
Research article

Adaptive discounted optimization for multidimensional random-walk targets in a two-stage stochastic transport model: Coverage and finiteness of first-encounter times

Mohamed Abd Allah El-Hadidy^{1,*} and Alaa Awad Alzulaibani²

¹ Mathematics Department, Faculty of Science, Tanta University, Tanta 31527, Egypt

² Department of Mathematics and Statistics, College of Science in Yanbu, Taibah University, Yanbu Governorate, Saudi Arabia

* **Correspondence:** Email: melhadidi@science.tanta.edu.eg.

Abstract: This paper developed an adaptive discounted two-stage stochastic transport model for locating and tracking a target governed by a multidimensional random walk. In the first stage, a discounted effort–reward criterion is incorporated into the regional search process to balance localization time, traveled distance, and detection probability. In the second stage, coordinated searchers track the target along intersecting trajectories while the target continues to move according to the multidimensional random-walk process. The tracking discount is adjusted according to the Stage-1 shell uncertainty, thereby linking the two stages within a unified framework. Coverage-preserving tracking distances are constructed to maintain recurrent exploration. Sufficient conditions are established for the existence of an optimal policy and complete coverage, while finiteness of the expected first-encounter time is obtained under an additional uniform positive finite-block encounter-probability assumption. The undiscounted model is recovered as the discount parameters approach 1. Numerical results illustrate the effects of discounting, uncertainty, and dimensionality on search effort and encounter performance.

Keywords: adaptive discounting; two-stage stochastic transport; multidimensional random walk; first-encounter time; coverage preservation

Mathematics Subject Classification: 49J20, 60G50, 60J10, 60J45, 60K35, 90B40

1. Introduction

Search problems under uncertainty appear in many settings, including environmental monitoring, hazardous-material detection, autonomous surveillance, rescue operations, and cooperative robotic sensing. The main difficulty is rarely limited to locating a stationary target. In practice, the target may move unpredictably, observations may be incomplete, and search resources may be limited. An effective strategy must therefore solve two related tasks: first, identify a likely initial location; second, track the target after it begins to move. This naturally leads to a two-stage formulation in which probabilistic localization is followed by stochastic tracking.

Random walks provide a standard framework for representing motion with uncertain increments, while first-passage and first-encounter times measure the delay until a searcher reaches or detects a target. Recent studies on stochastic resetting and finite-return mechanisms have shown that search performance depends on the target location, the reset distribution, and the admissible parameter regime, and that an optimal resetting strategy does not necessarily exist in every configuration [1–3].

Probability-based allocation provides a complementary approach. Rather than changing the stochastic dynamics through resetting, it directs search resources toward locations with larger prior probability. Recent multi-agent studies have used adaptive probability maps and stochastic optimization to improve spatial coverage and target localization under limited sensing resources [4,5]. These approaches are particularly relevant when target locations are uncertain and the available search resources are constrained.

Recent independent work has also developed complementary perspectives on stochastic search and first-passage phenomena. Boyer et al. studied optimization of capture probability and conditional first-passage time for finite-lived targets under random-walk and Lévy-flight search [6]. Méndez et al. analyzed Brownian first-passage times with resetting to randomly distributed positions [2]. Related work on moving-target search over complex networks used the mean first-encounter time to quantify capture performance for mobile searchers and targets [7]. These studies emphasize the broader role of first-passage optimization, stochastic motion, and encounter-time analysis beyond the specific transport framework considered here. Recent stochastic-control studies have also considered optimal tracking of stochastic dynamical systems, where control policies are designed to balance tracking accuracy and control effort under random disturbances [8]. This provides a complementary control-theoretic perspective to the stochastic search and encounter-time formulations considered here.

Recent methodological studies have also incorporated population-based optimization into complex dynamical models. For example, particle swarm optimization (PSO) has been used to tune flexible parameters in fractional-order octonion-valued neural-network synchronization problems [9]. Although that synchronization setting is different from the stochastic search problem considered here, it motivates the use of optimization algorithms as an additional numerical validation tool.

The theoretical foundations of stochastic search have been extensively developed through optimal search allocation, moving-target detection, and path-constrained search strategies. Stone [10] established necessary and sufficient conditions for optimal search plans involving moving targets, while Brown [11] investigated the allocation of search effort for targets evolving in discrete time and space. Eagle and Yee [12] subsequently addressed moving-target search under path constraints, providing an optimization procedure that accounts for restrictions on searcher movements. A broader treatment of probabilistic search planning, moving-target models, and associated optimization

methods was presented by Stone et al. [13]. Collectively, these studies provide a general framework for analyzing how target-location uncertainty, available search resources, and admissible search paths influence detection performance.

The analysis of first-passage and first-encounter times constitutes another important component of stochastic search theory. Redner [14] examined fundamental first-passage processes and their dependence on stochastic motion, spatial dimension, and boundary conditions. Bénichou et al. [15] investigated intermittent search strategies, emphasizing the influence of movement and detection phases on search efficiency. These studies highlight the importance of distinguishing eventual detection from the finiteness of the expected detection time. In particular, random target motion and repeated spatial exploration do not, by themselves, guarantee a finite expected first-encounter time. Additional assumptions concerning the search dynamics and encounter probabilities may be required. This distinction is essential in the present framework, where recurrent tracking coverage is combined with a uniform positive finite-block encounter-probability condition to establish finiteness of the expected encounter time.

Probabilistic localization and coordinated search have also received considerable attention in the literature. Optimal search-allocation models [10,11] relate detection performance to the spatial distribution of target-location probabilities and the allocation of available search effort. Path-constrained search formulations [12,13] additionally account for restrictions imposed by searcher trajectories and movement decisions. In a related direction, Alam and Bobadilla [16] investigated multi-robot coverage and persistent monitoring in sensing-constrained environments, demonstrating the relevance of coordinated and periodically repeated search trajectories. These contributions provide a broader methodological basis for probabilistic localization and cooperative exploration. In the present study, localization is performed over independent regions through coordinated search operations that preserve complete coverage and provide the information required to initialize the subsequent tracking stage.

The relationship between search effort, detection reward, and optimization criteria represents a further important consideration. Hohzaki and Iida [17] studied a path-constrained search problem in which route selection and observation decisions are optimized through a reward criterion incorporating detection benefits and search costs. Such formulations illustrate the importance of balancing resource expenditure against the expected benefit of successful detection. A specific discounted effort-reward formulation was subsequently introduced in [18], providing the methodological basis for assigning different weights to successive search stages. Under this formulation, discounting modifies the relative contributions of search effort and detection reward without necessarily reducing the physical search time or changing the underlying trajectories. Accordingly, the present study incorporates discounted effort-reward optimization while explicitly retaining the expected localization time and complete-coverage requirements as separate components of the optimization framework.

The closest predecessor to the present study is the two-stage stochastic transport model developed in [19]. That model combines probabilistic localization over independent regions with coordinated tracking of a multidimensional random-walk target along intersecting trajectories. The first stage concerns the physical localization process, whereas the second stage considers tracking distances, first-encounter times, and sufficient conditions for finiteness. The present study extends this framework in several respects. First, a discounted effort-reward criterion is incorporated into the localization stage while maintaining the actual localization time as a distinct performance measure.

Second, the information transferred between the stages includes not only the estimated target position but also the Stage-1 shell uncertainty, which determines the subsequent tracking discount. Third, coverage-preserving discounted tracking distances are constructed to prevent excessive contraction of the tracking trajectories. Fourth, existence and optimality properties are established, together with limiting consistency as the discount parameters approach unity. Finally, numerical comparisons with [19] are used to evaluate the additional effects of discounting and uncertainty-adaptive coupling.

A related investigation [20] considered discounted effort-reward optimization for tracking a multidimensional random-walk particle. However, that formulation addressed the tracking problem independently of a preceding probabilistic localization stage. Consequently, its tracking discount was not determined by information describing the uncertainty of the target's estimated initial location. The present study addresses this limitation by coupling discounted regional localization with an uncertainty-adaptive tracking mechanism. Specifically, the Stage-1 shell-probability distribution is summarized through a normalized uncertainty measure, which is subsequently used to determine the Stage-2 discount. The resulting tracking policy is constructed to preserve recurrent coverage while regulating the discounted contribution of successive tracking movements. This establishes an explicit connection between localization uncertainty and the subsequent tracking strategy within a unified two-stage stochastic transport framework.

The contribution of the present study therefore lies in extending established stochastic search and tracking principles rather than introducing a fundamentally different search geometry or encounter mechanism. The underlying two-stage structure and finite-block encounter framework are inherited from [19], while the discounted tracking formulation in [20] and the effort-reward principle in [18] provide the relevant methodological foundations. Building on these results, the present work introduces discounted Stage-1 optimization, an uncertainty-dependent Stage-2 discount, and coverage-preserving tracking distances within a coupled formulation. The theoretical analysis establishes the existence and optimality properties of the proposed discounted policies, preserves the required coverage conditions, and derives finiteness of the expected first-encounter time under an additional uniform positive finite-block encounter-probability assumption. The undiscounted formulation is recovered in the limiting case as the discount parameters approach unity.

Accordingly, the unresolved problem is not the construction of another two-stage search-and-tracking model, since this structure has already been established in [19]. The main issue is how to introduce discounting without confusing reduced weighted effort with reduced physical search time and without weakening the complete and recurrent coverage required for finite encounter times. The present study is designed to address precisely this difficulty.

The present paper addresses these issues through an adaptive discounted extension of the two-stage stochastic transport model. In Stage 1, a discounted effort-reward criterion is incorporated into the expanding regional search process, while the physical localization time, detection probability, and complete-coverage constraints remain explicit. This separation is important: Reducing a discounted cost is not automatically equivalent to reducing the actual detection time. The resulting constrained problem determines the search distances and the localization discount without allowing a trivial zero-effort solution.

The output of Stage 1 includes both an estimated initial position and the Stage-1 shell uncertainty, which measures the dispersion of the probability mass over the admissible search shells. In Stage 2, this shell uncertainty is used to adapt the tracking discount. Thus, the two stages are

connected not only through a geometric reference point but also through the information produced by the localization process. When the Stage-1 shell uncertainty is small, searchers can concentrate their tracking effort more strongly. When the shell uncertainty is large, the policy retains broader recurrent coverage.

Here, “adaptive” refers to an across-stage adjustment based on the localization information obtained before tracking begins, rather than to real-time state-feedback control. The tracking rule is therefore fixed during the tracking stage. A state-feedback extension could update the tracking policy using current observations and may improve responsiveness to changing target conditions. Related adaptive tracking-control studies have also addressed practical constraints such as state limitations and input saturation [21], which suggests a useful direction for extending the present tracking framework toward constrained real-time control.

To prevent the discounted paths from collapsing, the proposed weights are bounded away from zero. Coverage-preserving search and tracking distances are then constructed so that the admissible regions and trajectories continue to be visited as required. Within this setting, the paper studies the existence of optimal discounted policies and the preservation of complete and recurrent coverage and establishes finiteness of the expected first-encounter time conditional on a uniform positive encounter probability over finite time blocks. It also establishes consistency with the original formulation by recovering the undiscounted two-stage model as the discount parameters approach 1.

The main contributions of this paper are as follows:

- 1) A discounted Stage 1 optimization model is introduced that explicitly separates physical localization time from discounted movement effort and detection reward.
- 2) The Stage-1 shell uncertainty is coupled to the Stage-2 discount through an uncertainty-adaptive, coverage-preserving tracking rule.
- 3) The theoretical analysis establishes existence and optimality properties, preservation of complete and recurrent coverage, conditional finiteness of the first-encounter and total search times under Condition A4, and recovery of the undiscounted model as the discount parameters approach 1. Coverage-preserving discounted distances are constructed to prevent contraction of the regional search paths and tracking trajectories.
- 4) Numerical experiments and sensitivity analyses examine the effects of discounting, uncertainty, dimensionality, coverage period, and tracking geometry, including Monte Carlo evaluation of first-encounter behavior.

The proposed approach is distinct from stochastic resetting. Neither the target nor the searchers are instantaneously returned to a prescribed state, and the process is not restarted at an independent random time. Instead, discounting controls how effort is distributed over successive movements along structured search and tracking paths. This preserves the spatial interpretation of the two-stage model while providing an adaptive mechanism for balancing effort, detection, and coverage.

The remainder of the paper is organized as follows: Section 2 introduces the probabilistic setting, notation, and the baseline two-stage structure. Section 3 develops the discounted localization model and presents its existence and optimality properties. Section 4 defines the uncertainty measure produced by Stage 1 and the adaptive coupling rule for Stage 2. Section 5 formulates the coverage-preserving discounted tracking model. Section 6 establishes the main coverage, convergence, optimality, and finiteness results. Section 7 presents the numerical method and simulation study. Section 8 concludes the paper.

2. Problem formulation and two-stage search strategy

We consider a stochastic search problem in which the initial position of a moving target is uncertain, and its subsequent motion follows a multidimensional random walk. The search is carried out in two connected stages. In Stage 1, coordinated searchers explore independent regions to estimate the initial location of the target. In Stage 2, pairs of searchers patrol intersecting trajectories and attempt to encounter the target while it continues its random motion.

The two stages are linked through the information produced by the localization process. This information determines the adaptive tracking discount used in Stage 2. Figure 1 summarizes the complete strategy followed by the searchers, from regional localization to the first encounter with the moving target. Figure 1 is included as a schematic summary of the two-stage coupling mechanism and the principal notation, rather than as a workflow diagram only.




Figure 1. Two-stage search strategy for a multidimensional random-walk target. The figure illustrates the complete coupling mechanism between Stage 1 and Stage 2 and summarizes the principal notations used in both stages. Stage 1 uses expanding-square trajectories to localize the target, while the resulting estimate $\hat{O}$ and Stage-1 shell uncertainty U_1 determine the adaptive Stage-2 discount $\eta = \Psi(U_1)$. In Stage 2, paired searchers perform bounded outward-and-return motions along intersecting trajectories until one searcher enters the detection radius ρ of the target. Mathematical symbols in the figure were typeset using Microsoft Equation.

The baseline model preserves the two-stage structure of the earlier stochastic transport formulation, where probabilistic localization over independent regions is followed by coordinated first-encounter analysis for a multidimensional random-walk target. The present formulation extends that structure by introducing localization and tracking discounts, while keeping the physical search paths and coverage requirements explicit.

2.1. Search domain and prior distribution

Let $A = \bigcup_{i=1}^N A_i$ be a bounded search domain composed of N independent and non-overlapping square regions. Region A_i has center $O_i = (h_i, k_i)$ and half-side length $b_i > 0$. Thus,

$$A_i = [h_i - b_i, h_i + b_i] \times [k_i - b_i, k_i + b_i], \quad |A_i| = 4b_i^2. \quad (2.1)$$

The prior probability that the target initially lies in A_i is denoted by

$$q_i = P(\Xi \in A_i), \quad q_i > 0, \quad \sum_{i=1}^n q_i = 1, \quad (2.2)$$

where $\Xi = (X_0, Y_0)$ denotes the unknown initial position of the target during the localization stage. Some parts of a region may be inaccessible or excluded from the search. Let $D_i = \bigcup_{m=1}^{M_i} D_{im}$ be the deleted subset of A_i , and define the admissible part of the region by

$$A_i^{ad} = \frac{A_i}{D_i}. \quad (2.3)$$

Let $g_i(x, y)$ be the original bivariate density of the target location in region A_i . The corresponding truncated density is

$$f_i(x, y) = \frac{g_i(x, y) \mathbf{1}_{A_i^{ad}}(x, y)}{Z_i}, \quad Z_i = \int_{A_i^{ad}} g_i(x, y) dx dy, \quad (2.4)$$

where $\mathbf{1}_{A_i^{ad}}(x, y)$ denotes the indicator function of the admissible search region A_i^{ad} , defined by

$$\mathbf{1}_{A_i^{ad}}(x, y) = \begin{cases} 1, & (x, y) \in A_i^{ad} \\ 0, & (x, y) \notin A_i^{ad} \end{cases}.$$

Thus, the original density $g_i(x, y)$ is retained only on the admissible part of region A_i and is set equal to zero on the deleted or inaccessible subsets. The normalizing constant Z_i ensures that $\int_{A_i^{ad}} f_i(x, y) dx dy = 1$.

This representation allows prior uncertainty and inaccessible subsets to be handled within the same probabilistic model. The use of independent regions, truncated bivariate distributions, and structured expanding paths follows the geometric setting of the original two-stage formulation.

2.2. Stage 1: Probabilistic localization by expanding regional search

As illustrated in the left panel of Figure 1, one searcher is assigned to each region A_i . All regional searchers start simultaneously from their corresponding centers O_i . The searcher in region A_i follows a sequence of nested square contours, $C_{i1} \subset C_{i2} \subset \dots \subset C_{iK_i} = A_i$, whose half-side lengths satisfy

$$0 = r_{i0} < r_{i1} < \dots < r_{iK_i} = b_i. \quad (2.5)$$

The j -th square is defined by

$$C_{ij} = [h_i - r_{ij}, h_i + r_{ij}] \times [k_i - r_{ij}, k_i + r_{ij}], \quad j = 1, 2, \dots, K_i. \quad (2.6)$$

The path followed by the searcher during the j -th expansion is denoted by Δ_{ij} . Starting from O_i , the searcher moves outward, scans the boundary of C_{ij} , and then proceeds to the next expansion if the target has not yet been detected. The arrows shown in Figure 1 indicate the ordered movement along the expanding contours. The part of the region searched for the first time during cycle j is

$$B_{ij} = \left(\frac{C_{ij}}{C_{i(j-1)}} \right) \cap A_i^{ad}, \quad C_{i0} = \varnothing. \quad (2.7)$$

Accordingly, the conditional probability that the target is detected during cycle j , given that it lies in A_i , is

$$p_{ij} = \int_{B_{ij}} f_i(x, y) dx dy. \quad (2.8)$$

Complete regional coverage requires

$$\bigcup_{j=1}^{K_i} B_{ij} = A_i^{ad}, \quad \sum_{j=1}^{K_i} p_{ij} = 1. \quad (2.9)$$

The expanding-square strategy is therefore not merely a sequence of geometric loops. Each additional contour captures a new portion of the probability mass while preserving systematic coverage of the admissible region. The original formulation similarly uses expanding square trajectories initiated from each regional center and executes the regional searches in parallel.

Let $L_{ij} = \text{length}(\Delta_{ij})$ be the cumulative path length completed when cycle j ends. If the searcher in A_i moves at speed $v_i > 0$, the corresponding elapsed time is

$$t_{ij} = \frac{L_{ij}}{v_i}. \quad (2.10)$$

The expected localization time in region A_i is

$$E[T_{li}] = \sum_{j=1}^{K_i} t_{ij} p_{ij}, \quad (2.11)$$

and the total expected localization time is

$$\Omega_1(\mathbf{r}) = \sum_{i=1}^N \left(q_i \sum_{j=1}^{K_i} t_{ij} p_{ij} \right), \quad (2.12)$$

where $\mathbf{r} = \{r_{ij} : i = 1, 2, \dots, N; j = 1, 2, \dots, K_i\}$. The admissible set of regional expansion policies is

$$\Pi_1 = \{\mathbf{r} = 0 = r_{i0} < r_{i1} < \dots < r_{iK_i} = b_i, i = 1, 2, \dots, N\}. \quad (2.13)$$

The physical movement effort in Stage 1 is measured separately by

$$\tilde{\Omega}_1(\mathbf{r}) = \sum_{i=1}^N \left(q_i \sum_{j=1}^{K_i} L_{ij} p_{ij} \right). \quad (2.14)$$

This distinction is important because the actual localization time T_1 and the traveled distance do not represent the same quantity. The discount parameter γ will later be applied to the allocation of effort across expansion cycles, while the physical localization time remains explicitly represented.

At the end of Stage 1, the search procedure produces

$$I_1 = (\hat{O}, \tilde{\omega}_1, U_1, \gamma), \quad (2.15)$$

where $\hat{O}$ is the estimated target location, $\tilde{\omega}_1$ is the probability distribution over the Stage-1 search shells, U_1 is the corresponding normalized uncertainty index, and γ is the Stage-1 discount parameter. For each shell B_{ij} , we define $\tilde{\omega}_{ij} = P(\Xi \in B_{ij}) = q_i p_{ij}$. Thus, $\tilde{\omega}_1$ represents the probability mass distributed over the admissible search shells and is not a Bayesian posterior conditioned on a measurement observation.

2.3. Adaptive coupling between the two stages

The estimated point $\hat{O}$ provides the geometric center of the Stage-2 tracking system. Since Stage 1 is formulated in the planar search domain, its location estimate has the form $\hat{O} = (\hat{x}, \hat{y}) \in \mathbf{R}^2$. When Stage 2 is considered in $\mathbf{R}^d$, $d \geq 2$, this estimate is embedded canonically as $\hat{O}^{(d)} = (\hat{x}, \hat{y}, 0, \dots, 0) \in \mathbf{R}^d$. The point $\hat{O}^{(d)}$ is then used as the common intersection point of the Stage-2 tracking trajectories. This embedding changes neither the Stage-1 localization model nor the quantities U_1 and γ ; it only provides the geometric reference required to initialize the multidimensional tracking stage. However, the connection between the stages is not purely geometric. The Stage-1 shell uncertainty U_1 is also transferred to Stage 2 and is used to determine the tracking discount. Since $\eta = \Psi(U_1)$, where $\Psi : [0, 1] \rightarrow [\eta_{\min}, 1]$, $0 < \eta_{\min} < 1$. The function Ψ is selected so that the tracking strategy responds to the dispersion of the Stage-1 shell-probability distribution. A relatively large U_1 indicates that the probability mass is more dispersed across the search shells, so the tracking policy preserves broader coverage. A smaller U_1 indicates greater concentration of the

shell probability mass and permits stronger concentration of movement around the estimated location. Thus, the coupling rule is intended to satisfy

$$U_1^{(a)} > U_1^{(b)} \Rightarrow \Psi(U_1^{(a)}) > \Psi(U_1^{(b)}), \quad (2.16)$$

when larger values of η correspond to weaker discounting and, therefore, wider tracking coverage.

As shown in the central part of Figure 1, this adaptive rule passes information from the localization stage to the tracking stage. The searchers do not begin Stage 2 with a fixed discount chosen independently of Stage 1; rather, their tracking policy responds to the Stage-1 shell-probability dispersion summarized by U_1 .

2.4. Stage 2: Coordinated tracking on intersecting trajectories

Conditioned on I_1 , the embedded location $\hat{O}^{(d)}$ becomes the common intersection point of the admissible tracking trajectories. Let $L = \{L_1, L_2, \dots, L_n\}$ be a finite collection of straight trajectories embedded in $\mathbf{R}^d$, with $d \geq 2$, such that $\hat{O} \in L_\ell$, $\ell = 1, 2, \dots, n$.

Remark 2.1. For notational simplicity, we continue to write $\hat{O}$ for $\hat{O}^{(d)}$ throughout Stage 2.

Figure 1 displays three representative trajectories associated with the visible coordinate directions x_1 , x_2 , and x_d . The same construction extends to any finite number of intersecting trajectories. The target position at time t is $X(t) = (X_1(t), X_2(t), \dots, X_d(t))$, and its motion follows

$$X(t+1) = X(t) + \xi(t+1), \quad t \in N_0, \quad (2.17)$$

where $\{\xi(t)\}_{t \geq 1}$ is a sequence of independent and identically distributed random increments, and $N_0 = \{0, 1, 2, \dots\}$ denotes the set of non-negative integers.

If e_ℓ denotes a unit vector along trajectory L_ℓ , the target increments may be written as

$$\xi(t) \in \{\pm s_\ell e_\ell : \ell = 1, 2, \dots, n\}, \quad (2.18)$$

where s_ℓ is the corresponding random-walk step length. We assume $E[\xi(t)] = 0$ and $E[\|\xi(t)\|^2] < \infty$. The zero-mean assumption represents an unbiased random walk. It is used as a balance condition, but it does not by itself guarantee a finite encounter time. Finiteness also requires controlled searcher motion, recurrent trajectory coverage, finite tracking distances, and a uniformly positive probability of encounter.

For each trajectory L_ℓ , two coordinated searchers are deployed on opposite sides of $\hat{O}$. Their positions are denoted by $Z_\ell^+(t)$ and $Z_\ell^-(t)$. Here and throughout the paper, the index $\sigma \in \{+, -\}$ indicates the direction of motion along L_ℓ , where $\sigma = +$ denotes the positive direction, and $\sigma = -$ denotes the opposite direction. The paired searchers move outward from $\hat{O}$, reach prescribed turning points, and then return toward the intersection. If no encounter occurs, a new outward-and-return cycle begins.

During cycle k , the corresponding tracking distances are $\Theta_{\ell k}^+ > 0$ and $\Theta_{\ell k}^- > 0$. The collection

$$\Theta = \{\Theta_{\ell k}^+, \Theta_{\ell k}^- : \ell = 1, 2, \dots, n, k \geq 1\} \quad (2.19)$$

defines the Stage-2 tracking policy.

The arrows in the right panel of Figure 1 show how each pair repeatedly patrols its assigned trajectory. This outward-and-return pattern ensures that the searchers do not move indefinitely away from the estimated location. Instead, they repeatedly revisit the central and intermediate parts of the tracking domain.

The baseline paired deployment is consistent with the source model, where two coordinated searchers are assigned to each intersecting trajectory and move symmetrically with respect to the Stage-1 location estimate.

2.5. First-encounter time

Let $\rho > 0$ (the detection radius of the target); then, the distance between the target and the closest searcher at time $t > 0$ is

$$D(t) = \min_{\substack{1 \leq \ell \leq n \\ \sigma \in \{+, -\}}} \|X(t) - Z_{\ell}^{\sigma}(t)\|. \quad (2.20)$$

The first-encounter time is defined by

$$\tau_{\rho} = \inf\{t \geq 0 : D(t) \leq \rho\}. \quad (2.21)$$

Thus, an encounter occurs when at least one searcher enters the closed detection neighborhood of radius ρ around the target. This event is illustrated in Figure 1 by the final intersection between the random-walk path $X(t)$ and the detection neighborhood of one of the tracking searchers.

The expected Stage-2 encounter time is

$$\Omega_2(\Theta|I_1) = E[\tau_{\rho}|I_1]. \quad (2.22)$$

The same distance-based definition is used in the baseline framework, where the first encounter is the first time at which the Euclidean distance between the target and at least one searcher does not exceed the prescribed detection radius.

2.6. Admissible tracking conditions

Let Π_2 denote the class of admissible tracking policies. A policy $\Theta \in \Pi_2$ is required to satisfy the following conditions:

- 1) *Bounded tracking distances.* There exist finite constants $\bar{\Theta}_{\ell} \in \Pi_2$ such that $0 < \Theta_{\ell k}^{\sigma} \leq \bar{\Theta}_{\ell} < \infty$, $\sigma \in \{+, -\}$.
- 2) *Finite patrol-cycle duration.* There exists $T_{\max} < \infty$ such that every outward-and-return cycle is completed within $T_{\max}$.
- 3) *Recurrent coverage.* Every bounded admissible segment of each trajectory L_{ℓ} is revisited infinitely often.

4) *Uniform positive encounter probability.* There exist constants $T > 0$ and $\varepsilon > 0$ such that

$$P(\tau_\rho \leq (k+1)T | \tau_\rho > kT, \tilde{F}_{kT}) \geq \varepsilon, \quad (2.23)$$

for every $k > 0$, where $\{\tilde{F}_t\}$ is the natural filtration generated by the target and searcher motions.

Under (2.23),

$$P(\tau_\rho > kT) \leq (1-\varepsilon)^k, \quad (2.24)$$

and therefore,

$$E[\tau_\rho] \leq T \sum_{k=0}^{\infty} (1-\varepsilon)^k = \frac{T}{\varepsilon} < \infty. \quad (2.25)$$

This argument makes it clear why the recurrent outward-and-return strategy shown in Figure 1 is central to the model. The searchers repeatedly restore coverage of the admissible trajectories, while the uniform positive encounter probability prevents an indefinitely unsuccessful sequence of patrol cycles. The finiteness mechanism is therefore produced by the controlled search structure rather than by the unbiased target motion alone.

2.7. Baseline performance criterion

The physical total search time is $T_{tot} = T_1 + \tau_\rho$, with expectation $\Omega_{tot} = E[T_1] + E[\tau_\rho]$. A weighted baseline objective can be written as

$$J_0(\mathbf{r}, \Theta) = \alpha_1 \Omega_1(\mathbf{r}) + \alpha_2 \Omega_2(\Theta | I_1), \quad \alpha_1, \alpha_2 > 0. \quad (2.26)$$

This undiscounted criterion will serve as the reference case. In the following sections, the Stage-1 discount γ and the adaptive Stage-2 discount $\eta = \Psi(U_1)$ will be incorporated into the effort-reward criteria. The physical times T_1 and τ_ρ will remain explicit, while the discount factors will regulate how movement effort is distributed across the regional expansions and tracking cycles.

This separation is consistent with the main objective of the proposed model: discounting should improve resource allocation without causing the search paths to collapse or weakening complete and recurrent coverage.

3. Discounted localization model for Stage 1

The purpose of Stage 1 is to localize the target while controlling the effort spent over the successive regional expansions. The physical search structure introduced in Section 2 is retained without modification. In particular, the searcher assigned to A_i follows the nested square paths Δ_{ij} , and the final expansion reaches the boundary of the region through r_{iK_i} .

The discount parameter introduced below does not shorten these paths or remove any admissible shell from the search. It is used only to distinguish between the effort spent early in the localization process and the effort spent during later expansions.

3.1. Discounted effort and detection reward

Let $\gamma \in [\gamma_{\min}, 1]$, $0 < \gamma_{\min} < 1$ be the Stage-1 discount parameter. The weight assigned to the j -th expansion cycle is $w_j(\gamma) = \gamma^{j-1}$, $j = 1, 2, \dots, K_i$, [18]. Hence, the first cycle receives weight one, while the contribution of later cycles decreases when $\gamma < 1$. Since γ is bounded away from zero, $\gamma^{j-1} \geq \gamma_{\min}^{K_i-1} > 0$, $j = 1, 2, \dots, K_i$. Thus, every cycle remains represented in the discounted criterion.

Recall that L_{ij} is the cumulative path length completed by the end of cycle j , and p_{ij} is the conditional probability of detecting the target in the shell B_{ij} . The expected discounted movement effort is defined by

$$\tilde{\Omega}_{1\gamma}(\mathbf{r}) = \sum_{i=1}^N \left(q_i \sum_{j=1}^{K_i} \gamma^{j-1} L_{ij}(\mathbf{r}) p_{ij}(\mathbf{r}) \right). \quad (3.1)$$

The corresponding discounted detection reward is

$$\tilde{R}_{1\gamma}(\mathbf{r}) = \sum_{i=1}^N \left(q_i \sum_{j=1}^{K_i} \gamma^{j-1} p_{ij}(\mathbf{r}) \right). \quad (3.2)$$

The quantity $\tilde{R}_{1\gamma}(\mathbf{r})$ should not be interpreted as a new probability of eventual detection. It is a weighted reward that assigns greater value to detection during an earlier expansion. The physical detection probability remains determined by the complete-coverage condition $\sum_{j=1}^{K_i} p_{ij}(\mathbf{r}) = 1$, $i = 1, 2, \dots, N$. Consequently, $0 < \tilde{R}_{1\gamma}(\mathbf{r}) \leq 1$ and $\tilde{R}_{11}(\mathbf{r}) = 1$. Similarly, when $\gamma = 1$, the discounted movement effort reduces to the undiscounted effort introduced in Section 2: $\tilde{\Omega}_{1\gamma}(\mathbf{r}) = \tilde{\Omega}_1(\mathbf{r})$.

3.2. Stage-1 optimization criterion

The actual localization time and the discounted movement effort measure different aspects of the search. The first records the physical time required to localize the target, whereas the second describes how movement effort is distributed over the expansion cycles. For this reason, the actual expected localization time,

$$\Omega_1(\mathbf{r}) = \sum_{i=1}^N \left(q_i \sum_{j=1}^{K_i} t_{ij}(\mathbf{r}) p_{ij}(\mathbf{r}) \right), \quad (3.3)$$

is retained explicitly. We define the Stage-1 objective by

$$J_1(\mathbf{r}, \gamma) = \alpha_1 \Omega_1(\mathbf{r}) + \alpha_2 \tilde{\Omega}_{1\gamma}(\mathbf{r}) - \alpha_3 \tilde{R}_{1\gamma}(\mathbf{r}), \quad (3.4)$$

where $\alpha_1, \alpha_2, \alpha_3 > 0$.

The first term penalizes the expected physical localization time. The second term penalizes movement effort, with later expansions receiving smaller weights. The final term rewards early detection.

This form prevents a small discounted effort from being mistaken for a short localization time. A policy may have a small discounted cost because much of its movement occurs in later cycles, although the target is physically localized only after a long delay. The presence of Ω_1 rules out this interpretation.

3.3. Admissible discounted policies

The expansion distances must remain ordered and must preserve complete regional coverage. To obtain a closed admissible set, let $\delta_i > 0$ be the minimum permitted separation between two successive contours, with $K_i \delta_i \leq b_i$. Define

$$\Pi_1^\delta = \left\{ \mathbf{r} : \begin{array}{l} r_{i0} = 0, r_{iK_i} = b_i, \\ r_{ij} - r_{i(j-1)} \geq \delta_i, \\ i = 1, 2, \dots, N, j = 1, 2, \dots, K_i \end{array} \right\}. \quad (3.5)$$

The lower spacing condition excludes coincident contours and ensures that every cycle covers a nonempty shell. A lower bound on γ alone keeps the cycle weights positive, but it may not prevent the optimizer from selecting the smallest admissible discount. To exclude such a boundary solution when it gives insufficient value to delayed detection, we impose the operational constraint $\tilde{R}_{1\gamma}(\mathbf{r}) \geq \tilde{R}_{\min}$, $0 < \tilde{R}_{\min} < 1$. The admissible set is, therefore,

$$\bar{D}_1 = \{(\mathbf{r}, \gamma) : \mathbf{r} \in \Pi_1^\delta, \gamma \in [\gamma_{\min}, 1], \tilde{R}_{1\gamma}(\mathbf{r}) \geq \tilde{R}_{\min}\}. \quad (3.6)$$

The discounted localization problem is

$$\min_{(\mathbf{r}, \gamma) \in \bar{D}_1} J_1(\mathbf{r}, \gamma). \quad (3.7)$$

The value $\tilde{R}_{\min}$ is selected according to the minimum acceptable emphasis on timely detection. It does not replace the complete-coverage condition; it only restricts the discounted evaluation of the detection cycles.

3.4. Bounds for the discounted quantities

Let $K_{\max} = \max_{1 \leq i \leq N} K_i$, then the next result shows that discounting changes the relative weights of the cycles without eliminating their contributions.

Proposition 3.1. *For every $(\mathbf{r}, \gamma) \in \bar{D}_1$, the following inequalities hold: $\gamma_{\min}^{K_{\max}-1} \tilde{\Omega}_1(\mathbf{r}) \leq \tilde{\Omega}_{1\gamma}(\mathbf{r}) \leq \tilde{\Omega}_1(\mathbf{r})$ and $\gamma_{\min}^{K_{\max}-1} \leq \tilde{R}_{1\gamma}(\mathbf{r}) \leq 1$.*

Proof. For every i and j , we have

$$\gamma_{\min}^{K_{\max}-1} \leq \gamma^{j-1} \leq 1. \quad (3.8)$$

Since q_i , L_{ij} , and p_{ij} are nonnegative, multiplication of (3.8) by $q_i L_{ij} p_{ij}$, followed by summation over i and j , gives $\gamma_{\min}^{K_{\max}-1} \tilde{\Omega}_1(\mathbf{r}) \leq \tilde{\Omega}_{1\gamma}(\mathbf{r}) \leq \tilde{\Omega}_1(\mathbf{r})$.

Similarly, multiplication by $q_i p_{ij}$ gives

$$\gamma_{\min}^{K_{\max}-1} \sum_{i=1}^N q_i \sum_{j=1}^{K_i} p_{ij} \leq \tilde{R}_{1\gamma} \leq \sum_{i=1}^N q_i \sum_{j=1}^{K_i} p_{ij}.$$

Using $\sum_{i=1}^N q_i = 1$ and $\sum_{j=1}^{K_i} p_{ij} = 1$, we obtain $\gamma_{\min}^{K_{\max}-1} \leq \tilde{R}_{1\gamma}(\mathbf{r}) \leq 1$. $\square$

3.5. Preservation of complete regional coverage

The discount parameter appears in the objective function, not in the physical definition of the paths. This distinction ensures that the reduction in weighted effort is not achieved by shrinking the searched region.

Proposition 3.2. Every policy $(\mathbf{r}, \gamma) \in \bar{D}_1$ preserves complete coverage of A_i^{ad} for each $i = 1, 2, \dots, N$.

Proof. From the definition of Π_1^δ and (2.5), then $C_{i1} \subset C_{i2} \subset \dots \subset C_{iK_i} = A_i$. Also, from (2.7), the shells are disjoint apart from their boundaries and satisfy (2.9). The factor γ^{j-1} does not enter the definition of C_{ij} , B_{ij} , or Δ_{ij} . Hence, it cannot remove any admissible part of the region from the search. $\square$

3.6. Existence of an optimal Stage-1 policy

We impose the following regularity assumptions:

Assumption 1. For each i , the density $f_i(x, y)$ is continuous on A_i^{ad} .

Assumption 2. The path-length functions $L_{ij}(\mathbf{r})$ are continuous on Π_1^δ , and $v_i > 0$.

Assumption 3. The admissible set $\bar{D}_1$ is nonempty.

Under Assumption 1, the shell probabilities $p_{ij}(\mathbf{r}) = \int_{B_{ij}(\mathbf{r})} f_i(x, y) dx dy$ vary continuously with the expansion distances. Indeed, varying $\mathbf{r}$ changes only the boundaries of the square shells, and these boundaries have zero two-dimensional Lebesgue measure.

Theorem 3.1. Under Assumptions 1, 2, and 3, problem (3.7) admits at least one optimal solution $(\mathbf{r}^*, \gamma^*) \in \bar{D}_1$.

Proof. The endpoint and spacing constraints in (3.5) define a closed and bounded subset of a finite-dimensional Euclidean space. Hence, Π_1^δ is compact. The interval $[\gamma_{\min}, 1]$ is compact, and the continuity of $\tilde{R}_{1\gamma}(\mathbf{r})$ implies that $\{(\mathbf{r}, \gamma) : \tilde{R}_{1\gamma}(\mathbf{r}) \geq \tilde{R}_{\min}\}$ is closed. Therefore, $\bar{D}_1$ is a closed subset of the compact set $\Pi_1^\delta \times [\gamma_{\min}, 1]$, and is itself compact.

The functions p_{ij} , L_{ij} , and $t_{ij} = \frac{L_{ij}}{v_i}$ are continuous on the admissible set. It follows that Ω_1 , $\tilde{\Omega}_{1\gamma}$, $\tilde{R}_{1\gamma}$, and, hence, J_1 are continuous on $\bar{D}_1$. The conclusion now follows from the Weierstrass existence theorem: a continuous real-valued function defined on a nonempty compact set attains its minimum [22]. $\square$

Here, “optimal” means that $(\mathbf{r}^*, \gamma^*)$ globally minimizes the Stage-1 objective $J_1(\mathbf{r}, \gamma)$ over the

admissible set $\bar{D}_1$, subject to the coverage, spacing, discount, and reward constraints.

A recent optimization in [22] develops the existence theory of minimizers around precisely this form of the Weierstrass theorem, including its role in finite- and infinite-dimensional optimization.

3.7. Necessary optimality conditions

The existence result does not require differentiability. To characterize a local optimum, suppose in addition that J_1 and $\tilde{R}_{1\gamma}$ are continuously differentiable with respect to their decision variables. Write the reward constraint as,

$$\zeta_{\tilde{R}}(\mathbf{r}, \gamma) = \tilde{R}_{\min} - \tilde{R}_{1\gamma}(\mathbf{r}) \leq 0,$$

and the spacing constraints as,

$$\zeta_{ij}(\mathbf{r}) = \delta_i - r_{ij} + r_{i(j-1)} \leq 0.$$

The bounds on γ can be written as,

$$\zeta_{\gamma}^{-}(\gamma) = \gamma_{\min} - \gamma \leq 0, \quad \zeta_{\gamma}^{+}(\gamma) = \gamma - 1 \leq 0.$$

Let λ_{ij} , $\mu \geq 0$, $\nu_{-} \geq 0$, and $\nu_{+} \geq 0$ be the associated multipliers. The Lagrangian is

$$\tilde{L}(\mathbf{r}, \gamma) = J_1(\mathbf{r}, \gamma) + \sum_{i=1}^N \sum_{j=1}^{K_i} \lambda_{ij} g_{ij}(\mathbf{r}) + \mu \zeta_{\tilde{R}}(\mathbf{r}, \gamma) + \nu_{-} \zeta_{\gamma}^{-}(\gamma) + \nu_{+} \zeta_{\gamma}^{+}(\gamma).$$

Theorem 3.2. *Let $(\mathbf{r}^*, \gamma^*)$ be a local solution of (3.7). If a standard constraint qualification holds at $(\mathbf{r}^*, \gamma^*)$, then there exist nonnegative multipliers such that $\nabla_{\mathbf{r}} \tilde{L}(\mathbf{r}^*, \gamma^*) = 0$ and $\frac{\partial \tilde{L}(\mathbf{r}^*, \gamma^*)}{\partial \gamma} = 0$, together with primal feasibility, dual feasibility, and complementary slackness: $\lambda_{ij} \zeta_{ij}(\mathbf{r}^*) = 0$, $\mu \zeta_{\tilde{R}}(\mathbf{r}^*, \gamma^*) = 0$, and $\nu_{-} \zeta_{\gamma}^{-}(\gamma^*) = 0$, $\nu_{+} \zeta_{\gamma}^{+}(\gamma^*) = 0$.*

Proof. The result follows from the Karush–Kuhn–Tucker (KKT) necessary conditions for smooth nonlinear programs under a valid constraint qualification; see [23]. $\square$

Theorem 3.1 establishes the existence of a global minimizer by compactness and continuity. Theorem 3.2 does not serve as a verification theorem and does not by itself establish global optimality; it provides only necessary first-order conditions for a local solution under the stated differentiability and constraint-qualification assumptions. Thus, existence and first-order optimality characterization are kept logically distinct.

When γ^* lies in the interior of $[\gamma_{\min}, 1]$ and the reward constraint is inactive, $\frac{\partial \tilde{L}(\mathbf{r}^*, \gamma^*)}{\partial \gamma} = 0$ reduces to $\frac{\partial J_1(\mathbf{r}^*, \gamma^*)}{\partial \gamma} = 0$. Since Ω_1 does not depend explicitly on γ ,

$$\frac{\partial J_1}{\partial \gamma} = \alpha_2 \sum_{i=1}^N q_i \sum_{j=2}^{K_i} (j-1) \gamma^{j-2} L_{ij} p_{ij} - \alpha_3 \sum_{i=1}^N q_i \sum_{j=2}^{K_i} (j-1) \gamma^{j-2} p_{ij}.$$

Thus, an interior optimal discount satisfies

$$\sum_{i=1}^N q_i \sum_{j=2}^{K_i} (j-1) (\gamma^*)^{j-2} p_{ij} (\alpha_2 L_{ij} - \alpha_3) = 0. \quad (3.9)$$

Equation (3.9) balances the marginal discounted movement cost against the marginal reward

assigned to earlier detection.

3.8. Undiscounted limit

The proposed model must remain consistent with the original Stage-1 formulation.

Theorem 3.3. For every fixed $\mathbf{r} \in \Pi_1^\delta$, $\lim_{\gamma \rightarrow 1} \tilde{\Omega}_{1\gamma}(\mathbf{r}) = \tilde{\Omega}_1(\mathbf{r})$, $\lim_{\gamma \rightarrow 1} \tilde{R}_{1\gamma}(\mathbf{r}) = 1$, and $\lim_{\gamma \rightarrow 1} J_1(\mathbf{r}, \gamma) = \alpha_1 \Omega_1(\mathbf{r}) + \alpha_2 \tilde{\Omega}_{1\gamma}(\mathbf{r}) - \alpha_3$.

Proof. The sums in (3.1) and (3.2) contain finitely many terms. Since $\lim_{\gamma \rightarrow 1} \gamma^{j-1} = 1$, the limit may be taken term by term. Therefore, $\lim_{\gamma \rightarrow 1} \tilde{\Omega}_{1\gamma}(\mathbf{r}) = \tilde{\Omega}_1(\mathbf{r})$ and $\lim_{\gamma \rightarrow 1} \tilde{R}_{1\gamma}(\mathbf{r}) = 1$ follow immediately. Substitution into (3.4) gives $\lim_{\gamma \rightarrow 1} J_1(\mathbf{r}, \gamma) = \alpha_1 \Omega_1(\mathbf{r}) + \alpha_2 \tilde{\Omega}_{1\gamma}(\mathbf{r}) - \alpha_3$. $\square$

The term $-\alpha_3$ is constant with respect to $\mathbf{r}$. Therefore, it does not affect the optimal expansion distances, and the undiscounted Stage-1 problem is recovered as $\gamma \rightarrow 1$.

3.9. Output of the discounted localization stage

Let $(\mathbf{r}^*, \gamma^*)$ be an optimal solution of (3.7). The localization search is then carried out using the expansion distances $\mathbf{r}^*$. At the optimal policy $(\mathbf{r}^*, \gamma^*)$, the shell probabilities are $\tilde{\omega}_{ij}^* = q_i p_{ij}(\mathbf{r}^*)$. The information passed from Stage 1 is, therefore, $I_1 = (\hat{O}, \tilde{\omega}_1^*, U_1, \gamma^*)$.

4. Stage-1 shell uncertainty and adaptive coupling

The Stage-1 shell uncertainty U_1 , introduced in Section 2 and obtained from the optimized Stage-1 policy, is used here to determine the Stage-2 discount through $\eta = \Psi(U_1)$. This section specifies a normalized form of U_1 , defines the admissible properties of the coupling function Ψ , and establishes the stability and boundedness properties needed for the Stage-2 tracking model. The construction is designed to preserve bounded tracking distances, recurrent coverage, and finite patrol cycles, while the uniform positive encounter probability is imposed separately through Condition A4. The explicit discounted tracking distances are introduced in Section 5. A discounted tracking mechanism was previously applied to multidimensional random-walk targets in [20]. Here, the tracking discount is instead determined by the Stage-1 shell uncertainty, which provides a direct information-based link between the two stages.

4.1. Stage-1 shell uncertainty

Let $M = \sum_{i=1}^N K_i$ be the total number of admissible Stage-1 shells. For each shell B_{ij} , let $q_i = P(\Xi \in A_i)$ and $p_{ij} = P(\Xi \in B_{ij} | \Xi \in A_i)$. The corresponding unconditional shell probability is $\tilde{\omega}_{ij} = P(\Xi \in B_{ij}) = q_i p_{ij}$. Since $\sum_{i=1}^N q_i = 1$ and $\sum_{j=1}^{K_i} p_{ij} = 1$, the shell probabilities satisfy $\sum_{i=1}^N \sum_{j=1}^{K_i} \tilde{\omega}_{ij} = 1$. We define the Stage-1 shell uncertainty by

$$U_1 = -\frac{1}{\log M} \sum_{i=1}^N \sum_{j=1}^{K_i} \tilde{\omega}_{ij} \log \tilde{\omega}_{ij}, \quad (4.1)$$

with $0 \log 0 = 0$. The quantity U_1 measures the dispersion of the target-location probability over the Stage-1 search shells. It is not a posterior entropy obtained from a Bayesian observation update.

Proposition 4.1. *The uncertainty measure in (4.1) satisfies $0 \leq U_1 \leq 1$. Moreover, $U_1 = 0$ if and only if one shell probability equals one and all the others equal zero. The equality $U_1 = 1$ holds iff $\tilde{\omega}_{ij} = \frac{1}{M}$ for every admissible pair (i, j) .*

Proof. For a probability vector with M components, the entropy satisfies

$$0 \leq -\sum_{i=1}^N \sum_{j=1}^{K_i} \tilde{\omega}_{ij} \log \tilde{\omega}_{ij} \leq \log M.$$

The lower bound is attained when the shell-probability distribution is concentrated in a single shell, whereas the upper bound is attained when the probability mass is uniformly distributed over all M shells. Dividing by $\log M > 0$ proves the result. $\square$

The use of the complete shell-probability distribution is important because two Stage-1 search configurations may have the same estimated point $\hat{O}$ while exhibiting different levels of probability dispersion across the search shells.

4.2. Admissible adaptive coupling

Let $\eta \in [\eta_{\min}, 1]$, $0 < \eta_{\min} < 1$ be the Stage-2 discount parameter. It is determined from U_1 through $\eta = \Psi(U_1)$, where

$$\Psi : [0, 1] \rightarrow [\eta_{\min}, 1]. \quad (4.2)$$

The coupling function is assumed to satisfy

$$\Psi(0) = \eta_{\min}, \quad \Psi(1) = 1, \quad (4.3)$$

and

$$u_a \leq u_b \Rightarrow \Psi(u_a) \leq \Psi(u_b). \quad (4.4)$$

Thus, greater Stage-1 shell-probability dispersion, corresponding to a larger U_1 , produces a larger value of η . Under the convention adopted in Section 2, a larger value of η corresponds to weaker discounting and, therefore, broader tracking coverage. Conversely, a smaller value of U_1 permits stronger concentration around $\hat{O}$.

To prevent small perturbations in the Stage-1 shell-probability distribution from producing abrupt changes in the tracking policy, we additionally assume that Ψ is Lipschitz continuous. That is, there exists a finite constant $\tilde{L}_\Psi > 0$ such that

$$|\Psi(u_a) - \Psi(u_b)| \leq \tilde{L}_\Psi |u_a - u_b|, \quad u_a, u_b \in [0, 1]. \quad (4.5)$$

This is the standard Lipschitz condition described in [24]. It is stronger than ordinary continuity and provides a direct quantitative bound on the sensitivity of the Stage-2 discount to changes in the Stage-1 shell uncertainty. In particular, for two localization outcomes $U_1^{(a)}$ and $U_1^{(b)}$,

$$|\eta^{(a)} - \eta^{(b)}| \leq \tilde{L}_\Psi |U_1^{(a)} - U_1^{(b)}|, \text{ where } \eta^{(a)} = \Psi(U_1^{(a)}) \text{ and } \eta^{(b)} = \Psi(U_1^{(b)}).$$

4.3. Linear coupling rule

A simple rule satisfying (4.2)–(4.5) is $\Psi(u) = \eta_{\min} + (1 - \eta_{\min})u$, $0 \leq u \leq 1$. Accordingly,

$$\eta = \eta_{\min} + (1 - \eta_{\min})U_1. \quad (4.6)$$

This specification requires no additional tuning parameter beyond $\eta_{\min}$. It also preserves the limiting cases $U_1 = 0$, $\eta = \eta_{\min}$ and $U_1 = 1$, $\eta = 1$.

Proposition 4.2. *For every $U_1 \in [0, 1]$, the coupling rule in (4.6) satisfies $\eta_{\min} \leq \eta \leq 1$. Furthermore, η is strictly increasing in U_1 , and Ψ is Lipschitz continuous with constant $\tilde{L}_\Psi = 1 - \eta_{\min}$.*

Proof. Since $0 \leq U_1 \leq 1$, $0 \leq (1 - \eta_{\min})U_1 \leq 1 - \eta_{\min}$. Adding $\eta_{\min}$ gives $\eta_{\min} \leq \eta \leq 1$. Moreover,

$$\frac{d\eta}{dU_1} = 1 - \eta_{\min} > 0, \quad u_a, u_b \in [0, 1], \quad \text{so the rule is strictly increasing. Finally, for}$$

$$|\Psi(u_a) - \Psi(u_b)| \leq |(1 - \eta_{\min})(u_a - u_b)| = (1 - \eta_{\min})|u_a - u_b|, \text{ which completes the proof. } \square$$

The linear rule is adopted in the remainder of this paper because it gives a direct interpretation of the relation between uncertainty and tracking coverage. A nonlinear rule may be used without changing the analysis, provided that it satisfies (4.2)–(4.5).

4.4. Compatibility with the Stage-2 tracking conditions

Let $\Theta(\eta) = \{\Theta_{\ell k}^\sigma(\eta) : \ell = 1, 2, \dots, n; k \geq 1; \sigma \in \{+, -\}\}$ denote a family of Stage-2 tracking policies indexed by η . The precise formula for $\Theta_{\ell k}^\sigma(\eta)$ will be introduced in Section 5. The baseline Stage-2 framework requires bounded tracking distances, finite patrol-cycle duration, recurrent coverage, and a uniform positive probability of encounter. These conditions are essential because an unbiased multidimensional random walk does not, by itself, guarantee a finite expected first-encounter time [19].

We impose the following uniform requirements on the adaptive family $\Theta(\eta)$.

Condition A1. There exist constants $0 < \underline{\Theta}_\ell^\sigma \leq \overline{\Theta}_\ell^\sigma < \infty$ such that $\underline{\Theta}_\ell^\sigma \leq \Theta_{\ell k}^\sigma(\eta) \leq \overline{\Theta}_\ell^\sigma$ for every

$$\eta \in [\eta_{\min}, 1], \quad k \geq 1, \quad \ell = 1, 2, \dots, n, \quad \text{and} \quad \sigma \in \{+, -\}.$$

Condition A2. There exists a finite constant $T_{\max}$ such that every outward-and-return cycle generated by $\Theta(\eta)$ is completed within $T_{\max}$, uniformly in η .

Condition A3. Every bounded admissible segment of each trajectory L_ℓ is revisited infinitely often for every $\eta \in [\eta_{\min}, 1]$.

Condition A4. For every $\eta \in [\eta_{\min}, 1]$, there exist $T(\eta) > 0$ and $\varepsilon(\eta) > 0$ such that

$$P(\tau_\rho \leq (k+1)T(\eta) | \tau_\rho > kT(\eta), \mathfrak{T}_{kT(\eta)}) \geq \varepsilon(\eta), \text{ for every } k \geq 0.$$

Condition A4 is an additional probabilistic assumption on the joint target-searcher dynamics. It is not derived solely from bounded tracking distances, finite patrol-cycle duration, or recurrent geometric coverage. In particular, Proposition 5.1 and Theorem 5.1 establish the geometric properties of the adaptive policy, whereas A4 separately requires a uniformly positive probability of encounter over finite time blocks. Therefore, all finiteness results below are conditional on A4, or on any set of

model-specific assumptions that implies A4.

Conditions A1–A4 are the uncertainty-dependent counterparts of the admissibility conditions defining Π_2 in Section 2. They require the whole adaptive family, rather than only one fixed tracking policy, to remain inside the bounded controlled tracking framework of [19].

4.5. Uniform encounter probability under adaptive coupling

The encounter probability in Condition A4 must remain uniformly positive over the full admissible range of η . Positivity for each isolated value of η is not sufficient, because a sequence of admissible discounts could otherwise produce encounter probabilities converging to zero.

Proposition 4.3. *Suppose that $\varepsilon(\eta)$ in Condition A4 is continuous and strictly positive on $[\eta_{\min}, 1]$. Then there exists a constant $\varepsilon^* > 0$ such that $\varepsilon(\eta) \geq \varepsilon^*$ for every $\eta \in [\eta_{\min}, 1]$.*

Proof. The interval $[\eta_{\min}, 1]$ is compact. Since $\varepsilon(\eta)$ is continuous, the Weierstrass theorem [22] implies that it attains its minimum on this interval. Hence, $\varepsilon^* = \min_{\eta \in [\eta_{\min}, 1]} \varepsilon(\eta)$. Because $\varepsilon(\eta) > 0$ throughout the interval, its minimum is also strictly positive. Therefore, $\varepsilon(\eta) \geq \varepsilon^*$ for every $\eta \in [\eta_{\min}, 1]$ holds. $\square$

Similarly, if $T(\eta)$ is continuous on $[\eta_{\min}, 1]$, then $T^* = \max_{\eta \in [\eta_{\min}, 1]} T(\eta) < \infty$. Consequently, the adaptive tracking family satisfies the uniform estimate $P(\tau_\rho \leq (k+1)T^* | \tau_\rho > kT^*, \mathfrak{I}_{kT^*}) \geq \varepsilon^*$.

4.6. Transfer of the Stage-2 finiteness mechanism

The following result is conditional on Condition A4. It shows that, once the required positive finite-block encounter probability is available, the uncertainty-dependent coupling preserves the finiteness mechanism used in Section 2.

Theorem 4.1. *Suppose that Ψ satisfies (4.2)–(4.5), and that the family $\Theta(\eta)$ satisfies Conditions A1–A4. If $\varepsilon^* > 0$ and $T^* < \infty$ are defined by $\varepsilon^* = \min_{\eta \in [\eta_{\min}, 1]} \varepsilon(\eta)$ and $T^* = \max_{\eta \in [\eta_{\min}, 1]} T(\eta) < \infty$, respectively, then, for every $U_1 \in [0, 1]$,*

$$P(\tau_\rho > kT^* | U_1) \leq (1 - \varepsilon^*)^k, \quad k \geq 0 \quad (4.7)$$

and

$$E[\tau_\rho | U_1] \leq \frac{T^*}{\varepsilon^*} < \infty. \quad (4.8)$$

Proof. For every $U_1 \in [0, 1]$, the coupling rule gives $\eta = \Psi(U_1) \in [\eta_{\min}, 1]$. Hence, the corresponding policy $\Theta(\eta)$ satisfies Conditions A1–A4. Here, Conditions A1–A3 correspond respectively to bounded tracking distances, finite patrol-cycle duration, and recurrent trajectory coverage, while Condition A4 supplies the positive finite-block encounter probability. These are exactly the four ingredients required in the block argument below; no additional recurrence or encounter assumption is invoked. From the uniform encounter bound $P(\tau_\rho \leq (k+1)T^* | \tau_\rho > kT^*, \mathfrak{I}_{kT^*}) \geq \varepsilon^*$ and $P(\tau_\rho > (k+1)T^* | \tau_\rho > kT^*, U_1) \leq 1 - \varepsilon^*$. Iteration gives (4.7). Using the tail representation of a nonnegative stopping time,

$$E[\tau_\rho|U_1] \leq T^* \sum_{k=0}^{\infty} P(\tau_\rho > kT^*|U_1) \leq T^* \sum_{k=0}^{\infty} (1-\varepsilon^*)^k = \frac{T^*}{\varepsilon^*}.$$

Therefore, the conditional expected first-encounter time is finite for every admissible value of the Stage-1 shell uncertainty. $\square$

Theorem 4.1 does not claim that discounting alone guarantees a finite encounter time. It shows that uncertainty-dependent discounting preserves the finiteness result when the resulting tracking distances remain bounded, the patrol cycles remain finite, the trajectories are recurrently covered, and the encounter probability is uniformly positive. This is precisely the controlled mechanism underlying the Stage-2 model in [19]. In the original formulation, the geometric tail estimate also yields the bound $E[\tau_\rho] \leq \frac{T}{\varepsilon}$.

4.7. Information passed to the tracking model

After Stage 1, the shell-probability vector $\tilde{\omega}_1^*$ determines the uncertainty index U_1 , which, in turn, determines the Stage-2 discount through $\eta = \Psi(U_1)$. Hence, the information flow between the two stages is $\tilde{\omega}_1^* \rightarrow U_1 \rightarrow \eta \rightarrow \Theta(\eta)$. The embedded point $\hat{O}^{(d)}$, denoted simply by $\hat{O}$ in Stage 2, remains the common intersection point of the tracking trajectories. Section 5 then constructs the adaptive tracking distances $\Theta_{\ell k}^\sigma(\eta)$ while preserving the admissibility requirements of the tracking policy.

5. Coverage-preserving adaptive tracking distances

The discount parameter $\eta = \Psi(U_1)$ obtained in Section 4 is now incorporated into the Stage-2 tracking distances. The target process $X(t)$, the intersecting trajectories L_ℓ , $\ell = 1, 2, \dots, n$, and the first-encounter time τ_ρ remain as defined in Section 2.

For each trajectory L_ℓ , cycle $k \geq 1$, and direction $\sigma \in \{+, -\}$, let $\Theta_{\ell k}^\sigma(1)$ denote the corresponding distance under the undiscounted tracking policy. Assume that $0 < \underline{\Theta}_\ell^\sigma \leq \Theta_{\ell k}^\sigma(1) \leq \bar{\Theta}_\ell^\sigma < \infty$.

The lower bound prevents the tracking motion from collapsing to $\hat{O}$, whereas the upper bound confines each searcher to a bounded segment of L_ℓ . A direct definition of the form $\eta^{k-1}\Theta_{\ell k}^\sigma(1)$ is unsuitable when $\eta < 1$, because the resulting distances converge to zero. To reduce the tracking effort without losing recurrent coverage, we introduce periodic full-coverage cycles.

5.1. Adaptive distance rule

Fix an integer $m \geq 2$, and define

$$w_k(\eta) = \begin{cases} 1, & k = rm \text{ for some } r \geq 1 \\ \eta^{k-1}, & \text{Otherwise} \end{cases}.$$

The adaptive tracking distance is

$$\Theta_{\ell k}^\sigma(\eta) = \underline{\Theta}_\ell^\sigma + w_k(\eta)[\Theta_{\ell k}^\sigma(1) - \underline{\Theta}_\ell^\sigma]. \quad (5.1)$$

On an ordinary cycle, this becomes

$$\Theta_{\ell k}^{\sigma}(\eta) = \underline{\Theta}_{\ell}^{\sigma} + \eta^{k-1}[\Theta_{\ell k}^{\sigma}(1) - \underline{\Theta}_{\ell}^{\sigma}]. \quad (5.2)$$

On every m -th cycle, the searcher follows the undiscounted distance. In particular, when the nominal coverage cycle satisfies $\Theta_{\ell rm}^{\sigma}(1) = \bar{\Theta}_{\ell}^{\sigma}$, we obtain $\Theta_{\ell rm}^{\sigma}(\eta) = \bar{\Theta}_{\ell}^{\sigma}$. Thus, the Stage-1 shell uncertainty affects the ordinary cycles, while the periodic coverage cycles remain unchanged. This construction extends the discounted tracking idea used in [20]. Its difference is that the discount is determined by the Stage-1 shell uncertainty and is applied without allowing the tracking paths to contract permanently.

5.2. Basic properties

Proposition 5.1. *For every $\eta \in [\eta_{\min}, 1]$, $\underline{\Theta}_{\ell}^{\sigma} \leq \Theta_{\ell k}^{\sigma}(1) \leq \bar{\Theta}_{\ell}^{\sigma}$ for all $\ell = 1, 2, \dots, n$, $k \geq 1$, and $\sigma \in \{+, -\}$. Moreover, $\Theta_{\ell rm}^{\sigma}(\eta)$ is nondecreasing in η .*

Proof. Since $0 < w_k(\eta) \leq 1$ and $0 \leq \Theta_{\ell k}^{\sigma}(1) - \underline{\Theta}_{\ell}^{\sigma} \leq \bar{\Theta}_{\ell}^{\sigma} - \underline{\Theta}_{\ell}^{\sigma}$, then (5.1) immediately gives $\underline{\Theta}_{\ell}^{\sigma} \leq \Theta_{\ell k}^{\sigma}(1) \leq \bar{\Theta}_{\ell}^{\sigma}$. For $k \neq rm$, the weight $w_k(\eta) = \eta^{k-1}$ is non-decreasing in η . And, for $k = rm$, it is equal to one and is independent of η . Therefore, the adaptive distance is nondecreasing in η . $\square$

The result has a direct interpretation. A large value of U_1 gives a large value of η , so the ordinary patrols extend farther from $\hat{O}$. When the shell-probability distribution is more concentrated, corresponding to a smaller U_1 , the ordinary distances are reduced, but they remain above the positive baseline.

5.3. Dependence on the Stage-1 shell uncertainty

Substituting $\eta = \Psi(U_1)$ into (5.1) gives

$$\Theta_{\ell k}^{\sigma}(U_1) = \underline{\Theta}_{\ell}^{\sigma} + w_k(\Psi(U_1))[\Theta_{\ell k}^{\sigma}(1) - \underline{\Theta}_{\ell}^{\sigma}]. \quad (5.3)$$

For an ordinary cycle,

$$\Theta_{\ell k}^{\sigma}(U_1) = \underline{\Theta}_{\ell}^{\sigma} + [\Psi(U_1)]^{k-1}[\Theta_{\ell k}^{\sigma}(1) - \underline{\Theta}_{\ell}^{\sigma}]. \quad (5.4)$$

Assume that Ψ satisfies the Lipschitz condition (4.5), then, for each fixed ordinary cycle,

$$\left| \Theta_{\ell k}^{\sigma}(U_1^{(a)}) - \Theta_{\ell k}^{\sigma}(U_1^{(b)}) \right| \leq (k-1) \tilde{L}_{\Psi} [\Theta_{\ell k}^{\sigma}(1) - \underline{\Theta}_{\ell}^{\sigma}] |U_1^{(a)} - U_1^{(b)}|. \quad (5.5)$$

Indeed, the mean value theorem gives

$$\left| [\Psi(U_1^{(a)})]^{k-1} - [\Psi(U_1^{(b)})]^{k-1} \right| \leq (k-1) \left| \Psi(U_1^{(a)}) - \Psi(U_1^{(b)}) \right|,$$

and (5.5) follows from (5.4). On a full-coverage cycle, the distance is independent of U_1 . Hence, small changes in the Stage-1 shell uncertainty produce controlled changes in the ordinary tracking distances.

5.4. Recurrent coverage

Theorem 5.1. Suppose that $\Theta_{\ell rm}^\sigma(1) = \bar{\Theta}_\ell^\sigma$ holds for every $r \geq 1$. Then, the policy defined by (5.1) recurrently covers every admissible segment of L_ℓ between $\hat{O} - \bar{\Theta}_\ell^- e_\ell$ and $\hat{O} - \bar{\Theta}_\ell^+ e_\ell$.

Proof. Let S be any bounded admissible segment of L_ℓ contained between the two outer turning points determined by $\underline{\Theta}_\ell$ and $\bar{\Theta}_\ell$. For every $r \geq 1$, cycle $k = rm$ is a full-coverage cycle, and by (5.1), its adaptive distance coincides with the nominal full-coverage distance. Hence, during each such cycle, the two searchers move from the intersection point to the corresponding outer turning points and return, crossing every point of S . Since the sequence of full-coverage cycles $m, 2m, 3m, \dots$ is infinite, S is traversed infinitely many times. Because S was arbitrary, every bounded admissible segment of L_ℓ is revisited infinitely often. Therefore, Condition A3 holds and recurrent coverage is established. $\square$

The alternating outward-and-return construction is inherited from the Stage-2 framework in [19], where finite tracking distances and repeated traversal of the admissible paths form the geometric basis of the first-encounter analysis.

These geometric properties do not by themselves imply Condition A4. Condition A4 depends additionally on the transition law of the target process, the detection radius ρ , and the relative target-searcher geometry. Accordingly, the construction in this section verifies the coverage and boundedness components of admissibility, while the positive finite-block encounter probability remains a separate probabilistic requirement.

It follows from Proposition 5.1 and Theorem 5.1 that the adaptive policy preserves Conditions A1 and A3. Because the searchers move with unit speed, the duration of a completed cycle satisfies $T_{\ell k}^\sigma(\eta) = 2\Theta_{\ell k}^\sigma(\eta) \leq 2\bar{\Theta}_\ell^\sigma$. Thus, the finite-cycle condition is also preserved.

5.5. Admissible adaptive policies

Define $\Pi_2^{cp} = \{\Theta(\eta) \in \Pi_2 : \Theta_{\ell k}^\sigma(\eta) \text{ satisfies } \Theta_{\ell rm}^\sigma(1) = \bar{\Theta}_\ell^\sigma \text{ and (5.1)}\}$, where cp denotes coverage preservation. For every admissible policy in $\Theta(\eta) \in \Pi_2^{cp}$, the corresponding conditional expected first-encounter time is $E_{\Theta(\eta)}[\tau_\rho | I_{1\eta}]$. The expected tracking effort is defined by

$$\tilde{\Omega}_{2\eta} = E \left[\sum_{k=1}^{K_\rho} \sum_{\ell=1}^n \sum_{\sigma \in \{+, -\}} d_{\ell k}^\sigma(\eta) | I_{1\eta} \right], \quad (5.6)$$

where K_ρ is the cycle in which the first encounter occurs, and $d_{\ell k}^\sigma(\eta)$ is the actual distance traveled during cycle k . For a completed outward-and-return cycle, $d_{\ell k}^\sigma(\eta) = 2\Theta_{\ell k}^\sigma(\eta)$. The Stage-2 problem is

$$\min_{\Theta(\eta) \in \Pi_2^{cp}} \left\{ \beta_1 E_{\Theta(\eta)}[\tau_\rho | I_{1\eta}] + \beta_2 \tilde{\Omega}_{2\eta} \right\}, \quad (5.7)$$

where $\beta_1, \beta_2 > 0$. The first term represents the physical encounter delay, while the second represents the tracking distance. Restricting the problem to Π_2^{cp} prevents effort reduction from being achieved by abandoning the outer parts of the tracking trajectories. The discounted tracking-distance construction is not claimed to be optimal with respect to control energy or physical tracking time. It is used here as a coverage-preserving heuristic for reducing traveled distance within the admissible

tracking framework. Establishing energy- or time-optimal control laws would require a separate dynamic-control formulation and is left for future work.

5.6. Undiscounted limit

Proposition 5.2. *For every fixed ℓ , k , and σ , $\lim_{\eta \rightarrow 1} \Theta_{\ell k}^{\sigma}(\eta) = \Theta_{\ell k}^{\sigma}(1)$.*

Proof. For an ordinary cycle, $\lim_{\eta \rightarrow 1} \eta^{k-1} = 1$. Substitution into (5.2) gives $\lim_{\eta \rightarrow 1} \Theta_{\ell k}^{\sigma}(\eta) = \Theta_{\ell k}^{\sigma}(1)$. The equality already holds on every full-coverage cycle. $\square$

Therefore, the undiscounted Stage-2 policy is recovered as η approaches one. The resulting first-encounter and total-time properties are established in Section 6.

6. Main theoretical results

The previous sections have established the ingredients required for the complete two-stage policy. Stage 1 provides an optimal localization policy (r^*, γ^*) , while Sections 4 and 5 determine the adaptive discount $\eta = \Psi(U_1)$ and construct the corresponding tracking policy $\Theta(\eta)$. We now state the consequences for the first-encounter and total search times. Among the theoretical results, the principal results are the finiteness statements for the adaptive first-encounter and complete two-stage search times in Theorems 6.1 and 6.2. The remaining propositions and existence results provide the supporting admissibility, coverage, optimization, and consistency properties required for these main conclusions.

The theoretical results are organized sequentially. Theorems 3.1–3.3 establish the Stage-1 optimization and limiting-consistency properties, while Theorems 4.1 and 5.1 provide the Stage-2 finiteness and recurrent-coverage foundations. These results support Theorems 6.1 and 6.2, which give the principal finiteness results for the adaptive first-encounter time and the complete two-stage search time, respectively. Theorem 6.3 then establishes the existence of an optimal adaptive tracking policy for the finite-dimensional computational formulation.

6.1. Finiteness of the adaptive first-encounter time

Theorem 6.1. *Let $\Theta(\eta) \in \Pi_2^{cp}$ with $\eta = \Psi(U_1)$. Assume, in addition, that Condition A4 holds uniformly with constants $T^* < \infty$ and $\varepsilon^* > 0$, then*

$$P_{\Theta(\eta)}(\tau_{\rho} > kT^* | I_{1\eta}) \leq (1 - \varepsilon^*)^k, \quad k \geq 0, \quad (6.1)$$

and

$$E_{\Theta(\eta)}[\tau_{\rho} | I_{1\eta}] \leq \frac{T^*}{\varepsilon^*} < \infty. \quad (6.2)$$

Proof. By Proposition 5.1 and Theorem 5.1, every policy in Π_2^{cp} has bounded tracking distances, finite outward-and-return cycles, and recurrent coverage. Condition A4 supplies a uniform lower bound ε^* on the probability of encounter within a finite block of length T^* . The block argument used for the bounded controlled tracking model in [19] therefore applies directly. Conditional on no encounter before kT^* ,

$$P_{\Theta(\eta)}(\tau_\rho > (k+1)T^* | \tau_\rho > kT^*, I_{1\eta}) \leq 1 - \varepsilon^*.$$

Iteration gives (6.1). Hence,

$$E_{\Theta(\eta)}[\tau_\rho | I_{1\eta}] \leq T^* P_{\Theta(\eta)}(\tau_\rho > kT^* | I_{1\eta}) \leq T^* \sum_{k=0}^{\infty} (1 - \varepsilon^*)^k = \frac{T^*}{\varepsilon^*}.$$

Thus, the expected adaptive first-encounter time is finite. $\square$

The argument follows the finiteness mechanism established in [19], where the uniformly positive probability of entering the detection set over finite tracking blocks produces a geometric tail and therefore a finite mean encounter time. Importantly, this conclusion is a consequence of controlled recurrent tracking and not of the zero-mean random walk assumption alone.

6.2. Finiteness of the complete two-stage search time

Recall from Section 2 that $T_{tot} = T_1 + \tau_\rho$. The following result connects the two stages.

Theorem 6.2. *Let $(\mathbf{r}^*, \gamma^*)$ be an optimal Stage-1 policy and under the assumptions of Theorem 6.1, one can obtain $E[T_{tot}] \leq \Omega_1(\mathbf{r}^*) + \frac{T^*}{\varepsilon^*} < \infty$.*

Proof. Stage 1 contains finitely many regional expansion cycles. Since $t_{ij} = \frac{L_{ij}}{v_i}$, $v_i > 0$ and each admissible L_{ij} is finite, $E[T_1] \leq \Omega_1(\mathbf{r}^*) < \infty$. From (6.2), we obtain $E[T_{tot}] = E[T_1] + E[\tau_\rho] \leq \Omega_1(\mathbf{r}^*) + \frac{T^*}{\varepsilon^*}$. $\square$

Thus, the adaptive coupling does not compromise the finiteness of the original two-stage transport model. It changes how tracking effort is allocated after localization while retaining the finite first-encounter mechanism established in [19].

6.3. Existence of an optimal adaptive tracking policy

For the computational implementation Stage-2 problem (5.7), consider a finite number of tracking cycles and let the corresponding nominal tracking distances vary over their admissible bounds.

Theorem 6.3. *Suppose that the criterion in (5.7) is continuous in the nominal tracking distances and that the resulting finite-dimensional feasible set is nonempty and closed. Then, problem (5.7) admits at least one optimal adaptive tracking policy $\Theta^*(\eta) \in \Pi_2^{cp}$.*

Proof. The bounds imposed in Section 5 give $0 < \underline{\Theta}_\ell^\sigma \leq \Theta_{\ell k}^\sigma(1) \leq \overline{\Theta}_\ell^\sigma < \infty$. Hence, for a finite number of cycles, the feasible set is bounded. By assumption, it is also closed and nonempty, and therefore compact. The adaptive rule (5.1) is continuous in the nominal tracking distances. Consequently, the criterion in (5.7) is continuous on a compact feasible set. The Weierstrass theorem then guarantees the existence of a minimizer; see [22]. $\square$

This finite-dimensional result establishes the existence of an optimizer for the computational form of (5.7). The benchmark study in Section 7, however, evaluates the adaptive rule for prescribed nominal tracking distances rather than solving the full Stage-2 optimization problem.

6.4. Consistency of the complete adaptive model

The two discount mechanisms must recover the original two-stage formulation when discounting disappears.

Corollary 6.1. *For fixed admissible search and tracking policies, $\gamma \rightarrow 1$, $\eta \rightarrow 1$, imply $\tilde{\Omega}_{1\gamma}(\mathbf{r}) \rightarrow \tilde{\Omega}_1(\mathbf{r})$ and $\Theta_{\ell k}^\sigma(\eta) \rightarrow \Theta_{\ell k}^\sigma(1)$. Hence, the proposed adaptive discounted model converges to the undiscounted two-stage stochastic transport model.*

Proof. Equation $\tilde{\Omega}_{1\gamma}(\mathbf{r}) \rightarrow \tilde{\Omega}_1(\mathbf{r})$ follows from Theorem 3.3, while $\Theta_{\ell k}^\sigma(\eta) \rightarrow \Theta_{\ell k}^\sigma(1)$ follows from Proposition 5.2. No other component of the physical model depends on the limiting operation. In particular, the random-walk process $X(t)$, the detection radius ρ , and the definition of τ_ρ remain unchanged. $\square$

The limiting model is therefore consistent with the two-stage first-encounter framework of [19], which couples regional localization with coordinated multidimensional tracking. The published model establishes finite first-encounter behavior under bounded controlled tracking and recurrent coverage, precisely the structure retained here.

7. Numerical results

The numerical study is intended as a controlled illustration of the discounted and adaptive mechanisms under the benchmark geometry of [19], rather than as a full numerical solution of the general optimization problems (3.7) and (5.7). In Stage 1, the expansion distances are fixed at the benchmark values $r_{1j} = j$, and the discount parameter γ is evaluated over its admissible range. In Stage 2, the benchmark nominal tracking distances are retained and transformed through the adaptive rule (5.1). The general formulations in Sections 3 and 5 therefore establish the optimization framework, while Section 7 evaluates its behavior under the prescribed benchmark geometry. To broaden the numerical validation beyond the direct comparison with [19], we additionally examine whether a population-based optimization method can recover the Stage-1 discount obtained by the scalar benchmark optimization. Specifically, PSO is applied to the same constrained Stage-1 objective with the benchmark expansion distances fixed. This comparison concerns the numerical optimization procedure only and does not alter the benchmark geometry or imply a direct comparison with neural-network synchronization models.

The numerical example in [19] is used as the benchmark for evaluating the present adaptive discounted model. This benchmark is deliberately retained to provide a controlled comparison, so that the numerical differences reported below can be attributed to the newly introduced discounted and adaptive mechanisms rather than to changes in the underlying search geometry. The physical setting is kept unchanged: a single square region with $b_1 = 10$, center $O_1 = (0,0)$, unit search speed, and a symmetric truncated bivariate normal distribution with variance 9. The representative Stage-1 localization point is $\hat{O} = (3.2, -1.8)$, and the reported expected localization time is $\Omega_1 = 14.6$. To preserve the unit expansion used in [19], we take $K_1 = 10$, $\delta_1 = 1$, $r_{1j} = j$, $j = 1, 2, \dots, 10$. The parameters introduced by the present model are fixed at $\alpha_1 = 1$, $\alpha_2 = 0.25$, $\alpha_3 = 4$, $\gamma_{\min} = 0.5$, $\tilde{R}_{\min} = 0.5$, $\eta_{\min} = 0.5$, $m = 5$, $\underline{\Theta}_\ell^\sigma = 0.1$, and $\overline{\Theta}_\ell^\sigma = 1$. The objective weights are treated as relative preference parameters rather than universal constants. We set $\alpha_1 = 1$ as the reference scale and express the remaining weights relative to it, with $\alpha_2 = 0.25$ assigning a smaller penalty to discounted movement effort, and $\alpha_3 = 4$ assigning greater emphasis to early detection. Under the fixed benchmark

expansion policy, the physical localization-time term is independent of γ ; hence, α_1 does not affect the reported optimal discount, which is determined by the trade-off between the movement-effort and detection-reward terms. For other applications, these relative weights can be recalibrated according to the desired operational trade-off. Thus, the benchmark geometry is intentionally kept fixed, and the numerical study evaluates the proposed discounting and adaptive mechanisms within this prescribed geometry.

7.1. Stage 1 results

The probability mass over the 10 expanding shells is reported in Table 1.

Table 1. Probability mass of the Stage-1 search shells.

j	1	2	3	4	5	6	7	8	9	10
p_{1j}	0.0683	0.1772	0.2214	0.2027	0.1498	0.0933	0.0501	0.0237	0.0099	0.0037

The mass is concentrated mainly in the second to fifth shells. With the benchmark expansion policy $r_{1j} = j$ fixed, the Stage-1 criterion is minimized with respect to γ over its admissible interval. This gives $\gamma^* = 0.7650$ with $\tilde{\Omega}_{1\gamma^*} = 6.0603$ and $\tilde{R}_{1\gamma^*} = 0.5151$. Thus, $\gamma^* = 0.7650$ is the optimum for the fixed benchmark expansion policy, rather than for the unrestricted joint optimization in (r, γ) . The reward constraint is therefore satisfied because $\tilde{R}_{1\gamma^*} > \tilde{R}_{\min}$. The numerical effect of discounting is summarized in Table 2.

Table 2. Stage 1 comparison between the benchmark and the discounted model.

Quantity	Benchmark	Present model
Ω_1	14.6	14.6000
γ	1.0	0.7650
$\tilde{\Omega}_{1\gamma}$	14.6	6.0603
$\tilde{R}_{1\gamma}$	1.0	0.5151

The physical localization time remains unchanged at 14.6, whereas the discounted movement-effort functional decreases from 14.6 to 6.0603, corresponding to a reduction of 58.49%. This percentage refers only to the discounted objective component and should not be interpreted as an equivalent reduction in physical travel distance, localization time, or energy consumption.

As an additional numerical check, the Stage-1 discount was also examined using PSO, while keeping the benchmark expansion policy $r_{1j} = j$ and all other model parameters fixed. The same objective and admissibility constraints were retained. This algorithmic comparison is intended only as an independent verification of the reported Stage-1 optimum and does not alter the benchmark geometry or the formulation of the optimization problem. The PSO procedure converged to $\gamma^* = 0.7650$, consistent with the direct scalar optimization, thereby providing an independent numerical verification of the reported Stage-1 optimum.

7.2. Adaptive Stage-2 policy

For Stage 2, the full optimization problem (5.7) is not solved numerically in this benchmark study. Instead, the nominal tracking distances are kept fixed, and the adaptive rule (5.1) is applied to quantify the effect of the uncertainty-dependent discount. For the benchmark case, $N=1$ and $q=1$. Hence, $\tilde{w}_{1j} = q_1 p_{1j} = p_{1j}$, so the shell probabilities listed in Table 1 are directly used to evaluate the Stage-1 uncertainty:

$$U_1 = -\frac{1}{10} \sum_{j=1}^{10} p_{1j} \log p_{1j} = 0.8503.$$

Thus, the reported value measures the dispersion of the probability mass over the 10 search shells, rather than a Bayesian posterior uncertainty.

Figure 2 shows the benchmark operating point $(U_1, \eta) = (0.8503, 0.9251)$. The position of this point near the upper-right part of the coupling curve indicates that substantial Stage-1 shell uncertainty leads to a conservative Stage-2 discount. The first 10 tracking cycles are listed in Table 3.




Figure 2. Linear adaptive coupling between the Stage-1 shell uncertainty U_1 and the Stage-2 discount η .

Table 3. Undiscounted and adaptive tracking distances over the first 10 cycles.

k	$\Theta_{\ell k}^{\sigma}(1)$	$\Theta_{\ell k}^{\sigma}(\eta)$	$\frac{\Theta_{\ell k}^{\sigma}(\eta)}{\Theta_{\ell k}^{\sigma}(1)}$
1	1.0000	1.0000	1.0000
2	0.5000	0.4701	0.9402
3	0.3333	0.2997	0.8992
4	0.2500	0.2188	0.8752
5	1.0000	1.0000	1.0000
6	0.1667	0.1452	0.8712
7	0.1429	0.1269	0.8880
8	0.1250	0.1145	0.9160
9	0.1111	0.1060	0.9541
10	1.0000	1.0000	1.0000

The adaptive distances are smaller on the ordinary cycles, while the full-coverage cycles $k=5$ and 10 are unchanged.

Figure 3 makes this distinction visible: the adaptive trajectory lies below the undiscounted one between the full-coverage cycles, but both coincide at $k=5$ and 10 . Thus, the reduction in tracking distance is achieved without removing the periodic coverage required by the model. Also, Figure 4 provides the same comparison in normalized form. The ratio remains between approximately 0.87 and 1 , and returns to 1 at every full-coverage cycle. This confirms that the adaptive rule modifies the tracking distances gradually rather than collapsing them.




Figure 3. Cycle-by-cycle comparison of the undiscounted and adaptive tracking distances.




Figure 4. Cycle-wise contraction ratio $\frac{\Theta_{\ell k}^{\sigma}(\eta)}{\Theta_{\ell k}^{\sigma}(1)}$.

7.3. Overall comparison

For the first 10 cycles, the total tracking distance of the benchmark policy is 37.0317, whereas the adaptive policy gives 35.8485. The corresponding reduction is 3.20%. The main numerical results are collected in Table 4.

Table 4. Overall numerical comparison with the benchmark model.

Performance quantity	Benchmark	Present model	Change
Physical Stage-1 localization time Ω_1	14.6	14.6	0%
Discounted Stage-1 movement effort $\tilde{\Omega}_{1\gamma}$	14.6	6.0603	-58.49%
U_1	--	0.8503	--
η	1.0	0.9251	--
10-cycle tracking distance	37.0317	35.8485	-3.20%

Figure 5 summarizes the numerical comparison. The Stage-1 localization time is preserved, while the discounted Stage-1 effort is substantially reduced. The Stage-2 reduction is smaller because the relatively large value of U_1 keeps η close to one. This behavior is consistent with the purpose of the adaptive coupling: strong contraction is avoided when the Stage-1 shell probability mass is more dispersed.




Figure 5. Relative performance gains of the adaptive discounted model with respect to the benchmark policy.

Overall, the numerical experiment shows that the proposed model reduces weighted movement effort while retaining the physical localization time and the periodic Stage-2 coverage structure of the benchmark. The largest gain occurs in Stage 1, whereas Stage 2 responds more conservatively to the Stage-1 shell uncertainty.

7.4. Sensitivity analysis

To examine the behavior of the proposed model beyond the single benchmark setting, a compact sensitivity analysis was carried out for the parameters that directly control the adaptive mechanism. Unless otherwise stated, all remaining parameters were kept at their benchmark values.

First, the effect of the Stage-1 shell uncertainty was examined by varying U_1 while keeping $\eta_{\min}$ and $m = 5$. The coupling rule $\eta = \eta_{\min} + (1 - \eta_{\min})U_1$ was used in each case. Table 5 shows the resulting Stage-2 discount and the total tracking distance over the first 10 cycles.

Table 5. Sensitivity of the Stage-2 policy to the Stage-1 shell uncertainty.

U_1	η	10-cycle tracking distance
0.3000	0.6500	32.8986
0.6000	0.8000	34.2906
0.8503	0.9251	35.8485
0.9500	0.9750	36.6113

The results show the intended monotone response of the coupling mechanism. Smaller values of U_1 lead to stronger discounting and shorter ordinary tracking distances, whereas larger values of U_1 move η closer to one and preserve a wider tracking pattern.

The lower Stage-2 discount bound was also varied while keeping $U_1 = 0.8503$ and $m = 5$. The corresponding results are reported in Table 6.

Table 6. Sensitivity with respect to $\eta_{\min}$.

$\eta_{\min}$	η	10-cycle tracking distance
0.3	0.8952	35.4334
0.5	0.9251	35.8485
0.7	0.9551	36.2952

As expected, increasing $\eta_{\min}$ weakens the contraction of the ordinary cycles. Hence, the total tracking distance increases gradually as the admissible lower bound on the Stage-2 discount becomes more conservative.

To assess robustness with respect to the discounting level, we also compare several fixed discount settings while keeping the target dynamics, tracking geometry, and all other parameters unchanged. The undiscounted case gives a 10-cycle tracking distance of 37.0317, whereas discount values of 0.9750, 0.9251, 0.8, and 0.65 give 36.6113, 35.8485, 34.2906, and 32.8986, respectively. The monotone variation indicates that the reduction in tracking distance is stable across a broad range of discount levels rather than being tied to a single calibrated value.

To provide a direct comparison between the adaptive and undiscounted Stage-2 policies, we repeated the two-dimensional coordinate-direction Monte Carlo experiment under identical target dynamics, tracking geometry, initial condition, detection radius, searcher speed, simulation horizon, and random-number seed. The only change was the Stage-2 discount, with $\eta = 0.9251$ for the adaptive policy and $\eta = 1$ for the undiscounted policy. This comparison therefore isolates the effect of discounting on both traveled distance and first-encounter performance. The corresponding results are summarized in Table 7, which compares the adaptive and undiscounted policies in terms of

tracking distance and first-encounter performance.

Table 7. Direct comparison of the adaptive and undiscounted Stage-2 policies under identical target dynamics and tracking geometry.

Policy	η	10-cycle tracking distance	$\hat{P}(\tau_\rho \leq T_{\max})$	$\hat{E}[\tau_\rho \tau_\rho \leq T_{\max}]$	$\hat{E}[\tau_\rho \wedge T_{\max}]$
Adaptive	0.9251	35.8485	0.8793	200.17	779.51
Undiscounted	1.0000	37.0317	0.8860	210.47	756.47

Under the same target dynamics and tracking geometry, the adaptive policy reduces the 10-cycle traveled distance by approximately 3.20%. The undiscounted policy shows a slightly higher finite-horizon encounter proportion, while the adaptive policy yields a slightly lower conditional mean encounter time among successful realizations. These results illustrate the trade-off between reduced tracking distance and first-encounter performance introduced by discounting.

The influence of the full-coverage period m was examined next. The first 10 cycles were retained so that the lower tracking bound $\Theta = 0.1$ remained compatible with the nominal sequence used in the benchmark. Table 8 compares three representative values of m .

Table 8. Effect of the full-coverage period m .

m	Benchmark distance	Adaptive distance	Reduction
2	54.2984	53.6464	1.20%
5	37.0317	35.8485	3.20%
10	30.6317	29.2346	4.56%

More frequent full-coverage cycles increase the total traveled distance because the searchers repeatedly return to the undiscounted outer distance. In contrast, larger values of m allow more ordinary discounted cycles between successive coverage-restoration cycles and therefore produce a larger reduction in tracking effort. In all cases, periodic full coverage is retained.

The sensitivity with respect to $\gamma_{\min}$ was examined using $\gamma_{\min} \in \{0.4, 0.6, 0.8\}$. The remaining Stage-1 parameters were kept at their benchmark values. For $\gamma_{\min} = 0.4$ and 0.6 , the lower bound is inactive because the benchmark optimum $\gamma^* = 0.765$ remains feasible. Hence, the optimizer remains unchanged and $\tilde{R}_{1\gamma^*} = 0.5151$. When $\gamma_{\min} = 0.8$, the previous optimum is no longer feasible, and the solution moves to the lower admissible boundary, $\gamma^* = 0.8$, for which $\tilde{R}_{1\gamma^*} = 0.5677$. The results are summarized in Table 9.

Table 9. Sensitivity of the Stage-1 solution with respect to $\gamma_{\min}$.

$\gamma_{\min}$	γ^*	$\tilde{R}_{1\gamma^*}$
0.4	0.765	0.5151
0.6	0.765	0.5151
0.8	0.8	0.5677

Thus, $\gamma_{\min}$ has no effect as long as it remains below the unconstrained optimal value. Once the lower bound exceeds 0.765, it becomes active and directly determines the admissible Stage-1

discount. For the dimension, geometry, and first-encounter experiments, the Stage-2 simulation parameters were specified explicitly. The coordinate system was translated so that $\hat{O}$ is the origin of the tracking system. For the d -dimensional coordinate configuration, the target was initialized at $X(0) = \hat{O} + \frac{0.5}{\sqrt{d}} \sum_{\ell=1}^d e_{\ell}$, so that the initial Euclidean displacement from $\hat{O}$ is the same in all tested dimensions, and the simulation does not begin with an immediate encounter. The random-walk step lengths were fixed at $s_{\ell} = 0.1$, and the detection radius was set to $\rho = 0.1$. At each integer time step, the target increment was selected uniformly from $\{\pm 0.1e_1, \pm 0.1e_2, \dots, \pm 0.1e_d\}$, so that each increment has probability $1/(2d)$, and the zero-mean condition in (2.18) is preserved. The adaptive parameters were retained at their benchmark values $U_1 = 0.8503$, $\eta = 0.9251$, and $m = 5$, with the tracking bounds retained at their benchmark values, $\Theta_{\ell}^{\sigma} = 0.1$ and $\bar{\Theta}_{\ell}^{\sigma} = 1$. The searchers moved with unit speed, as in the benchmark model. For cycles beyond those listed in Table 3, the nominal sequence was continued by

$$\Theta_{\ell k}^{\sigma}(1) = \begin{cases} 1, & k = rm, \ r \geq 1, \\ \max\left\{0.1, \frac{1}{k}\right\} & \text{Otherwise,} \end{cases}$$

and the adaptive distances were then obtained from (5.1). This extension preserves the lower tracking bound while retaining the periodic full-coverage cycles.

The effect of target-motion dimension was examined for $d \in \{2, 3, 4\}$. For each dimension, the tracking trajectories were taken along the corresponding coordinate directions through $\hat{O}$. The random-walk transition law, detection radius, adaptive discount, full-coverage period, and tracking bounds were otherwise kept unchanged. This allows the effect of dimensionality to be examined without altering the adaptive coupling mechanism.

The effect of tracking geometry was examined separately in two dimensions. The coordinate-direction configuration was compared with three intersecting trajectories through $\hat{O}$, having unit direction vectors $e'_1 = (1, 0)$, $e'_2 = \left(\frac{1}{2}, \frac{\sqrt{3}}{2}\right)$, and $e'_3 = \left(-\frac{1}{2}, \frac{\sqrt{3}}{2}\right)$. Importantly, the target random-walk law was kept unchanged in this comparison, with increments still selected from $\{\pm 0.1e_1, \pm 0.1e_2\}$. Thus, only the search geometry was changed, while the target dynamics and all adaptive parameters were held fixed.

A Monte Carlo experiment with $N_{MC} = 10^4$ independent realizations was performed for each configuration. The random-number seed was fixed at 20,260,919 for reproducibility. Because the target process is defined at discrete time steps, encounters were checked at the target-update times $t = 0, 1, 2, \dots$, using the distance criterion in (2.21). A maximum simulation horizon of $T_{\max} = 5000$ was used.

For each configuration, we recorded the empirical encounter proportion $\hat{P}(\tau_{\rho} \leq T_{\max})$, the mean encounter time among realizations that encountered the searchers before $T_{\max}$, and the restricted mean first-encounter time $\hat{E}[\tau_{\rho} | \tau_{\rho} \leq T_{\max}]$, and the restricted mean first-encounter time

$$\hat{E}[\tau_{\rho} \wedge T_{\max}] = \frac{1}{N_{MC}} \sum_{r=1}^{N_{MC}} \min\{\tau_{\rho}^{(r)}, T_{\max}\}.$$

The restricted mean is reported because the Monte Carlo experiment is finite-horizon and is not

intended to establish Condition A4. The theoretical finiteness result therefore remains conditional on A4.

As shown in Table 10, the simulations show a clear deterioration in finite-horizon encounter performance as the target-motion dimension increases under the coordinate-direction geometry. The encounter proportion decreases from 0.8793 in two dimensions to 0.4166 in three dimensions and 0.1581 in four dimensions. In the two-dimensional geometry comparison, adding a third intersecting trajectory improves the encounter proportion from 0.8793 to 0.9398 and reduces the restricted mean first-encounter time from 779.51 to 411.89. These results illustrate the sensitivity of first-encounter behavior to dimensionality and tracking geometry. They do not, however, constitute a proof of finite expected encounter time; that conclusion remains conditional on Condition A4.

Table 10. Monte Carlo first-encounter statistics for different target-motion dimensions and tracking geometries.

d	Tracking geometry	N_{MC}	$\hat{P}(\tau_\rho \leq T_{\max})$	$\hat{E}[\tau_\rho \tau_\rho \leq T_{\max}]$	$\hat{E}[\tau_\rho \wedge T_{\max}]$
2	Coordinate directions	10000	0.8793	200.17	779.51
3	Coordinate directions	10000	0.4166	378.11	3074.52
4	Coordinate directions	10000	0.1581	226.36	4245.29
2	Alternative intersecting geometry	10000	0.9398	117.99	411.89

The present simulations also provide a controlled comparison with two related multi-agent tracking configurations. Under the two-dimensional coordinate-direction strategy, the encounter proportion is 0.8793, and the restricted mean first-encounter time is 779.51, whereas the alternative three-trajectory configuration increases the encounter proportion to 0.9398 and reduces the restricted mean to 411.89. At the same time, the adaptive discounted policy reduces the 10-cycle tracking distance from 37.0317 to 35.8485 while preserving periodic full coverage. These comparisons quantify the trade-off between encounter performance, traveled distance, and coverage preservation within the common stochastic-target setting. They are not intended as direct numerical comparisons with the different multi-UAV and multi-robot formulations in [4,5], whose objectives and motion models are not identical to those considered here.

7.5. Preliminary extension to fractional-order multidimensional systems

A preliminary extension of the present framework to fractional-order multidimensional dynamical systems can also be envisioned. In such a setting, the random-walk target $X(t)$ would be replaced by the state trajectory of a fractional-order multidimensional system, while the coordinated searchers could be interpreted as controllers whose trajectories are adjusted through the uncertainty-dependent discounting mechanism. The encounter condition would then correspond to driving the controlled trajectories into a prescribed neighborhood of the system state, leading to a synchronization-type objective. This interpretation is conceptually related to recent fractional-order neural-network synchronization studies such as [9]. However, the fractional dynamics, stability conditions, and synchronization analysis are outside the scope of the present paper; establishing these properties is left for future work.

8. Conclusions and future work

This paper developed an adaptive discounted two-stage stochastic transport model for the localization and subsequent tracking of a multidimensional random-walk target. The main extension of the baseline formulation is that discounting is introduced in both stages without altering the physical coverage requirements. In Stage 1, the objective combines the actual localization time, discounted movement effort, and discounted detection reward. In Stage 2, the tracking discount is determined by the Stage-1 shell uncertainty through $\eta = \Psi(U_1)$, so that the tracking policy responds directly to the Stage-1 shell-probability dispersion summarized by U_1 .

The analysis established the existence of optimal discounted policies under the stated compactness and continuity assumptions and the preservation of complete regional and recurrent Stage-2 coverage. The finiteness of the first-encounter and total search times is obtained additionally under the uniform positive finite-block encounter-probability assumption. The limiting case $\gamma \rightarrow 1$ and $\eta \rightarrow 1$ recovers the undiscounted two-stage model, which confirms consistency with the original framework.

The numerical study based on the benchmark in [19] showed that the proposed model reduces the discounted Stage-1 movement-effort functional while preserving the physical localization time and the underlying search geometry. For the selected parameters, the discounted movement-effort functional decreased from 14.6 to 6.0603, corresponding to a reduction of 58.49%. This reduction concerns the discounted performance measure only and does not represent an equivalent reduction in physical distance, search time, or energy consumption. The Stage-1 shell uncertainty $U_1 = 0.8503$ produced $\eta = 0.9251$, leading to a more moderate reduction of 3.20% in the 10-cycle Stage-2 tracking distance. This behavior reflects the intended adaptive mechanism: stronger contraction is avoided when the Stage-1 shell probability mass remains broadly dispersed, while the prescribed full-coverage cycles are preserved.

These results indicate that the proposed framework provides a practical balance between search effort, localization information, and recurrent tracking coverage. Further research may consider nonuniform and state-dependent random-walk dynamics, noisy sensing, data-driven calibration of the adaptive coupling rule, and energy- or time-optimal control formulations for the tracking stage. Another promising direction is to investigate dual-control strategies within the proposed two-stage framework. In particular, dual control could provide a systematic mechanism for balancing exploration of uncertain target locations with exploitation of the information already obtained during localization and tracking. Recent k -step look-ahead dual-control approaches combine online learning with exploration-exploitation decisions for auto-optimization [25], suggesting a possible basis for developing a real-time adaptive extension of the present model in which search and tracking actions are updated as new information becomes available. The proposed fractional-order extension also warrants further investigation, particularly with respect to stability conditions and synchronization analysis.

Author contributions

Mohamed Abd Allah El-Hadidy: Conceptualization, methodology, formal analysis, investigation, supervision, writing-original draft, writing-review & editing; Alaa Awad Alzulaibani: Methodology, formal analysis, software, validation, visualization, writing-review & editing. All authors have read and approved the final version of the manuscript for publication.

Use of Generative-AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The authors declare that there are no conflicts of interest.

References

1. M. R. Evans, S. N. Majumdar, Diffusion with stochastic resetting, *Phys. Rev. Lett.*, **106** (2011), 160601. <https://doi.org/10.1103/PhysRevLett.106.160601>
2. V. Méndez, R. Flaquer-Galmés, D. Campos, First-passage time of a Brownian searcher with stochastic resetting to random positions, *Phys. Rev. E*, **109** (2024), 044134. <https://doi.org/10.1103/PhysRevE.109.044134>
3. A. Biswas, A. Dubey, A. Kundu, A. Pal, Drift-diffusive resetting search process with stochastic returns: Speedup beyond optimal instantaneous return, *Phys. Rev. E*, **111** (2025), 014142. <https://doi.org/10.1103/PhysRevE.111.014142>
4. D. Zhang, H. Duan, Variational-resolution probability maps for multi-UAV search tasks with changeable searching behaviors, *Aerosp. Sci. Technol.*, **155** (2024), 109669. <https://doi.org/10.1016/j.ast.2024.109669>
5. P. Xin, P. Dames, Comparison of stochastic optimization strategies in multi-robot multi-target tracking scenarios, *Swarm Intell.*, **19** (2025), 215–243. <https://doi.org/10.1007/s11721-025-00249-y>
6. D. Boyer, G. Mercado-Vásquez, S. N. Majumdar, G. Schehr, Optimizing the random search of a finite-lived target by a Lévy flight, *Phys. Rev. E*, **109** (2024), L022103. <https://doi.org/10.1103/PhysRevE.109.L022103>
7. S. Feng, T. Weng, Y. Wang, Y. Xu, Z. Ren, L. Zhan, Random search processes on complex networks: From a static target to a moving object, *Physica A*, **636** (2024), 129544. <https://doi.org/10.1016/j.physa.2024.129544>
8. J. Lin, M. Wang, H. Yan, W. Yang, Prescribed-time optimal tracking control for a class of stochastic systems using reinforcement learning, *J. Franklin Inst.*, **362** (2025), 107881. <https://doi.org/10.1016/j.jfranklin.2025.107881>
9. J. Xiao, K. Shi, Y. Teng, J. Qi, H. Fan, Relaxed research on synchronization problem of fractional-order fuzzy octonion-valued BAM neural networks by the non-decomposition method on the high-dimension oblique field, *Fractal Fract.*, **10** (2026), 414. <https://doi.org/10.3390/fractalfract10060414>
10. L. D. Stone, Necessary and sufficient conditions for optimal search plans for moving targets, *Math. Oper. Res.*, **4** (1979), 303–477. <https://doi.org/10.1287/moor.4.4.431>
11. S. S. Brown, Optimal search for a moving target in discrete time and space, *Oper. Res.*, **28** (1980), 1259–1453. <https://doi.org/10.1287/opre.28.6.1275>
12. J. N. Eagle, J. R. Yee, An optimal branch-and-bound procedure for the constrained path, moving target search problem, *Oper. Res.*, **38** (1990), 2–184. <https://doi.org/10.1287/opre.38.1.110>

13. L. D. Stone, J. O. Royset, A. R. Washburn, *Optimal search for moving targets*, Cham: Springer, 2016. <https://doi.org/10.1007/978-3-319-26899-6>
14. S. Redner, *A guide to first-passage processes*, Cambridge: Cambridge University Press, 2012. <https://doi.org/10.1017/CBO9780511606014>
15. O. Bénichou, C. Loverdo, M. Moreau, R. Voituriez, Intermittent search strategies, *Rev. Mod. Phys.*, **83** (2011), 81–129. <https://doi.org/10.1103/RevModPhys.83.81>
16. T. Alam, L. Bobadilla, Multi-robot coverage and persistent monitoring in sensing-constrained environments, *Robotics*, **9** (2020), 47. <https://doi.org/10.3390/robotics9020047>
17. R. Hohzaki, K. Iida, Optimal strategy of route and look for the path constrained search problem with reward criterion, *Eur. J. Oper. Res.*, **100** (1997), 236–249. [https://doi.org/10.1016/S0377-2217\(96\)00167-1](https://doi.org/10.1016/S0377-2217(96)00167-1)
18. M. A. El-Hadidy, On the maximum discounted effort-reward search problem, *Asia-Pac. J. Oper. Res.*, **33** (2016), 1650019. <https://doi.org/10.1142/S0217595916500196>
19. M. A. El-Hadidy, A. A. Alzulaibani, A multidimensional first-encounter time model for random-walk search based on a two-stage stochastic transport process, *AIMS Math.*, **11** (2026), 17722–17765. <https://doi.org/10.3934/math.2026723>
20. M. A. El-Hadidy, A. A. Alzulaibani, A study on the finiteness of a tracking method with reduction in the collision time between a D-dimensional random-walk particle and one of multiple nano-sensors, *J. Comput. Theor. Transp.*, **52** (2023), 506–534. <https://doi.org/10.1080/23324309.2024.2303506>
21. L. Wang, W. Sun, S. F. Su, Adaptive asymptotic tracking control for nonlinear systems with state constraints and input saturation, *Appl. Math. Comput.*, **431** (2022), 127342. <https://doi.org/10.1016/j.amc.2022.127342>
22. M. Galewski, *Basics of nonlinear optimization: Around the Weierstrass theorem*, Cham: Birkhäuser, 2024. <https://doi.org/10.1007/978-3-031-77160-6>
23. A. Beck, *Introduction to nonlinear optimization: Theory, algorithms, and applications with Python and MATLAB*, 2 Eds., Philadelphia: SIAM, 2023. <https://doi.org/10.1137/1.9781611977622>
24. Ş. Cobzaş, R. Miculescu, A. Nicolae, *Lipschitz functions*, Cham: Springer, 2019. <https://doi.org/10.1007/978-3-030-16489-8>
25. Y. Yu, J. Jiang, W. H. Chen, Y. Zuo, k-Step look-ahead active concurrent learning-based dual control of exploration and exploitation for auto-optimization, *IEEE Trans. Cybern.*, **56** (2026), 3804–3817. <https://doi.org/10.1109/TCYB.2026.3660400>



AIMS Press

©2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)