s mission budget and a 90 m minimum target separation. It evaluated four clustering levels with five paired seeds per level and strategy, resulting in 80 trials.

The primary metric in Experiment 1 was the fraction of the 12 targets detected within the mission budget. Detection time was used as a secondary metric and summarized using the censored mean detection time and T_{50} . Detection time was the primary metric in Experiment 2. The T_{90} milestone was not reached in any trial: 0 of 10 trials in every Experiment 1 cell and 0 of 5 trials in every Experiment 2 cell reached 90% detection. Thus, no strategy detected 90% of the targets within the available mission time, confirming that both experiments operated within the intended resource-limited regime.

3.2. Experiment 1: detection yield and the clustering crossover

Table 2 reports the fraction of targets detected within the 120 s mission budget. The paired difference is calculated as the adaptive strategy score minus the Hilbert strategy score for each seed and is reported in targets per seed, together with a bootstrap 95% confidence interval and the corresponding win/tie/loss counts.

Table 2
Fraction of targets found within budget, Experiment 1

Cell (σ , m)	Lawnmower	Spiral	Hilbert	Adaptive (proposed)	Paired Δ targets [95% CI] (W/T/L)	Wilcoxon p
C1 (8)	0.433 $\pm$ 0.192	0.158 $\pm$ 0.114	0.550 $\pm$ 0.119	0.675 $\pm$ 0.121	+1.50 [+0.80, +2.20] (7/3/0)	0.016
C2 (20)	0.450 $\pm$ 0.177	0.200 $\pm$ 0.098	0.567 $\pm$ 0.135	0.658 $\pm$ 0.154	+1.10 [+0.20, +2.10] (5/4/1)	0.094
C3 (40)	0.483 $\pm$ 0.123	0.217 $\pm$ 0.090	0.508 $\pm$ 0.133	0.550 $\pm$ 0.153	+0.50 [-0.40, +1.40] (4/4/2)	0.438
C4 (80)	0.442 $\pm$ 0.125	0.300 $\pm$ 0.081	0.533 $\pm$ 0.137	0.400 $\pm$ 0.086	-1.60 [-2.40, -0.80] (1/1/8)	0.012
C5 (160)	0.475 $\pm$ 0.104	0.300 $\pm$ 0.148	0.533 $\pm$ 0.172	0.400 $\pm$ 0.110	-1.60 [-2.40, -0.70] (1/2/7)	0.016
C6 (400)	0.475 $\pm$ 0.176	0.275 $\pm$ 0.111	0.450 $\pm$ 0.131	0.350 $\pm$ 0.141	-1.20 [-2.10, -0.40] (1/3/6)	0.047

At the tightest clustering level, $\sigma = 8$ m, the adaptive strategy achieved the highest detection rate. It detected a mean of 0.675 of the available targets, corresponding to 8.1 of the 12 targets per trial. The best systematic baseline was the Hilbert strategy, with a mean detection fraction of 0.550, or 6.6 targets per trial *Figure 3*. Thus, the adaptive strategy detected approximately 23% more targets than Hilbert at this clustering level. Because the layouts were paired across strategies, the per-seed comparison provides a more direct measure of the difference. The adaptive strategy outperformed Hilbert on 7 of 10 seeds and tied on the remaining 3, with no losses. The mean paired advantage was +1.50 targets, meaning that the adaptive strategy found, on average, 1.5 more targets than Hilbert on the same layouts. The bootstrap 95% confidence interval for this advantage was [+0.80, +2.20] targets, indicating that the estimated improvement remained positive across the reported interval. A Wilcoxon signed-rank test on the paired differences gave $p = 0.016$, providing evidence that the observed advantage was not likely to arise from random variation alone under the test's null hypothesis of no systematic paired difference.




Figure 3

Adaptive sweep advantage in found target ratio

The adaptive advantage did not decrease smoothly to zero as the clusters became more dispersed. Instead, performance changed from a positive to a negative advantage near the scale of the camera footprint. The advantage remained clear at $\sigma = 20$ m, where the adaptive strategy achieved a mean detection fraction of 0.658 compared with 0.567 for Hilbert, with a paired advantage of +1.10 targets. At $\sigma = 40$ m, the difference was smaller, with a paired advantage of +0.50 targets, but this difference was not statistically significant. From $\sigma = 80$ m onward, the adaptive strategy performed worse than Hilbert, with significant paired deficits of -1.60 targets at C4, -1.60 targets at C5, and -1.20 targets at C6. Thus, the paired advantage changes sign between C3 and C4. Interpolation on the logarithmic σ scale places the crossover at approximately $\sigma = 47$ m. With a sensor footprint radius of 36 m at 30 m altitude, this corresponds to a cluster-to-footprint ratio of approximately 1.3.

Total fleet distance traveled was nearly identical for the adaptive and Hilbert strategies across all clustering levels, with differences of only a few meters. This indicates that the observed performance difference was not caused by substantially greater overall flight distance, but by how the two strategies allocated that flight effort between systematic coverage and detection-triggered detours.

3.3. Experiment 2: separating the search mechanism from frame-level co-detection

The co-detection confound in Experiment 1 - The target-count advantage observed in Experiment 1 is partly confounded by the sensor geometry. As target clustering becomes stronger, the probability that multiple targets fall within the same camera footprint increases. When two targets are separated by less than the 89 m footprint diagonal, an overflight can detect both targets within the same frame. This means that a higher detection count may partly reflect the ability of the sensor to detect multiple nearby targets simultaneously, rather than the ability of the search strategy to exploit spatial clustering. This effect was therefore measured directly. Table 3 reports three quantities for each strategy and clustering level: the proportion of target pairs separated by less than the 89 m footprint diagonal, the percentage of targets with at least one co-detectable neighbor, and the percentage of detected targets that were co-detected with at least one neighbor.

Table 3 shows frame-level co-detectability and realized co-detection in Experiment 1. The four right-hand columns report the percentage of found targets that arrived through a co-detection event.

Table 3
Realized co-detection in Experiment 1

Cell (σ , m)	Pairs < 89 m (%)	Targets w/ neighbor (%)	Lawnmower	Spiral	Hilbert	Adaptive (proposed)
C1 (8)	18.2	100.0	33.4	19.0	42.1	31.1
C2 (20)	18.0	100.0	18.7	14.8	15.6	11.1
C3 (40)	15.0	95.0	13.2	0.0	9.6	5.8
C4 (80)	10.5	79.2	8.3	0.0	2.9	3.3
C5 (160)	12.4	70.8	0.0	2.9	4.0	0.0
C6 (400)	9.8	68.3	2.0	5.0	0.0	0.0

The proportion of target pairs separated by less than the 89 m footprint diagonal decreases from 18.2% at C1 to 9.8% at C6, although the decrease is not strictly monotonic. In particular, C5 has a higher proportion than C4, at 12.4% compared with 10.5%. The proportion of targets with at least one co-detectable neighbor provides a more stable measure of co-detectability and decreases monotonically from 100.0% at C1 to 68.3% at C6. Across the 60 seed layouts, this measure is strongly negatively correlated with $\log \sigma$, with a Pearson correlation of $r=-0.77$ ($p < 0.001$). This indicates that increasing spatial dispersion reduces the likelihood of co-detectable neighbors and confirms that σ systematically controls the extent of the co-detection effect.

The realized co-detection rates are substantially lower than the geometric availability would suggest. At C1, 19.0% to 42.1% of detected targets were obtained through a co-detection event, whereas at C6 the highest observed rate was only 5.0%. The effect therefore reflects the spatial structure of the tighter layouts rather than a systematic preference of one search strategy. In particular, the adaptive strategy did not produce more co-detections than Hilbert. Its realized co-detection rate was lower at C1, at 31.1% compared with 42.1%, and was also lower at C2, C3, and C5. It was marginally higher only at C4, where the rates were 3.3% for the adaptive strategy and 2.9% for Hilbert. At C6, both strategies had a co-detection rate of 0.0%.

The confounding effect nevertheless remains in another form. Across the 60 paired seeds, layout co-detectability is positively correlated with the paired advantage in detected targets, defined as adaptive minus Hilbert, with a Pearson correlation of $r=+0.60$ ($p < 0.001$). Thus, layouts with greater co-detectability also tend to be layouts in which the adaptive strategy achieves a larger detection advantage. However, this relationship is largely explained by differences between clustering levels. After removing the mean effect of σ within each cell, the correlation decreases to ($r=+0.14$) and is no longer significant ($p = 0.29$). Therefore, most of the observed ($r=+0.60$) relationship comes from the between-cell variation in σ , rather than from differences between layouts at the same clustering level. Since increasing σ simultaneously changes both the spatial opportunity for halo-based exploitation and the probability of co-detection, Experiment 1 cannot cleanly separate these two effects. Experiment 2 was therefore designed to remove the co-detection effect while retaining variation in the spatial clustering of targets.

Elimination of co-detection by construction. Experiment 2 imposes a minimum target separation of 90 m, which is greater than the 89 m camera-footprint diagonal. Consequently, no two targets can occur within the same camera footprint, and simultaneous detection of multiple targets is excluded by construction. The co-detection audit confirms this result, with every experimental cell reporting 0.0% of detected targets as co-detection events. Any performance difference observed in Experiment 2 therefore cannot be attributed to the ability of the sensor to detect multiple nearby targets in the same frame.

Detection counts - Table 4 reports the target-count results for Experiment 2. The difference between the adaptive and Hilbert strategies is substantially smaller than in Experiment 1 and does not reach statistical significance.

Table 4

Fraction of targets found within budget in Experiment 2

Cell (σ , m)	Lawnmower	Spiral	Hilbert	Adaptive (proposed)	Paired Δ [95% CI]	W/T/L	Wilcoxon p
C1 (60)	0.450 ± 0.075	0.217 ± 0.126	0.550 ± 0.075	0.633 ± 0.095	$+1.00 [+0.20, +1.80]$	3/2/0	0.250
C2 (120)	0.483 ± 0.091	0.417 ± 0.118	0.517 ± 0.124	0.533 ± 0.095	$+0.20 [+0.00, +0.60]$	1/4/0	1.000
C3 (250)	0.550 ± 0.126	0.533 ± 0.151	0.567 ± 0.124	0.583 ± 0.102	$+0.20 [-1.20, +1.40]$	3/0/2	1.000
C4 (1000)	0.567 ± 0.109	0.700 ± 0.075	0.467 ± 0.075	0.433 ± 0.124	$-0.40 [-1.00, +0.40]$	1/1/3	0.625

Pooled across the two clustered conditions, C1 and C2, the adaptive strategy detected an average of 0.60 more targets than Hilbert on the same target layouts (Figure 4). The 95% confidence interval for the paired difference was $[+0.10, +1.10]$ targets, so the estimated advantage remained positive across the reported interval. However, the Wilcoxon signed-rank test gave $p = 0.0625$, which is above the conventional 0.05 significance threshold. The clustered conditions therefore show a positive trend in favor of adaptive search, but the evidence is insufficient to establish a statistically significant advantage at the 5% level. At the approximately uniform extreme, C4 with $\sigma = 1000$ m, the spiral strategy achieved the highest mean detection fraction at 0.700. The adaptive strategy detected 0.40 fewer targets than Hilbert on average, although this difference was not statistically significant. Thus, the count results suggest a shift from a possible adaptive advantage in clustered environments toward no clear advantage, and potentially a disadvantage, as the target distribution becomes approximately uniform.




Figure 4

Detection count advantage of adaptive strategy in Experiment 2

Detection timing - Table 5 reports the censored mean detection time and the T50 milestone for Experiment 2.

Table 5

Detection timing in Experiment 2

Cell (σ , m)	Method	Censored mean detect (s)	T50 attained	T50 mean (s)	Paired Δ vs Hilbert (s)
C1 (60)	Hilbert	496 $\pm$ 33	5/5	482.8	—
C1 (60)	Adaptive	447 $\pm$ 49	5/5	275.3	-48 [-79, -27]
C2 (120)	Hilbert	515 $\pm$ 109	4/5	493.4	—
C2 (120)	Adaptive	495 $\pm$ 65	4/5	397.0	-20 [-57, +33]
C3 (250)	Hilbert	561 $\pm$ 67	4/5	623.7	—
C3 (250)	Adaptive	515 $\pm$ 58	4/5	498.6	-46 [-96, +4]
C4 (1000)	Hilbert	572 $\pm$ 76	4/5	544.5	—
C4 (1000)	Adaptive	588 $\pm$ 76	2/5	536.4	+16 [-0, +40]

Pooled across the clustered cells, the adaptive strategy reduced the censored mean detection time by 34.19 s compared with Hilbert. The 95% confidence interval for the paired difference was [-59.45, -3.66] s, with the entire interval below zero, indicating a consistent reduction in detection time across the reported uncertainty range. A Wilcoxon signed-rank test gave $p=0.0322$, providing evidence of a statistically significant reduction at the conventional 5% level. The paired per-seed comparison also shows a clear pattern: the adaptive strategy was faster on 9 of the 10 seeds, tied on none, and was slower on only 1 seed. The improvement is also visible in the T_{50} milestone. At C1, the adaptive strategy reached 50% target detection in a mean of 275.3 s, compared with 482.8 s for Hilbert, while both strategies reached the T_{50} milestone in all five seeds. This shows that the adaptive strategy did not merely increase the final number of detected targets; under clustered conditions, it also reached the same intermediate detection level substantially earlier.




Figure 5
Mean detection time to target comparison in Experiment 2

At the approximately uniform extreme, the timing advantage of the adaptive strategy disappears, consistent with the detection-count results. At C4, the adaptive strategy reached the T_{50} milestone in only 2 of 5 seeds, compared with 4 of 5 seeds for Hilbert. Its censored mean T_{50} was also 16 s longer than that of Hilbert, indicating that the adaptive detours provided no timing benefit when targets were approximately uniformly distributed (Figure 5).

3.4. Ablation

With detours disabled, the adaptive planner reduces to the Hilbert sweep by construction. On the shared test seed at $\sigma = 8$ m in Experiment 1, the ablated planner produced exactly the same result as the Hilbert baseline, detecting 6 of 12 targets. A separate mechanism-identity check in the 1000 m arena recorded zero halo-ring detours when the maximum detour count was set to zero. These checks confirm that the exploration layers are identical and that any performance difference between the full adaptive strategy and Hilbert is attributable to the detection-triggered detour mechanism.

4. Discussion

4.1. A crossover set by the sensor, not by the strategy

The most informative feature of the Experiment 1 results is not the magnitude of the gain, but the location of the sign change. The adaptive advantage does not decrease gradually to zero as the clusters become more dispersed. Instead, it changes from positive to negative at a cluster scale close to the reliable sensor-footprint radius, with a cluster-to-footprint ratio of approximately 1.3. This suggests that the sensor geometry plays an important role in determining the boundary between the two search regimes. When the cluster scale is small relative to the region that can be resolved by the sensor, detecting one target provides useful information about the likely location of other targets. A bounded halo lap can then use this information to detect nearby targets before the systematic sweep reaches the same area again. As the cluster scale becomes larger than the footprint, a detection provides less useful information within the spatial range that the detour can cover, and the same lap is more likely to become an unproductive excursion. The crossover should therefore depend on the relationship between the target-cluster scale and the sensor's effective search scale. Changes in altitude, field of view, or lane spacing could consequently shift the crossover, rather than the crossover being determined only by the specific detour heuristic. This interpretation is consistent with regime crossovers reported in random-search studies, where the efficient movement pattern changes as the spatial scale of the targets changes relative to the searcher (Santos et al., 2004).

The distance measurements provide further support for this interpretation. The adaptive and Hilbert strategies traveled nearly identical total fleet distances at every clustering level, so the poorer performance of the adaptive strategy in dispersed layouts cannot be explained by substantially greater travel. Instead, the adaptive strategy used part of the same available flight effort for detours around detected targets, including targets that did not provide useful information about nearby anomalies. Under the fixed mission-time budget, this represents an allocation of limited search effort rather than a reduction in total travel capacity. The result is consistent with classical search theory, where concentrating search effort is beneficial when the target-location prior is concentrated, but provides less benefit when the prior approaches a uniform distribution (Koopman, 1957).

4.2. Why the advantage moves from count to time

The two experiments examine the same adaptive mechanism under different constraints, and the difference in their results is informative. In Experiment 1, a substantial part of the observed signal came from targets that could be detected within the same camera footprint. This effect was recorded mainly through an integer detection count under a short mission budget. Experiment 2 removed this channel by imposing a 90 m minimum separation between targets. The remaining effect was the ability of the adaptive detour to visit the region around a detected target before the systematic sweep would reach it. The systematic sweep eventually reaches the same region, so the detour does not necessarily increase the final number of detected targets when sufficient mission time is available. Its main effect is to change the order in which targets are detected.

This difference explains why the detection-count advantage in Experiment 2 is small and does not reach statistical significance, whereas the latency advantage does. The adaptive mechanism primarily changes when targets are detected rather than how many are eventually detected. Final detection counts under a fixed time budget therefore provide a relatively low-resolution measure of this effect, while paired detection latency captures the timing improvement directly. The paired latency results are also consistent across seeds: the adaptive strategy was faster on 9 of the 10 clustered pairs and slower on only one. This consistency provides evidence that the observed timing advantage reflects the search mechanism rather than a particularly favorable target layout.

4.3. Operational implications

The results can be translated into a decision rule that can be applied before a mission begins. The key quantity is the expected spatial correlation length of the anomalies relative to the reliable ground footprint of the sensor at the planned flight altitude. When the expected correlation length is smaller than or comparable to the footprint scale, a detected anomaly provides useful information about the likely location of nearby anomalies, and bounded detours can improve search performance. Examples include connected oil-slick spills and spatially extended corrosion zones. When anomalies are more isolated and their correlation length is large relative to the footprint, a detection provides less information about nearby targets, and a systematic sweep is more appropriate. Sporadic third-party damage and isolated manufacturing defects are examples of this regime.

The mission objective is equally important. When the priority is to minimize the time required to provide responders with confirmed anomaly locations, as in emergency spill response or search-and-rescue missions, detection latency is the relevant performance measure. Under clustered target conditions, the results support using the adaptive strategy in this setting. When the objective is instead to maximize the total number of anomalies detected during a routine patrol with a longer time budget, the benefit of adaptation is smaller because the systematic sweep eventually reaches regions explored by the detours. In this case, adaptation may provide little additional benefit and can introduce a small time cost. The bounded detour design also limits the cost of using adaptation when the target distribution is uncertain. A concurrency cap, a fixed lap count, and a minimum sweeper floor restrict the amount of flight effort that can be diverted from systematic coverage. Thus, an incorrect assumption about target clustering results in a known and bounded cost rather than an uncontrolled loss of coverage.

4.4. Limitations

Several limitations should be considered when interpreting these results. First, all targets are static and present from the beginning of the mission. The findings therefore do not directly extend to moving targets, such as drifting oil slicks, or to anomalies that emerge during the patrol. Second, target detection is modeled using a deterministic color segmenter. This provides controlled and repeatable detections but does not reproduce the false positives and missed detections that can occur with a learned computer-vision model. A noisier detector could trigger detours from false confirmations and could therefore reduce the measured benefit of the adaptive strategy. Third, the clustering generator uses a fixed structure of four parent locations with three offspring per parent. Consequently, the number of clusters and the number of targets within each cluster are not varied independently of σ . Fourth, the number of independent seeds is modest, with ten seeds per cell in Experiment 1 and five in Experiment 2. This limits statistical power and contributes to the wide confidence intervals and non-significant p -values observed in several Experiment 2 cells, even when the per-seed differences show a consistent direction. Finally, both experiments use square survey regions and a fleet of three UAVs. The results therefore do not establish how the detour allocation mechanism behaves with larger fleets or under different region geometries, where the interaction between adaptive detours and spatial partitioning may change.

5. Conclusion

This study investigated the level of spatial clustering at which detection-driven adaptive multi-UAV search begins to outperform systematic geometric coverage under a fixed time budget. Spatial clustering was treated as a controlled experimental variable rather than as an uncontrolled property of the target layouts. Three systematic sweep geometries were compared with a minimal adaptive strategy using identical target layouts. The adaptive strategy used the Hilbert sweep as its exploration layer and differed from it only through bounded detours triggered by confirmed detections. The comparison revealed a crossover between search regimes rather than a single strategy that was consistently superior. Under the tightest clustering condition in the 400 m arena, the adaptive strategy detected an average of 8.1 of 12 targets, compared with 6.6 for the best systematic sweep, corresponding to a 23% improvement. It outperformed Hilbert on 7 of 10 paired layouts and did not lose on any layout. The advantage decreased as the clusters became more dispersed, with the paired difference approaching parity at $\sigma = 40$ m and becoming negative at larger clustering scales, reaching deficits of 1.20 to 1.60 targets. Fleet distance remained nearly identical across the strategies at every clustering level, indicating that the reversal was caused by how the available flight time was allocated rather than by a difference in total travel. Interpolation placed the crossover at approximately $\sigma = 47$ m, or about 1.3 times the reliable cross-track sensing distance of the camera. Because this criterion is expressed as a ratio between the spatial scale of the target distribution and the effective sensing scale, it can be adapted to different flight altitudes and sensor configurations.

The second experiment imposed a minimum target separation greater than the sensor footprint diagonal. As a result, a single camera frame could no longer detect two targets simultaneously, removing the co-detection ambiguity present in the smaller arena. Under this constraint, the effect of adaptation was expressed more clearly through detection timing than through final target count. The pooled count advantage narrowed to 0.60 targets and did not reach statistical significance, whereas the censored mean detection time was reduced by 34.19 s, with a Wilcoxon test ($p = 0.0322$). The adaptive strategy was faster on 9 of the 10 paired layouts and slower on only one. At the tightest clustering level, the adaptive strategy reached 50% target detection in a mean of 275.3 s, compared with 482.8 s for Hilbert, and both strategies reached this milestone on all five layouts. These results indicate that the main benefit of exploitation is earlier detection rather than a consistent increase in the final number of targets detected. This distinction is particularly important in applications where the value of a detection decreases with time. Under approximately uniform placement, the timing advantage disappeared and the bounded detour became a cost, making systematic coverage preferable. Finally, the ablation test reproduced the Hilbert detection count exactly when detours were disabled, confirming that the observed difference between the two methods is attributable to the detection-triggered detour layer.

Three practical rules follow from these results.

- **Deployment.** Enable adaptive detours when the operational objective is to provide responders with confirmed anomaly locations as quickly as possible, particularly when the anomalies are expected to be spatially clustered.
- **Dispersion.** Prefer systematic coverage when anomalies are expected to be distributed over spatial scales larger than the reliable sensor footprint. If adaptive detours are retained under uncertain conditions, their maximum effort should be explicitly bounded so that the cost of an incorrect assumption about clustering remains limited.
- **Reporting.** Report target spacing relative to the effective sensor footprint. Without controlling or reporting this relationship, an observed advantage of an adaptive planner may be partly attributable to frame-level co-detection rather than to the planner itself.

More broadly, the choice between systematic and adaptive search should not be viewed simply as a choice between two planner families. The results indicate that it depends on a measurable relationship

between the spatial scale of the target distribution and the effective sensing scale of the UAV. Both quantities can be estimated before deployment, allowing the search strategy to be selected according to the expected operating regime. Whether the same relationship remains valid when targets move or appear during the mission is not established by the present experiments and provides a natural direction for future work. These scenarios will place greater emphasis on real-time detection latency, which is the setting in which adaptive coordination is expected to provide the greatest benefit. The strategies will also be tested in non-square and obstructed survey regions to determine whether the observed crossover depends on the geometry of the survey area. The framework will be scaled to larger and heterogeneous UAV swarms to examine how localized detour allocation interacts with spatial partitioning as the fleet size and platform capabilities increase. Finally, the deterministic color segmenter will be replaced with a learned detector that produces realistic false positives and missed detections. This will allow the sensitivity of the adaptive detour mechanism to detection quality to be evaluated under more realistic perception conditions.

References

- Allen, J., & Walsh, B. (2008). Enhanced oil spill surveillance, detection and monitoring through the applied technology of unmanned air systems. *International Oil Spill Conference Proceedings*, 2008(1), 113–120. <https://doi.org/10.7901/2169-3358-2008-1-113>
- Asadzadeh, S., de Oliveira, W. J., & de Souza Filho, C. R. (2022). UAV-based remote sensing for the petroleum industry and environmental monitoring: State-of-the-art and perspectives. *Journal of Petroleum Science and Engineering*, 208, Article 109633. <https://doi.org/10.1016/j.petrol.2021.109633>
- Bailon-Ruiz, R., Bit-Monnot, A., & Lacroix, S. (2022). Real-time wildfire monitoring with a fleet of UAVs. *Robotics and Autonomous Systems*, 152, Article 104071. <https://doi.org/10.1016/j.robot.2022.104071>
- Bourgault, F., Furukawa, T., & Durrant-Whyte, H. F. (2003). Coordinated decentralized search for a lost target in a Bayesian world. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)* (Vol. 1, pp. 48–53). IEEE. <https://doi.org/10.1109/IROS.2003.1250604>
- Cabreira, T. M., Di Franco, C., Ferreira, P. R., Jr., & Buttazzo, G. C. (2018). Energy-aware spiral coverage path planning for UAV photogrammetric applications. *IEEE Robotics and Automation Letters*, 3(4), 3662–3668. <https://doi.org/10.1109/LRA.2018.2854967>
- Cabreira, T. M., Brolora, L. B., & Ferreira, P. R., Jr. (2019). Survey on coverage path planning with unmanned aerial vehicles. *Drones*, 3(1), Article 4. <https://doi.org/10.3390/drones3010004>
- Chiu, S. N., Stoyan, D., Kendall, W. S., & Mecke, J. (2013). *Stochastic geometry and its applications* (3rd ed.). Wiley. <https://doi.org/10.1002/9781118658222>
- Choset, H. (2000). Coverage of known spaces: The boustrophedon cellular decomposition. *Autonomous Robots*, 9(3), 247–253. <https://doi.org/10.1023/A:1008958800904>
- Choset, H. (2001). Coverage for robotics – a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31(1), 113–126. <https://doi.org/10.1023/A:1016639210559>
- Choset, H., & Pignon, P. (1998). Coverage path planning: The boustrophedon cellular decomposition. In A. Zelinsky (Ed.), *Field and service robotics* (pp. 203–209). Springer. https://doi.org/10.1007/978-1-4471-1273-0_32
- Chung, T. H., & Burdick, J. W. (2012). Analysis of search decision making using probabilistic search strategies. *IEEE Transactions on Robotics*, 28(1), 132–144. <https://doi.org/10.1109/TRO.2011.2170333>
- Facinelli, D., Larcher, M., Brunelli, D., & Fontanelli, D. (2019). Cooperative UAVs gas monitoring using distributed consensus. In *2019 IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC)* (pp. 463–468). IEEE. <https://doi.org/10.1109/COMPSAC.2019.00072>

- Galceran, E., & Carreras, M. (2013). A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12), 1258–1276. <https://doi.org/10.1016/j.robot.2013.09.004>
- Hou, Y., Zhao, J., Zhang, R., Cheng, X., & Yang, L. (2024). UAV swarm cooperative target search: A multi-agent reinforcement learning approach. *IEEE Transactions on Intelligent Vehicles*, 9(1), 568–578. <https://doi.org/10.1109/TIV.2023.3316196>
- International Civil Aviation Organization, & International Maritime Organization. (2006). *International aeronautical and maritime search and rescue manual: Volume I. Organization and management* (5th ed.).
- Joshi, A. A., Bhatt, M. C., & Sinha, A. (2019). Modification of Hilbert’s space-filling curve to avoid obstacles: A robotic path-planning strategy. In *2019 Sixth Indian Control Conference (ICC)* (pp. 338–343). IEEE. <https://doi.org/10.1109/ICC47138.2019.9123166>
- Koopman, B. O. (1957). The theory of search: III. The optimum distribution of searching effort. *Operations Research*, 5(5), 613–626. <https://doi.org/10.1287/opre.5.5.613>
- Lyu, M., Zhao, Y., Huang, C., & Huang, H. (2023). Unmanned aerial vehicles for search and rescue: A survey. *Remote Sensing*, 15(13), Article 3266. <https://doi.org/10.3390/rs15133266>
- Miller, L. M., Silverman, Y., MacIver, M. A., & Murphey, T. D. (2016). Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1), 36–52. <https://doi.org/10.1109/TRO.2015.2500441>
- Mohanty, N., Gadde, M. S., Sundaram, S., Sundararajan, N., & Sujit, P. B. (2020). *Context-aware deep Q-network for decentralized cooperative reconnaissance by a robotic swarm* [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2001.11710>
- Palyulin, V. V., Chechkin, A. V., & Metzler, R. (2014). Lévy flights do not always optimize random blind search for sparse targets. *Proceedings of the National Academy of Sciences*, 111(8), 2931–2936. <https://doi.org/10.1073/pnas.1320424111>
- Rekleitis, I., New, A. P., Rankin, E. S., & Choset, H. (2008). Efficient boustrophedon multi-robot coverage: An algorithmic approach. *Annals of Mathematics and Artificial Intelligence*, 52(2–4), 109–142. <https://doi.org/10.1007/s10472-009-9120-2>
- Santos, M. C., Raposo, E. P., Viswanathan, G. M., & da Luz, M. G. E. (2004). Optimal random searches of revisitable targets: Crossover from superdiffusive to ballistic random walks. *Europhysics Letters (EPL)*, 67(5), 734–740. <https://doi.org/10.1209/epl/i2004-10114-9>
- Shakhatreh, H., Sawalmeh, A. H., Al-Fuqaha, A., Dou, Z., Almaita, E., Khalil, I., Othman, N. S., Khreishah, A., & Guizani, M. (2019). Unmanned aerial vehicles (UAVs): A survey on civil applications and key research challenges. *IEEE Access*, 7, 48572–48634. <https://doi.org/10.1109/ACCESS.2019.2909530>
- Shen, G., Lei, L., Zhang, X., Li, Z., Cai, S., & Zhang, L. (2023). Multi-UAV cooperative search based on reinforcement learning with a digital twin driven training framework. *IEEE Transactions on Vehicular Technology*, 72(7), 8354–8368. <https://doi.org/10.1109/TVT.2023.3245120>
- Spires, S. V., & Goldsmith, S. Y. (1998). Exhaustive geographic search with mobile robots along space-filling curves. In A. Drogoul, M. Tambe, & T. Fukuda (Eds.), *Collective robotics* (pp. 1–12). Springer. <https://doi.org/10.1007/BFb0033369>
- Stone, L. D. (1975). *Theory of optimal search*. Academic Press.
- Thomas, M. (1949). A generalization of Poisson’s binomial limit for use in ecology. *Biometrika*, 36(1–2), 18–25. <https://doi.org/10.1093/biomet/36.1-2.18>
- Viswanathan, G. M., Buldyrev, S. V., Havlin, S., da Luz, M. G. E., Raposo, E. P., & Stanley, H. E. (1999). Optimizing the success of random searches. *Nature*, 401(6756), 911–914. <https://doi.org/10.1038/44831>