Chapter 7

A Stigmergy-Based Multi-Robot Search Strategy for Post-Earthquake Rubble Environments

Mehmet Dinçer Erbaş[1]

Abstract

Post-earthquake search and rescue operations require rapid exploration under uncertain conditions, high obstacle density and limited communication availability. In such environments multi-robot systems provide advantages in scalability and parallel exploration, yet effective coordination without centralized control remains a critical challenge. This study proposes a **Multi-Component Stigmergic Search (MCSS)** approach that extends conventional stigmergy-based coordination by integrating multiple environmental decision factors, including pheromone intensity, target-generated indirect signals, robot density and visitation history. The proposed approach was evaluated in a grid-based simulation environment representing rubble conditions and compared with a non-stigmergic exploration strategy and a fully random search method under different obstacle densities. Performance was assessed in terms of target discovery over time, cost per target, and average route length per target across repeated simulation runs. The results demonstrate that MCSS consistently achieves faster exploration, lower search cost, and shorter route lengths than the comparison approaches while maintaining stable performance under increasing environmental complexity. These findings suggest that combining stigmergic indirect communication with multi-component environmental guidance can improve the efficiency and robustness of autonomous multi-robot search strategies for disaster response scenarios.

1 Doç. Dr., Bolu Abant İzzet Baysal Üniversitesi, dincer.erbas@ibu.edu.tr, ORCID ID: 0000-0003-1762-0428

https://doi.org/10.58830/ozgur.pub1351.c5540

1. Introduction

Search and rescue operations conducted in post-earthquake rubble environments involve highly uncertain and physically hazardous conditions. In such scenarios, reaching victims within the shortest possible time is of critical importance. However, factors such as irregular debris structures, narrow passages, limited visibility, and the risk of secondary collapses significantly complicate conventional human-based search efforts. For this reason, increasing attention has been directed toward the use of robotic systems in search and rescue operations in recent years [1], [2], [3]. In particular, multi-robot systems offer significant advantages for search and rescue tasks due to their ability to perform parallel exploration, scale to large operational areas, and maintain robustness against individual robot failures [2].

Despite these advantages, achieving effective coordination among multiple robots without centralized control remains a major challenge [4]. A considerable portion of conventional coordination approaches relies on direct communication and global information sharing. However, in disaster environments, communication infrastructures may become unreliable or partially unavailable, thereby limiting the effectiveness of communication-dependent systems.

In this context, stigmergy provides an alternative coordination mechanism based on indirect communication through the environment [5]. In stigmergy-based approaches, robots can indirectly influence each other's behaviors through traces left in the environment, enabling collective search behaviors to emerge without centralized control. These characteristics make stigmergy a particularly promising approach for complex and communication-constrained search and rescue environments.

In this study, a stigmergy-based multi-robot search approach is investigated within a grid-based simulation environment representing post-earthquake rubble areas. In the proposed method, robot movement decisions are determined not only by pheromone intensity but also by multiple environmental factors, including target-generated signals, robot density, and visitation history. Furthermore, the performance of the proposed approach is comparatively evaluated against a non-stigmergic exploration strategy and a fully random search method. Within the scope of the study, performance metrics such as target discovery rate, cost per target, and route length are analyzed under different obstacle densities. The proposed framework also extends an earlier prototype simulation environment and stigmergy-based search structure previously investigated in a master's thesis study [6].

From a broader computational intelligence perspective, the proposed framework can also be interpreted within the context of emerging adaptive decision-making paradigms in artificial intelligence. Rather than relying on centralized planning or explicit communication, the proposed approach utilizes distributed environmental information and simple local interaction rules to generate collective exploration behavior. In this respect, the study reflects key principles of computational intelligence, including decentralized optimization, adaptive information processing, and environment-driven decision mechanisms. Furthermore, the integration of multiple environmental decision factors into a unified search strategy aligns with contemporary perspectives that emphasize scalable intelligent systems capable of extracting useful behavioral patterns from large and dynamically changing information spaces.

The main contributions of this study can be summarized as follows:

- A Multi-Component Stigmergic Search (MCSS) approach is proposed for multi-robot search in post-earthquake rubble environments by extending conventional stigmergy-based coordination with additional environmental decision components.
- The proposed decision mechanism integrates multiple environmental factors, including pheromone intensity, target-generated indirect signals, robot density, and visitation history, within a unified probabilistic movement-selection framework.
- A comparative evaluation framework is established by comparing the proposed MCSS approach against both a non-stigmergic exploration strategy and a fully random search method, enabling the isolated assessment of the contribution of stigmergic coordination.
- The proposed approach is evaluated under different obstacle-density scenarios using multiple performance indicators, including target discovery speed, cost per target, and average route length per target.
- The study extends an earlier prototype simulation framework through redesigned search dynamics, expanded environmental decision components, and a more comprehensive performance evaluation methodology for disaster-oriented multi-robot exploration.

2. Related Work

Structural collapse environments caused by disasters such as earthquakes, floods, explosions, and similar catastrophic events represent some of the most challenging operational settings for search and rescue missions. In such environments, reaching victims rapidly is critical for increasing survival probability. However, rubble areas pose significant dangers to human rescue teams due to narrow passages, irregular obstacle distributions, limited visibility, dust, toxic gases, extreme temperatures, and the risk of secondary collapses. For this reason, the use of robotic systems as supportive tools in post-disaster search and rescue operations has been extensively investigated for many years. Studies in the field of Robotic Urban Search and Rescue (RUSAR) have demonstrated that robots can effectively be employed for exploration and information gathering in hazardous, inaccessible, or high-risk environments [1], [7]. In their review of urban search and rescue robots from a control perspective, Liu and Nejat identified mobility, sensing, mapping, autonomy, and human–robot interaction as the major research challenges in this domain [1]. Similarly, Drew emphasized that multi-robot systems provide substantial advantages for search and rescue applications in terms of scalability, robustness, and parallel exploration capabilities [2].

Teleoperation played a central role in early search and rescue robotic applications [1]. In teleoperated systems, robots are remotely controlled by human operators based on data obtained from onboard cameras and sensors [8]. Although this approach has demonstrated practical applicability in real disaster environments, it also presents several limitations. Restricted operator field of view, insufficient situational awareness, communication delays, and complex rubble structures may cause robots to become trapped or lose orientation. Experiences gained from robot-assisted search and rescue efforts following the September 11 attacks further highlighted the critical importance of human–robot interaction, situational awareness, and communication reliability in disaster operations [7]. Consequently, later studies focused not on completely abandoning teleoperation, but rather on increasing the autonomy level of robots and reducing operator workload [9], [10]. Nevertheless, single and operator-controlled robotic systems remain limited in their ability to rapidly scan large areas, simultaneously search for multiple targets, and maintain robustness against hardware failures.

These limitations have increased interest in the use of multi-robot systems for search and rescue problems. In multi-robot systems, tasks are distributed among multiple relatively simple robots instead of relying on a single highly complex platform [2]. This structure enables parallel exploration of the search area and improves system resilience against individual robot failures. Furthermore, increasing the number of robots allows the system to scale more effectively to larger or more complex operational environments. These characteristics make multi-robot systems particularly advantageous in uncertain, fragmented, and communication-constrained environments such as rubble fields [2]. However, one of the fundamental challenges in multi-robot systems is achieving effective coordination without dependence on a centralized control unit [4]. Such coordination can be achieved through various mechanisms, including direct communication, global information sharing, or indirect communication through environmental traces.

Swarm intelligence approaches are widely employed in multi-robot exploration and search problems [11], [12]. Swarm intelligence is based on the emergence of complex collective behaviors from the interactions of numerous individuals operating according to relatively simple rules. Among the most well-known methods in this field are Particle Swarm Optimization and Ant Colony Optimization [13], [14]. Particle Swarm Optimization (PSO) is an optimization approach in which individuals update their movement directions using both their own experiences and the best experience of the group. Although this structure can be adapted to multi-robot search problems, it generally requires global information sharing or direct inter-robot communication. In rubble environments, however, the reliability of direct wireless communication may deteriorate due to concrete debris, metallic structural elements, irregular geometries, and signal attenuation. Consequently, methods that heavily depend on direct messaging and centralized information sharing may have limited applicability in disaster scenarios.

In this context, stigmergy offers an important alternative coordination mechanism for multi-robot search systems [5]. Stigmergy refers to indirect communication among individuals through traces left in the environment. This mechanism is commonly observed in nature, particularly in ant and termite colonies. Individuals deposit pheromone-like traces in the environment, while the behaviors of other individuals are influenced by the intensity and distribution of these traces. Theraulaz and Bonabeau described stigmergy as one of the fundamental mechanisms underlying self-organizing behaviors observed in social insects [15]. This structure enables the emergence of collective behavior without centralized control. From a robotics perspective, stigmergy is particularly attractive because it reduces communication overhead while allowing the environment itself to function as a form of shared memory.

The stigmergic approach also constitutes the foundation of the Ant Colony Optimization algorithm. Systematically formalized by Dorigo and colleagues, Ant Colony Optimization (ACO) is a swarm intelligence method in which artificial ants reinforce promising solutions over time by depositing pheromone trails throughout the solution space [13]. In this algorithm, pheromone accumulation increases the likelihood of reselecting previously successful paths, whereas pheromone evaporation reduces the influence of outdated or ineffective information. The combined use of these two mechanisms helps establish a balance between exploration and exploitation. These characteristics have encouraged the application of stigmergy-based approaches to robotic problems such as target search, path planning, and area exploration.

In robotics literature, stigmergy is generally modeled through the concept of virtual pheromones rather than physical chemical substances. The "pheromone robotics" study conducted by Payton and colleagues is considered a pioneering work demonstrating that robot swarms can achieve coordination through virtual pheromone messages without centralized control [16]. In this approach, pheromones are represented not as chemical materials, but as numerical information maintained either by robots or within an environmental representation. The concept of virtual pheromones enables indirect coordination among robots in tasks such as surveillance, exploration, hazard detection, and path finding. Therefore, it provides a suitable modeling framework for simulation-based studies and digital map-based robotic applications that do not require physical pheromone emission.

In stigmergy-based robotic systems, pheromones can be defined as either attractive or repulsive. Attractive pheromones encourage robots to revisit successful paths or regions close to targets, whereas repulsive pheromones can discourage movement toward previously visited or potentially unproductive areas. Fossum and colleagues investigated the use of repulsive pheromones for efficient swarm robotic search in unknown environments and demonstrated that this approach can help reduce unnecessary revisits [17]. Similarly, Hamann and Wörn showed that swarm robotic search behaviors based on virtual pheromones can be modeled analytically and spatially [18]. These studies indicate that pheromone-based indirect communication is not merely a biologically inspired analogy, but also a computational mechanism that can be utilized to model and regulate the collective behaviors of robot swarms[16-19].

Nevertheless, the literature also demonstrates that stigmergy-based approaches do not always guarantee superior performance under all conditions. Hunt and colleagues examined the limitations of pheromone-based stigmergy in high-density robot swarms and reported that, under certain conditions, simple stigmergic avoidance behaviors may lose their advantage compared to random walk strategies [20]. This finding suggests that parameters such as robot density, environment size, obstacle distribution, pheromone evaporation rate, and target distribution must be carefully considered in stigmergy-based systems. Therefore, stigmergy should not be regarded as a standalone solution, but rather as a coordination mechanism that must be integrated with appropriate problem formulations, carefully selected parameters, and additional decision-making components.

From a search and rescue perspective, one of the most significant advantages of stigmergy-based approaches is the reduction of direct communication requirements. In rubble environments, wireless communication may become unreliable, and maintaining continuous connectivity among robots can be difficult. Under such conditions, decision-making mechanisms based on environmental traces or virtual environmental memory may provide a more robust coordination framework. Tang and colleagues proposed a stigmergy-based strategy for dynamic target search and tracking using swarm robots and demonstrated that a vector-based pheromone model can effectively support search and tracking behaviors [21]. Such studies reveal that stigmergy is not limited to static target search problems, but can also function as a flexible coordination mechanism under dynamic targets and changing environmental conditions.

A review of the existing literature reveals three major trends in search and rescue and multi-robot exploration research. The first trend is the transition from teleoperated and semi-autonomous systems toward more autonomous multi-robot structures [2], [9-10], [22-23]. The second trend is the shift from coordination mechanisms based on direct communication toward environment-mediated indirect communication approaches [4-5], [15]. The third trend involves extending simple pheromone-based guidance mechanisms with additional decision-making components such as visitation history, robot density, target signals, and environmental costs [17], [20-21], [24]. These trends highlight the increasing importance of multi-component decision mechanisms for complex environments characterized by uncertainty and high obstacle density, such as post-earthquake rubble search scenarios.

The simulation-based approach presented in this study can be positioned within this line of research. In the proposed method, robot movement decisions are not based solely on random selection or a single pheromone component. Instead, pheromone intensity, target-generated indirect signals, robot density, and cell visitation history are jointly considered within the decision-making process. This structure extends conventional pheromone-based search behavior by incorporating a more balanced and adaptive exploration mechanism. Furthermore, the study evaluates not only the performance of the stigmergy-based approach, but also compares it with a version of the same decision-making mechanism in which stigmergic components are disabled, as well as with a fully random search strategy. This comparative framework enables the contribution of stigmergic indirect communication to be analyzed separately against visitation-history-based exploration behavior and purely random movement strategies.

While many studies in the literature evaluate success primarily in terms of target acquisition or area coverage, the present study jointly considers target discovery speed, cost per target, and route length per target. This choice is particularly important in the context of search and rescue operations, since practical effectiveness depends not only on whether a target is found, but also on how quickly, efficiently, and with how little movement cost the target can be reached. In addition, the use of different obstacle densities allows the proposed method to be evaluated not only in relatively simple environments, but also in terms of its robustness under increasing environmental complexity. From this perspective, the present study extends existing stigmergy-based multi-robot search approaches within the specific context of post-earthquake rubble search problems at the simulation level. The primary contribution of the study is the investigation of a communication-independent, multi-component, probabilistic decision-making mechanism under different obstacle densities, together with a comparative evaluation of its performance against both random search and a non-stigmergic exploration strategy. A preliminary version of the simulation framework used in this study was previously explored in a master's thesis focusing on an ant-system-based control approach for rubble search problems [6]. In that earlier work, a basic stigmergy-based search structure and a limited set of environmental decision components were investigated. The present study substantially redesigns and extends this framework through the introduction of the proposed Multi-Component Stigmergic Search (MCSS) algorithm, additional environmental decision parameters, obstacle-density-based evaluations, expanded performance metrics, and a more comprehensive comparative analysis framework.

3. Multi-Component Stigmergic Search Algorithm

3.1 Problem Definition and Environment

Studies in the literature have demonstrated that stigmergy-based approaches can produce effective results in multi-robot exploration and search problems. However, in scenarios involving high obstacle density and severe mobility constraints, such as post-earthquake rubble environments, simple pheromone-based guidance mechanisms alone may not be sufficient. In particular, jointly considering additional components such as robot density, visitation history, and target-related environmental signals may contribute to the emergence of more balanced and efficient search behaviors. For this reason, this study investigates a simulation-based multi-robot search approach that extends stigmergy-based indirect communication with a multi-component decision-making structure. This approach will hereafter be referred to as the **Multi-Component Stigmergic**

Search (MCSS) algorithm. The simulation environment employed in this study was developed based on an earlier prototype framework introduced in a previous master's thesis study [6]. However, the current work significantly extends the original structure through redesigned search dynamics, multi-component probabilistic movement selection, obstacle-density-based scenario generation, and additional performance evaluation criteria including search cost and route efficiency.

The problem addressed in this study is a multi-robot target search problem representing search and rescue operations conducted in post-earthquake rubble environments. The problem is defined as the rapid detection of targets (victims) located within a structurally complex environment containing obstacles, using autonomous robots. Within this framework, the robots are expected to exhibit collective search behavior without relying on a centralized control mechanism, instead coordinating through environmental information and indirect communication mechanisms.

The simulation environment is modeled as a two-dimensional discrete structure. The search area is represented using a grid-based environment composed of equally sized cells. This representation was selected to abstract the irregular, fragmented, and mobility-constrained characteristics of post-earthquake rubble fields. Each cell may contain free space, an obstacle, or a target. An example simulation environment is illustrated in Fig. 1. The figure presents the irregular obstacle distribution within the grid-based search area, together with the initial positions of the robots and the spatial placement of the targets. This structure visually demonstrates both the mobility constraints encountered by the robots and the complexity of the operational environment.

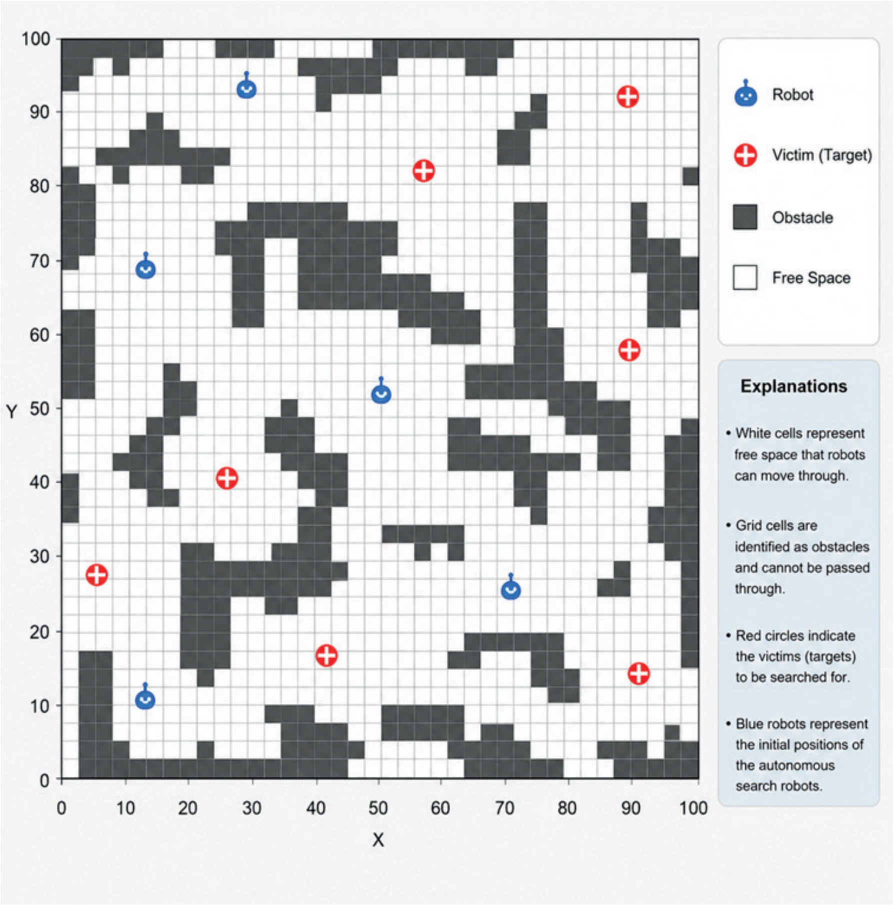

Fig 1. Example simulation environment containing obstacles. In the grid-based structure, white cells represent free spaces through which robots can move, while gray cells represent obstacles that cannot be traversed. Blue markers indicate the current positions of the robots, whereas red markers denote the locations of the search targets (victims).

Obstacles within the environment are modeled as non-traversable cells that restrict robot movement. These obstacles represent physical constraints such as rubble fragments, collapsed structural components, and narrow passageways commonly encountered in post-earthquake environments. The obstacle distribution was designed to increase environmental complexity, thereby preventing robots from moving freely and directly across the search area. This configuration enables the search problem to incorporate more realistic operational challenges.

Multiple targets are distributed throughout the simulation environment at different spatial locations. Rather than being directly observable by the robots, these targets are represented through indirect signals that can be perceived within a limited influence range. This approach abstracts real-world search and rescue scenarios in which victims trapped beneath rubble cannot be directly seen, but may instead be detected through indirect indicators such as sound, movement, or similar signals. The robots operating within the environment are initially positioned at different starting locations. Robots move over discrete time steps and, at each step, are allowed to transition only between neighboring cells. Movement decisions are determined using local environmental information rather than a centralized control mechanism. Robots are not permitted to enter cells containing obstacles. This framework enables the investigation of both multi-robot coordination behavior and the effects of the stigmergy-based indirect communication mechanism employed in the MCSS approach under complex search conditions.

In the MCSS approach, robots leave virtual traces within the environment. Other robots perceive these traces and use them to shape their movement decisions. The influence of these environmental traces gradually decreases over time, thereby creating a dynamic information structure within the environment. This mechanism provides an important advantage in disaster scenarios where direct communication is unavailable or severely limited. The simulation process is repeated multiple times under different initial conditions. These repetitions help reduce the influence of randomness and enable more reliable performance evaluations. The defined problem structure reflects the primary challenges encountered in post-earthquake rubble search and rescue scenarios, including obstacle density, mobility constraints, uncertainty, and multi-target search. This framework therefore enables a comparative evaluation of different search strategies under complex operational conditions.

3.2 Stigmergy-Based Search Strategy

In this study, the **Multi-Component Stigmergic Search (MCSS)** approach is employed, in which robots navigate by utilizing environmental traces and indirect signals associated with targets. Within the MCSS framework, robots do not move purely randomly. Instead, the decision-making process incorporates both environmental traces and indirect target-related signals. The proposed approach is inspired by the fundamental principles of Ant Colony Optimization (ACO) [13]. However, the classical structure has been extended and adapted to suit the specific characteristics of the problem domain.

At each time step, every robot evaluates the neighboring cells that are reachable from its current position. Directions containing obstacles are eliminated, and the remaining feasible movement options are identified. For each candidate cell, a preference weight is calculated. This weight is determined by the combination of four main components:

- Pheromone intensity
- Target signal (sound)
- Robot density
- Visitation history

The movement tendency of a robot from position i to a neighboring cell j is expressed as follows:

$$w_{i,j} = \left(1+\tau_{i,j}\right)^{\alpha}\left(1+S_{i,j}\right)^{\beta}\left(1+R_{i,j}\right)^{\gamma}\cdot\frac{1}{1+V_{i,j}} \tag{1}$$

In the above formulation:

$\tau_{i,j}$: pheromone intensity associated with the corresponding cell

$S_{i,j}$: intensity of the indirect signal originating from targets

$R_{i,j}$: term representing robot density

$V_{i,j}$: number of previous visits to the corresponding cell

α , β , γ : coefficients determining the influence of the respective components

This structure ensures that robots do not rely on a single source of information. Instead, multiple environmental factors are jointly evaluated to establish a more balanced decision-making mechanism.

Pheromone intensity enables robots to preferentially follow paths that were previously explored successfully. When a robot reaches a target, it deposits pheromones along the route it followed. This process can be expressed as follows:

$$\tau_{i,j} \leftarrow \left(1-\rho\right)\tau_{i,j} + \Delta\tau_{i,j} \tag{2}$$

Here, ρ represents the pheromone evaporation rate. This parameter controls the gradual reduction of pheromone intensity over time and prevents outdated information from becoming dominant within the system. The sound generated by targets is not directly observable. Instead, it is represented through a signal that can be perceived within a limited influence range. This signal guides robots toward potential target locations and supports the search process under conditions where direct target observation is not possible:

$$S_{i,j} = \frac{1}{d_{i,j} + \epsilon} \tag{3}$$

Here, $d_{i,j}$ denotes the distance between the corresponding cell and the target. The parameter ϵ is a small positive constant introduced to ensure numerical stability and to prevent division-by-zero conditions. Accordingly, the perceived signal intensity increases as the robot approaches the target. As a result, cells located closer to the target are assigned higher preference values.

The presence of a large number of robots within the same region may reduce search efficiency due to overcrowding and redundant exploration. For this reason, robot density is incorporated into the decision-making mechanism. Robot density is defined as follows:

$$R_{i,j} = \frac{N_{i,j}}{N_{\max}} \tag{4}$$

Here, $N_{i,j}$ represents the number of robots located within the corresponding cell or its surrounding neighborhood. The parameter $N_{\max}$ denotes the maximum robot density used for normalization purposes. This term reduces excessive clustering of robots within the same region and encourages a more balanced distribution across the search area.

The tendency of robots to repeatedly revisit frequently explored cells is intentionally reduced. For this purpose, a penalty term based on visitation history is incorporated into the preference weight. This effect is modeled as follows:

$$F_{i,j} = \frac{1}{1 + V_{i,j}} \tag{5}$$

Here, $V_{i,j}$ denotes the number of times the corresponding cell has previously been visited. As the visitation count increases, the value of this term decreases. Consequently, the probability of selecting frequently visited cells becomes lower. This mechanism limits unnecessary repeated exploration within the same regions and encourages the discovery of previously unexplored areas.

The movement decision is determined probabilistically according to the transition tendency values associated with traversable neighboring cells. The preference weights calculated for each neighboring cell are normalized to obtain a transition probability:

$$P_{i,j} = \frac{w_{i,j}}{\sum_{k \in \mathcal{N}_{\square}} w_{i,k}} \tag{6}$$

Here, $\mathcal{N}_{\square}$ denotes the set of neighboring cells that are reachable from the robot's current position. Through this structure, robots are more likely to select directions associated with higher weight values. However, the probability of selecting other feasible directions is not completely eliminated. As a result, the system maintains a balance between directed search behavior and exploratory behavior. The pseudocode of the proposed MCSS algorithm is presented in **Algorithm 1**.

Algorithm 1. Multi-Component Stigmergic Search (MCSS) Algorithm

1. **For each robot, determine the neighboring cells that are reachable from the current position.**
2. **Remove obstacle-containing or non-traversable cells from the set of candidate movements.**
3. **For each candidate cell, calculate the environmental attractiveness value by considering:**
    - **the pheromone intensity in the corresponding direction,**
    - **the influence of target-generated signals,**
    - **the robot density within the same region,**
    - **the reduced selection probability of previously visited cells.**
4. **Convert the attractiveness values of the candidate cells into transition probabilities.**
5. **Select the robot's next movement according to these probabilities.**
6. **Move the robot to the selected neighboring cell.**
7. **Update the robot's visitation information and traversed path.**
8. **If the robot reaches a target:**
    8.1 Mark the target as found.
    8.2 Retrieve the route followed by the robot while reaching the target.

8.3 Deposit pheromones along the transitions on this route.
8.4 Return the robot to its initial position.

9. **Apply pheromone evaporation at each time step by reducing pheromone values.**
10. **Repeat the process until all targets are found or the search procedure terminates.**

When a robot reaches a cell containing a target, the corresponding target is marked as found. The value associated with the target cell is then reset to zero, preventing the same target from being counted again in subsequent time steps. Within the MCSS approach, pheromone traces are reinforced along the route followed by the robot while reaching the target. This process enables successful search routes to be propagated throughout the environment. After reaching a target, the robot is returned to its initial position and reintroduced into the search process. This structure allows robots to continue operating continuously in environments containing multiple targets.

The random search strategy illustrates the behavior of the multi-robot system in the absence of environmental information and indirect coordination mechanisms. In contrast, the MCSS approach investigates how robots can exhibit more directed search behavior by utilizing environmental traces and indirect target-related signals. By comparing these two approaches, the effects of the stigmergic coordination mechanism on target discovery, search cost, and robot movement efficiency can be evaluated. In this way, the study examines how the stigmergy-based indirect communication mechanism discussed in the literature can contribute to search and rescue problems within a representative simulation environment.

3.3 Experimental Setup

To evaluate the performance of the MCSS approach, simulation experiments based on a multi-robot exploration problem were conducted. The MCSS approach was compared against two alternative strategies: (i) an exploration approach without a stigmergic coordination mechanism and (ii) a fully random search strategy. The experiments were performed in a two-dimensional discrete environment consisting of a $100\times100100\times100$ grid structure. Within this environment:

- 10 targets were randomly distributed throughout the search area,
- 12 robots were initially positioned at random locations,

- different numbers of obstacles were introduced to control environmental complexity.

At the beginning of each experiment, robots, targets, and obstacles were randomly distributed without overlap. At each time step, robots were allowed to move to one of the eight neighboring cells. Every movement decision was subject to obstacle checking, and the next position was selected only from feasible directions. Movements toward obstacle-containing cells were not permitted. If a robot had no valid movement option, it remained stationary for that time step.

When a robot reached a cell containing a target, the route followed by the robot to reach that target was recorded, and the robot was returned to its initial position. In the MCSS approach, robots deposit pheromone traces into the environment during this return process. As described in Section 3.2, the selection value for each possible movement is calculated probabilistically based on pheromone intensity, target-generated signals, robot density, and the visitation count of the corresponding cell. The resulting selection probabilities are normalized, and the next movement is determined through probabilistic sampling.

As previously noted, whenever a robot discovers a target, pheromone traces are reinforced along the traversed route. In addition, pheromone intensity is gradually reduced over time through an evaporation mechanism applied at each time step. The parameter values used in all experiments are presented in Table 1.

Table 1. Simulation Parameters

Parameter	**Value**
Environment size	100 x 100 cells
Number of robots	12
Number of targets	10
Number of obstacles	250 / 500 / 750
Maximum time steps	10,000
Number of experiment repetitions	100
Pheromone evaporation rate	0.02
σ (pheromone influence)	1
γ (target signal influence)	1
β (robot density influence)	1

The parameter values presented in Table 1 were selected to provide a controlled and computationally feasible evaluation environment while maintaining comparable influence among the decision components. The environment size (100 × 100 cells), number of robots (12), and number of targets (10) were chosen to create a sufficiently large search space that allows collective behaviors to emerge without introducing excessive computational cost. Obstacle densities of 250, 500, and 750 cells were used to represent progressively increasing environmental complexity and mobility constraints. The pheromone evaporation rate was set to 0.02 to preserve environmental memory while preventing outdated information from dominating the search process over long simulation periods. The influence coefficients of pheromone intensity ($\sigma = 1$), target-generated signals ($\gamma = 1$), and robot density ($\beta = 1$) were initialized with equal values to avoid introducing a priori preference toward any individual decision component and to allow the combined effect of the multi-component mechanism to be evaluated more transparently. These parameter values were determined through preliminary simulation trials aimed at obtaining stable search behavior and ensuring meaningful comparisons across different search strategies rather than performing exhaustive parameter optimization.

In this study, three different search approaches were compared. The first approach, MCSS, was described in detail in Section 3.2. The second approach, used for comparative purposes, was derived from the same decision-making framework with the stigmergic components disabled. In this configuration, the overall algorithmic structure was preserved, while the pheromone influence (σ), target-generated signal influence (γ), and robot density influence (β) parameters used during movement selection were set to zero. As a result, environmental traces and target-oriented guidance were removed, while the visitation-history component of the decision mechanism was retained. Consequently, robots exhibited an exploration-oriented behavior by preferentially moving toward less frequently visited cells. The third approach was a fully random search strategy. In this strategy, robots selected among feasible movements with equal probability and did not utilize any memory or guidance mechanism during the search process.

When these three approaches are evaluated together, the contribution of the stigmergic coordination mechanism can be analyzed comparatively against both visitation-history-based exploration behavior and a fully random movement strategy. The primary performance metrics considered in the experiments were the average number of targets found over time, robot movement cost per target, and route length per target. All results were obtained by averaging 100 independent simulation runs and are presented in the figures together with

95% confidence intervals. Although a dedicated component-wise ablation study was not originally designed, the comparative evaluation framework adopted in this study can also be interpreted as a functional ablation analysis of the proposed MCSS architecture. Specifically, comparison between the full MCSS configuration and the exploration strategy with disabled stigmergic components enables the isolated contribution of stigmergy-based indirect coordination to be evaluated while preserving the remaining exploration mechanism. Furthermore, comparison with the fully random search strategy provides an additional lower-bound reference in which all environmental guidance mechanisms are removed. This structure allows the incremental contribution of the proposed coordination mechanism to be analyzed without requiring a separate ablation protocol.

To evaluate whether the observed performance differences among the search strategies were statistically meaningful, additional statistical analyses were performed using the results obtained from 100 independent simulation runs for each experimental configuration. Prior to statistical comparison, normality assumptions were evaluated for the collected performance metrics and confirmed to be satisfied. Accordingly, pairwise comparisons among the search approaches were conducted using independent-samples t-tests. Since multiple pairwise comparisons were performed among the three evaluated approaches, p-values were adjusted using the Holm–Bonferroni correction procedure to control the family-wise error rate. Statistical significance was evaluated at $\alpha = 0.05$. The statistical analysis was conducted separately for each obstacle-density scenario and for each performance metric considered in the study.

4. Results

In this study, the MCSS approach was compared with the exploration approach in which stigmergic components were disabled and with the fully random search strategy using three different performance metrics: (i) the number of targets found over time, (ii) cost per target, and (iii) average route length per target. The experiments were conducted in three different environments containing 250, 500, and 750 obstacles, respectively, and each configuration was evaluated over 100 independent simulation runs. In all figures, average values are presented together with 95% confidence intervals.

4.1 Average Number of Targets Found Over Time

The results presented in Fig. 2 demonstrate that the three approaches exhibit clearly different exploration speeds and convergence behaviors. Across all obstacle densities, the MCSS approach achieved the fastest overall performance by a substantial margin. This approach was able to discover the majority of targets during the early stages of the simulation and reached the maximum number of targets within approximately 1000–1500 time steps.

In contrast, the exploration approach with disabled stigmergic components exhibited a slower but more gradual and stable increase in performance. Although this approach performed significantly better than the fully random search strategy, it was unable to achieve the early exploration advantage demonstrated by the MCSS approach. The random search strategy produced the lowest performance overall, requiring considerably longer times to reach targets and remaining notably behind the other approaches, particularly during the early stages of the simulations.

As the number of obstacles increased, a decrease in performance was observed for all approaches, as expected. However, the MCSS approach was the least affected by the increase in environmental complexity. Notably, even in the scenario containing 750 obstacles, the approach maintained its rapid early convergence behavior, indicating strong capabilities in navigation and indirect information sharing within complex environments. Furthermore, an examination of the confidence intervals reveals that the MCSS approach exhibited lower variance compared to the alternative methods. This finding suggests that the proposed approach provides not only fast but also stable and consistent performance across different simulation runs.

To further evaluate whether the observed differences were statistically meaningful, statistical comparisons were performed using pairwise independent-samples t-tests with Holm–Bonferroni correction. The results supported the visual trends presented in Fig. 2. Compared with the non-stigmergic exploration approach, the MCSS approach achieved statistically significant improvements in the average number of targets found during approximately the first 6000 time units (adjusted $p < 0.05$). After this stage, the performance difference gradually decreased as both approaches approached convergence. In contrast, the MCSS approach maintained statistically significant superiority over the random search strategy throughout the entire simulation period (adjusted $p < 0.05$), indicating that stigmergic coordination contributed substantially to faster and more effective exploration.

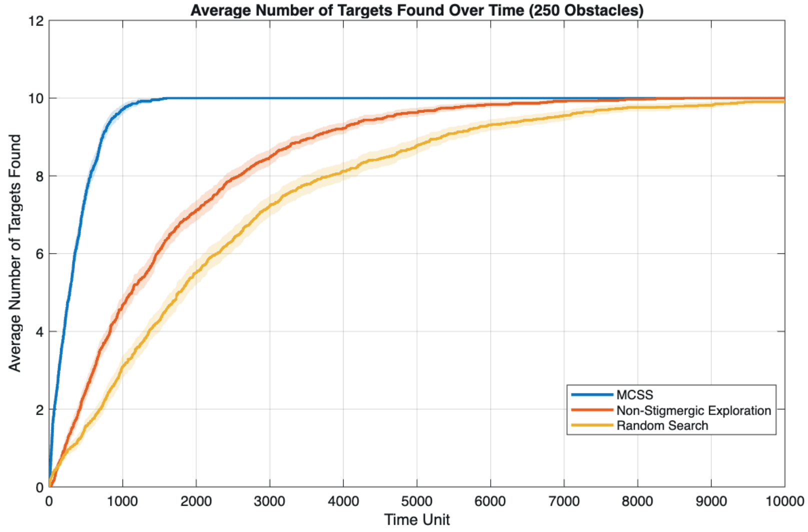
Average Number of Targets Found Over Time (250 Obstacles)
Average Number of Targets Found
Time Unit
MCSS
Non-Stigmergic Exploration
Random Search

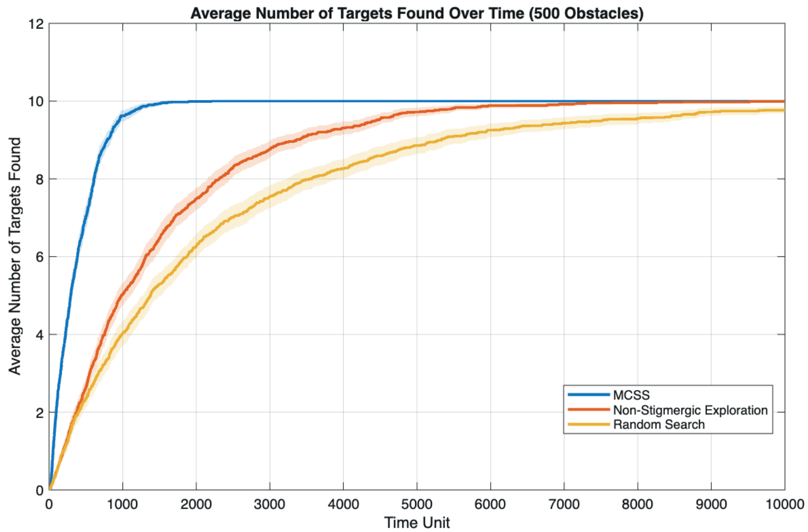
Average Number of Targets Found Over Time (500 Obstacles)
Average Number of Targets Found
Time Unit
MCSS
Non-Stigmergic Exploration
Random Search

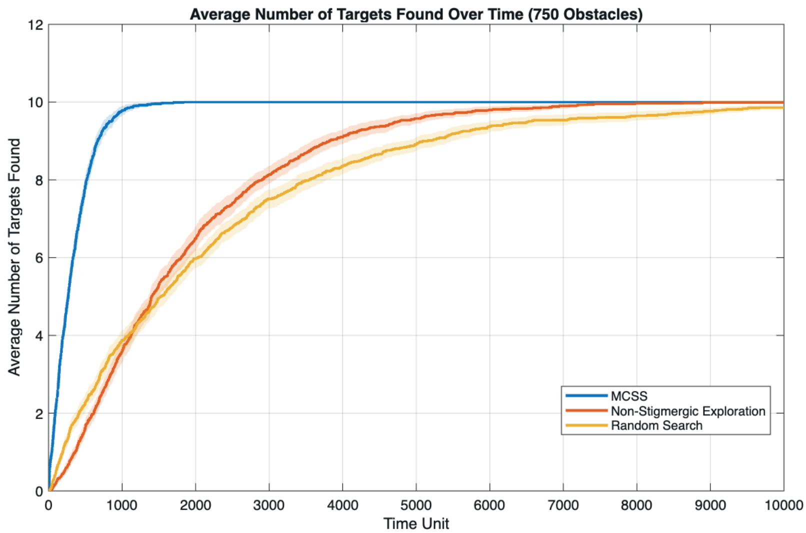

Fig. 2. Comparison of the average number of targets found over time for the three search methods under different obstacle densities (250, 500, and 750 obstacles). The curves represent the average values obtained from 100 independent simulation runs, while the shaded regions indicate the 95% confidence intervals. The Multi-Component Stigmergic Search (MCSS) approach exhibits faster convergence and earlier target discovery across all scenarios, the non-stigmergic exploration method demonstrates intermediate performance, and the random search strategy produces the lowest exploration speed.

4.2 Cost per Target

The boxplots presented in Fig. 3 clearly reveal the differences among the approaches in terms of cost per target. Across all experimental environments, the MCSS approach achieved the lowest cost values. The combination of low median values and narrow distributions indicates that this approach provides both efficient and consistent performance. The exploration approach with disabled stigmergic components demonstrates a noticeable improvement compared to the random search strategy. By prioritizing less frequently visited cells, this approach increases exploration efficiency and thereby reduces search cost. However, because it lacks the indirect information-sharing mechanism provided by stigmergy, its cost values remain higher than those of the MCSS approach, particularly in complex environments. The random search strategy exhibits the highest cost values overall. In this approach, the routes followed to reach targets are largely inefficient, and unnecessary revisits occur frequently. Examination of the boxplots further shows that this strategy has a considerably wider distribution and contains a large number of outliers. This observation indicates that random search produces highly unpredictable and unstable performance. As the number of obstacles increases, cost values rise for all approaches. However, the increase remains relatively limited in the MCSS approach, whereas it becomes substantially more pronounced in the alternative methods. These results suggest that the MCSS approach offers a more scalable structure for operation in complex environments.

To assess whether the observed differences in search efficiency were statistically meaningful, pairwise comparisons were conducted using independent-samples t-tests with Holm–Bonferroni correction. The statistical analysis confirmed that the MCSS approach achieved significantly lower cost-per-target values than both the non-stigmergic exploration approach and the random search strategy across all evaluated obstacle-density scenarios (adjusted $p < 0.05$). In contrast, the difference between the non-stigmergic exploration approach and random search gradually decreased as environmental complexity increased and became statistically non-significant under higher obstacle-density conditions. This finding suggests that the visitation-history-based exploration mechanism alone provides limited improvements in search efficiency and that the indirect coordination enabled by stigmergic information becomes increasingly important in more constrained environments.

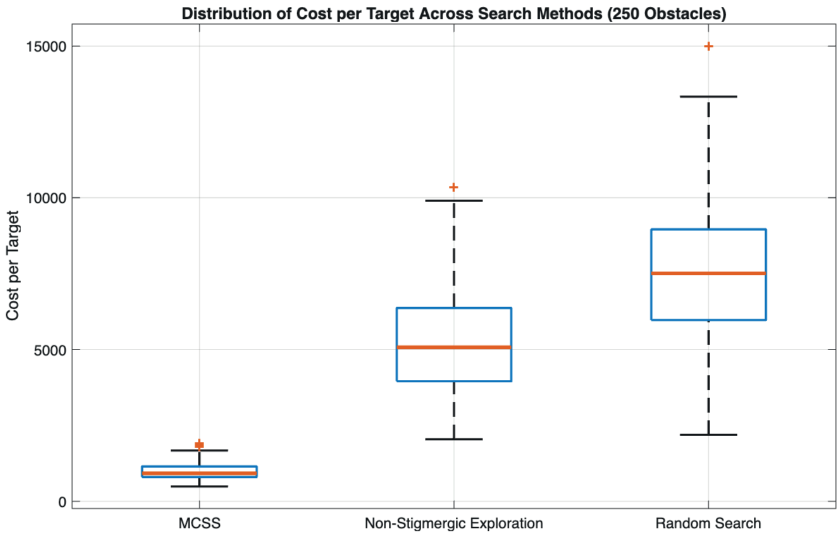
Distribution of Cost per Target Across Search Methods (250 Obstacles)
Cost per Target
15000
10000
5000
0
MCSS
Non-Stigmergic Exploration
Random Search

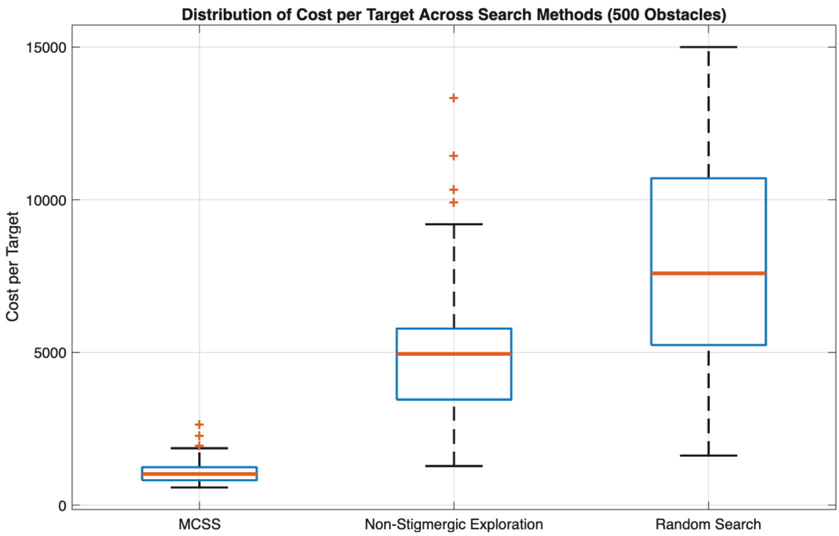
Distribution of Cost per Target Across Search Methods (500 Obstacles)
Cost per Target
15000
10000
5000
0
MCSS
Non-Stigmergic Exploration
Random Search

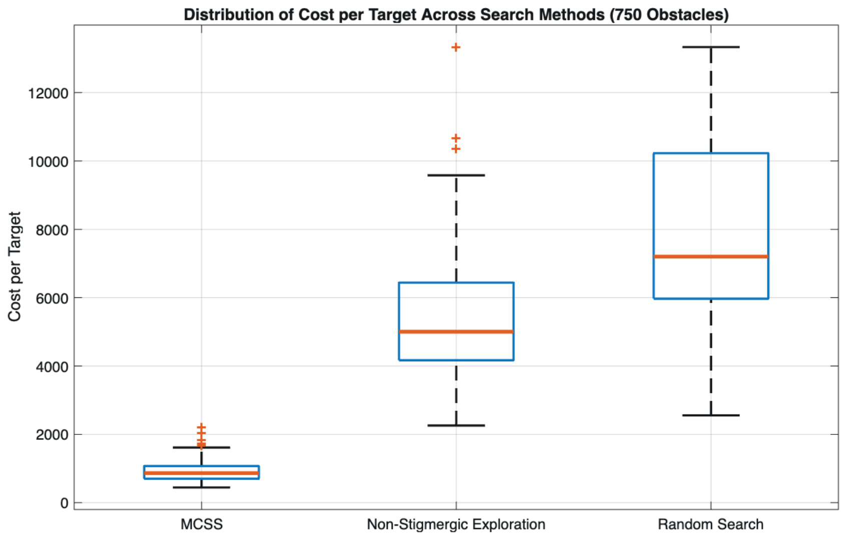

Fig. 3. Comparison of the cost-per-target distributions for the three search approaches under different obstacle densities (250, 500, and 750 obstacles). The boxplots illustrate the median values, interquartile ranges, and outliers obtained from 100 independent simulation runs. The Multi-Component Stigmergic Search (MCSS) approach achieves the lowest cost and the narrowest distribution across all scenarios, indicating the most efficient and stable performance. The exploration approach with disabled stigmergic components produces lower costs than the random search strategy, but still remains behind the MCSS approach. The random search strategy exhibits the highest costs and the widest variance, demonstrating the lowest overall efficiency.

4.3 Average Route Length per Target

The results presented in Fig. 4 demonstrate the effectiveness of the routes followed by the approaches in reaching targets. The findings clearly show that the MCSS approach produces the shortest route lengths among all methods. Through the use of environmental traces, robots in the MCSS framework avoid repeatedly exploring previously visited regions and instead follow more direct and efficient paths. This behavior reduces unnecessary wandering and minimizes route lengths. Examination of the boxplots further indicates that the MCSS approach exhibits both low median values and a narrow distribution.

The exploration approach with disabled stigmergic components performs better than the random search strategy in terms of route length. The strategy of favoring less frequently visited cells provides a certain degree of exploration efficiency and contributes to shorter routes. However, because this approach does not benefit from the collective information-sharing mechanism provided by stigmergy, its route optimization capability remains limited compared to the MCSS approach.

The random search strategy generates the longest route lengths overall. In this approach, the paths followed to reach targets are largely indirect and contain frequent repetitive movements. Moreover, the wide distribution observed in the boxplots indicates unstable and inconsistent behavior. As the number of obstacles increases, route lengths rise for all approaches. Nevertheless, the MCSS approach is the method least affected by this increase. This finding suggests that the proposed approach maintains effective path-planning capability even in environments with high obstacle density.

To evaluate whether the observed differences in route efficiency were statistically meaningful, pairwise comparisons were conducted using independent-samples t-tests with Holm–Bonferroni correction. The statistical analysis confirmed that the MCSS approach achieved significantly shorter average route lengths per target than both the non-stigmergic exploration approach and the random search strategy across all evaluated obstacle-density scenarios (adjusted $p < 0.05$). In contrast, the difference between the non-stigmergic exploration approach and random search progressively diminished as obstacle density increased and became statistically non-significant under more complex environmental conditions. These findings indicate that reducing route inefficiency in constrained search environments depends not only on avoiding previously visited regions but also on the indirect coordination and collective guidance provided by the stigmergic mechanism.

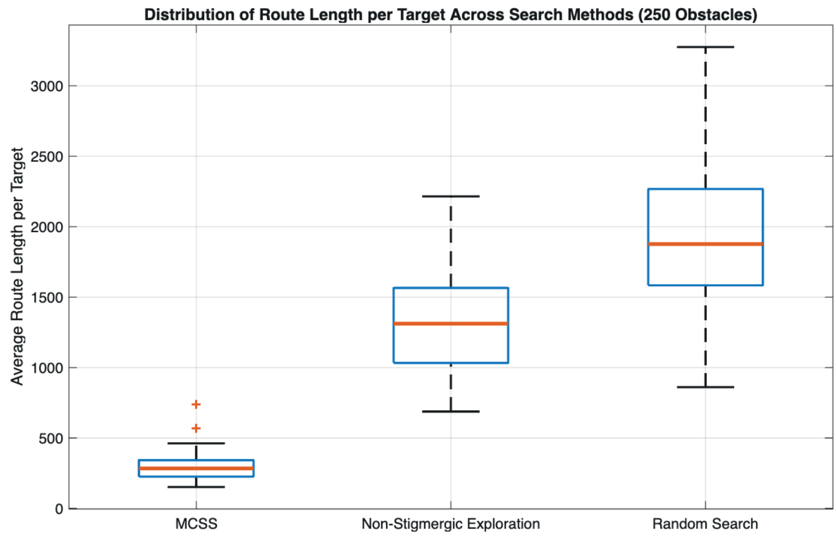
Distribution of Route Length per Target Across Search Methods (250 Obstacles)
Average Route Length per Target
0
500
1000
1500
2000
2500
3000
MCSS
Non-Stigmergic Exploration
Random Search

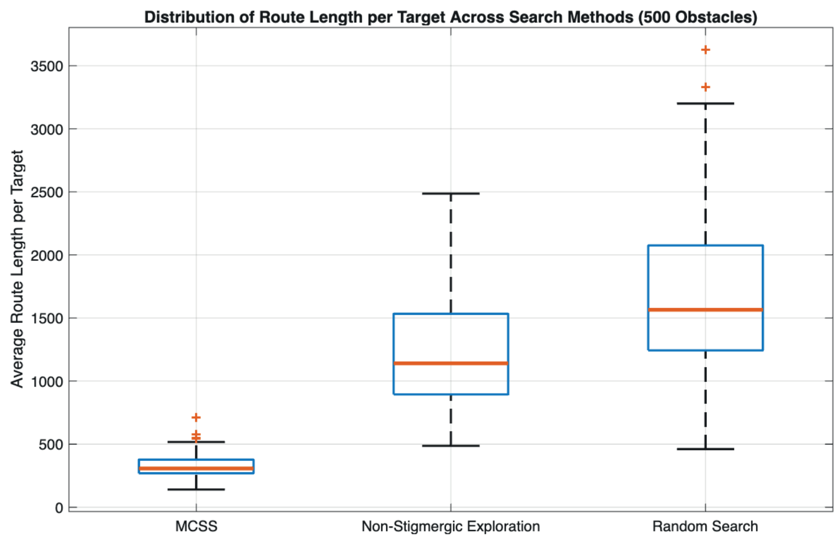
Distribution of Route Length per Target Across Search Methods (500 Obstacles)
Average Route Length per Target
0
500
1000
1500
2000
2500
3000
3500
MCSS
Non-Stigmergic Exploration
Random Search

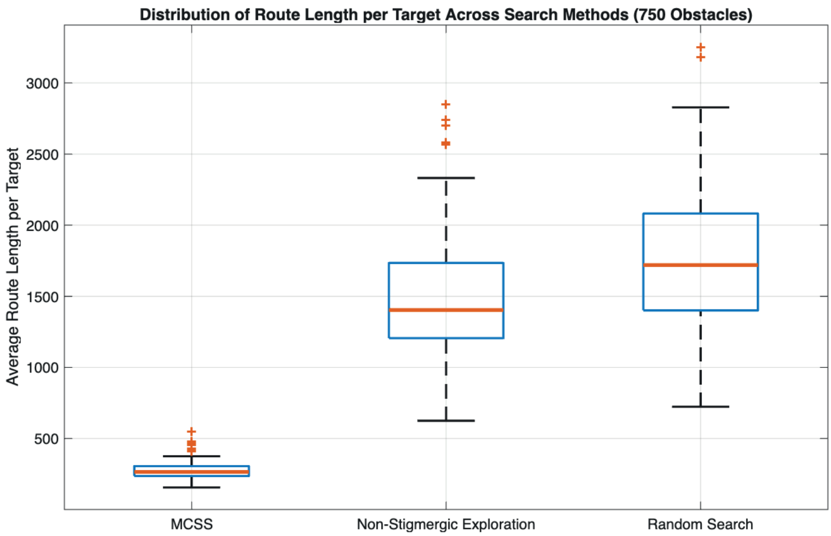

Fig. 4. Comparison of the average route-length distributions per target for the three search approaches under different obstacle densities (250, 500, and 750 obstacles). The boxplots present the median values, interquartile ranges, and outliers obtained from 100 independent simulation runs. The Multi-Component Stigmergic Search (MCSS) approach achieves the shortest route lengths and the narrowest distributions across all scenarios, demonstrating the most efficient path-planning performance. The exploration approach with disabled stigmergic components produces shorter routes than the random search strategy, but still remains behind the MCSS approach. The random search strategy exhibits the longest and most variable route lengths, resulting in the lowest overall performance.

When all results are evaluated collectively, both the graphical observations and the statistical analyses consistently indicate that the MCSS approach provides clear advantages in three major aspects:

- Faster exploration performance (early convergence)

- Lower search cost

- Shorter and more efficient routes

Pairwise comparisons performed using independent-samples t-tests with Holm–Bonferroni correction confirmed that the MCSS approach achieved statistically significant improvements over random search across all evaluated performance metrics and obstacle-density scenarios. Furthermore, compared with the non-stigmergic exploration approach, the MCSS approach also demonstrated statistically significant advantages in target discovery performance, cost per target, and average route length under the evaluated conditions.

The non-stigmergic exploration approach produced measurable improvements over random search, particularly in terms of exploration efficiency during the earlier stages of the search process. However, the statistical analysis further revealed that these advantages diminished as obstacle density increased and became non-significant for some efficiency-related metrics under more complex environmental conditions. This finding suggests that visitation-history-based guidance alone contributes positively to exploration behavior but remains insufficient to maintain robust search efficiency in constrained environments.

Overall, the combined experimental and statistical findings demonstrate that the MCSS approach offers an effective and scalable solution for multi-robot exploration problems and provides particularly strong advantages in environments with high obstacle density through stigmergy-based indirect coordination.

5. Conclusion

In this study, stigmergy-based multi-robot search approaches developed for post-earthquake search and rescue problems were investigated, and the performance of the proposed Multi-Component Stigmergic Search (MCSS) approach was evaluated in a simulation environment. Within the scope of the study, the effects of extending conventional stigmergy-based guidance mechanisms with additional environmental components such as robot density, visitation history, and target-generated indirect signals were examined in the context of multi-robot exploration behavior. For this purpose, the MCSS approach was compared against an exploration approach with disabled stigmergic components and a fully random search strategy.

The obtained results demonstrate that the MCSS approach provides clear advantages over the alternative methods, particularly in terms of early exploration performance, low movement cost, and short route generation. Analysis of the number of targets found over time showed that the MCSS approach achieved faster convergence and reached targets more rapidly across all obstacle densities. Furthermore, the cost-per-target and route-length results revealed that stigmergic indirect communication through environmental traces not only accelerates exploration but also improves the efficiency of robot movements. The limited performance degradation observed in high-obstacle-density scenarios further indicates that the proposed approach offers a scalable and robust structure for operation in complex environments.

The exploration approach with disabled stigmergic components produced better results than the random search strategy. This finding suggests that the visitation-history-based guidance mechanism contributes positively to the exploration process. However, due to the absence of indirect information sharing through environmental traces, this approach could not fully achieve the collective guidance advantages provided by the MCSS framework. The random search strategy, on the other hand, demonstrated the lowest performance across all evaluation metrics and exhibited highly inefficient behavior, particularly in environments with high obstacle density.

From a practical performance perspective, the proposed MCSS approach demonstrated substantial quantitative gains across different environmental conditions. Relative to the non-stigmergic exploration approach, MCSS reduced the average cost per target by approximately 78–83% and shortened the average route length per target by approximately 73–82%. Compared with random search, the improvements became even more pronounced, corresponding to approximately 87–88% lower search cost and 80–85% shorter average route lengths. These improvements were further supported by statistical analysis, which confirmed statistically significant differences with consistently large effect sizes across the evaluated efficiency-related metrics. Collectively, these findings indicate that the proposed stigmergic coordination mechanism provides not only statistically significant improvements but also practically meaningful gains with large effect sizes, indicating that the observed performance differences are unlikely to be attributable to random variation alone.

Overall, the obtained findings indicate that stigmergy-based indirect communication mechanisms can be effectively utilized in multi-robot search and rescue problems. The fact that the MCSS approach does not require centralized control, is not dependent on direct communication, and supports collective exploration behavior makes it particularly attractive for disaster scenarios in which communication infrastructure is limited or unreliable. In addition, the modular structure of the approach allows different environmental components to be incorporated into the decision-making mechanism, thereby providing a flexible foundation for the development of more advanced multi-robot systems in future studies. From a broader computational intelligence perspective, the findings further illustrate how adaptive collective behavior can emerge through distributed information processing and environment-driven decision mechanisms, supporting the development of scalable intelligent systems for complex search environments.

Nevertheless, the study has several limitations. The simulation environment was modeled as a two-dimensional discrete structure and therefore does not fully capture the physical complexity of real-world disaster environments. Robot movements were evaluated under idealized assumptions, while factors such as sensor errors, mechanical failures, communication delays, and dynamic environmental changes were not included in the model. Furthermore, the virtual stigmergy mechanism employed in this study does not directly incorporate perception and environmental interaction challenges that may arise in real physical environments.

In addition, several practical constraints commonly encountered in real-world robotic search operations were not explicitly modeled in the present study. The proposed framework assumes reliable localization, synchronized robot operation, and ideal environmental state updates, whereas real deployments may involve positioning uncertainty, asynchronous robot behavior, limited onboard computation, sensor noise, actuator inaccuracies, and intermittent communication availability. Furthermore, the virtual stigmergic representation used in the simulation assumes consistent environmental memory and instantaneous information propagation, while real implementations may require distributed map synchronization and introduce additional latency and uncertainty. These factors may influence both exploration efficiency and coordination quality and should therefore be incorporated into future validation studies involving more realistic robotic platforms and physical environments.

Future work should therefore focus on evaluating the proposed approach on real robotic platforms. In particular, experiments involving physical robot swarms would enable the practical applicability of the MCSS approach under real-world conditions to be investigated more comprehensively. In addition, extending the study to three-dimensional environment models could facilitate the analysis of more realistic search scenarios involving multilayer rubble structures and varying elevation levels. Incorporating factors such as dynamic obstacles, moving targets, sensor uncertainty, and energy consumption into the model would also allow the behavior of the approach in realistic disaster environments to be examined in greater detail.

Another important future research direction involves investigating the performance of the MCSS approach in heterogeneous robot swarm systems. Understanding how robots with different sensing capabilities, mobility characteristics, or task responsibilities can be coordinated within the same stigmergic framework represents a significant research challenge. Moreover, integrating stigmergy mechanisms with adaptive parameter update methods, learning-based movement strategies, and reinforcement learning techniques may contribute to the development of more advanced and adaptive multi-robot search systems.

Compared to the earlier thesis-based prototype framework [19], the proposed MCSS approach introduces a substantially expanded decision-making structure and evaluation methodology. In particular, the integration of multiple environmental guidance components, comparative analysis under varying obstacle densities, and the inclusion of cost- and route-based performance metrics provide a more comprehensive assessment of stigmergy-based search behavior in complex rubble environments.

In conclusion, this study demonstrates the potential of stigmergy-based multi-robot search approaches for post-earthquake search and rescue problems within a simulation environment and shows that the proposed Multi-Component Stigmergic Search approach can provide an effective solution in complex search environments. The findings obtained in this work are expected to provide a useful foundation for the future development of autonomous multi-robot search and rescue systems.

References

[1] Liu, Y., & Nejat, G. (2013). Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems, 72*, 147–165.

[2] Drew, D. S. (2021). Multi-agent systems for search and rescue applications. *Current Robotics Reports, 2*, 189–200.

[3] Sinha, S., Lee, S., & Singh, S. (2026). Survey on Reconnaissance Autonomous Robotic Systems for Disaster Management. *Sensors (Basel, Switzerland)*, *26*(5), 1659.

[4] Yan, Z., Jouandeau, N., & Cherif, A. A. (2013). A survey and analysis of multi-robot coordination. *International Journal of Advanced Robotic Systems, 10*(12), 399.

[5] Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life, 5*(2), 97–116.

[6] Karabulut, Z. (2025). *Application of a Control Algorithm Developed Based on the Ant System to the Debris Search Problem Together with a Stigmergy-Based Multi-Robot System*. M.Sc. Thesis, Bolu Abant Izzet Baysal University.

[7] Casper, J., & Murphy, R. R. (2003). Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B, 33*(3), 367–385.

[8] Moniruzzaman, M. D., Rassau, A., Chai, D., & Islam, S. M. S. (2022). *Teleoperation methods and enhancement techniques for mobile robots: A comprehensive survey*. Robotics and Autonomous Systems, 150, 103973.

[9] Finzi, A., & Orlandini, A. (2005). Human-robot interaction through mixed-initiative planning for rescue and search rovers. In ***AI***IA 2005: Advances in Artificial Intelligence* (pp. 483–494). Springer.

[10] Tang, H., Cao, X., Song, A., Guo, Y., & Bao, J. (2009). Human-robot collaborative teleoperation system for semi-autonomous reconnaissance robot. In *2009 International Conference on Mechatronics and Automation* (pp. 1934–1939). IEEE.

[11] Couceiro, M. S., Vargas, P. A., Rocha, R. P., & Ferreira, N. M. F. (2014). Benchmark of swarm robotics distributed techniques in a search task. *Robotics and Autonomous Systems, 62*(2), 200–213.

[12] Kumar, A. S., Manikutty, G., Bhavani, R. R., & Couceiro, M. S. (2017). Search and rescue operations using robotic Darwinian particle swarm optimization. *2017 International Conference on Advances in Computing, Communications and Informatics*, 1839–1843.

[13] Dorigo, M., Birattari, M., & Stützle, T. (2006). Ant colony optimization. *IEEE Computational Intelligence Magazine, 1*(4), 28–39.

[14] Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of ICNN'95 - International Conference on Neural Networks*, 4, 1942–1948.

[15] Heylighen, F. (2016). Stigmergy as a universal coordination mechanism II: Varieties and evolution. *Cognitive Systems Research, 38*, 50–59.

[16] Payton, D., Daily, M., Estowski, R., Howard, M., & Lee, C. (2001). Pheromone robotics. *Autonomous Robots, 11*, 319–324.

[17] Fossum, F., Montanier, J. M., & Haddow, P. C. (2014). Repellent pheromones for effective swarm robot search in unknown environments. *2014 IEEE Symposium on Swarm Intelligence*, 1–8.

[18] Hamann, H., & Wörn, H. (2007). An analytical and spatial model of foraging in a swarm of robots. In *Swarm Robotics: Second International Workshop, SAB 2006* (pp. 43–55). Springer.

[19] Kegeleirs, M., & Birattari, M. (2025). Towards applied swarm robotics: current limitations and enablers. *Frontiers in Robotics and AI*, *12*, 1607978.

[20] Hunt, E. R., Jones, S. W., & Hauert, S. (2019). Testing the limits of pheromone stigmergy in high-density robot swarms. *Royal Society Open Science, 6*(11), 190225.

[21] Tang, Q., Xu, Z., Yu, F., Zhang, Z., & Zhang, J. (2019). Dynamic target searching and tracking with swarm robots based on stigmergy mechanism. *Robotics and Autonomous Systems, 120*, 103251.

[22] Chiun, J., Zhang, S., Wang, Y., Cao, Y., & Sartoretti, G. (2025, May). MARVEL: Multi-Agent Reinforcement Learning for constrained field-of-View multi-robot Exploration in Large-scale environments. In *2025 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 11392-11398). IEEE.

[23] Li, L., Yang, B., Chen, C., Yan, Z., Sun, S., Xu, Y., ... & Zhao, L. (2025). Intelligent multi-robot exploration in non-exposed spaces: methods and challenges. *Artificial Intelligence Review*, *58*(12), 394.

[24] Wang, R., Lyu, M., & Zhang, J. (2025). A multi-robot collaborative exploration method based on deep reinforcement learning and knowledge distillation. *Mathematics*, *13*(1), 173.