Proceedings of the Thirty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2026)

H-ANTS: Hierarchical Ant System with Insert-and-Prune Charging for Capacitated Electric Vehicle Routing

Chu-Yin Peng, Feng-Feng Wei, Wei-Neng Chen

School of Computer Science and Engineering, South China University of Technology, China
pcy_in@163.com, fengfeng_scut@163.com, cwnraul634@aliyun.com

Abstract

The rising popularity of electric vehicles (EVs) has led to the emergence of the capacitated electric vehicle routing problem (CEVRP), where routing decisions and charging schedules are tightly coupled. Existing single-level models suffer from complex search spaces, while decoupled models frequently yield suboptimal solutions due to myopic decision-making and weak coordination. To address these limitations, this study proposes the hierarchical ant system (H-ANTS), a hierarchical algorithm that systematically enhances the decoupled framework. At the high-level abstract planning stage, an improved max-min ant system (MMAS) is combined with a batch-improvement variable neighborhood search (BI-VNS) to generate high-quality abstract customer sequences. Subsequently, at the low-level plan refinement stage, an insert-prune heuristic (IPH) is employed, utilizing a saturate-and-prune mechanism to globally optimize charging schedules, effectively overcoming the limitations of traditional greedy repairs. The cost of the refined solution acts as a topology-aware feedback signal to guide high-level pheromone updates, enabling routing decisions to adapt effectively to energy constraints. Experimental results on benchmark instances demonstrate that H-ANTS achieves superior solution quality compared with several representative algorithms, confirming the effectiveness and practicality of the proposed method.

Introduction

The vehicle routing problem (VRP) and its capacitated variant (CVRP) are canonical NP-hard challenges at the intersection of combinatorial optimization and automated planning (Gendreau et al. 2008). They serve as fundamental models for multi-agent planning tasks, requiring a fleet of agents (vehicles) to traverse routes and service customers to satisfy all customer demands under cargo capacity constraints.

With the growing emphasis on sustainability, the Electric Vehicle Routing Problem (EVRP), and specifically the capacitated EVRP (CEVRP) addressed in this study, introduces new dimensions of complexity. CEVRP extends CVRP by incorporating battery capacity limits, energy consumption along arcs, and charging operations at a limited set of stations. This creates the central challenge of CEVRP: the tight coupling between two core decisions: 1) combinatorial path planning (determining the customer service sequence) and 2) numerical resource scheduling (deciding when and where to recharge) (Yang et al. 2025).

This coupling significantly increases problem difficulty, and existing approaches can be broadly classified into two categories. Exact methods, such as branch-and-price (Nafstad, Desaulniers, and Stålhane 2025), can find optimal solutions for small-scale instances but struggle with scalability. Consequently, metaheuristic approaches are dominant (Wang et al. 2024). To apply them, researchers typically adopt one of the two main modeling strategies. The first is the single-level model. This strategy formulates the problem on a single, unified graph where customers and charging stations are treated as interchangeable nodes. In practice, this requires expanding the search space representation, such as the construction graph in ant colony optimization (ACO) (Lin, Liu, and Mei 2022) or the chromosome in genetic algorithms (GA) (Niu, Li, and Yi 2024)—to include both customers and charging decisions. This requires the algorithm to select from both types of nodes during route construction, creating an enormous and complex search space in which the combinatorial (routing) and numerical (charging) decisions constantly interfere, hindering efficient exploration. The second strategy is the decoupled model. This method, often presented as a bilevel or two-stage algorithm, splits the problem: a high-level planner generates an abstract customer-only sequence, and a low-level heuristic then repairs this sequence by inserting charging stations. While this simplifies the search, existing implementations often suffer from critical limitations at both levels. At the low-level, the repair strategy is frequently a myopic constructive heuristic —–e.g., a reactive "charge-only-when-necessary" policy, which often leads to feasible but highly suboptimal routes (Nie et al. 2022). Meanwhile, high-level planners such as classical ACO often rely on incremental neighborhood updates and can suffer from oscillations, failing to find a sufficiently high-quality abstract plan (Dorigo and Stützle 2018). Furthermore, most decoupled designs lack a feedback mechanism, creating a one-way information flow in which routing decisions cannot adapt to the true energy constraints revealed during the repair phase.

To address these limitations, this study proposes the hierarchical ant system (H-ANTS). Inspired by the hierarchical planning philosophy from the AI planning community (Kuroiwa and Beck 2023), H-ANTS employs a hybrid algorithm that decomposes the problem into two synergistic levels: high-level abstract planning and low-level plan refinement. The main contributions of this paper are as follows:

1. A closed-loop hierarchical framework that overcomes the myopic limitations of traditional two-stage models. H-ANTS establishes a bidirectional feedback mechanism where the refined low-level costs guide the high-level pheromone updates. This closed-loop interaction enables routing decisions to adapt to true energy constraints, avoiding the one-way, short-sighted behavior of existing decoupled methods.
2. A hybrid MMAS–BI-VNS planner that stabilizes global exploration and accelerates convergence. The high level integrates an improved MMAS with a batch-improvement VNS operator. BI-VNS performs batch-level refinements, providing a stabilizing corrective influence that counteracts the MMAS oscillation and improves the quality and robustness of abstract customer sequences.
3. A global insert–prune heuristic (IPH) that replaces traditional greedy charging repairs. The low-level IPH uses a two-phase saturate-and-prune strategy, i.e., feasibility is first ensured by saturation, followed by global pruning for optimization. This non-myopic mechanism produces well-optimized charging schedules and complements the high-level combinatorial search.

Across public benchmark instances, H-ANTS achieves competitive and often superior performance compared with well-established methods, demonstrating its practical effectiveness for planning problems with tightly coupled constraints.

Related Work

Research on the CEVRP can be viewed from two complementary perspectives: the traditional operations research (OR) perspective, which treats it as a combinatorial optimization problem, and the automated planning and scheduling perspective, which is the focus of this study.

From the OR perspective, the CEVRP is NP-hard, and existing solution methods can be broadly grouped into two categories. Exact methods primarily solve the EVRP by transforming it into a mixed-integer linear programming (MILP) model. For example, Tahami et al. proposed a branch-and-cut algorithm (Tahami, Rabadi, and Haouari 2020), Ceselli et al. introduced a branch-and-price algorithm based on path generation (Ceselli et al. 2021), and Lee developed a branch-and-price method for small-scale EVRPs (Lee 2021). Wu and Zhang addressed a bilevel EVRP using branch-and-price methods, further extending its applicability (Wu and Zhang 2023). While these methods can produce high-quality solutions for small-scale problems, they are inefficient for large-scale instances due to high computational complexity. Consequently, the research focus has predominantly been on metaheuristics for their ability to efficiently explore large solution spaces and avoid being trapped in local optima, such as genetic algorithms (GA) (**?**), and iterated local search

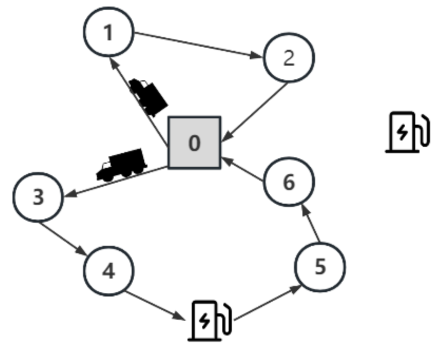

Figure 1: An example solution of CEVRP with two routes

(ILS) (Chen et al. 2024b). As a form of multi-agent swarm intelligence, ant colony optimization (ACO) has been a particularly popular choice for this class of problems (Jia, Mei, and Zhang 2021). This aligns with broader trends in the confluence of evolutionary computation and multi-agent systems (Chen et al. 2025), where recent work has focused on enhancing cooperative mechanisms within swarms (Chen et al. 2024a). More recently, advanced approaches have combined metaheuristics with machine learning, such as imitation improvement learning, which has shown strong performance on large-scale instances (Bui and Mai 2023). A common thread in these methods is the design of efficient neighborhood operators. Large neighborhood search (LNS) and its variants (He, de Weerdt, and Yorke-Smith 2020), for example, are often viewed as an efficient "destroy and repair" strategy for iterative improvement (Huang et al. 2022).

From an automated planning perspective, CEVRP is a complex problem of managing numeric fluents evolving over time (e.g., vehicle capacity and battery charge) (Cuchý, Vokřínek, and Jakob 2024). To tackle enormous state space, hierarchical planning is a powerful strategy from the planning community. This approach, seen in frameworks like hierarchical task networks (HTN) and other general-purpose methods (Kuroiwa and Beck 2023), decomposes abstract tasks (e.g., serve_all_customers) into concrete actions.

Definition of CEVRP

The CEVRP extends the classical Capacitated VRP by incorporating battery capacity limits and charging requirements specific to electric vehicles. The objective is to determine a set of routes that minimize the total travel distance while satisfying a series of constraints. The problem can be represented by a fully connected weighted graph $G = (V, A)$, where $V = \{0\} \cup I \cup F'$ denotes the set of nodes and $A = \{(i, j) \mid i, j \in V, i \neq j\}$ denotes the set of directed arcs. Node 0 represents the central depot, set I denotes the customers and F' the set of charging-station replica nodes. Since each physical charging station $i \in F$ may be visited multiple times, we create β_i replica nodes for station i, where β_i is a sufficiently large upper bound on the maximum number of feasible visits. The total number of replica nodes therefore satisfies $|F'| = \sum_{i \in F} \beta_i$. Each arc (i, j) has a non-negative value d_{ij} representing the Euclidean distance between nodes i and j. The battery consumption rate is h, and for each arc (i, j), the electric vehicle consumes an amount of energy $h \cdot d_{ij}$ when traversing the arc. Additionally, each customer i has a fixed demand c_i. The maximum cargo capacity for each EV is C, and the maximum battery capacity is Q. The CEVRP problem can be mathematically formulated as follows:

$$\min \sum_{i \in V} \sum_{\substack{j \in V \\ j \neq i}} d_{ij} x_{ij} \tag{1}$$

s.t.

$$\sum_{\substack{j \in V \\ j \neq i}} x_{ij} = 1, \forall i \in I \tag{2}$$

$$\sum_{\substack{j \in V \\ j \neq i}} x_{ij} \leq 1, \forall i \in F' \tag{3}$$

$$\sum_{\substack{j \in V \\ j \neq i}} x_{ij} - \sum_{\substack{j \in V \\ j \neq i}} x_{ji} = 0, \forall i \in V \tag{4}$$

$$u_j \leq u_i - c_i x_{ij} + C(1 - x_{ij}), \forall i \in V, \forall j \in V, i \neq j \tag{5}$$

$$0 \leq u_i \leq C, \forall i \in V \tag{6}$$

$$y_j \leq y_i - h d_{ij} x_{ij} + Q(1 - x_{ij}), \forall i \in A, \forall j \in V, i \neq j \tag{7}$$

$$y_j \leq Q - h d_{ij} x_{ij}, \forall i \in F' \cup \{0\}, \forall j \in V, i \neq j \tag{8}$$

$$0 \leq y_i \leq Q, \forall i \in V \tag{9}$$

$$x_{ij} \in \{0, 1\}, \forall i \in V, \forall j \in V, i \neq j \tag{10}$$

In this formulation, equation (1) defines the objective function of the CEVRP, which is to minimize the total travel distance across all traversed arcs. Constraint (2) ensures that each customer is visited exactly once. Constraint (3) ensures that each replica node in F' is visited at most once, which allows a physical charging station to be visited multiple times. Constraint (4) enforces flow conservation at every node. Constraints (5)–(6) model the load evolution along each route and ensure that vehicle capacity C is not exceeded. These two constraints ensure that the remaining capacity is non-negative when the EV reaches any node. Constraints (7) through (9) represent the energy constraints, where y_i denotes the remaining battery capacity of the EV upon arrival at node i. These three constraints ensure energy feasibility along every traversed arc. Finally, constraint (10) introduces a binary variable, which takes value 1 if arc (i, j) is included in a route, and 0 otherwise. Fig. 1 illustrates two example routes, where one vehicle performs a charging stop and the other does not.

The Proposed H-ANTS

This section describes the proposed H-ANTS algorithm designed to solve the CEVRP.

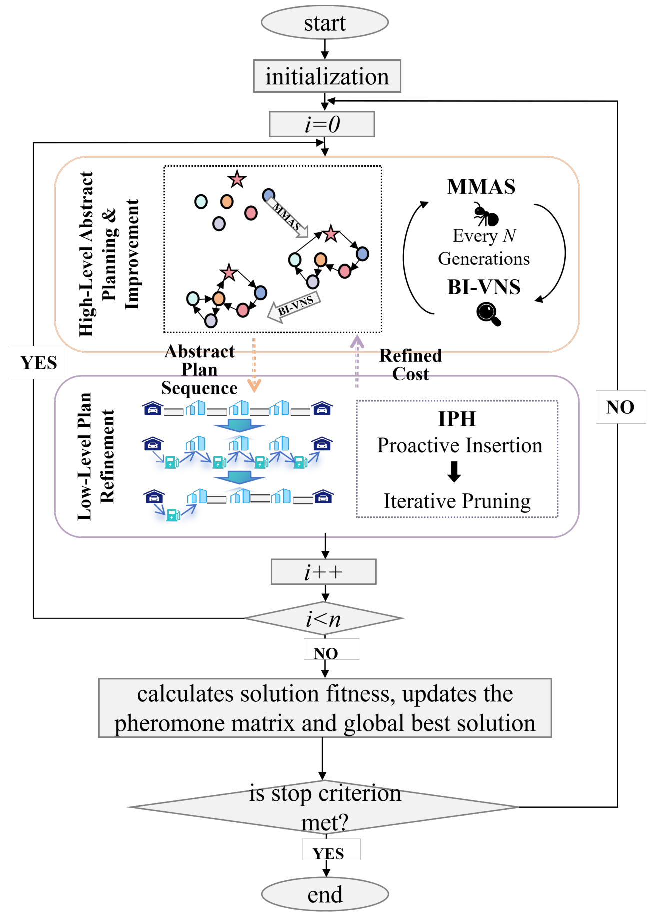

Figure 2: The framework of H-ANTS

Overall Algorithm Framework

The overall architecture of the proposed H-ANTS algorithm is illustrated in Fig. 2. The design integrates a hierarchical planning philosophy within a main iterative optimization loop and decomposes the complex CEVRP solving process into two synergistic levels: high-level abstract planning and low-level plan refinement. Following parameter initialization, the algorithm enters its primary loop, proceeding generation by generation. Structurally, the high level (MMAS + BI-VNS) aligns with the memetic algorithm (MA) paradigm (Neri and Cotta 2012). By guiding this MA structure with a hierarchical planning philosophy, the proposed approach effectively addresses the problem's coupled combinatorial-numerical nature.

Within each generation, the process begins at the high-level abstract planning and improvement stage. Here, the improved MMAS is responsible for constructing a customer visiting sequence, which serves as an abstract plan. This constructive process is complemented by the BI-VNS, which is triggered periodically (e.g., every N generations) to perform substantial, batch-wise improvements on the current best sequence, enhancing both convergence and global exploration.

The abstract plan generated at the high level is then transferred to the low-level plan refinement stage. At this stage, IPH is applied. It first ensures energy feasibility through a feasibility saturation phase and then systematically prunes redundant charging stops via an iterative pruning phase to optimize the route's total cost.

Crucially, the final cost of the complete, feasible plan refined at the low level is then used as feedback for the high level. This evaluation guides the update of the pheromone matrix and the global best solution for the next generation. This hierarchical, closed-loop process repeats until a maximum number of generations is reached, at which point the algorithm terminates and outputs the best-found solution.

High-Level Abstract Planning

Pheromone Setting MMAS is used to construct the abstract routes without considering charging constraints. It mimics the foraging behavior of real ants, where a colony of artificial agents constructs solutions probabilistically based on pheromone trails. These pheromones serve as a collective memory, reinforcing the edges that belong to high-quality solutions to guide subsequent searches. The size of the pheromone matrix is set to $n \times n$, where n equals the number of customers plus the depot node, $i.e., n = |\tilde{I}| + 1$. Each entry τ_{ij} represents the pheromone value from node i to node j. The pheromone values are bounded and satisfy $\tau_{min} \leq \tau_{ij} \leq \tau_{max}$. The boundaries are calculated as follows:

$$\tau_{max} = \frac{1}{(1-\rho) \cdot L_{opt}} \tag{11}$$

$$\tau_{min} = \tau_{max} \cdot \frac{1 - \sqrt[n]{pr}}{(n/2 - 1) \cdot \sqrt[n]{pr}} \tag{12}$$

where ρ is the parameter controlling the pheromone evaporation rate, and L_{opt} is the length of the global best solution. pr is typically set to 0.05.

Route Construction The route construction process can be divided into two subprocesses: 1) initial route generation and 2) route splitting. First, an empty route list R is initialized, and a set U containing all customers is created. A random initial node is selected, added to the route list R, and removed from set U. At each step, the next node is selected based on a probabilistic rule. The probability is calculated as follows:

$$p_{ij} = \frac{[\tau_{ij}]^{\alpha} \cdot [\eta_{ij}]^{\beta}}{\Sigma_{l \in U} [\tau_{ij}]^{\alpha} \cdot [\eta_{ij}]^{\beta}}, \quad \text{if } j \in U \tag{13}$$

where α and β are parameters balancing the relative influence of the pheromone trail and heuristic information. τ_{ij} denotes the pheromone value on edge (i, j), and $\eta_{ij} = \frac{1}{d_{ij}}$ represents the heuristic information, where d_{ij} is the length of edge (i, j). Consequently, nodes with higher pheromone intensity and shorter distances are prioritized. The selected node is appended to R and removed from U. This process repeats until U is empty, completing the construction of a giant tour.

In the route splitting phase, the initial route is partitioned into capacity-compliant subroutes. Nodes from the sequence are assigned to a vehicle until capacity constraints are violated, at which point the current route is finalized, and a new route is initiated. Following the splitting process, the 2-opt local search algorithm is applied to refine the solution. Finally, the iteration's best feasible solution is compared with the global best solution, updating the latter if an improvement is found.

Pheromone Update After the ant colony completes the route construction, the abstract plan is transferred to the low-level IPH to generate a final, feasible solution. This solution is evaluated, and its true cost is recorded. If this cost improves upon the global best solution, a pheromone update is performed. The pheromone matrix is updated based on the abstract plan that led to this new best solution:

$$\tau_{ij}(t) = (1-\rho) \cdot \tau_{ij} + \Delta\tau_{ij}^{best} + \Delta\tau^{\text{cloud}} \tag{14}$$

where $\Delta\tau_{ij}^{best} = \frac{1}{C_{\text{best}}}$ represents the pheromone deposited by the shortest path found in the current iteration, and C_{best} denotes the final, feasible cost derived from the low-level IPH. Crucially, this feedback mechanism enables the pheromone matrix to implicitly encode the underlying energy topology. Since C_{best} accounts for detour distances caused by charging, edges (i, j) that are topologically compatible with the charging network (i.e., permitting minimal-detour insertions) yield a lower C_{best} and thus receive stronger pheromone reinforcement. Over time, this transforms the pheromone matrix into a hybrid potential field, guiding high-level agents toward customer sequences that are not only spatially clustered but also EV-friendly, effectively bridging the gap between combinatorial routing and numerical resource scheduling. The $\Delta\tau^{\text{cloud}}$ term is not simple noise, but a structured increment generated by a cloud model (Li, Liu, and Gan 2009). It is designed to intelligently balance exploitation and exploration by generating dynamic, non-Gaussian perturbations based on three parameters: expectation (Ex), entropy (En), and hyper-entropy (He). By incorporating this dynamic perturbation, the algorithm's diversity and global search capability are significantly enhanced.

Iterative Plan Improvement with BI-VNS

At the high-level of H-ANTS, the BI-VNS serves as the primary engine for iterative plan improvement. While conceptually related to the LNS strategy, which is often employed for plan repair, BI-VNS introduces a significant innovation over standard VNS implementations. In contrast to traditional local search strategies (e.g., standard best-improvement) that typically commit to a single move per iteration, our BI-VNS employs a batch-improvement strategy. This mechanism identifies and applies multiple non-conflicting high-quality moves across disjoint routes simultaneously, thereby significantly accelerating the convergence toward high-quality solutions.

Neighborhood Structures BI-VNS leverages a diverse set of neighborhood structures, serving as the basis for both the shaking and Llocal search phases. These operators are grouped into two strategic categories. Intra-route operators, Algorithm 1: Local Search Phase of BI-VNS

Input: Initial plan Γ_{best}, Neighborhoods $\{N_1, ..., N_{k_{\max}}\}$
Output: Improved plan Γ^*

$\Gamma^* \leftarrow \text{InitiaLocalSearch}(\Gamma_{\text{best}})$ //Initialization
$k \leftarrow 1$
while $k \le k_{\max}$ **do**
  $Moves \leftarrow \text{FindBestMoves}(\Gamma^*, N_k)$
  //Systematically evaluate neighborhood
  if $ImprovementMoves$ is not empty **then**
    $\Gamma' \leftarrow \text{ApplyBatch}(\Gamma^*, ImprovementMoves)$ //Apply all found moves
    if $f(\Gamma') < f(\Gamma^*)$ **then**
      $\Gamma^* \leftarrow \Gamma'$
      $k \leftarrow 1$ //Reset to the first neighborhood
    else
      $k \leftarrow k + 1$ //Explore the next neighborhood
    end if
  else
    $k \leftarrow k + 1$
  end if
end while
return Γ^*

such as inversion (2-opt), relocation (Or-opt), and Exchange, refine a single vehicle's route by reversing subsequences, repositioning customers, or swapping nodes, thereby intensifying the search within promising regions. Inter-route operators, including cross-route relocation, exchange of customers or sub-routes, and the more sophisticated generalized insertion (GENI), restructure relationships between different routes, promoting diversification and enabling escape from local optima. This dual toolkit provides a robust balance between exploitation and exploration, ensuring both solution refinement and broad search coverage.

The BI-VNS Procedure The BI-VNS follows the standard VNS framework, alternating between a shaking phase (to escape local optima) and a local search phase (to find a new local optimum). The local search phase, which contains our core batch-improvement innovation, is detailed in Algorithm 1. This procedure is an iterative process that systematically explores neighborhoods to refine the current solution. Its distinguishing feature lies in how improvements are identified and applied. For each neighborhood structure, the algorithm evaluates all possible moves within each route and selects the best improvement. These best moves are then collected into a batch and applied simultaneously to the current plan in a single step, which constitutes the essence of BI-VNS. The resulting plan is immediately evaluated: if the total cost improves, the solution is updated and the search returns to the first neighborhood structure; otherwise, the algorithm proceeds to the next structure. This "search–collect–apply" cycle enables synergistic and accelerated improvements across the solution, in contrast to the incremental sequential single-move strategy used in standard VNS.

Low-Level Plan Refinement

Transforming abstract sequences into feasible schedules requires addressing the tight coupling between routing and charging. Traditional decoupled methods typically rely on a reactive, constructive strategy, often termed greedy repair (GR), that inserts stations only upon constraint violation. This forward-greedy paradigm is fundamentally myopic. It fails to anticipate future energy deficits, often forcing the agent into costly emergency detours to avoid depletion.

To address this, we propose the IPH. Unlike GR, IPH adopts a subtractive optimization strategy that operates in two phases:

1. Insertion Phase (Feasibility Saturation): Rather than reacting to deficits, this phase establishes a maximal feasibility envelope by proactively inserting minimal-detour stations on every route leg. This initializes the search from a globally feasible upper bound, shifting the problem from finding feasibility to optimizing redundancy.
2. Pruning Phase (Monotonic Descent): The algorithm iteratively removes the station yielding the maximal marginal distance savings while maintaining energy constraints. Analogous to sub-modular optimization, this reverse greedy process exploits diminishing marginal gains to ensure monotonic convergence to a high-quality local optimum, effectively avoiding the dead-end traps inherent to constructive repairs.

Experiments and Analyses

Experimental Setup

1) Benchmark Instances: To verify the effectiveness of the proposed solution, we conducted experiments using seven test instances from the IEEE WCCI-2020 Evolutionary Computation Benchmark for the electric vehicle routing problem (Mavrovouniotis et al. 2020). The brief information of the selected instances is shown in Table 1.

2) Algorithm Settings: To assess the performance of H-ANTS, we set a fixed execution time on each instance that allows H-ANTS to converge:

$$ExeTime = \frac{|I| + |F|}{100}(\text{hr}) \tag{15}$$

All the parameters of H-ANTS follow the canonical settings of MMAS (Stützle and Hoos 2000). The ant colony size is set to $n = |I| + 1$. The pheromone evaporation rate ρ is set to 0.98, and the parameters α and β are set to 1 and 2, respectively. The BI-VNS activation interval is set to $N = 20$. Empirical observations show that more frequent triggers (e.g., $N < 10$) disrupt MMAS pheromone accumulation, hindering effective landscape learning. Conversely, sparser intervals (e.g., $N > 50$) fail to correct structural stagnation. The chosen frequency balances exploration and exploitation, allowing sufficient pheromone accumulation before BI-VNS injects diversified structural improvements. Regarding the cloud model used for pheromone updates, we set Ex=0.5, En=0, He=0.1, and the number of cloud drops to 200. These values were determined based on preliminary experiments to ensure a stable distribution

Case	Customer	Station	Route	Index	BKS	H-ANTS		MMAS		SLMMAS	GA	EACO-F
						GR	IPH	GR	IPH			
E22	**21**	**8**	**4**	min mean std	384.67	**384.67=** **384.67=** 0.00	**384.67** **384.67** 0.00	**384.67=** 386.15- 2.56	**384.67=** 385.78- 0.98	385.44- 385.44- 0.00	**384.67=** **384.67=** 0.00	384.68- 384.68- 0.00
E23	**22**	**9**	**3**	min mean std	573.13	**571.94=** 572.16- 0.68	**571.94** **571.94** 0.00	581.07- 593.32- 13.06	**571.94=** 575.03- 4.73	582.61- 582.94- 0.60	**571.94=** **571.94=** 0.00	579.07- 584.03- 3.40
E30	**29**	**6**	**4**	min mean std	511.25	512.71- 512.94- 0.81	**509.47** **509.47** 0.00	513.11- 516.35- 5.12	511.25= 513.09- 1.85	514.42- 516.11- 0.61	**509.47=** **509.47=** 0.00	513.63- 513.63- 0.00
E33	**32**	**6**	**4**	min mean std	869.89	849.30- 855.09- 5.41	**843.3** **844.68** 0.63	856.66- 872.05- 13.61	849.69- 863.85- 12.23	859.25- 860.08- 0.74	844.25- 845.62- 0.92	847.07- 850.98- 2.22
E51	**50**	**5**	**5**	min mean std	570.17	533.61- 561.71- 11.59	**529.69** **541.89** 7.60	558.63- 579.35- 11.23	551.49- 574.85- 15.5	549.81- 564.93- 8.70	529.90- 542.08- 8.57	573.26- 575.14- 1.82
E76	**75**	**7**	**7**	min mean std	723.36	719.30- 733.86- 13.95	699.10 **714.01** 8.34	766.93- 806.34- 12.75	756.42- 791.65- 15.9	724.24- 762.09- 13.23	**697.27+** 717.30 9.58	746.79- 752.54- 8.62
E101	**100**	**9**	**8**	min mean std	899.88	853.27- 871.14- 11.93	**840.25** **850.43** 9.81	943.51- 981.15- 15.68	919.04- 949.72- 22.72	891.77- 904.62- 4.64	852.69- 872.69- 9.58	904.38- 908.39- 4.47

Table 1: Performance Comparison on Benchmark Instances

* Bold values indicate the best solution found among the compared methods. +, -, and = mean the solution is significantly better than, significantly worse than, and equal to H-ANTS-IPH, respectively.

of perturbations. The algorithm is implemented in Python and executed on an Intel (R) Core (TM) i5-10210U 1.6GHz CPU running Windows 11. Each instance was independently executed 20 times.

3) Complexity Analysis: The computational complexity of H-ANTS is dominated by its hierarchical components. The high-level MMAS operates in $O(K \cdot n^2)$, where K is the number of ants and n is the number of nodes. The low-level IPH, utilizing a subtractive strategy, is highly efficient: the feasibility saturation phase is linear $O(n)$, and the iterative pruning phase calculates savings for at most n stations over n iterations, bounded by $O(n^2)$. Consequently, the overall complexity per generation remains polynomial $O(K \cdot n^2)$, ensuring scalability comparable to standard metaheuristics while providing superior solution quality.

We compared H-ANTS with three algorithms: GA, EACO-F, and a single-level max-min ant system (SLMMAS). Among these, GA is one of the winners of the EVRP competition in IEEE WCCI-2020. EACO-F is an extended ant colony optimization (EACO) algorithm that adopts the single-level modeling strategy discussed in the introduction, using a pheromone matrix to record routing information for both customers and charging stations (Lin, Liu, and Mei 2022). SLMMAS, which similarly follows the single-level model, is implemented as described in (Jia, Mei, and Zhang 2021). In addition to these external competitors, we also implemented a baseline MMAS algorithm for our ablation study. This baseline algorithm is identical to our proposed H-ANTS but with the BI-VNS component disabled. Its purpose is to specifically isolate and validate the performance contribution of the BI-VNS component. To observe significant differences, the Wilcoxon rank-sum test was performed with a significance level of 0.05.

Performance of the Proposed H-ANTS

Table 1 compares the performance of H-ANTS with GA, EACO-F and SLMMAS, highlighting its superior solution quality. On small-scale datasets (E22, E23), H-ANTS achieved near-optimal results, significantly outperforming SLMMAS. For more complex datasets (e.g., E33, E101), H-ANTS-IPH consistently delivered higher-quality solutions, surpassing GA, EACO-F and SLMMAS. It demonstrates that our reinforced decoupled approach is a more effective modeling strategy than the single-level approach employed by EACO-F and SLMMAS. By separating the combinatorial and numerical decisions, H-ANTS avoids the complex, entangled search spaces that hinder those single-level algorithms.

These results are expected for several reasons. First, the use of BI-VNS as a high-level plan improver allows the algorithm to escape local optima and find superior abstract customer sequences. Second, the IPH provides a non-myopic, saturate-and-prune refinement strategy, capable of finding charging schedules that are significantly more efficient than those produced by simple greedy repairs. Third, the hierarchical structure connects these components effectively in a

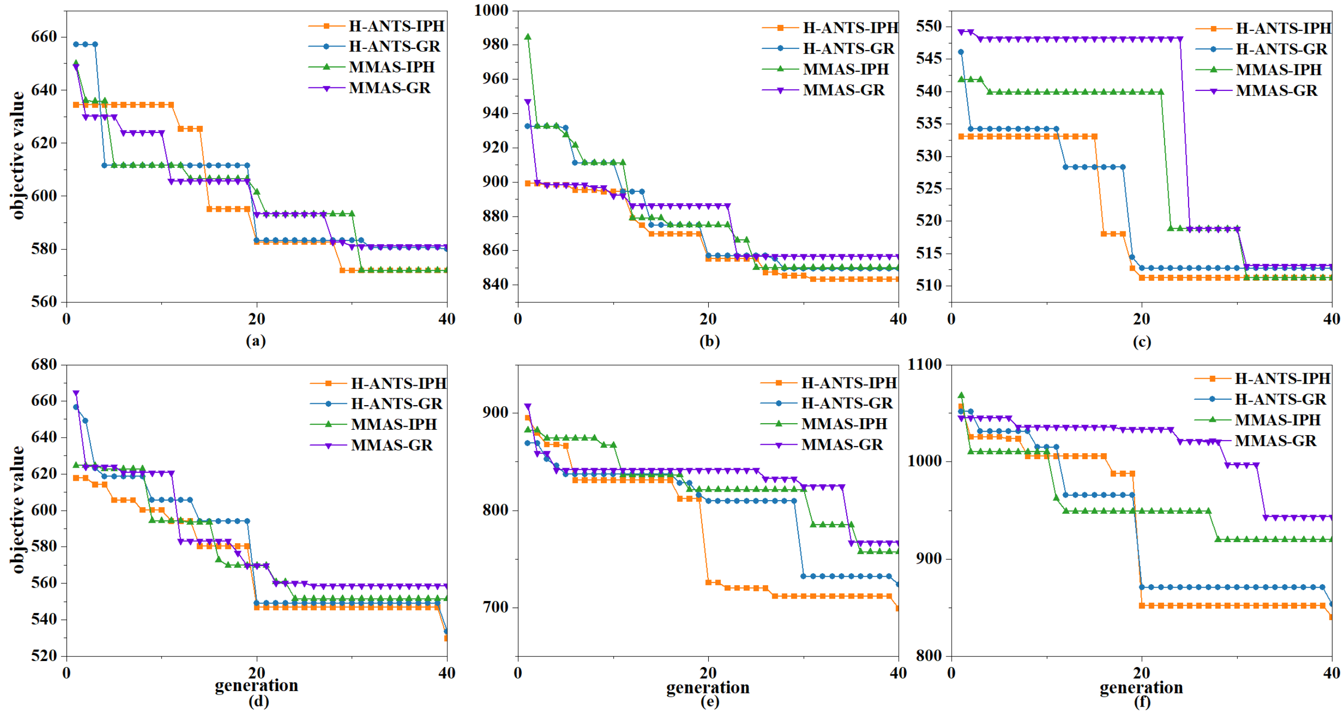

Figure 3: Convergence curves of H-ANTS-IPH, H-ANTS-GR, MMAS-IPH and MMAS-GR (a) E23. (b) E33. (c) E30. (d) E51. (e) E76. (f) E101. x-axis represents the generation. y-axis represents the objective value.

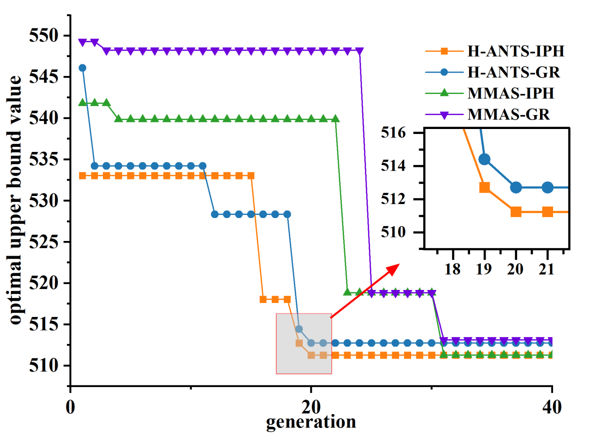

Figure 4: Impact of BI-VNS optimization at the 20th generation

closed-loop process: the IPH refines an abstract plan and its final, feasible cost is used as feedback to update both the global best solution and the high-level pheromone matrix. Collectively, these synergistic mechanisms underpin the robustness and effectiveness of H-ANTS.

Influence Investigation of BI-VNS

The contribution of the BI-VNS component is validated by comparing H-ANTS against the baseline MMAS (which has BI-VNS disabled). The convergence curves in Fig. 3 clearly illustrate this impact: while the baseline MMAS exhibits prolonged exploration with persistent oscillations, H-ANTS demonstrates significantly faster convergence. Specifically, the application of BI-VNS optimization triggers a rapid improvement in solution quality and attenuates fluctuations, thereby stabilizing convergence (Fig. 4). This stabilization effect is further confirmed by the boxplots in Fig. 6. The H-ANTS variants exhibit significantly tighter solution distributions and smaller interquartile differences, indicating high stability. In contrast, the baseline MMAS shows much greater variability and solution dispersion.

This combined evidence directly validates our motivation for the high-level planner. It confirms that the BI-VNS (with its batch-improvement strategy) successfully overcomes the limitations of traditional planners (like the baseline MMAS), which, as previously identified, often rely on incremental updates and suffer from oscillations. The synergy between ACO's intensification and BI-VNS's diversification ensures sustained convergence without premature stagnation.

Effectiveness of the IPH

To visually demonstrate the advantages and disadvantages of these two heuristic algorithms, we analyze a specific solution generated by H-ANTS on instance E-n30-k4. Initially, all charging stations in the solution are removed. Subsequently, two distinct solutions are reconstructed using the respective heuristic algorithms, as shown in Fig. 5. Regarding GR, the insertion positions are inherently suboptimal because the algorithm only recharges the EV upon imminent

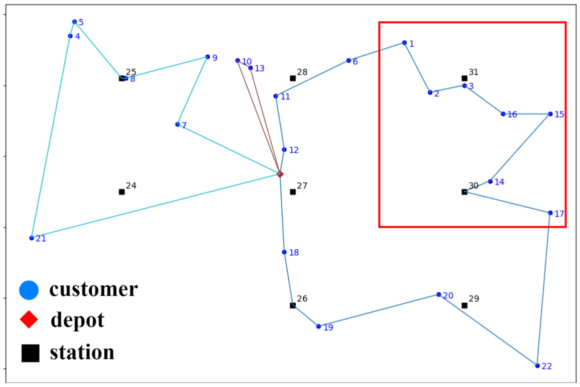

(a) GR (Constructive): Reactive insertion leads to suboptimal detour; the objective value is 579.92.

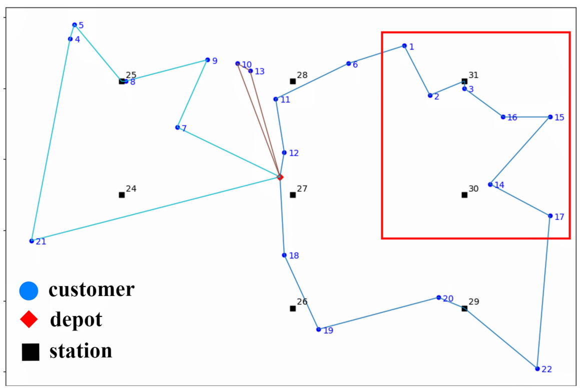

(b) IPH (Subtractive): Global pruning identifies the optimal charging node, the objective value is 571.94.

Figure 5: Comparison of Refinement Strategies

depletion. Examining the rightmost route in Fig. 5a, it is observed that the EV serves the customer directly rather than preemptively charging at the optimal node 31. Conversely, as seen in the corresponding route in Fig. 5b, IPH generates a superior recharging schedule compared to GR.

Beyond empirical case analysis, Fig. 6 and Fig. 3 comparatively evaluate performance of IPH and GR within H-ANTS. Fig. 3 shows that GR suffers from larger quality fluctuations and slower convergence, while IPH demonstrates superior stability. This robustness is confirmed by Fig. 6, which shows IPH maintaining concentrated solution distributions with minimal variability, the heuristic deficiency of GR leads to dispersed and unpredictable results. This verifies that the IPH (with its saturate-and-prune process) is a superior refinement strategy that directly addresses the critical flaw of existing decoupled models: their reliance on myopic greedy heuristics.

Conclusion

This paper proposes H-ANTS, a hierarchical planning framework for solving the Capacitated Electric Vehicle Routing Problem (CEVRP). By decomposing the problem into high-level abstract planning and low-level plan refinement, H-ANTS avoids the tightly coupled search space of

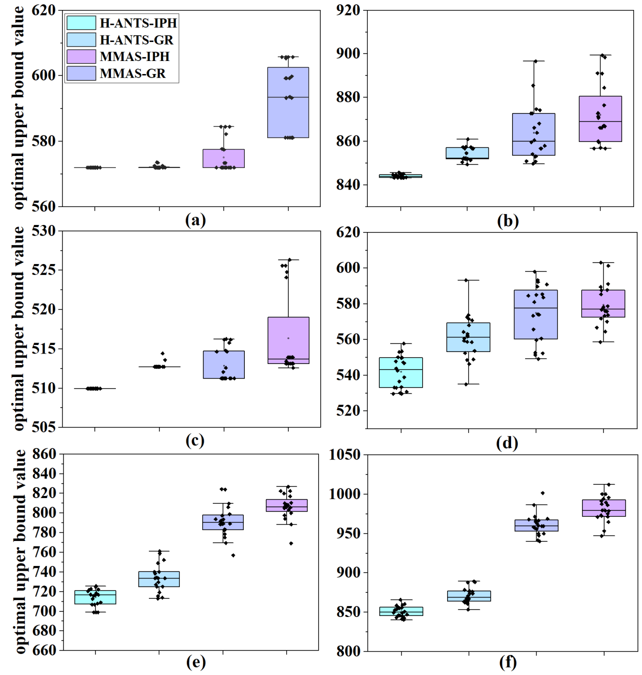

Figure 6: Boxplot comparison of all four algorithm variants (H-ANTS-IPH, H-ANTS-GR, MMAS-IPH, MMAS-GR) (a) E23. (b) E33. (c) E30. (d) E51. (e) E76. (f) E101. In (a)–(f), the median is represented by the horizontal line internal to the box. The boundary of the lower whisker is the minimum value, and the boundary of the upper whisker is the maximum value.

single-level models and enables more effective exploration. The closed-loop feedback mechanism further bridges the two levels, allowing the high-level planner to adapt according to refined execution costs. By integrating BI-VNS and the Insert-and-Prune Heuristic (IPH), the proposed method achieves competitive performance on benchmark CEVRP instances. Experimental results demonstrate the effectiveness of both components, highlighting their ability to overcome the limitations of incremental planners and greedy repair strategies. Future work will investigate the integration of machine learning techniques to better estimate execution costs and support more robust online decision-making in dynamic environments.

Acknowledgements

This work was supported by the National Key Research and Development Project No. 2023YFE0206200, the National Natural Science Foundation of China under Grant 62376097, the Guangdong Provincial Natural Science Foundation for Outstanding Youth Team Project (No. 2024B1515040010) and the Guangdong Basic and Applied Basic Research Foundation under Grant 2025A1515012028. (*Corresponding author: Feng-Feng Wei*)

References

Bui, V.; and Mai, T. 2023. Imitation improvement learning for large-scale capacitated vehicle routing problems. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, 551–559.

Ceselli, A.; Felipe, Á.; Ortuño, M. T.; Righini, G.; and Tirado, G. 2021. A branch-and-cut-and-price algorithm for the electric vehicle routing problem with multiple technologies. In *Operations Research Forum*, volume 2, 8. Springer.

Chen, T.-Y.; Chen, W.-N.; Wei, F.-F.; Guo, X.-Q.; Song, W.-X.; Zhu, R.; Lin, Q.; and Zhang, J. 2025. The Confluence of Evolutionary Computation and Multi-Agent Systems: A Survey. *IEEE/CAA Journal of Automatica Sinica*, 12: 1–19.

Chen, T.-Y.; Chen, W.-N.; Wei, F.-F.; Hu, X.-M.; and Zhang, J. 2024a. Multi-agent swarm optimization with adaptive internal and external learning for complex consensus-based distributed optimization. *IEEE Transactions on Evolutionary Computation*.

Chen, X.-L.; Liao, X.-C.; Wei, F.-F.; and Chen, W.-N. 2024b. An Order-aware Adaptive Iterative Local Search Metaheuristic for Multi-depot UAV Pickup and Delivery Problem. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 1183–1191.

Cuchỳ, M.; Vokřínek, J.; and Jakob, M. 2024. Multi-objective electric vehicle route and charging planning with contraction hierarchies. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, 114–122.

Dorigo, M.; and Stützle, T. 2018. Ant colony optimization: overview and recent advances. *Handbook of metaheuristics*, 311–351.

Gendreau, M.; Potvin, J.-Y.; Bräumlaysy, O.; Hasle, G.; and Løkketangen, A. 2008. *Metaheuristics for the Vehicle Routing Problem and Its Extensions: A Categorized Bibliography*, 143–169. Boston, MA: Springer US. ISBN 978-0-387-77778-8.

He, L.; de Weerdt, M.; and Yorke-Smith, N. 2020. Time/sequence-dependent scheduling: the design and evaluation of a general purpose tabu-based adaptive large neighbourhood search algorithm. *Journal of Intelligent Manufacturing*, 31(4): 1051–1078.

Huang, T.; Li, J.; Koenig, S.; and Dilkina, B. 2022. Anytime multi-agent path finding via machine learning-guided large neighborhood search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 9368–9376.

Jia, Y.-H.; Mei, Y.; and Zhang, M. 2021. A bilevel ant colony optimization algorithm for capacitated electric vehicle routing problem. *IEEE Transactions on Cybernetics*, 52(10): 10855–10868.

Kuroiwa, R.; and Beck, J. C. 2023. Domain-independent dynamic programming: Generic state space search for combinatorial optimization. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, 236–244.

Lee, C. 2021. An exact algorithm for the electric-vehicle routing problem with nonlinear charging time. *Journal of the Operational Research Society*, 72(7): 1461–1485.

Li, D.; Liu, C.; and Gan, W. 2009. A new cognitive model: Cloud model. *International journal of intelligent systems*, 24(3): 357–375.

Lin, B.-C.; Liu, X.-F.; and Mei, Y. 2022. Efficient extended ant colony optimization for capacitated electric vehicle routing. In *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, 504–511. IEEE.

Mavrovouniotis, M.; Menelaou, C.; Timotheou, S.; Panayiotou, C.; Ellinas, G.; and Polycarpou, M. 2020. Benchmark set for the IEEE WCCI-2020 competition on evolutionary computation for the electric vehicle routing problem. *KIOS COE*.

Nafstad, G. M.; Desaulniers, G.; and Stålhane, M. 2025. Branch-Price-and-Cut for the Electric Vehicle Routing Problem with Heterogeneous Recharging Technologies and Nonlinear Recharging Functions. *Transportation Science*, 59(3): 628–646.

Neri, F.; and Cotta, C. 2012. Memetic algorithms and memetic computing optimization: A literature review. *Swarm and Evolutionary Computation*, 2: 1–14.

Nie, Z.-H.; Yang, Q.; Zhang, E.; Liu, D.; and Zhang, J. 2022. Ant colony optimization for electric vehicle routing problem with capacity and charging time constraints. In *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 480–485. IEEE.

Niu, B.; Li, W.; and Yi, W. 2024. An improved genetic algorithm for vehicle routing problem with time window requirements. In *International Conference on Swarm Intelligence*, 13–25. Springer.

Stützle, T.; and Hoos, H. H. 2000. MAX–MIN ant system. *Future generation computer systems*, 16(8): 889–914.

Tahami, H.; Rabadi, G.; and Haouari, M. 2020. Exact approaches for routing capacitated electric vehicles. *Transportation Research Part E: Logistics and Transportation Review*, 144: 102126.

Wang, L.; Ding, Y.; Chen, Z.; Su, Z.; and Zhuang, Y. 2024. Heuristic Algorithms for Heterogeneous and Multi-Trip Electric Vehicle Routing Problem with Pickup and Delivery. *World Electric Vehicle Journal*, 15(2).

Wu, Z.; and Zhang, J. 2023. A branch-and-price algorithm for two-echelon electric vehicle routing problem. *Complex & Intelligent Systems*, 9(3): 2475–2490.

Yang, M.-C.; Chen, W.-N.; Wei, F.-F.; and Zhang, J. 2025. Gene Expression Programming-Based Ride Insert Policy for Online Electric Vehicle Ride-Hailing Optimization. *IEEE Transactions on Intelligent Transportation Systems*.