Human-Computer Interaction & Emerging Technologies, Vol. 233, 2026, 393–403
https://doi.org/10.54941/ahfe1007539

Multi-Level Natural Language Interaction for Eliciting Implicit Preferences in Interactive Vehicle Routing

Mingwei Chen and Liang Ma

National Key Laboratory of Human Factors Engineering, Department of Industrial Engineering, Tsinghua University, Beijing, China

ABSTRACT

Interactive vehicle routing enables planners to express preferences during optimization, yet existing interaction methods—such as scoring, ranking, and weight adjustment—constrain the richness of preference expression. In practice, planners naturally form multi-level judgments about routing solutions: from specific task assignments, to regional route patterns, to overall solution quality. However, current systems lack mechanisms to capture such semantically rich, multi-level feedback. This study proposes an LLM-enhanced multi-level interaction approach that enables planners to express preferences through natural language at three levels: node-level, regional-level, and solution-level. A large language model serves as a semantic bridge, interpreting natural language feedback and translating it into adjustments to a probability-based preference matrix that guides algorithmic search. The approach is illustrated through scenario-based demonstrations using a road cleaning problem, showing how natural language feedback at different levels can be interpreted and integrated into the optimization process. This work represents a step toward more intuitive and expressive human–algorithm collaboration in vehicle routing.

Keywords: Interactive vehicle routing, Large language models, Implicit preference, Human-algorithm collaboration, Natural language interaction

INTRODUCTION

The vehicle routing problem (VRP), which seeks optimal routes for a fleet of vehicles serving spatially distributed tasks, is a fundamental challenge in transportation and logistics (Vidal et al., 2020). While numerous algorithms have been developed to optimize formally defined objectives such as cost and workload balance, a persistent gap remains between algorithmically optimal solutions and practically acceptable ones (Chen, Wang, et al., 2025). In real-world operations, human planners bring domain knowledge and experiential judgment that resists full mathematical formulation (Canoy et al., 2023). Their preferences—such as tradeoffs among competing objectives (Wenger & Geiger, 2008)—are difficult to articulate a priori and tend to surface only when planners evaluate concrete solutions rather than during the modeling phase (Chen, Ma et al., 2025). Shaped by experience and context, these implicit preferences can decisively influence whether a solution is deemed acceptable in practice.

Received March 24, 2026; Revised April 28, 2026; Accepted May 13, 2026; Available online July 20, 2026

Human-in-the-loop optimization approaches, particularly in interactive vehicle routing, have been developed to address this gap by incorporating planners into the optimization loop (Baker & Carreto, 2003; Meignan et al., 2015). Chen et al. (2025) proposed an omnidimensional human–algorithm collaboration approach that supports planner interaction across four dimensions—decision variables, constraints, objectives, and solutions—demonstrating its effectiveness in capturing implicit preferences such as task-binding patterns and trade-off tendencies. Other studies have explored interaction through direct manipulation of routes, parameter adjustment, and solution comparison (Tezcaner & Köksalan, 2011; Tütüncü et al., 2009). However, across these approaches, interaction modalities remain constrained to narrow, predefined formats. Users can indicate which solution they prefer, but conveying why or how a solution should be improved remains difficult within such channels. Much of the planner's nuanced evaluative reasoning—particularly the kind that naturally emerges during solution inspection—is lost in this translation.

A key observation motivating this work is that planners, when evaluating routing solutions, naturally form judgments at multiple levels of granularity. They may notice that a specific task appears misassigned, that a geographic area has too many overlapping routes, or that the overall workload distribution feels uneven. These multi-level reactions reflect how humans perceive and reason about complex spatial arrangements—and they carry rich information about underlying preferences. Crucially, such judgments are naturally communicated through ordinary language: expressions like "this area looks too crowded" or "that task should go to another vehicle" convey nuanced, context-dependent preferences.

Recent advances in large language models (LLMs) make it feasible to use natural language as a direct interaction channel between planners and optimization algorithms. While LLMs have recently been applied in optimization—primarily for generating heuristic algorithms or search operators (Liu et al., 2024; Romera-Paredes et al., 2024)—their use as a continuous interaction medium within an optimization loop remains largely unexplored. Building on the algorithmic foundation of Chen et al. (2025), this study proposes an LLM-enhanced multi-level interaction approach for interactive vehicle routing. The approach enables planners to express preferences through natural language at three levels—node, regional, and solution—while an LLM serves as a semantic bridge that interprets the feedback and maps it to a probability-based preference matrix guiding the algorithmic search.

The main contributions of this work are twofold. First, we propose a multi-level natural language interaction approach for interactive vehicle routing, enabling planners to express preferences at node, regional, and solution levels, with an LLM translating feedback into structured guidance for the search algorithm. Second, we illustrate the approach through scenario-based demonstrations on a road cleaning problem, showing how different types of implicit preferences can be captured through natural language interaction.

PROBLEM DESCRIPTION

This study uses road cleaning operations as the application context, building on Chen et al. (2025). Each task is a road segment with a cleaning demand. Vehicles depart from a shared depot, service assigned road segments, and return, subject to capacity and working time constraints. Regardless of the specific problem variant, two types of decisions are involved: assigning tasks to vehicles, and constructing each vehicle's route. Following Chen et al. (2025), this study adopts a sequential "task assignment first, route construction second" approach, reflecting real-world practice where planners focus on allocating tasks among vehicles while drivers determine their own travel sequences. The interactive optimization operates on the task assignment layer—the decision most meaningful to the planner—while route construction is handled algorithmically. The primary explicit objectives are total cost (the sum of workloads across all routes) and workload balance (the difference between the maximum and minimum vehicle workloads).

Beyond these explicit objectives, planners hold implicit preferences that shape their evaluation of solutions but resist formalization as optimization objectives. Three representative types are relevant to this study:

- **Task-binding preferences** reflect tendencies to assign specific tasks to particular vehicles due to practical considerations such as vehicle capabilities, driver familiarity, or unmodeled operational constraints (Liu et al., 2020).
- **Regional service preferences** concern how tasks are spatially distributed among vehicles. Planners often prefer geographically coherent service territories, avoiding excessive route overlap (Chen, Wang, et al., 2025; Rossit et al., 2019). Notably, planners' perceptions of well-organized spatial distributions are subjective and context-dependent.
- **Tradeoff preferences** arise from the multi-objective nature of the problem (Xin et al., 2018). Different planners may prioritize cost versus balance differently, and these tendencies may shift as the planner develops a better understanding of achievable outcomes.

These categories represent typical implicit preferences identified in practice, though they are not exhaustive. Since implicit preferences are, by nature, difficult to enumerate in advance, the proposed approach is designed to allow unanticipated preferences to surface naturally through the interaction process.

LLM-ENHANCED MULTI-LEVEL INTERACTION APPROACH

This section presents the proposed approach, as illustrated in Figure 1. We first describe the underlying algorithm, then introduce the multi-level interaction design and the role of the LLM in translating natural language expression into algorithmic operation.

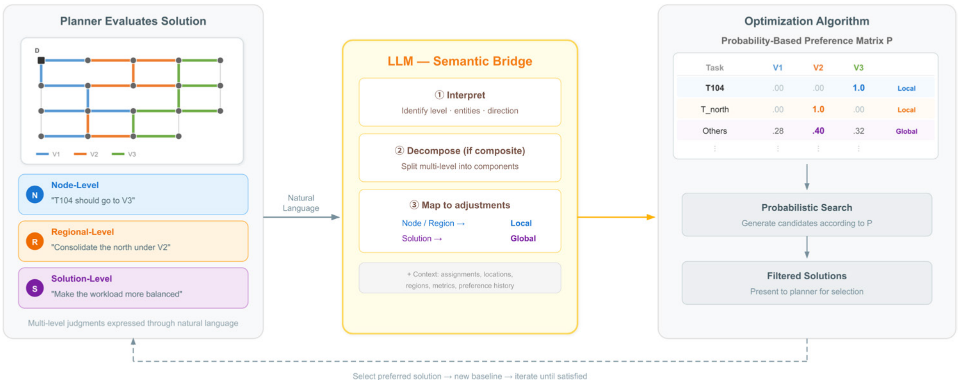

Figure 1: Illustration of the LLM-enhanced multi-level interaction approach.

Optimization Algorithm

The optimization builds on the probability-based neighbourhood search algorithm (PBNSA) of Chen et al. (2025). A task assignment is encoded as a vector where position i indicates the vehicle assigned to task i. The core mechanism is a probability-based preference matrix P, where $P(i, k)$ denotes the probability of assigning task i to vehicle k. In the baseline configuration, each task retains a high probability (e.g., $P_f = 0.95$) of staying with its current vehicle, with the remaining probability distributed among nearby vehicles. New candidate solutions are generated by probabilistically reassigning tasks according to P. For each candidate assignment, routes are constructed using the improved-merge heuristic (Willemse & Joubert, 2016). The algorithm retains Pareto non-dominated solutions and presents them for planner evaluation. The selected solution becomes the baseline for the next iteration.

Multi-Level Natural Language Interaction

Our approach enables planners to express preferences through natural language at three levels, corresponding to how planners naturally evaluate routing solutions. At the **node level**, the planner focuses on individual tasks, expressing preferences about specific assignments. For example: “Task T8 needs deep cleaning—assign it to Vehicle 3.” Such feedback addresses the same concern as direct task reassignment, but in a form that can convey conditional or relational reasoning. At the **regional level**, the planner evaluates spatial patterns within a geographic area. For example: “The western area has too many vehicles overlapping—one vehicle should cover that zone.” Regional preferences are visually apparent to planners but laborious to express through task-by-task manipulation. At the **solution level**, the planner assesses overall solution quality and expresses global preferences. For example: “The workload is too unbalanced—I’d accept slightly higher cost for better balance.” Compared to selecting “balance” as a priority, natural language conveys the degree of desired improvement and the tradeoffs the planner is willing to accept. In practice, planners may combine levels in a single statement—for instance, “Overall the balance is acceptable, but the western area looks crowded, and Task T8 should move to Vehicle 3.” The approach handles such mixed-level feedback by processing each component and integrating the adjustments.

LLM-Based Preference Mapping

The LLM serves as intermediary between the planner's natural language and the algorithm's probability-based search. When the planner provides feedback, the LLM receives the text together with contextual information about the current solution: task locations, current assignments, spatial region definitions, and performance metrics. The LLM interprets the feedback to identify referenced entities, preference direction, and interaction level.

The LLM then adjusts P through two types of operations. **Local adjustments** modify assignment probabilities of specific tasks: if the planner wants Task T8 assigned to Vehicle 3, $P(T_8,V_3)$ is increased while other probabilities decrease. For regional feedback, the LLM identifies all tasks in the referenced area and applies local adjustments collectively. **Global adjustments** shift the probability distribution across all tasks toward a particular optimization direction—for instance, biasing all tasks toward less-loaded vehicles when the planner requests better balance.

Interaction Workflow

As shown in Figure 1, the overall interaction follows an iterative process: the planner inspects the current solution and optionally provides natural language feedback at any combination of levels. The LLM processes the feedback, updates the preference matrix, and the algorithm generates a set of new candidate solutions accordingly. The planner then selects a preferred solution from these candidates, which serves as the basis for evaluation in the next iteration. This cycle repeats until the planner is satisfied with the solution.

SCENARIO-BASED DEMONSTRATIONS

To illustrate how the proposed approach works in practice, we present four scenarios based on a real-world road cleaning network. The network comprises 50 nodes and 156 task arcs (road segments requiring cleaning), served by three vehicles from a shared depot. In the initial solution, the total cost is 13.24 hours and the workload gap is 3.05 hours, with Vehicle 1 at 5.26 hours, Vehicle 2 at 2.47 hours, and Vehicle 3 at 5.51 hours. Each scenario demonstrates how a particular type of implicit preference is expressed through natural language and incorporated into the optimization. The four scenarios address node-level, regional-level, solution-level, and composite multi-level interaction, respectively.

Scenario A: Task-Binding Preference (Node-Level)

A new pollution source has been identified in the southeast of the network, near the road segments corresponding to tasks T104 and T105. These segments require deep cleaning, for which only Vehicle 3 is equipped. In the current solution, both tasks are assigned to Vehicle 1. Upon inspecting the map, the planner notices that these two segments are adjacent to Vehicle 3's existing service area, and reassigning them would create a more connected route for Vehicle 3 while addressing the deep-cleaning requirement. The planner provides the following feedback: "These two road segments (T_{104} and T_{105}) should be assigned to Vehicle 3—they connect nicely to Vehicle 3's existing route, and Vehicle 3 has deep-cleaning capability."

The LLM identifies this as node-level feedback referencing two specific tasks and a target vehicle. It generates a local adjustment: both $P(T_{104}, V_3)$ and $P(T_{105}, V_3)$ and are set to 1.0, effectively binding these tasks to Vehicle 3. The algorithm then searches with the updated probability matrix, generating a set of Pareto non-dominated candidate solutions. The planner reviews the search results and selects solution #94. In this solution, T_{104} and T_{105} are assigned to Vehicle 3 as requested, along with six other minor task reassignments—a natural consequence of the probabilistic search process, which explores the neighborhood of the preference-adjusted baseline. The total cost decreases slightly to 13.14 hours (–0.8%), and the workload balance improves to 2.52 hours (–17%). Figure 2 shows the current solution with the highlighted tasks, the planner's feedback and LLM interpretation, the search results, and the selected solution.

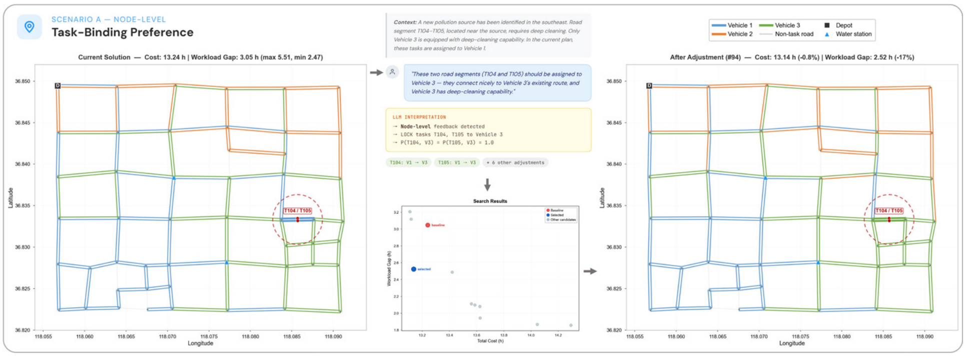

Figure 2: Illustration of node-level interaction.

Scenario B: Regional Service Preference (Regional-Level)

The planner reviews the current solution and observes that the northern region of the network has tasks scattered across all three vehicles, producing a visually chaotic pattern with overlapping routes. The planner believes that consolidating these tasks under a single vehicle would reduce travel inefficiency and produce a clearer, more organized routing pattern. The planner provides the following feedback: "The northern region looks chaotic—tasks are split among all three vehicles. I'd like to consolidate all northern-area tasks under Vehicle 2 for better efficiency."

The LLM identifies this as regional-level feedback. Using the spatial context provided in the prompt, it resolves "the northern region" to 32 tasks in that zone. It generates local adjustments for all 32 tasks, binding them to Vehicle 2. Of these, 11 tasks require actual reassignment—5 from Vehicle 1 and 6 from Vehicle 3—while the remaining 21 are already assigned to

Vehicle 2 and need no change. After search, the planner selects solution #86. The northern region is now served exclusively by Vehicle 2, with a clear, coherent service territory visible on the map. The cost increases marginally to 13.28 hours (+0.3%), while the workload balance improves to 1.22 hours (–60%). Figure 3 shows the before-and-after comparison with the northern region highlighted.

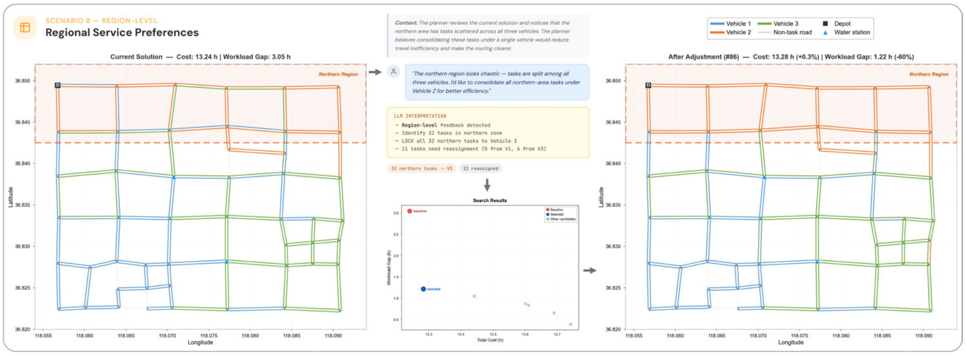

Figure 3: Illustration of regional-level interaction.

Scenario C: Trade-off Preference (Solution-Level)

The current solution has a workload gap of 3.05 hours, with Vehicle 1 at 5.26 hours, Vehicle 2 at 2.47 hours, and Vehicle 3 at 5.51 hours. The planner considers this imbalance too large and wants a more equitable distribution across the three vehicles. The planner provides the following feedback: "The workload feels quite uneven—there's a big gap between the busiest and the least busy vehicle. I'd like the assignment to be more balanced overall."

The LLM identifies this as solution-level feedback and applies a global balance adjustment. This shifts the assignment probabilities of all tasks toward less-loaded vehicles without targeting any specific task or region. After search, the planner selects solution #8, which achieves a substantially improved workload balance of 0.91 hours (–70%) at a cost of 13.87 hours (+4.7%). The route map shows a different task distribution, with Vehicle 2 now serving a much larger share of the network. Figure 4 shows the workload redistribution results.

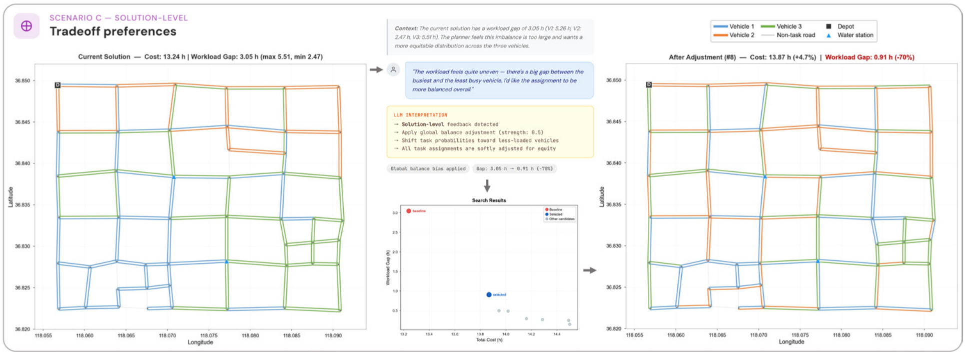

Figure 4: Illustration of solution-level interaction.

Scenario D: Composite Preference (Multi-Level)

In this final scenario, the planner expresses all three types of preferences simultaneously in a single statement, combining node-level, regional-level, and solution-level feedback: "Assign T104–T105 to Vehicle 3 for deep cleaning. Consolidate northern tasks under Vehicle 2. And try to keep the overall workload balanced."

The LLM decomposes this composite statement into three components. At the node level, it detects the request to bind T104 and T105 to Vehicle 3. At the regional level, it identifies the 32 northern tasks and binds them to Vehicle 2. At the solution level, it applies a global balance bias. The system resolves potential conflicts through priority ordering: the global balance adjustment is applied first to establish the overall direction, then the regional and node-level bindings overlay it to enforce specific preferences. After search, the planner selects solution #86, which satisfies all three preferences: T104 and T105 are assigned to Vehicle 3, the northern region is consolidated under Vehicle 2, and the workload gap is reduced to 0.97 hours (–68%) with a total cost of 13.21 hours (–0.2%). Figure 5 shows the complete multi-level interaction and its result.

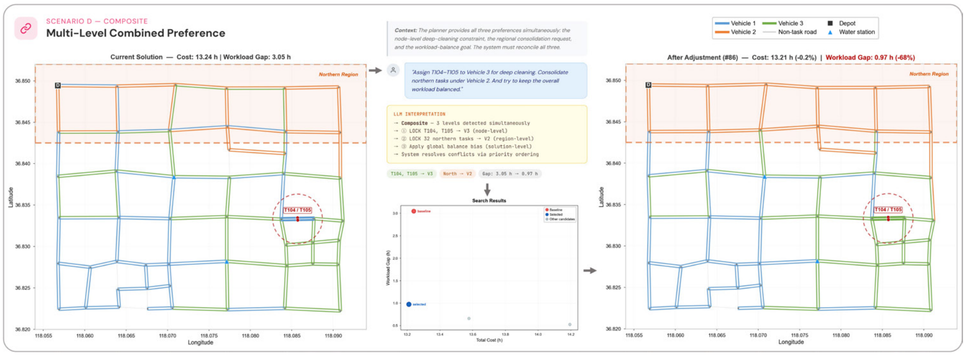

Figure 5: Illustration of multi-level interaction.

DISCUSSION

The scenario-based demonstrations illustrate how multi-level natural language interaction can capture implicit preferences that are difficult to express through conventional interaction mechanisms. Three observations emerge from the demonstrations. First, the three-level structure—node, regional, and solution—aligns with how planners naturally evaluate routing solutions. Rather than forcing feedback into a single predefined format, the approach allows planners to communicate at whatever granularity their observations naturally take. Second, natural language enables richer expression than conventional interaction methods: planners convey not only what should change but also why (e.g., deep-cleaning capability, route connectivity) and to what degree (e.g., "a bit more balanced" versus a binary priority toggle). As shown in Scenario B, a single sentence about "the northern region" triggered the identification and adjustment of 32 tasks—an operation that would require individually reassigning 11 tasks through manual clicking. Third, as demonstrated in Scenario D, natural language supports composite, multi-level preference expression within a single utterance. This reflects a qualitative difference in the kind of feedback that can be captured: the planner need not decompose their holistic judgment into discrete interface actions, as the LLM handles this decomposition automatically.

The LLM's role as a semantic bridge is worth noting. In this approach, the LLM neither replaces the algorithm nor replaces the human evaluator. Instead, it amplifies the planner's ability to communicate preferences to the algorithm. The planner remains the decision-maker throughout—inspecting solutions, forming judgments, and making final selections. The LLM translates the planner's natural language into structured probability adjustments that the search algorithm can act on. This design preserves human agency while expanding the bandwidth of human–algorithm communication.

Several limitations should be acknowledged. The current mapping relies on structured prompt templates that may not fully capture every subtlety of natural language—for instance, conditional or comparative preferences are not yet supported at a fine-grained level. The approach has been illustrated through designed scenarios using a single problem instance, rather than validated through controlled user experiments. Additionally, the LLM may occasionally misinterpret ambiguous feedback, which has not been systematically evaluated.

Future work will pursue several directions. We are developing an interactive platform for real-time use, and plan comparative user experiments evaluating natural language interaction against conventional methods in terms of preference capture richness, task performance, and user experience. We also intend to explore learning-based mapping mechanisms that adapt to individual planners' language patterns over time, improving mapping fidelity through observation of how planners' verbal feedback corresponds to their subsequent solution selections. Finally, we are investigating the extension of the LLM's role to include proactive interaction behaviors (Chen et al., 2026)—where the system initiates clarifying questions when feedback is ambiguous, suggests potential improvements based on detected patterns, or alerts planners when expressed preferences conflict with each other—moving the system from a passive translator toward a more active collaborative partner.

CONCLUSION

This study proposed a multi-level natural language interaction approach for interactive vehicle routing, in which an LLM serves as a semantic bridge between human planners and optimization algorithms. The approach enables planners to express implicit preferences at three levels—node, regional, and solution—through natural language, which the LLM translates into local and global adjustments to a probability-based preference matrix guiding the algorithmic search. Scenario-based demonstrations on a real-world road cleaning network illustrated how task-binding preferences, regional service preferences, tradeoff preferences, and their combinations can be captured and integrated through this approach. This work takes a step toward more intuitive human–algorithm collaboration in vehicle routing, where the richness of human judgment can be more fully communicated to and utilized by optimization algorithms.

ACKNOWLEDGMENT

This work was supported by the National Natural Science Foundation of China (Grant Nos. 72192824 and 72192822).

REFERENCES

Baker, B. M., & Carreto, C. A. (2003). A visual interactive approach to vehicle routing. Computers & Operations Research, 30(3), 321–337. https://doi.org/10.1016/S0305-0548(01)00099-5

Canoy, R., Bucarey, V., Mandi, J., & Guns, T. (2023). Learn and route: Learning implicit preferences for vehicle routing. Constraints, 28(3), 363–396. https://doi.org/10.1007/s10601-023-09363-2

Chen, M., Ma, L., & Wang, X. (2025). Vehicle Routing Incorporating Implicit Preferences: An Omnidimensional Human–Algorithm Collaboration Approach. IEEE Transactions on Human-Machine Systems, 55(5), 834–843. https://doi.org/10.1109/THMS.2025.3594577

Chen, M., Wang, X., Cai, L., & Ma, L. (2025). Designing visually and operationally attractive routes to improve driver acceptance in road cleaning vehicle routing problem. Computers & Operations Research, 177, 106973. https://doi.org/10.1016/j.cor.2025.106973

Chen, M., Zhou, X., & Ma, L. (2026). Designing Proactive Behaviors in LLM-Based Systems: A Scoping Review. Available at SSRN 5982734. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5982734

Liu, J., Dwyer, T., Tack, G., Gratzl, S., & Marriott, K. (2020). Supporting the problem-solving loop: Designing highly interactive optimisation systems. IEEE Transactions on Visualization and Computer Graphics, 27(2), 1764–1774. https://doi.org/10.1109/TVCG.2020.3030364

Liu, F., Tong, X., Yuan, M., Lin, X., Luo, F., Wang, Z., Lu, Z., & Zhang, Q. (2024). Evolution of Heuristics: Towards Efficient Automatic Algorithm Design Using Large Language Model (arXiv:2401.02051). arXiv. https://doi.org/10.48550/arXiv.2401.02051

Meignan, D., Knust, S., Frayret, J.-M., Pesant, G., & Gaud, N. (2015). A Review and Taxonomy of Interactive Optimization Methods in Operations Research. ACM Transactions on Interactive Intelligent Systems, 5(3), 1–43. https://doi.org/10.1145/2808234

Romera-Paredes, B., Barekatain, M., Novikov, A., Balog, M., Kumar, M. P., Dupont, E., Ruiz, F. J., Ellenberg, J. S., Wang, P., & Fawzi, O. (2024). Mathematical discoveries from program search with large language models. Nature, 625(7995), 468–475. https://doi.org/10.1038/s41586-023-06924-6

Rossit, D. G., Vigo, D., Tohmé, F., & Frutos, M. (2019). Visual attractiveness in routing problems: A review. Computers & Operations Research, 103, 13–34. https://doi.org/10.1016/j.cor.2018.10.012

Tezcaner, D., & Köksalan, M. (2011). An Interactive Algorithm for Multi-objective Route Planning. Journal of Optimization Theory and Applications, 150(2), 379–394. https://doi.org/10.1007/s10957-011-9838-y

Tütüncü, G. Y., Carreto, C. A., & Baker, B. M. (2009). A visual interactive approach to classical and mixed vehicle routing problems with backhauls. Omega, 37(1), 138–154. https://doi.org/10.1016/j.omega.2006.11.001

Vidal, T., Laporte, G., & Matl, P. (2020). A concise guide to existing and emerging vehicle routing problem variants. European Journal of Operational Research, 286(2), 401–416. https://doi.org/10.1016/j.ejor.2019.10.010

Wenger, W., & Geiger, M. J. (2008). Hybrid interactive planning under many objectives: An application to the vehicle routing problem. 2008 Eighth International Conference on Hybrid Intelligent Systems, 302–307. https://ieeexplore.ieee.org/abstract/document/4626646/

Willemse, E. J., & Joubert, J. W. (2016). Constructive heuristics for the mixed capacity arc routing problem under time restrictions with intermediate facilities. Computers & Operations Research, 68, 30–62. https://doi.org/10.1016/j.cor.2015.10.010

Xin, B., Chen, L., Chen, J., Ishibuchi, H., Hirota, K., & Liu, B. (2018). Interactive multiobjective optimization: A review of the state-of-the-art. IEEE Access, 6, 41256–41279. https://doi.org/10.1109/ACCESS.2018.2856832