Boosting Neural Combinatorial Optimization For Large-Scale Vehicle Routing Problems

Fu Luo[1], Xi Lin[2], Yaoxin Wu[3], Zhenkun Wang[1†], Tong Xialiang[4], Mingxuan Yuan[4], Qingfu Zhang[2]

[1] Southern University of Science and Technology,
[2] City University of Hong Kong,
[3] Eindhoven University of Technology,
[4] Huawei Noah's Ark Lab

†Corresponding author

Vehicle Routing Problem

- Vehicle Routing Problem (VRP)
  - A classic combinatorial optimization problem
  - Has wide applications in logistics, circuit design and other fields.

- Classical VRPs
  - Travelling Salesman Problem (TSP)
  - Capacitated Vehicle Routing Problem (CVRP)

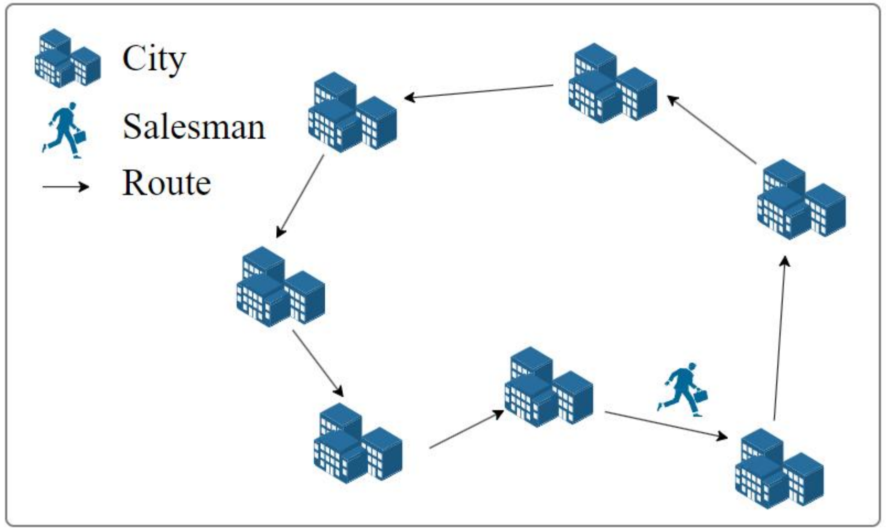

A TSP instance

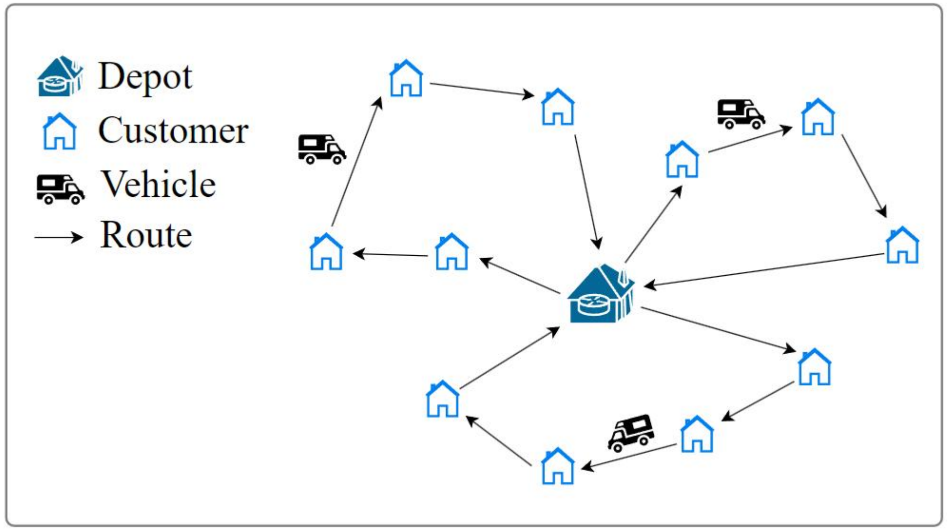

A CVRP instance

Neural Combinatorial Optimization

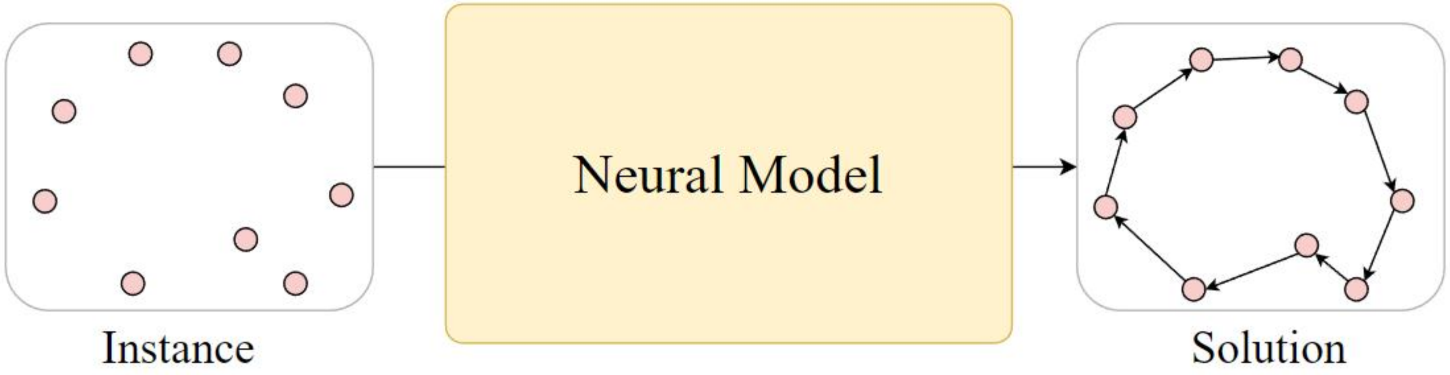

- Neural Combinatorial Optimization (NCO) Methods:
  - Advantage: learn problem-solving policies from data, significantly mitigate the need for costly manual algorithm design.
  - Limitation: struggle when applied to large-scale VRPs (limited scalability).

Obstacles in terms of scalability

- High computational complexity from the self-attention mechanism.

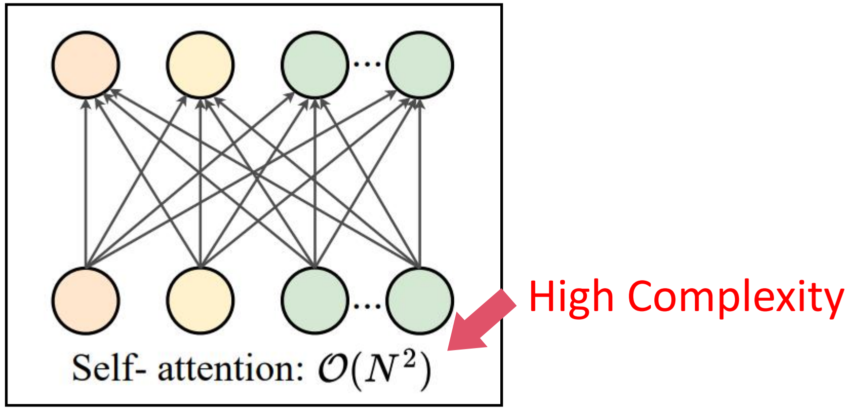

- Difficult to directly train the model on large-scale VRPs
  - Supervised Learning: hard to obtain high-quality labels.
  - Reinforcement Learning: sparse rewards, high GPU memory usage.

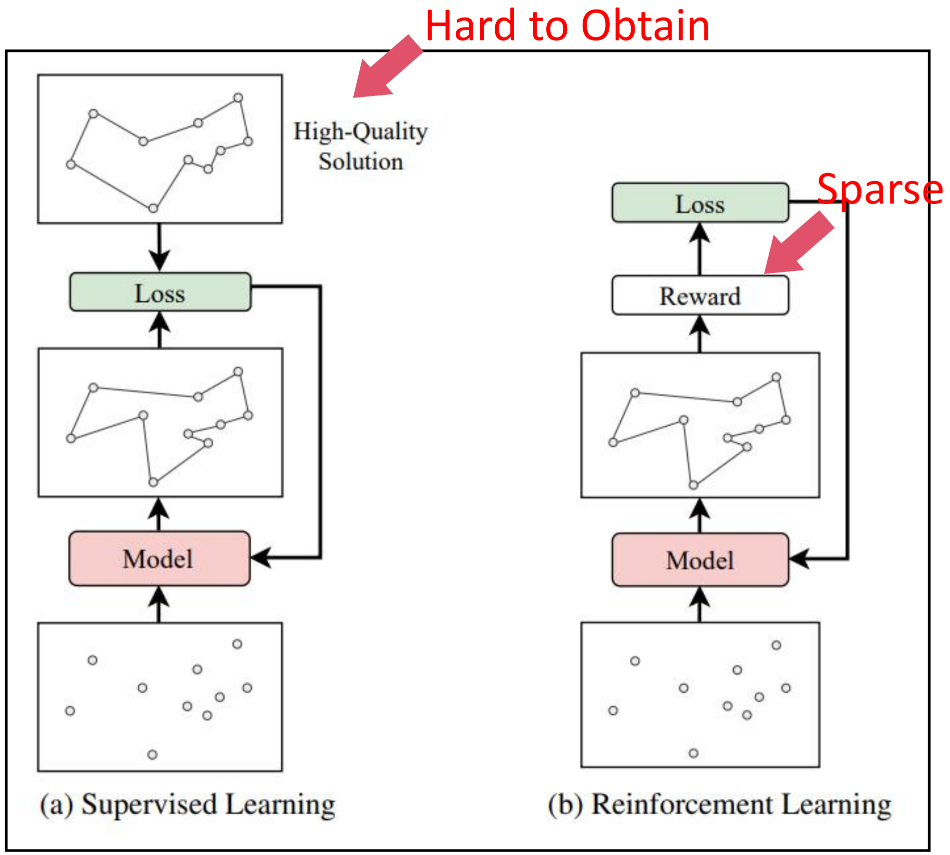

This Work

- Cross-attention mechanism: Linear computational complexity

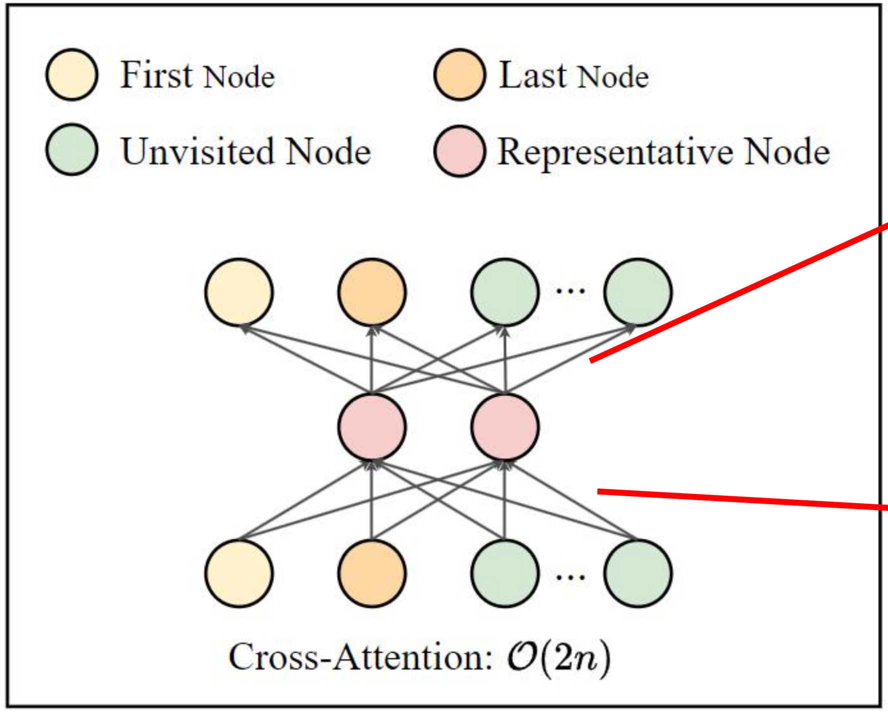

- Step2: update the embeddings of the other nodes by attending them to the representative nodes.
- Step1: update the embeddings of the representative nodes by attending them to the other nodes.

This Work

■ Self-Improved Training: Direct model training on large-scale VRP instances

➢ Local Reconstruction:
Generate enhanced solutions as pseudo-labels by the model itself.

➢ Model Training:
Improve the model's efficiency in enhancing solutions using the pseudo-labels.

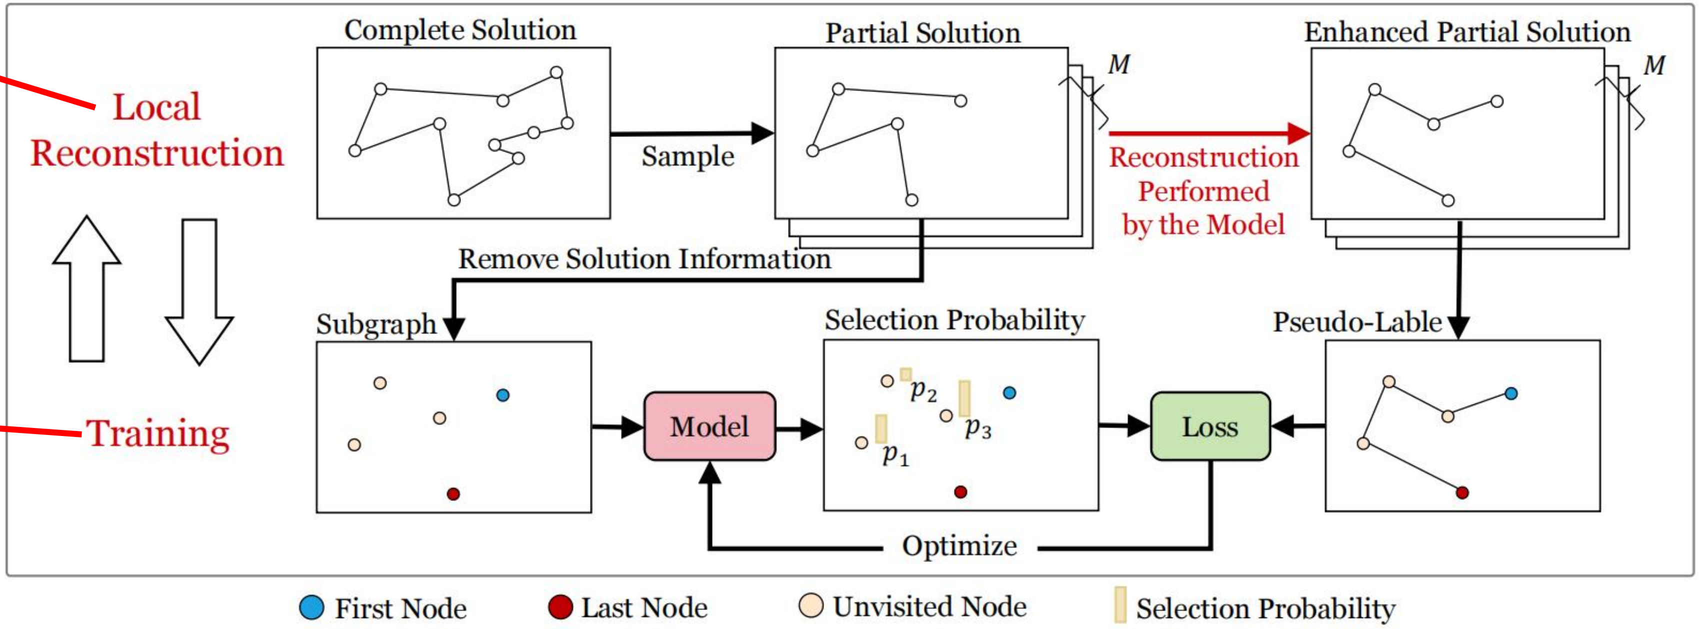

Performance On Uniformly Distributed Instances

Table 1: Comparative results on synthetic TSP and CVRP instances. *: results are cited directly from original publications. N/A: the method exceeds the time limit (e.g., seven days) or produces infeasible solutions. OOM: the method exceeded memory limits.

	TSP1K		TSP5K		TSP10K		TSP50K		TSP100K	
Method	Obj. (Gap)	Time	Obj. (Gap)	Time	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	159.93 (0.00%)	10h	225.99 (0.00%)	25h
Concorde	23.12 (0.00%)	1m	50.95 (-0.05%)	31m	72.00 (0.15%)	1.4h	N/A	N/A	N/A	N/A
Random Insertion	26.11 (12.9%)	<1s	58.06 (13.9%)	<1s	81.82 (13.9%)	<1s	182.65 (14.2%)	15.4s	258.13 (14.2%)	1.7m
DIFUSCO*	23.39 (1.17%)	11.5s	–	–	73.62 (2.58%)	3.0m	–	–	–	–
H-TSP	24.66 (6.66%)	48s	55.16 (8.21%)	1.2m	77.75 (8.38%)	2.2m	OOM		OOM	
GLOP	23.78 (2.85%)	10.2s	53.15 (4.26%)	1.0m	75.04 (4.39%)	1.9m	168.09 (5.10%)	1.5m	237.61 (5.14%)	3.9m
POMO aug×8	32.51 (40.6%)	4.1s	87.72 (72.1%)	8.6m	OOM		OOM		OOM	
ELG aug×8	25.738 (11.33%)	0.8s	60.19 (18.08%)	21s	OOM		OOM		OOM	
LEHD RRC1,000	23.29 (0.72%)	3.3m	54.43 (6.79%)	8.6m	80.90 (12.5%)	18.6m	OOM		OOM	
BQ bs16	23.43 (1.37%)	13s	58.27 (10.7%)	24s	OOM		OOM		OOM	
SIGD bs16	23.36 (1.03%)	17.3s	55.77 (9.42%)	30.5m	OOM		OOM		OOM	
INViT-3V greedy	24.66 (6.66%)	9.0s	54.49 (6.90%)	1.2m	76.85 (7.07%)	3.7m	171.42 (7.18%)	1.3h	242.26 (7.20%)	5.0h
LEHD greedy	23.84 (3.11%)	0.8s	58.85 (15.46%)	1.5m	91.33 (27.24%)	11.7m	OOM		OOM	
BQ greedy	23.65 (2.30%)	0.9s	58.27 (14.31%)	22.5s	89.73 (25.02%)	1.0m	OOM		OOM	
SIGD greedy	23.573 (1.96%)	1.2s	57.19 (12.20%)	1.8m	93.80 (30.68%)	15.5m	OOM		OOM	
Ours greedy	23.569 (1.95%)	0.2s	52.59 (3.17%)	5.2s	74.69 (4.05%)	20.1s	168.50 (5.36%)	7.7m	239.84 (6.13%)	33.0m
Ours PRC10	23.396 (1.20%)	0.9s	52.36 (2.73%)	5.1s	73.99 (3.08%)	10.0s	166.69 (4.22%)	1.33m	235.38 (4.16%)	3.0m
Ours PRC50	23.279 (0.69%)	4.6s	51.92 (1.85%)	23.4s	73.41 (2.27%)	49.0s	165.01 (3.17%)	4.9m	233.13 (3.16%)	9.2m
Ours PRC100	23.254 (0.58%)	9.4s	51.82 (1.67%)	52.0s	73.29 (2.11%)	1.7m	164.59 (2.91%)	8.6m	232.55 (2.90%)	17m
Ours PRC500	23.217 (0.43%)	46s	51.70 (1.43%)	4.6m	73.12 (1.87%)	8.5m	164.09 (2.60%)	42.2m	231.75 (2.55%)	1.4h
Ours PRC1,000	**23.207 (0.38%)**	1.5m	**51.67 (1.36%)**	9.4m	**73.08 (1.81%)**	17.0m	**163.95 (2.51%)**	1.38h	**231.52 (2.45%)**	2.6h

Performance On Uniformly Distributed Instances

Method	CVRP1K Obj. (Gap)	CVRP1K Time	CVRP5K Obj. (Gap)	CVRP5K Time	CVRP10K Obj. (Gap)	CVRP10K Time	CVRP50K Obj. (Gap)	CVRP50K Time	CVRP100K Obj. (Gap)	CVRP100K Time
HGS	36.29 (0.00%)	2.5m	89.74 (0.00%)	2.0h	107.40 (0.00%)	5.0h	267.73 (0.00%)	8.1h	476.11 (0.00%)	24h
LKH3	37.09 (2.21%)	3.3m	93.71 (5.19%)	1.33h	118.76 (10.6%)	1.74h	399.12 (49.1%)	15.8h	N/A	N/A
Random Insertion	57.42 (58.2%)	<1s	154.38 (72.0%)	<1s	191.80 (78.6%)	<1s	490.56 (83.2%)	<1s	943.87 (98.3%)	2s
GLOP-G (LKH3)	39.50 (8.83%)	1.3s	98.90 (10.2%)	6.8s	116.28 (8.27%)	11.2s	OOM		OOM	
POMO aug×8	84.89 (134%)	4.8s	393.27 (338%)	11m	OOM		OOM		OOM	
ELG aug×8	41.57 (14.56%)	1.1s	109.54 (22.06%)	30s	OOM		OOM		OOM	
LEHD RRC1,000	37.43 (3.15%)	3.4m	101.07 (12.6%)	31m	138.73 (29.2%)	41m	OOM		OOM	
BQ bs16	38.17 (5.17%)	14s	104.40 (16.3%)	2.6m	OOM		OOM		OOM	
SIGD bs16	39.15 (7.91%)	17.3s	103.46 (15.3%)	1.91m	131.48 (22.4%)	3.97m	477.43 (78.3%)	25.9m	OOM	
INViT-3V greedy	42.75 (17.8%)	11.4s	109.85 (22.41%)	1.4m	141.41 (31.66%)	4.2m	402.05 (50.17%)	2.9h	688.80 (44.67%)	8.3h
LEHD greedy	38.91 (7.23%)	0.8s	105.61 (17.69%)	1.56m	146.24 (36.16%)	11.85m	OOM		OOM	
BQ greedy	39.28 (8.23%)	1.03s	108.09 (20.48%)	8.1s	196.44 (82.9%)	1.2m	OOM		OOM	
SIGD greedy	40.18 (10.7%)	1.2s	106.14 (18.3%)	7.9s	135.12 (25.8%)	45s	493.64 (84.4%)	4.3m	OOM	
Ours greedy	38.11 (5.01%)	0.2s	92.44 (3.01%)	5.49s	109.02 (1.50%)	20.62s	269.34 (0.60%)	8.06m	475.06 (-0.22%)	33.1m
Ours PRC10	37.93 (4.52%)	0.7s	93.92 (4.65%)	3.9s	112.17 (4.43%)	6.8s	285.20 (6.52%)	28s	496.24 (4.23%)	59s
Ours PRC50	37.57 (3.54%)	3.5s	92.06 (2.58%)	19.9s	108.79 (1.29%)	34s	271.77 (1.51%)	2.3m	476.71 (0.13%)	4.8m
Ours PRC100	37.49 (3.31%)	8.0s	91.58 (2.05%)	46s	108.04 (0.59%)	1.3m	268.02 (0.11%)	5.49m	471.35 (-1.00%)	11.5m
Ours PRC500	37.33 (2.88%)	44.6s	91.00 (1.41%)	4.4m	106.85 (-0.51%)	7.6m	263.56 (-1.56%)	31.1m	465.18 (-2.30%)	1.1h
Ours PRC1,000	**37.28 (2.72%)**	1.5m	**90.81 (1.19%)**	8.8m	**106.69 (-0.66%)**	15.2m	**262.82 (-1.83%)**	1.04h	**463.95 (-2.55%)**	2.17h

Performance On Real-World Instances

Table 2: Results on TSPLIB and CVRPLIB. OOM: The method is inapplicable due to the memory limit. †: There exist instances that are not solvable by the NCO method due to the OOM issue.

	TSPLIB				CVRPLIB			
Method	$1K < n \leq 5K$	$n > 5K$	All	Solved #	$1K < n \leq 7K$	$n > 7K$	All	Solved #
GLOP	5.017%	6.870%†	5.495%	31/33	15.335%	21.317%	17.898%	14/14
ELG aug×8	11.34%	OOM	11.34%	23/33	10.57%†	OOM	10.57%	6/14
BQ bs16	10.648%	30.579%†	12.948%	26/33	13.918%	OOM	13.918%	8/14
LEHD RRC1,000	3.996%	18.458%†	7.371%	30/33	8.423%	21.525%†	11.043%	10/14
SIGD greedy	12.369%	152.879%†	48.630%	31/33	14.733%	49.491%	29.629%	14/14
BQ greedy	11.640%	162.116%†	64.649%	32/33	16.923%	52.267%	32.071%	14/14
INViT greedy	11.492%	9.996%	11.038%	33/33	15.873%	26.637%	20.486%	14/14
LEHD greedy	11.139%	39.343%†	17.720%	30/33	15.203%	32.797%†	18.722%	10/14
Ours greedy	6.767%	10.697%	8.244%	33/33	15.806%	15.504%	15.677%	14/14
Ours PRC1,000	**1.576%**	**4.043%**	**2.556%**	33/33	**8.347%**	**11.209%**	**9.574%**	14/14

Summary

- Cross-Attention:
  Linear complexity

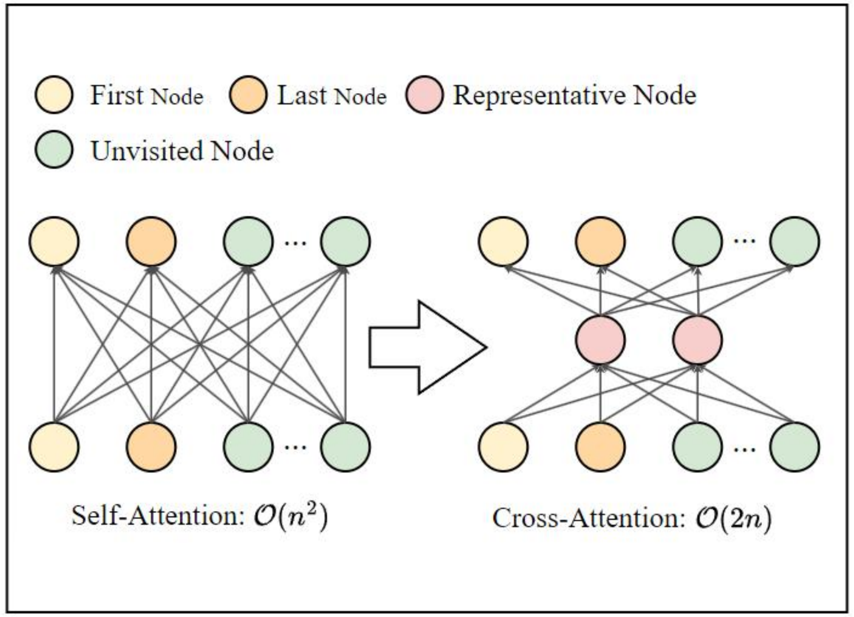

- Self-Improved Training:
  Direct model training on large-scale VRP instances