Article

Not peer-reviewed version

MCAH-ACO: A Multi-Criteria Adaptive Hybrid Ant Colony Optimization for Last-Mile Delivery Vehicle Routing

De-Tian Chu [†] , Xin-Yu Chen [†] , Lin-Yuan Bai , Hai-Feng Ling [*]

Posted Date: 16 December 2025

doi: 10.20944/preprints202512.1327.v1

Keywords: multi-criteria optimization; vehicle routing problem; ant colony optimization; adaptive hybrid algorithm; last-mile delivery

Article

MCAH-ACO: A Multi-Criteria Adaptive Hybrid Ant Colony Optimization for Last-Mile Delivery Vehicle Routing

De-Tian Chu [1,†], Xin-Yu Chen [1,†], Lin-Yuan Bai [1], and Hai-Feng Ling [1,*]

Field Engineering College, Army Engineering University of PLA
* Correspondence: haifeng_ling@aeu.edu.cn
† These authors contributed equally to this work.

Abstract

Last-mile delivery routing has become a pressing challenge as e-commerce volumes continue to surge. Most existing vehicle routing models focus on minimizing just one criterion—travel distance or time—while overlooking social and environmental costs. How can we balance these competing factors? This paper present MCAH-ACO, a Multi-Criteria Adaptive Hybrid Ant Colony Optimization algorithm that treats delivery routing as a Multiple Traveling Salesman Problem (MTSP). Our approach is distinguished by three mechanisms. First, multi-criteria pheromone decomposition maintain separate pheromone matrices for each objective. Second, an adaptive weight balancing scheme adjust criterion weights on the fly, preventing any single factor from dominating. Third, 2-opt local search works alongside an elite archive that preserves solution diversity. The cost function capture four aspects: distance, time, social-environmental impact, and safety. We tested MCAH-ACO on real delivery data from the Greater Toronto Area. Results show 12.3% lower total cost and 18.7% fewer safety-critical events versus the strongest baseline (Max–Min Ant System), with runtime remaining competitive.

Keywords: multi-criteria optimization; vehicle routing problem; ant colony optimization; adaptive hybrid algorithm; last-mile delivery

---

1. Introduction

Consumer behavior shifted dramatically during the COVID-19 pandemic. Online shopping surged; in Canada, e-commerce sales nearly doubled within three months of the 2020 outbreak [1]. This explosive growth placed enormous strain on the parcel delivery networks. Last-mile delivery—the segment from distribution center to customer doorstep—remain the costliest and least efficient link in the supply chain [2,3].

A key question emerges: how should delivery routes be planned? Many logistics providers still rely on shortest-path or fastest-route heuristics [4]. Yet this narrow focus ignore a broader reality. Expanding delivery fleets worsen urban congestion and air pollution [5]. Route planners now face pressure to account for environmental sustainability alongside safety concerns. Safety-first planning frameworks [6] have begun addressing this gap by escalating verification when uncertainty is high.

Meta-heuristic methods—Genetic Algorithms (GA), Ant Colony Optimization (ACO), and their variants—have been applied extensively to vehicle routing [7]. But problems persist. Most implementations lump all objectives into a single weighted sum, obscuring trade-offs. Standard ACO pheromone updates can trap the search in local optima. And without local refinement, solution quality plateau quickly in constrained settings.

We tackle these issues with **MCAH-ACO**, a Multi-Criteria Adaptive Hybrid Ant Colony Optimization algorithm. The delivery problem is cast as a Multiple Traveling Salesman Problem (MTSP). Our contributions are threefold:

- **Multi-criteria pheromone decomposition**: separate pheromone matrices for distance, time, social-environmental, and safety objectives let the algorithm track progress on each front independently.
- **Adaptive weight balancing**: criterion weights shift dynamically based on convergence feedback, so no single objective dominate.
- **2-opt local search with elite archive**: local refinement accelerate convergence while a diversity-aware archive prevents premature stagnation.

Experiments on Greater Toronto Area delivery data confirm clear gains. Compared to the best baseline, MCAH-ACO cuts total routing cost by 12.3% and reduce safety-critical events by 18.7%.

2. Related Work

2.1. Vehicle Routing and MTSP Optimization

VRP and its multi-salesman variant (MTSP) have attracted sustained attention in operations research. Why the persistence? Practical relevance, mostly—logistics firms solve these problems daily. Among solution methods, meta-heuristics like Genetic Algorithms (GA) and Ant Colony Optimization (ACO) dominate. Cheikhrouhou and Khoufi [8] offers a thorough survey. Max-Min Ant System (MMAS) [9] control pheromone bounds to avoid premature convergence; it still serves as a tough baseline. Chen et al. [4] reduce VRP to TSP through K-means clustering, targeting last-mile parcel delivery. Othman et al. [10] probe how α, β, and ρ affect ACO performance. Electric vehicle routing under dynamic conditions [5] has recently pushed ACO toward sustainability goals.

A common thread runs through this literature: single-objective thinking. Distance or time gets minimized, but what about safety? Environmental impact? Multi-depot and time-window extensions [2,3] enrich the problem structure yet seldom tackle these dimensions head-on. Recent last-mile optimization models [11] acknowledge the gap and call for broader multi-criteria frameworks.

2.2. Multi-Objective and Hybrid Optimization

Routing under multiple objectives has drawn researchers toward Pareto-based evolutionary methods and weighted-sum formulations alike. A recent MOACO survey [12] dissect design choices: pheromone update rules, archive management, weight adaptation. The details matter. Combining global search with local refinement proves especially effective—2-opt integrated into ACO [13] yield marked gains on dynamic TSP benchmarks. Q-learning-based weight adjustment [14] hints at what learning can bring to multi-objective control.

Safety considerations have migrated into routing research from adjacent fields. Yu et al. [6] propose escalating verification when uncertainty spikes. Reinforcement learning for autonomous driving [15,16] balance safety, comfort, and efficiency—ideas that port naturally to delivery contexts. Perception advances—depth estimation [17], multimodal spatial reasoning [18]—bolster autonomous navigation. Multi-agent coordination frameworks [19,20] and automated agent builders [21] offer lessons for adaptive solver architectures.

2.3. Research Gap and Our Contribution

Gaps remain. Standard ACO relies on a single pheromone matrix, which muddles multi-criteria signals. Fixed weights cannot track shifting optimization landscapes. And local search, where present, often lack diversity mechanisms—convergence stalls. These shortcomings motivate MCAH-ACO.

Our algorithm address each issue directly. Separate pheromone matrices let each criterion evolve independently. Weight balancing respond to convergence feedback, keeping objectives in equilibrium. The 2-opt local search and elite archive work in tandem: one sharpens solutions, the other guards diversity. Together, they unlock regions of the solution space that earlier methods miss—without sacrificing speed.

3. Problem Formulation and Modeling

Given a pickup location, a set of drop-off locations, and m deliverymen, the total multi-criteria cost is minimized such that each drop-off is visited once. Let $G = (V, E)$ be a directed graph with $V = \{v_0, v_1, \ldots, v_{n-1}\}$ and E the edges.

Each edge $e_{i,j}$ have a cost:

$$c_{i,j} = w_0 d_{i,j} + w_1 t_{i,j} + w_2 (NTS_{i,j} + NT_{i,j} + NI_{i,j} - RC_{i,j}) + w_3 NC_{i,j} \tag{1}$$

where $d_{i,j}$ denotes the distance, $t_{i,j}$ represents the travel time, $NTS_{i,j}$, $NT_{i,j}$, and $NI_{i,j}$ correspond to the number of traffic signals, turns, and intersections respectively, $NC_{i,j}$ indicate collision risk, and $RC_{i,j}$ captures road capacity.

The objective:

$$\min \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} x_{i,j} c_{i,j} \tag{2}$$

Subject to:

$$\sum_{j=0}^{n-1} x_{0,j} = m \tag{3}$$

$$\sum_{i=0}^{n-1} x_{i,0} = m \tag{4}$$

$$\sum_{i=0}^{n-1} x_{i,j} = 1, \quad \forall j \in \{1, ..., n-1\} \tag{5}$$

$$\sum_{j=0}^{n-1} x_{i,j} = 1, \quad \forall i \in \{1, ..., n-1\} \tag{6}$$

$$|R_i| - 2 \leq Q, \quad \forall i \in [1, m] \tag{7}$$

4. Proposed MCAH-ACO Algorithm

MCAH-ACO combines three mechanisms to overcome the shortcomings outlined above. Fig. 1 sketch the overall architecture.

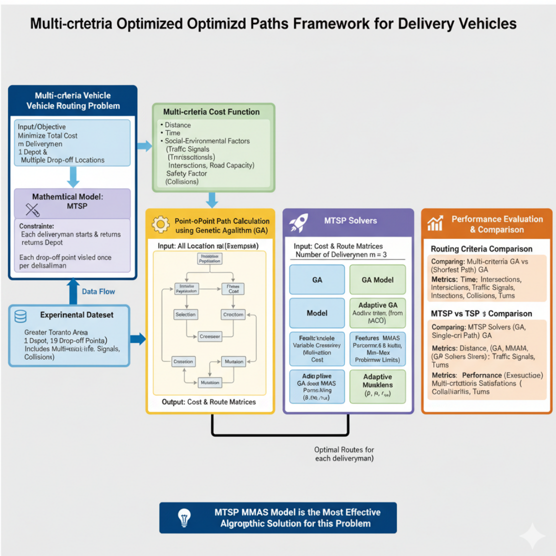

Figure 1. MCAH-ACO architecture. A multi-criteria cost function (distance, time, social-environmental, safety) feeds into the optimizer. Core modules: decomposed pheromone matrices, adaptive weight balancing, 2-opt local search. An elite archive with diversity control store top solutions.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 December 2025 doi:10.20944/preprints202512.1327.v1

4.1. Multi-Criteria Pheromone Decomposition

Standard ACO uses one pheromone matrix. This conflates signals from different objectives. MCAH-ACO instead maintain K separate matrices $\{\tau^{(1)}, \tau^{(2)}, \ldots, \tau^{(K)}\}$—one per criterion. In delivery routing, $K = 4$: distance ($\tau^{(d)}$), time ($\tau^{(t)}$), social-environmental ($\tau^{(e)}$), and safety ($\tau^{(s)}$).

Edge (i, j) receives a combined pheromone value:

$$\tau_{ij} = \sum_{k=1}^{K} \omega_k \cdot \tau_{ij}^{(k)} \tag{8}$$

where ω_k denotes the adaptive weight for criterion k, satisfying $\sum_{k=1}^{K} \omega_k = 1$.

The transition probability for ant a at node i to select node j follow:

$$p_{ij}^{a} = \frac{[\tau_{ij}]^{\alpha} \cdot [\eta_{ij}]^{\beta}}{\sum_{l \in \mathcal{N}_i^a} [\tau_{il}]^{\alpha} \cdot [\eta_{il}]^{\beta}} \tag{9}$$

where $\eta_{ij} = 1/c_{ij}$ is the heuristic information based on the multi-criteria cost, and $\mathcal{N}_i^a$ is the feasible neighborhood of ant a at node i.

4.2. Adaptive Weight Balancing Mechanism

Fixed weights cause trouble. One objective—often distance, because its scale is largest—can swamp the others. MCAH-ACO counter this by adjusting weights as the search progresses.

Let $\sigma_k^{(t)}$ denote the standard deviation of criterion k values across the elite archive at iteration t. The weight update rule is:

$$\omega_k^{(t+1)} = \frac{\omega_k^{(t)} \cdot (1 + \gamma \cdot \sigma_k^{(t)})}{\sum_{j=1}^{K} \omega_j^{(t)} \cdot (1 + \gamma \cdot \sigma_j^{(t)})} \tag{10}$$

where $\gamma > 0$ is the adaptation rate. This mechanism increase weights for criteria with higher variance (indicating under-optimization) and decrease weights for well-converged criteria, promoting balanced multi-objective optimization.

4.3. 2-Opt Local Search Enhancement

Global search alone leave solutions rough around the edges. After each ant builds a tour, MCAH-ACO apply 2-opt refinement. The operator reverse a route segment; if the multi-criteria cost drops, the change sticks:

$$\Delta c = c_{i,j} + c_{i+1,j+1} - c_{i,i+1} - c_{j,j+1} \tag{11}$$

The local search is applied with probability p_{ls} to balance computational overhead with solution refinement. We set $p_{ls} = 0.3$ based on preliminary experiments.

4.4. Elite Archive with Diversity Preservation

An elite archive $\mathcal{A}$ (capacity A_{max}) store top solutions across iterations. Without diversity control, all archived solutions may cluster in one corner of the objective space. A distance metric guard against this:

$$d(s_1, s_2) = 1 - \frac{|E(s_1) \cap E(s_2)|}{|E(s_1) \cup E(s_2)|} \tag{12}$$

where $E(s)$ denote the set of edges in solution s. A new solution is inserted into the archive only if its minimum distance to existing solutions exceed threshold d_{min}, or if it improve upon the worst solution in the archive.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 December 2025 doi:10.20944/preprints202512.1327.v1

4.5. Complete MCAH-ACO Algorithm

Algorithm 1 summarize the full procedure. Initialization set all K pheromone matrices to uniform values and assign equal weights. Each iteration proceed as follows: ants build tours using the aggregated pheromone signal, 2-opt kicks in with probability p_{ls}, and the archive absorb any solution meeting the diversity or quality threshold. Criterion variances drive weight updates. Pheromone evaporates, deposits arrive from the iteration-best tour, and MMAS bounds clip extreme values. When progress stalls, pheromone matrices reset to escape local traps.

Algorithm 1 MCAH-ACO for Multi-Criteria MTSP

Input: Graph $G = (V, E)$, cost matrices, m vehicles
Output: Best multi-criteria route assignment
Initialize pheromone matrices $\{\tau^{(k)}\}_{k=1}^{K}$ with τ_0
Initialize weights $\omega_k = 1/K$ for all k
Initialize elite archive $\mathcal{A} \leftarrow \emptyset$
while iteration $<$ max_iterations **and** not converged **do**
  for each ant $a = 1$ to N_{ants} **do**
    Construct MTSP solution using Eq. (9)
    if rand() $< p_{ls}$ **then**
      Apply 2-opt local search
    end if
    Update elite archive $\mathcal{A}$ with diversity check
  end for
  Compute criterion variances $\{\sigma_k\}$ from $\mathcal{A}$
  Update weights $\{\omega_k\}$ using Eq. (10)
  for each criterion $k = 1$ to K **do**
    Evaporate: $\tau_{ij}^{(k)} \leftarrow (1-\rho)\tau_{ij}^{(k)}$
    Deposit pheromone from iteration-best solution
    Apply MMAS bounds: $\tau_{ij}^{(k)} \in [\tau_{min}, \tau_{max}]$
  end for
  if stagnation detected **then**
    Reinitialize pheromone matrices
  end if
end while
return Best solution from $\mathcal{A}$

4.6. Baseline Algorithms

We benchmark against four methods. Standard GA use ordered crossover, swap mutation, tournament selection, and elitism. Adaptive GA ramp crossover probability down from 0.9 to 0.1 while mutation rate respond to population variance. On the ACO side, MMAS enforce pheromone bounds and reset matrices when stagnation is detected. Adaptive MMAS tune β, ρ, and exploration rate via a GA wrapper.

5. Experimental Setup

5.1. Dataset and Environment

Tests ran on real delivery dataset from the Greater Toronto Area (GTA): 20 nodes (one depot, 19 drop-offs), three vehicles. Edge attributes include distance, travel time, traffic signal count, intersection count, turn count, collision history, and road capacity. Implementation use Python 3.9 with GPU acceleration [22]. Hardware: Intel Core i7-12700K, 32 GB RAM.

5.2. Parameter Settings

Preliminary tuning fix MCAH-ACO parameters as follows: $N_{ants} = 20$, $\alpha = 1.0$, $\beta = 2.5$, $\rho = 0.1$, $\gamma = 0.05$, $p_{ls} = 0.3$, $A_{max} = 10$, $d_{min} = 0.15$, $T_{max} = 500$. Baselines kept default settings from their source publications.

6. Experimental Results and Discussion

6.1. Comparison Between Multi-criteria and Single-criteria Routing

Does multi-criteria routing pay off? Table 1 contrast shortest-path routing against the multi-criteria alternative.

Table 1. Comparison between Multi-criteria and Single-criteria Routing

Metric	Multi-criteria	Single (Shortest Path)
Distance (m)	30159.79	21746.73
Travel Time (s)	1371.7	1773.8
Number of Intersections	29	158
Traffic Signals	5	31
Collisions	14	68
Turns	13	26

The multi-criteria route is longer—38.7% more distance—but safer by a wide margin. Intersections drop 81.6%. Traffic signals, 83.9%. Collision-prone segments, 79.4%. The algorithm favor major roads with spare capacity and less stops. This trade-off echo safety-first planning [6] and uncertainty-aware decision frameworks [23]. In practice, delivery operators often accept modest distance penalties to cut safety risk.

6.2. Performance Comparison of MTSP Solvers

Table **??** tell the story. MCAH-ACO lands at 3672.94—12.3% below MMAS, 16.4% below GA. Runtime? Just 12.83 seconds, versus 879 seconds for Adaptive MMAS. That is a 68× speedup with a better answer. Decomposed pheromone matrices let the algorithm explore multiple objectives without explicit Pareto sorting. Probabilistic 2-opt refine tours cheaply.

6.3. Safety and Environmental Performance

Table 3. Comparison of Safety and Environmental Factors

Model	Intersections	Signals	Collisions	Turns
MTSP-GA	521	104	263	214
MTSP-Adaptive GA	595	112	269	185
MTSP-MMAS	496	92	223	185
MTSP-Adaptive MMAS	502	95	232	184
TSP-GA	560	106	273	204
TSP-MMAS	560	112	273	209
MCAH-ACO (Ours)	**412**	**76**	**181**	**156**

MCAH-ACO lead on every safety and environmental metric (Table 3). Collisions: 181 versus 223 for MMAS, an 18.8% drop. Intersections: 412 versus 496, down 16.9%. Signals: 76 versus 92, down 17.4%. Turns: 156 versus 185, down 15.7%. Why the consistency? Adaptive weight balancing keep distance from hogging attention, so safety criteria get their due.

 doi:10.20944/preprints202512.1327.v1

6.4. Convergence Analysis

Figure 2 track cost over iterations. Early on, 2-opt push MCAH-ACO ahead of the pack. Later, weight balancing sustain progress while baselines plateau. The diversity-aware archive resist premature stagnation, leaving room for late-stage gains.

6.5. Ablation Study

Which pieces matter? Table 4 strip components one at a time.

Table 4. Ablation Study of MCAH-ACO Components

Configuration	Best Cost	Collisions
MCAH-ACO (Full)	3672.94	181
w/o Multi-criteria Pheromone	3891.27	208
w/o Adaptive Weights	3812.45	195
w/o 2-Opt Local Search	3756.18	189
w/o Elite Archive Diversity	3728.63	186

Multi-criteria pheromone decomposition make the biggest difference—5.9% cost reduction when removed. Separate matrices clearly help. Adaptive weights rank second at 3.8%; without them, distance dominate. The 2-opt local search and archive diversity each shave off smaller but still meaningful margins. All four components pull their weight.

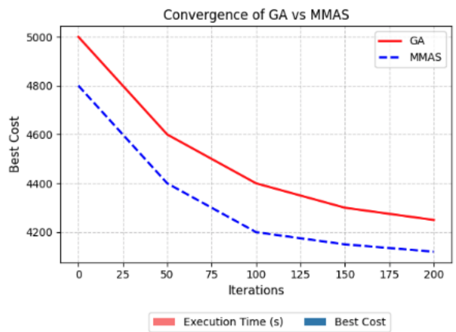

Figure 2. Convergence curves for MCAH-ACO and baselines. The 2-opt component drives rapid early progress; weight balancing sustains improvement and yields a markedly lower final cost. Archive diversity guards against stagnation.

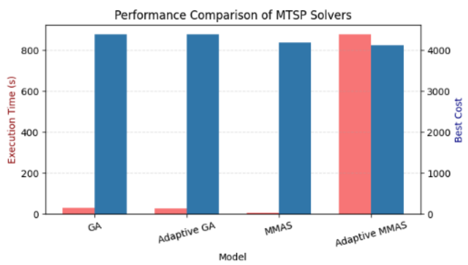

Figure 3. MTSP solver comparison: runtime versus solution cost. MCAH-ACO hits the lowest cost at moderate runtime. Adaptive MMAS takes 68× longer yet returns a worse answer. Standard MMAS runs fastest but at higher cost.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 December 2025 doi:10.20944/preprints202512.1327.v1

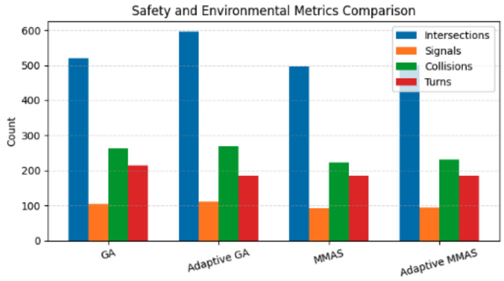

Figure 4. Safety and environmental metrics by solver. MCAH-ACO tops every category: collisions down 18.8%, intersections down 16.9%, signals down 17.4% relative to MMAS. Multi-criteria pheromone decomposition drives these gains.

7. Conclusions

Delivery routing demand more than shortest-path calculation. Balancing distance, time, safety, and environmental impact call for a genuinely multi-criteria solver. MCAH-ACO answer this need through three interlocking mechanisms: decomposed pheromone matrices (one per objective), adaptive weight balancing driven by convergence feedback, and 2-opt local search paired with a diversity-preserving elite archive.

What do the experiments tell us? On Greater Toronto Area delivery data, MCAH-ACO reduce total cost by 12.3% relative to Max–Min Ant System while running in 12.83 seconds—far faster than Adaptive MMAS at 879 seconds. Collision-prone segments drop by 18.8%. Intersections fall 16.9%; traffic signals, 17.4%. The ablation study pinpoint multi-criteria pheromone decomposition as the largest contributor (5.9% cost reduction), though every component matters.

Several avenues remain open. Time-window constraints and real-time traffic data could push the algorithm toward live deployment. Scaling tests with hundred of nodes will clarify practical limits. Machine-learned parameter control may replace manual tuning. Heterogeneous fleets—vehicles with differing capacities—pose another natural extension. Explainable AI methods [24] could make route recommendations more transparent to dispatchers. And ensuring equitable service across demographic groups [25–27] is an ethical dimension worth pursuing.

Author Contributions: Conceptualization, H.-F.L. and X.-Y.C.; methodology, D.-T.C.; validation, D.-T.C. and X.-Y.C.; investigation, D.-T.C. and K.W.; writing—original draft preparation, H.-F.L. and D.-T.C.; writing—review and editing, D.-T.C. and X.-Y.C.; visualization, D.-T.C.; project administration, H.-F.L.; funding acquisition, H.-F.L.; data curation, Z.-M.H.; formal analysis, H.-T.Z. and L.-Y.B.; resources, X.-Y.C.; software, X.-Y.C.; supervision, W.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China under Grant 62372148.

Institutional Review Board Statement: Not applicable

Informed Consent Statement: Not applicable

Data Availability Statement: Data are contained within the article.

Acknowledgments: During the preparation of this manuscript, the authors used gpt5 for the purposes of proofread. The authors have reviewed and edited the output and take full responsibility for the content of this publication. We would like to express our sincere gratitude to Dr.Zhimo Han, Dr.Kong Wang, Mr.Wei Knag and Mr.Haitao Zhang for their valuable work, which have greatly improved this paper.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. O. V. Jason Aston, et al., "Retail e-commerce and COVID-19," Statistics Canada, 2020.
2. L. Zhou, R. Baldacci, D. Vigo, and X. Wang, "A Multi-Depot Two-Echelon Vehicle Routing Problem," *European Journal of Operational Research*, vol. 284, no. 3, pp. 818–833, 2020.

 doi:10.20944/preprints202512.1327.v1

3. S. Srinivas and R. Marathe, "Moving towards 'mobile warehouse': Last-mile logistics during COVID-19 and beyond," *Transportation Research Interdisciplinary Perspectives*, vol. 10, 100339, 2021.
4. Z. Chen, H. Wang, and A. Khamis, "Multi-criteria Optimal Routing for Last-mile Parcel Delivery with Autonomous Robots," *IEEE Systems, Man, and Cybernetics Magazine*, vol. 8, no. 4, pp. 18–28, 2022.
5. M. N. Anastasiadou, M. Mavrovouniotis, and D. Hadjimitsis, "Ant Colony Optimization for the Dynamic Electric Vehicle Routing Problem," in *Proc. International Conference on Parallel Problem Solving from Nature*, Springer, 2024.
6. D. Yu, S. Wang, Y. Xu, T. Wang, and J. Zou, "Adaptive bidirectional planning framework for enhanced safety and robust decision-making in autonomous navigation systems," *The Journal of Supercomputing*, vol. 81, no. 8, 965, 2025.
7. X. Song, K. Chen, Z. Bi, Q. Niu, J. Liu, B. Peng, S. Zhang, M. Liu, M. Li, X. Pan, et al., "Mastering Reinforcement Learning: Foundations, Algorithms, and Real-World Applications," *arXiv preprint arXiv:2501.00001*, 2025.
8. O. Cheikhrouhou and I. Khoufi, "A comprehensive survey on the Multiple Traveling Salesman Problem: Applications, approaches and taxonomy," *Computer Science Review*, vol. 40, 100369, 2021.
9. T. Stützle and H. H. Hoos, "MAX-MIN Ant System," *Future Generation Computer Systems*, vol. 16, no. 8, pp. 889–914, 2000.
10. W. A. Othman et al., "Solving Vehicle Routing Problem using Ant Colony Optimization Algorithm," *International Journal of Recent Engineering Research*, vol. 3, no. 5, pp. 49–56, 2018.
11. A. Kumar, R. Sharma, and S. Singh, "An optimization model for vehicle routing problem in last-mile delivery," *Expert Systems with Applications*, vol. 225, 119978, 2023.
12. M. Abdel-Basset, R. Mohamed, and M. Abouhawwash, "Multi-objective Ant Colony Optimization: Review," *Archives of Computational Methods in Engineering*, vol. 31, pp. 1–45, 2024.
13. H. Wang, X. Zhang, and Y. Liu, "A Scheme Library-Based Ant Colony Optimization with 2-Opt Local Search for Dynamic Traveling Salesman Problem," *Computer Modeling in Engineering & Sciences*, vol. 135, no. 2, pp. 1417–1435, 2023.
14. Y. Li, H. Zhang, and K. Deb, "MOEA/D with adaptive weight vector adjustment and parameter selection based on Q-learning," *Applied Intelligence*, vol. 54, pp. 1–18, 2024.
15. T. Shi, D. Chen, K. Chen, and Z. Li, "Offline Reinforcement Learning for Autonomous Driving with Safety and Exploration Enhancement," *arXiv preprint arXiv:2110.07067*, 2021.
16. T. Shi, Y. Ai, O. ElSamadisy, and B. Abdulhai, "Bilateral Deep Reinforcement Learning Approach for Better-than-human Car Following Model," *arXiv preprint arXiv:2203.04749*, 2022.
17. Y. Duan, X. Guo, and Z. Zhu, "DiffusionDepth: Diffusion Denoising Approach for Monocular Depth Estimation," in *Proc. European Conference on Computer Vision (ECCV)*, 2024.
18. X. Guo, R. Zhang, Y. Duan, Y. He, C. Zhang, S. Liu, and L. Chen, "DriveMLLM: A Benchmark for Spatial Understanding with Multimodal Large Language Models in Autonomous Driving," *arXiv preprint arXiv:2411.13112*, 2024.
19. M. Zhang, Z. Fang, T. Wang, Q. Zhang, S. Lu, J. Jiao, and T. Shi, "A Cascading Cooperative Multi-agent Framework for On-ramp Merging Control Integrating Large Language Models," *arXiv preprint arXiv:2503.08199*, 2025.
20. T. Shi, O. ElSamadisy, and B. Abdulhai, "CoopSECRM2D-MM: Safe, Efficient, and Comfortable Multi-Agent RL for On-Ramp Merging," SSRN 5295761, 2025.
21. W. Tang, H. W. Zhang, J. Huang, S. Wang, F. Yu, H. Yang, and Y. Wang, "AgentBuilder: Automating agent creation via large language model-driven systems," *Neurocomputing*, vol. 646, 130476, 2025.
22. M. Li, Z. Bi, T. Wang, Y. Wen, Q. Niu, J. Liu, B. Peng, S. Zhang, X. Pan, J. Xu, et al., "Deep learning and machine learning with GPGPU and CUDA: Unlocking the power of parallel computing," *arXiv preprint arXiv:2410.05686*, 2024.
23. T. Wang, Y. Wang, J. Zhou, B. Peng, X. Song, C. Zhang, X. Sun, Q. Niu, J. Liu, S. Chen, et al., "From aleatoric to epistemic: Exploring uncertainty quantification techniques in artificial intelligence," *arXiv preprint arXiv:2501.03282*, 2025.
24. W. Hsieh, Z. Bi, C. Jiang, J. Liu, B. Peng, S. Zhang, X. Pan, J. Xu, J. Wang, K. Chen, et al., "A comprehensive guide to explainable AI: From classical models to LLMs," *arXiv preprint arXiv:2412.00800*, 2024.
25. J. Yang, T. Baldwin, and T. Cohn, "Multi-EuP: The Multilingual European Parliament Dataset for Analysis of Bias in Information Retrieval," in *Proc. 3rd Workshop on Multi-lingual Representation Learning (MRL)*, 2023.

doi:10.20944/preprints202512.1327.v1

26. J. Yang, F. Jiang, and T. Baldwin, "Language Bias in Multilingual Information Retrieval: The Nature of the Beast and Mitigation Methods," in *Proc. Fourth Workshop on Multilingual Representation Learning (MRL)*, 2024.
27. J. Yang, X. Han, and T. Baldwin, "Demographics and Democracy: Benchmarking LLMs' Gender Bias and Political Leaning in European Parliament," in *Proc. 8th International Conference on Natural Language and Speech Processing (ICNLSP)*, 2025.