The Iterative Chainlet Partitioning Algorithm for the Traveling Salesman Problem with Drone and Neural Acceleration

Jae Hyeok Lee[1], Minjun Kim[2], Jinkyoo Park[1,3], and Changhyun Kwon[*1,3]

[1]Department of Industrial and Systems Engineering, KAIST, Daejeon, 34141, Republic of Korea
[2]Samsung Advanced Institute of Technology, Suwon-si, Gyeonggi-do, 16678, Republic of Korea
[3]Omelet, Inc., Daejeon, 34051, Republic of Korea

Abstract

This study introduces the Iterative Chainlet Partitioning (ICP) algorithm and its neural acceleration for solving the Traveling Salesman Problem with Drone (TSP-D). The proposed ICP algorithm decomposes a TSP-D solution into smaller segments called chainlets, each optimized individually by a dynamic programming subroutine. The chainlet with the highest improvement is updated and the procedure is repeated until no further improvement is possible. The number of subroutine calls is bounded linearly in problem size for the first iteration and remains constant in subsequent iterations, ensuring algorithmic scalability. Empirical results show that ICP outperforms existing algorithms in both solution quality and computational time. Tested over 1,059 benchmark instances, ICP yields an average improvement of 2.75% in solution quality over the previous state-of-the-art algorithm while reducing computational time by 79.8%. The procedure is deterministic, ensuring reliability without requiring multiple runs. The subroutine is the computational bottleneck in the already efficient ICP algorithm. To reduce the necessity of subroutine calls, we integrate a graph neural network (GNN) to predict incremental improvements. We demonstrate that the resulting Neuro ICP (NICP) achieves substantial acceleration while maintaining solution quality. Compared to ICP, NICP reduces the total computational time by 49.7%, while the objective function value increase is limited to 0.12%. The framework's adaptability to various operational constraints makes it a valuable foundation for developing efficient algorithms for truck-drone synchronized routing problems.

Keywords: vehicle routing; traveling salesman problem; drones; deep learning; cost prediction

1 Introduction

Technological advancements have significantly changed various industries in recent years, including the logistics and supply chain sectors. As the demand for faster and more efficient delivery methods grows, traditional delivery systems, which rely heavily on trucks and other ground vehicles, face

---

*Corresponding author

increasing challenges, particularly in densely populated urban areas and remote rural regions. This has led to the exploration and adoption of innovative solutions, among which drones, or unmanned aerial vehicles (UAVs), have emerged as a promising alternative.

To enable drones to become an integral part of the logistics network, we face new challenges and opportunities in optimizing their integration with traditional delivery methods. It is rather unlikely that drones alone will make most deliveries due to short flying ranges. To utilize the full potential of drones, the concept of *drone-assisted delivery* has attracted researchers' attention instead to emphasize the coordination between traditional delivery vehicles and drones. In drone-assisted delivery, trucks are responsible for some of the customer demands, while drones launch from and land on the trucks to make deliveries while preserving the battery.

We emphasize that the significance goes beyond the truck-drone coordination. Autonomous robots are gradually being introduced in various markets to solve logistics problems such as goods deliveries, warehouse operations, security, and other tasks. Such autonomous robots will cooperate with the human workforce, human-driven vehicles, and other autonomous robots. The truck-drone synchronization problem provides a basis for developing computational methods for coordinated operations between heterogeneous entities whose functions depend on and complement each other.

Truck-drone combined routing problems vary widely, encompassing different fleet compositions, operational rules, and optimization objectives. Within this broad domain, our work focuses on the synchronized operation of a truck and a drone serving customers while minimizing completion time. For comprehensive reviews of other problem configurations in truck-drone routing, we refer readers to surveys: Chung et al. (2020), Macrina et al. (2020), and Liang and Luo (2022).

The Flying Sidekick Traveling Salesman Problem (FS-TSP) first formalized single truck-drone synchronized delivery (Murray and Chu, 2015). The problem introduces the concept of a *sortie*, where a unit-capacity drone launches from the truck at one customer location, serves an eligible customer and returns to the truck at a different customer location. Each sortie requires fixed service times for launch and rendezvous, and the drone's flight duration is constrained by battery life until rendezvous with the truck. The problem prohibits revisiting customer locations, and a drone cannot be redeployed once it returns to the depot. The objective is to minimize the total time required to serve all customers and return both vehicles to the depot.

In contrast, the Traveling Salesman Problem with Drone (TSP-D), a variant of the FS-TSP, captures the essence of the problem's collaborative nature among heterogeneous vehicles (Agatz et al., 2018). The TSP-D defines each drone delivery as an *operation*—analogous to sortie of FS-TSP—where a drone departs from the truck to serve a customer before returning for rendezvous. The TSP-D differs from FS-TSP by eliminating all drone-related service times and allowing more flexible operations: Trucks may revisit customer locations, and drones may execute *loops* by returning to their departure location. Additionally, drones may land while awaiting truck rendezvous. Various operational constraints exist; the variant without truck revisitation, loops, customer eligibility restrictions, and drone flying range limitations is called the basic TSP-D (Roberti and Ruthmair, 2021; Blufstein et al., 2024), which we will refer to simply as TSP-D throughout this paper.
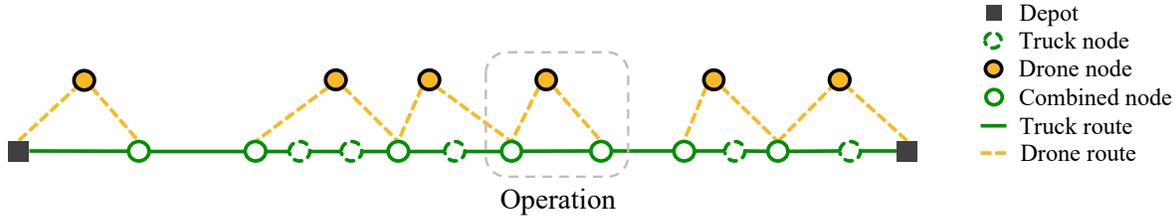
Figure 1: Illustration of Node Types in TSP-D Solution

The TSP-D is defined on a complete graph $\mathcal{G}$ with node set $\mathcal{N} = \{1, 2, \cdots, N\}$, where node 1 represents the depot. The objective is to determine synchronized truck and drone routes that serve all $N-1$ customer nodes exactly once and return to the depot in a minimal time, where the drone operates at a relative speed α to the truck. The drone operates under several assumptions:

1. The drone can visit any customer location;

2. The drone has unlimited flying range and has unit capacity requiring return to the truck after each delivery;

3. Parcel pickup and drone landings are restricted to customer locations or the depot;

4. Loops are prohibited. That is, launch and landing locations are required to be distinct; and

5. All drone-related service times are negligible.

Figure 1 illustrates a TSP-D solution where nodes serve three distinct roles: *truck nodes* visited exclusively by the truck, *drone nodes* serviced solely by the drone, and *combined nodes* serving as synchronization points for both vehicles. Each *operation* consists of exactly two combined nodes (start and end), at most one drone node, and zero or more truck nodes.

Despite simplification, coordinating the truck and drone on time remains challenging. Its inherent NP-hardness poses significant scalability challenges for exact algorithms, driving the need for heuristic solution methods. Researchers have explored a variety of approaches to tackle these challenges, ranging from dynamic programming (Agatz et al., 2018) to advanced metaheuristic algorithms (de Freitas and Penna, 2020; Mahmoudinazlou and Kwon, 2024) and machine learning-based strategies (Bogyrbayeva et al., 2023). While previous approaches have significantly progressed, challenges remain in balancing solution quality and computational efficiency.

In this paper, we aim to develop an approach that enhances both the quality and efficiency of solutions, addressing the trade-offs that have limited the performance of existing methods. We make three main methodological contributions. *First*, we propose the Iterative Chainlet Partitioning (ICP) algorithm for the TSP-D that finds solutions superior to the best-known solutions so far for many instances while reducing computational time significantly. In ICP, the TSP-D route is divided into smaller segments called *chainlets*, each optimized individually by a precise subroutine, namely TSP-EP-ALL of Agatz et al. (2018). During each iteration, newly generated chainlets are

3

evaluated for their potential improvement. The chainlet with the highest potential improvement is selected to update the corresponding part of the route. *Second*, we derive the upper bound on the number of subroutines required for ICP in a closed form. Our theoretical result explains the computational efficiency of the proposed ICP algorithm. *Third*, we implement a machine-learning method to enhance the algorithm's computational efficiency further. Specifically, we employ a graph neural network (GNN) that predicts the potential improvement of each chainlet. This reduces the need for direct executions of subroutines, accelerating the process while maintaining solution quality.

We emphasize two distinctive strengths of ICP. First, unlike many high-quality solution methodologies, which incorporate randomness and necessitate multiple trials to achieve reliable results, ICP is fully *deterministic*. This deterministic nature ensures stable algorithm performance, yielding consistent high-quality solutions in a single execution. Second, ICP offers flexibility regarding subroutines. While we demonstrate its effectiveness using TSP-EP-ALL, the framework readily accommodates any TSP-D optimization subroutine. This adaptability ensures that future algorithmic advances can seamlessly integrate into the ICP framework. Furthermore, while ICP inherently supports customer eligibility constraints and drone range limitations, it can be adapted to other operational settings with minimal modifications by incorporating targeted subroutines. These characteristics establish ICP as a robust foundation for developing efficient algorithms across the spectrum of truck-drone synchronized routing problems.

The rest of this paper is organized as follows. Section 2 briefly summarizes the related literature. Section 3 proposes the ICP algorithm and provides a theoretical analysis of the algorithm, while Section 4 further accelerates ICP using a graph neural network. The performance of ICP and its neural acceleration is tested through extensive numerical experiments in Section 5, and we conclude the paper in Section 6.

2   Literature Review

This section positions our work within the research landscape by examining three key areas: solution methodologies for the TSP-D, decomposition-based metaheuristics for large-scale routing problems, and machine-learning approaches for combinatorial optimization. We first review existing exact and heuristic approaches for solving the TSP-D, highlighting their capabilities and limitations. We then discuss how the ICP algorithm relates to an existing method, POPMUSIC, and explain our neural acceleration techniques in the context of machine learning for routing problems.

2.1   Solution Methodologies for TSP-D

We begin by examining the capabilities of exact methods. A compact formulation without big-M constraints is known to solve instances with up to 14 customers within an hour using an MILP solver, while a branch-and-price algorithm with dynamic programming and *ng*-route relaxation solves instances of up to 29 customers within an hour (Roberti and Ruthmair, 2021). The algorithm was recently advanced, extending solvability to instances with up to 59 customers (Blufstein et al., 2024).

Other exact solution approaches for FS-TSP and TSP-D variants include dynamic programming (Bouman et al., 2018), branch-and-cut (Van Dijck et al., 2018), branch-and-bound (Poikonen et al., 2019), Benders decomposition (Vásquez et al., 2021), and a combination of branch-and-cut with column generation (Boccia et al., 2021). To our best knowledge, the computational capability is limited to 59 customers regardless of which variant is solved.

Given these computational challenges, several heuristic approaches have been developed. Agatz et al. (2018) introduced an exact partitioning algorithm, TSP-EP, that uses dynamic programming to find a minimal TSP-D solution where truck and drone routes are subsequences of an initial TSP route. They extended TSP-EP to TSP-EP-ALL by incorporating local search methods sequentially to perturb the initial TSP route. Although efficient and effective for instances with fewer than 20 customers, TSP-EP-ALL's $O(N^5)$ complexity—combining $O(N^3)$ dynamic programming and $O(N^2)$ local search—limits its scalability, making it more valuable as a subroutine.

Building on this strength, Bogyrbayeva et al. (2023) proposed the Divide-Partition-and-Search (DPS), inspired by the divide-and-conquer heuristic (DCH) from Poikonen et al. (2019). The DPS_{25} partitions the problem into subgroups containing 25 nodes each and applies TSP-EP-ALL to optimize individual subgroups. It demonstrates moderate solution quality while maintaining computational efficiency for large-scale instances. Bogyrbayeva et al. (2023) also developed a deep reinforcement learning approach, proposing a Markov Decision Process formulation and a hybrid model (HM) that combines an attention encoder with a Long Short-Term Memory decoder. HM_{4800}, leveraging GPU to sample 4800 solutions in parallel, efficiently and effectively solves large instances, though it is restricted to predefined problem sizes, distributions, and truck-drone speed ratios.

Metaheuristic approaches have also been developed. The Hybrid General Variable Neighborhood Search (HGVNS) creates an optimal TSP solution with a MIP solver (de Freitas and Penna, 2020). It implements General Variable Neighborhood Search to obtain truck and drone routes, addressing FS-TSP and TSP-D. The Hybrid Genetic Algorithm with type-aware chromosomes (HGA-TAC$^+$) proposed type-aware chromosome encoding that distinguishes truck and drone nodes and incorporates specialized local search methods (Mahmoudinazlou and Kwon, 2024). It demonstrates strong performance for FS-TSP and TSP-D with loops, while readily adaptable to TSP-D.

While existing approaches have improved TSP-D solution capabilities, they face inherent limitations. Metaheuristic approaches deliver high-quality solutions, but their probabilistic nature requires substantial computational time and multiple runs, which introduces variability in solution quality. An end-to-end learning method offers computational efficiency but is limited to problems that match their training parameters and still inherits stochastic characteristics.

Our proposed ICP algorithm addresses these limitations by combining deterministic procedures with high solution quality and computational efficiency. Unlike existing methods, ICP requires no multiple trials to achieve consistent results and functions effectively across various problem configurations, even with neural acceleration. This positions ICP as a significant advancement in solving TSP-D instances, particularly for large-scale applications.
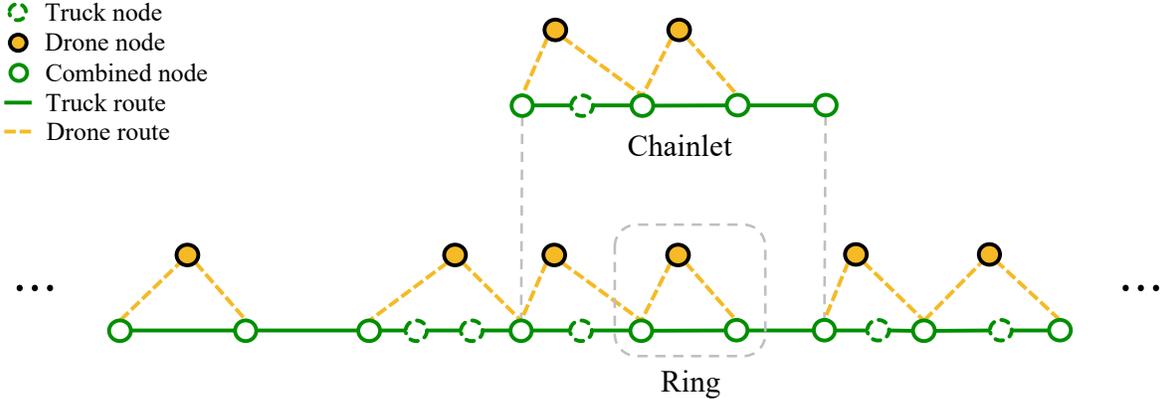
Figure 2: Conceptual Framework for Understanding TSP-D Tour as a Chain

2.2 Comparison with POPMUSIC

ICP shares methodological similarities with the Partial OPtimization Metaheuristic Under Special Intensification Conditions (POPMUSIC) framework (Ribeiro et al., 2002). POPMUSIC is a general framework designed to tackle large-scale combinatorial optimization problems. It has been successfully applied to various routing problems, including the Location Routing Problem (LRP) by Alvim and Taillard (2013), Capacitated Vehicle Routing Problem (CVRP) by Li et al. (2021), and TSP by Taillard and Helsgaun (2019).

POPMUSIC partitions a solution into *parts* and iteratively refines them through *subproblems*. Concretely, one first obtains a feasible solution and decomposes it into multiple parts, each representing a subset of the solution. Then, subproblems are formed by aggregating a chosen seed part with other related parts. Because these subproblems are constructed so that any improvement remains compatible with the original solution, each local improvement directly translates into a valid improvement of the global solution. The process repeats until no subproblem can be further improved.

While the POPMUSIC framework provides a general template, adapting it to the TSP-D requires addressing specific challenges. Unlike the TSP, which allows natural decomposition into consecutive segments of cities while preserving feasibility, the TSP-D requires synchronizing truck and drone routes, necessitating valid launch and rendezvous points for any partial solution. In ICP, we address this challenge through the concept of a *chainlet*. Figure 2 illustrates our framework for structuring TSP-D routes. We define the entire TSP-D tour as a *chain*, wherein each operation is viewed as a single *ring*. By grouping multiple rings, we define a route segment as a *chainlet*, representing a "small chain" within the entire route. Since TSP-EP-ALL naturally preserves the starting and ending combined nodes (treating them as depots), with careful initial tour construction, we can leverage this subroutine to optimize each chainlet without compromising the chain's integrity. Consequently, improvements in individual chainlets directly translate to improvements in the overall solution.

ICP differs from traditional POPMUSIC implementations in its selection strategy for subproblems.

6

While POPMUSIC can be viewed as a gradient method that improves the initial solution until reaching a local optimum (Taillard and Voss, 2018), ICP implements a steepest descent variant that always moves toward maximum improvement at each iteration. Specifically, whereas many POPMUSIC implementations randomly select subproblems or employ simple shifting strategies after modifications, ICP evaluates all generated chainlets and selects them greedily based on their potential for improvement.

Research into the influence of selection strategies within the POPMUSIC framework has been limited (Taillard and Voss, 2018). We identify several advantages of our greedy selection approach: First, when using only deterministic subroutines, the algorithm becomes deterministic, ensuring performance stability. Second, we provide theoretical explanation of the algorithm's scalability, analogous to POPMUSIC's empirically observed near-linear scaling when subproblem sizes remain fixed (Ostertag et al., 2009; Alvim and Taillard, 2013; Taillard and Helsgaun, 2019). Since we optimize only the chainlet with the highest improvement potential at each iteration, large portions of the solution remain unchanged between iterations. By storing previous optimization results, we effectively optimize all subproblems in the first iteration, and then, for subsequent iterations, we only need to optimize subproblems affected by the previous update. Using this fact with a fixed subproblem size, we can establish that a constant bounds the number of subproblems to be optimized after the first iteration. More precisely, we derive an upper bound on the number of subproblems to be solved that is linear in N for the first iteration and constant thereafter. Consequently, ICP inherits POPMUSIC's hallmark of scalability while offering more concrete theoretical foundations, providing an efficient local search algorithm for large-scale TSP-D instances.

2.3   Learning Approaches for Routing Problems

Based on deep neural networks, machine-learning methods have been applied to solve combinatorial optimization problems, especially the TSP and other vehicle routing problems. In early developments, deep reinforcement learning approaches were dominant in solving routing problems in an end-to-end fashion (Bello et al., 2016; Khalil et al., 2017; Kool et al., 2018; Nazari et al., 2018). While such approaches have demonstrated that neural networks could learn heuristics for solving NP-hard combinatorial optimization problems and produce quality solutions quickly, a well-crafted, human-developed algorithm often outperforms end-to-end learning-based approaches by a large margin. For instance, the Lin–Kernighan–Helsgaun (LKH) algorithm (Helsgaun, 2017) is reported to outperform end-to-end neural methods in both solution quality and, in some cases, computational time on randomly generated TSP instances (Sui et al., 2025).

Consequently, a new strand of *hybrid* methodologies has emerged to bridge the gap between purely handcrafted methods and end-to-end learned approaches, Some approaches replace human-designed subroutines with machine-learned heuristics, such as neural repair in large neighborhood search (Hottung and Tierney, 2020), neural separation in cutting-plane methods (Kim et al., 2024), and Language Hyper-Heuristics using large language models for heuristic generations (Ye et al., 2024). Others enhance search procedures, including neural scoring in the LKH algorithm (Xin et al.,

2021) and neural fitness prediction in genetic algorithms (Sobhanan et al., 2025). These hybrid methods have proven highly effective, maintaining the sophisticated structure of classical heuristics while strategically enhancing specific components through learning.

Our work reinforces this pattern: while the purely end-to-end learning-based approach, HM_{4800}, performs reasonably well, ICP, a carefully constructed heuristic, surpasses both solution quality and speed. Therefore, we adopt a hybrid approach that accelerates ICP using a learning model to predict subproblem solution costs, similar to the strategy used for POPMUSIC in CVRP (Li et al., 2021). As numerical experiments confirm, our hybrid approach demonstrates that integrating a learning model can accelerate an already efficient computational method without compromising solution quality.

3   Solution Methodology

This section presents the methodology developed for solving TSP-D. The first subsection introduces the Iterative Chainlet Partitioning (ICP) algorithm. The second subsection presents theoretical insights into the structural properties of the ICP algorithm, establishing formal bounds on the number of TSP-EP-ALL runs required per iteration, which demonstrates the scalability and computational efficiency of ICP.

3.1   The Iterative Chainlet Partitioning Algorithm

We first present the formal definition of *ring*, *chainlet*, and *chain*.

Definition 1. A *ring*, equivalent to an operation, consists of exactly two combined nodes serving as the start and end, at most one drone node, and zero or more truck nodes. A *chainlet* is a sequence of consecutive rings, and a *chain* is the complete set of rings forming the TSP-D solution. The size of each structure is defined as the number of distinct nodes it contains.

Built on this conceptual framework, Figure 3 illustrates the overall iteration process of the ICP algorithm. The algorithm begins by dividing the current chain into individual rings, which are then grouped into overlapping chainlets of manageable size. Each chainlet is evaluated using the precise TSP-EP-ALL heuristic to determine its potential for improvement. The chainlet with the highest potential improvement is selected to update the corresponding segment of the chain, and the process iterates until no further improvements are possible.

The process of the ICP algorithm is detailed in Algorithm 1. The algorithm begins by constructing an optimal TSP tour $\mathcal{T}$ using the Concorde solver (Applegate et al., 1998) (Line 3), which is then transformed into the initial chain $\mathcal{C}$ using the TSP-EP heuristic (Line 4). Rather than grouping a fixed number of rings into each chainlet, the algorithm dynamically adjusts the number of rings based on the chainlet's size, ensuring that each chainlet $\mathcal{C}_j$ remains within a maximum size of 20 (Line 7). The GROUP process begins by sequentially adding rings to the current chainlet $\mathcal{C}_j$ until adding another ring causes the chainlet to exceed the size limit. Once the chainlet reaches this limit,
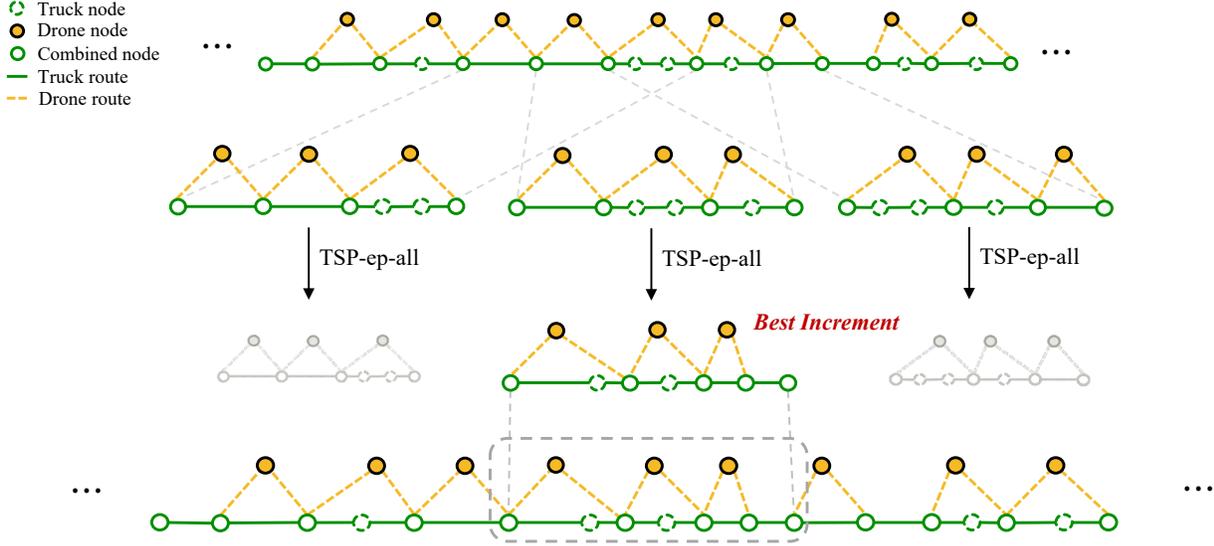
Figure 3: Illustrative Example of ICP Iteration

Algorithm 1 Iterative Chainlet Partitioning (ICP)

1: **Input:** Customer Node Set $\mathcal{N}$
2: **Output:** Optimized Chain $\mathcal{C}$
3: $\mathcal{T} \leftarrow \text{CONCORDE}(\mathcal{N})$
4: $\mathcal{C} \leftarrow \text{TSP-EP}(\mathcal{T})$
5: Initialize `cache` for chainlets
6: **repeat**
7:      $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_m \leftarrow \text{GROUP}(\mathcal{C})$ ▷ Group rings into chainlets
8:      **for** $j \in \{1, 2, \ldots, m\}$ **do**
9:          **if** $\text{HASH}(\mathcal{C}_j) \in \texttt{cache}$ **then**
10:              $\mathcal{C}'_j, \Delta_j \leftarrow \texttt{cache}[\text{HASH}(\mathcal{C}_j)]$
11:          **else**
12:              $\mathcal{T}_j \leftarrow \text{FARTHESTINSERTION}(\mathcal{C}_j)$ ▷ Farthest insertion with fixed end
13:              $\mathcal{C}'_j \leftarrow \text{TSP-EP-ALL}(\mathcal{T}_j)$
14:              $\Delta_j \leftarrow \texttt{Cost}(\mathcal{C}_j) - \texttt{Cost}(\mathcal{C}'_j)$
15:              $\texttt{cache}[\text{HASH}(\mathcal{C}_j)] \leftarrow (\mathcal{C}'_j, \Delta_j)$
16:              $\texttt{cache}[\text{HASH}(\mathcal{C}'_j)] \leftarrow (\mathcal{C}'_j, 0)$
17:      $k \leftarrow \arg\max_{i \in \{1, 2, \ldots, m\}} \Delta_i$
18:      **if** $\Delta_k > 0$ **then**
19:          Replace $\mathcal{C}_k$ in $\mathcal{C}$ with $\mathcal{C}'_k$
20: **until** $\Delta_i \leq 0 \ \forall i \in \{1, \ldots, m\}$
21: **return** $\mathcal{C}$

---

a new chainlet is initiated. If the newly generated chainlet is a subset of the previously generated chainlet, it is omitted from further consideration. The process then advances the starting ring to the next in sequence and continues until all rings have been grouped. Adjusting the number of rings per chainlet ensures consistent algorithmic performance across various settings, as it controls the maximum chainlet size regardless of how α affects individual ring sizes. The selection of this maximum node size balances several factors: grouping too many rings in a chainlet can slow the algorithm due to the exponential complexity of TSP-EP-ALL, while grouping too few may reduce the subroutine's effectiveness.

We leverage a `cache` that stores previous TSP-EP-ALL results to avoid redundant calculations. Since only the chainlet with the most significant improvement is used to update the corresponding segment of the chain in each iteration (Lines 17–19), the other non-overlapping chainlets remain unchanged in subsequent iterations. Therefore, for previously encountered chainlets, the optimized chainlet $\mathcal{C}'_j$ and its corresponding improvement Δ_j are retrieved directly from the `cache`, significantly enhancing the algorithm's efficiency (Lines 9–10). For chainlets not found in the cache, we construct an input tour $\mathcal{T}_j$ using the FARTHESTINSERTION (FI) heuristic (Line 12). Since altering the chainlet's end node would disrupt the chain's continuity by breaking its connection to subsequent segments, the FI heuristic preserves the integrity of the overall chain by first connecting the start node to the end node before inserting the remaining nodes. With the constructed input tour, each chainlet is then optimized using the TSP-EP-ALL subroutine (Line 13), and the optimization result is stored in the cache (Lines 15–16). We hash each chainlet using its node sequence, specifically the intersection of truck and drone routes. Within each ring, drone nodes always precede truck nodes. While this hashing approach does not produce a unique identifier for every conceivable chainlet configuration, it is sufficient to guarantee consistent optimization results when the same hash is encountered.

The following proposition formalizes the behavior of the ICP algorithm, highlighting its deterministic structure, monotonic decrease in solution cost, and termination within a finite number of iterations. These properties are direct consequences of the algorithm's design, which ensures a strict reduction in cost at each step through a greedy selection strategy and deterministic subroutines, with revisitation of prior configurations inherently ruled out by the update mechanism.

Proposition 1. The ICP algorithm terminates in finite time, and the cost of the solution monotonically decreases throughout the iterations. In addition, if the instance has a unique TSP optimal solution, the ICP algorithm is deterministic; that is, every run of ICP for any given TSP-D instance results in the same solution.

We provide all proofs in the Appendix A. The choice of the maximum node size and the exclusive employment of the Concorde solver and FI heuristic for input tour construction is based on the experimental results outlined in Appendix B. The procedures of the GROUP and FARTHESTINSERTION heuristics are detailed in Appendix C.
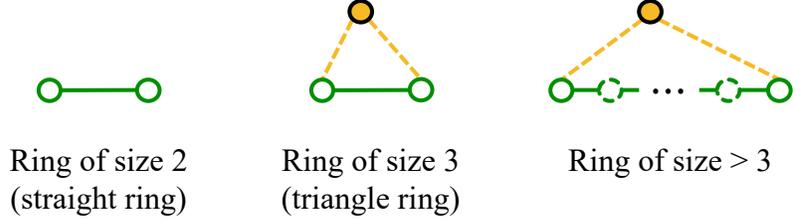
Figure 4: Illustration of Different Rings in a Chain

3.2 Theoretical Insights

This section presents theoretical results that quantify the efficiency of the ICP algorithm by formalizing its structural properties, particularly in the formation and partitioning of rings and chainlets. These results demonstrate ICP's capability to handle large-scale TSP-D problems effectively, with specific bounds established for the number of TSP-EP-ALL runs required per iteration.

Definition 2. A ring of size 2 is called a *straight ring*. A ring of size 3 is called a *triangle ring*.

First, we distinguished different rings based on the number of drone and truck nodes. A ring of size 2, which does not utilize the drone and consists solely of two combined nodes, is called a *straight ring*. A ring of size 3, which utilizes the drone but contains no intermediate truck nodes—meaning the truck directly traverses to the rendezvous point without visiting other nodes—is termed a *triangle ring*. Naturally, as the effectiveness of deploying the drone decreases with lower α, larger rings containing additional truck nodes become more prevalent. This consequently reduces the total number of rings within a TSP-D chain.

Lemma 1. Suppose a TSP-D instance where $\alpha > 1$ and no nodes are collinear. Then, no two consecutive straight rings can appear in an optimal TSP-D solution. The same holds for exact partitioning solutions for any given TSP tour.

Lemma 1 shows that for $\alpha > 1$ and non-collinear nodes, two consecutive straight rings do not appear in an optimal solution, as a single triangular ring provides a lower-cost alternative. The result also holds for exact partitioning solutions, as exact partitioning guarantees optimality within the current TSP sequence.

The structural property given by Lemma 1 enables the calculation of the maximum number of rings in a chain. Each ring must contain minimal nodes, achieved by maximizing the alternation between straight and triangular rings as shown in Figure 5. The following lemma formalizes this idea.

Lemma 2. Suppose a chain is constructed from a TSP-D instance where $\alpha > 1$ and no nodes are collinear. Then, for any chainlet with $n \geq 3$ nodes, the maximum number of rings is

$$\frac{2n - 3 + (n \bmod 3)}{3}. \tag{1}$$

Figure 5: Chainlet with 20 nodes exhibiting the maximum number of rings

Lemma 2 provides the maximum number of rings in any chainlet. Together with the caching mechanism—which only treats chainlets containing at least one ring from a modified chainlet as new—this result bounds the number of TSP-EP-ALL subroutine runs necessary per ICP iteration.

Proposition 2. Suppose a chain constructed from a TSP-D instance with N total nodes where $\alpha > 1$, no nodes are collinear, and N is sufficiently large. The maximum node size per chainlet in ICP is denoted as ℓ. Then, in the first iteration of ICP, the maximum number of TSP-EP-ALL runs is

$$\frac{2N - 1 + \left((N+1) \bmod 3\right)}{3}. \tag{2}$$

In each subsequent iteration, the maximum number of TSP-EP-ALL runs is

$$\frac{4\ell - 9 + 2(\ell \bmod 3)}{3}. \tag{3}$$

Proposition 2 indicates that in the initial iteration, the number of TSP-EP-ALL runs is linear in N. However, this number remains constant in subsequent iterations, regardless of the instance size. With $\ell = 20$ in ICP, the necessary TSP-EP-ALL runs do not exceed 25 in later iterations.

Our analysis in Section 5.1 substantiates and extends these theoretical bounds. We establish that TSP-EP-ALL constitutes the computationally dominant component of ICP's runtime, particularly at higher α values. Our empirical results further confirm that the number of iterations grows linearly with N. Note that each TSP-EP-ALL execution requires constant time due to the fixed maximum chainlet size of 20 nodes. Combining the linearity of iterations with our theoretical bound on subroutine calls per iteration, we conclude that the bound of the total number of TSP-EP-ALL executions scales as $O(N)$. This linear scaling property is instrumental in ensuring ICP's computational efficiency for large-scale instances.

4 Neural Acceleration

This section details constructing a neural network model to predict the outcomes of the TSP-EP-ALL heuristic and its integration into the ICP algorithm to enhance computational efficiency.

Despite the efficiency gained from caching in the ICP algorithm, a significant computational burden persists due to the need to compute TSP-EP-ALL for multiple chainlets in each iteration. To address this issue, we propose the integration of a neural cost predictor. This neural network estimates the cost increment, allowing the algorithm to bypass the direct execution of TSP-EP-ALL except for the chainlet with the highest predicted improvement. By leveraging GPU parallelization,
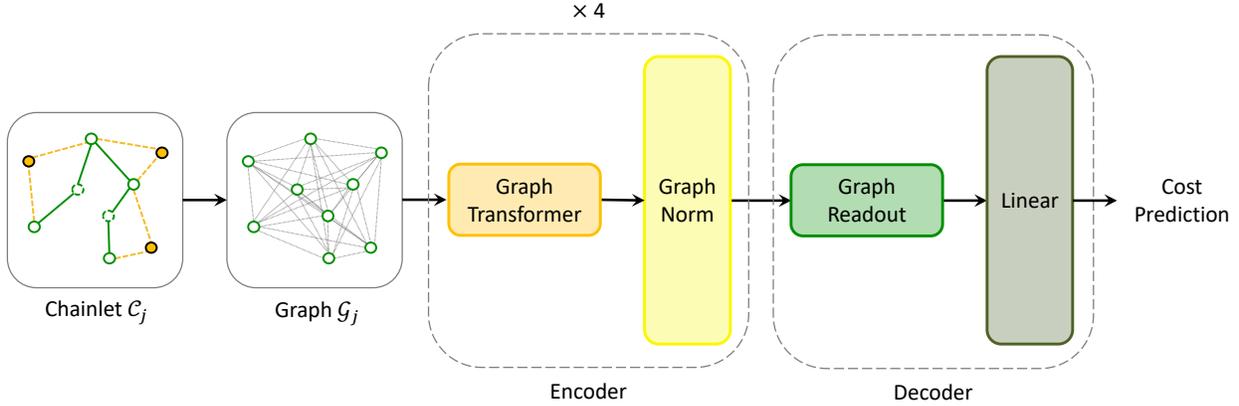
Figure 6: The neural cost predictor

the neural network can rapidly process multiple chainlets simultaneously, ensuring efficient prediction of potential improvements. Consequently, this approach effectively reduces computational demands while preserving the integrity of the solution.

4.1 Neural Cost Predictor

Figure 6 depicts the overall neural network architecture. A GNN is utilized, specifically chosen for its capability to effectively process input sizes that vary, such as the chainlets generated within the ICP algorithm. Let graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes with $|\mathcal{V}| = n$, and $\mathcal{E}$ is the set of edges. Each node $v \in \mathcal{V}$ is associated with an initial feature vector $\mathbf{X}_v$, and each edge $e \in \mathcal{E}$ has a corresponding feature vector $\mathbf{X}_e$. In a GNN, the representation of each node $\mathbf{h}_v$ evolves through the network layers, starting with $\mathbf{h}_v = \mathbf{X}_v$. The GNN iteratively updates these representations into $\mathbf{h}'_v$ through message passing, wherein each node aggregates information from its neighbors, thereby capturing k-hop neighborhood interactions within the graph. The task is to learn the parameters $\boldsymbol{\theta}$ of the cost prediction GNN f_θ, which estimates the TSP-EP-ALL cost $\widehat{y}$ for a given chainlet $\mathcal{C}_j$ based on its graph representation $\mathcal{G}_j$.

$$\mathcal{G}_j \xleftarrow[\text{construction}]{\text{graph}} \mathcal{C}_j$$

$$\widehat{y} = f_\theta(\mathcal{G}_j) := \texttt{Decoder}(\texttt{Encoder}(\mathcal{G}_j)),$$

The details of the graph construction, $\texttt{Encoder}(\cdot)$, and $\texttt{Decoder}(\cdot)$ are described in the following paragraphs.

Graph Construction The graph construction encapsulates the TSP tour sequence and truck and drone costs, precisely embedding the key inputs TSP-EP-ALL needs. We construct a complete graph $\mathcal{G}_j$ for each chainlet $\mathcal{C}_j$, with a node for each node in the chainlet, ensuring all possible connections are considered. From the initial TSP tour generated by the FI heuristic, node features $\mathbf{X}_v$ are

derived using sinusoidal positional encoding, which embeds the positional sequence of nodes within the tour. Each node v in the graph is assigned a feature vector $\mathbf{X}_v$ that captures its position in the sequence, using trigonometric functions to encode relative and absolute positions in a manner that preserves order information. Edge features $\mathbf{X}_e$ include truck and drone costs, normalized by dividing each cost by the maximum truck cost across all edges, ensuring the costs are scaled between 0 and 1.

Encoder   The $\texttt{Encoder}(\cdot)$ transforms the input graph $\mathcal{G}_j$ into latent vectors $\mathbf{H'} := \{\mathbf{h'}_1, \mathbf{h'}_2, \ldots, \mathbf{h'}_n\} \in \mathbb{R}^{n \times d}$, where d is the hidden dimension. It consists of a stack of alternating graph transformer layers and graph normalization layers. The graph transformer layer, based on (Shi et al., 2020), is a graph adaptation of the multi-head attention mechanism of the transformer (Vaswani et al., 2017), chosen for its ability to incorporate edge features into the attention process and its support for parallel multi-head processing for improved computational efficiency. Each GT layer computes multiple heads to transform the node feature $\mathbf{h}_v$ to $\mathbf{h'}_v$, updating the representation through a combination of residual connections and message passing from connected nodes. Specifically, each head computes the following transformation:

$$\mathbf{h'}_v = \mathbf{W}^r \mathbf{h}_v + \underbrace{\left(\sum_{u \in \mathcal{V} \backslash \{v\}} a_{uv} \left(\mathbf{W}^v \mathbf{h}_u + \mathbf{W}^e \mathbf{e}_{uv} \right) \right)}_{\mathbf{m}_v},$$

where the aggregated message $\mathbf{m}_v$ reflects the combined influence of connected nodes within the graph. Matrices $\mathbf{W}^r$, $\mathbf{W}^v$, and $\mathbf{W}^e$ are learnable, with $\mathbf{W}^r$ handling the residual connection, $\mathbf{W}^v$ performing value transformations to messages from other nodes and $\mathbf{W}^e$ processing the edge features $\mathbf{e}_{uv}$. The attention coefficient a_{uv} quantifies the importance of the message from node u to node v and is computed as:

$$a_{uv} = \mathrm{softmax} \left(\frac{(\mathbf{W}^q \mathbf{h}_v)^\top (\mathbf{W}^k \mathbf{h}_u + \mathbf{W}^e \mathbf{e}_{uv})}{\sqrt{\bar{d}}} \right),$$

where $\mathbf{W}^q$ and $\mathbf{W}^k$ are learnable weight matrices for the query and key transformations, respectively. Each head's results are independently calculated and concatenated, enabling the model to capture diverse interactions and enhance its expressive capabilities.

After passing through each GT layer, the node states are processed by a graph normalization (GraphNorm) layer (Cai et al., 2021). Normalization layers improve training stability and convergence by shifting and scaling input features, with a specific set of features for normalization depending on the type. Instance normalization for graph data treats each graph as an individual instance, normalizing the node features based on their mean and variance across all nodes within that graph. However, this standard mean shift can inadvertently discard structural information, as the aggregated mean of node features often captures unique patterns intrinsic to each graph. GraphNorm addresses this issue by introducing a learnable vector $\boldsymbol{\lambda}$, which directly controls the extent of the mean shift applied during normalization. By allowing the model to adjust the degree to which the

mean is subtracted, $\boldsymbol{\lambda}$ enables the preservation of critical structural information embedded in the mean of the node features while still achieving the benefits of normalization. GraphNorm for $\mathbf{h}'_v$ can be expressed as:

$$\mathrm{GraphNorm}(\mathbf{h}'_v) = \mathrm{diag}\left(\sqrt{\mathrm{Var}[\mathbf{h}'_v - \boldsymbol{\lambda}^\top \mathbb{E}[\mathbf{H}']]}\right)^{-1} \left(\mathbf{h}'_v - \boldsymbol{\lambda}^\top \mathbb{E}[\mathbf{H}']\right) \mathrm{diag}(\boldsymbol{\gamma}) + \boldsymbol{\beta},$$

where $\mathbb{E}[\mathbf{H}']$ and $\mathrm{Var}[\mathbf{H}']$ denote the element-wise mean and variance of the features across all nodes in the graph, $\boldsymbol{\gamma}$ is a learnable scaling vector represented as a diagonal matrix, and $\boldsymbol{\beta}$ is a learnable shifting vector.

After the GraphNorm layer, an Exponential Linear Unit (ELU) activation function is applied to introduce non-linearity, mitigating the vanishing gradient problem and accelerating learning by pushing the mean activations closer to zero and providing noise-robust representations (Clevert, 2015). To summarize, the `Encoder`$(\cdot)$ consists of a stack of four components, each sequentially applying a GT layer followed by a GraphNorm and ELU activation, thereby transforming the input graph representation $\mathcal{G}_j$ into its latent representation.

$$\mathbf{h}'_v = \mathrm{ELU}(\mathrm{GraphNorm}(\mathrm{GT}(\mathbf{h}_v))).$$

Decoder   The `Decoder`$(\cdot)$ maps the latent vectors $\mathbf{H}'$ to the predicted TSP-ep-all cost $\widehat{y}$ as:

$$\widehat{y} = \mathbf{W}\left(\underbrace{\sum_{v \in \mathcal{V}} \mathrm{softmax}\left(\mathbf{h}'_v\right) \mathbf{h}'_v}_{\text{graph readout}}\right) + b.$$

The graph readout module calculates attention weights for each node's features using a softmax function, assigning higher weights to more critical features. The weights are applied element-wise to the node features, amplifying significant information while diminishing less relevant features. The weighted features are then summed across all nodes to obtain the global graph representation $\mathbf{h}_{\mathcal{G}} \in \mathbb{R}^d$, capturing the key characteristics of the graph. Finally, $\mathbf{h}_{\mathcal{G}}$ passes through a linear layer with a learnable weight matrix $\mathbf{W}$ and bias term b, producing the predicted TSP-EP-ALL cost $\widehat{y}$.

4.2   The Neuro Iterative Chainlet Partitioning Algorithm

We denote the ICP integrated with the neural cost predictor as Neuro Iterative Chainlet Partitioning (NICP). The process of NICP is detailed in Algorithm 2. The initial steps—constructing the optimal TSP tour $\mathcal{T}$ using the Concorde solver (Line 3), transforming it into the chain $\mathcal{C}$ using the TSP-EP heuristic (Line 4), and grouping rings into chainlets $\mathcal{C}_j$ with a maximum size of 19 (Line 7)—are identical to those in ICP. NICP diverges from ICP by incorporating a neural network to predict the TSP-EP-ALL cost for each chainlet before applying the heuristic. While substituting direct TSP-EP-ALL executions with neural predictions could diminish the benefit of caching, NICP maintains efficiency through a two-level caching mechanism that uses separate stores for optimized results

Algorithm 2 Neuro Iterative Chainlet Partitioning (NICP)

1: **Input:** Customer Node Set $\mathcal{N}$
2: **Output:** Optimized Chain $\mathcal{C}$
3: $\mathcal{T} \leftarrow \textsc{Concorde}(\mathcal{N})$
4: $\mathcal{C} \leftarrow \text{TSP-EP}(\mathcal{T})$
5: Initialize `cache.optimized` and `cache.predicted`
6: **repeat**
7:      $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_m \leftarrow \textsc{Group}(\mathcal{C})$ $\triangleright$ Group rings into chainlets
8:      **for** $j \in \{1, 2, \ldots, m\}$ **do**
9:          **if** $\textsc{Hash}(\mathcal{C}_j) \in$ `cache.optimized` **then**
10:              $\mathcal{C}'_j, \Delta_j \leftarrow$ `cache.optimized`$[\textsc{Hash}(\mathcal{C}_j)]$
11:              $\widehat{\Delta}_j \leftarrow \Delta_j$
12:          **else if** $\textsc{Hash}(\mathcal{C}_j) \in$ `cache.predicted` **then**
13:              $\mathcal{T}_j, \widehat{\Delta}_j \leftarrow$ `cache.predicted`$[\textsc{Hash}(\mathcal{C}_j)]$
14:          **else**
15:              $\mathcal{T}_j \leftarrow \textsc{FarthestInsertion}(\mathcal{C}_j)$ $\triangleright$ Farthest insertion with fixed end
16:              $\widehat{\Delta}_j \leftarrow$ `Predictor`$(\mathcal{T}_j)$
17:              `cache.predicted`$[\textsc{Hash}(\mathcal{C}_j)] \leftarrow (\mathcal{T}_j, \widehat{\Delta}_j)$
18:      $k \leftarrow \arg\max_{i \in \{1, 2, \ldots, m\}} \widehat{\Delta}_i$
19:      **if** $\textsc{Hash}(\mathcal{C}_k) \notin$ `cache.optimized` **then**
20:          $\mathcal{C}'_k \leftarrow \text{TSP-EP-ALL}(\mathcal{T}_k)$
21:          $\Delta_k \leftarrow \texttt{Cost}(\mathcal{C}_k) - \texttt{Cost}(\mathcal{C}'_k)$
22:          `cache.optimized`$[\textsc{Hash}(\mathcal{C}_k)] \leftarrow (\mathcal{C}'_k, \Delta_k)$
23:          `cache.optimized`$[\textsc{Hash}(\mathcal{C}'_k)] \leftarrow (\mathcal{C}'_k, 0)$
24:      **if** $\Delta_k > 0$ **then**
25:          Replace $\mathcal{C}_k$ in $\mathcal{C}$ with $\mathcal{C}'_k$
26: **until** $\forall i \in \{1, \ldots, m\} : \Delta_i \leq 0$ and $\textsc{Hash}(\mathcal{C}_i) \in$ `cache.optimized`
27: **return** $\mathcal{C}$

(`cache.optimized`) and prediction data (`cache.predicted`).

The caching mechanism in NICP operates as follows: First, the algorithm checks if the results from a previous TSP-EP-ALL run, Δ_j and $\mathcal{C}'_j$, are already available (Line 9). These actual results take precedence when available, bypassing neural prediction (Lines 10–11). Otherwise, the algorithm searches for existing neural predictions. If found, it retrieves the cached input tour $\mathcal{T}_j$ and prediction $\widehat{\Delta}_j$ (Lines 12–13). For chainlets without any cached information, $\mathcal{T}_j$ is constructed using the FI heuristic. The neural network predicts the TSP-EP-ALL cost $\widehat{y}_j$ using GPU parallelization for efficient processing of multiple chainlets, and the predicted cost is rescaled by multiplying the maximum truck cost across all edges to restore it to the original scale. The expected improvement $\widehat{\Delta}_j$ is subsequently derived from this rescaled cost, and cached along with the input tour (Lines 14–17).

The chainlet $\mathcal{C}_k$ with the highest improvement, whether actual or predicted, is selected for updating (Line 18). If TSP-EP-ALL has not yet been applied to $\mathcal{C}_k$, it is executed using the cached input tour $\mathcal{T}_k$, and both the actual improvement Δ_k and the optimized chainlet $\mathcal{C}'_k$ are cached (Lines 19–23). The chain is updated only if an improvement is confirmed (Lines 24–25). The iterative process continues until all chainlets have been processed by TSP-EP-ALL and no further improvements are possible, ensuring comprehensive exploration of the chain.

Proposition 3. The NICP algorithm terminates in finite time, and the cost of the solution monotonically decreases throughout the iterations. Given that the instance has a unique TSP optimal solution, the NICP algorithm is deterministic.

The key arguments in Proposition 1 for the ICP algorithm also apply to NICP. Although the neural predictor may yield imprecise estimates, every update is finalized only after TSP-EP-ALL confirms a strict improvement, ensuring finite termination. Moreover, for a given set of parameters θ, neural inference is deterministic, and hence NICP remains deterministic.

5   Computational Study

This section presents a computational study, including a performance evaluation of the ICP algorithm, neural network training, and a comparison between the ICP and NICP algorithms. All experiments were conducted on a workstation featuring an AMD Ryzen 9 5900X 12-Core Processor (3.7 GHz, 24 logical processors), 64 GB of DDR4 RAM at 3200 MT/s, and an NVIDIA GeForce RTX 4070 GPU with 12 GB of memory, running Ubuntu 22.04.3 LTS. The neural network was implemented and trained in Python 3.10, and the algorithms were developed in Julia 1.10.0, utilizing PyCall for cost predictions. The code is publicly available at `https://github.com/0505daniel/TSPDroneICP.jl`.

The evaluation utilized two benchmark sets. *Set A* comprises 759 instances selected from Agatz et al. (2018), divided into three subsets based on distribution: uniform, 1-center, and 2-center. Each subset includes node sizes $N \in \{50, 75, 100, 175, 250, 375, 500\}$ and truck–drone speed ratios $\alpha \in \{1, 2, 3\}$. In the uniform subset, each customer's (x, y) coordinate is independently sampled from $\{0, 1, \ldots, 100\}$, and the depot is chosen from $[0, 1]$. The 1-center subset simulates a circular city center by drawing points from a radial distribution with mean 0 and standard deviation 50

around the origin. In contrast, the 2-center subset shifts each point by 200 units along the x-axis with probability 0.5, thus creating two clusters centered at $(0,0)$ and $(200,0)$. For the 1-center and 2-center subsets, the first generated location serves as the depot.

Set B, generated by Bogyrbayeva et al. (2023), consists of 300 instances at $\alpha = 2$, split into two subsets based on distribution: random and Amsterdam. The random subset contains 200 instances, with 100 each at $N = 50$ and $N = 100$. Each coordinate is drawn from $[1, 100]$, and the first sampled location is designated as the depot. The Amsterdam subset comprises 100 instances with $N = 50$, constructed using real-world electric vehicle (EV) parking locations in Amsterdam (Haider et al., 2019), thus reflecting more realistic spatial conditions. All instances used have at least 50 nodes, ensuring sufficient chainlets for evaluation. This comprehensive composition covers a range of node sizes, distributions, and truck-drone speed ratios, allowing thorough investigation of ICP and NICP under diverse problem settings.

Collectively, these benchmark sets enable comprehensive evaluation: Set A provides systematic coverage of varying problem dimensions (node sizes, truck-drone speed ratios, and spatial distributions), while Set B incorporates realistic urban spatial patterns. All instances contain at least 50 nodes, ensuring sufficient chainlet generation for thorough algorithmic evaluation.

We evaluated our algorithms against four benchmark methods: TSP-EP-ALL (dynamic programming), DPS_{25} (divide-and-conquer), HM_{4800} (end-to-end learning), and HGA-TAC$^+$ (metaheuristic)[1]. Due to computational constraints, TSP-EP-ALL is applied only to Set B. HM_{4800}, though its high solution quality and computational efficiency, was specifically trained for instances with $\alpha = 2$ and $N \in \{11, 15, 20, 50, 100\}$ in random and Amsterdam distributions, limiting its applicability to these particular problem configurations. Consequently, it is applied exclusively for Set B. All benchmark algorithms were implemented in Julia and executed in the same computational environment. For Set A, we compare against DPS_{25} and HGA-TAC$^+$, while for Set B, we include all benchmark algorithms. The following subsections present summary results, with detailed outcomes provided in Appendix E.

5.1 ICP performance

This subsection presents ICP's performance from multiple perspectives, offering a comprehensive view of its behavior across the benchmark instances in Set A and Set B. For Set A, we used DPS_{25} as the baseline algorithm, while for Set B, we employed HM_{4800} as the baseline algorithm. We select HGA-TAC$^+$ as our primary benchmark due to its established record of high-quality solutions in the literature. For HGA-TAC$^+$, we report the Best, Mean, and Worst values from 10 independent runs. Objective values are presented as relative gaps to the corresponding baseline algorithm.

Overall, ICP achieves an average solution quality improvement of 2.75% compared to mean performance of HGA-TAC$^+$ while reducing computation time by 79.8%. It consistently outperforms existing algorithms on both the parameter variations in Set A and the realistic urban distributions

---

[1]Implementations for these benchmark algorithms are publicly available at `https://github.com/chkwon/TSPDrone.jl` (TSP-EP-ALL, DPS_{25}, HM_{4800}) and `https://github.com/Sasanm88/TSPDroneHGATAC.jl` (HGA-TAC$^+$).

Table 1: A summary of results for ICP for all sets of instances. † is reported by Bogyrbayeva et al. (2023). Times are in seconds.

		Baseline	HGA-TAC$^+$				ICP	
Instance Set	Total	Time	Best	Mean	Worst	Time	Gap	Time
Set A		DPS$_{25}$						
- Uniform	255	3.81	-0.64%	0.49%	1.70%	69.02	-2.50%	10.87
- 1-center	252	3.10	-1.32%	0.07%	1.58%	65.91	-3.39%	10.28
- 2-center	252	3.04	-0.98%	0.32%	1.71%	66.62	-2.87%	9.56
Set B		HM$_{4800}$						
- Random	200	8.96	-1.46%	0.30%	2.11%	7.59	-1.32%	2.60
- Amsterdam	100	1.41†	-1.54%	0.64%	3.07%	4.02	-0.74%	1.18

Table 2: A summary of results for ICP with varying parameters α and N for Set A (Agatz et al., 2018) instances. Times are in seconds.

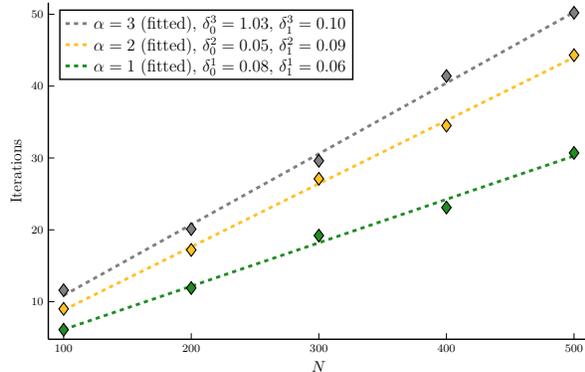|			DPS$_{25}$	HGA-TAC$^+$				ICP	
Set A		Total	Time	Best	Mean	Worst	Time	Gap	Time
	1	253	2.42	-0.51%	-0.11%	0.22%	49.12	-1.02%	3.52
α	2	254	3.37	-0.07%	1.38%	3.02%	67.42	-2.97%	10.19
	3	252	4.17	-2.74%	-0.44%	2.08%	85.10	-5.53%	17.03
	50	90	0.47	-3.51%	-1.78%	0.04%	4.63	-1.96%	1.25
	75	90	0.68	-3.24%	-1.36%	0.66%	8.65	-2.32%	2.37
N	100	90	0.85	-1.82%	-0.13%	1.55%	13.12	-2.56%	3.28
	175	90	2.06	-1.76%	-0.13%	1.52%	31.97	-2.91%	6.59
	250	180	2.73	-0.62%	0.53%	1.78%	60.56	-3.00%	10.32
	375	127	5.73	-0.48%	0.56%	1.71%	122.63	-3.32%	17.29
	500	92	10.17	0.10%	1.18%	2.44%	209.45	-3.25%	27.20

Table 3: Computational profile of ICP across main subroutines, averaged over 10 instances generated under Set A's (Agatz et al., 2018) uniform distribution scheme, for varying α and N. Times are in seconds.

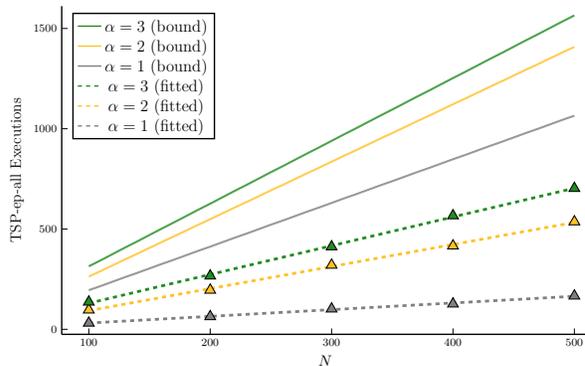|		CONCORDE		TSP-EP				TSP-EP-ALL				
α	N	Time	%	Time	%	$|\mathcal{C}_j|$	$\mathcal{I}$	#	Avg.	Time	%	%
	100	0.08	12.42	0.01	0.82	18.25	6.1	31.2	0.016	0.49	76.59	89.83
	200	0.51	32.60	0.08	5.28	18.18	11.9	63.5	0.015	0.95	60.92	98.80
1	300	2.24	54.10	0.42	10.09	18.10	19.2	101.8	0.014	1.47	35.54	99.74
	400	3.45	52.59	1.33	20.20	18.07	23.1	126.7	0.014	1.77	26.98	99.77
	500	8.39	59.86	3.23	23.07	18.04	30.7	166.4	0.014	2.37	16.88	99.81
	100	0.16	4.30	0.01	0.14	19.37	9.0	96.7	0.037	3.56	95.47	99.91
	200	0.36	4.88	0.08	1.13	19.35	17.2	195.4	0.035	6.85	93.79	99.80
2	300	2.66	18.49	0.42	2.93	19.37	27.1	320.1	0.035	11.25	78.34	99.76
	400	5.49	25.92	1.32	6.25	19.35	34.5	416.0	0.034	14.30	67.55	99.72
	500	8.68	28.59	3.23	10.64	19.37	44.3	536.3	0.034	18.35	60.46	99.69
	100	0.14	2.15	0.01	0.08	19.43	11.6	136.2	0.048	6.58	97.69	99.92
	200	0.63	4.67	0.08	0.61	19.44	20.1	266.3	0.048	12.70	94.57	99.85
3	300	0.82	3.90	0.42	2.00	19.45	29.6	411.8	0.048	19.69	93.87	99.77
	400	3.29	10.54	1.33	4.24	19.45	41.4	566.6	0.047	26.52	84.95	99.73
	500	7.11	16.31	3.24	7.44	19.46	50.2	702.9	0.047	33.11	75.95	99.70

in Set B, requiring significantly less computation time.

Table 1 summarizes the results by distribution type. In Set A, ICP achieves average improvements of 2.50–3.39% over DPS_{25} in objective value. When compared against HGA-TAC$^+$, ICP not only outperforms its average solutions but also surpasses the *best* solutions obtained across multiple runs, with an average improvement of 1.94%. These superior results are achieved while reducing computation time by 84.8%. In Set B, ICP demonstrates the advantages of a well-structured heuristic approach over end-to-end learning, achieving superior results in both solution quality and computational efficiency compared to HM_{4800}. Compared to HGA-TAC$^+$, ICP improves by 1.54% over HGA-TAC$^+$ mean while maintaining competitive performance against its best solutions, with 66.8% less computation time.

Table 2 examines ICP's performance with respect to the truck–drone speed ratio α and the instance size N for Set A. ICP consistently outperforms even the best results from HGA-TAC$^+$ across all α values. The algorithm's effectiveness becomes particularly evident with increasing instance sizes. The gap compared to HGA-TAC$^+$ best widens from 1.55% at $N = 50$ to -3.35% at $N = 500$, while computational efficiency shows remarkable scalability. While HGA-TAC$^+$ exhibits steep runtime growth as N increases, ICP maintains more moderate scaling, with time savings increasing from 73.0% to 87.0%. These results validate the robustness and scalability of ICP. Its deterministic nature, superior solution quality, and efficient computation times make it well-suited for large-scale TSP-D instances.

(a) Number of iterations

(b) Number of TSP-EP-ALL executions with bounds

Figure 7: Linear growth patterns in ICP.

We conclude our analysis of ICP by examining its computational characteristics through Table 3. The table presents a detailed breakdown of runtime across the algorithm's three principal subroutines: CONCORDE, TSP-EP, and TSP-EP-ALL. The data represent averages from ten instances generated under Set A's uniform distribution scheme. The columns detail the runtime components (Time) and the percentage contributions (%) for each subroutine. Additionally, for TSP-EP-ALL, we report the average chainlet size ($|\mathcal{C}_j|$), the number of iterations ($\mathcal{I}$), the total number of executions (#), and the mean execution time per call (*Avg.*). The final column (%) presents the total percentage contributions of these three subroutines, confirming that they account for nearly all of ICP's total runtime.

The results reveal several key patterns. Higher α values correspond to increased average chainlet sizes and more frequent executions of TSP-EP-ALL, as faster drone speeds produce smaller rings and thus generate more chainlets per iteration. Moreover, since TSP-EP-ALL is inherently faster for lower α, the average time per call and its share of total runtime rise significantly at higher α. Finally, as N increases, CONCORDE and TSP-EP exhibit rapidly growing runtimes, reflecting their exponential and $O(N^3)$ complexity.

In Figure 7a, we plot the average number of iterations versus N for each value of α and perform linear regression analysis. The results demonstrate a strong linear relationship between the number of iterations and N. Let $\mathcal{I}^\alpha(N)$ denote a function representing the number of iterations over N for a given α. Then, by Proposition 2, the bound on the total number of TSP-EP-ALL executions is:

$$\frac{2N - 1 + \left((N+1) \bmod 3\right)}{3} + 25(\mathcal{I}^\alpha(N) - 1) \tag{4}$$

Now we substitute $\mathcal{I}^\alpha(N) = \delta_0^\alpha + \delta_1^\alpha N$, which represents the linear regression equation of the number of iterations over N for a given α, into Equation (4). This yields:

$$\left(\frac{2}{3} + 25\delta_1^\alpha\right)N + \frac{\left((N+1) \bmod 3\right) - 1}{3} + 25(\delta_0^\alpha - 1) \tag{5}$$

21

which is linear in N. We visualize this theoretical bound in Figure 7b as solid lines. Notably, not only the theoretical bound but also the actual number of TSP-EP-ALL executions, shown as dashed lines, exhibits linear growth in N. Consequently, as N increases, the asymptotic complexity of ICP is dominated by the initial chain generation using CONCORDE and TSP-EP. Despite this favorable scaling, TSP-EP-ALL remains the dominant component of total runtime for large instances, particularly at higher α values, underscoring the potential benefits of replacing TSP-EP-ALL with efficient neural predictions.

5.2 Neural Network Training

Figure 8: Distribution of nodes within chainlets from ICP algorithm runs.

Figure 9: Training progress of the neural network model.

Creating training data that accurately reflects the structure of the chainlets in ICP is challenging. Chainlets are segments partitioned from a chain during the optimization process, which typically causes the nodes within each chainlet to be close to one another. Therefore, we generated training data by collecting samples during actual ICP algorithm runs, ensuring the data reflects the distributional characteristics that arise in practice.

In this context, individual instances were generated with N randomly chosen from $[50, 500]$ matching the uniform distribution scheme of Set A. The ICP algorithm was then run on these instances to collect training data, including the initial TSP tour, the truck and drone costs for each edge, and the corresponding TSP-EP-ALL cost. 384,000 data points were generated, with 128,000 for each $\alpha \in \{1, 2, 3\}$, and split 80:20 for training and validation. The neural network model was trained for 100 epochs to minimize the mean squared error (MSE) loss between the predicted and normalized TSP-EP-ALL costs:

$$L(\theta) = (f_\theta(\mathcal{G}_j) - \text{TSP-EP-ALL}(\mathcal{C}_j))^2.$$

The trained neural network models were evaluated using an independently generated test set of 38,400 data points. To assess the model's out-of-distribution capabilities, test sets were generated for the uniform distribution and additional sets for the 1-center, 2-center, and Amsterdam distributions.

22

Figure 10: Scatter plot showing predicted versus actual TSP-ep-all costs for test set with uniform distribution.

α	Uniform	1-Center	2-Center	Ams.
1	3.33	4.64	4.96	4.24
2	2.95	3.34	3.65	3.74
3	3.75	5.01	5.40	4.80
Avg.	3.34	4.33	4.67	4.26

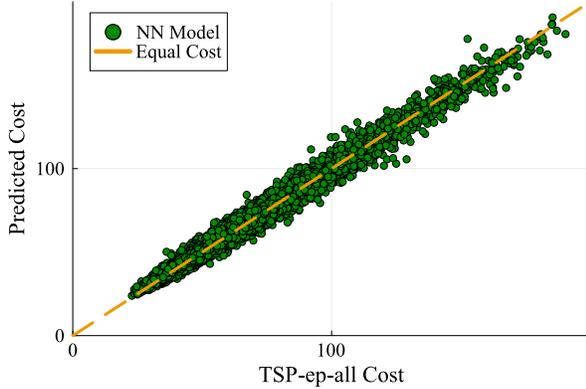Table 4: MAPE (%) for test sets with different distributions.

Test data for the uniform distribution was generated following the same methodology as in the training phase. For the 1-center and 2-center distributions, while the coordinate generation adhered to their respective schemes as defined in Set A, the subsequent processes mirrored those of the training phase. For the Amsterdam distribution, coordinates were utilized from the dataset generated by Bogyrbayeva et al. (2023), derived from the probability density functions of customer node coordinates estimated using Gaussian kernel density estimation (KDE). To generate test data, the process of randomly selecting 100 coordinates to create individual instances and executing the ICP algorithm on these instances was repeated until the desired data size was achieved. For all distributions, 12,800 data points were generated for each value of $\alpha \in \{1, 2, 3\}$, resulting in 38,400 test data points.

The scatter plot in Figure 10 compares predicted costs against actual TSP-ep-all costs for the uniform distribution test set. The x-axis represents the actual TSP-ep-all cost, while the y-axis shows the predicted cost. The dashed orange line represents equal predicted and actual costs, serving as a reference for perfect prediction. The close alignment of points along this line demonstrates the model's high accuracy in predicting TSP-ep-all costs for uniform instances.

Table 4 shows the Mean Absolute Percentage Error (MAPE) across different distributions and α values. The model exhibits robust performance for the uniform distribution with an average MAPE of 3.34%. Notably, it demonstrates solid out-of-distribution performance, yielding average MAPEs of 4.33%, 4.67%, and 4.26% for the 1-center, 2-center, and Amsterdam distributions, respectively. The slight increase in MAPE for non-uniform distributions indicates a minor impact of distributional shifts on prediction accuracy. Additionally, examining performance across α values reveals consistent results, with average MAPEs of 4.29%, 3.42%, and 4.74% for α values of 1, 2, and 3, respectively, suggesting that the model effectively captures the relationship between the relative speeds of trucks and drones and the TSP-ep-all costs. The overall low MAPE values, with an average of 4.15%, demonstrate the model's effectiveness in diverse scenarios, ensuring reliable cost predictions for the ICP algorithm across varying distributions and truck-drone speed configurations. Detailed results of

23

Table 5: A summary of results for ICP and NICP for all sets of instances. † is reported by Bogyrbayeva et al. (2023). Times are in seconds.

Instance Set	Total	Baseline Time	ICP Gap	ICP Time	NICP Gap	NICP Time	Comparison (%) Obj	Comparison (%) Time
Set A		DPS$_{/25}$						
-Uniform	255	3.81	-2.50%	10.87	-2.43%	6.48	0.07%	-40.33%
-1-center	252	3.10	-3.39%	10.28	-3.22%	6.00	0.17%	-41.68%
-2-center	252	3.04	-2.87%	9.56	-2.71%	5.42	0.17%	-43.28%
Set B		HM (4800)						
-Random	200	8.96	-1.32%	2.60	-1.26%	1.22	0.07%	-52.97%
-Amsterdam	100	1.41†	-0.74%	1.18	-0.46%	0.55	0.29%	-53.33%

the neural predictor's performance on these out-of-distribution datasets are provided in Appendix D.

5.3 Comparison of ICP and NICP

In this subsection, we compare the performance of NICP with ICP, using the same baselines as before: DPS$_{25}$ for Set A and HM$_{4800}$ for Set B. As in prior subsections, objective values are reported as gaps relative to these baselines. NICP demonstrates the strength of neural acceleration in improving efficiency without compromising performance. Overall, NICP reduces computation time by 49.7% while increasing objective values by only 0.12% compared to ICP. NICP achieves this acceleration by replacing direct execution of TSP-EP-ALL with efficient neural predictions leveraging GPU parallelization while maintaining ICP's core structure and characteristics.

Table 5 presents NICP's performance across different distribution types compared to ICP. The comparison columns quantify percentage differences between NICP and ICP, with positive values indicating NICP's increase over ICP. The results demonstrate that NICP successfully maintains solution quality while substantially reducing computation time. Notably, NICP's performance remains robust across various spatial distributions, preserving ICP's advantage in handling clustered and realistic scenarios while significantly improving computational efficiency. Across all distributions, the degradation in objective values remains minimal, ranging from 0.07% to 0.29%. NICP achieves consistent computational savings ranging from 40.33% to 53.33%.

Table 6 examines NICP's performance across varying values of α and N. The slight increase in objective gap with α from 0.06% to 0.36% suggests a minor trade-off between computation time and solution quality at higher drone speeds. The impact of neural acceleration becomes more pronounced as α increases, with time reductions growing from 9.84% at $\alpha = 1$ to 47.58% at $\alpha = 3$. This pattern aligns with our earlier analysis where TSP-EP-ALL's computational burden increases significantly with α, making neural prediction particularly effective for higher drone speeds. Regarding instance size, NICP maintains stable solution quality across all values of N, with objective gaps ranging from 0.03% to 0.27%. The computational advantage decreases slightly as N grows, with time reductions

Table 6: A summary of results for ICP and NICP with varying parameters α and N for Set A (Agatz et al., 2018) instances. Times are in seconds.

			DPS$_{/25}$	ICP		NICP		Comparison (%)	
Set A		Total	Time	Gap	Time	Gap	Time	Obj	Time
	1	253	2.42	-1.02%	3.52	-0.96%	3.17	0.06%	-9.84%
α	2	254	3.37	-2.97%	10.19	-2.89%	5.82	0.08%	-42.89%
	3	252	4.17	-5.53%	17.03	-5.19%	8.93	0.36%	-47.58%
	50	90	0.47	-1.96%	1.25	-1.91%	0.61	0.05%	-51.06%
	75	90	0.68	-2.32%	2.37	-2.30%	1.10	0.03%	-53.64%
	100	90	0.85	-2.56%	3.28	-2.40%	1.53	0.16%	-53.40%
N	175	90	2.06	-2.91%	6.59	-2.65%	3.41	0.27%	-48.27%
	250	180	2.73	-3.00%	10.32	-2.81%	5.64	0.20%	-45.39%
	375	127	5.73	-3.32%	17.29	-3.25%	10.25	0.08%	-40.72%
	500	92	10.17	-3.25%	27.20	-3.08%	17.56	0.18%	-35.43%

(a) Objective Value over Time (b) Objective Value over Iteration

Figure 11: Comparison of average ICP and NICP results for 10 randomly generated uniform instances with $N = 100$.

declining from 53.64% at $N = 75$ to 35.43% at $N = 500$. This trend reflects the diminishing share of TSP-EP-ALL in total runtime as N increases, where the polynomial complexity of TSP-EP and exponential nature of CONCORDE become more dominant factors. Nevertheless, NICP consistently achieves substantial time savings while maintaining solution quality across all instance sizes.

We conclude by illustrating the characteristics of neural acceleration through Figure 11, which compares ICP and NICP across ten uniformly generated instances with $N = 100$. The left figure presents the gap between ICP's final objective value over computation time, where the shaded regions represent standard deviations. NICP exhibits faster convergence, achieving greater reductions within the same time frame while closely tracking ICP's solution quality. Complementing this time-based analysis, the right figure shows the gap over iterations, demonstrating that NICP follows ICP's solution trajectory per iteration despite its faster execution. Together, these results validate that

NICP achieves actual neural acceleration - maintaining ICP's solution patterns while significantly reducing computation time through efficient neural prediction.

6   Conclusion

In this work, we developed the Iterative Chainlet Partitioning (ICP) algorithm for solving the TSP-D and its neural acceleration framework (NICP). ICP proposes a framework for tackling large-scale TSP-D instances. It decomposes a TSP-D solution into groups of consecutive rings called chainlets and optimizes them using the precise TSP-EP-ALL subroutine. By greedily selecting the chainlet with the highest improvement and caching previous optimization results, ICP achieves both deterministic behavior and theoretical tractability. Its deterministic nature ensures reliability, eliminating the need for repeated runs to achieve high-quality solutions. The number of subroutine calls is bounded linearly in N for the first iteration and remains constant, ensuring computational efficiency. Extensive computational experiments showed that ICP outperforms the existing state-of-the-art algorithm, improving solution quality by 2.75% while reducing computation time by 79.84%. For further acceleration, we integrated a GNN to predict improvement, allowing ICP to prioritize promising chainlets without exhaustive evaluation. Our targeted application of neural prediction preserves the algorithmic guarantees while reducing computational time by 49.7%. The framework's adaptability to various operational constraints makes it a valuable foundation for developing efficient algorithms for truck-drone synchronized routing problems.

Future work will focus on improving methods for initial chain generation, which currently dominates computational complexity for larger instances. Additionally, extending neural prediction capabilities to handle various operational constraints such as customer eligibility and drone range limitations offers promising avenues for expansion.

Acknowledgments

References

Agatz, N., Bouman, P., and Schmidt, M. (2018). Optimization approaches for the traveling salesman problem with drone. *Transportation Science*, 52(4):965–981.

Alvim, A. C. and Taillard, É. D. (2013). POPMUSIC for the world location-routing problem. *EURO Journal on Transportation and Logistics*, 2(3):231–254.

Applegate, D., Bixby, R., Chvátal, V., and Cook, W. (1998). On the solution of traveling salesman problems. *Documenta Mathematica*, pages 645–656.

Bello, I., Pham, H., Le, Q. V., Norouzi, M., and Bengio, S. (2016). Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*.

Blufstein, M., Lera-Romero, G., and Soulignac, F. J. (2024). Decremental state-space relaxations for the basic traveling salesman problem with a drone. *INFORMS Journal on Computing*.

Boccia, M., Masone, A., Sforza, A., and Sterle, C. (2021). A column-and-row generation approach for the flying sidekick travelling salesman problem. *Transportation Research Part C: Emerging Technologies*, 124:102913.

Bogyrbayeva, A., Yoon, T., Ko, H., Lim, S., Yun, H., and Kwon, C. (2023). A deep reinforcement learning approach for solving the traveling salesman problem with drone. *Transportation Research Part C: Emerging Technologies*, 148:103981.

Bouman, P., Agatz, N., and Schmidt, M. (2018). Dynamic programming approaches for the traveling salesman problem with drone. *Networks*, 72(4):528–542.

Cai, T., Luo, S., Xu, K., He, D., Liu, T.-y., and Wang, L. (2021). Graphnorm: A principled approach to accelerating graph neural network training. In *International Conference on Machine Learning*, pages 1204–1215. PMLR.

Chung, S. H., Sah, B., and Lee, J. (2020). Optimization for drone and drone-truck combined operations: A review of the state of the art and future directions. *Computers & Operations Research*, 123:105004.

Clevert, D.-A. (2015). Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*.

de Freitas, J. C. and Penna, P. H. V. (2020). A variable neighborhood search for flying sidekick traveling salesman problem. *International Transactions in Operational Research*, 27(1):267–290.

Haider, Z., Charkhgard, H., Kim, S. W., and Kwon, C. (2019). Optimizing the relocation operations of free-floating electric vehicle sharing systems. *Available at SSRN 3480725*.

Helsgaun, K. (2017). An extension of the Lin-Kernighan-Helsgaun TSP solver for constrained traveling salesman and vehicle routing problems. *Roskilde: Roskilde University*, 12:966–980.

Hottung, A. and Tierney, K. (2020). Neural large neighborhood search for the capacitated vehicle routing problem. In *ECAI 2020*, pages 443–450. IOS Press.

Khalil, E., Dai, H., Zhang, Y., Dilkina, B., and Song, L. (2017). Learning combinatorial optimization algorithms over graphs. *Advances in Neural Information Processing Systems*, 30.

Kim, H., Park, J., and Kwon, C. (2024). A neural separation algorithm for the rounded capacity inequalities. *INFORMS Journal on Computing*, 36(4):987–1005.

Kool, W., Van Hoof, H., and Welling, M. (2018). Attention, learn to solve routing problems! *arXiv preprint arXiv:1803.08475*.

Li, S., Yan, Z., and Wu, C. (2021). Learning to delegate for large-scale vehicle routing. *Advances in Neural Information Processing Systems*, 34:26198–26211.

Liang, Y.-J. and Luo, Z.-X. (2022). A survey of truck–drone routing problem: Literature review and research prospects. *Journal of the Operations Research Society of China*, 10(2):343–377.

Macrina, G., Pugliese, L. D. P., Guerriero, F., and Laporte, G. (2020). Drone-aided routing: A literature review. *Transportation Research Part C: Emerging Technologies*, 120:102762.

Mahmoudinazlou, S. and Kwon, C. (2024). A hybrid genetic algorithm with type-aware chromosomes for traveling salesman problems with drone. *European Journal of Operational Research*.

Murray, C. C. and Chu, A. G. (2015). The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. *Transportation Research Part C: Emerging Technologies*, 54:86–109.

Nazari, M., Oroojlooy, A., Snyder, L., and Takác, M. (2018). Reinforcement learning for solving the vehicle routing problem. *Advances in Neural Information Processing Systems*, 31.

Ostertag, A., Doerner, K. F., Hartl, R. F., Taillard, E., and Waelti, P. (2009). POPMUSIC for a real-world large-scale vehicle routing problem with time windows. *Journal of the Operational Research Society*, 60(7):934–943.

Poikonen, S., Golden, B., and Wasil, E. A. (2019). A branch-and-bound approach to the traveling salesman problem with a drone. *INFORMS Journal on Computing*, 31(2):335–346.

Ribeiro, C. C., Hansen, P., Taillard, E. D., and Voss, S. (2002). POPMUSIC—partial optimization metaheuristic under special intensification conditions. *Essays and surveys in metaheuristics*, pages 613–629.

Roberti, R. and Ruthmair, M. (2021). Exact methods for the traveling salesman problem with drone. *Transportation Science*, 55(2):315–335.

Shi, Y., Huang, Z., Feng, S., Zhong, H., Wang, W., and Sun, Y. (2020). Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509*.

Sobhanan, A., Park, J., Park, J., and Kwon, C. (2025). Genetic algorithms with neural cost predictor for solving hierarchical vehicle routing problems. *Transportation Science*, Article in Advance.

Sui, J., Ding, S., Huang, X., Yu, Y., Liu, R., Xia, B., Ding, Z., Xu, L., Zhang, H., Yu, C., et al. (2025). A survey on deep learning-based algorithms for the traveling salesman problem. *Frontiers of Computer Science*, 19(6):1–30.

Taillard, É. D. and Helsgaun, K. (2019). POPMUSIC for the travelling salesman problem. *European Journal of Operational Research*, 272(2):420–429.

Taillard, É. D. and Voss, S. (2018). Popmusic. *Handbook of heuristics.*

Van Dijck, E., Bouman, P., and Spliet, R. (2018). *A branch-and-cut algorithm for the traveling salesman problem with drone.* PhD thesis, Master's thesis, Erasmus University Rotterdam.

Vásquez, S. A., Angulo, G., and Klapp, M. A. (2021). An exact solution method for the tsp with drone based on decomposition. *Computers & Operations Research*, 127:105127.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

Xin, L., Song, W., Cao, Z., and Zhang, J. (2021). Neurolkh: Combining deep learning model with lin-kernighan-helsgaun heuristic for solving the traveling salesman problem. *Advances in Neural Information Processing Systems*, 34:7472–7483.

Ye, H., Wang, J., Cao, Z., Berto, F., Hua, C., Kim, H., Park, J., and Song, G. (2024). Reevo: Large language models as hyper-heuristics with reflective evolution. *arXiv preprint arXiv:2402.01145.*

Appendices

A Proofs of Propositions and Lemmas

Proof of Proposition 1. The ICP algorithm ensures that in each iteration, the chainlet with the highest positive improvement is selected and updated. This guarantees that the cost of the solution monotonically decreases over iterations, as no update can increase the cost. Moreover, cycling cannot occur in ICP. Once a chainlet is updated, each modified ring within the chainlet is an exact partitioning solution for the current tour sequence. For any chainlet containing these modified rings in subsequent iterations, the result of applying TSP-EP-ALL will remain unchanged unless the tour sequence itself is modified. Thus, rings modified during previous updates cannot appear in their original configuration in subsequent iterations. As a result, the algorithm progresses without revisiting prior configurations. Since the number of possible configurations of chainlets is finite, and each iteration guarantees a strict cost improvement, the ICP algorithm must terminate in a finite number of iterations.

Additionally, ICP is a deterministic procedure. This follows from the deterministic nature of subroutines: the CONCORDE solver is deterministic given that the TSP instance has a unique optimal solution, and both TSP-EP and TSP-EP-ALL are deterministic. Therefore, with the consistent greedy selection of chainlets, the sequence of chainlet updates depends only on the input instance and the deterministic subroutines, ensuring the final solution is uniquely determined. $\square$

Figure 12: Two Consecutive *Straight Rings* and an alternative *Triangular Ring*

Proof of Lemma 1. Consider two consecutive straight rings with arc costs a and b connecting three consecutive combined nodes, illustrated in Figure 12. The corresponding triangular ring can be formed where a drone serves the middle node, and the truck directly connects the first and third combined nodes. Let c denote the cost of the arc directly connecting these two combined nodes in the triangular ring. The total cost of the two consecutive straight rings is $a + b$, while the total cost of the triangular ring is $\max\left\{\frac{a+b}{\alpha}, c\right\}$. If $\frac{a+b}{\alpha} \geq c$, the cost of the triangular ring is $\frac{a+b}{\alpha}$, and since $\alpha > 1$, it follows that $\frac{a+b}{\alpha} < a + b$. If $\frac{a+b}{\alpha} < c$, the cost of the triangular ring is c, and by the triangle inequality, $c < a + b$. Thus, the cost of the triangular ring is strictly less than the cost of the two consecutive straight rings when $\alpha > 1$. $\square$

Proof of Lemma 2. The maximum number of rings in a chainlet occurs when there is a minimal number of nodes in each ring, which is achieved by maximizing the alternation between a straight

ring and a triangular ring. Three cases are considered:

(i) If $n \equiv 1 \pmod 3$, the maximum number of rings results from alternating between straight and triangular rings. Starting from a start node, a straight ring followed by a triangular ring yields 2 rings for every 3 nodes, so the total number of rings is $\frac{2(n-1)}{3}$. An equivalent number of rings is obtained by initiating the sequence with a triangular ring.

(ii) For a chainlet with one additional node than case (i), *i.e.*, $n \equiv 2 \pmod 3$, adding a straight ring maximizes the number of rings, giving $\frac{2(n-2)}{3} + 1$.

(iii) For a chainlet with one fewer node than case (i), *i.e.*, $n \equiv 0 \pmod 3$, attempting to change a triangular ring into a straight ring would always result in two consecutive straight rings. Thus, the only choice is to delete a straight ring, yielding $\frac{2n}{3} - 1$.

Altogether, the maximum number of rings in a chainlet can be universally stated as $\frac{2n - 3 + (n \bmod 3)}{3}$.

$\square$

Proof of Proposition 2. The maximum number of rings in a chain sets an upper bound on the number of chainlets generated, as each chainlet created by GROUP must contain at least one ring that has not been grouped into any chainlet. Unlike the count established in Lemma 2 for chainlets, a chain includes a dummy node representing the depot, which increases the overall node count by one. Therefore, plugging in $N + 1$ into the formula from Lemma 2, we obtain

$$\frac{2(N + 1) - 3 + ((N + 1) \bmod 3)}{3}$$

or (2) as the maximum number of rings in the chain. This number also bounds the maximum number of chainlets in the first iteration. It is also the maximum number of TSP-EP-ALL runs required in the first iteration where every chainlet is initially new.

In subsequent iterations, TSP-EP-ALL is executed only on chainlets containing rings modified by the updated chainlet. Consider an updated chainlet $\mathcal{C}_k$ with r rings. Define the prefix set as the collection of all consecutive subsequences starting from the first ring, and the suffix set as the collection of all consecutive subsequences ending with the last ring. Each set contains exactly r sequences, with the full sequence appearing in both sets. The union of the chainlet's prefix and suffix sets contains at most $2r - 1$ unique sequences. The maximum number of new chainlets arises when $\mathcal{C}_k$ contains the maximum number of rings and each element in this union forms a distinct chainlet.

By Lemma 2, the maximum number of rings in the updated chainlet is $\frac{2\ell - (3 - \ell \bmod 3)}{3}$. Consequently, the maximum number of new chainlets is $2\left(\frac{2\ell - (3 - \ell \bmod 3)}{3}\right) - 1$ or (3), which is also the maximum number of TSP-EP-ALL runs required in subsequent iterations. $\square$

B   Selection of ICP Configurations

Figure 13: Linear regression analysis of TSP versus TSP-D costs across different α values.

Method	TSP	TSP-D_1	TSP-D_2	TSP-D_3
Concorde	769.89	633.84	524.22	478.77
FI	828.44	657.75	535.60	486.99
NN	954.80	731.43	608.23	565.05
CI	917.94	688.24	559.63	511.17
Random	5125.81	1280.58	939.69	830.64

Table 7: Average costs for initial TSP construction methods. Subscript denotes α values.

This section presents an empirical analysis for determining optimal configurations of the ICP algorithm. We first investigate various approaches for constructing the initial TSP tour for chain generation, presented in Figure 13 and Table 7. The analysis compares five methods: Concorde, an exact TSP solver that guarantees optimal solutions; Farthest Insertion (FI), a construction heuristic that iteratively inserts the node farthest from the current tour; Nearest Neighbor (NN), a greedy construction heuristic that sequentially adds the closest unvisited node; Closest Insertion (CI), a construction heuristic that iteratively inserts nodes with minimal insertion cost; and Random, which generates a random permutation of nodes.

For this analysis, we generated 10 uniform instances with 100 nodes and executed ICP with each TSP construction method for chain initialization. The results, averaged across all instances, demonstrate that the quality of the initial TSP tour significantly influences the final TSP-D solution quality. Concorde consistently produces superior results across all values of α, achieving improvements of 3.6–7.7% over FI, the second-best performing method. As shown in Figure 13, where Random method is excluded from the regression analysis, TSP and TSP-D costs demonstrate an apparent linear relationship. However, this correlation diminishes as α increases, evidenced by the decreasing R^2 values from 0.743 at $\alpha = 1$ to 0.504 at $\alpha = 3$. Based on these results, we selected Concorde for the initial chain construction in ICP despite its higher computational overhead.

We then analyze methods for constructing chainlet input tours, with results presented in Table 8. For this analysis, we used LKH-3 (Helsgaun, 2017) instead of Concorde, as LKH-3 empirically achieves optimal solutions for instances of chainlet size while requiring less computational time. The other methods (FI, NN, CI, and Random) remain identical to those used in the initial chain construction analysis but are slightly modified to avoid altering the chainlet's end node.

Table 8: Performance comparison of chainlet input tour construction methods with varying maximum node sizes. Results are averaged over 10 uniform instances with 100 nodes. Times are in seconds.

α	Size	LKH-3 Obj	LKH-3 Time	FI Obj	FI Time	NN Obj	NN Time	CI Obj	CI Time	Random Obj	Random Time
1	18	680.76	0.96	678.44	0.52	680.39	0.54	680.45	0.44	681.41	1.16
	19	677.86	0.56	676.08	0.57	678.43	0.73	677.03	0.72	679.87	1.64
	20	677.20	0.70	676.60	0.61	676.79	0.96	677.21	0.77	675.71	2.49
	21	675.32	0.92	675.19	0.95	676.42	1.04	676.51	0.92	676.27	3.31
	22	674.76	1.12	674.77	1.16	675.83	1.46	675.71	1.20	677.49	4.09
	23	673.57	1.58	673.72	1.35	674.21	2.30	677.06	1.71	672.84	6.52
2	18	560.62	2.50	562.37	2.17	561.61	2.68	562.13	2.08	560.57	6.48
	19	561.83	2.91	561.93	2.61	560.29	3.53	559.49	3.04	555.93	8.25
	20	560.31	3.91	559.54	3.60	560.25	4.57	562.71	3.81	555.75	11.80
	21	557.59	5.78	560.59	4.82	558.05	6.78	558.96	5.31	555.61	16.14
	22	557.66	6.65	557.02	6.70	557.38	8.01	558.47	7.11	554.11	19.81
	23	553.85	9.59	558.55	8.16	558.63	9.74	557.39	9.91	554.50	27.18
3	18	519.30	4.22	518.73	3.51	517.57	4.76	513.36	4.29	511.81	8.54
	19	515.54	5.87	515.56	5.15	515.95	5.91	515.83	5.60	506.45	12.58
	20	515.71	7.47	512.53	7.04	518.18	7.58	517.18	7.04	504.16	17.59
	21	516.96	8.99	513.36	8.99	518.18	9.90	515.03	9.43	503.07	23.24
	22	516.30	11.71	512.79	12.09	511.50	15.00	512.90	13.76	501.83	29.06
	23	516.36	14.79	513.43	15.45	513.12	18.92	509.41	19.98	502.03	40.96

For the experiment, we again generated 10 uniform instances with 100 nodes. The experimental results demonstrate that FI consistently produces competitive solutions while maintaining efficient computation times across all tested configurations. While the Random method occasionally discovers good solutions, it exhibits significantly longer computation times and introduces undesirable non-deterministic behavior. Regarding the maximum node size for chainlets, our analysis indicates that size 20 provides an optimal balance between solution quality and computational efficiency. While larger sizes occasionally yield marginally better solutions, they incur substantially increased computation times due to the exponential complexity of TSP-EP-ALL. Conversely, though computationally faster, smaller sizes often result in inferior solutions as they constrain the optimization capabilities of TSP-EP-ALL. Based on these comprehensive results, we configured ICP to employ FI for input tour construction with a maximum node size of 19, ensuring consistent, high-quality solutions while maintaining computational efficiency.

C Pseudo Codes for ICP Modules

This appendix provides the pseudocode for two auxiliary procedures used in ICP: (i) grouping rings into chainlets (GROUP), and (ii) the Farthest Insertion (FI) heuristic with fixed end nodes (FARTHESTINSERTION).

Algorithm 1-1 Group rings into chainlets (GROUP)

 1: **Input:** Chain $\mathcal{C}$ with r rings
 2: **Output:** Sequence of chainlets
 3: $\mathcal{S} \leftarrow \emptyset$ ▷ Empty sequence of chainlets
 4: $i \leftarrow 1$ ▷ Starting position of the chainlet
 5: **repeat**
 6:     $\mathcal{C}_{\text{curr}} \leftarrow \emptyset$ ▷ Empty sequence of rings
 7:     $j \leftarrow 0$ ▷ Number of rings in the chainlet
 8:     **while** $|\mathcal{C}_{\text{curr}}| \leq 19$ and $i + j < r$ **do**
 9:         $\mathcal{C}_{\text{curr}}.\texttt{append}(\mathcal{C}_{\text{curr}})$ ▷ Add next ring
10:         $j \leftarrow j + 1$
11:     **if** $j > 0$ and not $\mathcal{C}_{\text{curr}} \subset \mathcal{S}[i-1]$ **then**
12:         $\mathcal{S}.\texttt{append}(\mathcal{C}_{\text{curr}})$ ▷ Add new chainlet
13:     $i \leftarrow i + 1$
14: **until** $i + j = r$
15: **return** $\mathcal{S}$

Algorithm 1-2 Farthest Insertion with fixed endpoints (FARTHESTINSERTION)

 1: **Input:** chainlet $\mathcal{C}_j$
 2: **Output:** tour $\mathcal{T}$
 3: $\mathcal{N} \leftarrow \bigcup_{i:\mathcal{R}_i \in \mathcal{C}_j} R_i$ ▷ Identify node set
 4: $s \leftarrow \mathcal{C}_j[1][start], t \leftarrow \mathcal{C}_j[end][end]$
 5: $\mathcal{T} \leftarrow [s, t]$ ▷ Initialize tour with endpoints
 6: $\mathcal{U} \leftarrow \mathcal{N} \setminus \{s, t\}$ ▷ Set of unvisited nodes
 7: **while** $\mathcal{U} \neq \emptyset$ **do**
 8:     $v^* \leftarrow \arg\max_{v \in \mathcal{U}} \min_{u \in \mathcal{T}} \texttt{cost}(v, u)$ ▷ Find farthest unvisited node
 9:     $(i^*, \text{cost}^*) \leftarrow (0, \infty)$
10:     **for** $i \in \{1, \ldots, |\mathcal{T}| - 1\}$ **do**
11:         $\text{cost} \leftarrow \texttt{cost}(\mathcal{T}_i, v^*) + \texttt{cost}(v^*, \mathcal{T}_{i+1}) - \texttt{cost}(\mathcal{T}_i, \mathcal{T}_{i+1})$
12:         **if** $\text{cost} < \text{cost}^*$ **then**
13:            $(i^*, \text{cost}^*) \leftarrow (i, \text{cost})$
14:     Insert v^* into $\mathcal{T}$ after position i^*
15:     $\mathcal{U} \leftarrow \mathcal{U} \setminus \{v^*\}$
16: **return** $\mathcal{T}$

D Results of Out-of-distribution Data

(a) Scatter plot for 1-center

(b) Error analysis for 1-center

(c) Scatter plot for 2-center

(d) Error analysis for 2-center

(e) Scatter plot for Amsterdam

(f) Error analysis for Amsterdam

Figure 14: Neural cost predictor performance on out-of-distribution data

Figures 14a–14f illustrate the neural cost predictor's performance on out-of-distribution data. The scatter plots compare predicted versus actual TSP-ep-all costs, while the corresponding figures show prediction error percentages relative to actual costs. Complementing Table 4, these visualizations confirm the neural model's robustness across diverse spatial distributions not represented in the training data.

E   Detailed Results of Computational Experiments

Table 9: (Set A, uniform) Results of ICP and NICP for Agatz et al. (2018) uniform instances. Comparison (%) represents the relative difference of NICP compared to ICP. Times are in seconds.

	N	DPS$_{25}$ Obj	Time	HGA-TAC$^+$ Best	Mean	Worst	Time	ICP Obj	Time	Gap Over (%) DPS	HGA	NICP Obj	Time	Gap Over (%) DPS	HGA	Comparison (%) Obj	Time
	50	498.07	0.25	493.89	496.86	500.10	4.52	494.46	0.31	-0.73	-0.48	494.84	0.22	-0.65	-0.41	0.08	-29.10
	75	569.16	0.32	571.55	574.71	577.13	8.12	573.29	0.40	0.73	-0.25	573.82	0.30	0.82	-0.16	0.09	-24.43
	100	646.48	0.44	650.58	655.64	658.95	13.08	649.27	0.56	0.43	-0.97	649.18	0.43	0.42	-0.99	-0.01	-22.40
$\alpha = 1$	175	837.70	1.62	837.84	840.99	843.45	24.14	830.40	2.14	-0.87	-1.26	832.10	1.69	-0.67	-1.06	0.20	-20.76
	250	993.68	2.11	993.01	997.25	1000.40	49.48	984.93	2.34	-0.88	-1.24	985.40	2.08	-0.83	-1.19	0.05	-10.85
	375	1191.77	5.76	1194.45	1197.97	1201.18	83.83	1181.00	7.08	-0.90	-1.42	1181.22	5.71	-0.89	-1.40	0.02	-19.34
	500	1369.08	9.91	1373.84	1377.19	1380.63	147.90	1355.57	14.09	-0.99	-1.57	1355.62	12.26	-0.98	-1.57	0.00	-12.98
	Avg.	872.28	2.91	873.59	877.23	880.26	47.30	866.99	3.84	-0.61	-1.17	867.45	3.24	-0.55	-1.11	0.05	-15.64
	50	417.55	0.52	405.51	411.32	418.69	3.84	409.26	1.60	-1.99	-0.50	408.23	0.84	-2.23	-0.75	-0.25	-47.64
	75	481.55	0.72	469.76	478.32	488.55	7.65	474.19	2.56	-1.53	-0.86	474.53	1.28	-1.46	-0.79	0.07	-50.17
	100	552.96	0.89	547.03	557.58	567.52	10.58	539.22	3.66	-2.49	-3.29	540.99	1.53	-2.16	-2.98	0.33	-58.29
$\alpha = 2$	175	708.88	2.51	704.67	713.47	724.89	29.63	693.79	6.91	-2.13	-2.76	692.52	3.96	-2.31	-2.94	-0.18	-42.67
	250	839.82	2.96	841.12	851.06	861.29	56.93	815.75	10.53	-2.87	-4.15	817.38	5.68	-2.67	-3.96	0.20	-46.05
	375	1013.79	6.39	1020.97	1031.93	1045.06	132.25	985.78	18.39	-2.76	-4.47	984.71	10.93	-2.87	-4.58	-0.11	-40.58
	500	1163.42	12.02	1172.79	1183.38	1196.83	217.18	1124.62	29.50	-3.33	-4.97	1125.18	19.42	-3.29	-4.92	0.05	-34.17
	Avg.	739.71	3.72	737.41	746.72	757.55	65.44	720.37	10.45	-2.61	-3.53	720.51	6.23	-2.60	-3.51	0.02	-40.35
	50	391.64	0.66	368.30	377.67	387.06	3.88	369.52	2.90	-5.65	-2.16	368.65	1.18	-5.87	-2.39	-0.23	-59.26
	75	451.68	0.82	428.26	441.22	454.44	7.25	430.14	5.22	-4.77	-2.51	432.25	2.01	-4.30	-2.03	0.49	-61.51
	100	521.06	1.14	502.36	514.75	529.76	12.55	497.57	6.35	-4.51	-3.34	499.29	2.72	-4.18	-3.00	0.35	-57.17
$\alpha = 3$	175	666.64	3.71	643.46	656.85	673.25	37.53	638.12	12.09	-4.28	-2.85	640.12	6.57	-3.98	-2.55	0.31	-45.63
	250	792.11	3.62	776.54	790.88	808.40	72.61	756.38	18.09	-4.51	-4.36	757.17	8.83	-4.41	-4.26	0.10	-51.22
	375	951.19	7.90	947.01	962.95	977.51	169.22	908.52	29.14	-4.49	-5.65	908.73	17.11	-4.46	-5.63	0.02	-41.29
	500	1093.30	13.44	1088.59	1107.08	1125.27	284.97	1038.23	44.22	-5.04	-6.22	1040.29	26.23	-4.85	-6.03	0.20	-40.68
	Avg.	695.38	4.47	679.22	693.06	707.96	84.00	662.64	16.86	-4.71	-4.39	663.79	9.24	-4.54	-4.22	0.17	-45.22

Table 10: (Set A, 1-center) Results of ICP and NICP for Agatz et al. (2018) 1-center instances. Comparison (%) represents the relative difference of NICP compared to ICP. Times are in seconds.

	N	DPS$_{25}$		HGA-TAC$^+$				ICP		Gap Over (%)		NICP		Gap Over (%)		Comparison (%)	
		Obj	Time	Best	Mean	Worst	Time	Obj	Time	DPS	HGA	Obj	Time	DPS	HGA	Obj	Time
$\alpha=1$	50	657.31	0.23	646.59	651.98	656.09	6.76	654.15	0.13	-0.48	0.33	654.71	0.09	-0.40	0.42	0.09	-30.19
	75	888.22	0.43	876.09	880.90	884.57	9.74	880.45	0.40	-0.88	-0.05	880.52	0.38	-0.87	-0.04	0.01	-6.37
	100	1058.68	0.45	1052.06	1058.34	1062.76	14.98	1059.42	0.41	0.07	0.10	1058.14	0.33	-0.05	-0.02	-0.12	-18.79
	175	1421.35	1.01	1410.81	1418.24	1424.51	27.13	1406.90	1.05	-1.02	-0.80	1407.08	0.90	-1.00	-0.79	0.01	-14.19
	250	1656.60	1.65	1645.01	1650.35	1655.24	51.58	1635.99	7.15	-1.24	-0.87	1636.93	6.83	-1.19	-0.81	0.06	-4.49
	375	2062.26	4.34	2050.70	2055.59	2059.87	81.79	2024.05	5.41	-1.85	-1.53	2027.52	4.84	-1.68	-1.37	0.17	-10.53
	500	2407.56	6.40	2418.79	2425.96	2431.75	126.34	2384.15	10.03	-0.97	-1.72	2385.36	9.60	-0.92	-1.67	0.05	-4.23
	Avg.	1450.28	2.07	1442.86	1448.77	1453.54	45.47	1435.02	3.51	-1.05	-0.95	1435.75	3.28	-1.00	-0.90	0.05	-6.53
$\alpha=2$	50	524.30	0.51	502.17	514.49	526.36	3.97	507.24	0.99	-3.25	-1.41	507.50	0.44	-3.20	-1.36	0.05	-55.95
	75	709.39	0.76	696.37	709.96	724.82	7.67	686.93	2.27	-3.17	-3.24	684.36	1.04	-3.53	-3.61	-0.37	-54.46
	100	867.45	0.88	847.06	866.86	885.33	12.12	825.84	3.40	-4.80	-4.73	827.82	1.52	-4.57	-4.50	0.24	-55.44
	175	1165.13	1.75	1142.29	1169.39	1190.52	29.75	1115.37	6.71	-4.27	-4.62	1121.20	3.07	-3.77	-4.12	0.52	-54.27
	250	1362.44	2.83	1359.25	1376.64	1397.85	59.30	1315.32	9.25	-3.46	-4.45	1315.67	4.92	-3.43	-4.43	0.03	-46.80
	375	1702.34	5.63	1712.47	1730.77	1759.36	118.96	1646.24	17.65	-3.30	-4.88	1650.12	10.30	-3.07	-4.66	0.24	-41.68
	500	1998.51	9.05	2028.21	2052.34	2090.54	218.78	1930.14	26.20	-3.42	-5.95	1932.69	16.59	-3.29	-5.83	0.13	-36.70
	Avg.	1189.94	3.06	1183.97	1202.92	1224.97	64.36	1146.73	9.50	-3.63	-4.67	1148.48	5.41	-3.48	-4.53	0.15	-43.05
$\alpha=3$	50	460.63	0.64	428.69	445.88	463.91	3.90	445.65	2.00	-3.25	-0.05	445.97	0.95	-3.18	0.02	0.07	-52.49
	75	649.63	1.08	601.89	626.00	656.16	8.31	605.71	4.11	-6.76	-3.24	607.30	1.86	-6.52	-2.99	0.26	-54.66
	100	779.61	1.20	741.45	765.61	790.62	12.68	728.89	6.24	-6.50	-4.80	734.35	2.47	-5.80	-4.08	0.75	-60.36
	175	1055.32	2.41	1002.35	1040.68	1075.74	37.56	988.79	11.50	-6.30	-4.99	994.78	5.42	-5.74	-4.41	0.61	-52.85
	250	1252.08	3.62	1213.35	1242.29	1274.78	71.08	1175.36	17.32	-6.13	-5.39	1178.72	7.94	-5.86	-5.12	0.29	-54.14
	375	1582.35	7.49	1542.09	1574.58	1605.24	158.38	1487.23	29.74	-6.01	-5.55	1489.69	15.66	-5.86	-5.39	0.17	-47.33
	500	1863.35	10.70	1846.44	1885.92	1934.25	274.46	1737.63	42.31	-6.75	-7.86	1747.16	23.95	-6.24	-7.36	0.55	-43.40
	Avg.	1091.86	3.88	1053.75	1082.99	1114.39	80.91	1024.18	16.17	-6.20	-5.43	1028.28	8.32	-5.82	-5.05	0.40	-48.54

Table 11: (Set A, 2-center) Results of ICP and NICP for Agatz et al. (2018) 2-center instances. Comparison (%) represents the relative difference of NICP compared to ICP. Times are in seconds.

	N	DPS$_{25}$ Obj	DPS$_{25}$ Time	HGA-TAC$^+$ Best	HGA-TAC$^+$ Mean	HGA-TAC$^+$ Worst	HGA-TAC$^+$ Time	ICP Obj	ICP Time	Gap Over (%) DPS	Gap Over (%) HGA	NICP Obj	NICP Time	Gap Over (%) DPS	Gap Over (%) HGA	Comparison (%) Obj	Comparison (%) Time
$\alpha=1$	50	1010.83	0.26	988.61	998.38	1005.84	6.49	1003.95	0.20	-0.68	0.56	1005.32	0.17	-0.54	0.70	0.14	-12.88
	75	1246.14	0.35	1225.86	1235.82	1242.53	11.53	1239.83	0.25	-0.51	0.32	1240.31	0.23	-0.47	0.36	0.04	-8.36
	100	1401.25	0.47	1378.95	1388.68	1396.80	15.06	1391.07	0.38	-0.73	0.17	1392.14	0.36	-0.65	0.25	0.08	-5.74
	175	1903.41	1.25	1891.53	1902.59	1911.47	33.51	1897.55	1.00	-0.31	-0.26	1897.83	0.89	-0.29	-0.25	0.01	-11.31
	250	2262.14	1.53	2248.26	2255.10	2260.93	52.09	2233.64	1.80	-1.26	-0.95	2236.04	1.74	-1.15	-0.85	0.11	-3.44
	375	2842.35	3.47	2826.84	2833.70	2840.19	78.46	2797.78	4.02	-1.57	-1.27	2798.29	4.04	-1.55	-1.25	0.02	0.55
	500	3308.88	7.38	3300.02	3309.00	3315.30	130.59	3267.43	10.57	-1.25	-1.26	3269.46	9.43	-1.19	-1.19	0.06	-10.78
	Avg.	1996.43	2.10	1980.01	1989.04	1996.15	46.82	1975.89	2.60	-1.03	-0.66	1977.06	2.41	-0.97	-0.60	0.06	-7.47
$\alpha=2$	50	819.73	0.48	791.23	806.81	823.82	4.32	811.35	0.96	-1.02	0.56	805.72	0.53	-1.71	-0.14	-0.69	-44.38
	75	1021.36	0.74	976.31	1005.04	1032.75	8.75	1000.69	2.18	-2.02	-0.43	1000.41	0.88	-2.05	-0.46	-0.03	-59.58
	100	1128.40	0.92	1129.89	1148.79	1169.96	12.75	1105.91	2.97	-1.99	-3.73	1108.29	1.53	-1.78	-3.53	0.21	-48.39
	175	1572.17	1.79	1563.52	1587.98	1618.15	29.89	1520.66	6.03	-3.28	-4.24	1524.29	2.74	-3.05	-4.01	0.24	-54.55
	250	1843.45	2.77	1870.08	1895.13	1919.31	58.19	1799.38	9.18	-2.39	-5.05	1799.82	4.64	-2.37	-5.03	0.02	-49.50
	375	2347.51	4.70	2354.47	2386.57	2421.52	124.52	2265.22	16.24	-3.51	-5.08	2264.11	9.01	-3.55	-5.13	-0.05	-44.51
	500	2723.67	9.73	2771.37	2801.04	2839.97	205.81	2651.61	26.91	-2.65	-5.33	2659.25	16.57	-2.37	-5.06	0.29	-38.41
	Avg.	1636.61	3.02	1636.70	1661.62	1689.36	63.46	1593.55	9.21	-2.63	-4.10	1594.55	5.13	-2.57	-4.04	0.06	-44.30
$\alpha=3$	50	755.96	0.68	716.57	734.25	756.19	3.95	731.94	2.16	-3.18	-0.31	739.17	1.08	-2.22	0.67	0.99	-49.84
	75	928.67	0.91	874.92	899.11	930.49	8.80	893.24	3.91	-3.81	-0.65	892.68	1.91	-3.88	-0.72	-0.06	-51.24
	100	1039.73	1.27	1000.70	1028.62	1057.56	14.30	994.15	5.54	-4.38	-3.35	993.47	2.86	-4.45	-3.42	-0.07	-48.40
	175	1443.36	2.53	1388.12	1430.19	1475.94	38.56	1369.12	11.93	-5.14	-4.27	1378.94	5.46	-4.46	-3.58	0.72	-54.26
	250	1692.75	3.49	1669.52	1703.91	1742.72	73.77	1597.10	17.21	-5.65	-6.27	1610.94	8.05	-4.83	-5.46	0.87	-53.19
	375	2174.17	5.80	2141.77	2181.41	2228.38	155.59	2044.64	27.88	-5.96	-6.27	2047.97	14.61	-5.80	-6.12	0.16	-47.62
	500	2555.13	12.80	2499.33	2559.20	2620.79	284.39	2389.57	42.07	-6.48	-6.63	2396.71	24.36	-6.20	-6.35	0.30	-42.09
	Avg.	1512.83	3.92	1470.13	1505.24	1544.58	82.77	1431.39	15.82	-5.38	-4.91	1437.13	8.33	-5.00	-4.53	0.40	-47.31

Table 12: (Set B) Results of ICP and NICP for Bogyrbayeva et al. (2023) instances. Average cost and time values are reported on 100 problem instances for each size N. $\dagger$ is reported by Bogyrbayeva et al. (2023). Comp. (%) represents the relative difference of NICP compared to ICP. Times are in seconds.

		TSP-ep-all		DPS$_{25}$		HM$_{4800}$		HGA-TAC$^+$				ICP		NICP		Comp. (%)	
Dataset	N	Obj	Time	Obj	Time	Obj	Time	Best	Mean	Worst	Time	Obj	Time	Obj	Time	Obj	Time
Random	50	397.59	18.05	404.99	0.46	396.26	3.80	391.49	398.64	406.71	3.85	394.32	1.50	394.85	0.72	0.13	-51.56
	100	535.94	2088.34	548.04	1.14	544.58	14.13	535.57	545.02	554.03	11.34	534.07	3.70	534.18	1.72	0.02	-53.54
	Avg.	466.77	1053.19	476.51	0.80	470.42	8.96	463.53	471.83	480.37	7.59	464.20	2.60	464.52	1.22	0.08	-52.55
Amsterdam	50	3.26	18.00	3.37	0.46	3.31†	1.41†	3.26	3.33	3.41	4.02	3.29	1.18	3.29	0.55	0.29	-53.33