
Learning to Search, Searching to Learn: A Closed-Loop Framework for Large-Scale Vehicle Routing

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Scaling neural solvers to large Vehicle Routing Problems (VRPs) runs into two
2 recurring difficulties. To stay within compute budgets, scalable solvers tend to rely
3 on partitioning, local candidate restriction, or staged decisions, which limits the
4 global structure that ultimately drives solution quality. To improve the final solution,
5 many neural pipelines fall back on a classical search, but use it as a one-shot post-
6 processor. The model predicts, search repairs, and no sustained feedback is formed
7 between them. We propose **LSSL** (Learning to Search, Searching to Learn), a
8 closed-loop learning-search framework for large-scale VRP that addresses both
9 difficulties in one framework. LSSL predicts a search-friendly guidance heatmap
10 on a sparse candidate graph. A classical heuristic backend then uses this heatmap
11 to refine the incumbent solution, and the structural state returned by the backend is
12 fed back to recondition the model. As a result, learning and search interact within a
13 closed loop, rather than operating as two detached stages. Built on a sparse diffusion
14 transformer over an $O(NK)$ candidate graph, the proposed method scales to ultra-
15 large instances on a single consumer-grade GPU without hard decomposition.
16 Experiments on large-scale TSP and CVRP benchmarks demonstrate that LSSL
17 achieves superior scalability, efficiency, and solution quality. Ablation studies
18 further reveal that the closed learning-search loop architecture accounts for a
19 substantial portion of the observed performance gains.

20 1 Introduction

21 The Vehicle Routing Problem (VRP) is a central problem in combinatorial optimization, with appli-
22 cations in logistics and supply-chain planning [1–4]. Because VRP is NP-hard, exact optimization is
23 impractical at large scale [5]. Strong heuristics and metaheuristics such as Lin-Kernighan-Helsgaun
24 (LKH) for the Traveling Salesman Problem (TSP) and Hybrid Genetic Search (HGS) for the Ca-
25 pacitated VRP (CVRP) remain the strong practical baselines, but rely on engineered candidate sets,
26 neighborhoods, and search control [6, 7]. Neural Combinatorial Optimization (NCO) learns solving
27 policies or structural priors directly from data [8–14]. While this idea has worked well on small and
28 medium instances, making it work reliably on ultra-large VRPs remains difficult.

29 A core difficulty is scaling learning-based solvers to very large instances (e.g., 100K nodes) un-
30 der a bounded compute budget while still capturing global structure. Recent work has improved
31 cross-scale generalization through stronger constructive solvers, better inductive biases, or test-time
32 adaptation [15–21]. Some works decompose ultra-large instances into subproblems and solve them
33 hierarchically or in stages [22–28]. Both lines share the same trade-off. Tight budgets localize or
34 partition decisions, weakening long-range couplings across regions and routes that determine overall

35 quality. This is especially critical in ultra-large VRP, because many local connections can only be
36 judged properly in the context of a full route, or even the coordination among multiple routes.

37 Another core difficulty is that neural pipelines still use classical heuristic search almost exclusively
38 as a single post-processing pass on the model output. Many works show that edge heatmaps,
39 candidate-edge scores, or initial solutions produced by neural models can substantially improve final
40 results when combined with downstream search or refinement [29–37]. For example, NeuroLKH
41 demonstrates the effectiveness of combining neural guidance with classical heuristic search, showing
42 that a learned module can provide valuable candidate information to an LKH backend and thereby
43 improve final solution quality [38]. However, this coupling is usually still one-way, with only a
44 single prediction followed by a single refinement, while trajectory structure is neither fed back during
45 inference nor mined for training. Search improves the initial solution but rarely participates as a
46 sustained, closed-loop part of learning. Although recent work has started to question the limitation of
47 treating search only as post-hoc refinement [39–42], how to unify learning with a classical heuristic
48 search backend into a closed-loop system that operates in both inference and training remains an
49 open question for large-scale VRP.

50 To address the above difficulties, we propose **LSSL (Learning to Search, Searching to Learn)**, a
51 closed-loop framework that connects learning and search during training and inference. *Learning to*
52 *Search* means the model does not calculate a full route in one pass; it predicts a sparse edge heatmap
53 on a sparse candidate graph, and heuristic search uses that heatmap as guidance. In the *Searching*
54 *to Learn* step, LSSL retains multiple high-quality route sets encountered during search on each
55 training instance, projects them onto the shared candidate graph, and aggregates these projections
56 into row-wise soft targets for model updates. These targets place higher mass on edges repeatedly
57 supported by good solutions, while keeping uncertainty where several local connections remain
58 plausible. At inference time, LSSL closes the complementary loop on the same instance with multiple
59 learning–search cycles. After each search round, the incumbent route set is projected back to a
60 structural heatmap that conditions the next model pass, and search spends another budget refining
61 the incumbent under the refreshed guidance. As a result, training targets distill consensus from
62 search trajectories, while inference propagates search feedback through repeated reconditioning. All
63 quantities exchanged between the model and the search backend are defined on a sparse candidate
64 graph with $O(NK)$ edges for fixed K . Prediction and feedback are therefore updated row-wise over
65 only K neighbours, yielding near-linear scaling in N rather than the quadratic cost of dense $O(N^2)$
66 models. This makes large VRP instances feasible under limited computational resources.

67 Experiments on large-scale TSP and CVRP benchmarks show that LSSL simultaneously delivers
68 strong solution quality, inference efficiency, and scalability, attaining state-of-the-art results in
69 multiple settings while remaining tractable on a single consumer GPU at node N up to 100K. The
70 contributions of this paper are as follows:

- 71 • **LSSL framework.** A sparse diffusion Transformer D_θ and a heuristic backend S are coupled
72 through a shared sparse candidate graph. At inference, they alternate to refine the same instance; at
73 training, search-refined solutions are aggregated into row-wise soft targets that supervise D_θ . This
74 coupling is backend-agnostic.
- 75 • **$O(NK)$ scaling.** Block-streamed online-softmax attention keeps memory and compute near-
76 linear in N , letting the closed loop run end-to-end on one consumer GPU at 100K nodes without
77 divide-and-conquer.
- 78 • **Empirical findings.** LSSL matches or outperforms strong heuristics and neural-guided baselines on
79 large-scale TSP and CVRP under matched compute; ablations attribute distinct gains to inference
80 reconditioning and training-time consensus supervision.

81 **2 LSSL: Closed-loop Learning–Search Framework**

82 **2.1 Large-scale VRP problem formulation**

83 Given a VRP instance on a node set \mathcal{V} with $N = |\mathcal{V}|$ nodes, we seek a feasible solution \mathbf{x}^* that
84 minimizes the total tour or route length. TSP and CVRP differ in their feasibility constraints: TSP
85 solutions form a single Hamiltonian tour that visits every node exactly once, while CVRP solutions
86 decompose into depot-rooted routes that jointly serve all customers under a vehicle-capacity constraint.
87 At the scales we target (N up to 100K), materialising the full $O(N^2)$ edge set is prohibitive in memory

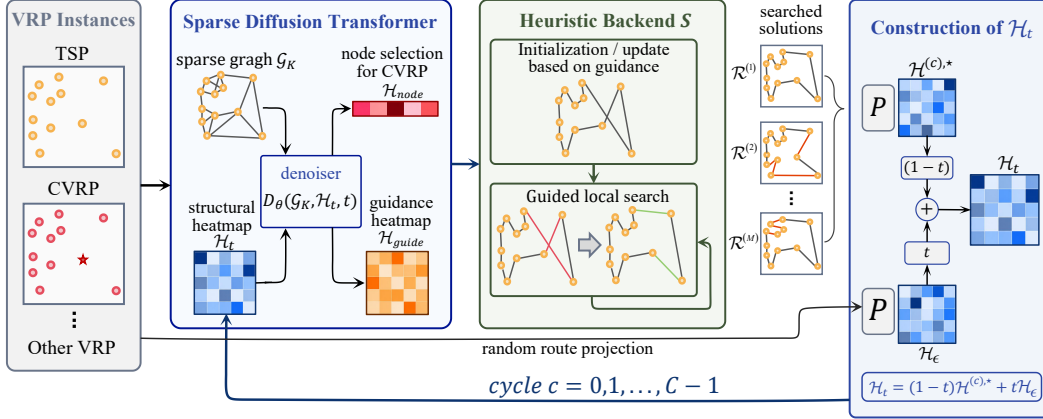


Figure 1: **Overview of LSSL.** The Sparse Diffusion Transformer takes the current structural state of the solution and the instance graph as input, and outputs a guidance heatmap $\mathcal{H}_{\text{guide}}$ over promising local connections. For CVRP, it additionally outputs a node-selection heatmap $\mathcal{H}_{\text{node}}$ to anchor route boundaries. The backend uses this guidance first to construct an initial solution and then to locally improve it; the improved solutions are used to update the structural state and fed back as the starting point of the next cycle. The same loop applies unchanged to TSP, CVRP, and other VRP variants.

88 and time. However, for Euclidean instances, the optimal solution rarely uses long edges, so most of
 89 those $O(N^2)$ candidates are not actually useful. LSSL therefore operates entirely on a sparse K -NN
 90 candidate subgraph of size $O(NK)$ with $K \ll N$ (Section 2.3). This sparse representation is the
 91 foundation on which both LSSL’s neural backbone and its search backend operate.

92 2.2 Framework overview

93 The LSSL framework (Figure 1) closes a loop between a sparse diffusion Transformer D_θ and a
 94 classical heuristic backend S on a shared sparse candidate graph: D_θ produces guidance heatmaps
 95 that S consumes, and S ’s refined output is projected back as the input for the next call to D_θ .
 96 Particularly, the framework rests on three backend-agnostic interfaces. The *denoiser* D_θ is a sparse
 97 diffusion Transformer that consumes the instance graph \mathcal{G}_K , a structural heatmap \mathcal{H}_t over candidate
 98 edges, and a noise level $t \in [0, 1]$, and produces a guidance heatmap $\mathcal{H}_{\text{guide}}$. The *search backend* S
 99 consumes $\mathcal{H}_{\text{guide}}$ and refines the incumbent route set under a per-cycle budget; we instantiate it as
 100 Guided-LKH [6] for TSP and Guided-HGS [43, 44] for CVRP (Appendices F and G). The *projection*
 101 P maps any candidate solution \mathcal{R} (a TSP tour or a CVRP route set) onto the sparse $N \times K$ adjacency
 102 by indicating which candidate edges are actually used in \mathcal{R} .

103 In our framework, training (Section 2.5) supplies D_θ with targets prepared by S : for each instance,
 104 several S -refined route sets are projected through P and aggregated into a consensus heatmap \mathcal{H}^* ,
 105 and D_θ is trained to recover \mathcal{H}^* from a noisy interpolation between \mathcal{H}^* and a random projection.
 106 Inference (Section 2.6) runs C cycles per instance, starting from a random feasible solution. Each
 107 cycle projects the current incumbent through P to obtain a structural heatmap, denoises it with D_θ
 108 along a noise schedule into a guidance heatmap, and calls S under that guidance to produce the next
 109 incumbent; the best solution across cycles is returned.

110 2.3 Guidance Heatmap on a Sparse Candidate Graph

111 Existing diffusion-style generative NCO methods [32–35] train a predictor to recover a reference
 112 solution \mathbf{x}^* from a perturbed state \mathbf{x}_t , where \mathbf{x}_t is itself a perturbed *solution* variable and is therefore
 113 pulled toward feasibility throughout the trajectory (Section A surveys the variants). This commits
 114 prematurely: the model fixes edge-level decisions before search begins, leaving search only local
 115 refinement around that commitment (as shown in Figure 2). The issue is especially severe for CVRP,
 116 where capacity and depot constraints aggressively drive \mathbf{x}_t onto the feasibility manifold.

117 LSSL instead parameterises a guidance heatmap over candidate edges. The model only expresses
 118 how likely each local connection is, leaving the assembly of a feasible solution entirely to search.

119 Concretely, the global graph \mathcal{G} is first mapped to a sparse candidate graph

$$\mathcal{G}_K = (\mathcal{V}, \mathcal{E}_K), \quad \mathcal{E}_K = \{(i, j) \mid j \in \mathcal{N}_K(i), |\mathcal{N}_K(i)| = K\}, \quad (1)$$

120 with per-node neighborhood $\mathcal{N}_K(i)$ which compresses the dominant cost from $O(N^2)$ to $O(NK)$.

121 On this sparse graph, we define the sparse diffusion Transformer D_θ as a per-step denoiser,

$$\mathcal{H}_{\text{guide}} = D_\theta(\mathcal{G}_K, \mathcal{H}_t, t), \quad (2)$$

122 where $\mathcal{H}_t \in [0, 1]^{N \times K}$ is a structural prior over candidate edges at noise level $t \in [0, 1]$. Each row
 123 i assigns a probability mass to each of node i 's K candidate neighbors, expressing how likely the
 124 edge (i, j) belongs to a high-quality solution. The two endpoints anchor opposite extremes. At
 125 $t = 0$, \mathcal{H}_t is a clean prior fully informative about high-quality solution regions; at $t = 1$, \mathcal{H}_t is a
 126 structureless random prior. The precise interpolation between them is given by Eq. (7) in Section 2.5.
 127 Crucially, $\mathcal{H}_{\text{guide}}$ is never required to satisfy combinatorial feasibility constraints, leaving full room
 128 for downstream search to adjust.

129 2.4 Sparse diffusion Transformer

130 The denoiser D_θ consists of a sparse encoder and a time-conditioned sparse decoder, both operating on the same $N \times K$
 131 sparse adjacency. The encoder takes per-
 132 node static features \mathbf{s}_i (coordinates for
 133 TSP, plus normalized demand for CVRP)
 134 and propagates them across \mathcal{G}_K through
 135 standard sparse self-attention, producing
 136 instance-aware node embeddings. The de-
 137 coder takes these embeddings as initial
 138 state and additionally injects two dynamic
 139 signals, i.e., the current heatmap \mathcal{H}_t and
 140 the noise level t , through the structural-bias
 141 channel below.
 142
 143

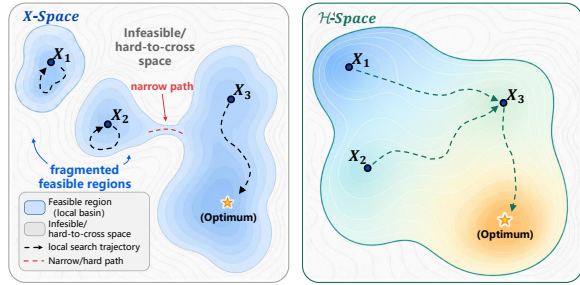


Figure 2: **x-space vs. \mathcal{H} -space.** In x-space, feasibility constraints fragment the landscape into disconnected basins, trapping local search. \mathcal{H} -space provides a smooth, connected representation that unifies all initialisations with the optimum, enabling global exploration.

144 **Decoder attention.** At decoder layer ℓ , head h , the attention score from node i to a candidate
 145 neighbor $j \in \mathcal{N}_K(i)$ is

$$e_{ij}^{(h)} = \frac{\langle \mathbf{q}_i^{(h)}, \mathbf{k}_j^{(h)} \rangle}{\sqrt{d_h}} + b_{ij}, \quad b_{ij} = u([\mathcal{H}_t]_{ij}) + v(\phi(t)), \quad (3)$$

146 where $\mathbf{q}_i^{(h)}, \mathbf{k}_j^{(h)}$ are the query and key projections of the current hidden state, d_h is the per-head
 147 dimension, $\phi(t)$ is a sinusoidal embedding of t , and u, v are learned scalar projections. The bias b_{ij}
 148 couples the structural belief on edge (i, j) with the global noise level where the first term modulates
 149 attention per-edge under \mathcal{H}_t , while the second lets the layer stay broad at high t and sharpen as
 150 $t \rightarrow 0$. The attention scores then drive a row-wise softmax over $\mathcal{N}_K(i)$, aggregate the value vectors
 151 $\mathbf{v}_j^{(h)}$, and feed into a standard FFN-residual block; stacking L such layers yields the final per-node
 152 representation $\mathbf{g}_i \in \mathbb{R}^{d_{\text{dec}}}$ at the decoder output.

153 To keep this layer numerically stable and $O(NK)$ at large K , the row-wise softmax over $\mathcal{N}_K(i)$ is
 154 never materialized; instead, $\mathcal{N}_K(i)$ is traversed block by block while maintaining an exact online-
 155 softmax statistic. Slot-level recursion and value-aggregation details are deferred to Appendix D.4.

156 **Edge-scoring head.** A lightweight head turns the final node representations \mathbf{g}_i into the guidance
 157 heatmap by a row-wise softmax over candidate-neighbor inner products,

$$[\mathcal{H}_{\text{guide}}]_{ij} = \frac{\exp(\langle \mathbf{g}_i, \mathbf{g}_j \rangle)}{\sum_{j' \in \mathcal{N}_K(i)} \exp(\langle \mathbf{g}_i, \mathbf{g}_{j'} \rangle)}, \quad j \in \mathcal{N}_K(i). \quad (4)$$

158 Two nodes whose final representations align in $\mathbb{R}^{d_{\text{dec}}}$ are scored as a likely edge, so the geometry
 159 encoded by the encoder, the structural belief carried by \mathcal{H}_t , and the time-dependent sharpening from
 160 $\phi(t)$ all flow through a single output channel.

161 **2.5 Search-feedback training**

162 The training procedure of LSSL specifies how search outputs are recycled into supervision on the
 163 sparse candidate graph. Unlike GenSCO [35], which supervises the model with a single high-quality
 164 tour per instance, LSSL aggregates multiple search-refined route sets per instance into a row-wise
 165 structural consensus and trains the denoiser as a direct endpoint predictor of this consensus.

166 For each training instance, the search backend S produces M high-quality route sets $\{\mathcal{R}^{(m)}\}_{m=1}^M$.
 167 The search-feedback target \mathcal{H}^* is the projection of these solutions onto the sparse candidate graph
 168 \mathcal{G}_K :

$$\mathcal{H}^* = P(\{\mathcal{R}^{(m)}\}_{m=1}^M; \mathcal{G}_K). \quad (5)$$

169 Here P extends the projection of Section 2.2 from a single solution to a set: it counts per-row edge
 170 usage across the M route sets and row-normalises. Specifically, for vertex i and candidate slot
 171 $j \in \mathcal{N}_K(i)$,

$$[P(\{\mathcal{R}^{(m)}\}_{m=1}^M; \mathcal{G}_K)]_{ij} = \frac{\sum_{m=1}^M \mathbb{1}\{(i, j) \in \mathcal{R}^{(m)}\}}{\sum_{j' \in \mathcal{N}_K(i)} \sum_{m=1}^M \mathbb{1}\{(i, j') \in \mathcal{R}^{(m)}\}}, \quad (6)$$

172 where $\mathbb{1}\{(i, j) \in \mathcal{R}^{(m)}\}$ equals 1 iff j is a tour-neighbour of i in $\mathcal{R}^{(m)}$. The result is a sparse
 173 structural consensus over multiple search route sets: edges repeatedly selected by good solutions
 174 receive higher probability mass, while alternative local connections remain plausible when search has
 175 not clearly resolved them.

176 We next define the source state of the structural transition. Let \mathcal{R}_ϵ be a randomly initialized route set
 177 generated without using the search-feedback target. Its sparse projection $\mathcal{H}_\epsilon = P(\mathcal{R}_\epsilon; \mathcal{G}_K)$ serves
 178 as a random route projection. Given the clean target \mathcal{H}^* and the random route projection \mathcal{H}_ϵ , we
 179 construct a continuous interpolation path on the sparse $N \times K$ adjacency:

$$\mathcal{H}_t = (1 - t)\mathcal{H}^* + t\mathcal{H}_\epsilon, \quad t \sim \mathcal{U}(0, 1). \quad (7)$$

180 Here $t = 0$ corresponds to the clean search-feedback target and $t = 1$ to the randomly initialized
 181 structural state. This interpolation should not be interpreted as a feasible VRP solution trajectory;
 182 it is instead a relaxed structural-prior trajectory over candidate edges, so neither the clean target
 183 \mathcal{H}^* nor the intermediate states \mathcal{H}_t are required to satisfy TSP or CVRP feasibility constraints. This
 184 relaxation lets D_θ focus on learning structural patterns rather than enforcing feasibility, which is left
 185 to downstream search at inference time.

186 Training minimizes a row-wise sparse cross-entropy on \mathcal{G}_K :

$$\mathcal{L}_{\text{LSSL}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_K(i)} \mathcal{H}_{ij}^* \log[D_\theta(\mathcal{G}_K, \mathcal{H}_t, t)]_{ij}. \quad (8)$$

187 CVRP additionally uses a per-customer node-selection loss. The decoder outputs a node-selection
 188 heatmap $\mathcal{H}_{\text{node}} \in (0, 1)^{N_c}$, supervised by a binary target $\mathcal{H}_{\text{node}}^* \in \{0, 1\}^{N_c}$ that marks customers
 189 starting or ending a sub-route in $\{\mathcal{R}^{(m)}\}_{m=1}^M$:

$$\mathcal{L}_{\text{node}} = -\frac{1}{N_c} \sum_{i=1}^{N_c} \left[[\mathcal{H}_{\text{node}}^*]_i \log[\mathcal{H}_{\text{node}}]_i + (1 - [\mathcal{H}_{\text{node}}^*]_i) \log(1 - [\mathcal{H}_{\text{node}}]_i) \right], \quad (9)$$

190 with N_c the number of customers. The total CVRP objective combines the edge and node losses:

$$\mathcal{L}_{\text{CVRP}} = \mathcal{L}_{\text{LSSL}} + \lambda_{\text{node}} \mathcal{L}_{\text{node}}. \quad (10)$$

191 **Self-bootstrapping across rounds.** The search-feedback mechanism supports iterative self-
 192 improvement at training time. In round 0, the search backend S is applied once to each training
 193 instance, and its M refined route sets are aggregated via Eq. (5) into the target \mathcal{H}^* ; D_θ is then trained
 194 on the resulting (instance, target) pairs. In each subsequent round, a subset of instances is re-sampled,
 195 and the current D_θ generates guidance heatmaps that direct S to produce new route sets. These are
 196 re-aggregated into an updated \mathcal{H}^* that replaces the previous target for the same instance, and D_θ is
 197 fine-tuned on the updated training set.

198 This loop is particularly valuable at scales where obtaining external high-quality labels is intractable.
 199 By coupling a smaller-scale-trained D_θ with S , self-bootstrapping generates supervision that improves
 200 with each round, without relying on a stronger external solver. Our main experiments use a single
 201 self-improvement round; details of the training schedule are deferred to Appendix C.

202 **2.6 Inference-time learning–search refinement loop**

203 The denoiser D_θ is trained as a direct endpoint predictor: given \mathcal{G}_K , an intermediate state \mathcal{H}_t on
 204 the interpolation path of Eq. (7), and the noise level t , it directly estimates the clean target \mathcal{H}^* .
 205 Unlike velocity-matching parameterizations, D_θ predicts this endpoint at every $t \in [0, 1]$, so the same
 206 network can later be unrolled along a noise-axis schedule at inference.

207 Inference runs C learning–search cycles on the same instance, alternating D_θ and the search backend
 208 S through the projection P . Algorithm 1 describes one full pass. Starting from a random feasible
 209 initialisation \mathcal{R}_ϵ as $\mathcal{R}^{(0)}$, each cycle (i) projects the current incumbent $\mathcal{R}^{(c)}$ to a structural heatmap
 210 $\mathcal{H}^{(c),*}$ on \mathcal{G}_K , (ii) unrolls D_θ along the noise-axis schedule $1 = t_0 > \dots > t_T = 0$ to obtain a guidance
 211 heatmap $\mathcal{H}_{\text{guide}}^{(c)}$, and (iii) invokes S under that guidance with budget B to refine the incumbent.
 212 The best feasible solution across the C cycles is returned. Backend-specific details are deferred to
 213 Appendices F and G.

214 The inner loop is a deterministic integrator that calls D_θ at every schedule step. Since D_θ is supervised
 215 at every $t \in [0, 1]$ (Eq. (8)), it can be queried at any t_r to obtain a current estimate of the clean target
 216 \mathcal{H}^* . The next state is then obtained by sliding along the linear interpolation path of Eq. (7) from t_r to
 217 t_{r+1} , with the model’s estimate playing the role of the clean endpoint:

$$\mathcal{H}_{t_{r+1}} = \frac{t_{r+1}}{t_r} \mathcal{H}_{t_r} + \left(1 - \frac{t_{r+1}}{t_r}\right) D_\theta(\mathcal{G}_K, \mathcal{H}_{t_r}, t_r), \quad (11)$$

218 a convex combination with mixing weight $1 - t_{r+1}/t_r$ on the prediction; the trajectory converges to
 219 \mathcal{H}^* as $t_r \rightarrow 0$.

220
221

222 **3 Experiments**

223 In this section, we systematically evaluate LSSL
 224 from four perspectives on large-scale TSP and
 225 CVRP. First, we compare it with classical
 226 solvers, constructive NCO methods, heatmap-
 227 based NCO methods, and two-stage large-scale
 228 routing methods on synthetic benchmarks to
 229 assess overall solution quality, inference effi-
 230 ciency, and scalability. Second, we evaluate
 231 cross-distribution generalization on real-world
 232 TSPLIB, VLSI, Art-TSP, and CVRPLIB in-
 233 stances. Third, we conduct ablations on the
 234 closed-loop mechanism and search-feedback su-
 235 pervision to verify whether the core design truly
 236 brings performance gains. Finally, we analyze
 237 the behavior of the framework under different in-
 238 ference budgets, cycle counts, candidate widths,
 239 and large-scale scaling curves.

240 **3.1 Experimental setup**

241 **Datasets.** A TSP instance contains N 2-D coordinates and a reference tour; a CVRP instance
 242 further specifies vehicle capacity. Following [11], synthetic instances are sampled uniformly
 243 from $[0, 1]^2$. Training sets are multi-scale: TSP uses 1.28M / 64K / 6.4K / 640 instances at
 244 $N=100/1K/10K/100K$, and CVRP uses 1.28M / 64K / 6.4K instances at $N=100/1K/10K$. For
 245 TSP100, TSP1K, and CVRP100, training labels are produced by LKH3 (TSP) and HGS (CVRP). At
 246 the larger scales TSP10K, TSP100K, CVRP1K, and CVRP10K, directly labelling the full training
 247 split with high-quality solver outputs is no longer feasible: the per-instance time required exceeds any
 248 uniform labelling pipeline that fits our offline training budget. We therefore use the self-bootstrapping
 249 mechanism of Section 2.5: a smaller-scale-trained LSSL model D_θ produces guidance heatmaps on
 250 freshly sampled target-scale instances and directs the search backend S to solve them; the resulting
 251 route sets are projected and aggregated via Eq. (5) into the search-feedback target \mathcal{H}^* , on which

Algorithm 1 LSSL inference (one instance).

Require: $\mathcal{G}_K; D_\theta$; backend S ; cycles C ; schedule
 $\{t_r\}_{r=0}^T, t_0 = 1, t_T = 0$; budget B
Ensure: best solution $\bar{\mathcal{R}}$

- 1: build candidate graph \mathcal{G}_K
- 2: $\mathcal{R}^{(0)} \leftarrow \mathcal{R}_\epsilon$ ▷ random feasible init
- 3: $\bar{\mathcal{R}} \leftarrow \mathcal{R}^{(0)}$
- 4: **for** $c = 0, \dots, C - 1$ **do**
- 5: $\mathcal{H}^{(c),*} \leftarrow P(\mathcal{R}^{(c)}; \mathcal{G}_K)$ ▷ (i)
- 6: $\mathcal{H}_{t_0} \leftarrow \mathcal{H}^{(c),*}$
- 7: **for** $r = 0, \dots, T - 1$ **do** ▷ (ii)
- 8: update $\mathcal{H}_{t_{r+1}}$ by Eq. (11)
- 9: **end for**
- 10: $\mathcal{H}_{\text{guide}}^{(c)} \leftarrow \mathcal{H}_{t_r}$
- 11: $\mathcal{R}_S \leftarrow S(\mathcal{R}^{(c)}, \mathcal{H}_{\text{guide}}^{(c)}, B)$ ▷ (iii)
- 12: **if** $\text{cost}(\mathcal{R}_S) < \text{cost}(\bar{\mathcal{R}})$ **then**
- 13: $\bar{\mathcal{R}} \leftarrow \mathcal{R}_S$
- 14: **end if**
- 15: $\mathcal{R}^{(c+1)} \leftarrow \bar{\mathcal{R}}$ ▷ best-so-far carries to next
cycle
- 16: **end for**
- 17: **return** $\bar{\mathcal{R}}$

Table 1: Results on synthetic TSP and CVRP instances. OOM: The method triggered an out-of-memory condition. N.T.: Not trained; the original publication did not train or report a model for this problem scale. C : The number of LSSL learning–search cycles; B_{trials} : The per-cycle LKH trial count. B_t : The per-cycle HGS time budget in seconds.

Method	TSP100			TSP1K			TSP10K		
	Length	Gap	Time	Length	Gap	Time	Length	Gap	Time
LKH3	7.763	0.000%	3.5m	23.119	0.000%	30.75m	71.778	0.000%	33.2m
Concorde	7.763	0.000%	2.125m	23.119	0.000%	29.25m	72	0.310%	1.4h
DIFUSCO	7.78	0.237%	12.4m	23.56	1.913%	12.1m	73.62	2.58%	14.5m
Fast-T2T	7.76	0.012%	7.1m	23.22	0.422%	15.6m	72.94	1.63%	8.7m
GenSCO ($C=10$) 2Opt	7.763	0.001%	16.2s	23.145	0.112%	30.2s	N.T.	N.T.	N.T.
GenSCO ($C=160$) 2Opt	7.763	0.000%	3.4m	23.132	0.057%	7.1m	N.T.	N.T.	N.T.
H-TSP	-	-	-	24.718	6.912%	34s	77.75	8.320%	2.2m
GLOP	7.767	0.048%	5.25m	23.84	3.119%	10.2s	75.04	4.545%	1.9m
UDC- α_{50} ($\alpha = 50$)	7.788	0.319%	1.38h	23.53	1.782%	6.21m	OOM	OOM	OOM
POMO aug $\times 8$	7.774	0.138%	21.3s	32.5	40.577%	2.1m	OOM	OOM	OOM
BQ bs16	7.764	0.010%	1.25m	23.432	1.354%	1.31m	OOM	OOM	OOM
LEHD RRC1000	7.763	0.002%	8.63m	23.288	0.731%	19.5m	80.9	12.709%	13.5m
L2C-Insert Greedy	7.812	0.628%	7.8s	23.963	3.808%	5.4s	77.299	7.727%	40.6s
L2C-Insert ($I=1000$)	7.692	0.000%	2.87m	23.185	0.438%	9.50m	73.227	2.052%	10.44m
LSSL ($C=1$, $B_{\text{trials}}=128$)	7.763	0.000%	1.3m	23.122	0.017%	17.1s	71.789	0.015%	32.2s
LSSL ($C=2$, $B_{\text{trials}}=256$)	-	-	-	23.120	0.007%	38.3s	71.783	0.007%	1.1m
LSSL ($C=4$, $B_{\text{trials}}=512$)	-	-	-	23.119	0.000%	3.6m	71.773	-0.006%	3.3m

Method	CVRP100			CVRP1K			CVRP10K		
	Length	Gap	Time	Length	Gap	Time	Length	Gap	Time
HGS	15.592	0.000%	63m	41.212	0.000%	3.5h	226.177	0.000%	1h
GLOP	-	-	-	45.753	11.019%	1.93m	276.252	22.140%	27.98m
UDC- α_{50} ($\alpha = 50$)	16.291	4.481%	2.69h	43.583	5.753%	11.96m	OOM	OOM	OOM
POMO aug $\times 8$	15.75	1.011%	17.2s	100.990	145.045%	1.76m	OOM	OOM	OOM
BQ bs16	15.81	1.396%	1.13m	43.121	4.632%	1.56m	275.494	21.805%	8.6m
LEHD RRC1000	15.63	0.242%	7.50m	42.408	2.902%	15.56m	248.604	9.916%	30.38m
L2C-Insert Greedy	16.580	4.001%	46s	43.076	7.691%	12.5s	254.990	12.739%	2.35m
L2C-Insert ($I=1000$)	16.011	0.433%	5.32m	41.704	4.261%	23.95m	245.940	8.738%	26.81m
LSSL ($C=1$, $B_t=4$)	15.591	-0.009%	4.2m	41.681	1.138%	25.0s	228.32	0.947%	4.7s
LSSL ($C=2$, $B_t=32$)	15.589	-0.021%	66.7m	41.328	0.281%	6.7m	224.69	-0.657%	64.1s
LSSL ($C=4$, $B_t=128$)	-	-	-	41.186	-0.063%	53.3m	222.49	-1.630%	8.5m

252 the target-scale LSSL model is then trained as a fixed dataset. Self-improved training has been
253 explored as a precedent for large-scale NCO [19]. The full recipe and per-scale time costs are
254 reported in Appendices C. Test sets follow [11, 32, 24, 21]: TSP100 has 1,280 instances, TSP1K 128,
255 CVRP100 1,280, CVRP1K 100 [24], and TSP/CVRP10K and TSP100K each 16; CVRP capacities
256 are 50/200/300 for $N=100/1K/\geq 10K$ [24]. For real-world generalization we use TSPLIB [45],
257 VLSI [46], Art-TSP [47], and CVRPLIB Set-X / Set-XL [48], with reference costs computed from
258 the tours published on each benchmark’s official site; per-instance results are in the Appendix B.

259 **Model Setting.** An LSSL configuration is specified by C (number of learning–search cycles) and
260 the per-cycle search size: the LKH trial count B_{trials} with backend Guided-LKH for TSP, and the
261 HGS time budget B_t with backend Guided-HGS for CVRP. The full baseline list is deferred to
262 Appendix C.5.

263 **Metrics.** (1) Length: the average tour length over the corresponding test instances; (2) Gap: the
264 relative performance gap with respect to the optimal solution or the reference solver; (3) Time:
265 the total computation time required to solve the entire test set, with each method running up to 16
266 instances in parallel whenever memory permits to reflect practical parallel throughput; otherwise, the
267 largest feasible parallelism is used, resulting in longer runtime.

268 3.2 Experimental Results

269 **Results on Synthetic Dataset.** Tables 1 and 2 report the solving performance comparison on
270 synthetic TSP and CVRP. LSSL is the only neural method that simultaneously surpasses LKH3 and
271 HGS across the entire 100–100K scale range. The advantage compounds at scale. On TSP10K, LSSL
272 ($C=4$, $B_{\text{trials}}=512$) drives the gap to -0.006% in just 3.3 m, surpassing LKH3 itself at roughly $10\times$

Table 2: Results on synthetic TSP100K.

Method	TSP100K		
	Length	Gap	Time
LKH3	225.997	0.000%	25h
GLOP	237.610	5.139%	3.90m
INViT greedy	242.260	7.196%	5.00h
L2C-Insert Greedy	242.757	7.420%	6.86m
L2C-Insert ($I=1000$)	237.221	4.971%	16.71m
L2C-Insert ($I=10000$)	230.132	1.834%	1.77h
LSSL ($C=1, B_{\text{trials}}=128$)	226.219	0.098%	16.9m
LSSL ($C=2, B_{\text{trials}}=512$)	226.028	0.014%	1.92h
LSSL ($C=4, B_{\text{trials}}=1024$)	225.997	0.000%	6.86h

Table 3: Results on real-world TSP benchmarks.

Method	TSPLIB 50-1K	VLSI		Art-TSP 100K-200K
		50-2K	2K-200K	
UDC- x_{250} ($\alpha = 50$)	3.072%	3.942%	-	-
GLOP	1.153%	3.881%	8.051%	3.238%
GenSCO	0.184%	1.588%	-	-
BQ bs16	1.557%	6.259%	-	-
L2C-Insert ($I=1000$)	1.499%	2.051%	4.142%	4.240%
NeuroLKH	0.020%	-0.061%	-	-
LSSL ($C=1, B_{\text{trials}}=128$)	0.002%	-0.264%	-0.273%	0.034%
LSSL ($C=2, B_{\text{trials}}=256$)	0.000%	-0.267%	-0.299%	0.026%

less computation than LKH3 (33.2 m). On TSP100K, LSSL ($C=4, B_{\text{trials}}=1024$) fully closes the gap to 0.000% in 6.86 h, roughly a $3.6\times$ speedup over LKH3 (25 h). On CVRP10K, LSSL ($C=4, B_t=128$) drives the gap to -1.630% in just 8.5 m, the first time a neural solver genuinely surpasses HGS at this scale, at roughly $7\times$ less computation than HGS (1 h). By contrast, the other neural baselines either OOM by $\geq 10K$, including GenSCO, Fast-T2T and UDC- x_{50} , or sit at gaps an order of magnitude wider, for example L2C-Insert ($I=10,000$) still lands at 1.834% on TSP100K after 1.77 h, and on CVRP10K LEHD RRC1000 and L2C-Insert ($I=1000$) remain at 9.916% and 8.738% respectively.

Results on Real-world Benchmark. Tables 3 and 4 report cross-distribution generalization on TSPLIB, VLSI, Art-TSP, and CVRPLIB Set-X / Set-XL. LSSL transfers cleanly from synthetic uniform $[0, 1]^2$ to all four distributions, and on the larger benchmarks the average tour is shorter than the published references themselves: at $C=2$ the gap is -0.267% on VLSI 50-2K and -0.299% on 2K-200K, despite never training on VLSI-style points; on TSPLIB 50-1K it closes to 0.000%, and on Art-TSP 100K-200K to 0.026%. Within each benchmark, raising the search budget tightens the gap monotonically, and the marginal improvement tracks how suboptimal the reference is: $0.002 \rightarrow 0.000\%$ (TSPLIB) and $-0.264 \rightarrow -0.267\%$ (VLSI 50-2K) where the reference is near-optimal, versus $1.067 \rightarrow 0.171\%$ (Set-X) and $4.149 \rightarrow 1.838\%$ (Set-XL) where it is far from optimal. Against the strongest neural baselines, the largest reductions appear where post-hoc heatmap and constructive methods struggle most: a $9.5\times$ smaller gap than L2C-Insert on Set-XL and a $124\times$ smaller gap than GLOP on Art-TSP.

Table 4: Results on real-world CVRP benchmarks. †: Several instances are not solvable due to exceeding time limits.

Method	CVRPLIB-setX 100-1K	CVRPLIB-setXL 1K-10K
UDC- x_{250} ($\alpha = 50$)	7.592%	-
BQ bs16	6.160%	-
L2C-Insert ($I=1000$)	3.223%	17.494%
NeuroLKH	2.496%	-
LSSL ($C=1, B_t=4$)	1.067%	4.149% [†]
LSSL ($C=2, B_t=32$)	0.353%	2.505% [†]
LSSL ($C=4, B_t=128$)	0.171%	1.838%

Takeaway. Two complementary regimes emerge across scales. At small to medium scale, the LSSL heatmap concentrates the backend on near-optimal edges from the first trial, bypassing the geometry-driven startup that vanilla LKH and HGS rely on; the corresponding anytime curves in Fig. 3 make this explicit, with \mathcal{H} -space already at 0.28% on CVRP10K at $t \approx 10$ s while Search-Only HGS still sits at 3.44%. At large scale, purely neural approaches still fall short of the 0.000% regime on 10K-100K instances, whereas LSSL fully closes or surpasses the reference at every synthetic scale up to 100K and on all four real-world benchmarks. Both regimes share the same mechanism, the closed-loop coupling between sparse generative inference and a discrete search backend, where the learned heatmap and the heuristic search refine each other across cycles under a fixed compute budget.

3.3 Ablation studies

Setup. On the same backend (LKH3 for TSP, HGS for CVRP) and under a matched total computation time per panel, we compare four variants tied to the primitives in Section 2. **Search-Only** runs S with no learned prior. **x-space** freezes a single prediction supervised by the hard label \mathbf{x}^* projected onto \mathcal{G}_K , as in NeuroLKH [38]. **\mathcal{H} -space** likewise freezes a single prediction but with the LSSL soft target \mathcal{H}^* from Eq. (5) as supervision, i.e. LSSL at $C=1$. **LSSL** with $C \in \{2, 4, 8\}$ uses the same supervision but refreshes the prior from the current incumbent at every cycle, splitting the total




Figure 3: **Performance of LSSL on TSP/CVRP at 1K/10K.** Anytime mean-gap under matched total time per panel (32/64/32/128 s), split into C equal cycles. x -space (shown on the 1K panels) and \mathcal{H} -space (1 cycle) are the frozen single-prediction variants from the Setup paragraph; LSSL refreshes the prior every cycle.

315 budget into C equal cycles. Figure 3 reports anytime mean-gap curves on TSP/CVRP \times 1K/10K,
 316 with x -space reported on TSP1K only as a low-resource sanity case.

317 **\mathcal{H} -space vs. x -space.** Comparing x -space with \mathcal{H} -space at $C=1$ isolates the supervision form
 318 alone: both variants make one prediction before search starts and keep that guidance frozen for the
 319 entire backend budget. On TSP1K, \mathcal{H} -space reaches a terminal gap of 0.00034% vs. 0.00066% for x -
 320 space, roughly halving it; on CVRP1K, it reaches 0.686% vs. 0.974%. Replacing the single-solution
 321 hard label with the multi-solution search-feedback target therefore consistently improves one-shot
 322 guidance, but as the next paragraph shows, the dominant contribution comes from refreshing this
 323 prior rather than from softening it.

324 **Closing the loop vs. freezing the prior.** This ablation keeps the soft supervision fixed and changes
 325 only whether the prior is frozen or refreshed: \mathcal{H} -space at $C=1$ predicts once, whereas LSSL at $C=2$
 326 refreshes the guidance after the first search cycle. With prior, backend, and total computation time all
 327 held fixed, refreshing the prior once already moves terminal gap from 0.686% to 0.590% on CVRP1K,
 328 from -0.657% to -0.878% on CVRP10K, and from -0.002% to -0.007% on TSP10K. The same
 329 refresh also accelerates convergence, pushing the anytime curves into the low-gap regime earlier
 330 within the fixed budget. This directly attributes the extra improvement to closing the learning–search
 331 loop rather than to changing the target or the backend. \mathcal{H} -space on its own already removes the bulk
 332 of the Search-Only gap, but the closed loop produces the final improvement.

333 **An intermediate refresh rate works best.** At fixed total computation time, we sweep the refresh
 334 count over $C = 1, 2, 4, 8$. The results show that more feedback is not always better. On CVRP1K,
 335 the terminal gap changes from 0.686 to 0.590, 0.595, and 0.731% as C increases; on CVRP10K,
 336 it changes from -0.657 to -0.878 , -0.671 , and -0.565% . The best setting is therefore $C=2$ on
 337 TSP10K, CVRP1K, and CVRP10K alike. This pattern reflects a balance between stale guidance and
 338 premature refresh: small C leaves the prior frozen for too long, while large C gives each backend
 339 cycle too little time to exploit the current prior. CVRP10K at $C=8$ being worse than $C=1$ makes this
 340 second failure mode especially clear.

341 4 Conclusion

342 This work presents a closed-loop perspective on neural combinatorial optimization for large-scale
 343 VRP. In this perspective, learning and classical search are two endpoints of a single feedback loop,
 344 rather than two stages of a one-shot prediction-and-repair pipeline. We instantiate this perspective
 345 with LSSL, which couples a globally aware sparse backbone with a heuristic search backend through
 346 a shared structural interface. Under this interface, search consumes priors from the model, while the
 347 model is in turn supervised by the structures that search produces. Inference iterates this loop within
 348 each instance, refreshing the guidance as the search progresses. Extensive experiments on large-scale
 349 TSP and CVRP show that LSSL outperforms recent neural solvers, and under matched compute it
 350 also improves on a strong pure-search backend. These results suggest that classical search remains
 351 an essential component of large-scale VRP solving, and that neural priors are most effective when
 352 coupled with such search rather than positioned as its replacement.

353 **References**

354 [1] Bruce L Golden, S Raghavan, and Edward A Wasil. *The Vehicle Routing Problem: Latest Advances and*
355 *New Challenges*, volume 43. Springer Science & Business Media, 2008.

356 [2] Paolo Toth and Daniele Vigo. *Vehicle routing: problems, methods, and applications*. SIAM, 2014.

357 [3] Victor Pillac, Michel Gendreau, Christelle Guéret, and Andrés L Medaglia. A review of dynamic vehicle
358 routing problems. *European Journal of Operational Research*, 225(1):1–11, 2013.

359 [4] Grigorios D Konstantakopoulos, Sotiris P Gayialis, and Evripidis P Kechagias. Vehicle routing problem and
360 related algorithms for logistics distribution: A literature review and classification. *Operational research*,
361 22(3):2033–2062, 2022.

362 [5] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela,
363 and Marco Protasi. *Complexity and approximation: Combinatorial optimization problems and their*
364 *approximability properties*. Springer Science & Business Media, 2012.

365 [6] Keld Helsgaun. An extension of the lin-kernighan-helsgaun tsp solver for constrained traveling salesman
366 and vehicle routing problems. *Roskilde: Roskilde University*, 12, 2017.

367 [7] Thibaut Vidal. Hybrid genetic search for the cvrp: Open-source implementation and swap* neighborhood.
368 *Computers & Operations Research*, 140:105643, 2022.

369 [8] Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial optimization: a
370 methodological tour d’horizon. *European Journal of Operational Research*, 290(2):405–421, 2021.

371 [9] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information*
372 *Processing Systems*, volume 28, 2015.

373 [10] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial
374 optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.

375 [11] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In
376 *International Conference on Learning Representations*, 2019.

377 [12] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoon Yoon, Youngjune Gwon, and Seungjai Min. Pomo:
378 Policy optimization with multiple optima for reinforcement learning. In *Advances in Neural Information*
379 *Processing Systems*, volume 33, pages 21188–21198, 2020.

380 [13] Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with heavy
381 decoder: Toward large scale generalization. In *Advances in Neural Information Processing Systems*, 2023.

382 [14] Darko Drakulic, Sofia Michel, Florian Mai, Arnaud Sors, and Jean-Marc Andreoli. BQ-NCO: Bisimulation
383 quotienting for efficient neural combinatorial optimization. In *Advances in Neural Information Processing*
384 *Systems*, 2023.

385 [15] Yan Jin, Yuandong Ding, Xuanhao Pan, Kun He, Li Zhao, Tao Qin, Lei Song, and Jiang Bian. Pointerformer:
386 deep reinforced multi-pointer transformer for the traveling salesman problem. In *Proceedings of the AAAI*
387 *Conference on Artificial Intelligence*, 2023.

388 [16] Jianan Zhou, Yaoxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Towards omni-generalizable neural
389 methods for vehicle routing problems. In *International Conference on Machine Learning*, pages 42769–
390 42789. PMLR, 2023.

391 [17] Yang Wang, Ya-Hui Jia, Wei-Neng Chen, and Yi Mei. Distance-aware attention reshaping: Enhance
392 generalization of neural solver for large-scale vehicle routing problems. *arXiv preprint arXiv:2401.06979*,
393 2024.

394 [18] Changliang Zhou, Xi Lin, Zhenkun Wang, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang. Instance-
395 conditioned adaptation for large-scale generalization of neural combinatorial optimization. *arXiv preprint*
396 *arXiv:2405.01906*, 2024.

397 [19] Fu Luo, Xi Lin, Yaoxin Wu, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang. Boosting
398 neural combinatorial optimization for large-scale vehicle routing problems. In *The Thirteenth International*
399 *Conference on Learning Representations*, 2025.

400 [20] Yuanyao Chen, Rongsheng Chen, Fu Luo, and Zhenkun Wang. Improving generalization of neural
401 combinatorial optimization for vehicle routing problems via test-time projection learning. In *Advances in*
402 *Neural Information Processing Systems*, 2025.

- 403 [21] Fu Luo, Xi Lin, Mengyuan Zhong, Fei Liu, Zhenkun Wang, Jianyong Sun, and Qingfu Zhang. Learning
404 to insert for constructive neural vehicle routing solver. In *Advances in Neural Information Processing*
405 *Systems*, 2025.
- 406 [22] Xuanhao Pan, Yan Jin, Yuandong Ding, Mingxiao Feng, Li Zhao, Lei Song, and Jiang Bian. H-tsp:
407 Hierarchically solving the large-scale traveling salesman problem. In *Proceedings of the AAAI Conference*
408 *on Artificial Intelligence*, February 2023.
- 409 [23] Hanni Cheng, Haosi Zheng, Ya Cong, Weihao Jiang, and Shiliang Pu. Select and optimize: Learning to
410 solve large-scale tsp instances. In *International Conference on Artificial Intelligence and Statistics*, pages
411 1219–1231. PMLR, 2023.
- 412 [24] Qingchun Hou, Jingwei Yang, Yiqiang Su, Xiaoqing Wang, and Yuming Deng. Generalize learned
413 heuristics to solve large-scale vehicle routing problems in real-time. In *The Eleventh International*
414 *Conference on Learning Representations*, 2023.
- 415 [25] Haoran Ye, Jiarui Wang, Helan Liang, Zhiguang Cao, Yong Li, and Fanzhang Li. Glop: Learning global
416 partition and local construction for solving large-scale routing problems in real-time. In *Proceedings of the*
417 *AAAI Conference on Artificial Intelligence*, 2024.
- 418 [26] Zhi Zheng, Changliang Zhou, Tong Xialiang, Mingxuan Yuan, and Zhenkun Wang. UDC: A unified neural
419 divide-and-conquer framework for large-scale combinatorial optimization problems. In *Advances in Neural*
420 *Information Processing Systems*, 2024.
- 421 [27] Chengrui Gao, Haopu Shang, Ke Xue, Dong Li, and Chao Qian. Towards generalizable neural solvers for
422 vehicle routing problems via ensemble with transferrable local policy. In *International Joint Conference*
423 *on Artificial Intelligence*, 2024.
- 424 [28] Han Fang, Zhihao Song, Paul Weng, and Yutong Ban. INVit: A generalizable routing problem solver with
425 invariant nested view transformer. In *International Conference on Machine Learning*, 2024.
- 426 [29] Chaitanya K Joshi, Thomas Laurent, and Xavier Bresson. An efficient graph convolutional network
427 technique for the travelling salesman problem. *arXiv preprint arXiv:1906.01227*, 2019.
- 428 [30] Zhang-Hua Fu, Kai-Bin Qiu, and Hongyuan Zha. Generalize a small pre-trained model to arbitrarily
429 large tsp instances. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages
430 7474–7482, 2021.
- 431 [31] Ruizhong Qiu, Zhiqing Sun, and Yiming Yang. DIMES: A differentiable meta solver for combinatorial
432 optimization problems. In *Advances in Neural Information Processing Systems*, 2022.
- 433 [32] Zhiqing Sun and Yiming Yang. DIFUSCO: Graph-based diffusion solvers for combinatorial optimization.
434 In *Advances in Neural Information Processing Systems*, 2023.
- 435 [33] Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. From distribution learning in training to gradient
436 search in testing for combinatorial optimization. In *Advances in Neural Information Processing Systems*,
437 2023.
- 438 [34] Yang Li, Jinpei Guo, Runzhong Wang, Hongyuan Zha, and Junchi Yan. Fast t2t: Optimization consistency
439 speeds up diffusion-based training-to-testing solving for combinatorial optimization. In *Advances in Neural*
440 *Information Processing Systems*, 2024.
- 441 [35] Yang Li, Lvda Chen, Haonan Wang, Runzhong Wang, and Junchi Yan. Generation as search operator
442 for test-time scaling of diffusion-based combinatorial optimization. In *Advances in Neural Information*
443 *Processing Systems*, 2025.
- 444 [36] Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible regions
445 of routing problems with flexible neural k-opt. In *Advances in Neural Information Processing Systems*,
446 volume 36, 2023.
- 447 [37] André Hottung, Paula Wong-Chung, and Kevin Tierney. Neural deconstruction search for vehicle routing
448 problems. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856.
- 449 [38] Liang Xin, Wen Song, Zhiguang Cao, and Jie Zhang. Neurokh: Combining deep learning model with
450 lin-kernighan-helsgaun heuristic for solving the traveling salesman problem. In *Advances in Neural*
451 *Information Processing Systems*, volume 34, pages 7472–7483, 2021.
- 452 [39] Yifan Xia, Xianliang Yang, Zichuan Liu, Zhihao Liu, Lei Song, and Jiang Bian. Position: Rethinking post-
453 hoc search-based neural approaches for solving large-scale traveling salesman problems. In *International*
454 *Conference on Machine Learning*, 2024.

- 455 [40] André Hottung, Yeong-Dae Kwon, and Kevin Tierney. Efficient active search for combinatorial optimization
456 problems. In *International Conference on Learning Representations, 2022*.
- 457 [41] Fu Luo, Xi Lin, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang. Self-improved
458 learning for scalable neural combinatorial optimization. *arXiv preprint arXiv:2403.19561*, 2024.
- 459 [42] Jonathan Pirnay and Dominik G. Grimm. Self-improvement for neural combinatorial optimization: Sample
460 without replacement, but improvement. *Transactions on Machine Learning Research*, 2024. ISSN
461 2835-8856. Featured Certification.
- 462 [43] Thibaut Vidal. Hybrid genetic search for the cvrp: Open-source implementation and swap* neighborhood.
463 *Computers & Operations Research*, 140:105643, 2022.
- 464 [44] Thibaut Vidal, Teodor Gabriel Crainic, Michel Gendreau, Nadia Lahrichi, and Walter Rei. A hybrid genetic
465 algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60(3):611–624,
466 2012.
- 467 [45] Gerhard Reinelt. TSPLIB—a traveling salesman problem library. *ORSA Journal on Computing*, 3(4):
468 376–384, 1991.
- 469 [46] André Rohe. VLSI data sets. <https://www.math.uwaterloo.ca/tsp/vlsi/index.html>, 2014.
470 Forschungsinstitut für Diskrete Mathematik, Universität Bonn.
- 471 [47] Robert Bosch and Adrienne Herman. Continuous line drawings via the traveling salesman problem.
472 *Operations Research Letters*, 32(4):302–303, 2004.
- 473 [48] Eduardo Uchoa, Diego Pecin, Artur Pessoa, Marcus Poggi, Thibaut Vidal, and Anand Subramanian.
474 New benchmark instances for the capacitated vehicle routing problem. *European Journal of Operational
475 Research*, 257(3):845–858, 2017.
- 476 [49] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks.
477 *Advances in neural information processing systems*, 27, 2014.
- 478 [50] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning
479 to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- 480 [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
481 Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing
482 Systems*, volume 30, 2017.
- 483 [52] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization
484 algorithms over graphs. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- 485 [53] Mohammadreza Nazari, Afshin Oroojlooy, Lawrence Snyder, and Martin Takác. Reinforcement learning
486 for solving the vehicle routing problem. In *Advances in Neural Information Processing Systems*, volume 31,
487 2018.
- 488 [54] Michel Deudon, Pierre Cournut, Alexandre Lacoste, Yossiri Adulyasak, and Louis-Martin Rousseau.
489 Learning heuristics for the tsp by policy gradient. In *International conference on the integration of
490 constraint programming, artificial intelligence, and operations research*, pages 170–181. Springer, 2018.
- 491 [55] Liang Xin, Wen Song, Zhiguang Cao, and Jie Zhang. Step-wise deep learning models for solving routing
492 problems. *IEEE Transactions on Industrial Informatics*, 17(7):4861–4871, 2020.
- 493 [56] Liang Xin, Wen Song, Zhiguang Cao, and Jie Zhang. Multi-decoder attention model with embedding
494 glimpse for solving vehicle routing problems. In *Proceedings of the AAAI Conference on Artificial
495 Intelligence*, volume 35, pages 12042–12049, 2021.
- 496 [57] Wouter Kool, Herke van Hoof, Joaquim Gromicho, and Max Welling. Deep policy dynamic programming
497 for vehicle routing problems. In *Integration of Constraint Programming, Artificial Intelligence, and
498 Operations Research: 19th International Conference, CPAIOR 2022, Los Angeles, CA, USA, June 20-23,
499 2022, Proceedings*, pages 190–213. Springer, 2022.
- 500 [58] Rui Sun, Zhi Zheng, and Zhenkun Wang. Learning encodings for constructive neural combinatorial
501 optimization needs to regret. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38,
502 pages 20803–20811, 2024.
- 503 [59] Sahil Manchanda, Sofia Michel, Darko Drakulic, and Jean-Marc Andreoli. On the generalization of neural
504 combinatorial optimization heuristics. In *Joint European Conference on Machine Learning and Knowledge
505 Discovery in Databases*, pages 426–442. Springer, 2022.

- 506 [60] Zhang-Hua Fu, Sipeng Sun, Jintong Ren, Tianshu Yu, Haoyu Zhang, Yuanyuan Liu, Lingxiao Huang,
507 Xiang Yan, and Pinyan Lu. A hierarchical destroy and repair approach for solving very large-scale
508 travelling salesman problem. *arXiv preprint arXiv:2308.04639*, 2023.
- 509 [61] Yong Liang Goh, Zhiguang Cao, Yining Ma, Jianan Zhou, Mohammed Haroon Dupty, and Wee Sun Lee.
510 Shield: Multi-task multi-distribution vehicle routing solver with sparsity and hierarchy. In *International
511 Conference on Machine Learning*, 2025.
- 512 [62] Dian Meng, Zhiguang Cao, Jie Gao, Yaoxin Wu, and Yaqing Hou. Uniteformer: Unifying node and edge
513 modalities in transformers for vehicle routing problems. In *Advances in Neural Information Processing
514 Systems*, 2025.
- 515 [63] Jiwoo Son, Minsu Kim, Hyeonah Kim, and Jinkyoo Park. Meta-sage: Scale meta-learning scheduled
516 adaptation with guided exploration for mitigating scale shift on combinatorial optimization. In *International
517 Conference on Machine Learning*, pages 32194–32210. PMLR, 2023.
- 518 [64] Chaitanya K Joshi, Quentin Cappart, Louis-Martin Rousseau, and Thomas Laurent. Learning tsp requires
519 rethinking generalization. In *27th International Conference on Principles and Practice of Constraint
520 Programming (CP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- 521 [65] Chaitanya K Joshi, Quentin Cappart, Louis-Martin Rousseau, and Thomas Laurent. Learning the travelling
522 salesperson problem requires rethinking generalization. *Constraints*, 27(1-2):70–98, 2022.
- 523 [66] Paulo R d O Costa, Jason Rhuggenaath, Yingqian Zhang, and Alp Akcay. Learning 2-opt heuristics for the
524 traveling salesman problem via deep reinforcement learning. In *Asian Conference on Machine Learning*,
525 pages 465–480. PMLR, 2020.
- 526 [67] Dongxiang Zhang, Ziyang Xiao, Yuan Wang, Mingli Song, and Gang Chen. Neural tsp solver with
527 progressive distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages
528 12147–12154, 2023.
- 529 [68] Fu Luo, Yaoxin Wu, Zhi Zheng, and Zhenkun Wang. Rethinking neural combinatorial optimization for
530 vehicle routing problems with different constraint tightness degrees, 2025.
- 531 [69] Shen Lin and Brian W Kernighan. An effective heuristic algorithm for the traveling-salesman problem.
532 *Operations research*, 21(2):498–516, 1973.
- 533 [70] Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible regions
534 of routing problems with flexible neural k-opt. In *Advances in Neural Information Processing Systems*,
535 volume 36, 2024.
- 536 [71] Yaoxin Wu, Wen Song, Zhiguang Cao, Jie Zhang, and Andrew Lim. Learning improvement heuristics for
537 solving routing problems.. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- 538 [72] André Hottung, Bhanu Bhandari, and Kevin Tierney. Learning a latent search space for routing problems
539 using variational autoencoders. In *International Conference on Learning Representations*, 2021.
- 540 [73] André Hottung and Kevin Tierney. Neural large neighborhood search for the capacitated vehicle routing
541 problem. In *European Conference on Artificial Intelligence*, 2020.
- 542 [74] Arthur Delarue, Ross Anderson, and Christian Tjandraatmadja. Reinforcement learning with combinatorial
543 actions: An application to vehicle routing. *Advances in Neural Information Processing Systems*, 33:
544 609–620, 2020.
- 545 [75] Jinho Choo, Yeong-Dae Kwon, Jihoon Kim, Jeongwoo Jae, André Hottung, Kevin Tierney, and Youngjune
546 Gwon. Simulation-guided beam search for neural combinatorial optimization. In *Advances in Neural
547 Information Processing Systems*, volume 35, pages 8760–8772, 2022.
- 548 [76] Changliang Zhou, Xi Lin, Zhenkun Wang, and Qingfu Zhang. L2r: Learning to reduce search space for
549 generalizable neural routing solver, 2025.
- 550 [77] Changliang Zhou, Xi Lin, Zhenkun Wang, and Qingfu Zhang. Learning to reduce search space for
551 generalizable neural routing solver. *arXiv preprint arXiv:2503.03137*, 2025.
- 552 [78] Changliang Zhou, Canhong Yu, Shunyu Yao, Xi Lin, Zhenkun Wang, Yu Zhou, and Qingfu Zhang. Urs: A
553 unified neural routing solver for cross-problem zero-shot generalization. *arXiv preprint arXiv:2509.23413*,
554 2025.

- 555 [79] Kai Li, Fei Liu, Zhenkun Wang, Xialiang Tong, Xiongwei Han, Mingxuan Yuan, and Qingfu Zhang. Ars:
556 Automatic routing solver with large language models. *arXiv preprint arXiv:2502.15359*, 2025.
- 557 [80] Quentin Cappart, Thierry Moisan, Louis-Martin Rousseau, Isabeau Prémont-Schwarz, and Andre A
558 Cire. Combining reinforcement learning and constraint programming for combinatorial optimization. In
559 *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- 560 [81] David Applegate, Robert Bixby, Vašek Chvátal, and William Cook. Concorde tsp solver, 2006.
- 561 [82] Éric D Taillard and Keld Helsgaun. Popmusic for the travelling salesman problem. *European Journal of*
562 *Operational Research*, 272(2):420–429, 2019.

563 **A Related work**

564 Neural combinatorial optimization (NCO) for VRPs differs along three axes that matter for LSSL:
565 how solutions or scores are produced, how computation scales to large N , and how the learned model
566 couples with classical search. We organise prior work along these axes; Table 5 contrasts LSSL with
567 the most closely related families.

568 **A.1 Scalable neural construction**

569 Constructive NCO, building on the sequence-to-sequence [49] and attention [50, 51] lineage adapted
570 to combinatorial outputs, autoregressively appends or inserts nodes to produce a feasible route.
571 Pointer Networks [9], REINFORCE-trained policies [10, 52–54], the Attention Model [11], step-
572 wise constructive variants [55, 56], POMO [12], and policy dynamic-programming hybrids [57]
573 establish the family; BQ-NCO [14], LEHD [13], Pointerformer [15], regret-aware encodings [58],
574 Boosting NCO [19] and L2C-Insert [21] strengthen decoding or scale supervised training, while
575 Manchanda et al. [59] characterise the out-of-distribution failure modes that motivate scale-aware
576 training. Three strategies extend this regime to ultra-large instances. Decomposition splits an instance
577 into subproblems with a learned partition or revision stage [22, 25, 26, 23, 24, 60]. Cross-scale
578 generalisation keeps a single end-to-end model and strengthens the inductive bias needed for scale
579 transfer [16, 17, 27, 28, 61–63]. Linear-complexity construction amortises a constructive solver to
580 100K nodes [19, 21]. Across this family the model parameterises the policy or its decoded route
581 directly; classical search, when present, runs once downstream, so search-discovered global structure
582 has no shared interface to feed back into the predictor, and the supervisory signal collapses any
583 plurality of high-quality solutions into a single argmax target.

584 **A.2 Heatmap, diffusion, and generative NCO**

585 Non-constructive NCO predicts edge heatmaps or tour distributions and decodes a feasible solution
586 post hoc. Early supervised GNN heatmaps decoded by beam search or 2-opt [29, 64, 65], learned
587 2-opt operators [66], Att-GCN+MCTS [30], and DIMES [31] initiated the line, with progressive-
588 distillation [67] variants extending the recipe to harder instances. Diffusion- and flow-based variants
589 are now dominant: DIFUSCO [32] casts edge selection as graph diffusion, T2T [33] adds objective-
590 guided test-time gradient search, Fast-T2T [34] compresses the trajectory via optimisation consistency,
591 and GenSCO [35] reframes generation as a search operator alternating disruption with rectified-flow
592 regeneration. Constraint-tightness analyses argue that single-shot heatmap decoders break when
593 feasibility manifolds are stiff [68], and Xia et al. [39] question whether post-hoc search alone scales to
594 very large TSPs. This is LSSL’s closest neural precursor (scoring and feasibility recovery are already
595 decoupled), but the time-conditioned state is a perturbed solution variable trained to converge to a
596 single reference x^* . Downstream search refines a sample without being looped back as a persistent
597 state, and training never aggregates multiple search outputs into a single supervisory target.

598 **A.3 Neural-guided search and one-sided learning–search loops**

599 A third line couples learning with classical search but closes only one side of the loop. The strongest
600 non-neural baselines remain Lin–Kernighan [69, 6] and HGS [7, 43]. NeuroLKH [38] is the most
601 direct precedent: a sparse GNN predicts candidate edges and node penalties consumed once by
602 LKH. Improvement-operator methods learn a policy that edits an existing tour [36, 70, 71, 37] or, in
603 latent-action variants, sample edits from a learned search space [72]; large-neighbourhood search is
604 similarly amenable to learning the destroy-and-repair operator [73, 74]. Search amplifiers such as
605 EAS [40] and SGBS [75] operate inside a constructive policy. On the training side, self-improvement
606 pipelines feed iteratively improved solutions back as relabelled data [42, 41, 19]. On the inference
607 side, Test-Time Projection Learning [20], instance-conditioned adaptation [18], and search-space-
608 reduction methods [76–79] adapt the predictor or prune the search space at test time. Generic
609 learning–search frameworks [8, 80] provide broader context. Each of these mechanisms closes at
610 most one side of the learning–search loop (training-time relabelling or inference-time adaptation),
611 but not both, and the interface is typically tied to a specific backend rather than the backend-agnostic
612 $N \times K$ heatmap that LSSL exchanges.

613 **A.4 Positioning summary**

614 LSSL differs from the families above on the axes summarised in Table 5. The exchanged object is a
 615 sparse, not-necessarily-feasible structural prior \mathcal{H}_t on an $O(NK)$ candidate graph rather than a route,
 616 a perturbed solution variable, or a one-shot LKH candidate set. The same heatmap interface is invoked
 617 in multiple inference cycles and aggregates row-wise soft targets from multiple backend-refined route
 618 sets at training, closing both sides of the loop with a backend-agnostic design instantiated here with
 619 LKH for TSP and HGS for CVRP.

Table 5: Comparison of LSSL with closely related method families along the main axes of learning–search coupling and scalability. “partial” indicates the property holds at moderate scale but not at the family’s largest tractable N under a 24 GB GPU budget.

Family (representatives)	Learned object	Search backend	Backend cycles	Search→training	Sparse $O(NK)$	Max N on 24 GB GPU
NeuroLKH [38]	candidate edges / π for LKH	LKH	×	×	partial	5K
Diffusion / generative NCO [32, 34, 35]	perturbed solution \mathbf{x}_t	2-opt / MCTS (post hoc)	×	×	partial	1K
Self-improvement [19, 42, 41]	constructed routes	none	×	✓ (relabel)	×	100K
Test-time / search-space reduction [20, 76, 18]	adapted projection / pruning	none	partial	×	partial	100K
Improvement operators [36, 37, 40, 75]	local repair policy	own learned	×	×	×	1K
LSSL (ours)	infeasible structural prior \mathcal{H}_t	backend-agnostic (LKH, HGS used)	✓	✓ (multi-soln soft)	✓	100K

620 **B Supplementary experiments**

621 **B.1 Real-world TSP results on TSPLIB**

622 Table 6 reports supplementary real-world TSPLIB TSP results. The DIFUSCO and Fast T2T baseline
 623 columns are transcribed from Tables 7 and 8 of [35]; the added L2C-Insert column comes from our
 624 local evaluation results, while the NeuroLKH column is transcribed from the supplementary TSPLIB
 tables of NeuroLKH [38].

Table 6: Results on real-world TSPLIB TSP instances with 50–1K nodes.

Instances	DIFUSCO	Fast T2T	GenSCO	L2C-Insert	NeuroLKH	LSSL Gap (Time)	LSSL Gap (Time)
						($C=1, B_{\text{trials}}=128$)	($C=2, B_{\text{trials}}=256$)
ei51	2.82%	0.00%	0.000%	0.000%	0.000%	0.000% (0.0s)	0.000% (0.2s)
berlin52	0.00%	0.00%	0.000%	1.584%	0.000%	0.000% (0.0s)	0.000% (0.0s)
st70	0.00%	0.00%	0.000%	0.000%	0.000%	0.000% (0.0s)	0.000% (0.2s)
ei76	0.34%	0.00%	0.000%	0.004%	0.000%	0.000% (0.1s)	0.000% (0.2s)
pr76	1.12%	0.00%	0.000%	0.000%	0.000%	0.000% (0.1s)	0.000% (0.4s)
rat99	0.09%	0.00%	0.000%	12.416%	0.000%	0.000% (0.1s)	0.000% (0.2s)
kroA100	0.10%	0.00%	0.000%	5.013%	0.000%	0.000% (0.1s)	0.000% (0.3s)
kroB100	2.29%	0.65%	0.000%	5.300%	0.000%	0.000% (0.2s)	0.000% (0.6s)
kroC100	0.00%	0.00%	0.000%	6.076%	0.000%	0.000% (0.1s)	0.000% (0.3s)
kroD100	0.07%	0.00%	0.000%	4.815%	0.000%	0.000% (0.1s)	0.000% (0.3s)
kroE100	3.83%	0.00%	0.000%	4.050%	0.000%	0.000% (0.1s)	0.000% (0.5s)
rd100	0.08%	0.00%	0.000%	0.000%	0.000%	0.000% (0.1s)	0.000% (0.3s)
ei101	0.03%	0.00%	0.000%	0.000%	0.000%	0.000% (0.0s)	0.000% (0.1s)
lin105	0.00%	0.00%	0.000%	4.337%	0.000%	0.000% (0.1s)	0.000% (0.2s)
pr107	0.91%	0.62%	0.000%	0.000%	0.000%	0.000% (0.2s)	0.000% (0.7s)
pr124	1.02%	0.08%	0.000%	0.000%	0.000%	0.000% (0.3s)	0.000% (1.0s)
bier127	0.94%	1.50%	0.000%	0.000%	0.016%	0.000% (0.1s)	0.000% (0.3s)
ch130	0.29%	0.00%	0.000%	0.000%	0.000%	0.000% (0.2s)	0.000% (0.6s)
pr136	0.19%	0.01%	0.000%	0.000%	0.000%	0.000% (0.5s)	0.000% (1.9s)
pr144	0.80%	0.39%	0.000%	0.000%	0.081%	0.089% (0.7s)	0.000% (2.0s)
ch150	0.57%	0.00%	0.000%	0.000%	0.000%	0.000% (0.4s)	0.000% (1.5s)
kroA150	0.34%	0.00%	0.000%	3.220%	0.000%	0.000% (0.3s)	0.000% (1.1s)
kroB150	0.30%	0.07%	0.340%	3.846%	0.000%	0.000% (0.5s)	0.000% (1.8s)
pr152	1.69%	0.19%	0.187%	0.005%	0.037%	0.000% (1.0s)	0.000% (3.4s)
u159	0.82%	0.00%	0.000%	0.025%	0.000%	0.000% (0.1s)	0.000% (0.2s)
rat195	1.48%	0.79%	0.194%	4.828%	0.000%	0.000% (0.4s)	0.000% (1.4s)
d198	3.32%	0.86%	0.751%	1.835%	0.285%	0.000% (3.6s)	0.000% (13.2s)
kroA200	2.28%	0.49%	0.160%	1.025%	0.000%	0.000% (0.3s)	0.000% (1.2s)
kroB200	2.35%	2.50%	0.098%	0.219%	0.000%	0.000% (0.2s)	0.000% (0.8s)
ts225	4.95%	1.37%	0.517%	0.000%	0.000%	0.000% (0.4s)	0.000% (1.7s)
tsp225	3.25%	0.81%	0.000%	2.860%	0.000%	0.000% (0.2s)	0.000% (0.6s)
pr226	4.22%	0.34%	0.211%	0.290%	0.016%	0.000% (0.4s)	0.000% (1.8s)
gil262	2.18%	0.18%	0.000%	0.089%	0.000%	0.000% (0.6s)	0.000% (2.0s)
pr264	0.92%	0.73%	0.000%	0.000%	0.000%	0.000% (0.6s)	0.000% (2.0s)
a280	1.39%	0.10%	0.000%	0.109%	0.000%	0.000% (0.4s)	0.000% (1.4s)
pr299	1.46%	1.40%	0.011%	2.390%	0.000%	0.000% (1.6s)	0.000% (4.6s)
lin318	2.95%	1.21%	0.025%	0.038%	0.000%	0.000% (1.1s)	0.000% (4.1s)
rd400	1.18%	0.08%	0.030%	0.162%	0.000%	0.000% (0.8s)	0.000% (2.6s)
fl417	3.30%	2.01%	0.596%	0.408%	0.056%	0.000% (11.8s)	0.000% (44.1s)
pr439	2.73%	0.50%	0.106%	0.387%	0.047%	0.000% (1.4s)	0.000% (6.0s)
pcb442	2.59%	0.61%	0.015%	0.598%	0.000%	0.000% (0.7s)	0.000% (2.7s)
d493	1.81%	1.43%	3.872%	0.469%	0.086%	0.000% (2.8s)	0.000% (10.3s)
u574	2.50%	0.94%	0.109%	0.630%	0.000%	0.000% (1.2s)	0.000% (4.4s)
rat575	2.32%	1.43%	0.000%	0.825%	0.001%	0.000% (1.1s)	0.000% (3.9s)
p654	7.49%	1.67%	1.023%	1.815%	0.354%	0.000% (19.6s)	0.000% (1.23m)
d657	4.86%	0.64%	0.457%	0.751%	0.001%	0.000% (1.5s)	0.000% (5.7s)
u724	2.05%	1.41%	0.040%	0.555%	0.000%	0.000% (1.4s)	0.000% (7.9s)
rat783	3.04%	1.03%	0.068%	0.995%	0.000%	0.000% (1.2s)	0.000% (3.5s)
Mean	1.735%	0.542%	0.184%	1.499%	0.020%	0.002% (1.2s)	0.000% (4.5s)

625

626 **B.2 Real-world TSP results on VLSI**

627 Tables 7 and 8 report supplementary real-world VLSI TSP results, both transcribed from the cor-
 628 responding real-world result summary file. The former covers instances with fewer than 2K nodes
 629 and reports GenSCO, L2C-Insert, NeuroLKH, BQ bs16, UDC-x₂₅₀ ($\alpha=50$), GLOP, LSSL Gap, and
 LSSL Time. The latter covers instances with 2K–200K nodes and reports L2C-Insert, LSSL

Table 7: Results on real-world VLSI TSP instances with fewer than 2K nodes.

Instance	GenSCO	L2C-Insert	NeuroLKH	BQ bs16	UDC-x ₂₅₀	GLOP	LSSL Gap (Time)	LSSL Gap (Time)
							(C=1, B _{trials} =128)	(C=2, B _{trials} =256)
xqf131	-0.138%	4.765%	-0.061%	-0.138%	-0.138%	1.494%	-0.138% (0.3s)	-0.138% (1.1s)
xqg237	0.043%	1.289%	-0.029%	-0.129%	-0.013%	1.783%	-0.183% (0.9s)	-0.183% (2.8s)
pma343	0.611%	15.229%	0.045%	2.934%	5.424%	2.167%	-0.175% (6.4s)	-0.175% (23.3s)
pka379	1.771%	19.875%	-0.039%	4.344%	5.434%	2.703%	-0.191% (5.8s)	-0.191% (21.2s)
bcl380	-0.193%	-0.134%	-0.006%	1.148%	3.840%	2.289%	-0.193% (1.6s)	-0.193% (2.5s)
pbl395	-0.364%	-0.356%	-0.141%	1.166%	2.286%	3.100%	-0.364% (0.7s)	-0.364% (3.0s)
pbk411	-0.460%	-0.376%	-0.365%	0.155%	0.754%	1.792%	-0.469% (1.5s)	-0.469% (5.4s)
pbn423	-0.383%	-0.340%	-0.030%	-0.023%	1.433%	2.977%	-0.394% (0.9s)	-0.394% (4.8s)
pbm436	-0.364%	0.198%	-0.045%	1.629%	3.754%	2.019%	-0.364% (2.9s)	-0.364% (6.9s)
xql662	-0.113%	0.028%	-0.097%	2.216%	2.471%	5.447%	-0.185% (4.2s)	-0.185% (15.6s)
rbx711	-0.129%	0.948%	-0.041%	2.346%	2.758%	5.247%	-0.174% (1.7s)	-0.174% (5.3s)
rbu737	0.223%	0.763%	0.034%	2.673%	1.466%	3.149%	-0.131% (1.9s)	-0.131% (6.7s)
dkg813	0.819%	0.666%	-0.086%	4.328%	6.442%	3.185%	-0.286% (4.0s)	-0.286% (12.8s)
lim963	-0.055%	0.785%	-0.016%	2.068%	2.251%	3.981%	-0.262% (2.7s)	-0.262% (12.6s)
pbd984	-0.225%	0.378%	-0.133%	2.413%	3.202%	4.463%	-0.309% (3.3s)	-0.340% (38.2s)
xit1083	-0.226%	0.674%	-0.005%	20.926%	1.618%	4.469%	-0.250% (5.5s)	-0.250% (19.9s)
dka1376	-0.123%	1.695%	-0.040%	3.455%	5.575%	6.088%	-0.170% (2.3s)	-0.170% (6.8s)
dca1389	0.411%	0.539%	-0.084%	8.239%	6.715%	4.553%	-0.258% (3.7s)	-0.195% (32.5s)
dja1436	0.437%	1.008%	0.045%	9.670%	5.111%	5.448%	-0.147% (5.4s)	-0.153% (15.5s)
icw1483	0.226%	1.245%	-0.077%	7.366%	8.725%	6.291%	-0.396% (3.0s)	-0.417% (9.9s)
fra1488	3.397%	0.637%	-0.161%	9.941%	7.098%	2.633%	-0.404% (3.8s)	-0.404% (12.7s)
rbv1583	0.598%	1.168%	0.013%	6.495%	3.695%	4.109%	-0.205% (11.0s)	-0.205% (25.4s)
rby1599	0.660%	0.967%	-0.094%	7.018%	4.133%	5.691%	-0.284% (5.4s)	-0.314% (14.7s)
fnb1615	10.120%	1.039%	-0.015%	11.194%	4.767%	4.808%	-0.278% (14.9s)	-0.278% (44.6s)
djc1785	2.269%	0.932%	-0.106%	5.945%	3.412%	6.520%	-0.253% (8.1s)	-0.264% (24.2s)
dec1911	0.761%	1.102%	-0.044%	8.892%	7.471%	5.441%	-0.200% (8.1s)	-0.226% (21.0s)
dkd1973	23.299%	0.653%	-0.078%	42.724%	6.744%	2.949%	-0.479% (15.1s)	-0.479% (45.3s)
Mean	1.588%	2.051%	-0.061%	6.259%	3.942%	3.881%	-0.264% (4.6s)	-0.267% (16.1s)

Table 8: Results on real-world VLSI TSP instances with 2K–200K nodes.

Instance	L2C-Insert	GLOP	LSSL Gap (Time)	LSSL Gap (Time)	Instance	L2C-Insert	GLOP	LSSL Gap (Time)	LSSL Gap (Time)
			(C=1, B _{trials} =128)	(C=2, B _{trials} =256)				(C=1, B _{trials} =128)	(C=2, B _{trials} =256)
djb2036	0.491%	6.445%	-0.383% (5.7s)	-0.387% (14.0s)	ido21215	3.486%	7.834%	-0.234% (3.75m)	-0.275% (10.20m)
dcb2086	1.069%	7.290%	-0.309% (3.7s)	-0.315% (14.8s)	fma21553	3.786%	8.634%	-0.162% (4.09m)	-0.230% (9.87m)
bva2144	0.854%	6.945%	-0.271% (3.9s)	-0.271% (10.7s)	lsb22777	4.204%	8.580%	-0.279% (4.42m)	-0.313% (12.44m)
xqe2175	1.589%	5.864%	-0.283% (4.9s)	-0.403% (14.7s)	xrh24104	5.330%	8.566%	-0.237% (6.23m)	-0.252% (13.37m)
bck2217	0.745%	6.051%	-0.400% (11.8s)	-0.407% (28.1s)	bbz25234	4.908%	7.387%	-0.314% (6.57m)	-0.320% (13.62m)
xpr2308	1.213%	8.387%	-0.222% (5.7s)	-0.295% (18.5s)	irx28268	4.258%	7.847%	-0.355% (9.64m)	-0.367% (18.09m)
ley2323	1.163%	8.366%	-0.210% (5.6s)	-0.246% (22.0s)	fyg28534	7.719%	10.375%	-0.274% (8.66m)	-0.288% (17.93m)
irw2802	1.547%	8.162%	-0.259% (6.9s)	-0.263% (20.0s)	icx28698	6.638%	9.307%	-0.271% (9.49m)	-0.295% (18.50m)
dbj2924	1.359%	7.232%	-0.125% (7.3s)	-0.125% (21.3s)	boa28924	6.409%	8.829%	-0.306% (9.73m)	-0.313% (20.25m)
pia3056	2.217%	6.826%	-0.378% (13.4s)	-0.383% (34.2s)	ird29514	5.461%	9.085%	-0.249% (8.64m)	-0.259% (21.18m)
xqe3891	1.246%	7.697%	-0.097% (9.6s)	-0.136% (30.5s)	pbb30440	4.929%	8.667%	-0.231% (9.62m)	-0.259% (20.66m)
bgb4355	1.453%	7.045%	-0.279% (13.4s)	-0.299% (40.7s)	xib32892	5.454%	8.935%	-0.231% (10.82m)	-0.246% (24.55m)
xsce6880	2.041%	8.126%	-0.317% (28.0s)	-0.361% (1.25m)	fry33203	7.195%	9.281%	-0.168% (14.87m)	-0.197% (27.18m)
lap7454	3.013%	8.450%	-0.344% (32.8s)	-0.354% (1.39m)	bby34656	6.315%	8.080%	-0.281% (14.18m)	-0.317% (34.10m)
ida8197	2.111%	8.127%	-0.262% (39.2s)	-0.297% (1.71m)	pba38478	6.531%	9.031%	-0.254% (17.54m)	-0.265% (37.21m)
xmc10150	2.449%	8.732%	-0.279% (52.7s)	-0.293% (3.61m)	ics39603	9.478%	9.699%	-0.262% (19.71m)	-0.319% (37.93m)
xvb13584	2.344%	6.654%	-0.300% (1.60m)	-0.307% (4.58m)	rbz43748	7.483%	8.412%	-0.282% (21.69m)	-0.323% (49.80m)
xrb14233	3.248%	9.089%	-0.184% (1.83m)	-0.198% (5.85m)	fht47608	8.644%	8.226%	-0.233% (27.85m)	-0.266% (54.20m)
xial6928	3.752%	7.383%	-0.244% (2.48m)	-0.262% (7.03m)	fna52057	7.334%	8.245%	-0.245% (35.21m)	-0.283% (1.13h)
pjh17845	3.007%	6.974%	-0.307% (2.69m)	-0.330% (7.70m)	bna56769	6.761%	7.740%	-0.298% (37.84m)	-0.320% (1.32h)
frh19289	3.965%	8.007%	-0.233% (3.07m)	-0.259% (10.91m)	dan59296	9.102%	8.465%	-0.229% (41.57m)	-0.263% (1.36h)
fncl9402	3.816%	9.094%	-0.211% (3.17m)	-0.260% (8.29m)	sral04815	6.114%	6.089%	-0.718% (1.18h)	-0.733% (1.67h)
Mean	4.142%	8.051%	-0.273% (9.36m)	-0.299% (18.78m)					

631 **B.3 Real-world TSP results on Art-TSP**

Table 9 reports results on the 6 large-scale real-world TSP instances from Art-TSP.

Table 9: Results on large-scale real-world TSP instances from Art-TSP.

Instance	Random Insertion	L2C-Insert	GLOP	LSSL Gap (Time)	
				($C=1, B_{\text{trials}}=128$)	($C=2, B_{\text{trials}}=256$)
mona-lisa100K	6.706%	3.764%	3.224%	0.034% (0.33h)	0.026% (1.37h)
vangogh120K	6.812%	4.022%	3.156%	0.038% (0.49h)	0.028% (1.91h)
venus140K	6.581%	4.102%	3.067%	0.027% (0.62h)	0.023% (2.54h)
pareja160K	6.830%	4.259%	3.158%	0.032% (0.87h)	0.024% (3.50h)
courbet180K	7.145%	4.777%	3.294%	0.035% (1.04h)	0.025% (4.00h)
earring200K	7.605%	4.514%	3.528%	0.039% (1.32h)	0.029% (5.18h)
Mean	6.947%	4.240%	3.238%	0.034% (0.78h)	0.026% (3.08h)

632

633 **B.4 Real-world CVRP results on CVRPLIB**

634 Tables 10 and 11 report supplementary real-world CVRPLIB results on the set-X (100 instances)
 635 and set-XL (100 instances) collections. The L2C-Insert numbers use $I=1000$ runs over all instances.

Table 10: Results on real-world CVRPLIB set-X CVRP instances.

Instance	L2C-Insert	Neuro LKH	BQ bs16	UDC-X ₂₅₀	LSSL Gap			Instance	L2C-Insert	Neuro LKH	BQ bs16	UDC-X ₂₅₀	LSSL Gap		
					($C=1, B_l=4$)	($C=2, B_l=32$)	($C=4, B_l=128$)						($C=1, B_l=4$)	($C=2, B_l=32$)	($C=4, B_l=128$)
X-n101-k25	2.55%	0.72%	11.240%	20.47%	0.00%	0.00%	0.00%	X-n336-k84	3.18%	3.11%	10.430%	9.35%	1.56%	0.65%	0.30%
X-n106-k14	0.75%	0.63%	2.121%	5.174%	0.99%	0.00%	0.00%	X-n344-k43	2.69%	2.30%	4.040%	7.42%	1.38%	0.19%	0.10%
X-n110-k13	0.04%	0.69%	0.849%	6.153%	0.00%	0.00%	0.00%	X-n351-k40	3.90%	4.02%	5.826%	7.13%	1.87%	0.76%	0.41%
X-n115-k10	0.62%	0.28%	11.065%	14.397%	0.00%	0.00%	0.00%	X-n359-k29	2.68%	3.65%	2.572%	3.57%	1.65%	0.66%	0.28%
X-n120-k6	0.038%	0.433%	0.077%	9.361%	0.00%	0.00%	0.00%	X-n367-k17	4.20%	2.37%	5.353%	6.70%	1.30%	0.05%	0.00%
X-n125-k30	2.90%	0.79%	12.173%	7.907%	0.18%	0.00%	0.00%	X-n376-k94	0.96%	0.560%	0.810%	5.65%	0.18%	0.05%	0.01%
X-n129-k18	0.92%	0.98%	1.833%	4.828%	0.00%	0.00%	0.00%	X-n384-k52	3.71%	4.35%	2.989%	7.94%	1.33%	0.88%	0.18%
X-n134-k13	0.708%	1.178%	5.152%	6.082%	0.19%	0.00%	0.00%	X-n393-k38	3.46%	2.18%	4.451%	7.75%	1.29%	0.43%	0.12%
X-n139-k10	0.127%	0.829%	1.398%	4.994%	0.00%	0.00%	0.00%	X-n401-k29	2.01%	1.36%	5.900%	4.53%	0.98%	0.27%	0.02%
X-n143-k7	1.91%	1.184%	1.860%	6.668%	0.00%	0.00%	0.00%	X-n411-k19	4.92%	2.12%	8.834%	8.33%	1.61%	0.12%	0.08%
X-n148-k46	1.78%	1.46%	13.26%	8.25%	0.18%	0.00%	0.00%	X-n420-k30	5.89%	3.68%	16.931%	7.96%	1.36%	0.57%	0.19%
X-n153-k22	3.81%	0.42%	25.813%	13.921%	0.60%	0.04%	0.02%	X-n429-k61	3.52%	2.46%	3.117%	8.01%	1.35%	0.32%	0.27%
X-n157-k13	0.588%	0.171%	2.829%	3.847%	0.00%	0.00%	0.00%	X-n439-k37	1.45%	2.16%	1.769%	9.20%	0.74%	0.05%	0.03%
X-n162-k11	0.721%	0.759%	1.633%	6.509%	0.00%	0.00%	0.00%	X-n449-k29	4.07%	6.11%	2.700%	4.96%	2.10%	0.67%	0.24%
X-n167-k10	1.69%	1.46%	1.811%	9.228%	0.18%	0.00%	0.00%	X-n459-k26	5.62%	3.74%	5.558%	5.28%	1.47%	0.46%	0.06%
X-n172-k51	2.11%	0.850%	11.831%	9.825%	0.41%	0.00%	0.00%	X-n469-k138	5.82%	4.69%	8.496%	9.30%	1.04%	0.59%	0.29%
X-n176-k26	4.53%	0.978%	20.355%	10.967%	0.37%	0.09%	0.00%	X-n480-k70	2.97%	1.65%	3.226%	6.06%	1.37%	0.50%	0.20%
X-n181-k23	0.088%	0.457%	1.170%	5.649%	0.14%	0.00%	0.00%	X-n491-k59	3.67%	3.37%	4.857%	9.38%	2.11%	1.00%	0.51%
X-n186-k15	0.70%	1.56%	2.546%	7.575%	0.04%	0.00%	0.00%	X-n502-k39	1.55%	0.57%	1.887%	4.19%	0.39%	0.18%	0.08%
X-n190-k48	1.37%	1.648%	3.097%	2.108%	0.48%	0.12%	0.00%	X-n513-k21	3.52%	2.60%	4.325%	4.34%	1.54%	0.29%	0.14%
X-n195-k51	1.718%	1.047%	16.678%	14.228%	0.49%	0.10%	0.03%	X-n524-k153	5.99%	2.59%	19.326%	9.44%	1.85%	0.36%	0.18%
X-n200-k36	2.96%	0.89%	6.738%	6.132%	0.15%	0.00%	0.00%	X-n536-k96	4.88%	4.94%	6.333%	7.70%	1.36%	0.49%	0.25%
X-n204-k19	1.47%	1.49%	3.916%	8.995%	0.19%	0.02%	0.00%	X-n548-k50	3.08%	1.20%	0.665%	5.93%	0.70%	0.07%	0.03%
X-n209-k16	1.67%	1.57%	1.728%	7.135%	0.15%	0.08%	0.00%	X-n561-k42	3.67%	2.80%	3.268%	4.26%	1.14%	0.48%	0.27%
X-n214-k11	2.40%	5.916%	4.394%	5.749%	0.59%	0.16%	0.04%	X-n573-k30	3.32%	1.92%	8.491%	4.87%	1.26%	0.62%	0.45%
X-n219-k73	0.715%	0.118%	5.239%	5.902%	0.01%	0.00%	0.00%	X-n586-k159	4.87%	3.26%	7.126%	9.20%	1.14%	0.32%	0.14%
X-n223-k34	1.93%	1.44%	3.567%	6.045%	0.38%	0.08%	0.02%	X-n599-k92	4.81%	12.39%	3.140%	8.89%	1.75%	0.57%	0.24%
X-n228-k23	4.17%	1.12%	22.614%	8.700%	0.61%	0.25%	0.04%	X-n613-k62	4.56%	4.90%	4.744%	8.45%	2.39%	0.76%	0.36%
X-n233-k16	1.748%	1.911%	2.670%	10.573%	0.59%	0.02%	0.00%	X-n627-k43	4.78%	3.30%	6.620%	7.10%	1.80%	0.85%	0.51%
X-n237-k14	1.730%	1.334%	1.120%	9.834%	0.25%	0.00%	0.00%	X-n641-k35	6.30%	3.08%	2.684%	8.72%	1.81%	0.61%	0.30%
X-n242-k48	2.01%	1.20%	2.976%	5.774%	0.85%	0.31%	0.07%	X-n655-k131	2.75%	0.90%	3.801%	4.28%	0.29%	0.11%	0.05%
X-n247-k50	5.31%	0.97%	23.825%	11.497%	0.82%	0.54%	0.11%	X-n670-k130	8.70%	6.48%	25.411%	11.56%	3.49%	0.67%	0.51%
X-n251-k28	1.86%	1.48%	4.115%	6.474%	0.79%	0.30%	0.00%	X-n685-k75	5.10%	3.84%	5.379%	9.49%	1.98%	0.52%	0.21%
X-n256-k16	1.23%	3.158%	4.254%	7.759%	0.21%	0.08%	0.00%	X-n701-k44	4.42%	2.85%	3.272%	3.56%	2.08%	0.96%	0.59%
X-n261-k13	1.74%	2.82%	2.430%	4.783%	0.24%	0.08%	0.00%	X-n716-k35	5.50%	3.96%	7.018%	7.05%	2.79%	0.74%	0.38%
X-n266-k58	2.44%	1.46%	3.22%	8.724%	1.06%	0.44%	0.12%	X-n733-k159	3.85%	2.23%	10.624%	7.38%	1.78%	0.70%	0.29%
X-n270-k35	1.610%	1.492%	3.494%	6.961%	0.49%	0.09%	0.04%	X-n749-k98	5.21%	4.33%	6.981%	7.86%	2.56%	1.17%	0.84%
X-n275-k28	0.95%	1.25%	1.958%	6.474%	0.52%	0.00%	0.00%	X-n766-k71	7.54%	3.37%	12.310%	8.19%	2.39%	0.79%	0.54%
X-n280-k17	3.18%	2.01%	7.707%	6.153%	0.90%	0.03%	0.00%	X-n783-k48	5.46%	4.27%	3.018%	7.50%	2.61%	0.94%	0.44%
X-n284-k15	2.69%	3.08%	3.708%	3.156%	1.27%	0.41%	0.23%	X-n801-k40	5.25%	3.08%	2.019%	12.45%	1.93%	0.60%	0.25%
X-n289-k60	2.99%	1.94%	7.40%	9.944%	1.03%	0.47%	0.13%	X-n819-k71	6.22%	4.42%	7.579%	7.42%	1.64%	0.77%	0.43%
X-n294-k50	2.34%	1.59%	6.054%	9.851%	0.17%	0.14%	0.11%	X-n837-k142	4.27%	2.26%	2.797%	7.45%	1.54%	0.61%	0.24%
X-n298-k31	1.87%	2.57%	1.794%	5.816%	0.64%	0.09%	0.00%	X-n856-k95	3.42%	2.17%	4.429%	6.33%	0.90%	0.23%	0.12%
X-n303-k21	2.40%	2.06%	3.755%	7.892%	0.73%	0.23%	0.04%	X-n876-k59	4.78%	3.63%	5.382%	4.48%	2.20%	1.06%	0.48%
X-n308-k13	1.85%	1.35%	4.046%	6.712%	1.03%	0.16%	0.02%	X-n895-k37	7.10%	9.62%	5.113%	9.00%	3.39%	1.17%	0.61%
X-n313-k71	3.754%	2.794%	7.849%	7.409%	1.18%	0.33%	0.04%	X-n916-k207	4.59%	3.09%	5.222%	7.50%	1.20%	0.60%	0.26%
X-n317-k53	0.784%	0.149%	2.662%	3.595%	0.17%	0.02%	0.00%	X-n936-k151	10.39%	9.45%	25.758%	11.79%	3.71%	1.87%	0.82%
X-n322-k28	2.88%	2.67%	3.370%	9.205%	0.67%	0.35%	0.04%	X-n957-k87	4.17%	2.24%	3.788%	7.78%	1.18%	0.54%	0.30%
X-n327-k20	3.32%	3.10%	3.436%	9.048%	1.01%	0.52%	0.15%	X-n979-k58	5.21%	2.61%	4.560%	5.98%	2.82%	0.80%	0.34%
X-n331-k15	1.99%	1.484%	1.909%	10.577%	0.83%	0.07%	0.03%	X-n1001-k43	5.96%	4.58%	4.209%	4.94%	1.07%	1.32%	0.59%
Mean	3.22%	2.49%	6.160%	7.592%	1.06%	0.35%	0.17%								

636

Table 11: Results on real-world CVRPLIB set-XL CVRP instances.

Instance	L2C-Insert Gap	L2C-Insert time	LSSL Gap (C=1, B _i =4)	LSSL Gap (C=2, B _i =32)	LSSL Gap (C=4, B _i =128)	Instance	L2C-Insert Gap	L2C-Insert time	LSSL Gap (C=1, B _i =4)	LSSL Gap (C=2, B _i =32)	LSSL Gap (C=4, B _i =128)
XL-n1048-k237	12.464%	657.17s	-	3.248%	1.914%	XL-n3408-k524	8.497%	2451.82s	2.237%	1.633%	1.415%
XL-n1094-k157	5.244%	680.48s	0.630%	0.369%	0.276%	XL-n3484-k436	4.041%	2420.72s	0.430%	0.262%	0.218%
XL-n1141-k112	11.177%	664.97s	1.494%	1.034%	0.636%	XL-n3561-k229	15.092%	2372.08s	4.240%	3.161%	2.171%
XL-n1188-k96	9.826%	668.40s	2.200%	1.183%	0.936%	XL-n3640-k211	18.468%	2385.89s	4.090%	2.823%	2.007%
XL-n1234-k55	9.397%	693.84s	2.871%	1.997%	1.486%	XL-n3721-k77	22.581%	2483.03s	5.279%	3.754%	2.076%
XL-n1281-k29	17.019%	698.57s	3.646%	1.990%	1.621%	XL-n3804-k29	36.146%	2593.52s	8.042%	5.228%	2.749%
XL-n1328-k19	22.963%	749.11s	4.633%	1.924%	0.899%	XL-n3888-k1010	11.575%	3326.79s	-	-	2.733%
XL-n1374-k278	10.076%	766.16s	1.928%	0.841%	0.841%	XL-n3975-k687	8.484%	3435.63s	2.271%	1.887%	1.506%
XL-n1421-k232	5.262%	790.87s	1.431%	1.253%	0.903%	XL-n4063-k347	10.341%	3243.75s	2.638%	1.935%	1.562%
XL-n1468-k151	7.438%	791.09s	1.490%	0.961%	0.684%	XL-n4153-k291	11.321%	3296.78s	2.396%	1.996%	1.415%
XL-n1514-k106	12.103%	797.49s	2.354%	1.180%	0.732%	XL-n4245-k203	23.355%	3320.98s	2.517%	1.435%	1.153%
XL-n1561-k75	13.303%	814.39s	1.649%	1.032%	0.795%	XL-n4340-k148	22.302%	3330.50s	6.047%	3.095%	2.453%
XL-n1608-k39	21.031%	803.25s	3.357%	1.718%	1.037%	XL-n4436-k48	39.119%	3896.42s	6.651%	4.263%	2.941%
XL-n1654-k11	19.225%	824.63s	7.646%	2.682%	1.333%	XL-n4535-k1134	6.746%	4288.59s	0.229%	0.167%	0.128%
XL-n1701-k562	14.417%	988.37s	1.515%	1.095%	1.095%	XL-n4635-k790	9.644%	4505.06s	2.777%	2.266%	2.122%
XL-n1748-k271	11.071%	1000.83s	2.142%	1.629%	1.307%	XL-n4738-k487	9.018%	4499.19s	2.194%	1.804%	1.327%
XL-n1794-k163	8.338%	918.05s	1.340%	0.777%	0.617%	XL-n4844-k321	11.555%	4467.05s	4.448%	3.097%	2.557%
XL-n1841-k126	9.424%	909.82s	2.903%	2.064%	1.633%	XL-n4951-k203	20.809%	4576.57s	5.696%	3.212%	2.221%
XL-n1888-k82	12.208%	954.52s	3.094%	2.175%	1.539%	XL-n5061-k184	30.736%	4812.85s	5.630%	3.480%	2.323%
XL-n1934-k46	18.939%	970.06s	4.029%	2.845%	2.194%	XL-n5174-k55	49.057%	5003.24s	11.200%	7.484%	5.526%
XL-n1981-k13	33.021%	1081.15s	6.154%	3.072%	2.020%	XL-n5288-k1246	10.234%	6199.65s	1.870%	1.053%	0.965%
XL-n2028-k617	12.630%	1127.05s	1.927%	1.185%	1.027%	XL-n5406-k783	5.849%	6475.50s	2.059%	1.507%	1.238%
XL-n2074-k264	6.360%	1150.32s	2.575%	2.178%	1.910%	XL-n5526-k553	12.169%	6109.14s	1.544%	0.998%	0.668%
XL-n2121-k186	7.456%	1144.95s	3.397%	1.636%	1.249%	XL-n5649-k401	12.063%	6107.04s	8.870%	3.495%	2.631%
XL-n2168-k138	13.529%	1170.37s	3.470%	2.328%	1.626%	XL-n5774-k290	22.527%	6181.67s	6.601%	4.240%	3.115%
XL-n2214-k131	11.267%	1186.25s	1.546%	0.836%	0.561%	XL-n5902-k122	29.987%	6373.03s	8.686%	5.191%	3.602%
XL-n2261-k54	14.323%	1186.05s	4.990%	3.523%	2.430%	XL-n6034-k61	51.171%	7074.99s	10.486%	7.228%	5.681%
XL-n2307-k34	21.714%	1203.58s	5.978%	3.119%	2.129%	XL-n6168-k1922	12.197%	8910.71s	1.870%	2.717%	1.686%
XL-n2354-k631	9.937%	1451.36s	1.116%	0.669%	0.469%	XL-n6305-k1042	8.921%	9297.80s	2.475%	1.670%	1.409%
XL-n2401-k408	9.309%	1484.36s	2.667%	2.192%	1.546%	XL-n6445-k628	8.995%	9019.53s	2.195%	1.388%	1.168%
XL-n2447-k290	14.727%	1428.11s	4.416%	1.994%	1.414%	XL-n6588-k473	18.703%	8934.79s	5.019%	3.351%	2.772%
XL-n2494-k194	7.813%	1416.22s	1.981%	1.247%	1.002%	XL-n6734-k330	19.550%	8968.28s	4.083%	2.503%	1.693%
XL-n2541-k121	18.555%	1393.17s	2.547%	1.401%	1.125%	XL-n6884-k148	38.515%	9268.05s	6.788%	4.688%	3.545%
XL-n2587-k66	15.555%	1398.07s	5.122%	3.138%	2.124%	XL-n7037-k38	65.663%	10034.09s	10.397%	8.444%	5.015%
XL-n2634-k17	43.241%	1518.16s	10.063%	6.618%	5.683%	XL-n7193-k1683	9.048%	12958.66s	2.262%	0.677%	0.555%
XL-n2681-k540	6.156%	1605.13s	2.996%	1.493%	1.349%	XL-n7353-k1471	5.214%	13772.80s	0.515%	0.330%	0.266%
XL-n2727-k546	6.721%	1749.10s	0.611%	0.411%	0.315%	XL-n7516-k859	10.655%	13359.74s	3.397%	2.278%	1.998%
XL-n2774-k286	7.854%	1735.59s	1.861%	1.117%	0.896%	XL-n7683-k602	14.891%	13321.26s	2.549%	1.719%	1.432%
XL-n2821-k208	11.590%	1613.14s	4.007%	2.489%	2.103%	XL-n7854-k365	12.304%	13259.23s	-	1.360%	1.044%
XL-n2867-k130	13.502%	1631.60s	3.645%	1.976%	1.469%	XL-n8028-k294	28.905%	13917.90s	6.700%	4.599%	3.520%
XL-n2914-k95	16.882%	1704.10s	5.334%	3.402%	2.132%	XL-n8207-k108	67.731%	15057.08s	8.474%	5.488%	4.539%
XL-n2961-k55	19.949%	1734.12s	7.598%	4.440%	3.150%	XL-n8389-k2028	6.546%	18839.21s	-	1.887%	-
XL-n3007-k658	8.106%	2078.38s	2.109%	1.658%	1.107%	XL-n8575-k1297	8.681%	19603.71s	2.971%	1.966%	1.751%
XL-n3054-k461	7.083%	2132.77s	1.506%	1.107%	0.808%	XL-n8766-k1032	10.699%	19641.99s	4.882%	1.683%	1.405%
XL-n3101-k311	10.965%	2046.84s	3.688%	1.964%	1.338%	XL-n8960-k634	12.830%	19361.49s	1.921%	1.651%	1.329%
XL-n3147-k232	14.902%	2032.63s	2.583%	1.899%	1.450%	XL-n9160-k379	29.890%	19941.54s	7.983%	5.090%	3.738%
XL-n3194-k161	16.659%	1997.66s	3.874%	2.292%	1.677%	XL-n9363-k209	43.011%	20924.73s	6.987%	4.667%	2.985%
XL-n3241-k115	20.739%	2032.68s	2.812%	1.617%	1.114%	XL-n9571-k55	83.047%	22431.75s	29.430%	10.432%	7.478%
XL-n3287-k30	42.752%	2236.79s	7.440%	4.502%	2.936%	XL-n9784-k2774	7.775%	29594.32s	-	0.904%	0.771%
XL-n3334-k934	7.518%	2384.53s	1.357%	1.157%	0.892%	XL-n10001-k1570	11.069%	23881.49s	2.007%	1.133%	0.804%
Mean	17.494%	5337.3s	4.149%	2.505%	1.838%						

637 **B.5 Scaling of inference and training cost with instance size**

638 We profile inference (4-step Euler flow-matching) and training (single-step forward+backward) wall-
 639 clock and GPU peak memory for the LSSL backbone D_θ , validating the near-linear-in- N claim of
 640 Section 2. The architecture matches the main experiments; we compare sparse $N \times K$ ($K=1000$)
 641 against a dense $N \times N$ baseline at batch size 1, with 1000 instances per scale.

Table 12: Inference (INF) and training (TR) wall-clock and GPU peak memory at batch size 1. Sparse uses $K=1000$; dense uses full $N \times N$ attention with the same backbone. ‘‘OOM’’ means out-of-memory on a 24 GB GPU.

N	K	Mode	INF total (s)	INF/iter (ms)	INF peak (MB)	TR total (s)	TR/step (ms)	TR peak (MB)	Notes
1,000	999	sparse	22.84	22.84	159	57.25	57.25	323	LSSL baseline
1,000	—	dense	4.36	4.36	143	14.22	14.22	1,105	dense fits at small N
10,000	1000	sparse	209.80	209.80	510	517.00	517.00	870	LSSL baseline
10,000	—	dense	337.76	337.76	4,100	—	—	OOM	TR needs ~ 37 GB
100,000	1000	sparse	2,414.17	2,414.17	3,338	6,224.82	6,224.82	6,051	LSSL baseline
100,000	—	dense	—	—	—	—	—	OOM	infeasible (bias ≥ 40 GB)

642 **B.6 Allocation of inference budget between refinement cycles and per-cycle search**

643 We probe how a fixed wall-budget B should be split between the number of LSSL refinement cycles
 644 C and the per-cycle LKH time $\tau=B/C$. All runs use the canonical Guided-LKH configuration
 645 (Appendix F.2) on the 16-instance TSP10000 test set, $K=1000$, 8-way LKH parallelism, identical
 646 seed. Mean gap% vs. ground-truth tours is reported in Table 13; negative values mean LSSL beat the
 reference.

Table 13: Mean gap% on TSP10000 (16 instances) for total wall-budget $B = C \cdot \tau$. Columns:
 refinement cycles C (per-cycle time $\tau=B/C$).

B (s)	$C=1$	$C=2$	$C=4$	$C=8$	$C=16$
8	+0.0329	-	-	-	-
16	+0.0152	+0.0059	-	-	-
32	+0.0065	-0.0029	-0.0021	-	-
64	-0.0016	-0.0067	-0.0066	-0.0048	-
128	-0.0054	-0.0083	-0.0094	-0.0101	-0.0097

647

648 **C Experimental Details**

649 **C.1 Hyperparameters**

650 Let L^{enc} and L^{dec} denote the number of encoder and decoder layers, d_{enc} and d_{dec} their hidden
 651 widths, H the number of attention heads (shared by encoder and decoder), $d_h = d_{\bullet}/H$ the per-head
 652 dimension, and $d_{\text{ffn}} = 3d_{\bullet}$ the gated feed-forward hidden width. For TSP we use $L^{\text{enc}} = 16$, $L^{\text{dec}} = 6$,
 653 $d_{\text{enc}} = d_{\text{dec}} = 256$, and $H = 8$. For CVRP, which carries additional depot, demand, and capacity
 654 features, we widen the backbone to $d_{\text{enc}} = d_{\text{dec}} = 384$ with $H = 12$ attention heads, keeping $L^{\text{enc}} = 16$
 655 and $L^{\text{dec}} = 6$. The same backbone is shared across all instance sizes within each task. For instances
 656 with $N > 1000$, the sparse neighborhood size is fixed to $K = 1000$.

657 Optimization uses Muon for all backbones and all instance sizes, with weight decay 1×10^{-4} , global-
 658 norm gradient clipping at 1.0, peak learning rate 1×10^{-3} , linear warmup over the first $\sim 1\%$ of
 659 update steps, and cosine decay to 1×10^{-5} . The closed-loop schedule uses $M = 3$ search-refined
 660 route sets per instance to construct the row-wise soft target \mathcal{H}^* , matching the parallelism of S at
 661 training time.

662 **C.2 Training-label construction**

663 Training labels are constructed in two regimes. At the small scales (TSP100, TSP1K, CVRP100),
 664 each label is the output of a single direct solver call, using LKH3 for TSP and HGS for CVRP. At the
 665 larger scales (TSP10K, TSP100K, CVRP1K, CVRP10K), where uniform direct labelling exceeds
 666 our offline budget, labels are constructed through the self-bootstrapping mechanism of Section 2.5:
 667 the smaller-scale-trained LSSL checkpoint $\theta^{(s)}$ produces guidance heatmaps on freshly sampled
 668 target-scale instances, the search backend S solves each instance under its standard parameter file,
 669 and the resulting M route sets are projected and aggregated via Eq. (5) into the search-feedback
 670 target \mathcal{H}^* , which is stored as the fixed training label. Table 14 summarises the per-scale recipe.

671 Self-improved training has been used to extend supervision for large-scale NCO. The SIT frame-
 672 work of Luo et al. [19] casts training as a multi-round iteration: at round r the current model $\theta^{(r)}$
 673 reconstructs a tour on each training instance, those tours are recorded as the pseudo-labels for round
 674 $r+1$, and $\theta^{(r+1)}$ is trained on this updated supervision; by repeatedly weaving the model’s own
 675 outputs back into the supervision signal, a single architecture is pushed past the level of its initial
 676 fixed-label set on TSP and CVRP. The label-construction recipe of this section inherits the same
 677 high-level idea (letting the model participate in constructing its own supervision) and instantiates
 678 it as a cross-scale bootstrap: the guidance signal $\theta^{(s)}$ comes from an LSSL checkpoint trained at a
 679 smaller scale, and every target-scale label is the actual solver output produced by the search backend
 680 S under that guidance. The wall-clock cost of label construction is folded into the per-scale training
 681 budget reported in Appendix C.

Table 14: Per-scale training-label recipe. “Direct” denotes a single solver call; “self-bootstrap” uses the smaller-scale checkpoint $\theta^{(s')}$ as guidance for the search backend S . The aggregation count M matches the search parallelism used at training time (cf. Section 2.5). “Rounds r ” is the number of multi-round iterations of label refresh (Section 2.5); direct rows do not iterate.

Scale	Label source	Guidance $\theta^{(s')}$	Search backend S	M	Rounds r	# instances
TSP100	direct	–	Concorde	–	–	1.28M
TSP1K	direct	–	LKH3	3	–	64K
TSP10K	self-bootstrap	$\theta^{(TSP1K)}$	LKH3	3	2	6.4K
TSP100K	self-bootstrap	$\theta^{(TSP10K)}$	LKH3	3	2	640
CVRP100	direct	–	HGS	3	–	1.28M
CVRP1K	self-bootstrap	$\theta^{(CVRP100)}$	HGS	3	2	64K
CVRP10K	self-bootstrap	$\theta^{(CVRP1K)}$	HGS	3	2	6.4K

682 C.3 Hardware

683 All training and inference results reported in the main paper were carried out on a single workstation:
 684 one NVIDIA GeForce RTX 4090 GPU (24 GB), two AMD EPYC 7K62 CPU (48 cores), and 128 GB
 685 of DDR4 system memory.

686 C.4 Computational Resources

687 All evaluations are run on a single NVIDIA RTX 4090 GPU paired with two AMD EPYC 7K62 CPU.
 688 On the same RTX 4090, end-to-end LSSL training takes ≤ 2 days / ≤ 1 day / ≤ 1 day for TSP-100 /
 689 TSP-1K / TSP-10K respectively; the TSP-10K/TSP-100K self-bootstrapping take ≤ 2 days and ≤ 1
 690 days respectively.

691 For CVRP, training takes ≤ 3 days / ≤ 2 days / ≤ 2 days for CVRP-100 / CVRP-1K / CVRP-10K,
 692 and CVRP-1K / CVRP-10K self-bootstrapping take ≤ 3 days and ≤ 3 days respectively.

693 For comparison, prior heatmap baselines already require this range or more for training alone on
 694 stronger hardware: DIFUSCO [32] reports 8.6 and 5.1 days for TSP-100 and TSP-1K training on
 695 a single A100 GPU, whereas the budgets above already cover both training and self-bootstrapped
 696 label construction on a weaker consumer-grade RTX 4090, placing LSSL’s total time in a comparable
 697 range.

698 C.5 Implementation details of baselines

699 This appendix details only those baselines that appear in Tables 1, 2, 3, and 4. We organise them
 700 into five families. **(1) Classical solvers:** Concorde [81], LKH3 [6], and HGS [43]. **(2) Constructive**
 701 **NCO methods:** POMO [12], BQ [14], LEHD [13], INVIT [28], and the three L2C-Insert variants
 702 (Greedy, $I=1000$, $I=10000$) [21]. **(3) Heatmap-based NCO methods:** Att-GCN+MCTS [30],
 703 DIFUSCO [32], T2T [33], Fast-T2T [34], and GenSCO [35]. **(4) Two-stage methods:** H-TSP [22],
 704 GLOP [25], and UDC [26]; Tables 1 and 2 use UDC- x_{50} , while Tables 3 and 4 use UDC- x_{250} . **(5)**
 705 **Neural-guided classical search:** NeuroLKH [38], reported in Tables 3 and 4. Methods that never
 706 appear in these tables are not discussed here.

- 707 • **Concorde** [81]: we use the PyConcorde Python wrapper at its default settings on TSP
 708 instances of every size.
- 709 • **LKH3** [6]: we use the runner shipped with AM [11]. On TSP100K we fol-
 710 low Fu et al. [60] and switch the preprocessing stage to POPMUSIC [82] (CANDI-
 711 DATE_SET_TYPE=POPMUSIC, INITIAL_PERIOD=1000) to keep the per-instance
 712 cost manageable; all other parameters are kept at their defaults.
- 713 • **HGS** [43]: we use the PyHygese Python wrapper of the HGS-CVRP solver.
- 714 • **LEHD** [13]: we cap the RRC maximum destruction size at 1000. We adopted this cap after
 715 verifying that lifting it leaves the solution quality essentially unchanged while substantially
 716 increasing inference time.

- 717 • **L2C-Insert** [21]: we report the greedy decoder together with the local-reconstruction
718 variants at $I=1000$ and $I=10000$; the corresponding rows are labelled “L2C-Insert Greedy”,
719 “L2C-Insert ($I=1000$)”, and “L2C-Insert ($I=10000$)”. The $I=10000$ setting is reported only
720 on TSP100K (Table 2).
- 721 • **UDC** [26]: we keep the variant naming of the original paper. The synthetic-instance main
722 tables use UDC- x_{50} ($\alpha=50$), and the real-world main tables use UDC- x_{250} ($\alpha=50$).
- 723 • **DIFUSCO** [32]: we re-run the official categorical-diffusion checkpoint with $T_s=50$ denois-
724 ing sampling steps, decode each instance by greedy heatmap decoding, and refine with 2-opt
725 local search.
- 726 • **Fast-T2T** [34]: we re-run the official consistency-trained checkpoint with $T_s=5$ denoising
727 sampling steps coupled with $T_g=5$ objective-guided rewrite cycles, generate candidate tours
728 by stochastic heatmap decoding, and refine with 2-opt local search.
- 729 • **GenSCO** [35]: on the synthetic TSP main table (Table 1) we re-run the official implemen-
730 tation under $C=10$ and $C=160$, both with the bundled 2-opt post-processing, and label
731 the corresponding rows “GenSCO ($C=10$) 2Opt” and “GenSCO ($C=160$) 2Opt”; both
732 configurations run out of memory at TSP10K.
- 733 • **NeuroLKH** [38]: the TSPLIB column is transcribed from the supplementary TSPLIB tables
734 of Xin et al. [38], and the VLSI column follows the real-world TSP summary used for
735 Table 7. For CVRPLIB Set-X, we combine the available supplementary NeuroLKH entries
736 for the smaller instances with our local rerun on the larger Set-X instances. The rerun uses
737 the released NeuroLKH checkpoints and 1024 LKH trials, a 7200-second per-instance time
738 limit.

739 D Network architecture details

740 Our backbone is inspired by the use of Transformer-based generative models for combinatorial opti-
741 mization in [35], but differs in being tailored to LSSL’s closed-loop search interface: all intermediate
742 representations are backend-facing heatmaps on the sparse K -NN graph rather than dense pairwise
743 node logits. This appendix gives the full architecture of the sparse diffusion Transformer D_θ (cf.
744 Section 2.4). All operators map $N \times K$ tensors on the K -NN subgraph \mathcal{G}_K to $N \times K$ tensors of the
745 same shape; no $N \times N$ object is ever materialised, so per-layer compute and memory stay at $O(NK)$
746 for any candidate width $K \ll N$.

747 D.1 Per-node features and slot indexing

748 TSP and CVRP share the same sparse candidate graph \mathcal{G}_K . We treat $\mathcal{N}_K(i)$ as an ordered K -tuple of
749 slots with the slot-to-node alias

$$j_{ik} := \mathcal{N}_K(i)[k] \in \mathcal{V}, \quad k \in \{1, \dots, K\}. \quad (12)$$

750 Per-node features are $\mathbf{s}_i = [\hat{x}_i, \hat{y}_i]$ for TSP and $[\hat{x}_i, \hat{y}_i, \rho_i]$ for CVRP, with $\rho_i = q_i/Q$ on customers
751 (using the same q_i for customer demand as in Appendix G) and $\rho_0 = -1$ on the depot; coordinates
752 (\hat{x}_i, \hat{y}_i) are centred and isotropically rescaled by the mean node-wise ℓ_2 norm.

753 D.2 Heatmap construction at training time

754 The decoder consumes two conditioning channels, each obtained by projecting two endpoints onto \mathcal{G}_K
755 via the operator $P(\cdot; \mathcal{G}_K)$ of Section 2.2. The *edge channel* uses the multi-solution search-feedback
756 target $\mathcal{H}^* = P(\{\mathcal{R}^{(m)}\}_{m=1}^M; \mathcal{G}_K)$ of Eq. (5) (the row-normalised edge-usage frequency over the M
757 retained route sets) together with an independent random route set \mathcal{R}_ϵ on the same instance, written
758 as $\mathcal{H}_\epsilon = P(\mathcal{R}_\epsilon; \mathcal{G}_K)$ (no perturbation of any $\mathcal{R}^{(m)}$ is involved). The *node-selection channel* (CVRP
759 only) uses the route-boundary frequency $\mathcal{H}_{\text{node}}^*$ and a matched random boundary indicator $\mathcal{H}_{\text{node}, \epsilon}$.
760 Both channels are interpolated independently by the same noise level t :

$$\mathcal{H}_t = (1-t)\mathcal{H}^* + t\mathcal{H}_\epsilon, \quad \mathcal{H}_{\text{node}, t} = (1-t)\mathcal{H}_{\text{node}}^* + t\mathcal{H}_{\text{node}, \epsilon}, \quad t \sim \mathcal{U}(0, 1). \quad (13)$$

761 The two channels are kept separate: \mathcal{H}^* supervises the row-stochastic edge head F_{out} via the row-
762 wise sparse cross-entropy of Eq. (8), while $\mathcal{H}_{\text{node}}^*$ supervises the per-customer node-selection head

763 via the binary cross-entropy of Eq. (9); the two losses are different functionals on different decoder
764 heads and are not collapsed into a single tensor at the input either. For TSP, $\mathcal{H}_{\text{node}}^* = \mathcal{H}_{\text{node},\epsilon} = 0$
765 so $\mathcal{H}_{\text{node},t} \equiv 0$, only the edge channel is active, and Eq. (13) reduces to a pure linear interpolation
766 between \mathcal{H}^* and \mathcal{H}_ϵ , recovering Eq. (7).

767 **CVRP boundary structure.** The node-selection heatmaps $\mathcal{H}_{\text{node}}^*, \mathcal{H}_{\text{node},\epsilon} \in [0, 1]^{N_c}$ are per-
768 customer scalars and live on a separate channel from the edge tensor \mathcal{H}^* : $\mathcal{H}_{\text{node}}^*$ is the route-boundary
769 frequency over the same M route sets $\{\mathcal{R}^{(m)}\}_{m=1}^M$ used for \mathcal{H}^* (matching the supervision target of
770 Eq. (9)), and $\mathcal{H}_{\text{node},\epsilon}$ is sampled by drawing customers from a small pool of those closest to the depot.
771 The decoder pairs the edge head F_{out} with a separate per-customer head producing $\mathcal{H}_{\text{node}} \in [0, 1]^{N_c}$.

772 D.3 Sparse diffusion Transformer backbone

773 **Macro structure.** $D_\theta(\mathcal{G}_K, \mathcal{H}_t, t)$ is a stack of five $N \times K$ operators:

$$774 \underbrace{F_{\text{in}}}_{\text{input projection}} \longrightarrow \underbrace{(F_{\text{enc}}^{(\ell)})_{\ell=1}^{L^{\text{enc}}}}_{\text{sparse encoder}} \longrightarrow \underbrace{F_{\text{mid}} \oplus \tau(t)}_{\text{time bridge}} \longrightarrow \underbrace{(F_{\text{dec}}^{(\ell)})_{\ell=1}^{L^{\text{dec}}}}_{\text{time-conditioned decoder}} \longrightarrow \underbrace{F_{\text{out}}}_{\text{edge head}}. \quad (14)$$

774 $F_{\text{in}}: \mathbb{R}^{d_s} \rightarrow \mathbb{R}^{d_{\text{enc}}}$ is a per-node affine map ($d_s = 2$ TSP, $d_s = 3$ CVRP). The encoder propagates $\{s_i\}$
775 on \mathcal{G}_K to instance-aware node states; the decoder additionally consumes \mathcal{H}_t and t to inject which
776 candidate edges are currently believed to belong to a high-quality solution and how informative that
777 belief is. The edge head F_{out} produces row-stochastic $\mathcal{H}_{\text{guide}} \in \mathbb{R}_+^{N \times K}$ in the same shape as \mathcal{H}_t .

778 **Time embedding.** The noise level t is encoded by a sinusoidal basis followed by a two-layer MLP,
779 with frequency width d_{freq} , max period $T_{\text{max}} = 10000$, and timestep scale T_{scale} :

$$780 \tau(t) = W_2^T \text{SiLU}(W_1^T [\cos(\omega t), \sin(\omega t)]) \in \mathbb{R}^{d_{\text{dec}}}, \quad (15)$$

$$781 \omega_k = \exp\left(-\frac{2k}{d_{\text{freq}}} \log T_{\text{max}}\right), \quad \bar{t} = T_{\text{scale}} t.$$

780 $\tau(t)$ is injected additively at the decoder entrance: $\mathbf{h}_i^{(0,\text{dec})} = F_{\text{mid}} \mathbf{h}_i^{(L^{\text{enc}}, \text{enc})} + \tau(t)$, where
781 $F_{\text{mid}}: \mathbb{R}^{d_{\text{enc}}} \rightarrow \mathbb{R}^{d_{\text{dec}}}$ aligns encoder/decoder widths.

782 **Pre-norm Transformer block.** Each encoder/decoder layer is a pre-norm Transformer block on
783 the $N \times K$ tensor:

$$784 \tilde{\mathbf{h}}_i^{(\ell)} = \mathbf{h}_i^{(\ell-1)} + \text{SparseMHA}^{(\ell)}(\text{LN}(\mathbf{h}_i^{(\ell-1)}); \mathcal{G}_K, \mathbf{b}_i^{(\ell)}), \quad (16)$$

$$785 \mathbf{h}_i^{(\ell)} = \tilde{\mathbf{h}}_i^{(\ell)} + \text{SwiGLU}^{(\ell)}(\text{LN}(\tilde{\mathbf{h}}_i^{(\ell)})), \quad (17)$$

785 with bias-free LayerNorm, the SparseMHA operator detailed in Section D.4, and the position-wise
786 gated FFN

$$787 \text{SwiGLU}^{(\ell)}(\mathbf{u}) = (\text{SiLU}(W_1^{(\ell)} \mathbf{u} + \mathbf{b}_1^{(\ell)}) \odot (W_2^{(\ell)} \mathbf{u} + \mathbf{b}_2^{(\ell)})) W_3^{(\ell)} + \mathbf{b}_3^{(\ell)}, \quad (18)$$

787 of hidden width $d_{\text{fn}} = 3d$ acting independently on each row i . Optionally, query and key activations
788 are normalised by a parameter-free RMSNorm on the head dimension before the dot product.

789 D.4 Sparse multi-head attention forward operator

790 For H heads of width d_h ($d = Hd_h$), with $\bar{\mathbf{h}}_\bullet^{(\ell)} = \text{LN}(\mathbf{h}_\bullet^{(\ell-1)})$, the per-head projections at layer ℓ ,
791 head h , row i , candidate $j \in \mathcal{N}_K(i)$ are

$$792 \mathbf{q}_i^{(\ell,h)} = W_Q^{(\ell,h)} \bar{\mathbf{h}}_i^{(\ell)}, \quad \mathbf{k}_j^{(\ell,h)} = W_K^{(\ell,h)} \bar{\mathbf{h}}_j^{(\ell)}, \quad \mathbf{v}_j^{(\ell,h)} = W_V^{(\ell,h)} \bar{\mathbf{h}}_j^{(\ell)}. \quad (19)$$

792 The softmax-input score combines the bilinear term with a per-row bias:

$$e_{ij}^{(\ell,h)} = \frac{\langle \mathbf{q}_i^{(\ell,h)}, \mathbf{k}_j^{(\ell,h)} \rangle}{\sqrt{d_h}} + b_{ij}^{(\ell)}, \quad (20)$$

793 with $b^{(\ell)} \in \mathbb{R}^{N \times K}$ per instance (broadcast across batch and heads). Consistent with the main-text
 794 decoder bias of Eq. (3), in the decoder $b_{ij}^{(\ell)}$ fuses the per-edge heatmap entry with a sinusoidal time
 795 embedding through learned scalar projections $u^{(\ell)}, v^{(\ell)}$, while in the encoder it vanishes:

$$b_{ij}^{(\ell)} = u^{(\ell)}([\mathcal{H}_t]_{ij}) + v^{(\ell)}(\phi(t)) \quad (\text{decoder}); \quad b_{ij}^{(\ell)} \equiv 0 \quad (\text{encoder}). \quad (21)$$

796 Here $\phi(t)$ is a sinusoidal embedding of t , and $u^{(\ell)}, v^{(\ell)}$ are layer-wise learned scalar linear maps; the
 797 second term is broadcast across (i, j) within a layer, so $b^{(\ell)}$ remains $N \times K$ per instance. Self-loops
 798 are excluded by construction since $\mathcal{N}_K(i)$ is a strict K -NN list. CVRP-specific structure (route bound-
 799 aries) enters attention through both channels of Eq. (13): the edge tensor \mathcal{H}_t contributes $u^{(\ell)}([\mathcal{H}_t]_{ij})$
 800 on every (i, j) slot, while the per-customer node-selection scalar $[\mathcal{H}_{\text{node}, t}]_i$ of Section D.2 contributes
 801 an additional learned scalar projection broadcast across all K slots of customer row i (the depot row
 802 receives no node-selection bias), so the per-edge edge channel and the per-row node channel act on
 803 different axes without being summed at the input.

804 **Block-streaming forward.** We never materialize the row-wise softmax over $\mathcal{N}_K(i)$. Instead, the
 805 operator visits each row’s neighborhood in a sequence of disjoint blocks $\mathcal{B}_1, \dots, \mathcal{B}_{N_b}$ that together
 806 cover $\mathcal{N}_K(i)$, and maintains an exact online-softmax statistic. After block b the maximum and
 807 partition terms are

$$m_i^{(\ell, h, b)} = \max\left(m_i^{(\ell, h, b-1)}, \max_{j \in \mathcal{B}_b} e_{ij}^{(\ell, h)}\right), \quad m_i^{(\ell, h, 0)} = -\infty, \quad (22)$$

$$\lambda_i^{(\ell, h, b)} = \lambda_i^{(\ell, h, b-1)} \exp(m_i^{(\ell, h, b-1)} - m_i^{(\ell, h, b)}) + \sum_{j \in \mathcal{B}_b} \exp(e_{ij}^{(\ell, h)} - m_i^{(\ell, h, b)}), \quad \lambda_i^{(\ell, h, 0)} = 0, \quad (23)$$

809 and the running value aggregation $\mathbf{o}_i^{(\ell, h, b)}$ obeys the same rebasing rule,

$$\mathbf{o}_i^{(\ell, h, b)} = \frac{\lambda_i^{(\ell, h, b-1)} \exp(m_i^{(\ell, h, b-1)} - m_i^{(\ell, h, b)})}{\lambda_i^{(\ell, h, b)}} \mathbf{o}_i^{(\ell, h, b-1)} + \frac{1}{\lambda_i^{(\ell, h, b)}} \sum_{j \in \mathcal{B}_b} \exp(e_{ij}^{(\ell, h)} - m_i^{(\ell, h, b)}) \mathbf{v}_j^{(\ell, h)}. \quad (24)$$

810 After the final block N_b , the head output is $\mathbf{o}_i^{(\ell, h)} = \mathbf{o}_i^{(\ell, h, N_b)}$ and the row-wise normalized attention
 811 weight is

$$\alpha_{ij}^{(\ell, h)} = \frac{\exp(e_{ij}^{(\ell, h)} - m_i^{(\ell, h, N_b)})}{\lambda_i^{(\ell, h, N_b)}}, \quad \sum_{j \in \mathcal{N}_K(i)} \alpha_{ij}^{(\ell, h)} = 1. \quad (25)$$

812 The heads are concatenated and projected by $W_O^{(\ell)} \in \mathbb{R}^{d \times d}$ to obtain the layer’s attention output.
 813 By construction the operator only ever accesses the gathered tensors $\mathbf{k}_{\mathcal{N}_K(i)}^{(\ell, h)}, \mathbf{v}_{\mathcal{N}_K(i)}^{(\ell, h)}, b_{i, \mathcal{N}_K(i)}^{(\ell)}$ for a
 814 single row i at a time, so its working set scales as $O(K d_h H)$ (independent of $|\mathcal{V}|$) and its compute
 815 as $O(|\mathcal{V}| K d)$ per layer, well below the dense $O(|\mathcal{V}|^2 d)$.

816 D.5 Edge-scoring head and row-wise sparse cross-entropy

817 Let \mathbf{g}_i be the projected decoder state at row i after the final norm. The edge head scores only the
 818 candidate edges of \mathcal{G}_K (logit scale γ):

$$[\mathcal{H}_{\text{guide}}]_{ij} = \frac{\exp(\gamma \langle \mathbf{g}_i, \mathbf{g}_j \rangle)}{\sum_{j' \in \mathcal{N}_K(i)} \exp(\gamma \langle \mathbf{g}_i, \mathbf{g}_{j'} \rangle)}, \quad j \in \mathcal{N}_K(i), \quad (26)$$

819 with padding slots masked before normalisation; the scoring reuses the decoder’s $N \times K$ sparse
 820 index, never materialising a dense $|\mathcal{V}| \times |\mathcal{V}|$ matrix.

821 Consistent with the main-text definition (Eq. (5)–(6)), the row-wise soft target is the multi-solution
 822 edge-usage frequency over the M retained route sets $\{\mathcal{R}^{(m)}\}_{m=1}^M$ projected onto \mathcal{G}_K :

$$\mathcal{H}_{ij}^* = \frac{\sum_{m=1}^M \mathbb{1}\{(i, j) \in \mathcal{R}^{(m)}\}}{\sum_{j' \in \mathcal{N}_K(i)} \sum_{m=1}^M \mathbb{1}\{(i, j') \in \mathcal{R}^{(m)}\}}, \quad j \in \mathcal{N}_K(i), \quad (27)$$

823 so $\mathcal{H}_{ij}^* > 0$ iff j is a tour-neighbour of i in at least one $\mathcal{R}^{(m)}$, and edges hit by more route sets receive
824 proportionally larger mass; for CVRP this also covers $j=0$ on route-boundary customers, since their
825 depot adjacency is counted by the same indicator. The LSSL training objective is the row-weighted
826 sparse cross-entropy

$$\mathcal{L}_{\text{LSSL}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_K(i)} \mathcal{H}_{ij}^* \log[\mathcal{H}_{\text{guide}}]_{ij}, \quad (28)$$

827 E Analysis of the computational complexity of LSSL

828 The complexity of LSSL has three sources: the sparse neural backbone, the projection/heatmap
829 interface, and the heuristic search backend. All three are restricted to the same $N \times K$ sparse
830 adjacency, where $K \ll N$ is the candidate width per node, so scaling along N stays linear rather
831 than $O(N^2)$. Throughout this section, N is the number of nodes, K the candidate width per row,
832 d the hidden dimension, L the number of sparse attention layers, S_{denoise} the number of denoising
833 calls along the noise schedule, C the number of learning-search cycles, M the number of refined
834 route sets retained per training instance, and B the per-cycle search budget (instantiated as B_{trials} ,
835 the LKH trial count, for TSP, and B_t , the HGS time budget, for CVRP).

836 **Sparse backbone and supervision cost** For a single network call, sparse attention, structural-bias
837 injection, candidate-edge scoring, and the row-wise softmax are all performed only over $\mathcal{N}_K(i)$,
838 while the per-node FFN and projections have fixed per-node width d . The dominant time complexity
839 is therefore

$$T_{D_\theta} = O(LN(Kd + d^2)),$$

840 with sparse storage $O(NK)$; the blockwise online-softmax used in this paper processes attention
841 block by block without ever materializing the full per-layer, per-head weight matrix. Full neural
842 inference across multiple cycles costs $O(C S_{\text{denoise}} LN(Kd + d^2))$, so the only factor scaling with
843 problem size is N . The projection P from a tour or a route set back to the same $N \times K$ sparse
844 adjacency is approximately linear in the number of solution edges (each TSP node contributes a
845 predecessor and a successor; CVRP boundaries add depot slots), with $O(NK)$ as the sparse-storage
846 upper bound. At training time, building the search-feedback target \mathcal{H}^* from M refined route sets
847 requires $O(MN)$ edge scanning plus $O(NK)$ storage, and the row-wise cross-entropy in Eq. (8) is
848 computed only over the NK candidate slots, never over a dense $N \times N$ label matrix.

849 F LKH algorithm for TSP

850 This appendix specifies how the Lin–Kernighan–Helsgaun heuristic [6] is integrated into the LSSL
 851 learning–search loop. Section F.1 formalises vanilla LKH around a single per-node candidate measure,
 852 the sole channel through which structural prior information enters k -opt search. Section F.2 introduces
 853 Guided-LKH, which retains the LKH search core but replaces this static, geometry-based measure
 854 with a trajectory-conditioned one refreshed at every LSSL cycle. Algorithm 2 provides a line-level
 855 diff between the two main loops.

856 F.1 Vanilla LKH

857 Let an instance be a complete graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ on $N = |\mathcal{V}|$ cities with non-negative edge weight
 858 $w : \mathcal{E} \rightarrow \mathbb{R}_{\geq 0}$, and let $\text{cost}(\mathcal{R}) = \sum_{e \in \mathcal{R}} w(e)$ denote the length of a Hamiltonian tour $\mathcal{R} \subset \mathcal{E}$.
 859 Vanilla LKH [69, 6] attacks $\min_{\mathcal{R}} \text{cost}(\mathcal{R})$ by sequential k -opt local search. Its efficiency rests on
 860 two static, instance-only objects fixed once at the start: a node-potential vector $\pi^* \in \mathbb{R}^N$ used to lift
 861 the cost matrix into reduced costs, and a per-vertex candidate measure $\mathcal{C}^{\text{geo}}(i)$ that lists, for each city
 862 i , the K neighbours that the k -opt recursion is allowed to consider when extending a partial move
 863 from i . We introduce the four ingredients in turn (k -opt, the minimum 1-tree, α -nearness, and the
 864 dual π^*) and then state the trial loop summarised in Algorithm 2(a).

865 **k -opt move and sequential search.** A k -opt move replaces k tour edges $\{x_1, \dots, x_k\} \subset \mathcal{R}$
 866 with k non-tour edges $\{y_1, \dots, y_k\}$ such that the resulting edge set is again a Hamiltonian tour
 867 \mathcal{R}' with $\text{cost}(\mathcal{R}') < \text{cost}(\mathcal{R})$. LKH restricts attention to sequential k -opt moves whose removed
 868 and added edges chain into a single alternating city sequence $(p_1, p_2, \dots, p_{2k}, p_{2k+1}=p_1)$, with
 869 $x_i = (p_{2i-1}, p_{2i})$ and $y_i = (p_{2i}, p_{2i+1})$, and explores them by a depth-first recursion of maximum
 870 depth $k_{\text{max}}=5$ (the LKH default). At each removed edge x_i the next vertex p_{2i} is one of the
 871 two tour neighbours of p_{2i-1} ; at each added edge y_i the next vertex p_{2i+1} is drawn from the
 872 candidate set $\mathcal{C}^{\text{geo}}(p_{2i})$ defined below. The recursion accepts the first y -sequence whose partial gain
 873 $\sum_{j \leq i} (w(x_j) - w(y_j))$ remains positive at every depth and whose closing edge (p_{2i}, p_1) yields a
 874 strictly improving full tour. Once accepted, the move is applied and search continues from the new
 875 tour; the trial terminates when one full sweep over \mathcal{V} uncovers no further improving move.

876 **Minimum 1-tree and α -nearness.** LKH constructs the candidate set from a sensitivity analysis of
 877 the Held–Karp lower bound. Fix an arbitrary special vertex $v_0 \in \mathcal{V}$. A 1-tree of \mathcal{G} is a spanning tree
 878 of $\mathcal{V} \setminus \{v_0\}$ together with two distinct edges of \mathcal{E} incident to v_0 ; a minimum 1-tree T is a 1-tree of
 879 minimum total cost. Its length $L(T)$ is a valid lower bound on $\min_{\mathcal{R}} \text{cost}(\mathcal{R})$, and equals it whenever
 880 every vertex has degree exactly 2 in T (in which case T is the optimal tour). The α -value of an edge
 881 $(i, j) \in \mathcal{E}$ measures by how much $L(T)$ would have to grow to forcibly contain (i, j) :

$$\alpha(i, j) = L(T^+(i, j)) - L(T), \quad T^+(i, j) = \arg \min_{T' \ni (i, j)} L(T'). \quad (29)$$

882 Edges with small α are precisely those the lower bound is unwilling to pay much to include, and are
 883 therefore strong candidates for an optimal tour.

884 **Held–Karp dual and subgradient ascent.** The raw statistic in (29) is sharp only when $L(T)$
 885 already approaches the optimum. To tighten it, LKH lifts the cost matrix by node potentials $\pi =$
 886 (π_1, \dots, π_N) via the reduced cost

$$\tilde{w}_{ij}(\pi) = w_{ij} + \pi_i + \pi_j. \quad (30)$$

887 This reduction is invariant on the optimal tour (every tour visits each vertex exactly twice in incident-
 888 edge degree, so its total cost merely shifts by the constant $2\sum_i \pi_i$) yet shifts the minimum 1-tree.
 889 Writing $L(T_\pi)$ for the minimum 1-tree length under $\tilde{w}(\pi)$, the Held–Karp lower bound becomes

$$w(\pi) = L(T_\pi) - 2\sum_{i \in \mathcal{V}} \pi_i, \quad (31)$$

890 and any maximiser $\pi^* \in \arg \max_{\pi} w(\pi)$ pushes every vertex’s degree in T_π towards 2. LKH
 891 approximates π^* by subgradient ascent (the LKH ASCENT routine): starting from $\pi^0 = 0$ it iterates

$$\pi^{\tau+1} = \pi^\tau + t^\tau (d^\tau - 2\mathbf{1}), \quad d_i^\tau = \deg_{T_{\pi^\tau}}(i), \quad (32)$$

892 with Polyak-style step size $t^\tau > 0$, until d^τ stabilises near 2 **1**. ASCENT is invoked exactly once per
 893 instance and dominates pre-search compute on large N . The post-ascent statistic

$$\alpha^{\pi^*}(i, j) = L(T_{\pi^*}^+(i, j)) - L(T_{\pi^*}) \quad (33)$$

894 is what LKH actually uses for candidate construction, and $\tilde{w}(\pi^*)$ is the cost matrix carried inside
 895 k -opt.

896 **Candidate measure and trial loop.** The geometric candidate measure of vertex i is the top- K
 897 slice of $-\alpha^{\pi^*}$:

$$\mathcal{C}^{\text{geo}}(i) = \text{argtop}_K\{-\alpha^{\pi^*}(i, \cdot)\}, \quad (34)$$

898 sorted in ascending α^{π^*} . This set serves a dual role inside k -opt: it is the only family from which
 899 y -edges may be drawn (filter), and its ordering is the priority along which the depth-first recursion
 900 expands children (search direction). Both π^* and \mathcal{C}^{geo} are deterministic functions of \mathcal{G} and remain
 901 fixed for the rest of the run.

902 The outer driver launches B_{trials} trials. Trial $r=1$ starts from a tour \mathcal{R} generated greedily under \mathcal{C}^{geo} ,
 903 runs sequential k -opt to a local optimum, and updates the global incumbent \mathcal{R} . Each subsequent
 904 trial perturbs the current incumbent by a non-sequential 4-opt-like segment swap (the LKH kick) and
 905 re-optimises. Vanilla LKH terminates once the trial budget is exhausted; the procedure is summarised
 906 in Algorithm 2(a). Crucially, \mathcal{C}^{geo} is the only channel through which any prior knowledge of likely-
 907 optimal edges enters search, and it never reacts to what search has discovered, an opening that LSSL
 908 exploits next.

909 F.2 Guided-LKH

910 LSSL retains every component of Algorithm 2(a), namely the Held–Karp dual π^* , the reduced cost
 911 $\tilde{w}(\pi^*)$, sequential k -opt with depth $k_{\text{max}}=5$, the non-sequential kick, and the multi-trial loop, and
 912 modifies a single object: the static geometric candidate measure \mathcal{C}^{geo} of Eq. (34) is replaced by a
 913 trajectory-conditioned measure $\mathcal{C}^{(c)}$, $c = 0, 1, \dots, C - 1$, that the sparse diffusion Transformer D_θ
 914 of Section 2 re-emits at every LSSL cycle in response to the running incumbent. The slots inside
 915 k -opt and their priority semantics are unchanged; only the contents of those slots are refreshed.

916 At cycle c , LSSL projects the cycle- c incumbent $\mathcal{R}^{(c)}$ (the same object Algorithm 1 carries be-
 917 tween D_θ and the search backend S) onto the K -NN sparse graph \mathcal{G}_K via the binary mask
 918 $\mathcal{H}^{(c)} = P(\mathcal{R}^{(c)}; \mathcal{G}_K) \in \{0, 1\}^{N \times K}$ of Section 2.2, where $\mathcal{H}_{ik}^{(c)}=1$ iff slot k at vertex i is one of
 919 the two tour neighbours of i in $\mathcal{R}^{(c)}$. D_θ ingests $(\mathcal{G}_K, \mathcal{H}^{(c)})$ along with the noise-axis schedule and
 920 emits a sparse, row-normalised edge-prior heatmap

$$\mathcal{H}_{\text{guide}}^{(c)} \in [0, 1]^{N \times K}, \quad [\mathcal{H}_{\text{guide}}^{(c)}]_{ik} \approx \Pr[(i, j_{ik}) \in \mathcal{R}^* \mid \mathcal{R}^{(c)}], \quad (35)$$

921 where $j_{ik} = \mathcal{N}_K(i)[k]$ is the slot-to-node alias of Eq. (12) and \mathcal{R}^* denotes the optimal tour. The
 922 cycle- c candidate measure consumed by the inner k -opt recursion is the ordered concatenation

$$\mathcal{C}^{(c)}(i) = \mathcal{C}^{\text{rsv}}(i) \circ \left(\{j_{ik} : k \in \text{argtop}_K[\mathcal{H}_{\text{guide}}^{(c)}]_{i, \cdot}\} \setminus \mathcal{C}^{\text{rsv}}(i) \right), \quad (36)$$

923 of total width K . The argtop_K ordering inherits the priority role that α^{π^*} played for \mathcal{C}^{geo} , so k -opt
 924 visits the model’s highest-confidence neighbours first. The reserved prefix $\mathcal{C}^{\text{rsv}}(i) \subseteq \mathcal{V} \setminus \{i\}$ collects
 925 slots that must remain in $\mathcal{C}^{(c)}(i)$ regardless of $\mathcal{H}_{\text{guide}}^{(c)}$; by default we populate it with the top heatmap
 926 slots from both the per-cycle and the cycle-averaged predictions, so that consensus edges across
 927 symmetry views (e.g. rotation and reflection augmentations of \mathcal{G}) are not displaced by within-cycle
 928 stochasticity. Geometric variants with $\mathcal{C}^{\text{rsv}} \subseteq \mathcal{C}^{\text{geo}}$ are equally admissible. The remaining slots are
 929 filled by the highest-scoring model neighbours.

930 The replacement of \mathcal{C}^{geo} by $\mathcal{C}^{(c)}$ is more than a candidate-set swap.

931 (i) Time variance: $\mathcal{C}^{(c)}$ depends on c and, through $\mathcal{R}^{(c)}$, on the entire search prefix; setting $\mathcal{C}^{(c)} \equiv \mathcal{C}^{\text{geo}}$
 932 for all c collapses Algorithm 2(b) back to (a).

933 (ii) Cumulative warm-start: cycle c is launched from $\mathcal{R}^{(c)}$, so earlier improvements both bias the next
 934 heatmap $\mathcal{H}_{\text{guide}}^{(c+1)}$ and seed k -opt inside the basin already discovered.

935 (iii) Amortised dual: π^* is computed by ASCENT at cycle 0 and then reused via a persistent
 936 LKH session that keeps the reduced cost $\tilde{w}(\pi^*)$ resident across cycles; only $\mathcal{C}^{(c)}$ is refreshed
 937 (concretely, after cycle 0 writes π^* to a PI_FILE, subsequent cycles consume it together with the new
 938 CANDIDATE_FILE and bypass ASCENT entirely). Guided-LKH thus preserves the Held–Karp dual
 939 setup and the sequential k -opt admissibility of vanilla LKH; the closed loop only changes which edges
 940 k -opt examines, through the dataflow $\mathcal{R}^{(c)} \rightarrow \mathcal{H}_{\text{guide}}^{(c)} \rightarrow \mathcal{C}^{(c)} \rightarrow \mathcal{R}^{(c+1)}$ shown in Algorithm 2(b).

Algorithm 2 Vanilla LKH (top) vs. Guided-LKH (bottom).

REUSED: dual π^* , reduced cost $\tilde{w}(\pi^*)$, k -opt admissibility, kick, multi-trial loop.

REPLACED: static geometric \mathcal{C}^{geo} becomes trajectory-conditioned $\mathcal{C}^{(c)}$.

CLOSED LOOP: $\mathcal{R}^{(c)} \rightarrow \mathcal{H}_{\text{guide}}^{(c)} \rightarrow \mathcal{C}^{(c)} \rightarrow \mathcal{R}^{(c+1)}$.

(a) Vanilla LKH.

Require: graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$, candidate width K , trial budget B_{trials}

Ensure: best tour $\bar{\mathcal{R}}$ found across all trials

```

1: compute dual  $\pi^* \leftarrow \arg \max_{\pi} w(\pi)$  via ASCENT ▷ Eqs. (31),(32)
2: build static candidate measure  $\mathcal{C}^{\text{geo}} = \{\mathcal{C}^{\text{geo}}(i)\}_{i \in \mathcal{V}}$  ▷ Eqs. (33),(34)
3:  $\bar{\mathcal{R}} \leftarrow \perp$ 
4: for  $r = 1, \dots, B_{\text{trials}}$  do
5:   if  $r=1$  then initialise  $\mathcal{R}$  greedily under  $\mathcal{C}^{\text{geo}}$ ; else kick:  $\mathcal{R} \leftarrow \bar{\mathcal{R}}$  and apply a non-sequential segment swap
6:   repeat
7:     perform one  $k$ -opt sweep over  $\mathcal{V}$  under  $\tilde{w}(\pi^*)$  with candidate measure  $\mathcal{C}^{\text{geo}}$ , accepting the first improving move
8:   until no improving move is found in a full sweep
9:   if  $\bar{\mathcal{R}} = \perp$  or  $\text{cost}(\mathcal{R}) < \text{cost}(\bar{\mathcal{R}})$  then  $\bar{\mathcal{R}} \leftarrow \mathcal{R}$ 
10:  end if
11: end for
12: return  $\bar{\mathcal{R}}$ 

```

(b) Guided-LKH (LSSL). blue block = line differs from (a); orange = structurally replaced symbol.

Require: graph \mathcal{G} , sparse diffusion Transformer D_{θ} , candidate width K , number of cycles C , trial budget per cycle B_{trials} ,
 denoising schedule $\{t_r\}_{r=0}^T$ with $t_0 = 1, t_T = 0$

Ensure: best tour $\bar{\mathcal{R}}$ found across all cycles

```

1: compute dual  $\pi^* \leftarrow \arg \max_{\pi} w(\pi)$  via ASCENT ▷ Eq. (31); amortised: solved once at cycle 0, then reused
2: build static candidate measure  $\mathcal{C}^{\text{geo}}$  ▷ geometric reference, retained as fallback
3:  $\bar{\mathcal{R}} \leftarrow \perp, \mathcal{R}^{(0)} \leftarrow \perp$ 
4: for  $c = 0, 1, \dots, C - 1$  do
5:    $\mathcal{H}^{(c)} \leftarrow P(\mathcal{R}^{(c)}; \mathcal{G}_K)$  if  $\mathcal{R}^{(c)} \neq \perp$  else random  $N \times K$  init ▷ condition on incumbent prefix
6:    $\mathcal{H}_{\text{guide}}^{(c)} \leftarrow \text{DENOISESCHEDULE}(D_{\theta}, \mathcal{G}_K, \mathcal{H}^{(c)}; \{t_r\}_{r=0}^T)$  ▷ Alg. 1 lines 6–10; Eq. (35)
7:   form trajectory-conditioned candidate measure  $\mathcal{C}^{(c)}$  from  $\mathcal{H}_{\text{guide}}^{(c)}$  ▷ Eq. (36)
▷ only  $\mathcal{C}^{(c)}$  refreshes per cycle
8:   for  $r = 1, \dots, B_{\text{trials}}$  do
9:     if  $r = 1$  and  $\bar{\mathcal{R}} \neq \perp$  then
10:      warm-start  $\mathcal{R} \leftarrow \bar{\mathcal{R}}$  ▷ cumulative incumbent passes through cycle boundary
11:     else
12:       kick: perturb  $\mathcal{R}$  by a non-sequential segment swap
13:     end if
14:     repeat
15:       perform one  $k$ -opt sweep over  $\mathcal{V}$  under  $\tilde{w}(\pi^*)$  with candidate measure  $\mathcal{C}^{(c)}$ , accepting the first improving
move
16:     until no improving move is found in a full sweep
17:     if  $\bar{\mathcal{R}} = \perp$  or  $\text{cost}(\mathcal{R}) < \text{cost}(\bar{\mathcal{R}})$  then  $\bar{\mathcal{R}} \leftarrow \mathcal{R}$ 
18:     end if
19:     end for
20:      $\mathcal{R}^{(c+1)} \leftarrow \bar{\mathcal{R}}$  ▷ snapshot incumbent for next cycle
21:   end for
22: return  $\bar{\mathcal{R}}$ 

```

941 **G HGS algorithm for CVRP**

942 This appendix specifies how the Hybrid Genetic Search heuristic [44, 43] is integrated into the
 943 LSSL learning–search loop. Section G.1 formalises vanilla HGS around its granular neighbourhood
 944 measure, the sole channel through which structural prior information enters education and the SWAP*
 945 neighbourhood. Section G.2 introduces Guided-HGS, which retains the HGS search core but replaces
 946 this static, geometry-based measure with a trajectory-conditioned counterpart refreshed at every
 947 LSSL cycle. Algorithm 3 provides a line-level diff between the two main loops.

948 **G.1 Vanilla HGS**

949 Let a CVRP instance specify a customer set \mathcal{V}_c , a depot 0, distances c_{ij} , customer demands q_i , and a
 950 vehicle capacity Q . A solution (route set) \mathcal{R} is encoded by a giant tour $\tau \in \text{Sym}(\mathcal{V}_c)$ on which the
 951 linear-time SPLIT dynamic program [44]

$$f^\phi(\tau) = \min_{\sigma \in \Sigma(\tau)} \sum_{r \in \sigma} (c(r) + \lambda_Q(Q(r) - Q)^+) \quad (37)$$

952 recovers the unique cost-optimal capacity-feasible route segmentation, where $\Sigma(\tau)$ is the set of all
 953 segmentations of τ , $Q(r)$ is the load of route r , and λ_Q is a current capacity-penalty coefficient. HGS
 954 evolves a population $\mathcal{P} = \mathcal{P}_{\text{feas}} \cup \mathcal{P}_{\text{inf}}$ of giant tours under a biased fitness combining a cost rank
 955 with a diversity rank,

$$\Phi_{\mathcal{P}}(\mathcal{R}) = f_{\mathcal{P}}^\phi(\mathcal{R}) + \left(1 - \frac{n_{\text{elite}}}{|\mathcal{P}|}\right) f_{\mathcal{P}}^{\text{div}}(\mathcal{R}), \quad (38)$$

956 where $f_{\mathcal{P}}^\phi$ is the cost rank and $f_{\mathcal{P}}^{\text{div}}$ the broken-pairs distance to the n_{closest} most similar individuals in
 957 \mathcal{P} ; the $1 - n_{\text{elite}}/|\mathcal{P}|$ weight preserves the top n_{elite} on quality alone. Capacity and duration violations
 958 V^Q, V^T are softly penalised by $\tilde{f} = f^\phi + \lambda_Q V^Q + \lambda_T V^T$, with λ_Q, λ_T periodically retuned toward
 959 a target feasible-ratio ξ^{ref} (controlled infeasibility). A single sparse object then dominates HGS, the
 960 granular neighbourhood measure

$$\Gamma^{\text{geo}}(i) = \underset{\gamma}{\text{argtop}}\{-c_{ij} : j \in \mathcal{V}_c \setminus \{i\}\}, \quad (39)$$

961 i.e. the γ spatial nearest neighbours of customer i , which cap each move neighbourhood at $\mathcal{O}(\gamma n)$
 962 instead of $\mathcal{O}(n^2)$ and form the single channel through which prior structural information enters
 963 search: every classical move (RELOCATE, SWAP, 2-OPT, 2-OPT*) and the SWAP* neighbourhood
 964 of Vidal [43] restrict their partner customers to Γ^{geo} , and the 4μ random initial giant tours are
 965 educated by first-improvement local search under the same restriction. Each generation then draws
 966 two binary-tournament parents under $\Phi_{\mathcal{P}}$, recombines them by Ordered Crossover on the depot-free
 967 permutation, recovers feasibility via SPLIT, and re-educates the offspring; offspring still infeasible
 968 are REPAIRED at $10\times$ penalty with probability $1/2$.

969 **G.2 Guided-HGS**

970 LSSL leaves the HGS search core untouched (SPLIT (37), biased fitness (38), controlled infeasibility,
 971 and the SWAP* structural theorem of Vidal [43]), and replaces only the static geometric prior Γ^{geo}
 972 with a trajectory-conditioned counterpart refreshed at every cycle $c = 0, 1, \dots, C - 1$ by the sparse
 973 diffusion Transformer D_θ of Section 2, with three structural consequences.

974 (i) Dynamic neighbourhood prior. At cycle c , D_θ ingests the $N \times K$ projection $\mathcal{H}^{(c)} = P(\mathcal{R}^{(c)}; \mathcal{G}_K)$
 975 of the cumulative-best route set $\mathcal{R}^{(c)}$ (the same carrier as in Algorithm 1, set to the best feasible
 976 solution found through cycle $c-1$) and emits a customer edge-prior heatmap $\mathcal{H}_{\text{guide}}^{(c)} \in [0, 1]^{|\mathcal{V}_c| \times K}$
 977 with $[\mathcal{H}_{\text{guide}}^{(c)}]_{ik} \approx \Pr[(i, j_{ik}) \in \mathcal{R}^* \mid \mathcal{R}^{(c)}]$, where \mathcal{R}^* denotes the optimal route set. The granular
 978 measure driving every classical move and SWAP* is then

$$\tilde{\Gamma}^{(c)}(i) = \Gamma_D^{(c)}(i) \circ (\Gamma^{\text{geo}}(i) \setminus \Gamma_D^{(c)}(i)), \quad \Gamma_D^{(c)}(i) = \{j_{ik} : k \in \arg \text{top}_\gamma[\mathcal{H}_{\text{guide}}^{(c)}]_{i,\cdot}\}, \quad (40)$$

979 with model-ranked neighbours leading and Γ^{geo} as backfill. The substitution preserves every HGS
 980 invariant, including the $\Theta(N^3) \rightarrow o(N^2)$ SWAP* compression.

981 (ii) Dynamic boundary bias. At cycle c , the decoder’s node-selection head additionally outputs a
 982 per-customer heatmap $\mathcal{H}_{\text{node}}^{(c)} \in [0, 1]^{N_c}$ (Section D.2), with $[\mathcal{H}_{\text{node}}^{(c)}]_i$ modelling the probability that
 983 customer i anchors a route boundary. Its rank induces a node-boundary bonus

$$\delta_{\text{node}}^{(c)}(i) = 1 - r_{\text{node}}^{(c)}(i) / |\mathcal{V}_c| \in [0, 1), \quad (41)$$

984 used as a tiebreaker on depot-attachment moves while leaving move evaluations unchanged.

985 (iii) Incumbent-carrying restart. The cumulative best $\mathcal{R}^{(c)}$ crosses cycle boundaries: it conditions
 986 $\mathcal{H}_{\text{guide}}^{(c)}$, enters $\mathcal{P}_{\text{feas}}$ as an elite when available, and seeds the rest of the population along $\tilde{\Gamma}^{(c)}$, so
 987 each cycle resumes inside the basin already discovered. Setting $C = 1$, $\tilde{\Gamma}^{(0)} \equiv \Gamma^{\text{geo}}$, $\delta_{\text{node}}^{(0)} \equiv 0$ with
 988 random seeding recovers vanilla HGS.

Algorithm 3 Vanilla HGS (top) vs. Guided-HGS (bottom).

REUSED: SPLIT, biased fitness, OX, controlled infeasibility, survivor selection, SWAP*.

REPLACED: $\Gamma^{\text{geo}} \rightarrow \tilde{\Gamma}^{(c)}$ plus node-boundary bonus $\delta_{\text{node}}^{(c)}$.

CLOSED LOOP: $\mathcal{R}^{(c)} \rightarrow \mathcal{H}_{\text{guide}}^{(c)} \rightarrow (\tilde{\Gamma}^{(c)}, \mathcal{H}_{\text{node}}^{(c)}, \delta_{\text{node}}^{(c)}) \rightarrow \mathcal{R}^{(c+1)}$.

(a) Vanilla HGS.

Require: instance \mathcal{G} , capacity Q , granularity γ , total time budget

Ensure: incumbent feasible route set $\bar{\mathcal{R}}$

- 1: build static granular measure Γ^{geo} ▷ Eq. (39)
 - 2: seed \mathcal{P} with random giant tours, SPLIT (37), educate on Γ^{geo}
 - 3: $\bar{\mathcal{R}} \leftarrow \arg \min_{\mathcal{R} \in \mathcal{P}_{\text{feas}}} f^\phi(\mathcal{R})$
 - 4: **while** time budget not exhausted **do**
 - 5: binary-tournament parents P_1, P_2 by $\Phi_{\mathcal{P}}$ ▷ Eq. (38)
 - 6: $O \leftarrow \text{OX}(P_1, P_2)$; SPLIT; educate on Γ^{geo} ; w.p. 1/2 REPAIR at 10× penalty if infeasible
 - 7: insert O into $\mathcal{P}_{\text{feas}}/\mathcal{P}_{\text{inf}}$; survivor selection on overflow (clones first, then by $\Phi_{\mathcal{P}}$)
 - 8: periodically adapt λ_Q, λ_T toward feasible-ratio target ξ^{ref} ▷ controlled infeasibility
 - 9: $\bar{\mathcal{R}} \leftarrow \arg \min_{\mathcal{R} \in \mathcal{P}_{\text{feas}} \cup \{\bar{\mathcal{R}}\}} f^\phi(\mathcal{R})$ ▷ update incumbent
 - 10: **end while**
 - 11: **return** $\bar{\mathcal{R}}$
-

(b) Guided-HGS (LSSL). ■ = line differs from (a); ■ = structurally replaced symbol.


Require: instance \mathcal{G} , capacity Q , sparse diffusion Transformer D_θ , granularity γ , cycles C , denois-
 ing schedule $\{t_r\}_{r=0}^T$ with $t_0 = 1, t_T = 0$

Ensure: incumbent feasible route set $\bar{\mathcal{R}}$


- 1: build static granular measure Γ^{geo} ▷ geometric backfill only
 - 2: $\bar{\mathcal{R}} \leftarrow \perp, \mathcal{R}^{(0)} \leftarrow \perp$
 - 3: **for** $c = 0, 1, \dots, C - 1$ **do**
 - 4: $\mathcal{H}^{(c)} \leftarrow \mathcal{P}(\mathcal{R}^{(c)}; \mathcal{G}_K)$ if $\mathcal{R}^{(c)} \neq \perp$ else random $N \times K$ init
 - 5: $\mathcal{H}_{\text{guide}}^{(c)} \leftarrow \text{DENOISESCHEDULE}(D_\theta, \mathcal{G}_K, \mathcal{H}^{(c)}; \{t_r\}_{r=0}^T)$ ▷ Alg. 1 lines 6–10
 - 6: form $\tilde{\Gamma}^{(c)}, \mathcal{H}_{\text{node}}^{(c)}, \delta_{\text{node}}^{(c)}$ from $\mathcal{H}_{\text{guide}}^{(c)}$ ▷ Eqs. (40)–(41)
 - 7: re-seed \mathcal{P} along $\tilde{\Gamma}^{(c)}$, with $\mathcal{R}^{(c)}$ as elite if available
 - 8: **while** cycle- c time budget not exhausted **do**
 - 9: binary-tournament parents P_1, P_2 by $\Phi_{\mathcal{P}}$ ▷ Eq. (38)
 - 10: $O \leftarrow \text{OX}(P_1, P_2)$; SPLIT; educate on $\tilde{\Gamma}^{(c)}$ (■ $\delta_{\text{node}}^{(c)}$ tiebreaks at depot); w.p. 1/2 REPAIR
 at 10× penalty if infeasible
 - 11: insert O into $\mathcal{P}_{\text{feas}}/\mathcal{P}_{\text{inf}}$; survivor selection on overflow (clones first, then by $\Phi_{\mathcal{P}}$)
 - 12: periodically adapt λ_Q, λ_T toward feasible-ratio target ξ^{ref} ▷ controlled infeasibility
 - 13: $\bar{\mathcal{R}} \leftarrow \arg \min_{\mathcal{R} \in \mathcal{P}_{\text{feas}} \cup \{\bar{\mathcal{R}}\}} f^\phi(\mathcal{R})$ ▷ update incumbent
 - 14: **end while**
 - 15: $\mathcal{R}^{(c+1)} \leftarrow \bar{\mathcal{R}}$ ▷ best-so-far carries to next cycle
 - 16: **end for**
 - 17: **return** $\bar{\mathcal{R}}$
-

989 **H Solution Visualization**

990 **H.1 Solution Visualization on Synthetic TSP Instances**




(a) Optimal tour
(BKS, length 71.7031).




(b) LSSL tour
($C=2$, $B_{\text{trials}}=256$, gap -0.0135%).

Figure 4: Solution visualization on a synthetic TSP10000 instance (sample #12, $N=10000$).

991 **H.2 Solution Visualization on Art-TSP Instances**




(a) Optimal tour
(BKS, length 5,757,191).




(b) LSSL tour
($C=4$, $B_{\text{trials}}=512$, gap 0.026%).

Figure 5: Solution visualization on Art-TSP instance *mona-lisa*100K ($N=100000$).

992 **H.3 Solution Visualization on VLSI Instances**

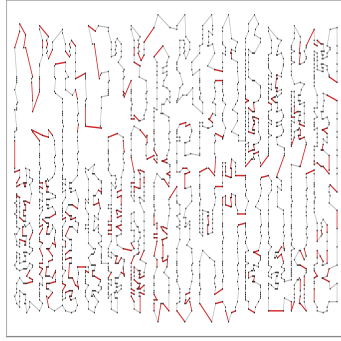


(a) Optimal tour
(BKS, length 6,197).




(b) LSSL tour
($C=2$, $B_{\text{trials}}=256$, gap -0.387%).

Figure 6: Solution visualization on VLSI instance *djb2036* ($N=2036$).




(a) Optimal tour
(BKS, length 6,764).



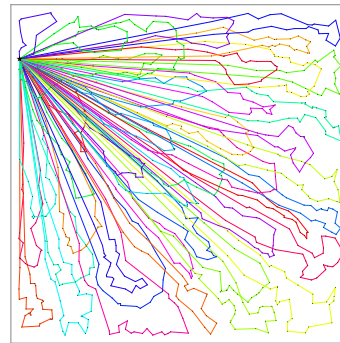
(b) LSSL tour
($C=2$, $B_{\text{trials}}=256$, gap -0.407%).

Figure 7: Solution visualization on VLSI instance bck2217 ($N=2217$).

993 **H.4 Solution Visualization on CVRPLIB Instances**




(a) Optimal tour
(BKS, length 72,355).

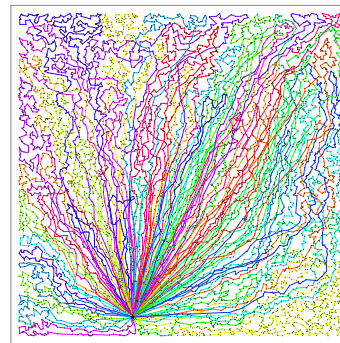


(b) LSSL tour
($C=4$, $B_t=128$, gap 1.093%).

Figure 8: Solution visualization on CVRPLIB Set-X instance X-n1001-k43 ($N=1001$, 43 routes).



(a) HGS tour
($B_t=8\text{h}$, length 111,115).



(b) LSSL tour
($C=4$, $B_t=128$, length 114,633, gap 3.166%).

Figure 9: Solution visualization on CVRPLIB Set-XL instance XL-n9571-k55 ($N=9,571$, 55 routes). CVRPLIB only publishes the optimal cost (106,791) for this instance and not the corresponding tour, so we substitute a long-time HGS run ($B_t=8\text{h}$, no LSSL candidates) as the reference tour against which the LSSL gap is computed. The gap shown is therefore measured against this HGS-8h reference rather than the unattainable BKS tour.

994 I Broader impacts, Limitations and Future Work

995 **Broader impacts.** This research advances neural combinatorial optimization by introducing LSSL,
 996 a closed-loop learning–search framework for large-scale Vehicle Routing Problems (VRPs). We
 997 hope this framework will provide useful insight for future work on scalable, search-aware neural
 998 optimization. As a general methodology for solving VRPs, LSSL is not anticipated to create specific
 999 negative societal impacts beyond those already associated with optimization systems in logistics and
 1000 routing.

1001 **Limitation and future work.** First, although LSSL already scales well, the sparse backbone and
 1002 the search backend could still be co-designed more tightly so that learned priors and downstream local
 1003 operators interact even more effectively. Second, divide-and-conquer methods such as GLOP [25]
 1004 and UDC [26] have been shown to dramatically shorten solving time on very large instances. For
 1005 methodological purity we deliberately keep LSSL free of any such decomposition in this paper;
 1006 combining LSSL with a divide-and-conquer scheme is a natural next step that may compound their
 1007 strengths. Finally, extending LSSL to a broader range of constrained combinatorial optimization
 1008 problems remains an important future direction.

1009 J Licenses

1010 The licenses for the codes and the datasets used in this work are listed in Table 15.

Table 15: List of licenses for the codes and datasets we used in this work

Resource	Type	Link	License
LKH3 [6]	Code	http://webhotel4.ruc.dk/keld/research/LKH-3/	Available for academic research use
HGS [43]	Code	https://github.com/chkwon/PyHygese	MIT License
Concorde [81]	Code	https://github.com/jvkersch/pyconcorde	BSD 3-Clause License
POMO [12]	Code	https://github.com/yd-kwon/POMO	MIT License
BQ [14]	Code	https://github.com/naver/bq-nco	CC BY-NC-SA 4.0
LEHD [13]	Code	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/LEHD	MIT License
INViT [28]	Code	https://github.com/Kasumigaoka-Utaha/INViT	Available for academic research use
L2C-Insert [21]	Code	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/L2C-Insert	MIT License
Att-GCN+MCTS [30]	Code	https://github.com/SaneLYX/TSP_Att-GCRN-MCTS	MIT License
DIFUSCO [32]	Code	https://github.com/Edward-Sun/DIFUSCO	MIT License
T2T [33]	Code	https://github.com/Thinklab-SJTU/T2TCO	MIT License
Fast-T2T [34]	Code	https://github.com/Thinklab-SJTU/Fast-T2T	MIT License
GenSCO [35]	Code	https://github.com/Thinklab-SJTU/GenSCO	MIT License
H-TSP [22]	Code	https://github.com/Learning4Optimization-HUST/H-TSP	MIT License
GLOP [25]	Code	https://github.com/henry-yeh/GLOP	MIT License
UDC [26]	Code	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/UDC-Large-scale-CO-master	MIT License
NeuroLKH [38]	Code	https://github.com/liangxinedu/NeuroLKH	Available for academic research use
TSPLib [45]	Dataset	http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/	Available for any non-commercial use
VLSI [46]	Dataset	http://www.math.uwaterloo.ca/tsp/vlsi/	Available for academic research use
Art-TSP [47]	Dataset	http://www.math.uwaterloo.ca/tsp/data/art/	Available for academic research use
CVRPLib [48]	Dataset	http://vrp.galgos.inf.puc-rio.br/index.php/en/	Available for academic research use

1011 **NeurIPS Paper Checklist**

1012 **1. Claims**

1013 Question: Do the main claims made in the abstract and introduction accurately reflect the
1014 paper’s contributions and scope?

1015 Answer: **[Yes]**

1016 Justification: The abstract and introduction state the closed-loop LSSL framework, its sparse
1017 scaling motivation, and its empirical scope, and these claims are matched by the method and
1018 experiment sections(Section 2 and 3).

1019 Guidelines:

- 1020 • The answer **[N/A]** means that the abstract and introduction do not include the claims
1021 made in the paper.
- 1022 • The abstract and/or introduction should clearly state the claims made, including the
1023 contributions made in the paper and important assumptions and limitations. A **[No]** or
1024 **[N/A]** answer to this question will not be perceived well by the reviewers.
- 1025 • The claims made should match theoretical and experimental results, and reflect how
1026 much the results can be expected to generalize to other settings.
- 1027 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
1028 are not attained by the paper.

1029 **2. Limitations**

1030 Question: Does the paper discuss the limitations of the work performed by the authors?

1031 Answer: **[Yes]**

1032 Justification: The paper discusses the limitations of the work in Appendix I.

1033 Guidelines:

- 1034 • The answer **[N/A]** means that the paper has no limitation while the answer **[No]** means
1035 that the paper has limitations, but those are not discussed in the paper.
- 1036 • The authors are encouraged to create a separate “Limitations” section in their paper.
- 1037 • The paper should point out any strong assumptions and how robust the results are to
1038 violations of these assumptions (e.g., independence assumptions, noiseless settings,
1039 model well-specification, asymptotic approximations only holding locally). The authors
1040 should reflect on how these assumptions might be violated in practice and what the
1041 implications would be.
- 1042 • The authors should reflect on the scope of the claims made, e.g., if the approach was
1043 only tested on a few datasets or with a few runs. In general, empirical results often
1044 depend on implicit assumptions, which should be articulated.
- 1045 • The authors should reflect on the factors that influence the performance of the approach.
1046 For example, a facial recognition algorithm may perform poorly when image resolution
1047 is low or images are taken in low lighting. Or a speech-to-text system might not be
1048 used reliably to provide closed captions for online lectures because it fails to handle
1049 technical jargon.
- 1050 • The authors should discuss the computational efficiency of the proposed algorithms
1051 and how they scale with dataset size.
- 1052 • If applicable, the authors should discuss possible limitations of their approach to
1053 address problems of privacy and fairness.
- 1054 • While the authors might fear that complete honesty about limitations might be used by
1055 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
1056 limitations that aren’t acknowledged in the paper. The authors should use their best
1057 judgment and recognize that individual actions in favor of transparency play an impor-
1058 tant role in developing norms that preserve the integrity of the community. Reviewers
1059 will be specifically instructed to not penalize honesty concerning limitations.

1060 **3. Theory assumptions and proofs**

1061 Question: For each theoretical result, does the paper provide the full set of assumptions and
1062 a complete (and correct) proof?

1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115

Answer: [N/A]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The experimental details and model settings are in Section 3 and Appendix C, D, F and G. We will make our source code publicly available upon acceptance.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

1116 Question: Does the paper provide open access to the data and code, with sufficient instruc-
1117 tions to faithfully reproduce the main experimental results, as described in supplemental
1118 material?

1119 Answer: [No]

1120 Justification: We will release our code and data once the paper is accepted.

1121 Guidelines:

- 1122 • The answer [N/A] means that paper does not include experiments requiring code.
- 1123 • Please see the NeurIPS code and data submission guidelines ([https://neurips.cc/
1124 public/guides/CodeSubmissionPolicy](https://neurips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 1125 • While we encourage the release of code and data, we understand that this might not
1126 be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not
1127 including code, unless this is central to the contribution (e.g., for a new open-source
1128 benchmark).
- 1129 • The instructions should contain the exact command and environment needed to run to
1130 reproduce the results. See the NeurIPS code and data submission guidelines ([https:
1131 //neurips.cc/public/guides/CodeSubmissionPolicy](https://neurips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 1132 • The authors should provide instructions on data access and preparation, including how
1133 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 1134 • The authors should provide scripts to reproduce all experimental results for the new
1135 proposed method and baselines. If only a subset of experiments are reproducible, they
1136 should state which ones are omitted from the script and why.
- 1137 • At submission time, to preserve anonymity, the authors should release anonymized
1138 versions (if applicable).
- 1139 • Providing as much information as possible in supplemental material (appended to the
1140 paper) is recommended, but including URLs to data and code is permitted.

1141 6. Experimental setting/details

1142 Question: Does the paper specify all the training and test details (e.g., data splits, hyperpa-
1143 rameters, how they were chosen, type of optimizer) necessary to understand the results?

1144 Answer: [Yes]

1145 Justification: The experimental details and model settings are in Section 3 and Ap-
1146 pendix C, D, F and G. We will make our source code publicly available upon acceptance.

1147 Guidelines:

- 1148 • The answer [N/A] means that the paper does not include experiments.
- 1149 • The experimental setting should be presented in the core of the paper to a level of detail
1150 that is necessary to appreciate the results and make sense of them.
- 1151 • The full details can be provided either with the code, in appendix, or as supplemental
1152 material.

1153 7. Experiment statistical significance

1154 Question: Does the paper report error bars suitably and correctly defined or other appropriate
1155 information about the statistical significance of the experiments?

1156 Answer: [No]

1157 Justification: In neural combinatorial optimization, we usually adopt the optimality gap and
1158 inference time as metrics to measure performance. Both of them are deterministic.

1159 Guidelines:

- 1160 • The answer [N/A] means that the paper does not include experiments.
- 1161 • The authors should answer [Yes] if the results are accompanied by error bars, confidence
1162 intervals, or statistical significance tests, at least for the experiments that support the
1163 main claims of the paper.
- 1164 • The factors of variability that the error bars are capturing should be clearly stated (for
1165 example, train/test split, initialization, random drawing of some parameter, or overall
1166 run with given experimental conditions).

- 1167 • The method for calculating the error bars should be explained (closed form formula,
1168 call to a library function, bootstrap, etc.)
- 1169 • The assumptions made should be given (e.g., Normally distributed errors).
- 1170 • It should be clear whether the error bar is the standard deviation or the standard error
1171 of the mean.
- 1172 • It is OK to report 1-sigma error bars, but one should state it. The authors should
1173 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
1174 of Normality of errors is not verified.
- 1175 • For asymmetric distributions, the authors should be careful not to show in tables or
1176 figures symmetric error bars that would yield results that are out of range (e.g., negative
1177 error rates).
- 1178 • If error bars are reported in tables or plots, the authors should explain in the text how
1179 they were calculated and reference the corresponding figures or tables in the text.

1180 8. Experiments compute resources

1181 Question: For each experiment, does the paper provide sufficient information on the com-
1182 puter resources (type of compute workers, memory, time of execution) needed to reproduce
1183 the experiments?

1184 Answer: [Yes]

1185 Justification: We provide the resource requirements for training and dataset self-
1186 bootstrapping in the Appendix C.

1187 Guidelines:

- 1188 • The answer [N/A] means that the paper does not include experiments.
- 1189 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
1190 or cloud provider, including relevant memory and storage.
- 1191 • The paper should provide the amount of compute required for each of the individual
1192 experimental runs as well as estimate the total compute.
- 1193 • The paper should disclose whether the full research project required more compute
1194 than the experiments reported in the paper (e.g., preliminary or failed experiments that
1195 didn't make it into the paper).

1196 9. Code of ethics

1197 Question: Does the research conducted in the paper conform, in every respect, with the
1198 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

1199 Answer: [Yes]

1200 Justification: Our research conducted in the paper conforms with the NeurIPS Code of
1201 Ethics.

1202 Guidelines:

- 1203 • The answer [N/A] means that the authors have not reviewed the NeurIPS Code of
1204 Ethics.
- 1205 • If the authors answer [No], they should explain the special circumstances that require a
1206 deviation from the Code of Ethics.
- 1207 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
1208 eration due to laws or regulations in their jurisdiction).

1209 10. Broader impacts

1210 Question: Does the paper discuss both potential positive societal impacts and negative
1211 societal impacts of the work performed?

1212 Answer: [Yes]

1213 Justification: We discuss the broader impacts in Appendix I.

1214 Guidelines:

- 1215 • The answer [N/A] means that there is no societal impact of the work performed.
- 1216 • If the authors answer [N/A] or [No], they should explain why their work has no societal
1217 impact or why the paper does not address societal impact.

- 1218
- 1219
- 1220
- 1221
- 1222
- 1223
- 1224
- 1225
- 1226
- 1227
- 1228
- 1229
- 1230
- 1231
- 1232
- 1233
- 1234
- 1235
- 1236
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
 - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
 - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
 - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

1237

1238 Question: Does the paper describe safeguards that have been put in place for responsible

1239 release of data or models that have a high risk for misuse (e.g., pre-trained language models,

1240 image generators, or scraped datasets)?

1241 Answer: [N/A]

1242 Justification: This paper has no such risks.

1243 Guidelines:

- 1244
- 1245
- 1246
- 1247
- 1248
- 1249
- 1250
- 1251
- 1252
- 1253
- The answer [N/A] means that the paper poses no such risks.
 - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
 - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
 - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

1254

1255 Question: Are the creators or original owners of assets (e.g., code, data, models), used in

1256 the paper, properly credited and are the license and terms of use explicitly mentioned and

1257 properly respected?

1258 Answer: [Yes]

1259 Justification: We have properly cited the assets (e.g., code, data, models) in section J.

1260 Guidelines:

- 1261
- 1262
- 1263
- 1264
- 1265
- 1266
- 1267
- 1268
- 1269
- 1270
- 1271
- The answer [N/A] means that the paper does not use existing assets.
 - The authors should cite the original paper that produced the code package or dataset.
 - The authors should state which version of the asset is used and, if possible, include a URL.
 - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
 - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
 - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- 1272
- 1273
- 1274
- 1275
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
 - If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

1276 **13. New assets**

1277 Question: Are new assets introduced in the paper well documented and is the documentation
1278 provided alongside the assets?

1279 Answer: [N/A]

1280 Justification: Currently, the paper does not release new assets. Upon acceptance, we will
1281 make the source code and LSSL-bootstrapped datasets publicly available, accompanied by
1282 comprehensive documentation.

1283 Guidelines:

- 1284
- 1285
- 1286
- 1287
- 1288
- 1289
- 1290
- 1291
- The answer [N/A] means that the paper does not release new assets.
 - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
 - The paper should discuss whether and how consent was obtained from people whose asset is used.
 - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

1292 **14. Crowdsourcing and research with human subjects**

1293 Question: For crowdsourcing experiments and research with human subjects, does the paper
1294 include the full text of instructions given to participants and screenshots, if applicable, as
1295 well as details about compensation (if any)?

1296 Answer: [N/A]

1297 Justification: The paper does not involve crowdsourcing nor research with human subjects.

1298 Guidelines:

- 1299
- 1300
- 1301
- 1302
- 1303
- 1304
- 1305
- 1306
- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
 - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
 - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

1307 **15. Institutional review board (IRB) approvals or equivalent for research with human
1308 subjects**

1309 Question: Does the paper describe potential risks incurred by study participants, whether
1310 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
1311 approvals (or an equivalent approval/review based on the requirements of your country or
1312 institution) were obtained?

1313 Answer: [N/A]

1314 Justification: This paper does not incur such risks.

1315 Guidelines:

- 1316
- 1317
- 1318
- 1319
- 1320
- 1321
- 1322
- 1323
- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
 - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
 - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

- 1324 • For initial submissions, do not include any information that would break anonymity (if
1325 applicable), such as the institution conducting the review.

1326 **16. Declaration of LLM usage**

1327 Question: Does the paper describe the usage of LLMs if it is an important, original, or
1328 non-standard component of the core methods in this research? Note that if the LLM is used
1329 only for writing, editing, or formatting purposes and does *not* impact the core methodology,
1330 scientific rigor, or originality of the research, declaration is not required.

1331 Answer: [N/A]

1332 Justification: The core methodological development of this work does not rely on LLMs as
1333 essential, original, or non-standard components.

1334 Guidelines:

- 1335 • The answer [N/A] means that the core method development in this research does not
1336 involve LLMs as any important, original, or non-standard components.
1337 • Please refer to our LLM policy in the NeurIPS handbook for what should or should not
1338 be described.