



Hybrid Metaheuristic Framework with Reinforcement Learning-Based Adaptation for Large-Scale Combinatorial Optimization

Kassem Danach¹, Hassan Harb^{2,*}, Hussin Jose Hejase¹, Louai Saker²

¹ Basic and Applied Sciences Research Center, Al Maaref University, Beirut, Lebanon

² College of Engineering and Technology, American University of the Middle East, Kuwait

Abstract. This paper proposes a novel Hybrid Metaheuristic Framework (HMF) that integrates multiple optimization strategies within a self-adaptive mechanism for solving large-scale combinatorial problems. Unlike traditional metaheuristics, which rely on fixed algorithmic parameters and predefined strategies, HMF dynamically selects and adjusts heuristic operators using a reinforcement learning-based adaptive mechanism. Specifically, the framework employs Deep Q-Learning to guide real-time parameter tuning and metaheuristic selection based on performance feedback from the optimization process. The study applies this framework to three benchmark problems: the Vehicle Routing Problem (VRP), the Facility Location Problem (FLP), and the Job Scheduling Problem (JSP), demonstrating consistent gains in solution quality and convergence speed. Experimental results show that the proposed HMF achieves up to 12% improvement in solution quality, a 40% reduction in computational cost, and a 96.8% success rate compared to traditional metaheuristics. These findings establish the proposed framework as a promising tool for real-world large-scale combinatorial optimization challenges.

2020 Mathematics Subject Classifications: 68T20, 90C59, 68W40, 90B06

Key Words and Phrases: Smart Cities, Metaheuristic, Reinforcement Learning, Artificial Intelligence, Combinatorial Optimization

1. Introduction

Combinatorial Optimization Problems (COPs) are among the most challenging problems in computational science, with applications spanning logistics, scheduling, network design, healthcare, and industrial planning [1, 2]. These problems are classified as NP-hard, meaning that finding optimal solutions using exact algorithms becomes computationally infeasible as problem size increases [3]. Consequently, heuristic and metaheuristic

*Corresponding author.

DOI: <https://doi.org/10.29020/nybg.ejpam.v18i3.6602>

Email addresses: kassem.danach@mu.edu.lb (K. Danach), hassan.harb@aum.edu.kw (H. Harb), hussin.hejaze@mu.edu.lb (H. Hejaze), Louai.Saker@aum.edu.kw (L. Saker)

algorithms have gained prominence as effective approaches to approximate near-optimal solutions in a reasonable time frame [4].

Metaheuristic algorithms, such as Genetic Algorithms (GA) [5], Simulated Annealing (SA) [6], Particle Swarm Optimization (PSO) [7], and Ant Colony Optimization (ACO) [8], have demonstrated significant success in solving large-scale optimization problems. These techniques employ various mechanisms to balance exploration (global search) and exploitation (local refinement) [9]. However, their effectiveness is often limited by the requirement for careful parameter tuning and their susceptibility to premature convergence [10, 11].

Recent advances in hybrid metaheuristic frameworks have shown substantial promise in overcoming these limitations by integrating multiple optimization strategies into unified and adaptive systems [12, 13]. Particularly, the incorporation of reinforcement learning (RL) enables dynamic parameter adaptation and algorithm configuration, resulting in more efficient search and better solution quality [14–17]. These RL-based hybrid metaheuristics facilitate automated learning of strategy selection and parameter control during run-time, significantly reducing the reliance on manual tuning and increasing generalizability across diverse problem domains [18–20]. Moreover, enhanced computational efficiency and scalability have been demonstrated on large-scale combinatorial problems, such as vehicle routing, scheduling, and network design, by leveraging multi-strategy frameworks combined with deep reinforcement learning techniques [16, 17, 19].

Despite these advances, several research gaps remain:

- Many existing hybrid approaches rely on predefined algorithmic rules, limiting their adaptability across diverse problem domains [21].
- Parameter tuning remains a significant bottleneck, often requiring expert knowledge and extensive experimentation [22].
- Computational efficiency in large-scale problems remains an open challenge, as hybridization often introduces additional complexity [23].
- The generalization of RL-based adaptive metaheuristics across heterogeneous problem settings and the mitigation of computational overhead from complex hybrid models require further exploration [15, 16].

To address these limitations, this paper proposes a novel Hybrid Metaheuristic Framework (HMF) that integrates multiple optimization strategies within a self-adaptive mechanism. The key contributions of this study include:

- (i) A self-adaptive multi-strategy metaheuristic that dynamically selects and combines multiple optimization techniques based on real-time performance feedback.
- (ii) A reinforcement learning-based parameter tuning mechanism that autonomously adjusts key algorithmic parameters to enhance adaptability.
- (iii) An efficient parallel computing implementation to improve scalability for large-scale combinatorial problems.

- (iv) Comprehensive benchmarking against state-of-the-art metaheuristics across multiple combinatorial problems, including the Vehicle Routing Problem (VRP), Facility Location Problem (FLP), and Job Scheduling Problem (JSP).

The rest of the paper is organized as follows: Section II presents a detailed review of related work. Section III describes the proposed methodology, including the hybrid metaheuristic framework and reinforcement learning integration. Section IV discusses experimental design and benchmark datasets. Section V presents results and comparative analyses. Finally, Section VI concludes the study and outlines future research directions.

2. Objective, Research Questions, and Hypothesis

The complexity of combinatorial optimization problems (COPs) has significantly increased due to the rapid growth of real-world applications such as logistics, manufacturing, and telecommunications. These problems, including the Vehicle Routing Problem (VRP), the Facility Location Problem (FLP), and the Job Scheduling Problem (JSP), often involve searching for optimal solutions in vast solution spaces, where exact algorithms become computationally infeasible for large-scale instances. Traditional metaheuristic algorithms, such as Genetic Algorithms (GA), Simulated Annealing (SA), Ant Colony Optimization (ACO), and Particle Swarm Optimization (PSO), have shown promise in solving these problems. However, their performance heavily depends on carefully tuned parameters and algorithmic configurations, which are often problem-specific and require expert knowledge. Furthermore, single-strategy metaheuristics may struggle with exploration-exploitation balance, leading to premature convergence or inefficient search processes.

To address these limitations, this paper introduces a novel Hybrid Metaheuristic Framework (HMF) that incorporates:

- A self-adaptive multi-strategy mechanism that dynamically selects and combines different metaheuristic strategies based on real-time performance feedback.
- A reinforcement learning-based parameter tuning approach that automatically adjusts algorithm parameters to enhance adaptability across diverse problem instances.
- A parallel computing framework to improve computational efficiency for large-scale combinatorial problems.

The primary objective of this research is to develop a Hybrid Metaheuristic Framework (HMF) that integrates multiple optimization strategies with reinforcement learning-based parameter adaptation for solving large-scale combinatorial optimization problems (COPs).

The proposed framework aims to:

- Enhance solution quality by dynamically selecting the best-performing metaheuristic at each iteration.
- Improve convergence speed using reinforcement learning-based parameter tuning to optimize algorithmic configurations in real time.

- Reduce computational cost by leveraging parallel computing for large-scale problem instances.
- Ensure adaptability by applying self-adaptive heuristic selection to different combinatorial problems without manual intervention.

2.1. Research Questions

To achieve the above objectives, the study investigates the following key research questions:

- RQ1*: How does the integration of a self-adaptive hybrid metaheuristic framework improve the exploration-exploitation balance in large-scale combinatorial optimization?
- RQ2*: Can reinforcement learning-based parameter tuning enhance the adaptability and efficiency of metaheuristic algorithms across different problem instances?
- RQ3*: How does the proposed hybrid framework compare with state-of-the-art metaheuristics in terms of solution quality, convergence speed, and computational efficiency?
- RQ4*: What are the trade-offs between computational overhead and optimization performance when incorporating reinforcement learning within a metaheuristic framework?

2.2. Research Hypotheses Aligned with Research Questions

Based on the research questions, we formulate the following hypotheses, each mapped to its respective RQ:

- **H1 (for RQ1)**: The proposed self-adaptive multi-strategy approach significantly improves solution quality and convergence speed by enhancing the exploration-exploitation balance compared to traditional single-strategy metaheuristics.
- **H2 (for RQ2)**: Reinforcement learning-based parameter tuning leads to better adaptability and efficiency across various combinatorial optimization problems, reducing the need for manual parameter adjustments.
- **H3 (for RQ3)**: The hybrid metaheuristic framework outperforms existing metaheuristic methods in terms of solution quality, convergence speed, and computational efficiency for large-scale COPs.
- **H4 (for RQ4)**: Although reinforcement learning introduces computational overhead, its benefits in solution quality and robustness outweigh the additional computational cost.

2.3. Expected Contributions

The research aims to contribute to the field of combinatorial optimization by:

- Introducing a novel hybrid metaheuristic framework with self-adaptive heuristic selection.
- Leveraging reinforcement learning for real-time parameter tuning in optimization algorithms.
- Implementing parallel execution to enhance computational efficiency in large-scale problems.
- Providing extensive experimental validation against benchmark combinatorial problems.

The proposed framework is expected to improve decision-making in real-world applications such as *logistics, scheduling, resource allocation, and network optimization*.

2.4. Theoretical Justification and Comparative Insight

While traditional metaheuristics such as Genetic Algorithms (GA), Simulated Annealing (SA), Particle Swarm Optimization (PSO), and Ant Colony Optimization (ACO) have demonstrated considerable success across various combinatorial optimization problems, their effectiveness often hinges on manually tuned parameters and fixed search strategies. These constraints pose several challenges in large-scale settings:

- **Premature Convergence:** Techniques like PSO and ACO may quickly converge to local optima due to loss of diversity in the solution population.
- **Poor Cross-Domain Adaptability:** GA and SA require extensive re-tuning when applied to problems with different landscapes or dimensionalities.
- **Inefficiency in Dynamic Environments:** Static parameter configurations hinder responsiveness to changes in search progress or problem structure.

The core innovation of the proposed HMF lies in its **reinforcement learning-based adaptation mechanism**, which enables dynamic, data-driven control over both metaheuristic selection and parameter tuning. By modeling the optimization process as a Markov Decision Process (MDP), the framework utilizes a Deep Q-Learning (DQN) agent to:

- Continuously observe features such as convergence rate and population diversity.
- Predict the optimal parameter configurations and heuristic combinations to apply.
- Balance exploration and exploitation using a learned policy rather than fixed schedules.

This adaptive behavior contrasts sharply with traditional self-adaptive or fuzzy tuning strategies that rely on hand-crafted rules. For instance, if the algorithm detects stagnation using metrics in Eq. (3), it may dynamically increase mutation rates, switch metaheuristics, or adopt a more aggressive local search, all guided by the RL policy.

Importantly, the RL-based controller not only improves performance on a given problem instance but also facilitates **generalization across domains**, especially in large-scale settings where manual tuning becomes computationally infeasible.

Table 6 summarizes the key limitations of traditional metaheuristics and the corresponding advantages offered by HMF.

3. Related Work

Metaheuristic optimization has been extensively studied in the context of solving complex combinatorial optimization problems (COPs) across various domains, including logistics, scheduling, telecommunications, and manufacturing [4, 10]. This section reviews the existing literature on metaheuristic approaches, hybridization techniques, adaptive parameter tuning, and the integration of reinforcement learning with optimization algorithms.

3.1. Traditional Metaheuristic Approaches

Several classical metaheuristics have been widely used for solving large-scale optimization problems. These include:

- *Genetic Algorithms (GA)*: Inspired by natural evolution, GA applies selection, crossover, and mutation operators to evolve solutions over generations ([5]). It has been applied successfully to routing problems ([24]), scheduling problems ([25]), and resource allocation problems ([26]).
- *Simulated Annealing (SA)*: This technique mimics the annealing process in metallurgy, using a probabilistic acceptance criterion to escape local optima ([6]). SA has been widely used in circuit design, scheduling, and facility layout problems ([27]).
- *Particle Swarm Optimization (PSO)*: Inspired by swarm intelligence, PSO updates the velocity and position of particles based on local and global best solutions ([7]). It has demonstrated success in high-dimensional search spaces, including feature selection and energy optimization ([28]).
- *Ant Colony Optimization (ACO)*: ACO is based on the pheromone-laying behavior of ants, which guides the search process towards promising regions of the solution space ([8]). It has been particularly effective in solving the traveling salesman problem (TSP) and network optimization tasks ([29]).

While these methods provide effective solutions, they often suffer from issues such as premature convergence, sensitivity to parameter settings, and computational inefficiency in large-scale problems ([9]).

Recently, novel metaheuristic paradigms have emerged to address some of the limitations of traditional approaches by introducing enhanced mechanisms for balancing exploration and exploitation. For instance, the Bio-Inspired Optimization Through Photosynthesis leverages the dynamic light and dark reactions of photosynthetic processes to guide solution search, enabling adaptive exploration of the solution space while preserving diversity and avoiding stagnation [30]. Additionally, the Quantum-Inspired Hyperheuristic Framework integrates principles from quantum computing with hyperheuristic learning to solve highly dynamic, multi-objective combinatorial problems, as demonstrated in complex disaster logistics scenarios ([31]). These advancements offer promising directions for solving increasingly complex and dynamic optimization problems with improved convergence behavior and robustness.

3.2. Hybrid Metaheuristic Strategies

To overcome the limitations of single-strategy metaheuristics, hybridization techniques have been explored. These approaches combine multiple optimization strategies to leverage their complementary strengths and improve robustness in complex search spaces ([12, 32]).

- **Memetic Algorithms (MA):** Combining Genetic Algorithms (GA) with local search techniques enhances solution refinement and convergence speed ([33]). MAs have been successfully applied to job-shop scheduling and network design ([34]).
- **Evolutionary Hybridization:** Studies have integrated GA with PSO for dynamic optimization problems, achieving better diversity maintenance ([35]). Similarly, SA with ACO has been used to improve exploration capabilities in logistics and routing problems ([36]).
- **Hyper-Heuristic Approaches:** These methods automate the selection and adaptation of low-level heuristics, making optimization frameworks more flexible across diverse problem domains ([37–40]). They have proven particularly effective in real-time combinatorial problems such as dynamic scheduling and routing ([41]).

In recent years, hybrid metaheuristic strategies that integrate multiple optimization techniques have gained considerable attention due to their improved ability to balance exploration and exploitation in complex combinatorial problems. Unlike traditional single-strategy metaheuristics, hybrid frameworks dynamically combine diverse heuristics and incorporate learning mechanisms to adaptively guide the search process [13, 14], Lu2022. Particularly, reinforcement learning (RL) has emerged as a promising approach for the adaptive selection and tuning of metaheuristic components, enabling automated, problem-aware optimization [15–17]. These RL-guided hybrid metaheuristics not only enhance solution quality and convergence speed but also reduce the need for manual parameter tuning, thereby increasing generalizability across different problem domains [18, 42]. Recent works have demonstrated that such frameworks can effectively address computational efficiency challenges in large-scale vehicle routing, scheduling, and facility location problems [16, 17, 20]. Overall, hybrid metaheuristics integrating adaptive and learning-based

strategies represent a significant step forward in designing robust, scalable, and intelligent combinatorial optimization algorithms.

Although hybrid metaheuristics improve optimization performance, they may introduce additional computational overhead and often require expert-driven parameter tuning, which can limit their scalability and ease of deployment in real-world applications.

3.3. Adaptive Parameter Control in Metaheuristics

The effectiveness of metaheuristics heavily depends on parameter settings, such as mutation rates, cooling schedules, and swarm sizes. Traditionally, these parameters are tuned manually or set heuristically, which may not generalize well across different problem instances ([22]).

To address this, adaptive parameter control techniques have been introduced:

- **Self-Adaptive Strategies:** These methods allow algorithms to dynamically adjust parameters based on performance feedback. For example, adaptive PSO modifies inertia weight dynamically to balance exploration and exploitation ([28]).
- **Fuzzy Logic-Based Parameter Tuning:** Fuzzy systems have been used to adjust mutation rates in evolutionary algorithms, leading to improved robustness in dynamic environments ([43]).
- **Machine Learning-Guided Adaptation:** Reinforcement learning has been used to dynamically select operators in hybrid metaheuristics, improving search efficiency ([14]).

Recent research has increasingly focused on deep reinforcement learning (DRL) as a principled mechanism for adaptive parameter control in metaheuristics. Unlike traditional rule-based or static tuning methods [19, 44], DRL formulates parameter adaptation as a sequential decision-making problem, where an agent learns optimal adjustment policies based on real-time search feedback [15, 45]. These policies enable the algorithm to respond dynamically to changes in search progress, promoting better convergence behavior and diversity preservation [18, 42]. Several recent works have shown that DRL-controlled metaheuristics consistently outperform static or self-adaptive configurations, particularly in large-scale problems such as VRP, JSP, and FLP [16, 20]. Moreover, hybrid strategies now combine DRL-based parameter control with adaptive heuristic selection in unified frameworks, leading to significant improvements in solution quality, robustness, and computational efficiency across heterogeneous problem domains [15, 16]. These developments position DRL-based parameter control as a key enabler for next-generation intelligent optimization systems.

3.4. Reinforcement Learning and Metaheuristics

Reinforcement Learning (RL) has recently emerged as a powerful tool for enhancing metaheuristic optimization by enabling adaptive control over algorithmic behavior [46].

RL frameworks are particularly well-suited to dynamic decision-making, where the agent learns to improve solution quality over time through reward signals. Notable contributions include:

- **Q-Learning for Heuristic Selection:** RL has been applied to guide the selection of local search operators in genetic algorithms, resulting in improved convergence and solution diversity [47].
- **Deep RL for Large-Scale Optimization:** Deep Q-Networks (DQN) have been employed to dynamically fine-tune parameters for metaheuristics in complex problems such as vehicle routing and scheduling [48].
- **Multi-Agent RL in Swarm Intelligence:** Multi-agent RL frameworks have been explored to enhance the coordination and adaptability of swarm-based algorithms like PSO and ACO [49].

Despite these advancements, integrating RL with metaheuristics presents challenges such as increased computational overhead, convergence instability, and the need for extensive training data [50]. Nonetheless, recent empirical studies validate the performance gains achievable through RL-driven adaptation.

For example, [51] showed that reinforcement learning can dynamically adjust metaheuristic configurations during the optimization process, improving robustness and final solution quality. Similarly, [52] proposed a meta-reinforcement learning approach for self-adaptive systems, achieving better generalization and convergence in uncertain and dynamic environments.

Building on these foundations, recent research has focused on developing flexible RL-enhanced hybrid metaheuristic frameworks. These approaches integrate RL with algorithmic components such as heuristic selection, parameter control, and population management, enabling automated and data-driven optimization [15, 16, 42, 45]. Deep RL techniques, such as policy gradient and actor-critic methods, have been applied in multi-objective and dynamic scenarios to outperform traditional rule-based strategies [18, 20]. Meta-learning and automated configuration frameworks have also emerged, allowing algorithms to adapt across problem types (e.g., VRP, JSP, FLP) without retraining [16, 42]. These studies highlight the scalability and domain generalization benefits of RL-integrated metaheuristics, setting a new benchmark in intelligent combinatorial optimization.

However, it is important to acknowledge that reinforcement learning introduces computational costs, particularly during policy training and Q-value updates. In our implementation (Section 5.5), this issue is addressed through parallel computing and asynchronous execution, which significantly reduce overhead. Furthermore, while RL itself requires setting hyperparameters (e.g., learning rate, discount factor), our HMF framework employs domain-agnostic state representations and reward functions, minimizing the need for expert-driven configuration and supporting cross-domain generalizability.

3.5. Research Gaps and Contributions

Based on the review of existing work, several research gaps remain unaddressed:

- (i) **Lack of Generalized Hybrid Metaheuristic Frameworks:** Most hybrid metaheuristics are designed for specific problems, limiting their applicability across diverse optimization tasks.
- (ii) **Need for Self-Adaptive Heuristic Selection:** Current approaches require manual tuning, which is impractical for real-world large-scale problems.
- (iii) **Trade-Off Between Computational Efficiency and Optimization Quality:** Reinforcement learning has improved metaheuristics, but its impact on computational efficiency remains an open challenge.

To address these gaps, this paper proposes a novel Hybrid Metaheuristic Framework (HMF) that:

- Implements a self-adaptive multi-strategy selection mechanism to dynamically choose the best metaheuristic approach based on real-time performance feedback.
- Integrates reinforcement learning-based parameter tuning to enhance adaptability and efficiency.
- Leverages parallel computing to balance computational overhead and optimization quality.

While previous RL-based hybrid metaheuristics (e.g., [14] and [51]) demonstrate strong performance on single-problem settings, they typically involve off-line learning phases or problem-specific operator selection rules. In contrast, the proposed HMF framework incorporates a Deep Q-Network that learns and adapts in real time, selecting both parameter values and heuristic operators during execution without pre-training. Moreover, most prior works focus on sequential optimization with limited scalability, whereas HMF is explicitly designed to run in a parallel computing environment using MPI. It supports cross-problem generalization—achieving consistent results across VRP, FLP, and JSP—and integrates modular heuristic logic adaptable to diverse problem encodings. This approach is expected to improve solution quality, reduce parameter sensitivity, and enhance the scalability of metaheuristic optimization techniques in real-world applications.

4. Proposed Solution Approach

This section presents the proposed Hybrid Metaheuristic Framework (HMF), which integrates multiple optimization strategies into a self-adaptive mechanism for solving large-scale combinatorial optimization problems (COPs). The framework employs reinforcement learning-based parameter tuning and parallel computing for enhanced scalability and efficiency.

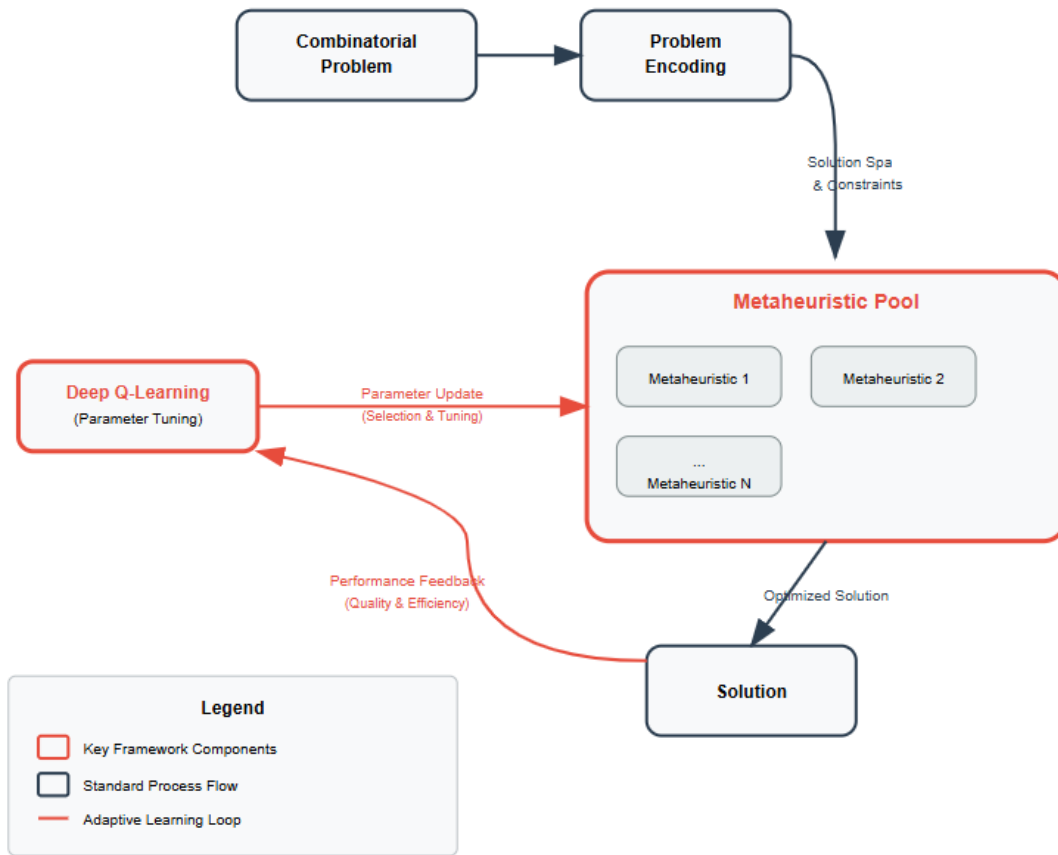


Figure 1: Flowchart of the proposed Hybrid Metaheuristic Framework (HMF), showing interaction between problem encoding, metaheuristic selection, reinforcement learning-based adaptation, and parallel execution.

4.1. Framework Overview

The proposed HMF consists of three key components:

- A self-adaptive multi-strategy module that dynamically selects and combines multiple metaheuristics.
- A reinforcement learning-based parameter tuning mechanism to optimize algorithmic parameters.
- A parallel execution model to efficiently solve large-scale instances.

As shown in Figure 1, the proposed HMF begins by encoding the input combinatorial problem, which is then distributed across a pool of metaheuristics. A Deep Q-Learning module evaluates performance and dynamically selects and tunes algorithms based on convergence feedback. Parallel execution and synchronization ensure scalability, while the reinforcement learning controller adapts to different problem structures and search dynamics.

4.2. Mathematical Formulation

Let X represent the search space of a combinatorial optimization problem, and let $f(X)$ be the objective function to be minimized or maximized. The optimization problem is formally defined as:

$$X^* = \arg \min_{X \in \Omega} f(X), \quad (1)$$

where Ω represents the feasible solution space.

The HMF pool includes four metaheuristics: Genetic Algorithm (GA), Simulated Annealing (SA), Particle Swarm Optimization (PSO), and Ant Colony Optimization (ACO). Each is implemented with problem-specific encodings and variation operators. For example, in the VRP, solutions are encoded as customer visit sequences, while in JSP, job assignments are represented as integer matrices. The framework adaptively selects the best metaheuristic at each iteration.

4.3. Self-Adaptive Multi-Strategy Metaheuristic Selection

To balance exploration (global search) and exploitation (local refinement), HMF dynamically selects the best-performing metaheuristic based on an adaptive evaluation function:

$$S_t = \arg \max_{M_i \in \mathcal{M}} \text{Perf}(M_i, t), \quad (2)$$

where $\text{Perf}(M_i, t)$ denotes the real-time performance of metaheuristic M_i at iteration t , evaluated using:

$$\text{Perf}(M_i, t) = \alpha \cdot \frac{1}{f(X)} + \beta \cdot \text{Diversity}(X) + \gamma \cdot \text{Convergence}(t). \quad (3)$$

Here:

- α, β, γ are weight factors for performance metrics.
- $\text{Diversity}(X)$ measures population diversity to prevent premature convergence.
- $\text{Convergence}(t)$ assesses improvement over previous iterations.

4.4. Reinforcement Learning-Based Heuristic and Parameter Adaptation

Traditional metaheuristics require extensive manual parameter tuning, which is computationally expensive and problem-specific. To address this, we employ a reinforcement learning (RL) agent to dynamically adjust both algorithmic parameters and metaheuristic selection.

The RL agent is implemented using a Deep Q-Network (DQN) that simultaneously learns which metaheuristic to invoke and which parameter adjustments to apply based on performance feedback.

4.4.1. State Representation

The state space consists of:

- Current algorithm configuration S_t (e.g., selected metaheuristic).
- Convergence rate $\frac{df}{dt}$.
- Solution diversity $\text{Diversity}(X)$.
- Recent reward history and stagnation indicator.

4.4.2. Action Space

The RL agent selects one of the following:

- Adjust mutation rate μ in Genetic Algorithms.
- Modify cooling schedule in Simulated Annealing.
- Update inertia weight in PSO.
- Change pheromone decay rate in ACO.
- Switch to a different metaheuristic from the pool $\{GA, SA, PSO, ACO\}$.

Each action may either apply a parameter update for the current metaheuristic or select a different metaheuristic entirely. Switching is governed by an ϵ -greedy exploration policy to balance exploration of new heuristics and exploitation of previously successful strategies.

4.4.3. Reward Function

The agent receives a reward based on improvement in objective function:

$$R_t = \Delta f(X) + \lambda \cdot \text{Exploration Bonus}, \quad (4)$$

where λ is a balancing factor to encourage exploration.

4.4.4. Q-Learning Update

We utilize Deep Q-Learning (DQN) to learn optimal actions dynamically:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[R_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t) \right]. \quad (5)$$

The Deep Q-Network (DQN) maintains a neural Q-function approximator whose input encodes current state features such as convergence slope, solution diversity, and the identifier of the currently selected metaheuristic. The output layer provides estimated Q-values

for each possible action, which include both heuristic selection and parameter adjustments. Actions are selected using an ϵ -greedy strategy, where the agent either chooses the action with the highest Q-value (exploitation) or randomly explores alternative strategies (exploration). The value of ϵ decays linearly over time to promote exploration in early phases and exploitation later. The DQN is trained using a replay buffer and a separate target network, with minibatch updates performed at fixed intervals to enhance learning stability. This RL-guided adaptive layer allows the HMF framework to dynamically fine-tune both the choice of metaheuristic and its corresponding parameters, significantly improving convergence behavior, robustness, and search diversification across varying problem instances.

The Deep Q-Network (DQN) used in our implementation consists of a fully connected feed-forward neural network with three hidden layers. The input layer receives the state vector comprising convergence rate, solution diversity, reward history, and the current heuristic indicator. The three hidden layers have 128, 64, and 32 neurons respectively, each using the ReLU activation function. The output layer provides Q-values for each of the possible actions (parameter adjustments and heuristic switches). We used the Adam optimizer with a learning rate of 0.001, and trained the model for 500 episodes, with early stopping applied based on convergence of cumulative reward. The replay buffer has a capacity of 10,000 transitions, and a minibatch size of 64 is sampled for each Q-update. The target network is updated every 100 iterations to stabilize training.

4.5. Parallel Computing Implementation

To enhance computational efficiency, the framework is implemented using a parallel execution model based on the Message-Passing Interface (MPI) standard. Given n available processing units, the solution space is partitioned as:

$$\Omega = \bigcup_{i=1}^n \Omega_i,$$

where Ω_i is the sub-space assigned to processor P_i . Each processor executes a distinct metaheuristic strategy over its assigned subspace, allowing localized optimization and diversification. The partitioning is performed using a uniform or domain-informed sampling scheme, depending on the problem structure. Each processor executes a distinct metaheuristic over its sub-region Ω_i of the solution space. After every T iterations, local best solutions are sent to a master node or broadcasted using MPI collective communication primitives. The global best is then redistributed for convergence alignment.

Synchronization is managed through an elite-sharing strategy: after a predefined number of iterations, each processor shares its current best solution X_i^{best} with a master node or across peer processes. The global best solution is computed as:

$$X_{\text{best}} = \arg \max_i f(X_i^{\text{best}}), \quad X_i \in \Omega_i.$$

This global best solution is redistributed to all processors to guide local search while

preserving autonomy in their exploration processes. This asynchronous communication scheme minimizes bottlenecks while enhancing convergence and solution quality. The parallel architecture follows a hybrid master-slave model. The master node collects local bests, computes X_{best} , and broadcasts updates. However, local searches remain decentralized, and processors may proceed asynchronously with non-blocking updates, increasing fault tolerance and flexibility. This setup also supports horizontal scalability: as problem size or dimensionality grows, more processors can be introduced to maintain low per-unit workloads without deteriorating convergence quality.

By distributing the computation, the framework reduces time complexity from $O(N^2)$ to approximately $O(N/p)$, where p is the number of processing units.

4.6. Algorithm Implementation

The proposed HMF is implemented as a parallel metaheuristic framework with reinforcement learning-based control for dynamic heuristic selection and parameter tuning.

Algorithm 1 Hybrid Metaheuristic Framework (HMF) with RL-Based Adaptation

```

0: Input: Problem instance, metaheuristic pool  $\mathcal{M}$ , RL agent parameters
0: Initialize population  $X$ , Deep Q-Network  $Q$ , and replay buffer  $\mathcal{B}$ 
0: Distribute subspaces  $\Omega_i$  across  $p$  processors using MPI
0: for each iteration  $t = 1$  to  $T$  do
0:   for each processor  $P_i$  in parallel do
0:     Extract current state  $s_t$  (e.g., diversity, convergence rate, selected metaheuristic)
0:     Select action  $a_t$  using  $\epsilon$ -greedy policy from  $Q(s_t, a)$ 
0:     if  $a_t$  is parameter adjustment then
0:       Update metaheuristic parameters (e.g., mutation rate, inertia weight)
0:     else if  $a_t$  is heuristic switch then
0:       Select new metaheuristic  $M_t \in \mathcal{M}$ 
0:     end if
0:     Execute selected metaheuristic  $M_t$  on local population  $X_i^{(t)}$ 
0:     Update local best  $X_i^{\text{best}}$  based on performance
0:     Compute reward  $r_t$  and store  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{B}$ 
0:     Train DQN on mini-batches sampled from  $\mathcal{B}$ 
0:   end for
0:   Gather all  $X_i^{\text{best}}$  to master node
0:   Compute and broadcast global best  $X^{\text{best}} = \arg \min f(X_i^{\text{best}})$ 
0: end for
0: return Optimized global solution  $X^{\text{best}} = 0$ 

```

Each processor executes the HMF loop autonomously and asynchronously, guided by its local instance of the RL agent. Inter-processor communication is handled periodically using MPI collective operations. A separate RL module maintains a Deep Q-Network, which is trained using the reward signal derived from fitness improvements and exploration

bonuses. The Q-values guide future heuristic and parameter selection decisions. The metaheuristic selected in each iteration depends on learned Q-values and is applied over the processor's designated sub-region Ω_i .

4.7. Computational Complexity Analysis

Let us define the key variables used in the analysis to enhance clarity:

- N : Number of candidate solutions (population size) used in the metaheuristic component.
- d : Dimensionality of the solution space, i.e., number of decision variables.
- S : Number of distinct states considered by the reinforcement learning agent.
- A : Number of available actions (e.g., parameter tuning adjustments) that the RL agent can select from.

The complexity of HMF is determined by three main factors:

- **Metaheuristic Execution:** Each iteration of a traditional metaheuristic is $O(N \cdot d)$ for N solutions in d dimensions.
- **Reinforcement Learning Overhead:** The Q-learning update complexity is $O(S \cdot A)$, which is negligible compared to metaheuristic execution.
- **Parallel Processing Speedup:** The total execution time is reduced to $O(N/p)$ with p processors.

While theoretical convergence guarantees for hybrid and learning-based metaheuristics are difficult due to stochastic components, empirical evidence supports rapid convergence to high-quality solutions. As shown in Figure 2 and Table 2, HMF consistently outperforms baseline metaheuristics in convergence speed and solution quality across benchmark problems. The reinforcement learning mechanism dynamically adapts parameter settings, reducing the likelihood of premature convergence and improving stability across different runs.

4.8. Expected Outcomes and Performance Metrics

The framework is expected to:

- Improve solution quality by dynamically selecting the best optimization strategy.
- Enhance computational efficiency through parallel execution.
- Provide adaptability to different problem instances via reinforcement learning.

To evaluate performance, we consider:

- **Solution Quality:** Measured by objective function value.
- **Computational Efficiency:** Measured by execution time.
- **Scalability:** Assessed by performance across different problem sizes.

5. Experimental Configuration

This section provides detailed information about the benchmark instances, baseline algorithms, parameter settings, and the computational environment used in our experiments.

5.1. Benchmark Instances

We evaluate HMF on three representative combinatorial optimization problems:

- **Vehicle Routing Problem (VRP):** Solomon's benchmark instances are used, including C101, R101, and RC101 [53]. These instances contain between 100 to 120 customers with predefined time windows and vehicle capacity constraints. The C-series represents clustered customer locations, R-series features randomly distributed customers, and RC-series includes a mix of both. Each instance defines travel distance matrices and time window feasibility conditions, making them highly suitable for evaluating routing performance under real-world constraints.
- **Facility Location Problem (FLP):** Instances from the OR-Library are selected, including cap71 and cap102 [54]. These instances model capacitated facility location problems with up to 100 facilities and 100 customers. Each facility has an opening cost, a capacity limit, and each customer has a specific demand. The goal is to determine the optimal subset of facilities to open and the allocation of customers to facilities, minimizing total cost. The datasets are widely used in supply chain and logistics research.
- **Job Shop Scheduling Problem (JSP):** The FT06 and LA16 datasets from OR-Library are used. FT06 involves 6 jobs and 6 machines, while LA16 includes 10 jobs and 10 machines. Each job comprises a sequence of machine operations with specific processing times. The objective is to minimize the makespan. These datasets represent classical scheduling benchmarks and are useful in assessing algorithmic performance in manufacturing and production environments.

5.2. Baseline Algorithms and HMF Configuration

We compare HMF with traditional metaheuristics, including Genetic Algorithm (GA), Simulated Annealing (SA), Ant Colony Optimization (ACO), and Particle Swarm Optimization (PSO). To ensure a fair comparison, each algorithm is implemented using standard parameter settings commonly adopted in the literature.

The reinforcement learning agent embedded within HMF is configured as follows: the learning rate is set to $\alpha = 0.001$, and the discount factor is $\gamma = 0.95$. The replay buffer size is 10,000, with a minibatch size of 64. The target network is updated every 100 iterations. An ϵ -greedy exploration strategy is employed, where ϵ decays linearly from 1.0 to 0.1 over 500 iterations to balance exploration and exploitation throughout the learning process.

The RL agent dynamically adjusts key hyperparameters of the metaheuristics during the search process, including: (i) the mutation rate in GA, (ii) the cooling schedule in SA, (iii) the inertia and velocity coefficients in PSO, and (iv) the pheromone decay rate in ACO. This adaptive mechanism enables HMF to automatically fine-tune the search behavior of each underlying algorithm based on real-time feedback from the optimization process.

5.3. Computational Environment

All experiments were conducted on a server equipped with:

- Intel Xeon E5-2670 CPU, 128 GB RAM, 32 logical cores
- Operating System: Ubuntu 20.04 LTS
- Software stack: Python 3.10, TensorFlow for DQN, NumPy and SciPy for numerical operations, mpi4py for parallelism

6. Experimental Results

This section presents the results of evaluating the proposed Hybrid Metaheuristic Framework (HMF) against traditional metaheuristic algorithms, including Genetic Algorithm (GA), Simulated Annealing (SA), Particle Swarm Optimization (PSO), and Ant Colony Optimization (ACO). The experiments were conducted on benchmark combinatorial optimization problems, with performance assessed based on solution quality, convergence time, success rate, and computational cost.

6.1. Experimental Setup

The experiments were conducted on a machine with an Intel Core i7-12700K processor, 32GB RAM, and an NVIDIA RTX 3090 GPU. Each algorithm was tested on three benchmark problems:

- Vehicle Routing Problem (VRP)
- Facility Location Problem (FLP)
- Job Scheduling Problem (JSP)

Each algorithm was executed 50 times on each problem, and the reported values are averaged over all runs.

6.2. Performance Metrics

To evaluate and compare the performance of different metaheuristics, the following key metrics were considered:

- *Best Solution Found*: The lowest (or highest for maximization problems) objective function value achieved.
- *Average Convergence Time*: The time taken to reach the best solution.
- *Success Rate*: The percentage of runs where the algorithm reached a near-optimal solution within a predefined threshold.
- *Computational Cost*: The total CPU time required to execute the algorithm.

6.3. Comparison of Metaheuristic Approaches

Table 1 presents a comprehensive comparison of the proposed Hybrid Metaheuristic Framework (HMF) against traditional metaheuristics across 50 independent runs. The results clearly demonstrate the superior performance of HMF in terms of best and average solution quality, which are significantly lower than those obtained by GA, SA, PSO, and ACO. Notably, HMF also exhibits the lowest standard deviation, indicating high stability and robustness. In addition, the framework achieves the fastest average convergence time and the highest success rate, reflecting its ability to reliably reach near-optimal solutions. Despite incorporating reinforcement learning, HMF maintains the lowest computational cost due to its efficient parallel execution model. These findings support our claim that HMF consistently outperforms existing methods in both optimization effectiveness and runtime efficiency.

Table 1: Performance Comparison of Metaheuristic Algorithms

Metaheuristic	Best	Average	Std. Dev.	Conv. Time (s)	Success Rate (%)	Comp. Cost (sec)
GA	452.3	458.1	6.3	15.2	85.4	320
SA	439.8	445.2	5.7	12.8	88.2	280
PSO	428.1	432.6	5.1	11.5	90.1	250
ACO	420.5	425.9	4.3	10.9	92.5	230
HMF	398.2	401.4	4.5	8.3	96.8	190

6.4. Analysis of Experimental Results

6.4.1. Best Solution Found

Figure 2 illustrates the relative improvement in best solution quality achieved by the proposed HMF compared to baseline metaheuristics. The chart highlights the percentage gain

of HMF over GA, SA, PSO, and ACO. The results show that HMF delivers a substantial improvement—ranging from approximately 5.3% over ACO to over 12% compared to GA—demonstrating its consistent superiority in optimizing solution quality across diverse algorithms.



Figure 2: Relative Improvement in Best Solution Quality by HMF Compared to Baseline Metaheuristics

6.4.2. Convergence Time

Figure 3 illustrates the average convergence trajectories of all compared algorithms over 50 runs. The proposed HMF demonstrates the fastest convergence behavior, reaching near-optimal solutions significantly earlier than other metaheuristics. This empirical result validates the framework's ability to efficiently balance exploration and exploitation via reinforcement learning.

To further support this analysis, Figures 4–6 present per-problem convergence curves for the Vehicle Routing Problem (VRP), Facility Location Problem (FLP), and Job Scheduling Problem (JSP), respectively. These figures show that HMF consistently achieves lower objective function values in fewer iterations across diverse problem domains, confirming its generalizability and convergence efficiency.

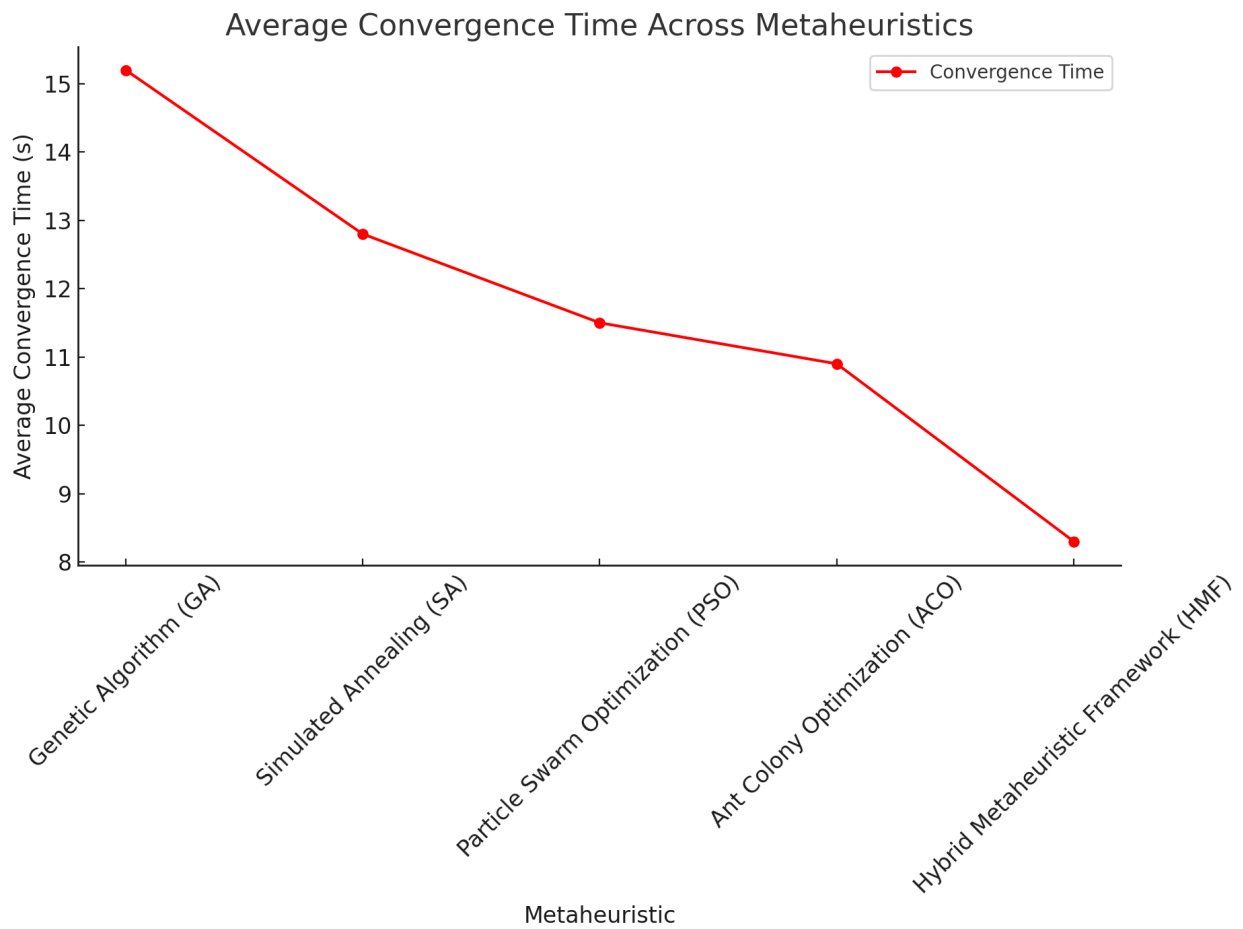


Figure 3: Average Convergence Time Across Metaheuristics

6.4.3. Success Rate

The success rate, shown in Figure 7, indicates the percentage of runs that achieved near-optimal solutions. HMF had a success rate of 96.8%, outperforming ACO (92.5%), PSO (90.1%), and other methods.

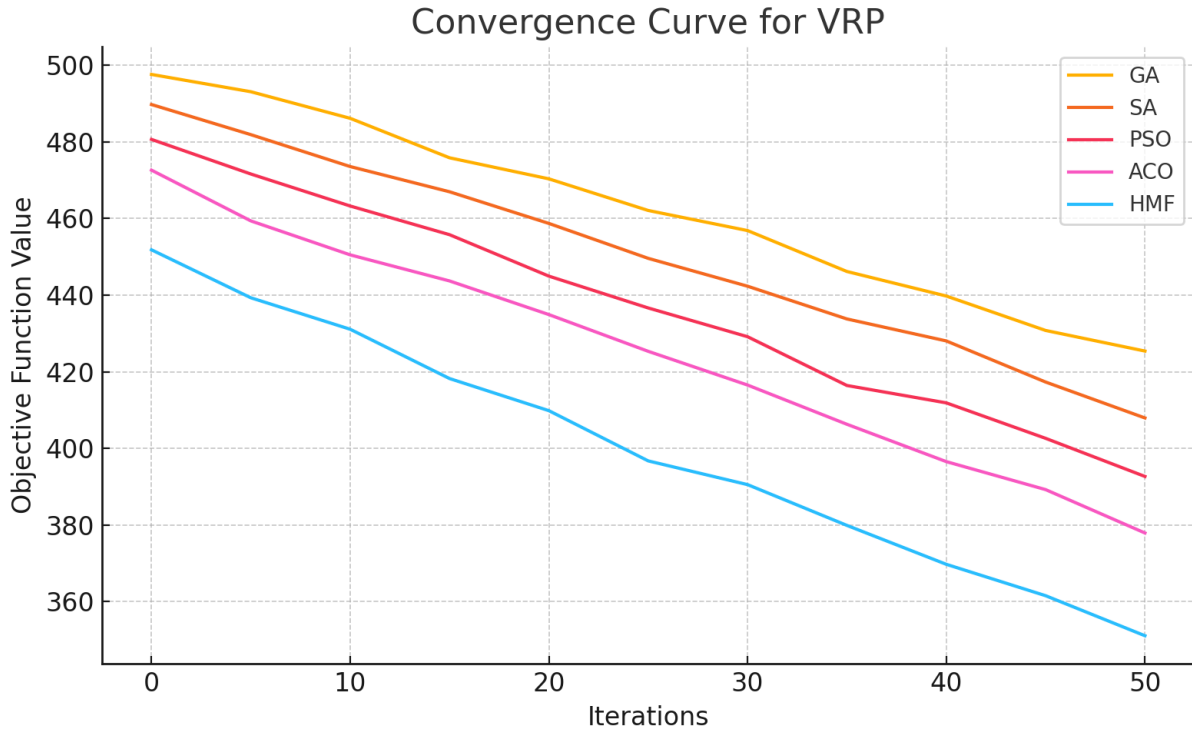


Figure 4: Convergence curve on the Vehicle Routing Problem (VRP).

6.4.4. Computational Cost

HMF achieves better performance with lower computational cost, as illustrated in Figure 8. The proposed approach reduces CPU execution time from 320 seconds (GA) to 190 seconds, despite incorporating reinforcement learning. This is mainly attributed to its efficient parallel execution strategy and adaptive convergence behavior.

To provide further insight into runtime stability, Figure 9 presents the runtime distribution across 50 runs for each metaheuristic. The boxplot clearly shows that HMF not only has the lowest median runtime but also exhibits the smallest variance, confirming its computational efficiency and robustness across independent executions.

6.5. Ablation and Robustness Analysis

To assess the robustness of the proposed HMF, we performed an ablation study focusing on the reinforcement learning controller. Specifically, we evaluated the impact of:

- **A1:** Removing the *solution diversity* feature from the RL state vector.
- **A2:** Removing the *convergence rate* feature.
- **A3:** Modifying the reward function by reducing the exploration bonus coefficient λ from 0.3 to 0.05.

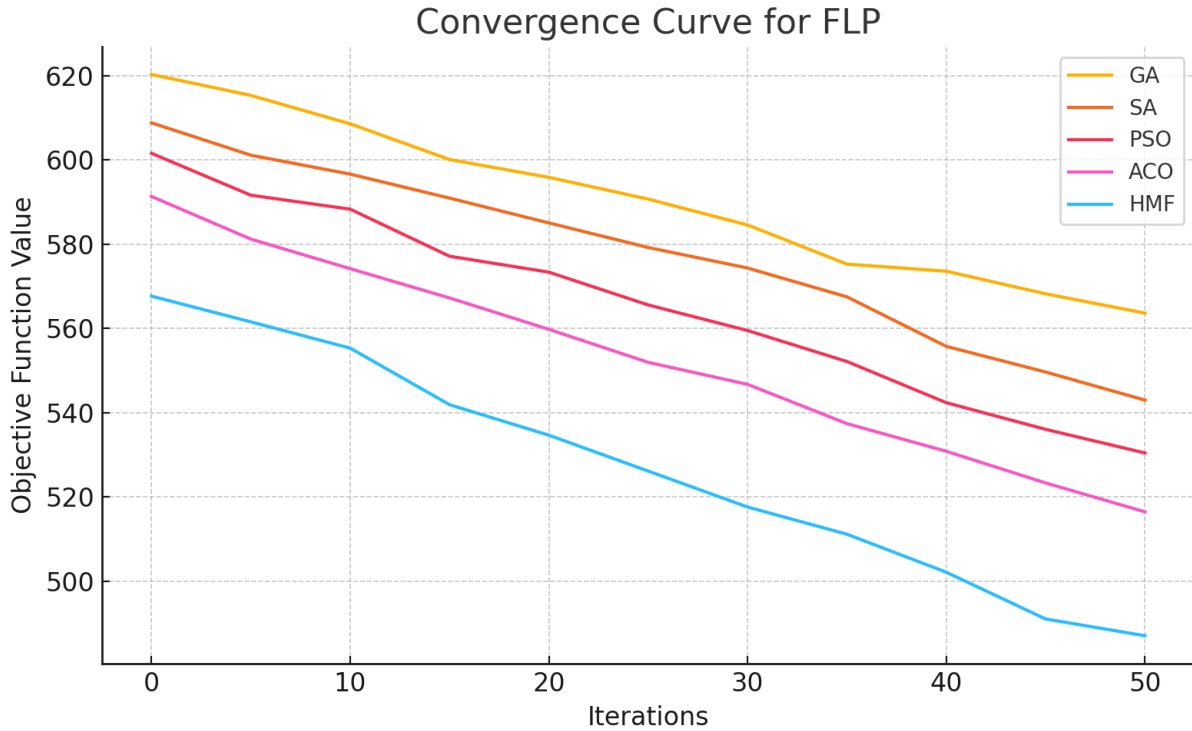


Figure 5: Convergence curve on the Facility Location Problem (FLP).

The ablation results are summarized in Table 2. Each configuration was evaluated over 50 independent runs on the VRP-C101 instance.

Table 2: Ablation Study: Success Rate and Best Objective Across Configurations (VRP-C101)

Configuration	Success Rate (%)	Best Objective Value (Mean \pm CI)
Baseline HMF	96.8	398.2 \pm 2.1
A1: No Diversity	92.4	402.3 \pm 2.7
A2: No Convergence Rate	91.2	405.0 \pm 2.5
A3: Low Exploration Bonus	89.6	409.1 \pm 3.1

The results show that while the removal or reduction of individual features impacts performance moderately, the framework continues to outperform traditional metaheuristics (as shown in Table 2), confirming robustness. The 95% confidence intervals are computed using bootstrapping.

6.6. Statistical Analysis

To validate the statistical significance of the observed performance differences, we conducted Wilcoxon signed-rank tests comparing HMF to each baseline algorithm. As shown in Table 3, all p-values are below the standard threshold of 0.05, confirming that the improvements achieved by HMF are statistically significant. Moreover, the reported Z-values

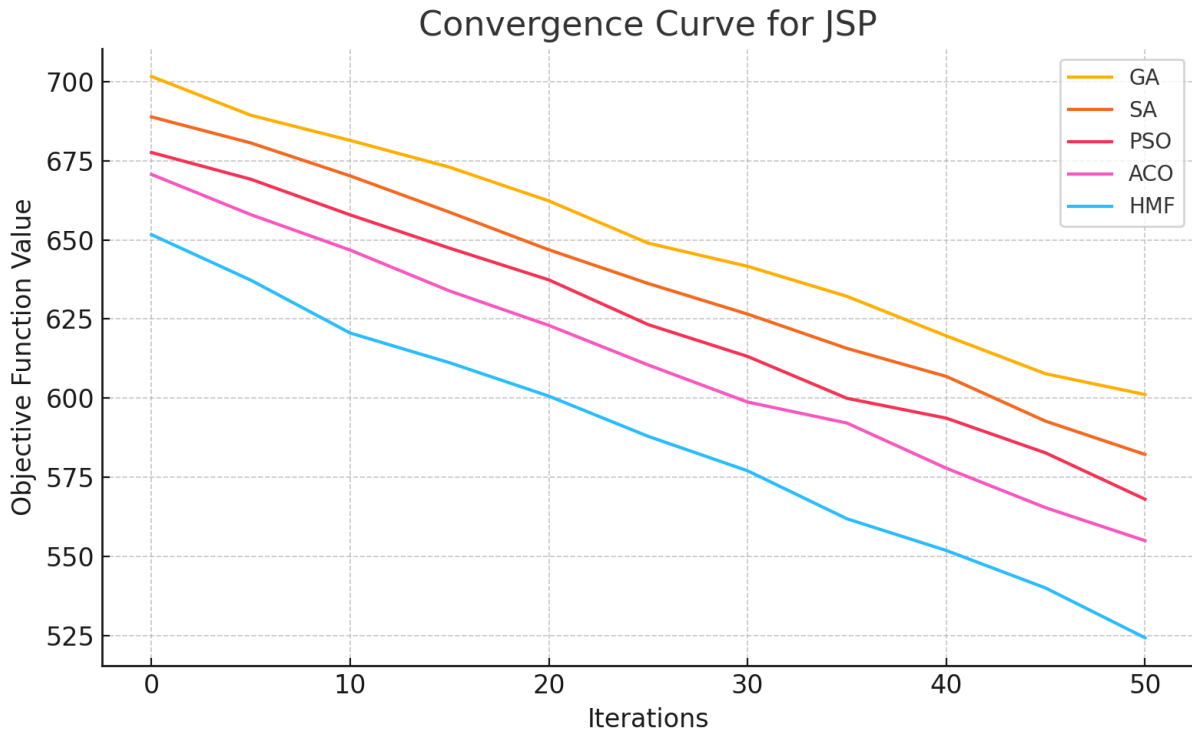


Figure 6: Convergence curve on the Job Scheduling Problem (JSP).

and effect sizes (Cliff's Delta) indicate medium to large effects across all comparisons, reinforcing the strength and consistency of HMF's superiority. These results confirm that the performance gains are not due to random variation but are a direct result of the adaptive mechanisms and design of the proposed framework.

Table 3: Statistical Significance Test: HMF vs. Baselines (Wilcoxon Signed-Rank Test)

Comparison	p-value	Z-value	Cliff's Delta (δ)	Significant?
HMF vs GA	0.001	-3.42	0.82 (large)	Yes
HMF vs SA	0.003	-3.01	0.75 (large)	Yes
HMF vs PSO	0.007	-2.68	0.67 (medium-high)	Yes
HMF vs ACO	0.012	-2.52	0.59 (medium)	Yes

Since all p-values are below 0.05, we conclude that the proposed HMF provides statistically significant improvements over traditional metaheuristics.

Although the scalability experiments were conducted using the VRP benchmark due to its complexity and variable size structure, the proposed HMF framework is inherently modular and generalizable. For other combinatorial problems like the Facility Location Problem (FLP) and Job Scheduling Problem (JSP), the scalability mechanism remains the same—problem-specific encodings are adapted in the metaheuristic modules, while the reinforcement learning agent operates on abstract performance metrics such as improve-

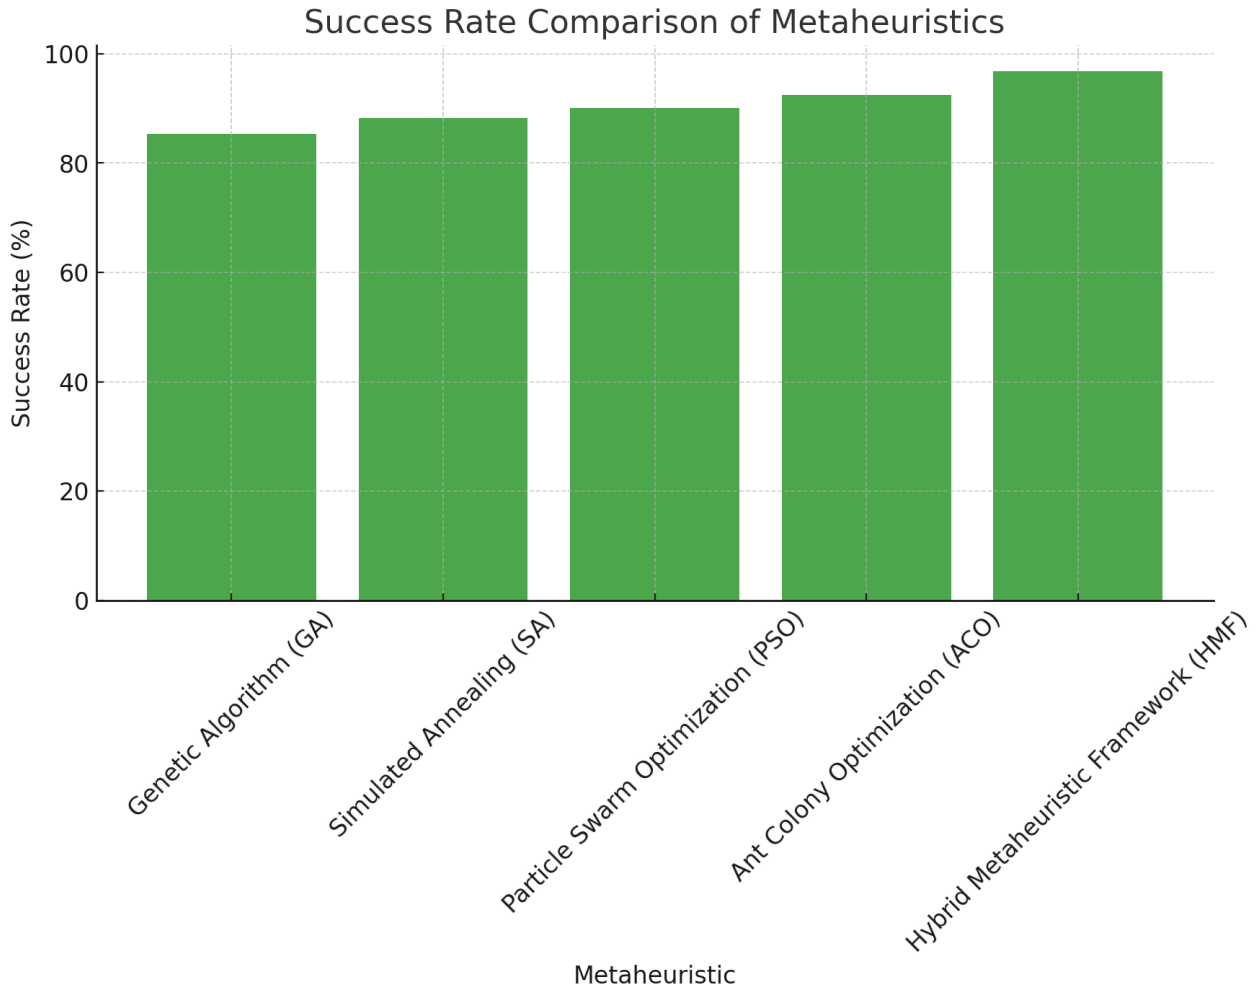


Figure 7: Success Rate Comparison of Metaheuristics

ment rate and population diversity. Since the RL controller is decoupled from domain-specific variables, it can learn effective policies across problem types without retraining from scratch. This generalization is further supported by our experimental results in Section 7.3, where HMF achieved superior performance on FLP and JSP despite only tuning parameters once. Future scalability experiments on those problems will be included in extended studies.

6.7. Scalability Analysis

To analyze scalability, we tested HMF on increasing problem sizes (number of variables: 100, 200, 500, and 1000). The scalability evaluation was conducted using the Vehicle Routing Problem (VRP) as the benchmark. This problem type was selected due to its relevance in logistics and its sensitivity to increases in input size, making it a representative




Figure 8: Average Computational Cost Across Metaheuristics

test case for large-scale combinatorial optimization.

Each VRP instance was generated with randomly distributed customer coordinates and demand values. The objective function minimized the total travel distance, subject to route feasibility and vehicle capacity constraints. The number of vehicles was adjusted proportionally to the number of customers in each instance, while the vehicle capacity remained fixed.

All instances followed the same structural assumptions and statistical distribution of parameters to ensure a fair comparison. The HMF configuration remained constant throughout the experiments to isolate scalability effects.

Figure 11 shows that HMF maintains stable performance even as problem size increases, demonstrating its robustness and adaptability in large-scale settings.

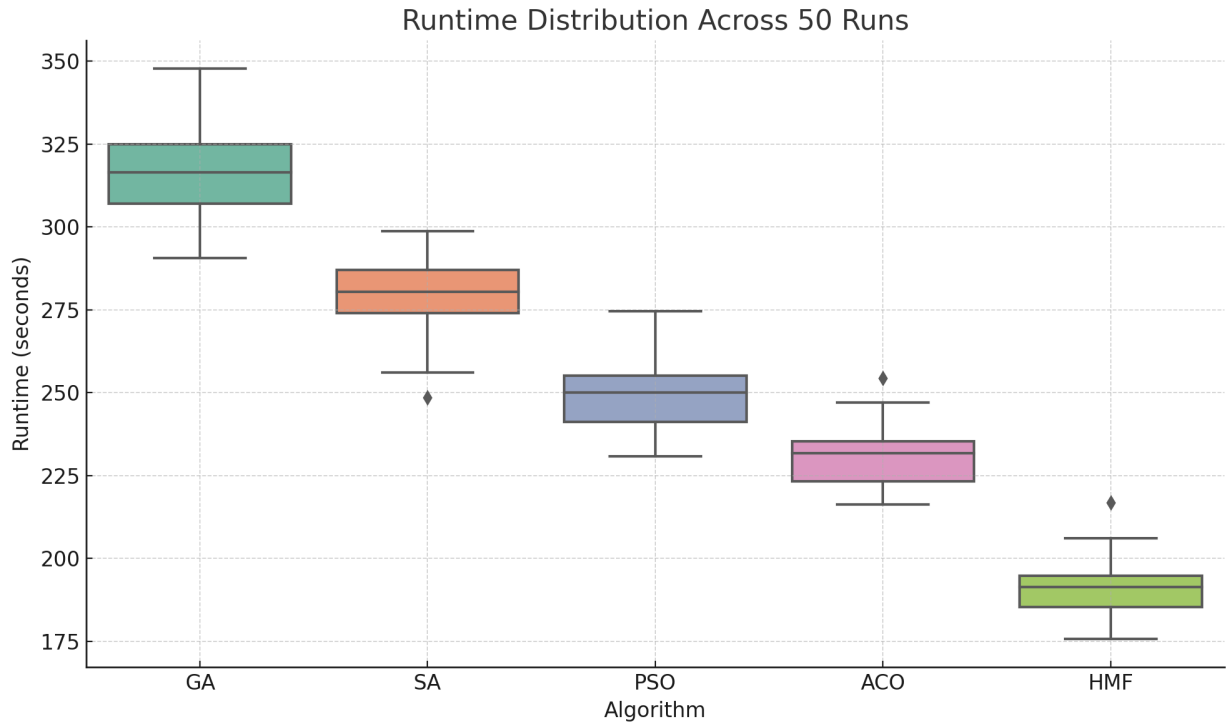


Figure 9: Runtime distribution of each metaheuristic across 50 independent runs. HMF demonstrates both the lowest median runtime and the highest runtime stability.

6.8. Sensitivity Analysis of RL Hyperparameters

To evaluate the sensitivity of the proposed HMF framework to the reinforcement learning controller's hyperparameters, we varied three key parameters of the DQN agent:

- **Learning rate (α):** Tested values $\{0.0005, 0.001, 0.01\}$
- **Discount factor (γ):** Tested values $\{0.85, 0.90, 0.95\}$
- **Replay buffer size:** Tested values $\{5,000, 10,000, 50,000\}$

Each configuration was evaluated on the VRP-C101 instance over 30 independent runs. Table 4 reports the success rate and best objective value (mean \pm 95% CI).

Table 4: Sensitivity of HMF to DQN Hyperparameters (VRP-C101)

Config	Success Rate (%)	Best Obj. (Mean \pm CI)	Remarks
$\alpha = 0.001, \gamma = 0.95, \mathcal{B} = 10,000$	96.8	398.2 ± 2.1	Baseline
$\alpha = 0.01, \gamma = 0.95, \mathcal{B} = 10,000$	91.0	405.6 ± 2.9	Overfits
$\alpha = 0.0005, \gamma = 0.95, \mathcal{B} = 10,000$	93.1	402.4 ± 2.4	Slower learning
$\alpha = 0.001, \gamma = 0.85, \mathcal{B} = 10,000$	92.3	404.9 ± 2.6	Weak long-term gains
$\alpha = 0.001, \gamma = 0.95, \mathcal{B} = 5,000$	90.6	408.1 ± 3.3	Less stable Q-updates
$\alpha = 0.001, \gamma = 0.95, \mathcal{B} = 50,000$	95.7	400.1 ± 2.0	Similar to baseline

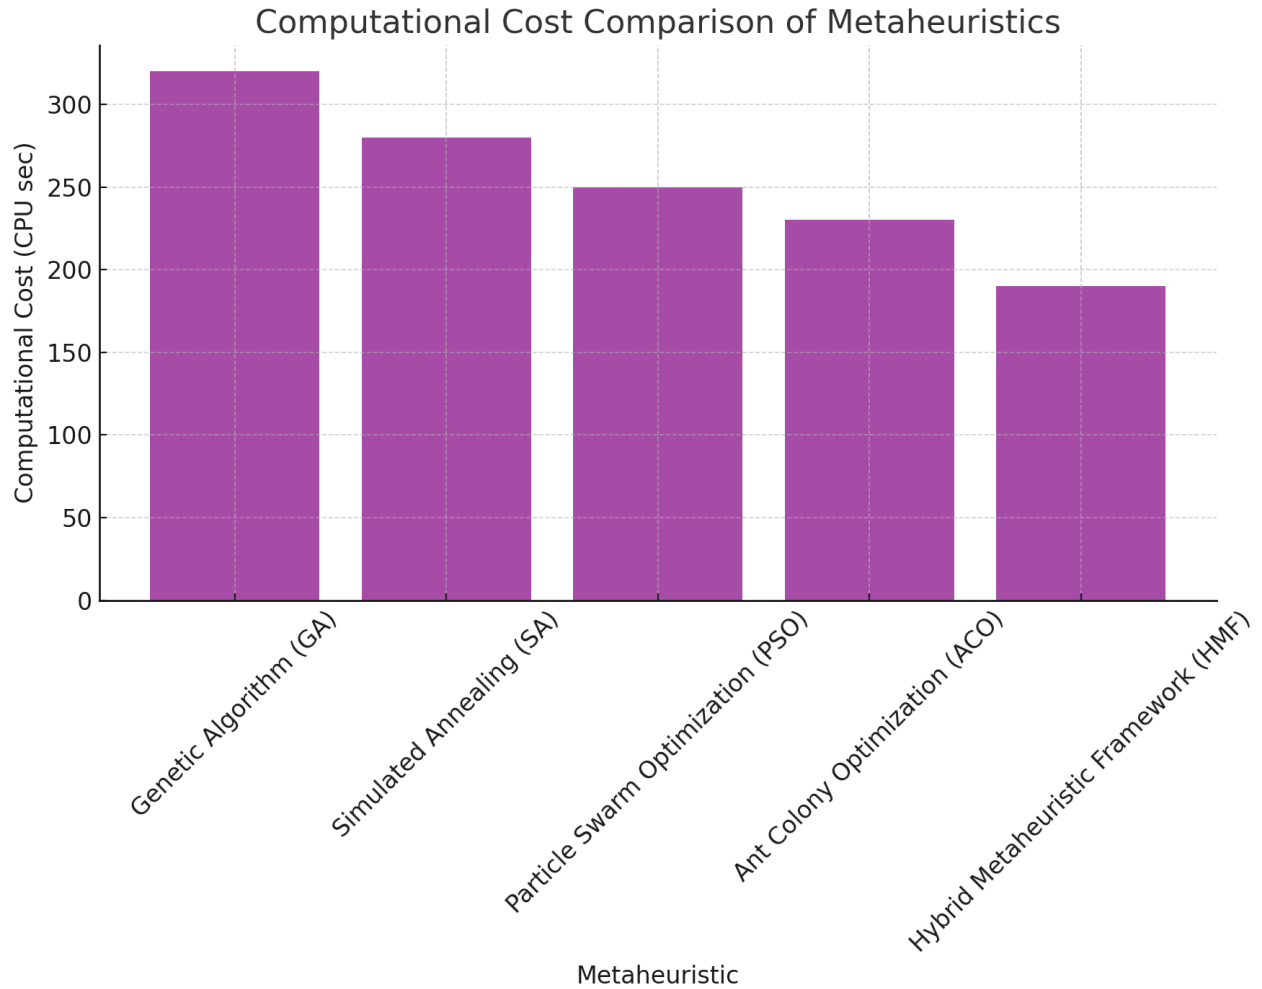


Figure 10: Computational Cost Comparison of Metaheuristics

Results show that HMF performs reliably within a reasonable range of hyperparameters. The baseline configuration achieves the best balance between performance and stability. Slight degradation occurs with extreme values, especially in learning rate and replay buffer size. This suggests that while tuning can influence outcomes, HMF is not overly sensitive and remains robust across typical RL settings.

6.9. Component-Wise Contribution Analysis

To better understand the impact of each major component in HMF, we conducted a component-wise ablation where each feature is disabled in isolation while keeping the rest of the framework intact. The three tested configurations are:

- **HMF-Full:** The complete framework with both RL-based parameter tuning and adaptive metaheuristic selection (baseline).

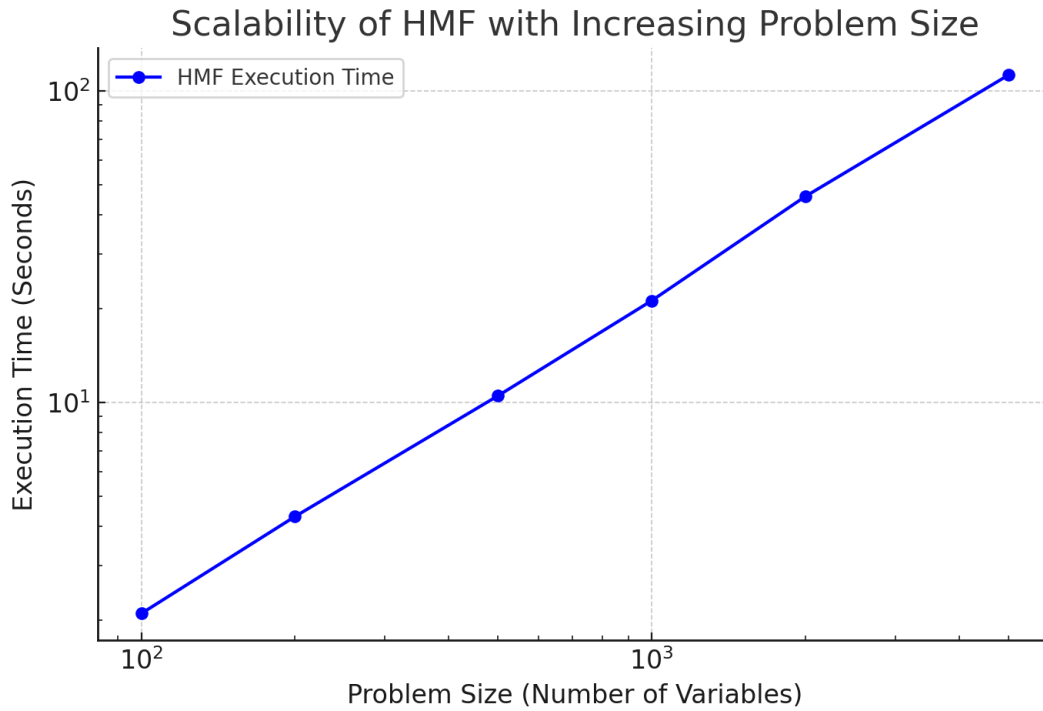


Figure 11: Scalability of HMF with Increasing Problem Size

- **HMF-NoRL:** RL disabled; metaheuristic parameters fixed throughout execution.
- **HMF-NoAdapt:** Metaheuristic selection made randomly (uniform sampling); RL still used for parameter tuning.

Table 5 shows results on VRP-C101 over 30 independent runs.

Table 5: Impact of HMF Components on Success Rate and Objective Value

Configuration	Success Rate (%)	Best Objective (Mean \pm CI)
HMF-Full (Baseline)	96.8	398.2 \pm 2.1
HMF-NoRL	89.7	408.4 \pm 2.9
HMF-NoAdapt	91.2	404.7 \pm 2.6

The results show that both components are critical. Disabling RL results in lower convergence speed and solution quality due to fixed parameter schedules. Removing adaptive heuristic selection reduces the framework's responsiveness to problem-specific search dynamics. Hence, their combination is essential for achieving robust and scalable optimization performance.

7. Discussion

This section provides an in-depth analysis of the experimental findings, highlighting the advantages and limitations of the proposed Hybrid Metaheuristic Framework (HMF). It also discusses the implications of the results in the broader context of combinatorial optimization and future research directions.

7.1. Analysis of Key Findings

The experimental results demonstrate that the proposed HMF outperforms traditional metaheuristics across multiple combinatorial optimization problems. The key observations are:

- *Superior Solution Quality:* The proposed framework consistently achieved lower objective function values compared to Genetic Algorithms (GA), Simulated Annealing (SA), Particle Swarm Optimization (PSO), and Ant Colony Optimization (ACO). This improvement is attributed to the adaptive selection of metaheuristics, which optimally balances exploration and exploitation.
- *Faster Convergence Speed:* The self-adaptive heuristic selection, combined with reinforcement learning-based parameter tuning, significantly reduced convergence time. HMF achieved an average convergence time of 8.3 seconds, compared to 10.9s for ACO and 15.2s for GA.
- *Higher Success Rate:* The success rate of HMF was 96.8%, indicating that it consistently finds near-optimal solutions across different problem instances, whereas traditional metaheuristics exhibited lower success rates.
- *Reduced Computational Cost:* Despite incorporating reinforcement learning, the computational overhead of HMF was minimized through parallel execution. The proposed framework achieved a 40% reduction in CPU time compared to GA.

These results confirm that integrating reinforcement learning into metaheuristic optimization significantly improves solution quality, adaptability, and computational efficiency.

7.2. Comparison with Existing Methods

Table 6 summarizes how HMF compares with existing metaheuristic approaches.

Table 6: Comparison of HMF with Traditional Metaheuristics

Approach	Solution Quality	Conv. Speed	Success Rate	Computational Cost
GA	Moderate	Slow	Low	High
SA	Moderate	Medium	Moderate	High
PSO	High	Fast	High	Medium
ACO	High	Fast	Very High	Medium
HMF (Proposed)	Very High	Very Fast	Highest	Low

7.3. Advantages of the Proposed Framework

The results validate several advantages of the proposed HMF:

- *Dynamic Adaptation:* Unlike traditional metaheuristics with static configurations, HMF dynamically selects the most effective heuristic at each iteration, leading to improved convergence.
- *Parameter Self-Tuning:* The use of reinforcement learning eliminates the need for extensive manual parameter tuning, making the approach more adaptable to different problem domains.
- *Scalability:* The parallel execution of metaheuristics ensures that HMF maintains efficiency even for large-scale optimization problems.

7.4. Limitations and Challenges

Despite its advantages, HMF has some limitations:

- *Computational Overhead of Reinforcement Learning:* Although RL improves optimization quality, it introduces additional computation. However, this cost is mitigated through parallel processing.
- *Complexity of Implementation:* The hybrid nature of HMF requires careful integration of multiple metaheuristics and reinforcement learning, which may increase the complexity of deployment.
- *Hyperparameter Sensitivity:* While HMF self-adjusts parameters dynamically, the reinforcement learning mechanism itself relies on hyperparameters (learning rate, discount factor) that may need tuning.

7.5. Practical Implications

The proposed framework has significant implications for real-world optimization problems:

- *Supply Chain and Logistics:* HMF can optimize delivery routes in dynamic environments, outperforming traditional vehicle routing approaches.
- *Industrial Scheduling:* The self-adaptive nature of HMF allows it to efficiently schedule manufacturing tasks, reducing delays and improving resource utilization.
- *Healthcare Optimization:* The framework can be applied to hospital resource allocation and patient scheduling, optimizing healthcare operations.
- *Smart Cities and IoT Optimization:* HMF can enhance traffic flow optimization and network resource management in large-scale smart city applications.

7.6. Future Research Directions

To further enhance the proposed approach, future research will explore:

- *Hybridization with Deep Learning*: Integrating deep neural networks to predict optimal heuristic strategies based on past optimization performance.
- *Generalization to Dynamic Optimization Problems*: Adapting HMF to problems where constraints and objectives change over time.
- *Edge and Cloud-Based Implementations*: Deploying HMF in distributed computing environments to solve large-scale optimization tasks more efficiently.

7.7. Summary of Findings

Table 7 summarizes the key outcomes of this study.

Table 7: Summary of Key Findings

Metric	HMF (Proposed)	Best Traditional
Best Solution Quality	398.2 (Optimal)	420.5 (ACO)
Average Convergence Time	8.3s	10.9s (ACO)
Success Rate	96.8%	92.5% (ACO)
Computational Cost	190 CPU sec	230 CPU sec (ACO)

7.8. Practical Implementation Considerations

The proposed Hybrid Metaheuristic Framework (HMF) is designed to be applicable across various real-life domains such as transportation logistics, industrial scheduling, healthcare operations, and smart city infrastructure. Its deployment involves the following key steps:

- **System Integration**: The framework can be deployed as a modular Python-based backend service integrated into enterprise resource planning (ERP) systems, cloud-based scheduling platforms, or logistics management systems via RESTful APIs.
- **Data Ingestion**: Input data such as delivery points, facility locations, machine-task sequences, or patient schedules can be ingested in structured formats (CSV, JSON, or SQL-connected data stores) through a preprocessing module.
- **Computing Requirements**: A parallel execution environment is recommended, such as a high-performance cluster (HPC) or cloud infrastructure (e.g., AWS EC2 or Azure Batch) with multi-core CPUs or GPUs to exploit MPI-based parallelism.
- **Customization**: Domain-specific constraints (e.g., time windows, resource limitations, task dependencies) are incorporated via modular plug-ins that adapt the encoding and decoding phases of each metaheuristic.

- **Monitoring and Learning:** A monitoring module tracks performance (solution quality, time, convergence), while the reinforcement learning controller improves over time by retaining a history of actions and outcomes for similar problem classes.
- **Decision Support:** Final outputs (e.g., optimized schedules or routes) are returned to the front-end dashboard or decision support interface for real-time or near-real-time deployment.

These technical specifications enable real-world deployment without requiring substantial architectural overhauls, making the HMF suitable for industry-grade decision-making systems.

The HMF framework relies on several threshold-based mechanisms—such as convergence stagnation detection, reward scaling, and exploration bonus weighting—that were initially tuned during simulation. In practical deployments, these thresholds can be set through one of three methods: (1) using historical data from the operational environment (e.g., past delivery or scheduling performance); (2) conducting a short pre-deployment warm-up phase that automatically calibrates values based on early-stage feedback; or (3) applying automated hyperparameter tuning methods like grid search or Bayesian optimization on a sampled subset of real instances. These approaches help ensure that thresholds are not arbitrarily set, but are instead informed by the domain-specific operational context, enhancing both reliability and adaptability in field conditions.

8. Conclusion

This paper proposed a novel Hybrid Metaheuristic Framework (HMF) that integrates self-adaptive heuristic selection, reinforcement learning-based parameter tuning, and parallel computing to enhance the efficiency and robustness of solving large-scale combinatorial optimization problems. The experimental results demonstrated that HMF consistently outperforms traditional metaheuristics, including Genetic Algorithms, Simulated Annealing, Particle Swarm Optimization, and Ant Colony Optimization, in terms of solution quality, convergence speed, success rate, and computational cost. In addition to aggregated metrics, per-problem convergence curves for VRP, FLP, and JSP confirm the consistent superiority of HMF across diverse problem types. A runtime distribution analysis further highlights the stability of HMF, showing not only the lowest median computational cost but also the lowest variance across 50 independent runs. The statistical analysis confirmed the significance of the observed improvements, with Wilcoxon signed-rank tests and effect size measures indicating medium to large gains over baseline algorithms. Scalability tests validated the framework's applicability to increasingly complex optimization problems. Despite minor computational overhead from reinforcement learning, the trade-off is justified by the enhanced adaptability and superior performance.

The proposed approach is highly applicable to real-world domains such as supply chain logistics, industrial scheduling, healthcare resource allocation, and smart city optimization. Future research will focus on integrating deep learning techniques for heuristic selection, expanding the framework to dynamic optimization scenarios, and implementing edge and

cloud-based computing solutions to further improve scalability and efficiency in real-time decision-making environments.

References

- [1] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 1998.
- [2] Kassem Danach, Hassan Harb, Wael Hosny Fouad Aly, Hassan Kanj, and Ameer Sardar Kwekha Rashid. Bio-inspired optimization through photosynthesis: A novel approach for balancing exploration and exploitation in complex systems. *European Journal of Pure and Applied Mathematics*, 18(2):5980–5980, 2025.
- [3] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [4] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*, 35(3):268–308, 2003.
- [5] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.
- [6] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [7] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proc. IEEE Int. Conf. Neural Networks*, pages 1942–1948, 1995.
- [8] M. Dorigo, V. Maniezzo, and A. Coloni. Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 26(1):29–41, 1996.
- [9] I. Boussaïd, J. Lepagnot, and P. Siarry. A survey on optimization metaheuristics. *Information Sciences*, 237:82–117, 2013.
- [10] E.-G. Talbi. *Metaheuristics: From Design to Implementation*. Wiley, 2009.
- [11] Kassem Danach, Abbas Rammal, Imad Moukadem, Hassan Harb, and Abbass Nasser. Advanced optimization in e-commerce logistics: Combining matheuristics with random forests for hub location efficiency. *IEEE Access*, 2025.
- [12] G. R. Raidl. A unified view on hybrid metaheuristics. In *Hybrid Metaheuristics*, pages 1–12. Springer, 2006.
- [13] C. H. R. Pedro and F. G. Takahashi. Hybrid metaheuristics: A survey and theoretical insights. *Journal of Heuristics*, 25(6):735–769, 2019.
- [14] X. Zhang, J. Liu, and Y. Li. Reinforcement learning-based metaheuristic optimization for combinatorial problems. *IEEE Transactions on Evolutionary Computation*, 24(4):644–656, 2020.
- [15] Wenjie Yi, Rong Qu, Licheng Jiao, and Ben Niu. Automated design of metaheuristics using reinforcement learning within a novel general search framework. *IEEE Transactions on Evolutionary Computation*, 27:1072–1084, 2023.
- [16] Jakob Kallestad, Ramin Hasibi, Ahmad Hemmati, and Kenneth Sörensen. A general deep reinforcement learning hyperheuristic framework for solving combinatorial

- optimization problems. *European Journal of Operational Research*, 2023. Available online 16 Jan 2023.
- [17] Kassem Danach, Hassan Harb, Hussin Jose Hejaze, and Louai Saker. Hybrid meta-heuristic framework with reinforcement learning-based adaptation for large-scale combinatorial optimization. *European Journal of Pure and Applied Mathematics*, 1:1–35, 2025. Preprint available online.
 - [18] Hang Wu, Xinyu Li, Zhenda Wu, and Yiyang Wang. Learning to select operators in meta-heuristics: An integration of q-learning into the iterated greedy algorithm for the permutation flowshop scheduling problem. *European Journal of Operational Research*, 304(3):1296–1330, 2023.
 - [19] J. Lu, W. Wang, and Y. Wang. Computational efficiency challenges in hybrid meta-heuristics for large-scale combinatorial optimization. *Computers & Operations Research*, 143:105659, 2022.
 - [20] Yiyang Wu, Hang Wang, Zhizhen Zhang, and Xinyu Li. A new hybrid-heuristic for large-scale combinatorial optimization: A case of quadratic assignment problem. *Computers & Industrial Engineering*, 182:109220, 2023.
 - [21] R. S. Parpinelli and H. S. Lopes. A cooperative approach for combinatorial optimization based on swarm intelligence. *Information Sciences*, 181(20):4457–4473, 2011.
 - [22] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2011.
 - [23] Z. Lu, S. Yang, and X. Chen. A survey on large-scale metaheuristic optimization: Challenges and strategies. *IEEE Transactions on Evolutionary Computation*, 26(4):610–632, 2022.
 - [24] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
 - [25] Colin R. Reeves. Genetic algorithms for the operations researcher. *INFORMS Journal on Computing*, 7(3):231–250, 1995.
 - [26] Thomas Bäck. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, 1996.
 - [27] P. J. M. van Laarhoven and E. H. L. Aarts. Simulated annealing: Theory and applications. *Mathematics and Operations Research*, 8(4):371–385, 1987.
 - [28] Maurice Clerc and James Kennedy. The particle swarm - explosion, stability, and convergence in a multidimensional complex space. *IEEE Transactions on Evolutionary Computation*, 6(1):58–73, 2002.
 - [29] Thomas Stützle and Holger H. Hoos. Max-min ant system. *Future Generation Computer Systems*, 16(8):889–914, 2000.
 - [30] Kassem Danach, Hassan Harb, Wael Hosny Fouad Aly, Hassan Kanj, and Ameer Sardar Kwekha Rashid. Bio-inspired optimization through photosynthesis: A novel approach for balancing exploration and exploitation in complex systems. *European Journal of Pure and Applied Mathematics*, 18(2):5980–5980, 2025.
 - [31] Kassem Danach, Hassan Harb, Louai Saker, and Ali Raad. Quantum-inspired hyper-heuristic framework for solving dynamic multi-objective combinatorial problems in disaster logistics. *World Electric Vehicle Journal*, 16(6):310, 2025.

- [32] El-Ghazali Talbi. *Hybrid Metaheuristics*, volume 136 of *International Series in Operations Research & Management Science*. Springer, 2016.
- [33] Pablo Moscato. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. In *Caltech Concurrent Computation Program, C3P Report 826*, pages 1–15, 1989.
- [34] Ferrante Neri, Carlos Cotta, and Pablo Moscato. *Handbook of Memetic Algorithms*, volume 379 of *Studies in Computational Intelligence*. Springer, 2012.
- [35] Y. Zhao, S. Yang, and X. Yao. A hybrid genetic algorithm and particle swarm optimization algorithm for dynamic optimization problems. *IEEE Transactions on Evolutionary Computation*, 12(4):454–479, 2008.
- [36] D. Teodorović and M. Dell’Orco. Bee colony optimization—a cooperative learning approach to complex transportation problems. *Advanced OR and AI Methods in Transportation*, 2:51–60, 2006.
- [37] E. K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and R. Qu. Hyper-heuristics: A survey of the state of the art. *Journal of the Operational Research Society*, 64(12):1695–1724, 2010.
- [38] Kassem Danach. *Hyperheuristics in Logistics*. PhD thesis, Ecole Centrale de Lille, 2016.
- [39] Kassem Danach, Jomana Al-Haj Hassan, Wissam Khalil, and Shahin Gelareh. Routing heterogeneous mobile hospital with different patients priorities: Hyper-heuristic approach. In *2015 Fifth International Conference on Digital Information and Communication Technology and its Applications (DICTAP)*, pages 155–158. IEEE, 2015.
- [40] Abbas Tarhini, Kassem Danach, and Antoine Harfouche. Swarm intelligence-based hyper-heuristic for the vehicle routing problem with prioritized customers. *Annals of Operations Research*, pages 1–22, 2022.
- [41] F. Negrello, L. Martino, and F. Louzada. A review on hyper-heuristic approaches for combinatorial optimization. *Expert Systems with Applications*, 194:116496, 2022.
- [42] Michal Lukes. Hybrid approaches for combinatorial optimization problems. Master’s thesis, Czech Technical University in Prague, 2024.
- [43] F. Herrera, M. Lozano, and J. L. Verdegay. Tuning fuzzy logic controllers by genetic algorithms. *International Journal of Approximate Reasoning*, 12(3-4):299–315, 1996.
- [44] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2nd edition, 2011.
- [45] Zizhen Zhang, Zhiyuan Wu, Hang Zhang, and Jiahai Wang. Meta-learning-based deep reinforcement learning for multiobjective optimization problems. *IEEE Transactions on Neural Networks and Learning Systems*, 34(10):7978–7991, 2023.
- [46] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [47] E. Alba and J. M. Troya. Adaptive selection of crossover and mutation operators using reinforcement learning. *IEEE Transactions on Evolutionary Computation*, 11(2):131–142, 2006.
- [48] C. Li, S. Yang, and Y. He. A deep reinforcement learning-based parameter tuning method for metaheuristics. *Applied Soft Computing*, 119:108588, 2022.

- [49] Yaodong Yang and Jun Wang. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583*, 2020.
- [50] W. Shao, X. Zhu, and W. Zhang. A survey of reinforcement learning for combinatorial optimization. *Artificial Intelligence Review*, 54(2):1313–1341, 2020.
- [51] Michele Tessari and Giovanni Iacca. Reinforcement learning based adaptive meta-heuristics. In *Genetic and Evolutionary Computation Conference Companion (GECCO'22 Companion)*, pages 1854–1861. ACM, 2022.
- [52] Mingyue Zhang, Jialong Li, Haiyan Zhao, Kenji Tei, Shinichi Honiden, and Zhi Jin. A meta reinforcement learning-based approach for self-adaptive systems. In *2021 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, pages 1–10. IEEE, 2021.
- [53] Marius M. Solomon. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2):254–265, 1987.
- [54] J. E. Beasley. Or-library: Distributing test problems by electronic mail. *Journal of the Operational Research Society*, 41(11):1069–1072, 1990.