

Destroy and Repair Using Hyper-Graphs for Routing

Ke Li^{1, 2}, Fei Liu², Zhengkun Wang^{1*}, Qingfu Zhang²

¹School of System Design and Intelligent Manufacturing, Southern University of Science and Technology

²Department of Computer Science, City University of Hong Kong

12250110@mail.sustech.edu.cn, fliu36-c@my.cityu.edu.hk, wangzhenkun90@gmail.com, qingfu.zhang@cityu.edu.hk

Abstract

Recent advancements in Neural Combinatorial Optimization (NCO) have shown promise in solving routing problems like the Traveling Salesman Problem (TSP) and Capacitated Vehicle Routing Problem (CVRP) without handcrafted designs. Research in this domain has explored two primary categories of methods: iterative and non-iterative. While non-iterative methods struggle to generate near-optimal solutions directly, iterative methods simplify the task by learning local search steps. However, existing iterative methods are often limited by restricted neighborhood searches, leading to suboptimal results. To address this limitation, we propose a novel approach that extends the search to larger neighborhoods by learning a destroy-and-repair strategy. Specifically, we introduce a Destroy-and-Repair framework based on Hyper-Graphs (DRHG). This framework reduces consecutive intact edges to hyper-edges, allowing the model to pay more attention to the destroyed part and decrease the complexity of encoding all nodes. Experiments demonstrate that DRHG achieves state-of-the-art performance on TSP with up to 10,000 nodes and shows strong generalization to real-world TSPLib and CVR-PLib problems.

Code — <https://github.com/CIAM-Group/DRHG>

1 Introduction

Routing problems are significant combinatorial optimization problems with broad real-world applications in logistics, transportation, and manufacturing. Their NP-hard nature poses a significant challenge to the application of exact methods. Heuristics sacrifice the optimality while can obtain near-optimal solutions in a reasonable time. However, the development of heuristics usually relies on human designs with domain expert knowledge, which hinders their practical applications.

Neural Combinatorial Optimization (NCO), which trains a neural network to learn heuristics to solve routing problems without handcraft design, has gained much attention. The existing NCO methods can be roughly classified into two categories: 1) non-iterative and 2) iterative methods.

In non-iterative methods, the neural solvers construct a solution in one shot (Vinyals, Fortunato, and Jaitly 2015; Kool, van Hoof, and Welling 2019; Kwon et al. 2020). Most of these works train neural solvers to determine the next node in an auto-regressive manner, i.e., nodes are selected one by one to be added to the end of a partial solution. Others learn to predict a heuristic, such as a heatmap, and then construct a solution based on the learned information. These works can generate reasonable solutions in a short time. Nevertheless, these non-iterative methods may lead to irreversible consequences if an error occurs in one of the construction steps, thus placing excessive demands on the model’s capability to narrow the optimality gap for large-scale problems.

Iterative methods adopt neural solvers to tackle a subproblem in each iteration rather than solve the entire problem at once. The iterative approach reduces the burden of neural solvers and increases the performance, leading to state-of-the-art results. Some existing iterative NCO methods (de O. da Costa et al. 2020; Wu et al. 2022; Ma et al. 2021) primarily focus on learning low-level operators within small neighborhoods, such as k-opt or swap. Others follow a destroy-and-repair manner, iteratively destroying the solution into a partial solution and then reconstructing the destroyed nodes, operating within a large neighborhood and excelling at producing high-quality solutions. However, the neural networks are trained in a conventional way, either using Reinforcement Learning (RL) (Kim, Park, and Kim 2021; Cheng et al. 2023; Ye et al. 2024; Zheng et al. 2024) or Supervised Learning (SL) (Luo et al. 2023, 2024) without being tailored for the destroy and repair framework. Therefore, they can only deal with the destruction of one segment, bringing challenges in reducing the optimality gap.

To address the issue, we propose a novel iterative NCO method for routing, termed Destroy and Repair by Hyper-Graphs (DRHG). We employ SL to train a model that approaches the best repair after the destruction. Specifically, after the destruction, the complete tour becomes some segments of consecutive edges and some isolated nodes. We reduce the segments to hyper-edges to build a hyper-graph, then fix them during the repair. The model learns to connect isolated nodes and fixed hyper-edges to form a reduced solution, which is restored to a complete solution later. Thanks to the condensed formulation of the hyper-graph, the scale of model input depends only on the degree of the destruction

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

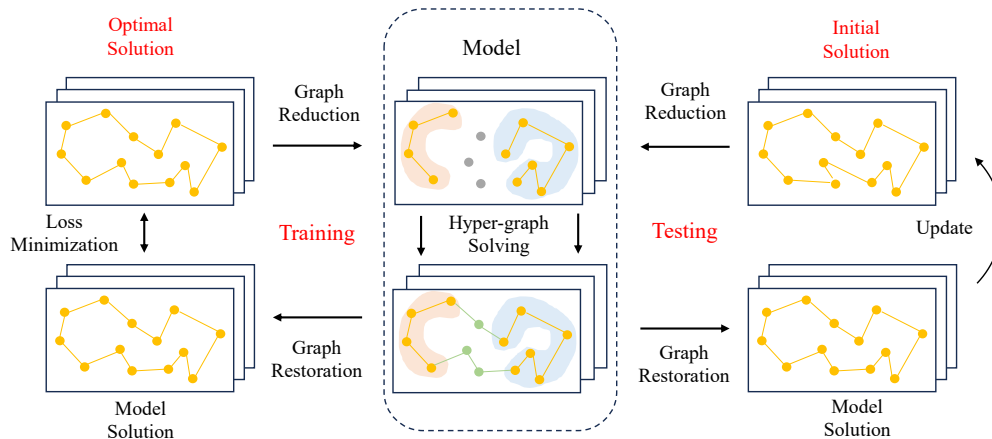


Figure 1: Pipeline of Destroy-and-Repair using Hyper-Graphs

but not the scale of the original problem. This ensures our method has a low computational complexity and allows our method to iterate on large-scale problems.

Our contributions can be summarized as follows:

- We propose a novel NCO framework of destroy-and-repair for routing problems. By learning to repair a destroyed problem in a supervised way, our model can search in large neighborhoods more efficiently.
- We adopt a condensed hyper-graph formulation of the destroyed problem by reducing consecutive edges to fixed hyper-edges, which decreases the computational complexity and enables the model to iterate on large-scale problems.
- The experiments show that our method achieves state-of-the-art performance on TSPs from 100 nodes to 10K nodes, and also competitive results on CVRP. Our method generalizes well to real-world instances as well.

2 Related Works

2.1 Non-iterative NCO Routing Solvers

One-shot Constructive Solvers One-shot constructive methods are one of the earliest lines of work that use NCO to solve routing problems. Pioneering works (Vinyals, Fortunato, and Jaitly 2015; Bello et al. 2017; Nazari et al. 2018) show that neural networks such as RNN can be trained to solve routing problems. Inspired by Vaswani et al. (2017), some works (Kool, van Hoof, and Welling 2019; Deudon et al. 2018) introduce the Transformer architecture to build more powerful NCO models and achieve promising performance. Following their works, various Transformer-based methods (Kwon et al. 2020; Drakulic et al. 2023; Luo et al. 2023) emerged. Although they have made progress in training methods or model structures, the one-shot approach can hardly further narrow the performance gap to the optimal results.

Heatmap-based Solvers Heatmap-based methods aim to predict an informative heatmap to expedite the search pro-

cess and enhance the quality of solutions. Joshi, Laurent, and Bresson (2019) train a Graph Neural Network (GNN) in SL to predict the probabilities of edges to be optimal, then use the beam search to generate feasible solutions. Kool et al. (2022) adopt dynamic programming and eliminate dominated partial solutions to reduce searching time. The most prominent works (Fu, Qiu, and Zha 2021; Sun and Yang 2023) in this category employ Monte Carlo Tree Search (MCTS) to construct solutions. Leveraging MCTS reduces the stringent requirements for the accuracy of edge score predictions. However, most heatmap-based methods are limited to TSPs, as their search strategies are incompatible with problems involving additional constraints, such as CVRPs.

2.2 Iterative NCO Routing Solvers

Most existing iterative NCO routing solvers focus on learning low-level operators searching within small neighborhoods. Chen and Tian (2019) employ a region-picking policy to identify a node for relocation and a rule-picking policy to determine the target position for the node’s movement. de O. da Costa et al. (2020); Sui et al. (2021) propose to learn 2-opt or 3-opt steps to improve the solution. Furthermore, Lu, Zhang, and Yang (2020); Wu et al. (2022) utilize a pool of operators from which the model selects, demonstrating superior performance compared to approaches that rely on a single operator. Ma et al. (2021) propose a Dual-Aspect Collaborative Transformer (DACT) with a Cyclic Positional Encoding (CPE) method and a Dual-Aspect Collaborative Attention (DAC-Att) to encode problems, which achieves pretty good performance. However, iterative NCOs with low-level operators are limited to solving small-size problems due to the extensive number of iterations required for convergence. Moreover, the overall quality of local optimal of small neighborhoods is inferior, which implies that the final solutions obtained by these methods are often sub-optimal.

Other iterative NCO routing solvers focus on reconstructing a partial solution of node sequence. Either trained with RL (Kim, Park, and Kim 2021; Cheng et al. 2023; Ye et al. 2024) or SL (Luo et al. 2023, 2024), the models learn to

reconstruct a segment given the starting node and the ending node. By operating within a large neighborhood, these methods outperform those using low-level operators. However, the neighborhoods that these methods can search in are still limited since the nodes outside the segment remain unaltered. Therefore, two nodes that are spatially close but far away in the solution may have no chance of being reconnected together. In contrast, our framework enables a more flexible neighborhood search by permitting arbitrary destruction and, subsequently, the repair of reconnecting the segments with isolated nodes.

3 Methodology

3.1 DRHG Framework

Schematically illustrated in Fig. 1, we reformulate our destroy-and-repair approach as a graph reduction, hyper-graph solving, and graph restoration process. For a graph representing the incomplete solution where a set of edges is destroyed, we reduce the graph by encoding the remaining consecutive edges as hyper-edges. As these edges remain unchanged during the repair, redefining them as fixed hyper-edges helps reduce the complexity of the problem for the model. Then, we train the model in a supervised way to solve the reduced problem on the hyper-graph. In the testing phase, we iteratively destroy the current solution to obtain a hyper-graph, solve the resulting hyper-graph, and recover the hyper-graph solution on that of the original problem.

3.2 Hyper-graph Representation

Mathematically, a hyper-graph is a special graph where an edge can join any number of vertices. Formally, a hyper-graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, in which \mathcal{V} is the vertex (node) set and \mathcal{E} is the hyper-edge set. Using hypergraph neural networks for embedding hypergraphs is intuitive, but challenging. Specifically, when constructing a solution sequentially, it becomes necessary to align the embeddings of nodes and edges in order to predict the subsequent node or hyper-edge, which may be hard for models. Even if we train an excellent model to predict the sequence, resolving the solution with respect to an undirected hyper-graph remains a non-trivial challenge. Since each hyper-edge has two possible directions, resolving the best complete solution may require a huge number of enumerations. Therefore, we propose to use two endpoints to represent a hyper-edge.

Note a TSP instance of n nodes by the node coordinates as $V = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$. After the destruction, h undestroyed segments constitute the directed hyper-edges set of size $2h$, i.e., $E = \{e^i = (i_1, i_2, \dots, i_{p_i}) \mid i = 1, 2, \dots, 2h\}$, where p_i is the number of nodes in the directed hyper-edge e^i .

For hyper-graph reduction, we remove the middle nodes inside the hyper-edges and keep only the endpoints to represent the hyper-edge, i.e., $e^i = (i_1, i_{p_i})$. Consequently, in the reduced graph, we have one set of isolated nodes A , one set of endpoint nodes B , and one set of reduced hyper-edges E_r . The hyper-graph size is $m = |A| + |B|$. Then, the input feature of the reduced graph is formulated as:

$$r_i = (x_i^a, y_i^a, x_i^b, y_i^b, flag_i), i = 1, 2, \dots, m, \quad (1)$$

$$(x_i^a, y_i^a) = (x_i, y_i), \quad (2)$$

$$(x_i^b, y_i^b) = \begin{cases} (x_i, y_i), & \text{if } i \in A, \\ (x_j, y_j), & \text{if } i \in B \text{ and } (i, j) \in E_r, \end{cases} \quad (3)$$

where r_i is the input feature for the model, and $flag_i$ is a binary variable to indicate whether a node is an endpoint node or an isolated node.

Similarly, for a CVRP instance of n customers and a depot noted as 0, we can define the problem by the node coordinates and the demands: $V = \{(x_0, y_0, 0), (x_1, y_1, d_1), \dots, (x_n, y_n, d_n)\}$, where d_i is the demand of node i . To simplify the problem, we destroyed all edges connected to the depot. Then, the input feature of the reduced graph for CVRP can be formulated as follows:

$$r_i = (x_i^a, y_i^a, x_i^b, y_i^b, flag_i, dr_i), i = 1, 2, \dots, m, \quad (4)$$

$$dr_i = \begin{cases} d_i, & \text{if } i \in A, \\ \sum_k d_k, & \text{if } i \in B \text{ and } k \in (i_1, i_2, \dots, i_{p_i}). \end{cases} \quad (5)$$

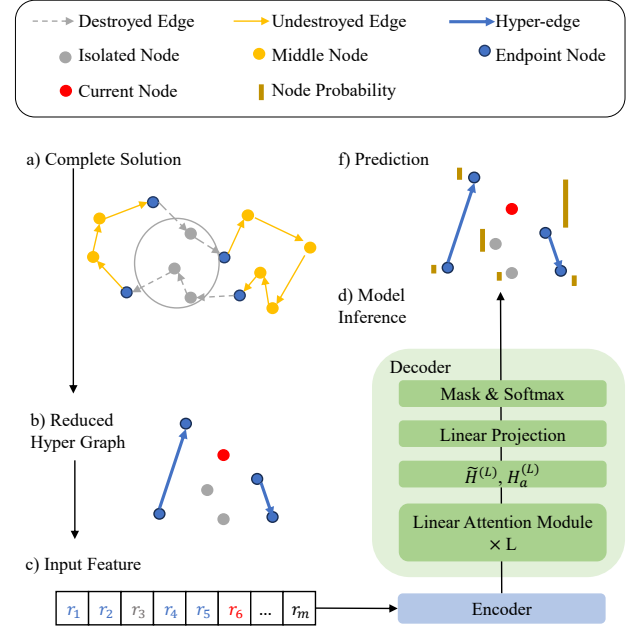


Figure 2: Model structure of DRHG

3.3 Model Structure

As shown in Fig. 2, given the input features of the reduced problems, our model yields a prediction of the next node through a light encoder and a heavy decoder.

Encoder The encoder consists of a single linear projection layer, which transforms the input $r_i \in \mathbb{R}^{d_i}$ into embedding $h_i^{(0)} \in \mathbb{R}^{d_h}$.

Decoder The decoder has a slightly changed linear attention module in Luo et al. (2024). At each step t , the decoder takes the node embeddings of the first node $h_f^{(0)}$, the current node $h_c^{(0)}$, and the remaining unselected nodes $H_a^{(0)} = \{h_i^{(0)} | i = 1, 2, \dots, m - t\}$ as inputs. Then, the first node $h_f^{(0)}$ and the current node $h_c^{(0)}$ are used to generate r virtual representative nodes embeddings $\tilde{H}^{(0)} = \{h_j^{(0)} | j = 1, 2, \dots, r\}$, which combined with $H_a^{(0)}$, form the input of the first linear attention module.

Then, we stack L linear modules as the main component of the decoder. A linear attention module is composed of an aggregating attention layer and a broadcasting attention layer. The aggregating layer aggregates information to the representative nodes, and then the broadcasting layer broadcasts gathered information to all nodes in the graph. The details of the linear attention module are provided in Appendix A. Note the l -th linear attention module as $L - Att^{(l)}$, we have

$$\tilde{H}^{(l)}, H_a^{(l)} = L - Att^{(l)}(\tilde{H}^{(l-1)}, H_a^{(l-1)}). \quad (6)$$

After L attention module, we obtain a hidden representation $\tilde{H}^{(L)}$ and $H_a^{(L)}$. Then we take only $H_a^{(L)}$ to calculate the probability of selecting the next node by a linear projection layer and the softmax function:

$$a_i = \phi(h_i^{(L)} W_o), \quad (7)$$

$$p_i = \frac{e^{a_i}}{\sum_j e^{a_j}}. \quad (8)$$

3.4 Training Scheme

We use SL to train our model. We apply the clustering destruction, as optimal edges are more likely to connect proximal nodes. Furthermore, the distributions of reduced problems after clustering destructions are more consistent across problems of different scales. We adopt the coordinate transformation in Ye et al. (2024) to enhance the distribution homogeneity and consistency. For hyper-edges, once one endpoint is selected, the subsequent node must be the other endpoint. This behavior is dictated by the constraint rather than the model. Correspondingly, we introduce a masking mechanism to block the associated gradients. Additionally, destroying the problem by k -nearest neighbors results in a variable number of segments and makes the hyper-graph size differ across instances. This variability introduces instability during the training process. To tackle this problem, we design a special destruction scheme to get fixed-size hyper-graphs. We detail this method in Appendix B.

4 Experiments

We compare our method with other representative learning-based and classical solvers on TSP and CVRP instances with different scales and the instances in the real world.

4.1 Experiment Setup

Implementing Details We set the embedding dimension of the encoder to 128. The decoder is composed of 6 linear attention modules, and each has 8 attention heads and 16 representative starting nodes. The hidden dimension of the feed-forward layer is set to 512.

For TSP, we train the model for 100 epochs on 1,000,000 TSP100 instances. We fine-tune 20 epochs on 10,000 TSP1000 instances for large-scale problems. For CVRP, we train the model for 100 epochs on 1,000,000 CVRP100 instances. We use a batch size of 1024 and sample the training sample size in $[20, 0.8n]$ where n is the problem size. As the fixed-size destruction scheme will discard a small part of the samples, the true batch size is around 800. We use the cross-entropy loss and the Adam optimizer (Kingma and Ba 2015). The initial learning rate is $1e-4$, and the decay rate is 0.97 per epoch. We train and test our model with a single NVIDIA GeForce RTX 3090 GPU with 24GB memory.

Baselines We compare our method with:

- 1) **Classical Solvers:** Concorde (Applegate et al. 2006), LKH3 (Helsgaun 2017), and HGS (Vidal 2022);
- 2) **Traditional Heuristic:** Random Insertion, Sweep;
- 3) **Construction-based Method:** POMO (Kwon et al. 2020), BQ (Drakulic et al. 2023);
- 4) **Heatmap-based Method:** Att-GCN+MCTS (Fu, Qiu, and Zha 2021), DIMES (Qiu, Sun, and Yang 2022), and DIFUSCO (Sun and Yang 2023);
- 5) **Segment-reconstruction Method:** LEHD (Luo et al. 2023), GLOP (Ye et al. 2024) and SIL (Luo et al. 2024);
- 6) **Operator-iteration Method:** Neural Rewriter (Chen and Tian 2019), Learning 2-Opt (de O. da Costa et al. 2020), Learning 3-Opt (Sui et al. 2021), and DACT (Ma et al. 2021).

For most baseline methods, we run their source code with default settings. The result of Att-GCN+MCTS (Fu, Qiu, and Zha 2021), DIMES (Qiu, Sun, and Yang 2022), and DIFUSCO (Sun and Yang 2023), SIL (Luo et al. 2024), Neural Rewriter (Chen and Tian 2019), and Learning 3-Opt (Sui et al. 2021) are taken from their original papers.

Metrics We use the average objective value (Obj.) and optimality gap (Gap) to evaluate the model performance and the inference time (Time) to evaluate the model efficiency. The ground truth labels of TSP are generated by Concorde for TSP100 to TSP1000 and by LKH for TSP above 1000. The ground truth labels of CVRP are generated by HGS.

Testing We test our method on TSPs from 100 to 10,000 and CVRPs from 100 to 1000. There are 10,000 instances for TSP100 and CVRP100, 128 instances for problems of size 200 to 5,000, and 16 instances for TSP10,000. We use random insertion to generate initial solutions for TSP and sweep for CVRP. Regarding to the k -nn destruction, we sample $k \in [20, \min(1000, n)]$ for TSP and $k \in [20, \min(200, n)]$ for CVRP, where n is the problem scale. For simplicity, we disconnect all nodes adjacent to the depot in CVRP. We set the number of iterations to 1000.

Method	TSP100			TSP200			TSP500		
	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time
LKH3	7.763	0.000%	0.34s	10.704	0.000%	1.88s	16.522	0.000%	15.0s
Concorde	7.763	0.000%	0.20s	10.704	0.000%	1.41s	16.522	0.000%	15.0s
Random Insertion	8.513	9.662%	0.00s	11.948	11.627%	<0.01s	18.546	12.252%	<0.1s
Att-GCN+MCTS*	7.764	0.037%	0.09s	10.814	0.884%	0.94s	16.966	2.537%	2.8s
DIMES*	-	-	-	-	-	-	16.840	1.760%	60.5s
DIFUSCO*	7.780	0.240%	-	-	-	-	16.800	1.490%	1.7s
POMO augx8	7.774 (± 0.231)	0.134% ($\pm 0.224\%$)	0.01s	10.868 (± 0.225)	1.534% ($\pm 0.523\%$)	0.04s	20.187 (± 0.251)	22.187% ($\pm 0.997\%$)	0.5s
BQ bs16	7.764 (± 0.229)	0.015% ($\pm 0.057\%$)	0.17s	10.717 (± 0.208)	0.129% ($\pm 0.149\%$)	0.94s	16.617 (± 0.212)	0.579% ($\pm 0.239\%$)	5.5s
GLOP (more revision)	7.767 (± 0.234)	0.046% ($\pm 0.126\%$)	0.79s	10.774 (± 0.213)	0.653% ($\pm 0.410\%$)	0.33s	16.883 (± 0.214)	2.186% ($\pm 0.474\%$)	0.8s
LEHD RRC1000	7.763 (± 0.229)	0.002% ($\pm 0.014\%$)	1.04s	10.706 (± 0.206)	0.0182% ($\pm 0.054\%$)	4.92s	16.550 (± 0.209)	0.167% ($\pm 0.128\%$)	33.8s
Learning 2-Opt (T=1000)	7.853	1.150%	0.09s	11.107	3.765%	0.20s	21.339	29.158%	0.5s
Learning 3-Opt (T=1000)*	7.850	1.060%	0.23s	-	-	-	-	-	-
DACT (T=1000)	7.892	1.653%	0.07s	12.870	20.252%	0.41s	20.846	26.171%	1.6s
DRHG (T=1000)	7.763 (± 0.229)	0.000% ($\pm 0.007\%$)	2.73s	10.705 (± 0.206)	0.010% ($\pm 0.036\%$)	9.05s	16.540 (± 0.211)	0.111% ($\pm 0.090\%$)	20.6s
Method	TSP1K			TSP5K			TSP10K		
	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time
LKH3	23.12	0.00%	1.7m	50.97	0.00%	12m	71.78	0.00%	33m
Concorde	23.12	0.00%	1m	50.95	-0.05%	31m	72.00	0.15%	1.4h
Random Insertion	26.11	12.90%	<1s	58.06	13.90%	<1s	81.82	13.90%	<1s
Att-GCN+MCTS*	23.86	3.20%	6s	-	-	-	74.93	4.39%	6.6m
DIMES*	23.69	2.46%	2.2m	-	-	-	74.06	3.19%	3m
DIFUSCO*	23.39	1.17%	11.5s	-	-	-	73.62	2.58%	3.0m
POMO augx8	32.51	40.60%	4.1s	87.72	72.10%	8.6m	OOM		
BQ bs16	23.43 (± 0.221)	1.37% ($\pm 0.284\%$)	13s	58.27 (± 0.951)	10.70% ($\pm 1.827\%$)	24s	OOM		
GLOP (more revision)	23.78 (± 0.218)	2.85% ($\pm 0.401\%$)	10.2s	53.15 (± 0.231)	4.26% ($\pm 0.289\%$)	1.0m	75.04 (± 0.215)	4.39% ($\pm 0.153\%$)	1.9m
LEHD RRC1000	23.29 (± 0.220)	0.72% ($\pm 0.176\%$)	3.3m	54.43 (± 0.394)	6.79% ($\pm 0.671\%$)	8.6m	80.90 (± 0.532)	12.50% ($\pm 0.663\%$)	18.6m
SIL PRC1000*	23.31	0.82%	1.2m	51.91	1.84%	7.6m	73.38	2.23%	13.7m
Learning 2-Opt (T=1000)	61.15	164.50%	1.3s	-	-	-	-	-	-
DACT (T=1000)	29.03	25.56%	7.8s	OOM			OOM		
DRHG (T=1000)	23.22	0.45%	1.72m	51.98	2.05%	1.79m	74.38	3.46%	3.63m
DRHG-FT (T=1000)	23.19 (± 0.210)	0.29% ($\pm 0.108\%$)	1.66m	51.39 (± 0.187)	0.88% ($\pm 0.087\%$)	1.82m	72.85 (± 0.217)	1.33% ($\pm 0.084\%$)	3.70m

Table 1: Results on TSP

Competitors	TSP100		TSP200		TSP500	
	Competitor's	Ours	Competitor's	Ours	Competitor's	Ours
POMO augx8	0.134%	1.311% (-)	1.534%	1.269%(+)	22.187%	1.113%(+)
BQ bs16	0.015%	0.025% (-)	0.129%	0.076%(+)	0.579%	0.322%(+)
GLOP (more revision)	0.046%	0.004% (+)	0.653%	0.206%(+)	2.186%	0.879%(+)
LEHD RRC1000	0.002%	0.003% (-)	0.018%	0.016%(+)	0.167%	0.111%(+)

Table 2: Comparison of different methods on TSPs with scale < 1,000 given the same running time

Method	CVRP100			CVRP200		
	Obj.	Gap	Time	Obj.	Gap	Time
LKH3	15.647	0.00%	4.32s	20.173	0.00%	59.06s
HGS	15.563	-0.53%	1.62s	19.946	-1.13%	39.38s
Sweep	20.606	32.14%	<0.01s	0.269	32.78%	<0.01s
POMO augx8	15.754 (±1.800)	0.69% (±0.649%)	0.01s	21.154 (±2.138)	4.87% (±1.133%)	<0.01s
BQ bs16	15.806 (±1.807)	1.02% (±1.015%)	0.11s	20.362 (±2.139)	0.94% (±0.950%)	0.56s
LEHD RRC 1000	15.629 (±1.793)	-0.11% (±0.616%)	1.01s	20.095 (±2.137)	-0.38% (±0.627%)	19.69s
Neural Rewriter*	16.100	-	-	-	-	-
DACT	16.202	3.55%	0.14s	23.230	14.71%	1.02s
DRHG (T=1000)	15.643 (±1.790)	-0.02% (±0.647%)	2.37s	20.233 (±2.133)	-0.16% (±0.648%)	8.91s
Method	CVRP500			CVRP1K		
Method	Obj.	Gap	Time	Obj.	Gap	Time
LKH3	37.229	0.00%	154.69s	37.090	0.00%	199.7s
HGS	36.561	-1.79%	112.50s	36.288	-2.16%	149.1s
Sweep	46.839	25.81%	<0.01s	49.166	32.56%	<0.1s
POMO augx8	44.638 (±3.112)	19.90% (±11.109%)	0.47s	84.898	128.89%	4.7s
BQ bs16	37.606 (±4.216)	1.01% (±0.825%)	3.47s	38.147 (±3.174)	2.88% (±1.258%)	8.7s
GLOP-G(LKH3)	-	-	-	39.651 (±3.779)	6.90% (±2.013%)	0.8s
LEHD RRC 1000	37.100 (±4.257)	-0.35% (±0.534%)	56.25s	37.432 (±3.237)	0.92% (±0.874%)	202.5s
SIL PRC1000*	-	-	-	36.810	-0.76%	78.8s
DACT	46.393	24.98%	3.83s	OOM		
DRHG (T=1000)	37.718 (±4.446)	1.31% (±1.035%)	12.60s	39.932 (±3.854)	7.66% (±2.226%)	12.8s

Table 3: Results on CVRP

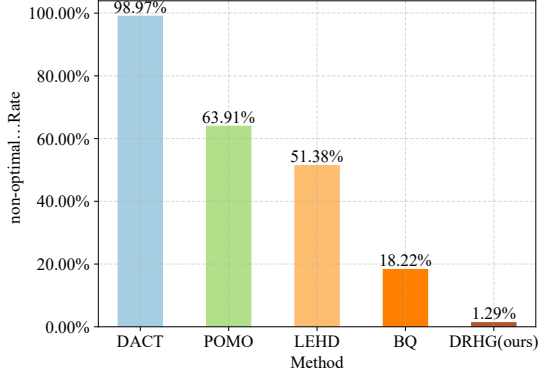


Figure 3: Non-optimal rate on 10k TSP100 instances

4.2 Experimental Results

The main experimental results on uniformly distributed TSP instances are reported in Table 1. All results are reported in terms of per-instance solving time. The values in parentheses represent the variance. Rank-sum tests are conducted on POMO augx8, BQ bs16, GLOP (more revision) and LEHD RRC1000 to assess whether the path length and the gap of given methods differ significantly from those of our method. Except for the LEHD RRC1000 on TSP100, our DRHG demonstrates statistically significant differences ($p < 0.05$) from other methods across all other test settings.

Particularly for TSP100, we track the number of non-optimal cases of the other representative NCO methods compared with our method, which is illustrated in Fig. 3. For TSPs with 100 to 500 nodes, we perform a comparative analysis of our method against other approaches under identical running time in Table 2, where + means our method outperforms its competitor, and vice versa. Since the running time of one-shot constructive solvers cannot be adjusted, we allocate the same running time to our DRHG as that of its competitors.

The results demonstrate that our proposed DRHG method can achieve very good performance on instances of all sizes. Notably, on TSP100, our method yields non-optimal solutions in only 129 out of 10,000 cases, reducing the non-optimality ratio by an order of magnitude. On TSP200 and TSP500, our method reduces the gap by approximately one-third and outperforms its competitors given the same running time. With a small fine-tuning budget, our method outperforms all other methods on large-scale problems, including SIL (Luo et al. 2024), which is separately trained for each problem scale.

For CVRP (Table 3), DRHG can also achieve pretty good performance. Our method outperforms the traditional heuristic method LKH3 on CVRP100 and 200. For CVRP200 and 500, our method outperforms most learning methods except for LEHD RRC1000, which, however, requires much more time. Overall, the performance of DRHG is slightly less dominant than that of TSP, but it is still promising.

Table 4 and Table 5 show the test results on real-world

TSPLib and CVRPLib instances with different sizes and distributions. The results show that our method is robust for different sizes and distributions. More results on TSPLib can be found in Appendix C.

The ablation studies are presented in Appendix D.

size	POMO augx8	BQ bs16	LEHD R. 1K	GLOP more r.	DRHG T=1K
< 100	0.79%	0.49%	0.48%	0.54%	0.48%
100-200	2.42%	1.66%	0.20%	0.79%	0.15%
200-500	13.41%	1.41%	0.38%	1.87%	0.36%
500-1k	31.68%	2.20%	1.21%	3.28%	0.26%
> 1k	63.71%	6.68%	4.14%	7.23%	2.09%
all	26.41%	2.95%	1.59%	3.58%	0.95%

Table 4: Results on TSPLib

size	POMO augx8	BQ bs16	LEHD R. 1K	GLOP more r.	DRHG T=1K
A	4.97%	1.62%	0.75%	26.18%	7.17%
B	4.75%	4.06%	1.09%	20.77%	5.55%
E	11.40%	1.91%	0.58%	18.25%	11.59%
F	15.97%	7.36%	1.36%	39.24%	33.18%
M	4.86%	3.43%	1.43%	22.60%	2.37%
P	15.53%	2.01%	0.93%	17.28%	7.53%
X	21.68%	7.09%	3.69%	18.48%	14.78%
All	15.45%	4.94%	2.36%	20.10%	11.49%

Table 5: Results on CVRPLib

5 Conclusion, Limitation, and Future Work

Conclusion This paper has proposed a novel destroy-and-repair framework using hyper-graphs (DRHG) for routing problems. By leveraging the condensed hyper-graph formulation of the destroyed problem, we have reduced the burden of model learning and constrained the input size of the model to the scale of destruction. Extensive experiments comparing our model with other representative NCO methods on both synthetic and real-world instances have demonstrated the superiority of DRHG across different problem scales and distributions.

Limitation and Future Work The DRHG shows great performance on TSP, but our current design for CVRP has not fully realized the potential of DRHG. It could be interesting to ameliorate the implementation of DRHG and extend it to other routing problems. Furthermore, future work could explore more sophisticated destruction methods other than clustering destruction.

Acknowledgments

This work was supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (GRF Project No. CityU 11215622), the National Natural

Science Foundation of China (Grant No. 62106096 and Grant No. 62476118), the Natural Science Foundation of Guangdong Province (Grant No. 2024A1515011759), the National Natural Science Foundation of Shenzhen (Grant No. JCYJ20220530113013031).

References

- Applegate, D.; Bixby, R.; Chvatal, V.; and Cook, W. 2006. Concorde TSP solver.
- Bello, I.; Pham, H.; Le, Q. V.; Norouzi, M.; and Bengio, S. 2017. Neural Combinatorial Optimization with Reinforcement Learning. arXiv:1611.09940.
- Chen, X.; and Tian, Y. 2019. Learning to Perform Local Rewriting for Combinatorial Optimization. In *Advances in Neural Information Processing Systems 32, December 8-14, 2019, Vancouver, BC, Canada*, volume 32, 6278–6289. Curran Associates, Inc.
- Cheng, H.; Zheng, H.; Cong, Y.; Jiang, W.; and Pu, S. 2023. Select and Optimize: Learning to solve large-scale TSP instances. In *International Conference on Artificial Intelligence and Statistics, 25-27 April 2023, Palau de Congressos, Valencia, Spain*, volume 206 of *Proceedings of Machine Learning Research*, 1219–1231. PMLR.
- de O. da Costa, P. R.; Rhuggenaath, J.; Zhang, Y.; and Akcay, A. 2020. Learning 2-opt Heuristics for the Traveling Salesman Problem via Deep Reinforcement Learning. In *Proceedings of The 12th Asian Conference on Machine Learning, ACML, 18-20 November 2020, Bangkok, Thailand*, volume 129 of *Proceedings of Machine Learning Research*, 465–480. PMLR.
- Deudon, M.; Cournut, P.; Lacoste, A.; Adulyasak, Y.; and Rousseau, L. 2018. Learning Heuristics for the TSP by Policy Gradient. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 15th International Conference, CPAIOR, volume 10848 of Lecture Notes in Computer Science*, 170–181. Springer.
- Drakulic, D.; Michel, S.; Mai, F.; Sors, A.; and Andreoli, J. 2023. BQ-NCO: Bisimulation Quotienting for Efficient Neural Combinatorial Optimization. In *Advances in Neural Information Processing Systems 36*, volume 36, 77416–77429. Curran Associates, Inc.
- Fu, Z.; Qiu, K.; and Zha, H. 2021. Generalize a Small Pre-trained Model to Arbitrarily Large TSP Instances. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI, 7474–7482*. AAAI Press.
- Helsgaun, K. 2017. An Extension of the Lin-Kernighan-Helsgaun TSP Solver for Constrained Traveling Salesman and Vehicle Routing Problems: Technical report. Technical report, Roskilde Universitet.
- Joshi, C. K.; Laurent, T.; and Bresson, X. 2019. An Efficient Graph Convolutional Network Technique for the Travelling Salesman Problem. arXiv:1906.01227.
- Kim, M.; Park, J.; and Kim, J. 2021. Learning Collaborative Policies to Solve NP-hard Routing Problems. In *Advances in Neural Information Processing Systems 34*, volume 34, 10418–10430. Curran Associates, Inc.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kool, W.; van Hoof, H.; Gromicho, J. A. S.; and Welling, M. 2022. Deep Policy Dynamic Programming for Vehicle Routing Problems. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 19th International Conference, CPAIOR, Los Angeles, CA, USA, June 20-23, 2022, Proceedings*, Lecture Notes in Computer Science, 190–213. Springer.
- Kool, W.; van Hoof, H.; and Welling, M. 2019. Attention, Learn to Solve Routing Problems! In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Kwon, Y.; Choo, J.; Kim, B.; Yoon, I.; Gwon, Y.; and Min, S. 2020. POMO: Policy Optimization with Multiple Optima for Reinforcement Learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, volume 33, 21188–21198. Curran Associates, Inc.
- Lu, H.; Zhang, X.; and Yang, S. 2020. A Learning-based Iterative Method for Solving Vehicle Routing Problems. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Luo, F.; Lin, X.; Liu, F.; Zhang, Q.; and Wang, Z. 2023. Neural Combinatorial Optimization with Heavy Decoder: Toward Large Scale Generalization. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, volume 36, 8845–8864. Curran Associates, Inc.
- Luo, F.; Lin, X.; Wang, Z.; Tong, X.; Yuan, M.; and Zhang, Q. 2024. Self-Improved Learning for Scalable Neural Combinatorial Optimization. arXiv:2403.19561.
- Ma, Y.; Li, J.; Cao, Z.; Song, W.; Zhang, L.; Chen, Z.; and Tang, J. 2021. Learning to Iteratively Solve Routing Problems with Dual-Aspect Collaborative Transformer. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 11096–11107. Curran Associates, Inc.
- Nazari, M.; Oroojlooy, A.; Snyder, L. V.; and Takác, M. 2018. Reinforcement Learning for Solving the Vehicle Routing Problem. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 9861–9871. Curran Associates, Inc.
- Qiu, R.; Sun, Z.; and Yang, Y. 2022. DIMES: A Differentiable Meta Solver for Combinatorial Optimization Problems. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, volume 35, 25531–25546. Curran Associates, Inc.

- Sui, J.; Ding, S.; Liu, R.; Xu, L.; and Bu, D. 2021. Learning 3-opt heuristics for traveling salesman problem via deep reinforcement learning. In *Asian Conference on Machine Learning, ACML 2021, 17-19 November 2021, Virtual Event*, volume 157 of *Proceedings of Machine Learning Research*, 1301–1316. PMLR.
- Sun, Z.; and Yang, Y. 2023. DIFUSCO: Graph-based Diffusion Solvers for Combinatorial Optimization. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, volume 36, 3706–3731. Curran Associates, Inc.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5998–6008. Curran Associates, Inc.
- Vidal, T. 2022. Hybrid genetic search for the CVRP: Open-source implementation and SWAP* neighborhood. *Comput. Oper. Res.*, 140: 105643.
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer Networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, 2692–2700. Curran Associates, Inc.
- Wu, Y.; Song, W.; Cao, Z.; Zhang, J.; and Lim, A. 2022. Learning Improvement Heuristics for Solving Routing Problems. *IEEE Trans. Neural Networks Learn. Syst.*, 33(9): 5057–5069.
- Ye, H.; Wang, J.; Liang, H.; Cao, Z.; Li, Y.; and Li, F. 2024. GLOP: Learning Global Partition and Local Construction for Solving Large-Scale Routing Problems in Real-Time. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI, February 20-27, 2024, Vancouver, Canada*, 20284–20292. AAAI Press.
- Zheng, Z.; Zhou, C.; Xialiang, T.; Yuan, M.; and Wang, Z. 2024. UDC: A Unified Neural Divide-and-Conquer Framework for Large-Scale Combinatorial Optimization Problems. arXiv:2407.00312.