

# 000 VRPAGENT: LLM-DRIVEN DISCOVERY OF HEURIS- 001 002 TIC OPERATORS FOR VEHICLE ROUTING PROBLEMS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT 009 010

011 Designing high-performing heuristics for vehicle routing problems (VRPs) is a  
012 complex task that requires both intuition and deep domain knowledge. Large lan-  
013 guage model (LLM)-based code generation has recently shown promise across  
014 many domains, but it still falls short of producing heuristics that rival those crafted  
015 by human experts. In this paper, we propose VRPAGENT, a framework that inte-  
016 grates LLM-generated components into a metaheuristic and refines them through  
017 a novel genetic search. By using the LLM to generate problem-specific operators,  
018 embedded within a generic metaheuristic framework, VRPAGENT keeps tasks  
019 manageable, guarantees correctness, and still enables the discovery of novel and  
020 powerful strategies. Across multiple problems, including the capacitated VRP,  
021 the VRP with time windows, and the prize-collecting VRP, our method discov-  
022 ers heuristic operators that outperform handcrafted methods and recent learning-  
023 based approaches while requiring only a single CPU core. To our knowledge,  
024 VRPAGENT is [among the first LLM-based paradigms](#) to advance the state-of-the-  
025 art in VRPs, highlighting a promising future for automated heuristics discovery.  
026

## 027 1 INTRODUCTION

028 Solving combinatorial optimization problems requires sophisticated solution approaches. This is  
029 especially true for vehicle routing problems (VRPs), where real-world instances often involve  
030 complex constraints and a large number of customers. Over the past decades, operations researchers have  
031 developed countless heuristics to address these problems ([Konstantakopoulos et al., 2022](#)). Designing  
032 a new method that meaningfully improves upon the state of the art across multiple problems  
033 is extremely challenging, requiring years of experience and a deep understanding of both general  
034 heuristics and the specific problem at hand. Practitioners face similar challenges when applying  
035 heuristics. Real-world applications often involve ever-changing requirements that are not supported  
036 by existing solution methods. Adapting approaches from the literature to such requirements is a  
037 time-consuming and challenging task, and is often considered impractical, even by large companies.  
038

039 In recent years, neural combinatorial optimization (NCO) has garnered increasing attention due to  
040 its potential [to automate the discovery of effective heuristics](#) ([Bello et al., 2017](#); [Bengio et al., 2021](#)).  
041 NCO approaches aim to solve optimization problems by training deep neural networks, typically  
042 with reinforcement learning. While NCO methods have demonstrated the ability to learn powerful  
043 solution strategies for various combinatorial problems, they also come with notable limitations.  
044 First, they require expensive GPUs at test time, which restricts their practical deployment. Second,  
045 scalability remains a major challenge as their reliance on attention mechanisms makes it difficult to  
046 apply these models to problems that involve processing full distance matrices. Finally, the learned  
047 strategies are often difficult for experts to interpret, which raises concerns about their safety and  
reliability in real-world applications.

048 The recent advent of performant large language models (LLMs) has enabled new opportunities for  
049 automation in general algorithmic design across domains ranging from code synthesis to symbolic  
050 planning and mathematical discovery ([Madaan et al., 2023](#)). LLMs have [proven to be](#) promising ap-  
051 proach for discovering new heuristics in combinatorial optimization problems: they can be used to  
052 design new heuristics from scratch or adapt existing ones to real-world requirements, enabling cus-  
053 tomized solutions at a fraction of the cost of an operations research (OR) expert. Among pioneering  
works, [Romera-Paredes et al. \(2024\)](#); [Liu et al. \(2024a\)](#); [Ye et al. \(2024a\)](#) propose evolutionary

frameworks that iteratively evolve general CO problem heuristics. Recent research has focused on increasingly sophisticated approaches for automating heuristic discovery (Dat et al., 2025; Zheng et al., 2025; Yang et al., 2025b; Novikov et al., 2025; Liu et al., 2025). Although these works provide valuable contributions, the discovered heuristics still fall short of those designed by human experts for VRPs. We identify key limitations in most of these works: design of end-to-end functions, absence of correctness guardrails and overall solution frameworks, and inefficient exploration of the search space, which leads to a failure in challenging the state-of-the-art.

We introduce VRPAGENT, a novel approach that uses LLMs to design heuristic operators for a large neighborhood search (LNS). The high-level LNS is designed to be largely problem-agnostic, allowing our framework to tackle new problems by creating new heuristic operators with minimal human input. To discover strong operators for the LNS, we employ a genetic algorithm (GA) with elitism and biased crossover that iteratively improves operator quality. By focusing on a metaheuristic framework where only problem-specific operators are generated via LLMs, we keep the generation task manageable and effective, while still enabling strong performance on complex problems. We evaluate our method on the capacitated vehicle routing problem (CVRP), the vehicle routing problem with time windows (VRPTW), and the prize-collecting VRP (PCVRP). Our approach discovers heuristic operators for all problems that significantly outperform those designed by human experts.

In summary, we make the following contributions with VRPAGENT:

- We propose an LNS-based metaheuristic in which the problem-specific heuristic operators are generated by an LLM.
- We introduce a simple GA for heuristic discovery featuring elitism with biased crossover for improved exploitation, and a code length penalty to reduce LLM inference costs.
- We show that VRPAGENT discovers strong heuristics across multiple tasks. To the best of our knowledge, it is the first LLM-based approach to advance the state-of-the-art in VRPs.

## 2 RELATED WORK

**Traditional Heuristics** VRPs are ubiquitous problems in logistics that have been studied for decades. Large and richly constrained instances remain difficult to solve to optimality within a practical time frame, and thus heuristics are commonly used in real-world settings (Santini et al., 2023). LNS and its adaptive variants are particularly influential, iteratively destroying and repairing parts of a solution (Shaw, 1998; Schrimpf et al., 2000; Christiaens & Vanden Berghe, 2020). Other well-known approaches include LKH3 (Helsgaun, 2017) and hybrid genetic search (HGS) (Vidal, 2022; Wouda et al., 2024). While capable of producing high-quality solutions, these approaches take significant expertise to design and implement, motivating the need for automating their design.

**Neural Combinatorial Optimization** NCO aims to automate heuristic design by training neural networks from data or via reinforcement learning (Bengio et al., 2021; Berto et al., 2025a; Li et al., 2025b). Approaches can be broadly divided into construction and improvement methods. Construction methods, such as pointer networks (Vinyals et al., 2015; Bello et al., 2017) and subsequent attention-based models for VRPs (Kool et al., 2019; Kwon et al., 2020; Kim et al., 2022; Berto et al., 2025b; Huang et al., 2025a), generate complete solutions quickly in an autoregressive fashion. Further works include enhancements for diversity (Grinsztajn et al., 2023; Hottung et al., 2025a) and out-of-distribution robustness (Drakulic et al., 2023; Luo et al., 2023). Improvement methods instead refine existing solutions at test time, for example, by learning local edits (Ma et al., 2021), guiding  $k$ -opt moves (Wu et al., 2019; da Costa et al., 2020; Ma et al., 2023), or integrating with metaheuristic approaches as LNS (Hottung & Tierney, 2020) or ant colony optimization (Ye et al., 2023; Kim et al., 2025). Divide-and-conquer frameworks further extend scalability to large instances (Kim et al., 2021; Li et al., 2021; Ye et al., 2024b; Ouyang et al., 2025). Hottung et al. (2025b) adopts a LNS approach with learned heuristics for deconstruction and ordering VRP nodes, showing competitive results against state-of-the-art solvers. Despite continuous progress, most NCO work still falls short of state-of-the-art handcrafted solvers and requires expensive GPU resources, motivating our use of LLM-generated operators as a lightweight alternative.

**Automated Heuristic Discovery** The goal of automatically discovering high-performing heuristics is a long-standing challenge in optimization (Muth, 1963). Early works include genetic pro-

gramming and hyper-heuristics, which construct new solution methods by combining or tuning a set of low-level heuristic components (Burke et al., 2006; 2013) and grammar-based generation (Mascia et al., 2014). The recent advent of LLMs has enabled a new wave of automation for algorithmic design across domains ranging from code synthesis to symbolic planning and mathematical discovery (Madaan et al., 2023; Shinn et al., 2023; Novikov et al., 2025). Early works in heuristic discovery with LLMs including Romera-Paredes et al. (2024); Liu et al. (2024a) employ evolutionary approaches that generate heuristic code snippets for simple heuristics in combinatorial problems, including VRPs, packing, and scheduling. Building on this trend, reflection-augmented evolution has been shown to discover more sophisticated heuristics during the refinement process (Ye et al., 2024a). Orthogonal search strategies further expand the design space: diversity-driven evolution (Dat et al., 2025), Monte Carlo tree search (Zheng et al., 2025), ensembling of different LLMs (Novikov et al., 2025), meta-prompt optimization (Shi et al., 2025), and portfolio-style discovery of sets of complementary heuristics (Yang et al., 2025b; Liu et al., 2025). More specifically for VRPs, Tran et al. (2025) design heuristics to enhance NCO model decoding. In parallel, finetuning and instruction specialization of LLMs for algorithmic synthesis have been proposed to improve reliability and sample efficiency (Šurina et al., 2025; Huang et al., 2025b; Chen et al., 2025b), and benchmark suites have begun to standardize evaluation protocols for LLM-driven heuristics (Liu et al., 2024b; Sun et al., 2025; Feng et al., 2025; Li et al., 2025a; Chen et al., 2025a).

Despite encouraging progress, LLM-generated heuristics still lag behind traditional solvers and NCO methods alike on VRPs, especially under tight time budgets and realistic constraints. We identify three recurring limitations: (i) weak “agentic playground” formulations that ask LLMs to design small heuristic snippets without an overall solution framework; (ii) weak or absent correctness guards around generated code; and (iii) inefficient exploration that drifts toward verbose, brittle implementations. Our approach follows the principle of *keeping AI agents on a leash*: we constrain the search to problem-specific operators nested within a robust, correctness-enforcing metaheuristic. VRPAGENT’s design keeps the synthesis task tractable, preserves feasibility, and still enables the discovery of novel operators that can advance the state-of-the-art for VRP solving.

### 3 VEHICLE ROUTING PROBLEMS

VRPs are a fundamental class of combinatorial optimization problems with the aim to minimize travel costs while respecting some constraints. Travel cost is usually measured by the total distance traveled. Formally, a VRP is defined on a graph  $G = (V, E)$ , where each node  $i \in V$  denotes a customer and each edge  $(i, j) \in E$  models traveling from  $i$  to  $j$  with an associated cost, e.g., the distance between  $i$  and  $j$ . All routes originate from and end at the depot node 0. In the **CVRP**, vehicles performing the routes have a limited capacity. The total demand on any route cannot exceed the vehicle’s capacity  $C$  at any time, and every customer is served exactly once. The **VRPTW** extends this setting, assigning a service time  $s_i$  and a time window  $[t_i^l, t_i^r]$  to each customer  $i$ . The service to any customer must start within their time window, i.e., if a vehicle arrives before a customer’s time window starts, it has to wait until  $t_i^l$ . The **PCVRP** relaxes the requirement of visiting all customers. Servicing a customer  $i$  is associated with a prize  $p_i$ . The objective is to maximize the total collected prize while minimizing travel cost.

### 4 VRPAGENT

VRPAGENT is a framework for solving VRPs that automatically discovers strong heuristic operators using LLMs. It is built on two main components. The first (Section 4.1) is an LNS (Shaw, 1998) variant for VRPs, which relies on heuristic operators to improve solutions iteratively. At test time, this LNS produces solutions for VRP instances on a single CPU core. The second component (Section 4.2) is a GA used in a discovery phase to generate these heuristic operators. In the discovery phase, operator implementations are iteratively created, modified, and refined with the help of an LLM. Classic genetic operations such as crossover and mutation are applied to operator implementations, with the LLM carrying out these transformations. Each generated operator is evaluated by inserting it into the LNS and testing the resulting search performance on a set of training instances. The performance on these instances defines the fitness value of the individual. Over successive generations, the GA produces increasingly effective heuristic operators. Fig. 1 shows a high-level overview of VRPAGENT.

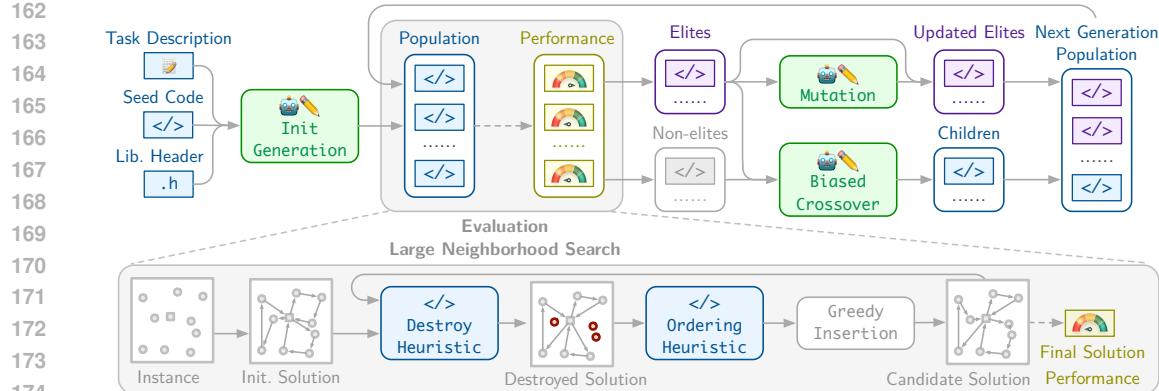


Figure 1: VRPAGENT overview.

#### 4.1 LARGE NEIGHBORHOOD SEARCH WITH LLM-GENERATED OPERATORS

VRPAGENT employs an LNS with LLM-generated operators to find solutions for VRPs. The high-level LNS guides the search process and ensures feasibility by acting as a safeguard around the LLM-generated code. In short, the LNS works by repeatedly removing a set of customers from their tours, ordering the removed customers, and reinserting them one by one at their locally optimal positions. The removal and ordering strategies are defined by LLM-written heuristic operators.

Algorithm 1 presents the pseudocode of our LNS. The algorithm begins by generating an initial solution  $s$  for a given instance  $l$ . For all routing problems, this initial solution is constructed with one tour per customer. The solution is then iteratively improved until a termination criterion is met. In each iteration,  $s$  is first destroyed by removing customers from their tours using the LLM-generated removal operator  $f_{\text{REMOVE}}$ . This yields an incomplete solution  $s'$  in which the removed customers are unassigned. Next, the removed customers are ordered by the LLM-generated operator  $f_{\text{ORDER}}$ . They are then reinserted one by one in that order, always placed at the locally best position. That is, the insertion that increases the objective value as little as possible. Finally, an acceptance decision is made:  $s'$  may replace  $s$  only if it is better, or it may be accepted under a simulated annealing rule. After all iterations, the algorithm returns the best solution  $s$ .

#### Algorithm 1 VRPAGENT-LNS

**Input:** CVRP Instance  $l$ , Destroy Operator  $f_{\text{REMOVE}}$ , Ordering Operator  $f_{\text{ORDER}}$

```

1: function LNS( $l, f_{\text{REMOVE}}, f_{\text{ORDER}}$ )
2:    $s \leftarrow \text{GENERATESTARTSOLUTION}(l)$ 
3:   while termination criteria not reached do
4:      $s' \leftarrow f_{\text{REMOVE}}(l, s)$                                  $\triangleright$  Remove some customers from their tours (LLM-Operator)
5:      $insertionOrder \leftarrow f_{\text{ORDER}}(l, s')$                    $\triangleright$  Order the removed customers (LLM-Operator)
6:     for  $c$  in  $insertionOrder$  do                                 $\triangleright$  Reinsert removed customers one by one
7:       Insert customer  $c$  at their locally optimal position in  $s'$ 
8:     end for
9:      $s \leftarrow \text{ACCEPT}(s, s')$ 
10:   end while
11:   return  $s$ 
12: end function

```

#### 4.2 HEURISTIC DISCOVERY

VRPAGENT discovers heuristic operators using a simple GA that is strongly geared toward exploitation. Given the very large search space and limited search budget, this bias toward exploitation proves highly beneficial, leading to significant improvements in our experiment. Each individual in our GA represents an implementation of the operator pair  $(f_{\text{REMOVE}}, f_{\text{ORDER}})$  as C++ code. During

216 the discovery phase, new individuals are created through the means of mutation and crossover. To  
 217 evaluate an individual, we run VRPAGENT-LNS using its operator pair on a set of training instances.  
 218

219 **Algorithm 2** outlines the core logic of our GA. It takes as input the initial population size  $M_{\text{init}}$ , the  
 220 number of elites  $M_E$ , and the number of offspring  $M_C$  generated in each iteration. The algorithm  
 221 begins by creating an initial population (i.e., a set) of heuristic operators and then enters the main  
 222 evolutionary loop, which runs until a termination criterion is reached. At the start of each iteration,  
 223 all individuals are evaluated, and the top  $M_E$  are placed in the set of elites  $\mathcal{P}_E$ , and the remainder in  
 224 the set of non-elites  $\mathcal{P}_{NE}$ . Next,  $M_C$  offspring are created by pairing one elite with one non-elite  
 225 individual and combining them using biased crossover. This crossover favors the elite parent while  
 226 still injecting diversity from the non-elite.

---

227 **Algorithm 2** VRPAGENT-GA

---

228 **Input:** Initial population size  $M_{\text{init}}$ , number of elites  $M_E$ , number of offspring  $M_C$

229 1: **function**  $\text{GA}(M_{\text{init}}, M_E, M_C)$

230 2:    $\mathcal{P} \leftarrow \text{GENERATESTARTPOP}(M_{\text{init}})$  ▷ Initialize population

231 3:   **while** termination criteria not reached **do**

232 4:      $(\mathcal{P}_E, \mathcal{P}_{NE}) \leftarrow \text{TOP-K-ELITE}(\mathcal{P}, M_E)$  ▷ Rank heuristics and take the top  $M_E$  as elite

233 5:      $C \leftarrow \emptyset$

234 6:     **while**  $|C| < M_C$  **do**

235 7:        $p_e \leftarrow \text{RANDOM}(\mathcal{P}_E)$  ▷ Select random elite

236 8:        $p_{ne} \leftarrow \text{RANDOM}(\mathcal{P}_{NE})$  ▷ Select random non-elite

237 9:        $p_c \leftarrow \text{BIASED-CROSSOVER}(p_e, p_{ne})$  ▷ Generate new heuristic via crossover

238 10:       $C \leftarrow C \cup \{p_c\}$

239 11:     **end while**

240 12:     **for all**  $p_e \in \mathcal{P}_E$  **do**

241 13:        $p_m \leftarrow \text{MUTATION}(p_e)$  ▷ Improve each elite via mutation

242 14:       **if**  $\text{FIT}(p_m) < \text{FIT}(p_e)$  **then**

243 15:            $p_e \leftarrow p_m$  ▷ Modify heuristic using mutation prompt

244 16:       **end if** ▷ Replace elite if improved

245 17:     **end for**

246 18:      $\mathcal{P} \leftarrow \mathcal{P}_E \cup C$  ▷ Create next generation

247 19:   **end while**

248 20:   **return**  $\text{BEST}(\mathcal{P})$

249 21: **end function**

---

250 Each elite is refined through mutation. Unlike standard GAs, where mutation is typically applied to  
 251 offspring to promote exploration, we apply it directly to elites. These mutations are small, making  
 252 the procedure more exploitation-focused. If a mutated elite achieves better fitness, it replaces the  
 253 original; otherwise, the original is kept. This replacement rule ensures that mutated elites do not  
 254 accumulate in the population, thereby preventing premature convergence. Finally, the next genera-  
 255 tion is formed by combining the elites with the newly generated offspring. The process repeats  
 256 until termination, after which the best heuristic operator discovered is returned. We describe the  
 257 initialization, crossover, mutation, and fitness evaluation steps in the following paragraphs.

258 **Initial Population Generation** The initial population is created by prompting the LLM to gener-  
 259 ate implementations of the two operators. For this, the LLM is provided with a global system context  
 260 that explains the overall problem and the LNS, a trivial example implementation of both operators,  
 261 and technical details of the LNS implementation in the form of C++ header files that allow operators  
 262 to access shared variables and methods efficiently (see [Prompts 1, 2](#) and [9](#) in [Appendix A](#)).

263 **Biased Crossover** The offspring are produced by combining two parents through crossover. We  
 264 use biased crossover, which pairs an elite individual with a non-elite one. The LLM is given both  
 265 their implementations and is prompted to take most ideas and concepts from the implementation of  
 266 the elite individuals and only a predefined % from the implementation of the non-elite individual.  
 267 This adaptation of standard crossover significantly increases exploitation. The complete crossover  
 268 instruction provided to the LLM corresponds to [Prompts 1](#) and [3](#) in [Appendix A](#).

269 **Mutation** VRPAGENT uses mutation to slightly modify elite implementations. Following [Liu](#)  
 270 et al. (2024a), we implement multiple mutation prompts that focus on different areas of improvement

270 for better exploration. More precisely, the LLM is given the implementation that should be modified  
 271 together with one randomly selected mutation prompt (Prompts 4 to 7 in Appendix A) and the  
 272 global system context Prompt 1. The four supported mutations include *Ablation* (i.e., removing  
 273 a random mechanic), *Extend* (i.e., adding a new mechanic), *Adjust-Parameters* (i.e., changing the  
 274 hyperparameter settings), and *Refactor* (i.e., modify the code so that the runtime is improved).  
 275

276 **Fitness Function with Code Length Penalty** We evaluate an individual  $i$  by running VRPA-  
 277 GENT-LNS with the operator pair defined by  $i$  on a set of training instances  $I^{\text{train}}$ . The resulting  
 278 solutions are then used to compute the fitness value. Specifically, the fitness of  $i$  is defined as the  
 279 average objective value across all training instances, plus a penalty proportional to the length (i.e.,  
 280 number of lines) of the corresponding implementation  $\mathcal{C}_i$ :

$$\text{Fit}(i) = \frac{1}{|I^{\text{train}}|} \sum_{j \in I^{\text{train}}} \text{Obj}(s_{i,j}) + \lambda \cdot \text{Len}(\mathcal{C}_i), \quad (1)$$

285 where  $s_{i,j}$  is the solution obtained by applying VRPAGENT-LNS with the operators of individual  $i$   
 286 to training instance  $j$ , and  $\lambda$  controls the strength of the code length penalty.

287 The penalty helps prevent uncontrolled growth in implementation size, which we observed when  
 288 no regularization was applied. More compact implementations are also easier for humans to inter-  
 289 pret and maintain, making the generated heuristics more useful in practice. Finally, a shorter code  
 290 reduces the number of tokens processed by the LLM during generation, which in our experiments  
 291 lowers token usage by more than 50%, and thus significantly reduces generation costs and latency.

## 293 5 EXPERIMENTS

295 We evaluate VRPAGENT on three vehicle routing problems: the CVRP, the VRPTW, and the  
 296 PCVRP. The best discovered heuristics are made publicly available in our online repository at  
 297 <https://anonymous.4open.science/r/vrpagent-submission>.

299 **VRPAGENT Hyperparameters** During the discovery phase we use the following parameters un-  
 300 less stated otherwise: elite size  $M_E = 10$ , offspring size  $M_C = 30$ , an initial population size of  
 301  $M_{\text{init}} = 100$ . The code length penalty factor is set to  $\lambda = 2 \cdot 10^{-4}$  and the discovery phase is ter-  
 302 minated after 40 iterations. Individuals are evaluated with VRPAGENT-LNS on a training set of 64  
 303 instances, each with 500 customers, using a runtime limit of 20s per instance. We employ Gemini  
 304 2.5 Flash (Comanici et al., 2025) as the LLM.

### 305 5.1 COMPARISON TO STATE-OF-THE-ART

307 **Benchmark Setup** We evaluate all methods on the three problems using instance sets of 500,  
 308 1000, and 2000 customers. To ensure consistency, we adopt the same test instances and baseline  
 309 configurations as described in Hottung et al. (2025b) using a single core of an AMD Milan 7763  
 310 processor and an additional single NVIDIA A100 for approaches that require a GPU. For a fair  
 311 comparison, we limit the search by runtime when possible. The operators used by VRPAGENT are  
 312 obtained from 10 discovery runs per problem (conducted on instances of size 500 only), with the  
 313 best operator selected based on performance on a separate validation set. Additionally, we provide  
 314 an evaluation of VRPAGENT on the more realistic X instances in Appendix C, which better reflect  
 315 the properties encountered in real-world scenarios.

316 **Baselines** We compare VRPAGENT to several established operations research solvers: HGS (Vi-  
 317 dal, 2022), SISRs (Christiaens & Vanden Berghe, 2020), LKH3 (Helsgaun, 2017), and Gurobi  
 318 (Gurobi Optimization, LLC, 2024). We also include PyVRP (Wouda et al., 2024) (version 0.9.0),  
 319 an open-source extension of HGS that supports additional VRP variants, and the recent GPU-based  
 320 NVIDIA cuOpt (NVIDIA Corporation, 2025). For the CVRP, we further consider learning-based  
 321 approaches that require GPUs at test time: BQ (Drakulic et al., 2023), LEHD (Luo et al., 2023),  
 322 UDC (Zheng et al., 2024b), and NDS (Hottung et al., 2025b). In addition, we compare against  
 323 LLM-based methods that learn a construction heuristic, including EoH (Liu et al., 2024a), MCTS-  
 324 AHD (Zheng et al., 2024a), and ReEvo (Ye et al., 2024a). These approaches only generate a single

324 solution with runtime values <1 second and do not benefit from additional search budget. We fur-  
 325 ther compare against three stronger LLM-based baselines that leverage search: ReEvo-ACO, which  
 326 combines LLM-generated heuristics with the Ant Colony Optimization (ACO) metaheuristic, NCO-  
 327 LLM (Tran et al., 2025), which enhances LEHD by automating the design of logit reshaping during  
 328 search, and LLM-LNS (Ye et al., 2025), a recent method using LLMs to design search heuristics.  
 329

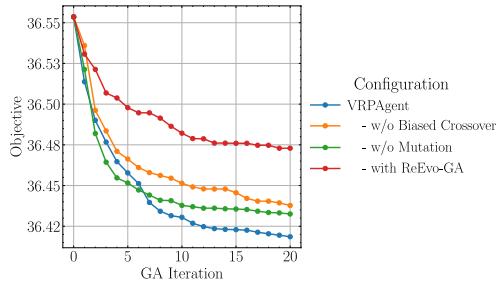
330 **Results** Table 1 presents the results of our experiments. On instances with 1000 and 2000  
 331 customers, VRPAGENT outperforms all other methods across all problem types, achieving  
 332 gaps of around  $-0.30\%$  relative to the state-of-the-art SISR. This represents a substan-  
 333 tial improvement that can translate into significant savings in large-scale, real-world scenarios.  
 334 On smaller instances, VRPAGENT approaches the  
 335 performance of NDS, which relies on an expensive  
 336 GPU at test time and is trained specifically for each  
 337 instance size. Compared to other LLM-based meth-  
 338 ods, VRPAGENT consistently demonstrates signif-  
 339 icantly better performance across all test cases. **It is**  
 340 **also noteworthy that VRPAGENT performs well on**  
 341 **all problem sizes, despite being trained only on in-**  
 342 **stances with 500 customers. This strong generaliza-**  
 343 **tion performance stems from the algorithmic op-**  
 344 **erators. These usually do not rely on scale-dependent**  
 345 **neural representations, which are poisoned to im-**  
 346 **plicitly encode all characteristics of the training distribution, including unwanted characteristics like**  
 347 **node count and spatial density.**

## 348 5.2 ANALYSES

349 **Ablation Studies** We analyze the contribution of key components in our GA by disabling or  
 350 replacing them. Specifically, we test three variants: (i) replacing our biased crossover with a standard  
 351 crossover, where the LLM is instructed to take roughly half the elements from each parent, (ii)  
 352 removing mutation while increasing offspring size to maintain a comparable population, and (iii)  
 353 replacing our entire GA with the GA of ReEvo (Ye et al., 2024a), which uses a reflection mech-  
 354 anism. Each variant is tested on the CVRP 10 times, and results are averaged. Fig. 2 reports  
 355 performance on the training set during the discovery phase. Across all cases, modifications lead to  
 356 reduced performance. Biased crossover is particularly important: by favoring elite solutions while  
 357 still incorporating elements from weaker parents, it balances exploitation and exploration and drives  
 358 faster convergence. Removing mutation lowers final solution quality, and replacing our GA with  
 359 ReEvo’s yields the weakest results, confirming that our combination of elitism, biased crossover,  
 360 and mutation is essential for discovering high-quality heuristics.

361 **Performance Across Different LLMs** We study the performance of VRPAGENT when paired  
 362 with different LLMs. We conduct discovery runs of 20 iterations each for the CVRP using six  
 363 models. We access Gemini 2.0 Flash and Gemini 2.5 Flash (Comanici et al., 2025) via API, while  
 364 Qwen3 (Yang et al., 2025a), Llama 3.3 (Grattafiori et al., 2024), Gemma 3 (Team et al., 2025),  
 365 and gpt-oss (Agarwal et al., 2025) are served locally via vLLM (Kwon et al., 2023). Fig. 3 reports  
 366 the average objective value on the training set during discovery (left) and the total computational  
 367 cost per run (right). All tested models substantially improve the heuristic operators throughout  
 368 the discovery process. Gemini 2.5 Flash and gpt-oss both discover heuristics that outperform the state-  
 369 of-the-art baseline. Gemini 2.5 Flash achieves the best overall results, but at a cost of nearly \$20 per  
 370 run. In contrast, the open-source gpt-oss model, run on two NVIDIA A100 (40GB) GPUs, achieves  
 371 nearly the same performance under \$2 per run.

372 **Performance Over the Discovery Process** We analyze the convergence rate of the discovery pro-  
 373 cess with Gemini 2.5 Flash on all three problems across 40 iterations. As a baseline, we report the  
 374 performance of VRPAGENT-LNS when used in combination with handcrafted operators. Specifi-  
 375 cally, we reimplement the operators from SISR (Christiaens & Vanden Berghe, 2020), which rep-  
 376 resent the state of the art in LNS-based routing methods. As shown in Fig. 4, VRPAGENT produces  
 377 heuristics that outperform the state-of-the-art (SOTA) handcrafted operators. In Appendix D, we ad-



378 Figure 2: Ablation results.

378  
379  
380  
381  
382  
383  
384385 Table 1: Performance on test data. The gap is calculated relative to SISRs. Runtime is reported on a per-  
386 instance basis in seconds. The best results (i.e., those with the lowest objective function value) are shown in  
387 **bold**, and the second-best are underlined. \* Indicates that a feasible solution was not found for all instances.  
388 **OOM** indicates that the corresponding solver yielded an out-of-memory error. **VRPAGENT** demonstrates state-  
389 of-the-art results among heuristics and NCO solvers at larger scales.

390	391	392	Method	393 N=500			394 N=1000			395 N=2000		
				396    Obj.↓	397 Gap.↓	398 Time	399 Obj.↓	400 Gap.↓	401 Time	402 Obj.↓	403 Gap.↓	404 Time
393	SISRs	CPU	36.65	-	60	41.14	-	120	56.04	-	240	
394	HGS	CPU	36.66	0.00%	60	41.51	0.84%	121	57.38	2.33%	241	
395	LKH3	CPU	37.25	1.66%	174	42.16	2.46%	408	58.12	3.70%	1448	
396	NVIDIA cuOpt	CPU+GPU	37.38	1.98%	60	42.71	3.78%	121	59.22	5.66%	241	
397	<b>Gurobi</b>	CPU	<b>47.85</b>	<b>30.53%</b>	<b>60</b>	<b>87.36</b>	<b>111.88%</b>	<b>120</b>	-	<b>OOM</b>	-	
398	CVRP	BQ (BS64)	CPU+GPU	37.51	2.34%	23	43.32	5.30%	164	-	-	-
399		LEHD (RRC)	CPU+GPU	37.04	1.06%	60	42.47	3.25%	121	60.11	7.25%	246
400		UDC	CPU+GPU	37.63	2.69%	60	42.65	3.68%	121	-	-	-
401		NDS	CPU+GPU	36.57	<b>-0.20%</b>	60	41.11	<b>-0.07%</b>	120	56.00	<b>-0.07%</b>	240
402		EoH	CPU	45.89	25.21%	<1	52.42	27.42%	<1	71.21	27.07%	<1
403		MCTS-AHD	CPU	45.51	24.17%	<1	52.49	27.59%	<1	71.15	26.96%	<1
404		ReEvo	CPU	44.21	20.63%	<1	52.23	26.96%	<1	70.01	24.93%	<1
405		ReEvo-ACO	CPU	40.25	9.83%	60	46.22	12.34%	120	63.76	13.77%	240
406		NCO-LLM	CPU+GPU	36.93	0.76%	60	41.96	1.99%	121	59.43	6.05%	246
407		<b>LLM-LNS</b>	CPU	<b>38.99</b>	<b>6.38%</b>	<b>60</b>	<b>45.00</b>	<b>9.38%</b>	<b>120</b>	<b>62.28</b>	<b>11.13%</b>	<b>240</b>
408	VRPTW	<b>VRPAGENT</b>	CPU	36.60	<b>-0.12%</b>	60	41.06	<b>-0.19%</b>	120	55.98	<b>-0.11%</b>	240
409		SISRs	CPU	48.09	-	60	87.68	-	120	167.49	-	240
410		PyVRP-HGS	CPU	49.01	1.91%	60	90.35	3.08%	120	173.46	3.62%	240
411		NVIDIA cuOpt	CPU+GPU	49.30	2.60%	61	90.31	3.11%	121	173.52	3.85%*	243
412		<b>Gurobi</b>	CPU	<b>69.01</b>	<b>43.83%</b>	<b>60</b>	<b>148.45</b>	<b>71.46%</b>	<b>120</b>	-	<b>OOM</b>	-
413		NDS	CPU+GPU	47.94	<b>-0.30%</b>	60	87.54	<b>-0.16%</b>	120	167.48	<b>-0.00%</b>	240
414		EoH	CPU	60.40	25.60%	<1	118.80	35.49%	<1	245.70	46.70%	<1
415		MCTS-AHD	CPU	58.31	21.25%	<1	113.72	29.70%	<1	231.11	37.98%	<1
416		ReEvo	CPU	58.01	20.63%	<1	110.55	26.08%	<1	218.90	30.69%	<1
417		ReEvo-ACO	CPU	52.91	10.03%	60	97.39	11.07%	120	193.13	15.31%	240
418	PCVRP	<b>VRPAGENT</b>	CPU	47.97	<b>-0.24%</b>	60	87.40	<b>-0.33%</b>	120	166.96	<b>-0.33%</b>	240
419		SISRs	CPU	43.22	-	60	81.12	-	120	158.17	-	240
420		PyVRP-HGS	CPU	44.97	4.10%	60	84.91	4.81%	120	165.56	4.78%	240
421		NVIDIA cuOpt	CPU+GPU	43.34	0.19%	60	81.89	0.84%	121	160.33	1.22%	241
422		<b>Gurobi</b>	CPU	<b>71.58</b>	<b>68.23%</b>	<b>60</b>	<b>147.08</b>	<b>86.03%</b>	<b>120</b>	-	<b>OOM</b>	-
423	<b>VRPAGENT</b>	NDS	CPU+GPU	43.12	<b>-0.23%</b>	60	80.99	<b>-0.17%</b>	121	158.09	<b>-0.06%</b>	241
424		<b>VRPAGENT</b>	CPU	43.18	<b>-0.09%</b>	60	80.95	<b>-0.21%</b>	120	157.69	<b>-0.32%</b>	240

426  
427  
428  
429  
430  
431

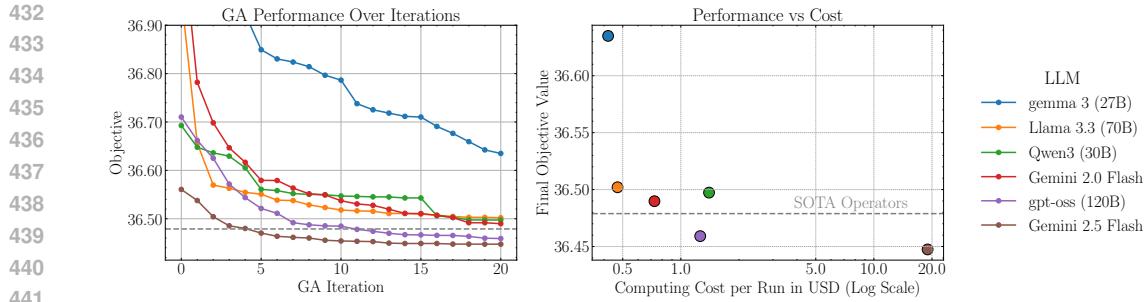


Figure 3: Performance on the CVRP for different LLMs. Detailed cost calculations are provided in Section B.1.

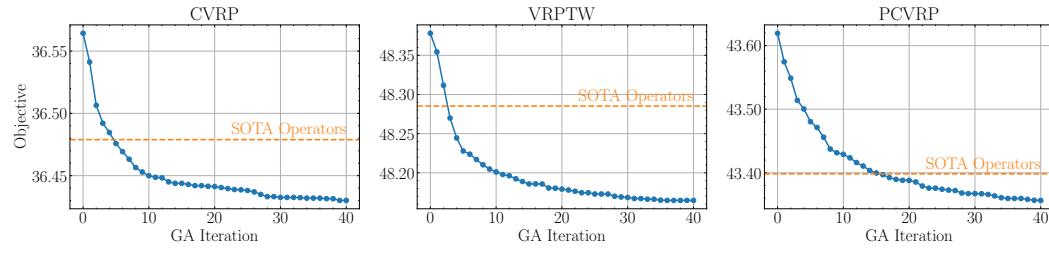


Figure 4: Performance over the course of the discovery process.

ditionally report the percentage of operators that compile successfully and produce valid solutions (success rate) throughout the runs.

**Crossover Bias and Elite Size** We evaluate the effect of the crossover bias and elite size  $M_E$  on the performance of our GA. As shown in Fig. 5a, elite sizes between 5 and 15 yield strong performance, whereas an elite size of 1 leads to drastically worse results by making the search overly greedy. Regarding the crossover bias, a strong preference for the better-performing operator is clearly beneficial, with an operator bias of 80% achieving the best results.

**Code Length Penalty** We investigate the impact of the code length penalty factor  $\lambda$  on both the quality and length of the discovered heuristics. Several discovery runs with varying  $\lambda$  values reveal that the penalty strongly controls implementation size without substantially degrading performance. For instance, increasing  $\lambda$  to  $4 \cdot 10^{-4}$  reduces the average length of generated heuristics by roughly 50%, while only causing a marginal drop in objective value (Fig. 5b). These results highlight several key insights. First, the penalty effectively discourages overly long implementations, producing heuristics that are easier for humans to read and maintain. Second, higher penalties reduce token usage during LLM generation, improving efficiency and lowering computational costs. Third, the fact that performance remains largely unaffected even for strong penalties demonstrates that simple, concise heuristics can perform just as well as more complex ones.

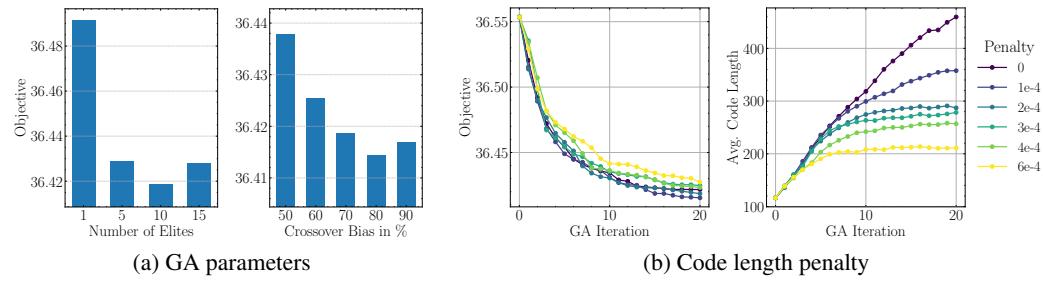


Figure 5: Results of the sensitivity analyses.

486 **Warm-starting Heuristic Discovery** We analyze whether heuristics discovered for one problem  
 487 can be adapted to others via warm-starts. Specifically, we initialize the discovery process for the  
 488 VRPTW and PCVRP using the best CVRP heuristic and task the model with adapting it to the  
 489 new problem. Warm-starting consistently accelerates convergence and improves final performance  
 490 compared to cold-start runs, making this approach promising for real-world settings where related  
 491 problems need to be solved efficiently. Detailed results and analyses are provided in Appendix E.  
 492

## 493 6 DISCOVERED HEURISTIC OPERATORS

494  
 495 A key advantage of LLM-driven heuristic discovery over deep learning-based approaches is that the  
 496 resulting strategies can be directly inspected and interpreted by human experts. To further facilitate  
 497 such analysis, we introduce a post-processing stage in which an LLM is used to improve the code  
 498 quality and readability of the generated implementations. This stage is necessary because code  
 499 quality is not considered during the discovery process itself: the fitness functions focus solely on  
 500 performance, and we explicitly instruct the LLMs not to include comments during search. While this  
 501 reduces interpretability, it substantially lowers token usage and accelerates the discovery procedure.  
 502

503 **Code quality improvement** To refine the quality of the discovered heuristics, we pass each im-  
 504 plementation to an LLM together with a prompt specifying the desired improvements. The returned  
 505 code is then evaluated on the validation instances to verify that performance remains unchanged.  
 506 We repeat this refinement–evaluation loop until we obtain an improved version that matches the  
 507 original performance. The improvement prompt emphasizes (1) adding comments that explain both  
 508 high-level heuristic logic and low-level implementation details, (2) restructuring code into clearer  
 509 functional units, (3) and eliminating unnecessarily convoluted uses of random numbers.  
 510

511 **Analysis** We give the task of analyzing the best discovered heuristic operators for each problem  
 512 to three coauthors of this paper who have years of experience writing OR heuristics in routing and  
 513 related fields. The goal of our analysis is to assess the (1) readability, (2) coherence and soundness,  
 514 (3) maintainability, (4) interpretability, (5) and novelty of the generated heuristic operators. We note  
 515 that our assessment of the heuristics is not meant to be a thorough scientific analysis that generalizes  
 516 to other LLMs or optimization problems. We offer a detailed analysis in Appendix F.  
 517

518 The removal and sorting mechanisms of all three analyzed heuristics can be described as ensem-  
 519 ble approaches that use random numbers to choose different (combinations of) heuristics in each  
 520 iteration. Given the popularity and success of such ensembles in well-known metaheuristics (e.g.,  
 521 adaptive LNS (Pisinger & Ropke, 2018)), it is perhaps not a surprise that we encounter ensembles in  
 522 the discovered heuristics. Overall, we find the generated heuristics to be relatively easy to interpret  
 523 and structurally coherent. The comments introduced during post-processing substantially enhance  
 524 comprehensibility, and the decomposition of the code into functions further improves readability.  
 525

526 All of the heuristics generated can be said to be novel. We are not aware of any heuristics in the  
 527 literature exactly matching these algorithms, however we note that the heuristics mainly consist of  
 528 recombinations of ideas existing in the literature, e.g., SISRs or simple greedy criteria related to  
 529 distance/demand/time/prizes. Given the complexity of the ensembles, with some having up to nine  
 530 different component heuristics, a detailed ablation study would be necessary to try to find out which  
 531 components or combinations of components lead to good performance.  
 532

## 533 7 CONCLUSION

534 In this work, we introduced VRPAGENT, a metaheuristic framework in which LLMs generate  
 535 problem-specific operators for a LNS. By focusing on operator generation rather than end-to-end  
 536 heuristics, VRPAGENT makes the discovery task more manageable while achieving strong perfor-  
 537 mance. Using a GA with elitism and biased crossover for algorithm discovery, VRPAGENT con-  
 538 sistently finds heuristic operators that outperform human-designed approaches on a range of vehicle  
 539 routing problems. Our results highlight a promising future for automated heuristic discovery, sug-  
 540 gesting that LLMs could play a key role in designing efficient and adaptable optimization methods  
 541 for complex, real-world problems. For future work, we will investigate how to further simplify the  
 542 generated heuristics to help increase the ease of using VRPAGENT generated code in practice.  
 543

540  
541

## REPRODUCIBILITY STATEMENT

542  
543  
544  
545

We have made every effort to ensure the reproducibility of our results. Detailed descriptions of configurations, prompts, the discovery pipeline, and overall experimental setups are provided in both the main paper and the appendix to enable independent reproducibility. All code to reproduce the experiments will be made open-source upon acceptance.

546

547

## REFERENCES

548  
549  
550  
551

Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*, 2025. URL <https://arxiv.org/abs/2508.10925>.

552  
553  
554  
555

Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy Bengio. Neural Combinatorial Optimization with Reinforcement Learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Bk9mx1SFx>.

556  
557  
558  
559  
560

Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research*, 290(2): 405–421, 2021. ISSN 0377-2217. doi: <https://doi.org/10.1016/j.ejor.2020.07.063>. URL <https://www.sciencedirect.com/science/article/pii/S0377221720306895>.

561  
562  
563  
564  
565  
566  
567

Federico Berto, Chuanbo Hua, Junyoung Park, Laurin Luttmann, Yining Ma, Fanchen Bu, Jiarui Wang, Haoran Ye, Minsu Kim, Sanghyeok Choi, Nayeli Gast Zepeda, André Hottung, Jianan Zhou, Jieyi Bi, Yu Hu, Fei Liu, Hyeonah Kim, Jiwoo Son, Haeyeon Kim, Davide Angioni, Wouter Kool, Zhiguang Cao, Jie Zhang, Kijung Shin, Cathy Wu, Sungsoo Ahn, Guojie Song, Changhyun Kwon, Lin Xie, and Jinkyoo Park. RL4CO: an Extensive Reinforcement Learning for Combinatorial Optimization Benchmark. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2025a. URL <https://github.com/ai4co/r14co>.

568  
569  
570  
571

Federico Berto, Chuanbo Hua, Nayeli Gast Zepeda, André Hottung, Niels Wouda, Leon Lan, Junyoung Park, Kevin Tierney, and Jinkyoo Park. RouteFinder: Towards Foundation Models for Vehicle Routing Problems. *Transactions on Machine Learning Research*, 2025b. ISSN 2835-8856. URL <https://openreview.net/forum?id=QzGLoaOPiY>.

572  
573  
574  
575

Edmund K Burke, Matthew R Hyde, and Graham Kendall. Evolving bin packing heuristics with genetic programming. In *International Conference on Parallel Problem Solving from Nature*, pp. 860–869. Springer, 2006.

576  
577  
578

Edmund K Burke, Michel Gendreau, Matthew Hyde, Graham Kendall, Gabriela Ochoa, Ender Özcan, and Rong Qu. Hyper-heuristics: A survey of the state of the art. *Journal of the Operational Research Society*, 64(12):1695–1724, 2013.

579  
580  
581  
582  
583  
584

Hongzheng Chen, Yingheng Wang, Yaohui Cai, Hins Hu, Jiajie Li, Shirley Huang, Chenhui Deng, Rongjian Liang, Shufeng Kong, Haoxing Ren, Samitha Samaranayake, Carla P. Gomes, and Zhiru Zhang. Heurigym: An agentic benchmark for LLM-crafted heuristics in combinatorial optimization. *arXiv preprint arXiv:2506.07972*, 2025a. URL <https://arxiv.org/abs/2506.07972>.

585  
586  
587

Yitian Chen, Jingfan Xia, Siyu Shao, Dongdong Ge, and Yinyu Ye. Solver-informed RL: Grounding large language models for authentic optimization modeling. In *Advances in Neural Information Processing Systems*, 2025b.

588  
589  
590

Jan Christiaens and Greet Vanden Berghe. Slack induction by string removals for vehicle routing problems. *Transportation Science*, 54(2):417–433, 2020.

591  
592  
593

Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

594 Paulo da Costa, Jason Rhuggenaath, Yingqian Zhang, and Alp Eren Akçay. Learning 2-opt Heuristics  
 595 for the Traveling Salesman Problem via Deep Reinforcement Learning. In *Asian Conference  
 596 on Machine Learning*, 2020.

597

598 Pham Vu Tuan Dat, Long Doan, and Huynh Thi Thanh Binh. Hsevo: Elevating automatic heuristic  
 599 design with diversity-driven harmony search and genetic algorithm using llms. In *Proceedings  
 600 of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 26931–26938, 2025. <https://github.com/datphamvn/HSEvo>.

601

602 Darko Drakulic, Sofia Michel, Florian Mai, Arnaud Sors, and Jean-Marc Andreoli. BQ-  
 603 NCO: Bisimulation Quotienting for Efficient Neural Combinatorial Optimization. In A. Oh,  
 604 T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-  
 605 ral Information Processing Systems*, volume 36, pp. 77416–77429. Curran Associates, Inc.,  
 606 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/f445ba15f0f05c26e1d24f908ea78d60-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/f445ba15f0f05c26e1d24f908ea78d60-Paper-Conference.pdf).

607

608 Shengyu Feng, Weiwei Sun, Shanda Li, Ameet Talwalkar, and Yiming Yang. A comprehen-  
 609 sive evaluation of contemporary ML-based solvers for combinatorial optimization. *ArXiv*,  
 610 abs/2505.16952, 2025. URL <https://arxiv.org/abs/2505.16952>.

611

612 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad  
 613 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd  
 614 of models. *arXiv preprint arXiv:2407.21783*, 2024.

615

616 Nathan Grinsztajn, Daniel Furelos-Blanco, Shikha Surana, Clément Bonnet, and Tom Barrett. Win-  
 617 ner Takes It All: Training Performant RL Populations for Combinatorial Optimization. In  
 618 A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-  
 619 ral Information Processing Systems*, volume 36, pp. 48485–48509. Curran Associates, Inc.,  
 620 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/97b983c974551153d20ddfabb62a5203-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/97b983c974551153d20ddfabb62a5203-Paper-Conference.pdf).

621

622 Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024. URL <https://www.gurobi.com>.

623

624 Keld Helsgaun. An extension of the Lin-Kernighan-Helsgaun TSP solver for constrained traveling  
 625 salesman and vehicle routing problems. *Roskilde: Roskilde University*, 2017.

626

627 André Hottung and Kevin Tierney. Neural Large Neighborhood Search for the Capacitated Vehicle  
 628 Routing Problem. In *European Conference on Artificial Intelligence*, pp. 443–450, 2020.

629

630 André Hottung, Mridul Mahajan, and Kevin Tierney. PolyNet: Learning diverse solution strategies  
 631 for neural combinatorial optimization. In *International Conference on Learning Representations*,  
 631a.

632

633 André Hottung, Paula Wong-Chung, and Kevin Tierney. Neural deconstruction search for vehicle  
 634 routing problems. *Transactions on Machine Learning Research*, 2025b. ISSN 2835-8856. URL  
 635 <https://openreview.net/forum?id=bCmEP1Ltwq>.

636

637 Ziwei Huang, Jianan Zhou, Zhiguang Cao, and Yixin Xu. Rethinking light decoder-based solvers  
 638 for vehicle routing problems. In *The Thirteenth International Conference on Learning Represen-  
 639 tations*, 2025a. URL <https://openreview.net/forum?id=4pRwkYpa2u>.

640

641 Ziyao Huang, Weiwei Wu, Kui Wu, Jianping Wang, and Wei-Bin Lee. Calm: Co-evolution of  
 642 algorithms and language model for automatic heuristic design, 2025b. URL <https://arxiv.org/abs/2505.12285>.

643

644 Minsu Kim, Jinkyoo Park, and Joungho Kim. Learning Collaborative Policies to  
 645 Solve NP-hard Routing Problems. In M. Ranzato, A. Beygelzimer, Y. Dauphin,  
 646 P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Pro-  
 647 cessing Systems*, volume 34, pp. 10418–10430. Curran Associates, Inc., 2021. URL  
[https://proceedings.neurips.cc/paper\\_files/paper/2021/file/564127c03caab942e503ee6f810f54fd-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/564127c03caab942e503ee6f810f54fd-Paper.pdf).

648 Minsu Kim, Junyoung Park, and Jinkyoo Park. Sym-NCO: Leveraging Symmetric-  
 649     ity for Neural Combinatorial Optimization. In S. Koyejo, S. Mohamed, A. Agar-  
 650     wal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Pro-  
 651     cessing Systems*, volume 35, pp. 1936–1949. Curran Associates, Inc., 2022. URL  
 652     [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/0cddb777d3441326544e21b67f41bdc8-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/0cddb777d3441326544e21b67f41bdc8-Paper-Conference.pdf).

654 Minsu Kim, Sanghyeok Choi, Hyeonah Kim, Jiwoo Son, Jinkyoo Park, and Yoshua Bengio. Ant  
 655     Colony Sampling with GFlowNets for Combinatorial Optimization. In Yingzhen Li, Stephan  
 656     Mandt, Shipra Agrawal, and Emtiyaz Khan (eds.), *Proceedings of The 28th Interna-  
 657     tional Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning  
 658     Research*, pp. 469–477. PMLR, 2025. URL <https://proceedings.mlr.press/v258/kim25a.html>.

660 Grigoris D Konstantakopoulos, Sotiris P Gayialis, and Evripidis P Kechagias. Vehicle routing  
 661     problem and related algorithms for logistics distribution: A literature review and classifica-  
 662     tion. *Operational research*, 22(3):2033–2062, 2022.

664 Wouter Kool, Herke van Hoof, and Max Welling. Attention, Learn to Solve Routing Problems! In  
 665     *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=ByxBFsRqYm>.

667 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.  
 668     Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model  
 669     serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating  
 670     Systems Principles*, 2023.

672 Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seung-  
 673     jai Min. POMO: Policy Optimization with Multiple Optima for Reinforcement Learn-  
 674     ing. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21188–21198,  
 675     2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/f231f2107df69eab0a3862d50018a9b2-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/f231f2107df69eab0a3862d50018a9b2-Paper.pdf).

677 Sirui Li, Zhongxia Yan, and Cathy Wu. Learning to delegate for large-scale vehicle routing. *Ad-  
 678     vances in Neural Information Processing Systems*, 34:26198–26211, 2021.

679 Xiaozhe Li, Jixuan Chen, Xinyu Fang, Shengyuan Ding, Haodong Duan, Qingwen Liu, and Kai  
 680     Chen. Opt-bench: Evaluating llm agent on large-scale search spaces optimization problems,  
 681     2025a. URL <https://arxiv.org/abs/2506.10764>.

683 Yang Li, Jiale Ma, Wenzheng Pan, Runzhong Wang, Haoyu Geng, Nianzu Yang, and Junchi Yan.  
 684     Unify ml4tsp: Drawing methodological principles for tsp and beyond from streamlined design  
 685     space of learning and search. In *The Thirteenth International Conference on Learning Represen-  
 686     tations*, 2025b.

687 Fei Liu, Xialiang Tong, Mingxuan Yuan, Xi Lin, Fu Luo, Zhenkun Wang, Zhichao Lu, and Qingfu  
 688     Zhang. Evolution of heuristics: Towards efficient automatic algorithm design using large language  
 689     model. In *International Conference on Machine Learning (ICML)*, 2024a. URL <https://arxiv.org/abs/2401.02051>.

691 Fei Liu, Rui Zhang, Zhuoliang Xie, Rui Sun, Kai Li, Xi Lin, Zhenkun Wang, Zhichao Lu, and  
 692     Qingfu Zhang. LLM4AD: A platform for algorithm design with large language model. 2024b.  
 693     URL <https://arxiv.org/abs/2412.17287>.

695 Fei Liu, Yilu Liu, Qingfu Zhang, Xialiang Tong, and Mingxuan Yuan. Eoh-s: Evolution of heuristic  
 696     set using llms for automated heuristic design. *arXiv preprint arXiv:2508.03082*, 2025.

697 Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial op-  
 698     timization with heavy decoder: Toward large scale generalization. In A. Oh, T. Nau-  
 699     mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-  
 700     ral Information Processing Systems*, volume 36, pp. 8845–8864. Curran Associates, Inc.,  
 701     2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/1c10d0c087c14689628124bbc8fa69f6-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1c10d0c087c14689628124bbc8fa69f6-Paper-Conference.pdf).

702 Yining Ma, Jingwen Li, Zhiguang Cao, Wen Song, Le Zhang, Zhenghua Chen, and Jing Tang.  
 703 Learning to Iteratively Solve Routing Problems with Dual-Aspect Collaborative Transformer.  
 704 In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.),  
 705 *Advances in Neural Information Processing Systems*, volume 34, pp. 11096–11107. Curran  
 706 Associates, Inc., 2021. URL [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/5c53292c032b6cb8510041c54274e65f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/5c53292c032b6cb8510041c54274e65f-Paper.pdf).

708 Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to Search Feasible and In-  
 709 feasible Regions of Routing Problems with Flexible Neural k-Opt. In A. Oh, T. Nau-  
 710 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural*  
 711 *Information Processing Systems*, volume 36, pp. 49555–49578. Curran Associates, Inc.,  
 712 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/9bae70d354793a95fa18751888cea07d-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/9bae70d354793a95fa18751888cea07d-Paper-Conference.pdf).

714 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegr-  
 715 effe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bod-  
 716 hisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and  
 717 Peter Clark. Self-refine: Iterative refinement with self-feedback. In A. Oh, T. Nau-  
 718 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural*  
 719 *Information Processing Systems*, volume 36, pp. 46534–46594. Curran Associates, Inc.,  
 720 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/91edff07232fb1b55a505a9e9f6c0ff3-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/91edff07232fb1b55a505a9e9f6c0ff3-Paper-Conference.pdf).

722 Franco Mascia, Manuel López-Ibáñez, Jérémie Dubois-Lacoste, and Thomas Stützle. Grammar-  
 723 based generation of stochastic local search heuristics through automatic algorithm configura-  
 724 tion tools. *Computers & operations research*, 51:190–199, 2014.

726 J Muth. Probabilistic learning combinations of local job-shop scheduling rules. *Industrial scheduling*, 1963.

728 Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt  
 729 Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian,  
 730 et al. AlphaEvolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint*  
 731 *arXiv:2506.13131*, 2025. URL <https://arxiv.org/abs/2506.13131>.

733 NVIDIA Corporation. NVIDIA cuOpt: Gpu-accelerated decision optimization. <https://github.com/NVIDIA/cuopt>, 2025. Accessed: 2025-09-25.

735 Wenbin Ouyang, Sirui Li, Yining Ma, and Cathy Wu. Learning to segment for vehicle routing  
 736 problems. *arXiv preprint arXiv:2507.01037*, 2025.

738 David Pisinger and Stefan Ropke. Large neighborhood search. In *Handbook of metaheuristics*, pp.  
 739 99–127. Springer, 2018.

740 Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog,  
 741 M Pawan Kumar, Emilien Dupont, Francisco JR Ruiz, Jordan S Ellenberg, Pengming Wang,  
 742 Omar Fawzi, et al. Mathematical discoveries from program search with large language models.  
 743 *Nature*, 625(7995):468–475, 2024.

744 Alberto Santini, Michael Schneider, Thibaut Vidal, and Daniele Vigo. Decomposition strategies for  
 745 vehicle routing heuristics. *INFORMS Journal on Computing*, 35(3):543–559, 2023.

747 Gerhard Schrimpf, Johannes Schneider, Hermann Stamm-Wilbrandt, and Gunter Dueck. Record  
 748 breaking optimization results using the ruin and recreate principle. *Journal of Computational*  
 749 *Physics*, 159(2):139–171, 2000.

750 Paul Shaw. Using constraint programming and local search methods to solve vehicle routing prob-  
 751 lems. In *International conference on principles and practice of constraint programming*, pp.  
 752 417–431. Springer, 1998.

754 Yiding Shi, Jianan Zhou, Wen Song, Jieyi Bi, Yaxin Wu, and Jie Zhang. Generalizable Heuris-  
 755 tic Generation Through Large Language Models with Meta-Optimization. *arXiv preprint*  
 756 *arXiv:2505.20881*, 2025. URL <https://arxiv.org/abs/2505.20881>.

756 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao.  
 757     Reflexion: Language Agents with Verbal Reinforcement Learning. In A. Oh, T. Naumann,  
 758     A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural  
 759     Information Processing Systems*, volume 36, pp. 8634–8652. Curran Associates, Inc.,  
 760     2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1b44b878bb782e6954cd888628510e90-Paper-Conference.pdf).

762 Weiwei Sun, Shengyu Feng, Shanda Li, and Yiming Yang. Co-bench: Benchmarking language  
 763 model agents in algorithm search for combinatorial optimization. *ArXiv*, abs/2504.04310, 2025.  
 764 URL <https://arxiv.org/abs/2504.04310>.

766 Anja Šurina, Amin Mansouri, Amal Seddas, Maryna Viazovska, Emmanuel Abbe, and Caglar  
 767 Gulcehre. Algorithm discovery with LLMs: Evolutionary search meets reinforcement learning.  
 768 In *Scaling Self-Improving Foundation Models without Human Supervision*, 2025. URL  
 769 <https://openreview.net/forum?id=1kAwyBpo01>.

770 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,  
 771 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical  
 772 report. *arXiv preprint arXiv:2503.19786*, 2025.

774 Cong Dao Tran, Quan Nguyen-Tri, Huynh Thi Thanh Binh, and Hoang Thanh-Tung. Large language  
 775 models powered neural solvers for generalized vehicle routing problems. In *ICLR 2025 Workshop  
 776 on Towards Agentic AI for Science: Hypothesis Generation, Comprehension, Quantification, and  
 777 Validation*, 2025. URL <https://openreview.net/forum?id=EVqlVjvlt8>.

778 Eduardo Uchoa, Diego Pecin, Artur Pessoa, Marcus Poggi, Thibaut Vidal, and Anand Subramanian.  
 779     New benchmark instances for the capacitated vehicle routing problem. *European Journal of  
 780     Operational Research*, 257(3):845–858, 2017. ISSN 0377-2217. doi: <https://doi.org/10.1016/j.ejor.2016.08.012>. URL <https://www.sciencedirect.com/science/article/pii/S0377221716306270>.

783 Thibaut Vidal. Hybrid genetic search for the CVRP: Open-source implementation and SWAP\*  
 784     Neighborhood. *Computers & Operations Research*, 140:105643, 2022.

786 Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer Networks. In C. Cortes,  
 787 N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (eds.), *Advances in Neural  
 788     Information Processing Systems*, volume 28, pp. 2692–2700. Curran Associates, Inc.,  
 789 2015. URL [https://proceedings.neurips.cc/paper\\_files/paper/2015/file/29921001f2f04bd3baee84a12e98098f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2015/file/29921001f2f04bd3baee84a12e98098f-Paper.pdf).

791 Niels A. Wouda, Leon Lan, and Wouter Kool. PyVRP: a high-performance VRP solver package.  
 792     *INFORMS Journal on Computing*, 36(4):943–955, 2024. doi: 10.1287/ijoc.2023.0055. URL  
 793 <https://doi.org/10.1287/ijoc.2023.0055>.

795 Yaxin Wu, Wen Song, Zhiguang Cao, Jie Zhang, and Andrew Lim. Learning Improvement Heuristics  
 796     for Solving Routing Problems. *IEEE Transactions on Neural Networks and Learning Systems*,  
 797 2019.

798 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,  
 799     Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint  
 800     arXiv:2505.09388*, 2025a.

802 Xianliang Yang, Ling Zhang, Haolong Qian, Lei Song, and Jiang Bian. HeurAgenix: Lever-  
 803     aging LLMs for Solving Complex Combinatorial Optimization Challenges. *arXiv preprint  
 804     arXiv:2506.15196*, 2025b. URL <https://arxiv.org/abs/2506.15196>.

805 Haoran Ye, Jiarui Wang, Zhiguang Cao, Helan Liang, and Yong Li. DeepACO: Neural-  
 806     enhanced Ant Systems for Combinatorial Optimization. In A. Oh, T. Naumann,  
 807     A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural  
 808     Information Processing Systems*, volume 36, pp. 43706–43728. Curran Associates, Inc.,  
 809 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/883105b282fe15275991b411e6b200c5-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/883105b282fe15275991b411e6b200c5-Paper-Conference.pdf).

810 Haoran Ye, Jiarui Wang, Zhiguang Cao, Federico Berto, Chuanbo Hua, Haeyeon Kim, Jinkyoo Park,  
 811 and Guojie Song. Reevo: Large language models as hyper-heuristics with reflective evolution.  
 812 *Advances in neural information processing systems*, 37:43571–43608, 2024a.

813

814 Haoran Ye, Jiarui Wang, Helan Liang, Zhiguang Cao, Yong Li, and Fanzhang Li. GLOP: Learning  
 815 Global Partition and Local Construction for Solving Large-Scale Routing Problems in Real-Time.  
 816 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 20284–20292,  
 817 2024b. doi: 10.1609/aaai.v38i18.30009. URL <https://ojs.aaai.org/index.php/AAAI/article/view/30009>.

818

819 Huigen Ye, Hua Xu, An Yan, and Yaoyang Cheng. Large language model-driven large neighbor-  
 820 hood search for large-scale milp problems. In *Forty-second International Conference on Machine  
 821 Learning*, 2025.

822 Zhi Zheng, Zhuoliang Xie, Zhenkun Wang, and Bryan Hooi. Monte carlo tree search for compre-  
 823 hensive exploration in LLM-based automatic heuristic design. In *International Conference on  
 824 Machine Learning (ICML)*, 2024a.

825

826 Zhi Zheng, Changliang Zhou, Xialiang Tong, Mingxuan Yuan, and Zhenkun Wang. UDC: A Unified  
 827 Neural Divide-and-Conquer Framework for Large-Scale Combinatorial Optimization Problems.  
 828 In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Ad-  
 829 vances in Neural Information Processing Systems*, volume 37, pp. 6081–6125. Curran Associates,  
 830 Inc., 2024b. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/0b8e4c8468273ee3bafb288229c0acbc-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/0b8e4c8468273ee3bafb288229c0acbc-Paper-Conference.pdf).

831

832 Zhi Zheng, Zhuoliang Xie, Zhenkun Wang, and Bryan Hooi. Monte carlo tree search for com-  
 833 prehensive exploration in LLM-based automatic heuristic design. In *Forty-second International  
 834 Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=DolOdZzYHr>.

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 

## A PROMPTS

866 The exact prompts used by VRPAGENT-GA are presented below. We distinguish between *general*  
 867 *prompts*, which remain the same across all problems, and *problem-specific prompts*, which must be  
 868 tailored to each task. The two are combined by substituting variables in the general prompts (e.g.,  
 869 replacing `{problem_desc}` with the corresponding problem-specific description). For crossover and  
 870 ablation prompts, the provided template is further extended by inserting the implementations of the  
 871 associated individuals at the designated positions.

872 For brevity, we report only the problem-specific prompts for the CVRP [here](#). The prompts for all  
 873 problems can be found in our [online repository](#).

875 

### A.1 GENERAL PROMPTS

```
877 You are an operations research expert. Your task is to design new heuristics for an
878 existing **Large Neighborhood Search (LNS)** framework applied to the
879 {problem_name_long}. The framework iteratively improves a given initial solution
880 through the following steps:
881 1. **Customer Removal**: Select a subset of customers to remove using a specified
882 heuristic.
883 2. **Solution Perturbation**: Remove the selected customers from their tours. This
884 results in an infeasible solution where the removed customers are no longer served.
885 3. **Customer Ordering**: Order the removed customers using another heuristic.
886 4. **Greedy Reinsertion**: Reinsert the removed customers one by one into the tours,
887 following the order defined in step 3.

888 Your job is to implement **new heuristics for**:
889 - **Step 1**: Customer selection (select_by_llm_1)
890 - **Step 3**: Ordering of the removed customers (sort_by_llm_1)

891 All other components of the LNS framework are fixed and **cannot be modified**.

892 # Routing Problem Description
893 {problem_desc}

894 # Other implementation notes and requirements:
895 - The framework is implemented in **C++**.
896 - The LNS targets **large instances** (e.g., more than 500 customers).
897 - Only a small number of customers should be removed in each iteration.
898 - The selected customers do **not need to form a single compact cluster**, but **each**
899 selected customer should be close to at least one or a few other selected customers**.
900 This encourages meaningful changes during greedy reinsertion.
901 - The heuristic must incorporate **stochastic behavior** to ensure sufficient diversity
902 over **millions of iterations**.
903 - The search is limited by runtime, meaning that the two new heuristics should be very
904 fast.

905 # Code style
906 - IMPORTANT: DO NOT ADD ***ANY*** COMMENTS unless asked
```

907 **Prompt 1: Global system context.**

```
908 [TASK]
909 Write high-quality heuristics for select_by_llm_1 and sort_by_llm_1 in the LNS
910 framework. Write the full code file in a ```cpp``` code block.

911 # Example implementation
912 {seed_code}

913 # Library context
914 You are also provided with some selected header function information with comments that
915 could be useful:
916 {LNS_headers}
```

917 **Prompt 2: Initial operator generation.**

```

918
919 [Better Code]
920 {code_parent_1}
921
922 [Worse Code]
923 {code_parent_2}
924
925 [Task]
926 Write new high-quality heuristics for `select_by_llm_1` and `sort_by_llm_1` in the LNS
927 framework. Your implementation
928 should be a crossover of the two implementations above, taking most ideas from the
929 better code (80%) and only some ideas from the worse code (20%).
930 Ensure that the new code maintains a comparable overall complexity and length to the
931 two implementations above.
932 Output code only and enclose your code with C++ code block: ```cpp ... ```.
933 Do not
934 comment your code.
935
936
937
938

```

Prompt 3: Crossover prompt with 80% bias.

```

939
940
941
942 [Code]
943 {code}
944
945 [Task]
946 To simplify the heuristics implemented in `select_by_llm_1` and `sort_by_llm_1` we want
947 to conduct an ablation study.
948 Choose a random mechanic/component from the code that you think might not be important
949 and remove any trace of it from the code. We will
950 then run your code to evaluate the impact of the removed component. Output code only
951 and enclose your code with C++ code block: ```cpp ... ```.
952
953
954
955
956
957

```

Prompt 4: Ablation (Mutation).

```

958
959
960
961 [Code]
962 {code}
963
964 [Task]
965 The goal is improve the heuristics implemented in `select_by_llm_1` and `sort_by_llm_1`.
966 Add a new mechanic/component to the code above. Be innovative. We will
967 then run your code to evaluate the impact of the new component. Output code only and
968 enclose your code with C++ code block: ```cpp ... ```.
969
970
971

```

Prompt 5: Extend (Mutation).

```

972
973
974
975 [Code]
976 {code}
977
978 [Task]
979 The goal is to find new parameter settings for heuristics implemented in
980 `select_by_llm_1` and `sort_by_llm_1`.
981 Modify the parameters of the code above to improve the effectiveness of the heuristic.
982 If there are magic numbers in the code, replace them with constants that are set at the
983 beginning of each function.
984 Do not make any other changes to the code.
985 Output code only and enclose your code with C++ code block: ```cpp ... ```.
986
987
988
989

```

Prompt 6: Adjust-Parameters (Mutation).

```

990
991
992
993 [Code]
994 {code}
995
996 [Task]
997 The goal is improve the runtime of the heuristics implemented in `select_by_llm_1` and
998 `sort_by_llm_1`.
999 Modify the code so that the runtime is reduced. It is ok to slightly change the logic
1000 of the heuristic to achieve this.
1001 Output code only and enclose your code with C++ code block: ```cpp ... ```.
1002
1003
1004
1005
1006
1007

```

Prompt 7: Refactor (Mutation).

```

1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021

```

972 A.2 PROBLEM-SPECIFIC PROMPTS

973

974 A.2.1 CVRP

975

976 The Capacitated Vehicle Routing Problem (CVRP) involves determining a set of delivery  
 977 routes from a depot to a group of customers, where each customer has a specific demand  
 978 and each vehicle has a fixed capacity. The objective is to design routes that minimize  
 979 the total distance traveled, while ensuring that:  
 980 Each route starts and ends at the depot.  
 981 Each customer is visited exactly once by a single vehicle.  
 982 The total demand on any route does not exceed the vehicle capacity.

983 There is no limit on the number of vehicles that can be used.

984

985 Prompt 8: Problem description CVRP).

986

987

```
988 From `Instance.h`:
989
990 ```cpp
991 struct Instance {
992     int numNodes; // Total number of nodes including depot
993     int numCustomers; // Total number of customers (excluding depot)
994     int vehicleCapacity; // Capacity of the vehicle (identical for all vehicles)
995     std::vector<int> demand; // Demand of each node (with the depot at index 0 having a
996     // demand of 0)
997     std::vector<std::vector<float>> distanceMatrix; // Distance matrix between nodes
998     std::vector<std::vector<float>> nodePositions; // Node positions in 2D space
999     std::vector<std::vector<int>> adj; // Adjacency list for each node, sorted by
1000     // distance
1001 }
1002
1003
1004 From `Solution.h`:
1005
1006 ```cpp
1007 struct Solution {
1008     const Instance& instance; // Reference to the instance to avoid copying
1009     float totalCosts; // Total cost of the solution
1010     std::vector<Tour> tours; // List of tours in the solution
1011     std::vector<int> customerToTourMap; // Map from each customer to its tour index. This
1012     // can be used to
1013     // quickly find which tour a customer belongs to, e.g. solution.tours[solution.
1014     // customerToTourMap[c]] returns the tour of customer c.
1015 }
1016
1017
1018 From `Tour.h`:
1019
1020 ```cpp
1021 struct Tour {
1022     std::vector<int> customers; // Customers in the tour, excluding depot
1023     int demand = 0; // Total demand of the tour
1024     float costs = 0; // Total cost of the tour including distance to and from the depot
1025 }
1026
1027
1028 From `Utils.h`:
1029
1030 ```cpp
1031 int getRandomNumber(int min, int max);
1032 float getRandomFraction(float min = 0.0, float max = 1.0);
1033 float getRandomFractionFast(); // Function to generate a random fraction (float) in the
1034     // range [0, 1] using a fast method
1035 std::vector<int> argsort(const std::vector<float>& values); // Function to perform
1036     // argsort on a vector of float values
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
20100
20101
20102
20103
20104
20105
20106
20107
20108
20109
20110
20111
20112
20113
20114
20115
20116
20117
20118
20119
20120
20121
20122
20123
20124
20125
20126
20127
20128
20129
20130
20131
20132
20133
20134
20135
20136
20137
20138
20139
20140
20141
20142
20143
20144
20145
20146
20147
20148
20149
20150
20151
20152
20153
20154
20155
20156
20157
20158
20159
201510
201511
201512
201513
201514
201515
201516
201517
201518
201519
201520
201521
201522
201523
201524
201525
201526
201527
201528
201529
201530
201531
201532
201533
201534
201535
201536
201537
201538
201539
201540
201541
201542
201543
201544
201545
201546
201547
201548
201549
201550
201551
201552
201553
201554
201555
201556
201557
201558
201559
201560
201561
201562
201563
201564
201565
201566
201567
201568
201569
201570
201571
201572
201573
201574
201575
201576
201577
201578
201579
201580
201581
201582
201583
201584
201585
201586
201587
201588
201589
201590
201591
201592
201593
201594
201595
201596
201597
201598
201599
2015100
2015101
2015102
2015103
2015104
2015105
2015106
2015107
2015108
2015109
2015110
2015111
2015112
2015113
2015114
2015115
2015116
2015117
2015118
2015119
2015120
2015121
2015122
2015123
2015124
2015125
2015126
2015127
2015128
2015129
2015130
2015131
2015132
2015133
2015134
2015135
2015136
2015137
2015138
2015139
2015140
2015141
2015142
2015143
2015144
2015145
2015146
2015147
2015148
2015149
2015150
2015151
2015152
2015153
2015154
2015155
2015156
2015157
2015158
2015159
20151510
20151511
20151512
20151513
20151514
20151515
20151516
20151517
20151518
20151519
20151520
20151521
20151522
20151523
20151524
20151525
20151526
20151527
20151528
20151529
20151530
20151531
20151532
20151533
20151534
20151535
20151536
20151537
20151538
20151539
20151540
20151541
20151542
20151543
20151544
20151545
20151546
20151547
20151548
20151549
20151550
20151551
20151552
20151553
20151554
20151555
20151556
20151557
20151558
20151559
20151560
20151561
20151562
20151563
20151564
20151565
20151566
20151567
20151568
20151569
20151570
20151571
20151572
20151573
20151574
20151575
20151576
20151577
20151578
20151579
20151580
20151581
20151582
20151583
20151584
20151585
20151586
20151587
20151588
20151589
20151590
20151591
20151592
20151593
20151594
20151595
20151596
20151597
20151598
20151599
201515100
201515101
201515102
201515103
201515104
201515105
201515106
201515107
201515108
201515109
201515110
201515111
201515112
201515113
201515114
201515115
201515116
201515117
201515118
201515119
201515120
201515121
201515122
201515123
201515124
201515125
201515126
201515127
201515128
201515129
201515130
201515131
201515132
201515133
201515134
201515135
201515136
201515137
201515138
201515139
201515140
201515141
201515142
201515143
201515144
201515145
201515146
201515147
201515148
201515149
201515150
201515151
201515152
201515153
201515154
201515155
201515156
201515157
201515158
201515159
201515160
201515161
201515162
201515163
201515164
201515165
201515166
201515167
201515168
201515169
201515170
201515171
201515172
201515173
201515174
201515175
201515176
201515177
201515178
201515179
201515180
201515181
201515182
201515183
201515184
201515185
201515186
201515187
201515188
201515189
201515190
201515191
201515192
201515193
201515194
201515195
201515196
201515197
201515198
201515199
201515200
201515201
201515202
201515203
201515204
201515205
201515206
201515207
201515208
201515209
201515210
201515211
201515212
201515213
201515214
201515215
201515216
201515217
201515218
201515219
201515220
201515221
201515222
201515223
201515224
201515225
201515226
201515227
201515228
201515229
201515230
201515231
201515232
201515233
201515234
201515235
201515236
201515237
201515238
201515239
201515240
201515241
201515242
201515243
201515244
201515245
201515246
201515247
201515248
201515249
201515250
201515251
201515252
201515253
201515254
201515255
201515256
201515257
201515258
201515259
201515260
201515261
201515262
201515263
201515264
201515265
201515266
201515267
201515268
201515269
201515270
201515271
201515272
201515273
201515274
201515275
201515276
201515277
201515278
201515279
201515280
201515281
201515282
201515283
201515284
201515285
201515286
201515287
201515288
201515289
201515290
201515291
201515292
201515293
201515294
201515295
201515296
201515297
201515298
201515299
201515300
201515301
201515302
201515303
201515304
201515305
201515306
201515307
201515308
201515309
201515310
201515311
201515312
201515313
201515314
201515315
201515316
201515317
201515318
201515319
201515320
201515321
201515322
201515323
201515324
201515325
201515326
201515327
201515328
201515329
201515330
201515331
201515332
201515333
201515334
201515335
201515336
201515337
201515338
201515339
201515340
201515341
201515342
201515343
201515344
201515345
201515346
201515347
201515348
201515349
201515350
201515351
201515352
20
```

```

1026
1027 #include "AgentDesigned.h"
1028 #include <random>
1029 #include <unordered_set>
1030 #include "Utils.h"
1031
1032 // Customer selection
1033 std::vector<int> select_by_llm_1(const Solution& sol) {
1034     // random selection of customers
1035     std::unordered_set<int> selectedCustomers;
1036
1037     int numCustomersToRemove = getRandomNumber(10, 20);
1038
1039     while (selectedCustomers.size() < numCustomersToRemove) {
1040         int randomCustomer = getRandomNumber(1, sol.instance.numCustomers);
1041         selectedCustomers.insert(randomCustomer);
1042     }
1043
1044     return std::vector<int>(selectedCustomers.begin(), selectedCustomers.end());
1045 }
1046
1047
1048 // Function selecting the order in which to remove the customers
1049 void sort_by_llm_1(std::vector<int>& customers, const Instance& instance) {
1050     // Placeholder for LLM-based sorting logic
1051     // This function should implement the logic to sort customers based on a learned
1052     // model
1053     // For now, we will just sort randomly as a placeholder
1054     // sort_randomly(customers, instance);
1055     static thread_local std::mt19937 gen(std::random_device{}());
1056     std::shuffle(customers.begin(), customers.end(), gen);
1057 }

```

Prompt 10: Seed heuristic ({seed\_code}).

## 1050 B ADDITIONAL DETAILS

### 1053 B.1 COSTS PER RUN

1055 Table 2: Comparison token usage and cost estimates across models per run (as of Sep. 2025) with inference  
1056 providers sources.

1058 Model	1059 Open Source	1060 Token Usage		1061 Costs (\$)		1062 Total Costs (\$)	1063 Source
		1064 Input	1065 Output	1066 Input	1067 Output		
Gemini 2.0 Flash	✗	2.5M	1.2M	0.10	0.40	0.73	Vertex AI
Gemini 2.5 Flash	✗	4.1M	7.1M	0.30	2.50	18.98	Vertex AI
gpt-oss (120B)	✓	4.0M	2.5M	0.09	0.36	1.26	Clarifai
gemma 3 (27B)	✓	2.5M	1.2M	0.09	0.16	0.42	DeepInfra
Qwen3 (30B)	✓	1.5M	4.4M	0.08	0.29	1.40	Clarifai
Llama 3.3 (70B)	✓	2.3M	1.0M	0.08	0.29	0.47	Clarifai

### 1066 B.2 USE OF LARGE LANGUAGE MODELS

1069 LLMs played an active role in this work. Beyond serving as general-purpose writing assistants for  
1070 improving clarity, style, and grammar and as coding assistants, LLMs were employed as heuristic  
1071 discovery tools during the optimization phase of our study. Importantly, the core research contribu-  
1072 tions, including the design of the framework, theoretical development, and validation of results,  
1073 were conceived, implemented, and verified exclusively by the authors. All outputs from LLMs  
1074 were critically assessed, refined, and integrated to ensure correctness and adherence to academic  
1075 standards.

## 1076 C EVALUATION ON CVRPLIB INSTANCES

1078 We also evaluate VRPAGENT on the X instance set (Uchoa et al., 2017) from CVRPLib, which  
1079 consists of 100 CVRP instances with sizes ranging from 100 to 1000 customers. This dataset is ex-

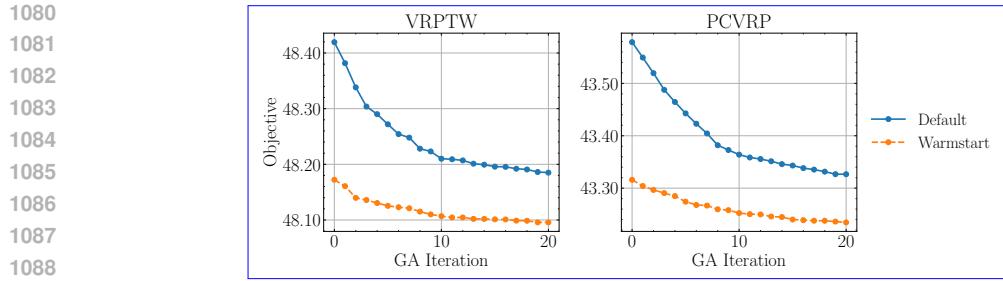


Figure 6: Objective values during heuristic discovery for VRPTW and PCVRP with and without warm-starts from the best CVRP heuristic.

licitly designed to encompass a wide variety of challenging problem structures, including instances with highly clustered customers and heterogeneous demands, providing a rigorous test of our agent's robustness in scenarios that closely resemble real-world complexity. To generate a heuristic for these instances, we train **VRPAGENT** on a separate set of 128 instances exhibiting similar characteristics. During the heuristic discovery process, the runtime for each instance is constrained to  $0.025n$  seconds, where  $n$  denotes the number of nodes.

The resulting heuristic is subsequently evaluated on the full  $X$  instance set with an increased runtime limit of  $1.0n$  seconds per instance. We compare the performance of the heuristic generated by VR-PAGENT against the state-of-the-art handcrafted method SISR<sub>s</sub> under the same runtime constraints. For each instance, we perform 10 independent runs and report the average results.

Table 3 summarizes the performance of both approaches on all individual instances. VRPAGENT achieves superior solutions compared to SISR on 67 instances, whereas SISR outperforms VRPAGENT on 24 instances, demonstrating the overall advantage of our approach in this setting. Notably, for larger instances with more than 500 customers, the heuristic produced by VRPAGENT consistently yields better solutions, with only a few exceptions, highlighting its effectiveness in scaling to more complex and larger-scale problems.

## D SUCCESS RATE

During the main training runs, we track the proportion of generated operators that function correctly, i.e., those that both compile successfully and produce solutions for all validation instances. We refer to this proportion as the success rate, and it provides insight into the difficulty of the generation tasks posed to the LLMs.

Fig. 7 reports the success rate over the course of the discovery runs. Across all problem settings, the success rate remains above 85%, indicating that the generation tasks are generally manageable for the LLM used. For the VRPTW, we also observe a gradual increase in success as the discovery process progresses. For the PCVRP, the success rate is consistently around 95%, likely because not all customers must be visited, making feasible solution generation comparatively easy (e.g., even a solution with no tours is technically feasible).

Note that the figure reflects the success rate only for operators generated within the main loop of VRPAGENT-GA (i.e., offspring and mutation operators). During the initial population generation, the success rates are lower: 73% for the CVRP, 68% for the VRPTW, and 90% for the PCVRP.

## E WARM-STARTING HEURISTIC DISCOVERY

We investigate whether operators discovered for one problem can be effectively adapted to other problems by conducting experiments in which we warm-start the discovery process for the VRPTW and PCVRP using the best heuristic found for the CVRP. For the warm-start, we provide the CVRP implementation as an example to the model and task it with adapting the heuristic to the new problem (i.e., VRPTW or PCVRP). Each configuration is run 10 times, and the results are averaged.

1134 Table 3: Comparisons on the X set of CVRPLib: best-known solutions (BKS), SISR<sub>s</sub> and VRPAgent costs, and  
 1135 their percentage gaps. Results show the average performance across 10 runs.

1136

1137  
1138  
1139

Instance	BKS	SISR <sub>s</sub>		VRPAgent		Instance	BKS	SISR <sub>s</sub>		VRPAgent	
		Obj.	Gap%	Obj.	Gap%			Obj.	Gap%	Obj.	Gap%
X-n101-k25	27591	27591	0.000	27591	0.000	X-n336-k84	139111	139515	0.290	<b>139377</b>	<b>0.190</b>
X-n106-k14	26362	26383	0.080	<b>26366</b>	<b>0.020</b>	X-n344-k43	42050	<b>42110</b>	<b>0.140</b>	42146	0.230
X-n110-k13	14971	14971	0.000	14971	0.000	X-n351-k40	25896	<b>25982</b>	<b>0.340</b>	25992	0.370
X-n115-k10	12747	12747	0.000	12747	0.000	X-n359-k29	51505	51575	0.140	<b>51570</b>	<b>0.130</b>
X-n120-k6	13332	<b>13332</b>	<b>0.000</b>	13332	0.010	X-n367-k17	22814	22866	0.230	<b>22829</b>	<b>0.070</b>
X-n125-k30	55539	55606	0.120	<b>55545</b>	<b>0.010</b>	X-n376-k94	147713	147814	0.070	<b>147737</b>	<b>0.010</b>
X-n129-k18	28940	<b>28952</b>	<b>0.040</b>	28955	0.050	X-n384-k52	65928	66128	0.300	<b>66066</b>	<b>0.210</b>
X-n134-k13	10916	10943	0.250	<b>10942</b>	<b>0.240</b>	X-n393-k38	38260	38393	0.350	<b>38355</b>	<b>0.250</b>
X-n139-k10	13590	13599	0.070	<b>13596</b>	<b>0.050</b>	X-n401-k29	66154	66258	0.160	<b>66251</b>	<b>0.150</b>
X-n143-k7	15700	15717	0.110	15716	0.110	X-n411-k19	19712	19785	0.380	<b>19764</b>	<b>0.270</b>
X-n148-k46	43448	<b>43478</b>	<b>0.070</b>	43492	0.100	X-n420-k130	107798	107891	0.090	107897	0.090
X-n153-k22	21220	<b>21227</b>	<b>0.030</b>	21316	0.460	X-n429-k61	65449	65609	0.240	<b>65590</b>	<b>0.220</b>
X-n157-k13	16876	16883	0.040	<b>16876</b>	<b>0.000</b>	X-n439-k37	36391	36485	0.260	<b>36461</b>	<b>0.190</b>
X-n162-k11	14138	<b>14153</b>	<b>0.110</b>	14153	0.120	X-n449-k29	55233	<b>55420</b>	<b>0.340</b>	55443	0.380
X-n167-k10	20557	<b>20557</b>	<b>0.000</b>	20590	0.160	X-n459-k26	24139	24247	0.450	<b>24226</b>	<b>0.360</b>
X-n172-k51	45607	45630	0.050	<b>45607</b>	<b>0.000</b>	X-n469-k138	221824	222339	0.230	<b>222249</b>	<b>0.190</b>
X-n176-k26	47812	<b>47841</b>	<b>0.060</b>	47889	0.160	X-n480-k70	89449	<b>89568</b>	<b>0.130</b>	89603	0.170
X-n181-k23	25569	25580	0.040	<b>25576</b>	<b>0.030</b>	X-n491-k59	66483	66708	0.340	<b>66608</b>	<b>0.190</b>
X-n186-k15	24145	<b>24163</b>	<b>0.080</b>	24180	0.140	X-n502-k39	69226	69275	0.070	69272	0.070
X-n190-k8	16980	16996	0.090	<b>16990</b>	<b>0.060</b>	X-n513-k21	24201	<b>24295</b>	<b>0.390</b>	24298	0.400
X-n195-k51	44225	44308	0.190	<b>44277</b>	<b>0.120</b>	X-n524-k153	154593	154905	0.200	<b>154786</b>	<b>0.120</b>
X-n200-k36	58578	58641	0.110	<b>58635</b>	<b>0.100</b>	X-n536-k96	94846	95209	0.380	<b>95153</b>	<b>0.330</b>
X-n204-k19	19565	<b>19631</b>	<b>0.340</b>	19675	0.570	X-n548-k50	86700	86806	0.120	<b>86783</b>	<b>0.100</b>
X-n209-k16	30656	<b>30670</b>	<b>0.050</b>	30679	0.080	X-n561-k42	42717	42878	0.380	<b>42820</b>	<b>0.240</b>
X-n214-k11	10856	10904	0.450	<b>10897</b>	<b>0.380</b>	X-n573-k30	50673	50826	0.300	<b>50773</b>	<b>0.200</b>
X-n219-k73	117595	117624	0.030	<b>117612</b>	<b>0.020</b>	X-n586-k159	190316	190699	0.200	<b>190598</b>	<b>0.150</b>
X-n223-k34	40437	40551	0.280	<b>40491</b>	<b>0.140</b>	X-n599-k92	108451	108706	0.240	<b>108685</b>	<b>0.220</b>
X-n228-k23	25742	<b>25794</b>	<b>0.200</b>	25797	0.220	X-n613-k62	59535	59765	0.390	<b>59679</b>	<b>0.240</b>
X-n233-k16	19230	19282	0.270	19281	0.270	X-n627-k43	62164	62335	0.280	<b>62320</b>	<b>0.250</b>
X-n237-k14	27042	<b>27103</b>	<b>0.230</b>	27134	0.340	X-n641-k35	63682	63877	0.310	<b>63842</b>	<b>0.250</b>
X-n242-k48	82751	82902	0.180	<b>82879</b>	<b>0.150</b>	X-n655-k131	106780	106884	0.100	<b>106825</b>	<b>0.040</b>
X-n247-k50	37274	37389	0.310	<b>37342</b>	<b>0.180</b>	X-n670-k130	146332	<b>146992</b>	<b>0.450</b>	147491	0.790
X-n251-k28	38684	38808	0.320	<b>38737</b>	<b>0.140</b>	X-n685-k75	68205	68401	0.290	<b>68340</b>	<b>0.200</b>
X-n256-k16	18839	18901	0.330	<b>18885</b>	<b>0.250</b>	X-n701-k44	81923	82131	0.260	<b>82080</b>	<b>0.190</b>
X-n261-k13	26558	26657	0.380	<b>26650</b>	<b>0.350</b>	X-n716-k35	43373	<b>43491</b>	<b>0.270</b>	43493	0.280
X-n266-k58	75478	75650	0.230	75651	0.230	X-n733-k159	136187	136462	0.200	<b>136401</b>	<b>0.160</b>
X-n270-k35	35291	<b>35347</b>	<b>0.160</b>	35362	0.200	X-n749-k98	77269	77584	0.410	<b>77493</b>	<b>0.290</b>
X-n275-k28	21245	<b>21275</b>	<b>0.150</b>	21285	0.190	X-n766-k71	114417	114761	0.300	<b>114688</b>	<b>0.240</b>
X-n280-k17	33503	33634	0.390	<b>33559</b>	<b>0.170</b>	X-n783-k48	72386	72645	0.360	<b>72568</b>	<b>0.250</b>
X-n284-k15	20215	20289	0.370	<b>20272</b>	<b>0.280</b>	X-n801-k40	73305	73446	0.190	<b>73438</b>	<b>0.180</b>
X-n289-k60	95151	95452	0.320	<b>95421</b>	<b>0.280</b>	X-n819-k171	158121	158513	0.250	<b>158408</b>	<b>0.180</b>
X-n294-k50	47161	47280	0.250	<b>47268</b>	<b>0.230</b>	X-n837-k142	193737	194047	0.160	<b>194018</b>	<b>0.150</b>
X-n298-k31	34231	34274	0.120	<b>34264</b>	<b>0.100</b>	X-n856-k95	88965	89152	0.210	<b>89119</b>	<b>0.180</b>
X-n303-k21	21736	<b>21788</b>	<b>0.240</b>	21819	0.380	X-n876-k59	99299	99544	0.250	<b>99494</b>	<b>0.200</b>
X-n308-k13	25859	26192	1.290	<b>25934</b>	<b>0.290</b>	X-n895-k37	53860	54138	0.520	<b>54104</b>	<b>0.450</b>
X-n313-k71	94043	94235	0.200	<b>94189</b>	<b>0.160</b>	X-n916-k207	329179	<b>329556</b>	<b>0.110</b>	329566	0.120
X-n317-k53	78355	78397	0.060	<b>78364</b>	<b>0.010</b>	X-n936-k151	132715	133387	0.510	<b>133201</b>	<b>0.370</b>
X-n322-k28	29834	29933	0.330	<b>29905</b>	<b>0.240</b>	X-n957-k87	85465	85640	0.200	85633	0.200
X-n327-k20	27532	27671	0.510	<b>27636</b>	<b>0.380</b>	X-n979-k58	118976	119147	0.140	<b>119118</b>	<b>0.120</b>
X-n331-k15	31102	<b>31141</b>	<b>0.130</b>	31164	0.200	X-n1001-k43	72355	72563	0.290	<b>72536</b>	<b>0.250</b>

1177

1178  
1179  
1180  
1181

1182 Fig. 6 shows the objective function values over the course of the discovery process for the VRPTW  
 1183 and PCVRP, comparing warm-start and cold-start runs. Warm-starting significantly improves per-  
 1184 formance, yielding better results from the initial population. All warm-start runs begin from the best  
 1185 CVRP implementation found across 10 discovery runs, which required a substantial amount of com-  
 1186 pute. The combination of high-quality initial heuristics and diversity introduced through the start  
 1187 population generation allows the method to discover implementations that outperform the cold-start  
 1188 runs on average.

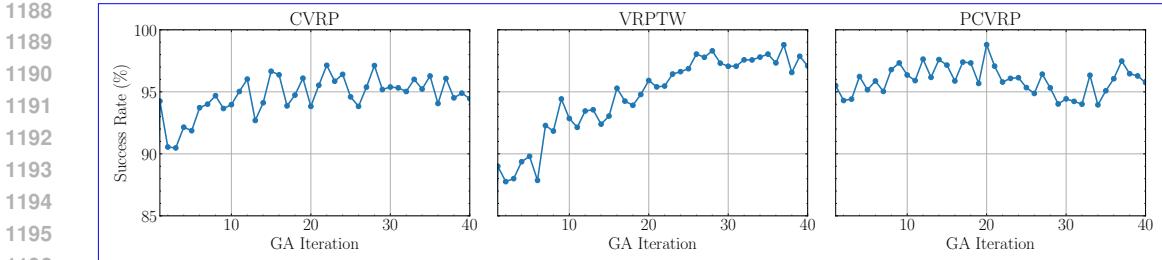


Figure 7: Percentage of operators that compile successfully and produce valid solutions (success rate) throughout the main discovery runs.

## F ANALYSIS OF DISCOVERED HEURISTICS

### F.1 HIGH-LEVEL DESCRIPTION

#### F.1.1 CVRP

**Customer Removal** The removal heuristic, iteratively constructs a set of customers to remove based on a target size drawn from a uniform distribution (e.g., between 15 and 28). The process initializes by either cutting a random contiguous tour segment (with probability 0.2) or selecting a single random seed. The set is subsequently expanded by choosing a pivot from the currently selected customers and applying one of three weighted strategies: adding a spatial neighbor from the pivot’s k-nearest graph (60%), adding a short contiguous segment from the pivot’s current tour (25%), or adding the pivot’s immediate tour predecessor or successor (15%). A random fallback ensures the target removal count is always met.

**Customer Ordering** The customer ordering operator stochastically selects one of eight sorting rules and a random sorting direction (ascending or descending). The scoring criteria include geometric metrics (distance to depot, distance to the removed set’s centroid, polar angle), problem-specific attributes (demand, combined demand-distance), and topological features (connectivity among removed nodes). Alternatively, the order may be determined by a stochastic Nearest-Neighbor chain or a simple random shuffle. To enhance diversity, minor additive noise is applied to the calculated scores during the sorting process.

#### F.1.2 VRPTW

**Customer Removal** The removal operator, constructs a cluster of customers to remove by growing a set from an initial random seed up to a target size drawn uniformly between 10 and 15. The expansion logic prioritizes connectivity: at each step, the algorithm attempts to select a tour neighbor (immediate predecessor or successor in the current route) with a probability of 0.45, or a geographical neighbor from the instance’s adjacency list with a probability of 0.35, where the latter is sampled with a power-law bias to favor closer nodes. If these connected expansion attempts fail, or with a 15% probability once the set is sufficiently large, the algorithm performs a “random jump” or falls back to a uniform random selection to ensure the target removal count is met.

**Customer Ordering** The customer ordering operator uniformly selects one of ten available scoring strategies. These strategies range from sorting by single raw attributes (e.g., time window width, start time, or negative demand) to complex composite scores. Specifically, one advanced strategy aggregates five normalized attributes (including service time and distance to depot) using randomly generated weights and polarities, while another assesses demand density relative to time window tightness. To maximize diversity, the final ordering process includes a microscopic noise factor for tie-breaking, a 25% probability of reversing the sorted list, and a final pass that applies a small number of random swaps.

1242 F.1.3 PCVRP  
1243

1244 **Customer Removal** The removal operator, constructs a cluster of customers to remove by growing  
1245 a set from an initial random seed up to a target size drawn uniformly between 10 and 20. The seed-  
1246 ing process prioritizes unassigned customers (25%) and those currently belonging to tours (60%)  
1247 before falling back to a uniform selection. The expansion logic is iterative: at each step, a subset  
1248 of "source" customers is chosen from the current selection, and candidates are gathered from both  
1249 their topological neighbors (via the adjacency list) and random samples of their current tour-mates.  
1250 A single candidate is uniformly selected from this pool, with a random safeguard applied if the  
1251 candidate pool is empty.

1252 **Customer Ordering** The customer ordering operator executes a pure random shuffle with proba-  
1253 bility 0.1; otherwise, it employs a score-based sort relative to a pivot node. The pivot is selected from  
1254 the removed set (80%) or set to the depot (20%). The scoring function calculates a weighted sum  
1255 involving the distance to the pivot, the distance to the depot, the customer's demand, and their prize  
1256 (if applicable). To enforce diversity, the specific weights for these components are modulated by  
1257 three distinct parameter regimes—varying the emphasis between prize collection and distance mini-  
1258 mization—and the final list is sorted in either ascending or descending order with equal probability.

1260 F.2 EXPERT EVALUATION  
1261

1262 LLM generated code raises many questions about its quality and maintainability. A further question  
1263 is how the code works and how it is able to achieve state-of-the-art performance. While we are  
1264 unable to fully answer these questions, we try to provide some initial insights into the quality of the  
1265 best heuristic generated for each problem. To do this, we have three co-authors of the paper with  
1266 many years of experience writing heuristics by hand analyze the code according to several criteria.  
1267 We acknowledge that this is not a scientific study and is not intended to draw generalizations about  
1268 the ability of LLMs to code heuristics for optimization problems. Rather, our goal is to give some  
1269 indications as to how the code generated compares to code written by humans and what kind of ideas  
1270 are present. We note that the code is generated without comments to avoid the LLM influencing the  
1271 analysis, however we note that variable names are present that do give some contextual information  
1272 about what the code does.

1273 The three heuristics experts have XX (expert 1), YY (expert 2) and ZZ (expert 3) years of experience  
1274 writing OR heuristics<sup>1</sup>. All experts have experience with routing problems in addition to other types  
1275 of OR problems. Each expert provides an evaluation of the generated code of the best performing  
1276 heuristic for each of the three problems examined in this work. The experts describe a consensus  
1277 description of how the heuristic works then write independent discussions of each heuristic. The  
1278 individual rating criteria are as follows:

1. Readability (noting that the assessments are not general statements about LLM code)
2. Coherence and soundness
3. Maintainability
4. Interpretability, i.e., do we know why this code works well?
5. Are there any new ideas in the heuristic?

1286 F.2.1 CVRP  
1287

1288 The removal and sorting mechanisms are best described as ensembles of heuristics in which the  
1289 heuristic applied at any given iteration is chosen at random according to a probability distribution  
1290 determined through the static parameters of the approach. For the selection of customers for re-  
1291 moval, the heuristics of the ensemble show a similarity to the SISRs heuristic. In the first, adjacent  
1292 customers are selected for removal and in the second, random segments of tours are chosen. Since  
1293

1294 <sup>1</sup>To avoid potentially violating the double blind submission policy, we do not indicate the years of experience  
1295 of the experts, as they are all coauthors of the work. These will be provided in the accepted version of this work,  
and this message will be removed.

1296 these segments can overlap, we also have a SISRs-like idea. The third heuristic, as best as we can  
 1297 determine, tries to expand a tour segment. For sorting, the heuristic first decides whether to sort de-  
 1298 scending or ascending according to one of seven different heuristics. We omit a detailed description  
 1299 of all the heuristics, but note that these include generally known ideas for sorting customers in a  
 1300 CVRP, e.g. using criteria such as the distance to the depot, the demand of the customers, weighted  
 1301 combinations of distance and demand, and greedy nearest-neighbor sequencing.

1302  
 1303 **Expert 1** The code is surprisingly well-written and is split into different functions in a logical  
 1304 fashion. The heuristics are coherent and reasonable. The ensemble contains many components and  
 1305 an ablation analysis would be necessary to determine how much each component actually contributes  
 1306 to the solution quality. The code looks relatively easy to maintain and follows good design principles  
 1307 for C++ code. There are no special constructs or libraries used, the memory management is very  
 1308 simple. While the heuristic does not present anything radically new, the scoring mechanism for  
 1309 removed neighbor connectivity seems looks novel. The removal heuristic is a SISRs variant, but not  
 1310 exactly the same as the original.

1311  
 1312 **Expert 2** The code is easy to understand and well-written. It contains both very meaningful high-  
 1313 level comments that make it easy to quickly assess the big picture of main functions, and low-level  
 1314 comments that facilitate understanding the details. In addition, the code is decomposed in a useful  
 1315 way into smaller functions with good names. As far as I see, this is good and standard C++ code that  
 1316 should be very maintainable and does not contain any weird parts. The created selection heuristics  
 1317 can be characterized as an ensemble of various known (and reasonably different) approaches that are  
 1318 applied in a random fashion. The sorting function randomly chooses between eight scoring schemes  
 1319 that involve use meaningful criteria; I assume that some combinations of criteria used for scoring  
 1320 are novel, although the criteria themselves are mostly not.

1321  
 1322 **Expert 3** The code is generally well-structured, with its components split into separate functions.  
 1323 This organization makes it easier for users to follow the logic without becoming overwhelmed by  
 1324 a single large code block. Additionally, the descriptive variable names and comments help readers  
 1325 quickly recognize the purpose of each section and provide useful guidance when deeper understand-  
 1326 ing is needed. Conceptually, the code combines several known ideas from the literature in a struc-  
 1327 tured ensemble; while none of the components are new by themselves, the way they are combined  
 1328 is somewhat novel. However, because it is not immediately clear which components are essential in  
 1329 practice, further analysis would be required to determine whether the code could be simplified.

### 1330 F.2.2 VRPTW

1331 The selection and sorting heuristics form portfolio approaches, that is, ensembles of different heuris-  
 1332 tics that are probabilistically selected and combined. The selection heuristic removes a randomly  
 1333 chosen number of customers (10-15). It starts with a randomly picked customer and then uses a  
 1334 main loop in which one additional customer is added per iteration, using spatial proximity and tour  
 1335 neighborhood to a selected reference customer as well as pure randomness to guide the choice. The  
 1336 sorting code involves ten different strategies, of which one is randomly chosen per call of the heuris-  
 1337 tic. Five strategies rely on simple smetrics such as demand or time window start time, four are based  
 1338 on combined metrics and one is a purely random shuffling of the customers. One of the combined  
 1339 metrics is highly complex and heavily parameterized.

1340  
 1341 **Expert 1** This heuristic is again very easy to read and the code is nicely structured. The functions  
 1342 are decently easy to understand and the code can be considered interpretable. While the sorting  
 1343 criteria are mostly standard and there are no big surprises, some combinations of metrics (e.g., the  
 1344 normalized combination in case 7) are potentially novel.

1345  
 1346 **Expert 2** Like in the CVRP case, the code is very well explained in both high and low-level  
 1347 comments and structured into functions in a meaningful way, avoiding deeply nested statements  
 1348 and making the code easy to follow and maintain. The code is reasonable and concise modern  
 1349 C++. The selection heuristic randomly applies a set of reasonable heuristics, and the sorting code  
 again is based on randomly choosing between a set of useful and known sorting criteria, some of  
 which are weighted combinations of multiple criteria. Although consisting of simple components,

1350 the big number of heuristics and the multitude of parameters are remarkable. A human expert may  
 1351 have come up with similar ideas, but selecting the components and tuning the parameters would have  
 1352 taken a lot of time.  
 1353

1354 **Expert 3** The code is overall readable and, as in the other case, reflects a consistent style the  
 1355 LLM uses. Splitting the implementation into several functions helps users understand the individual  
 1356 components more easily, while the comments and variable names further enhance readability. Again,  
 1357 the sorting section includes many variants, but the large number of options risks obscuring which  
 1358 criteria actually matter. Overall, the components used are standard for the VRPTW.  
 1359

### 1360 F.2.3 PCVRP

1361 The removal and sorting heuristics form an ensemble of heuristics, with the specific method chosen  
 1362 at random according to static parameters defined at the beginning. The removal operator first de-  
 1363 termines how many customers to remove. It starts the selection from either an unvisited customer,  
 1364 a customer based on a random tour, or a purely random choice. Customers are then added itera-  
 1365 tively based on adjacency and tour neighbors, with random selection used as a fallback. The sorting  
 1366 operator either shuffles customers randomly or scores them using a combination of properties such  
 1367 as prize, distance to the depot, and demand, with probabilities applied to vary the weights of these  
 1368 properties.  
 1369

1370 **Expert 1** The code is well structured as with the other problems. A weakness in this code is one  
 1371 code block in which many “magic numbers” are inserted into the code along with many random  
 1372 numbers. While this code is not hard to understand, it could be written in a cleaner way. The scoring  
 1373 mechanism is likely novel.

1374 **Expert 2** The code is mostly well-written and involves mostly meaningful explanations both in  
 1375 high-and low-level comments. Again, it is reasonably structured into simple and mostly easy-  
 1376 to-understand functions. Both the selection and sorting heuristics seem reasonable and involve  
 1377 mostly known component heuristics and criteria. Compared the CVRP and the VRPTW, the LLM-  
 1378 generated sorting function involves only two main strategies, one of which is a weighted combi-  
 1379 nation where the weights are selected using multiple nested random number draws. This is a bit hard to  
 1380 make sense of, I am sure that a similar behavior could have been achieved in an easier-to-understand  
 1381 way (e.g. by framing the highest level of the nested random selection as separate strategies instead  
 1382 of introducing what the LLM calls “sub-biases” in a single strategy).  
 1383

1384 **Expert 3** The code is well-structured and generally readable, though some functions are longer  
 1385 than in the previous heuristic and combine multiple components that could be separated into dis-  
 1386 tinct functions, as was done in the CVRP and VRPTW implementations. Nevertheless, the code  
 1387 remains understandable thanks to the accompanying comments. While the components themselves  
 1388 are familiar and standard for this problem domain, some of their combinations are unusual.  
 1389

## 1390 F.3 ABLATION STUDY

1391 We conduct ablation studies for the best-performing heuristics to better understand their internal  
 1392 mechanisms. To this end, we systematically remove individual subcomponents from the designed  
 1393 heuristics and assess the impact of each change. To ensure statistically meaningful results, we  
 1394 evaluate all variants on larger validation sets comprising 500 instances per problem, using the same  
 1395 60s per-instance runtime limit as in the main experiments. For every configuration, including the  
 1396 default version (which contains none of the ablated modifications), we run VRPAGENT 10 times  
 1397 on the same set of 500 test instances and record the average tour cost of each run. We report the  
 1398 mean and standard deviation of these run-level averages, together with a 95% confidence interval.  
 1399 Statistical significance relative to the default configuration is determined using the Wilcoxon signed-  
 1400 rank test applied to the paired run averages.  
 1401

### 1402 F.3.1 CVRP

1403 **Customer Removal** We ablate the following subcomponents that are used by the generated heuris-  
 1404 tic during the customer removal procedure:

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

- Initial Segment Removal:** This mechanism is employed as an alternative to selecting a single seed customer. It selects a fixed-length contiguous segment of customers for removal to initiate the destruction process. In the subsequent expansion steps, this tour segment is then extended by other customers.
- Nearest Neighbor Expansion:** Used in the iterative removal loop, this mechanism grows the removed set by spatial clustering around already-selected customers. It selects a candidate  $c$  from the  $k$ -nearest neighbors of a pivot  $p$  using a biased probability distribution, which strongly favors the nearest physical neighbor; this is intended to ensure the final removed set is geographically compact, creating a large, localized "hole".
- Adjacent Tour Node Expansion:** This alternative expansion strategy works by identifying the immediate predecessor and successor of a pivot customer  $p$  in its current tour and then uniformly selecting one of these adjacent (unselected) nodes. Its intention is to ensure the removal of customers that are tightly coupled in the solution tour.
- Route Segment Removal:** This alternative expansion strategy works by identifying a segment of fixed length that includes the pivot  $p$  in its tour and adding all nodes in that segment to the removed set.

Table 4 shows the results of the ablation experiments. Removing Route Segment Removal significantly or Nearest Neighbor Expansion leads to a significant increase in costs. Other components, such as Initial Segment Removal and Adjacent Tour Node Expansion, have only minor effects.

Table 4: Ablation of removal strategies for the CVRP.

Configuration	Mean	Std	CI95	p-value
Default Configuration	36.70378	0.003310	0.002368	—
- w/o Initial Segment Removal	36.70551	0.003138	0.002245	0.375000
- w/o Adjacent Tour Node Expansion	36.70705	0.002539	0.001816	0.064453
- w/o Route Segment Removal	36.71770	0.003249	0.002324	0.001953
- w/o Nearest Neighbor Expansion	37.72012	0.018678	0.013362	0.001953

**Customer Ordering** We ablate the following subcomponents that are used by the generated heuristic during the customer ordering procedure:

- Demand:** A simple ordering rule that orders removed customers based on the customer's demand value in descending order, ensuring high-demand customers are prioritized early during the re-insertion phase.
- Center Proximity:** This ordering rule first calculates the geometric center (centroid) of all removed customers, and then sorts them by their Euclidean distance from this centroid.
- Depot Proximity:** This ordering rule calculates the score as the raw Euclidean distance from the depot. The list is typically sorted in ascending order of this distance.
- Polar Angle:** This rule defines a radial sequencing of insertion priorities. The mechanism calculates the polar angle of each customer relative to the depot (origin) and then sorts the list ascending or descending by this angle to implement a sweep re-insertion strategy.
- Connectivity:** This rule prioritizes customers based on their spatial integration with the removed set. The mechanism calculates the score by measuring the connectivity, determined by the sum of inverse distances between each customer and all other customers in the removed set; this is intended to prioritize the re-insertion of customers that are highly integrated with the destroyed cluster.
- Weighted Demand-Distance Score:** This rule prioritizes customers based on a weighted sum of demand and distance to the depot. The mechanism calculates a linear combination score, where the weights are varied stochastically and a small noise value is applied to increase exploration.
- Nearest Neighbor Chain:** This rule defines a local, greedy sequence. It starts with a random customer and then iteratively finds the **closest unplaced customer** to the most

1458 recently selected one until the entire set is ordered; this generates a local, greedy sequence  
 1459 for the destroyed cluster.

1460 8. **Random Shuffle:** This rule simply orders all removed customers at random.

1461  
 1462 Table 5 shows the results. For most ordering rules, we do not observe a statistically significant  
 1463 impact, which is expected given that many strategies overlap and the effect of each individual rule  
 1464 is limited when numerous strategies are combined. Interestingly, the Connectivity rule stands out as  
 1465 a significant contributor to performance, despite being more computationally expensive than most  
 1466 other strategies.

1467  
 1468 Table 5: Ablation of ordering strategies for the CVRP.

Configuration	Mean	Std	CI95	p-value
Default Configuration	36.70378	0.003310	0.002368	—
- w/o Nearest Neighbor Chain	36.70083	0.001283	0.000918	0.037109
- w/o Demand	36.70116	0.003045	0.002178	0.130859
- w/o Weighted Demand-Distance Score	36.70196	0.001920	0.001374	0.173071
- w/o Random Shuffle	36.70390	0.002159	0.001545	0.625000
- w/o Polar Angle	36.70499	0.001673	0.001197	0.431641
- w/o Connectivity	36.70747	0.001244	0.000890	0.009766
- w/o Center Proximity	36.70984	0.003388	0.002424	0.003906
- w/o Depot Proximity	36.75184	0.002093	0.001498	0.001953

### F.3.2 VRPTW

1482 **Customer Removal** We ablate the following subcomponents that are used by the generated heuristic  
 1483 during the customer removal procedure:

1. **Tour Neighbor Expansion:** This mechanism expands the removed set by identifying the  
 1485 immediate predecessor or successor of a pivot customer in its current tour. It is intended to  
 1486 remove chains of customers that are in the same solution tour.
2. **Nearest Neighbor Expansion:** Used to maintain spatial locality, this mechanism selects  
 1488 a candidate from the closest customers of a pivot customer. It utilizes a biased probability  
 1489 distribution that strongly favors the nearest neighbors to ensure the destroyed set remains  
 1490 geographically compact.
3. **Random Jump Expansion:** This mechanism introduces a perturbation to the growth pro-  
 1492 cess by selecting a completely random unselected customer, regardless of proximity to the  
 1493 current set. It is only triggered once the removed set reaches a minimum size.

1495 Table 6 presents the results. Removing the Random Jump Expansion has no statistically signifi-  
 1496 cant effect on the VRPTW heuristic, while both of the other expansion strategies are critical for  
 1497 maintaining high solution quality.

1498  
 1499 Table 6: Ablation of selection strategies for the VRPTW.

Configuration	Mean	Std	CI95	p-value
Default Configuration	47.46744	0.005987	0.004283	—
- w/o Random Jump Expansion	47.46464	0.004139	0.002961	0.375000
- w/o Tour Neighbor Expansion	47.61712	0.003621	0.002590	0.001953
- w/o Nearest Neighbor Expansion	48.16588	0.005831	0.004171	0.001953

1507 **Customer Ordering** We ablate the following subcomponents that are used by the generated  
 1508 heuristic during the customer ordering procedure:

1. **Time Window Width:** This rule orders removed customers based on the width of their  
 1510 time window in ascending order, ensuring that customers with tighter time constraints are  
 1511 prioritized for re-insertion.

1512 2. **Time Window Start:** This rule orders customers strictly by the start time of their time  
 1513 window in ascending order.  
 1514 3. **Demand:** A rule that orders customers based on their demand value in descending order,  
 1515 prioritizing high-demand customers early in the insertion process.  
 1516 4. **Depot Distance:** This rule orders customers based on their Euclidean distance from the  
 1517 depot in descending order, prioritizing the re-insertion of outlying customers before those  
 1518 closer to the depot.  
 1519 5. **Weighted Start and Width:** This mechanism calculates a deterministic score as a linear  
 1520 combination of the time window start time and half the time window width; it prioritizes  
 1521 customers that start early and have tight constraints.  
 1522 6. **Service Time:** This rule orders customers based on their service duration in descending  
 1523 order, prioritizing those that consume the most route time.  
 1524 7. **Multi-Attribute Score:** This rule prioritizes customers using a linear combination of five  
 1525 normalized attributes: time window start, time window width, demand, service time, and  
 1526 distance. The weights for each attribute are randomized for every execution to maximize  
 1527 exploration of different sorting criteria.  
 1528 8. **Arrival Time Slack:** This rule prioritizes customers based on the maximum possible slack  
 1529 between the earliest arrival (determined by distance from depot) and the latest possible  
 1530 arrival (determined by the time window end). It orders customers with larger slack values  
 1531 first.  
 1532 9. **Weighted Density and Tightness:** This rule calculates a composite score combining de-  
 1533 mand density (demand divided by service time) and time window tightness (inverse of  
 1534 width). It weights these two components using a stochastic factor to alternate priority be-  
 1535 tween high-density and highly-constrained customers.  
 1536 10. **Pure Random Shuffle:** This rule simply orders all removed customers at random.  
 1537

1538 Table 7 shows the results. For most ordering strategies, there is no statistical evidence that they indi-  
 1539 videntally have a meaningful impact on the heuristic’s performance. This is expected, as 10 different  
 1540 strategies are used and selected uniformly at random during the search, so removing a single strategy  
 1541 has only a minor effect. Notably, the Weighted Start and Width heuristic stands out as an exception,  
 1542 significantly contributing to the overall performance.

1544 1545 Table 7: Ablation of ordering strategies for the VRPTW.  
 1546

Configuration	Mean	Std	CI95	p-value
Default Configuration	47.46744	0.005987	0.004283	—
- w/o Demand	47.46509	0.003528	0.002524	0.556641
- w/o Time Window Width	47.46617	0.002892	0.002069	0.431641
- w/o Arrival Time Slack	47.46631	0.006066	0.004339	0.695312
- w/o Multi-Attribute Score	47.46645	0.006211	0.004443	1.000000
- w/o Depot Distance	47.46671	0.007461	0.005337	1.000000
- w/o Time Window Start	47.46708	0.004177	0.002988	1.000000
- w/o Weighted Density and Tightness	47.46756	0.005045	0.003609	1.000000
- w/o Pure Random Shuffle	47.46765	0.006718	0.004806	0.845703
- w/o Service Time	47.46789	0.007746	0.005541	0.921875
- w/o Weighted Start and Width	47.49367	0.005458	0.003905	0.001953

1558 F.3.3 PCVRP  
 1559

1560 **Customer Removal** We ablate the following subcomponents that are used by the generated heuris-  
 1561 tic during the customer removal procedure:

1562 1. **Biased Seed Selection:** Instead of selecting the initial seed customer uniformly at random,  
 1563 this mechanism prioritizes the selection of customers that are currently visited by a vehicle  
 1564 over unvisited customers.

1566  
 1567 2. **Multi-Source Expansion:** Used in the iterative removal process, this strategy selects cus-  
 1568 tomers using multiple seed customer, allowing the destruction area to expand in clusters.  
 1569 3. **Tour-Based Neighbor Expansion:** This mechanism expands the removed set by iden-  
 1570 tifying customers that share the same route as a seed customer. It samples a subset of  
 1571 tour-neighbors to ensure that entire segments of existing routes are targeted are removed.

1572 Table 8 shows the results. All three ablated removal strategies contribute to the overall performance  
 1573 of the heuristic. Interestingly, even subtle algorithmic details, such as the Biased Seed Selection  
 1574 strategy, which starts the customer removal process from previously visited customers with higher  
 1575 probability rather than uniformly at random, can produce small but consistent improvements.

1576  
 1577 Table 8: Ablation of selection strategies for the PCVRP.

Configuration	Mean	Std	CI95	p-value
Default Configuration	42.73483	0.003547	0.002538	—
- w/o Multi-Source Expansion	42.73821	0.001865	0.001334	0.027344
- w/o Biased Seed Selection	42.75303	0.002260	0.001617	0.001953
- w/o Tour-Based Neighbor Expansion	42.97553	0.003856	0.002758	0.001953

1578  
 1579 **Customer Ordering** We ablate the following subcomponents that are used by the generated  
 1580 heuristic during the customer ordering procedure:

1581 1. **Pivot Proximity:** This mechanism orders customers based on their distance to a specific  
 1582 pivot node that is stochastically chosen as either the depot or a random customer from the  
 1583 removed set.

1584 2. **Prize-Focused Scoring:** This configuration corresponds to the first weighting profile. It  
 1585 calculates a weighted sum of the customer’s prize and their distance to the depot, assigning  
 1586 significantly higher weight to the prize value. A small noise factor is added to encourage  
 1587 slight exploration, but the primary intention is to prioritize the re-insertion of high-profit  
 1588 customers.

1589 3. **Prize-Distance Scoring:** This configuration corresponds to the second weighting profile.  
 1590 It assigns roughly equal weights to the customer’s prize and their distance to the depot.  
 1591 To promote solution diversity a significantly larger noise factor is applied compared to the  
 1592 other profiles.

1593 4. **Distance-Demand Scoring:** This configuration corresponds to the third weighting profile.  
 1594 It shifts the heuristic’s focus by heavily weighting the distance to the depot and introducing  
 1595 a weight for customer demand (which is ignored in the other profiles).

1596 5. **Random Shuffle:** This mechanism bypasses the scoring logic entirely and orders all re-  
 1597 moved customers at random. It is selected with a small probability to increase exploration.

1598 Table 9 shows the results. Interestingly, removing the Prize-Distance Scoring slightly improves per-  
 1599 formance, making it the only case in our ablation experiments where removal leads to a statistically  
 1600 significant improvement. In contrast, both the Distance-Demand Scoring and Pivot Proximity order-  
 1601 ing strategies contribute positively and improve performance in a statistically significant manner.

1602  
 1603 Table 9: Ablation of ordering strategies PCVRP.

Configuration	Mean	Std	CI95	p-value
Default Configuration	42.73483	0.003547	0.002538	—
- w/o Prize-Distance Scoring	42.73103	0.002071	0.001481	0.013672
- w/o Prize-Focused Scoring	42.73353	0.002825	0.002021	0.492188
- w/o Random Shuffle	42.73397	0.001956	0.001399	0.921875
- w/o Distance-Demand Scoring	42.74467	0.003440	0.002461	0.003906
- w/o Pivot Proximity	42.86265	0.002567	0.001837	0.001953