Position: Rethinking Post-Hoc Search-Based Neural Approaches for Solving Large-Scale Traveling Salesman Problems

Yifan Xia   Xianliang Yang   Zichuan Liu
Zhihao Liu   Lei Song   Jiang Bian

The 41th International Conference on Machine Learning (ICML 2024)

Current Paradigm for Solving Large-Scale TSP

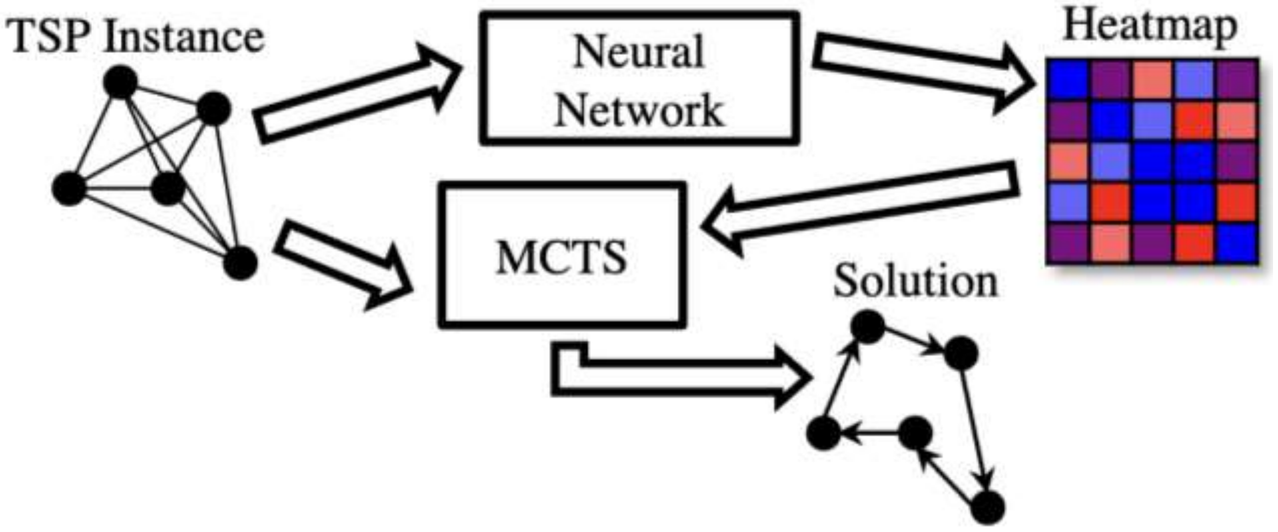

Current Paradigm for Solving Large-Scale TSP

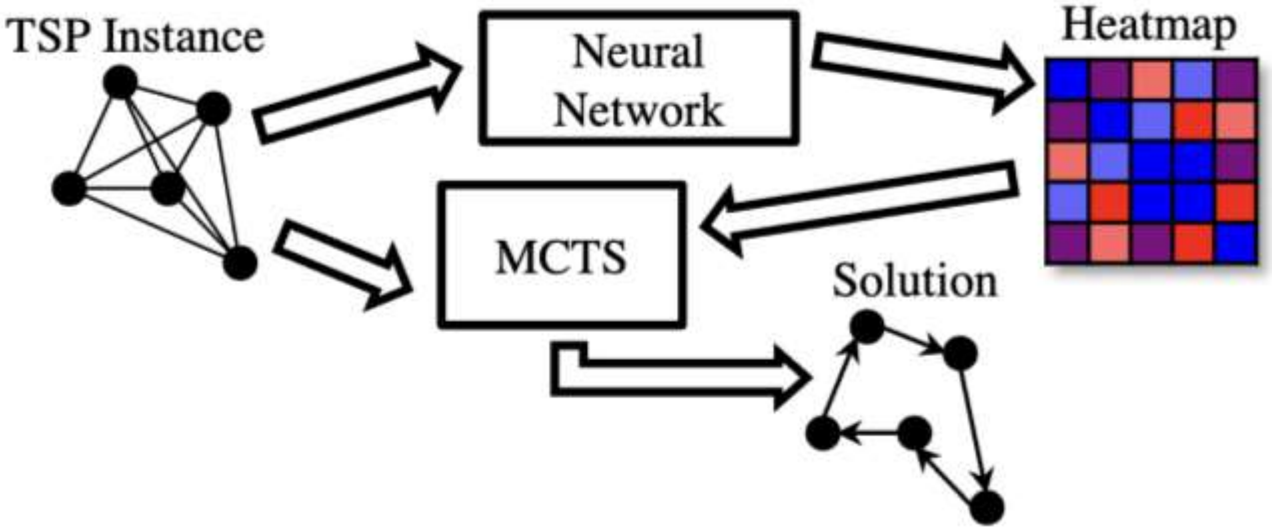

Current Paradigm for Solving Large-Scale TSP

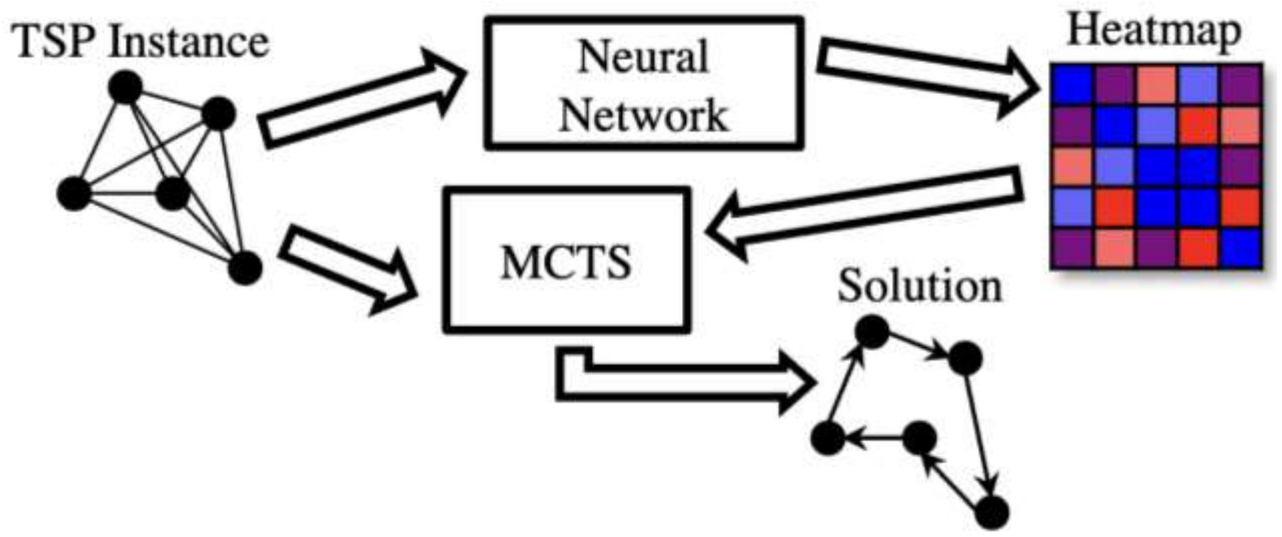

Current Paradigm for Solving Large-Scale TSP

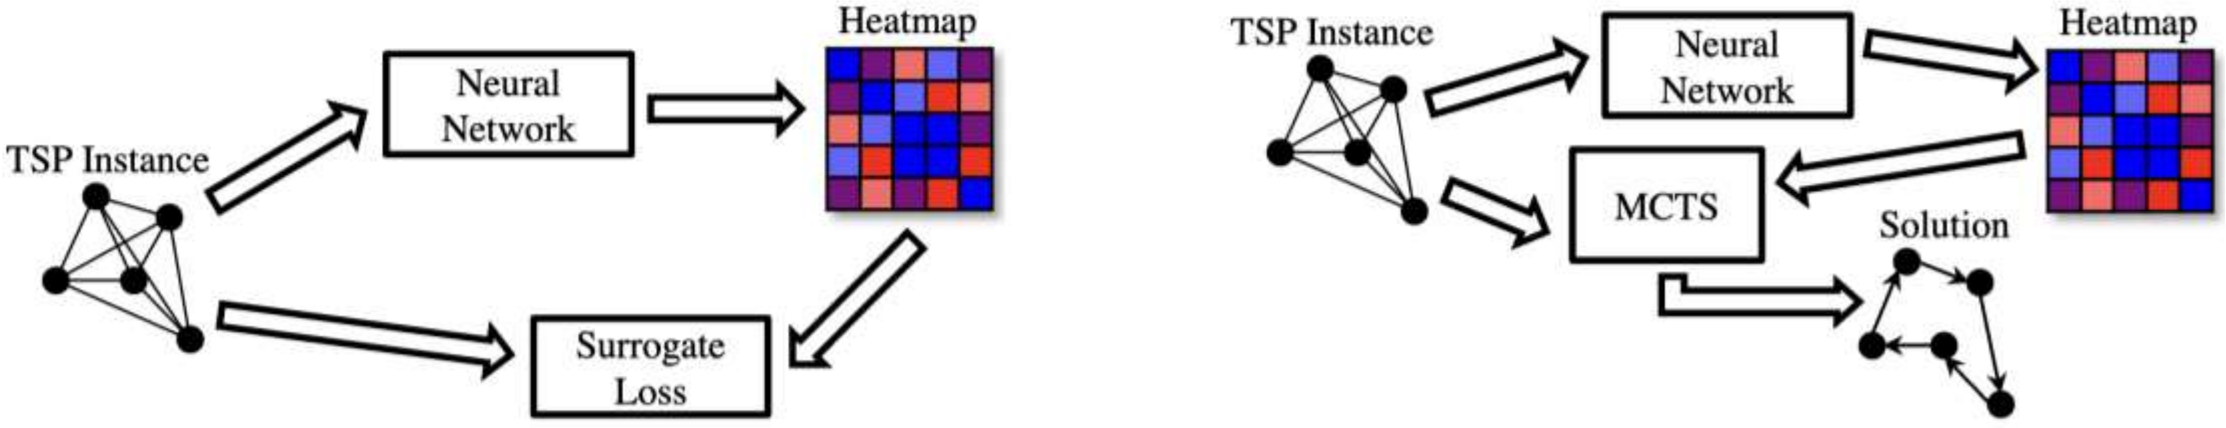

(a) Training Phase (b) Test Phase

What’s Wrong with the Current Paradigm?

1. Do those ML models really learn to generate heatmaps?
   - Misaligned training and test phases

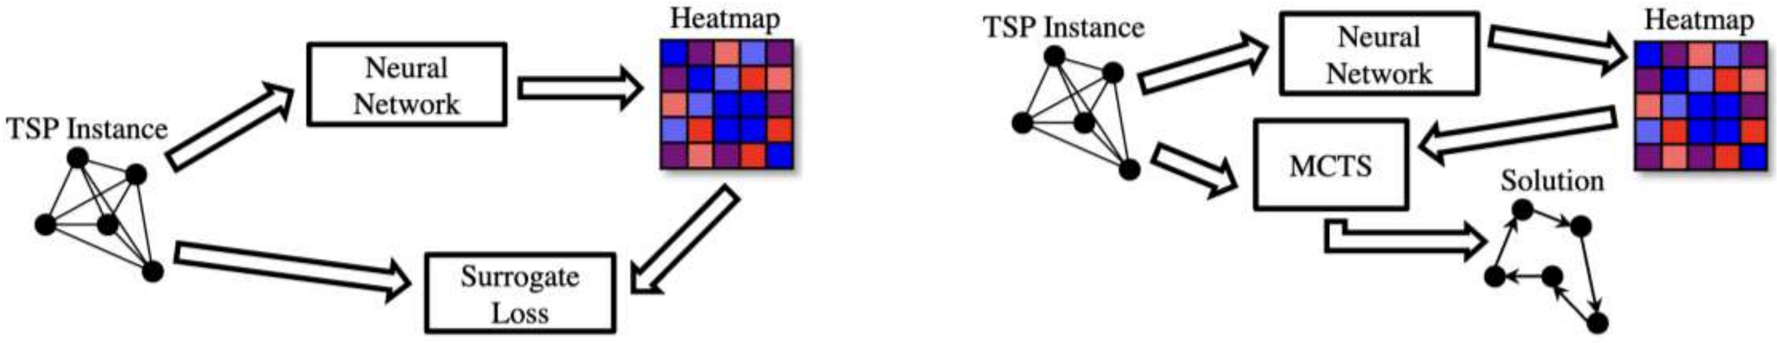

(a) Training Phase

(b) Test Phase

What’s Wrong with the Current Paradigm?

1. Do those ML models really learn to generate heatmaps?
    - Misaligned training and test phases

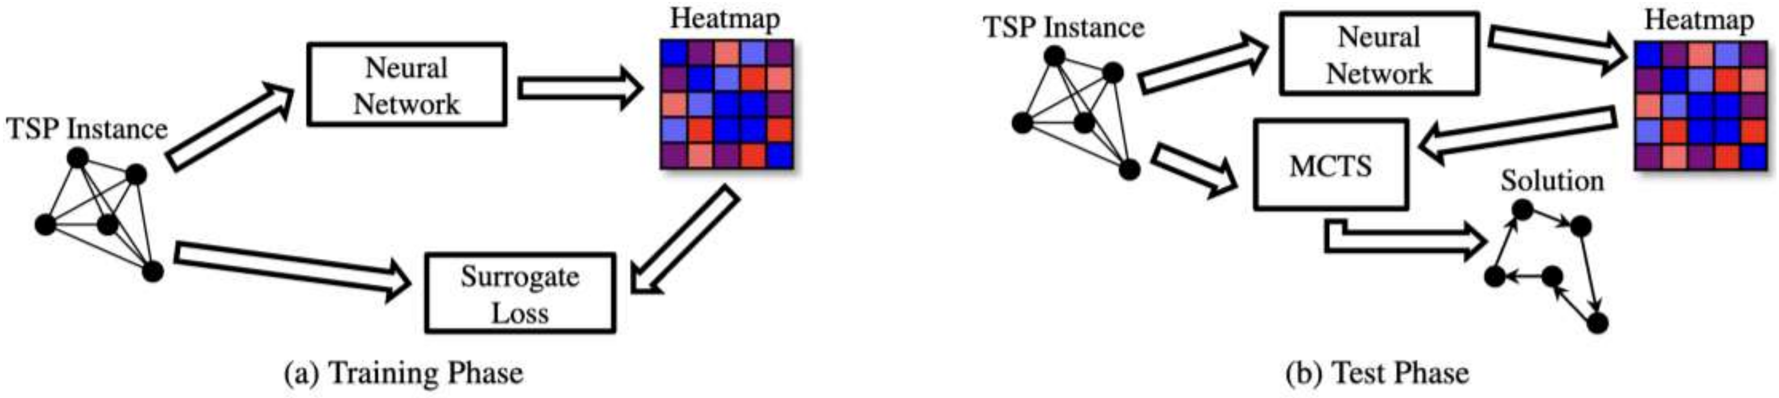

    - Original loss: $$\mathcal{L}(\theta) = \mathbb{E}_{s\sim\mathcal{S}}\left[\mathbb{E}_{\Phi\sim f_\theta(s)}\left[\mathbb{E}_{\boldsymbol{\pi}\sim g(s,\Phi)}\left[c\left(\boldsymbol{\pi}\right)\right]\right]\right]$$
    - Surrogate loss: $$\mathcal{L}_{surrogate}(\theta) = \mathbb{E}_{s\sim\mathcal{S}}\left[\mathbb{E}_{\Phi\sim f_\theta(s)}\left[\ell\left(s,\Phi\right)\right]\right]$$

What’s Wrong with the Current Paradigm?

1. Do those ML models really learn to generate heatmaps?
    - Misaligned training and test phases

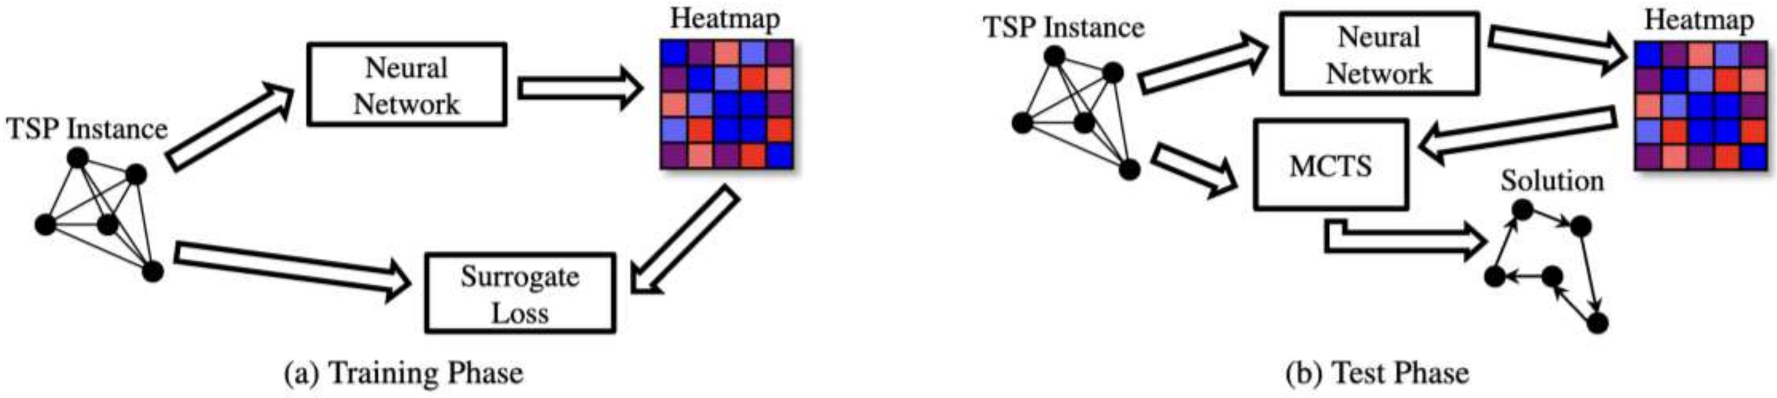

2. With reliance on MCTS algorithm, why not use LKH-3 instead?
    - Both are handcrafted heuristics, relying heavily on expert knowledge

What’s Wrong with the Current Paradigm?

- k-opt heuristic

(1) initial tour (2) breaking k edges (3) reconnecting

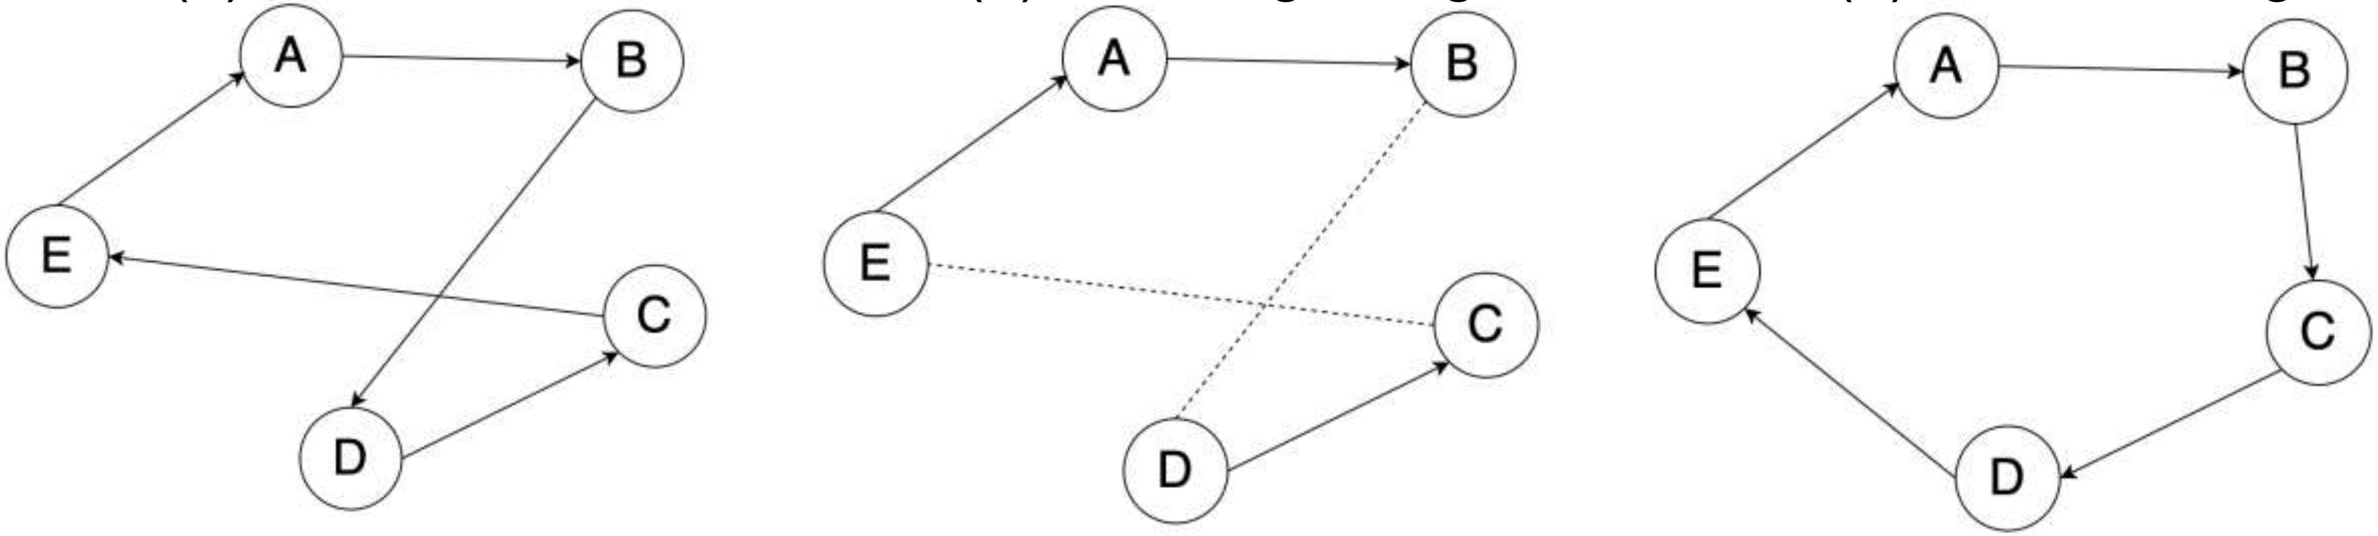

2. With reliance on MCTS algorithm, why not use LKH-3 instead?
   - Both are handcrafted heuristics, relying heavily on expert knowledge

Demonstrating ML's inefficiency in Generating Heatmaps

- Simplest heatmap generation method: $\Phi_{i,j} = \frac{e^{-d_{i,j}/\tau}}{\sum_{k \neq i} e^{-d_{i,k}/\tau}}$

Demonstrating ML's inefficiency in Generating Heatmaps

- Simplest heatmap generation method: $\Phi_{i,j} = \frac{e^{-d_{i,j}/\tau}}{\sum_{k \neq i} e^{-d_{i,k}/\tau}}$

```
heatmap = F.softmax(-distance_matrix / tau, dim=2)
```

Demonstrating ML's inefficiency in Generating Heatmaps

- Simplest heatmap generation method: $\Phi_{i,j} = \frac{e^{-d_{i,j}/\tau}}{\sum_{k \neq i} e^{-d_{i,k}/\tau}}$

- Align training and test phases: simple heatmap generation methods can outperform ML

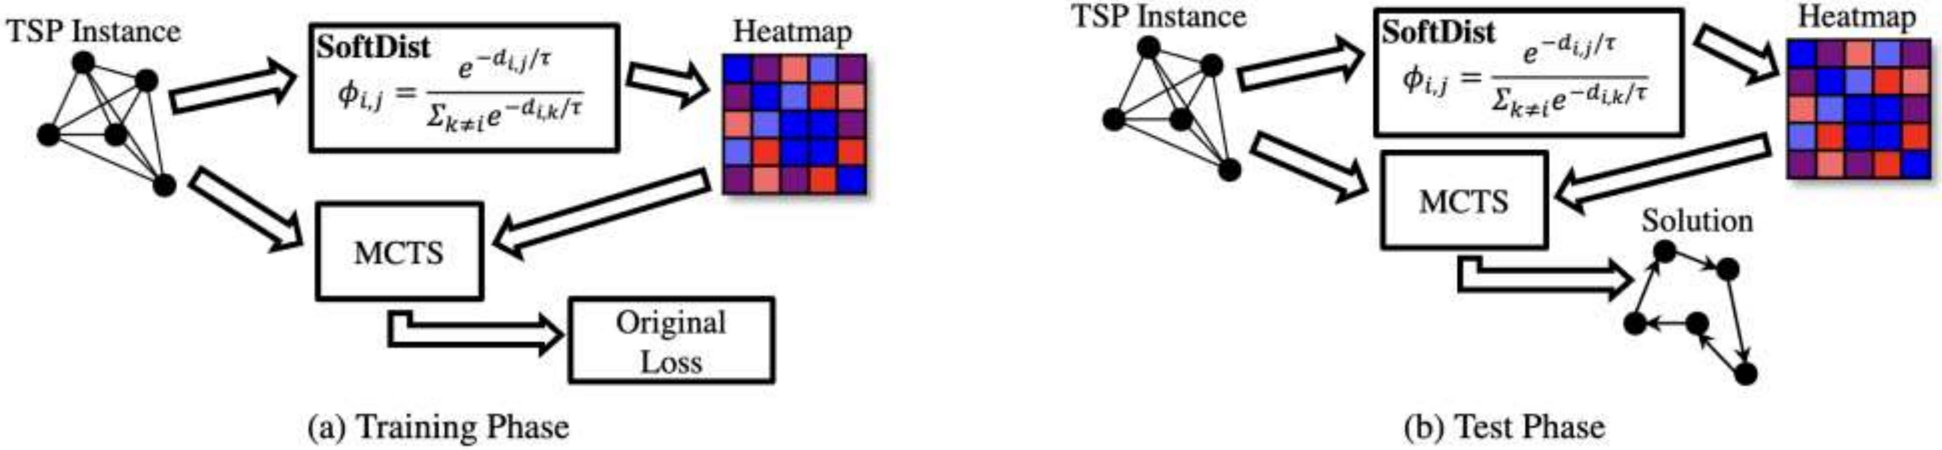

Figure 3. SoftDist-Guided MCTS Phases.

Demonstrating MCTS's Lack of Practicality Compared to LKH-3

- Fair comparison between MCTS and LKH-3
  - Same CPU, same thread count, same search time

Demonstrating MCTS's Lack of Practicality Compared to LKH-3

- Fair comparison between MCTS and LKH-3
  - Same CPU, same thread count, same search time
- What percentage of LKH-3's performance does MCTS achieve?

Demonstrating MCTS's Lack of Practicality Compared to LKH-3

- Fair comparison between MCTS and LKH-3
  - Same CPU, same thread count, same search time
- What percentage of LKH-3's performance does MCTS achieve?
  - A novel metric to measure the ratio of the performance gaps of LKH-3 and MCTS

$$Score = \frac{Gap_{\text{LKH-3}}}{Gap_{\text{MCTS}}}$$

Experiment 1

- Heatmaps generated by ML underperform those generated by simpler methods

Method	Type	TSP-500 Length ↓	TSP-500 Gap ↓	TSP-500 Time ↓	TSP-1000 Length ↓	TSP-1000 Gap ↓	TSP-1000 Time ↓	TSP-10000 Length ↓	TSP-10000 Gap ↓	TSP-10000 Time ↓
Concorde	OR(exact)	16.55*	—	37.66m	23.12*	—	6.65h	N/A	N/A	N/A
Gurobi	OR(exact)	16.55	0.00%	45.63h	N/A	N/A	N/A	N/A	N/A	N/A
LKH-3 (default)	OR	16.55	0.00%	46.28m	23.12	0.00%	2.57h	71.78*	—	8.8h
LKH-3 (less trails)	OR	16.55	0.00%	3.03m	23.12	0.00%	7.73m	71.79	—	51.27m
Nearest Insertion	OR	20.62	24.59%	0s	28.96	25.26%	0s	90.51	26.11%	6s
Random Insertion	OR	18.57	12.21%	0s	26.12	12.98%	0s	81.85	14.04%	4s
Farthest Insertion	OR	18.30	10.57%	0s	25.72	11.25%	0s	80.59	12.29%	6s
EAN	RL+S	28.63	73.03%	20.18m	50.30	117.59%	37.07m	N/A	N/A	N/A
EAN	RL+S+2-opt	23.75	43.57%	57.76m	47.73	106.46%	5.39h	N/A	N/A	N/A
AM	RL+S	22.64	36.84%	15.64m	42.80	85.15%	63.97m	431.58	501.27%	12.63m
AM	RL+G	20.02	20.99%	1.51m	31.15	34.75%	3.18m	141.68	97.39%	5.99m
AM	RL+BS	19.53	18.03%	21.99m	29.90	29.23%	1.64h	129.40	80.28%	1.81h
GCN	SL+G	29.72	79.61%	6.67m	48.62	110.29%	28.52m	N/A	N/A	N/A
GCN	SL+BS	30.37	83.55%	38.02m	51.26	121.73%	51.67m	N/A	N/A	N/A
POMO+EAS-Emb	RL+AS	19.24	16.25%	12.80h	N/A	N/A	N/A	N/A	N/A	N/A
POMO+EAS-Lay	RL+AS	19.35	16.92%	16.19h	N/A	N/A	N/A	N/A	N/A	N/A
POMO+EAS-Tab	RL+AS	24.54	48.22%	11.61h	49.56	114.36%	63.45h	N/A	N/A	N/A
DIFUSCO#	SL+MCTS	16.63	0.51%	3.61m+ 1.67m	23.39	1.18%	11.86m+ 3.34m	73.76	2.77%	28.51m+ 16.87m
ATT-GCN†	SL+MCTS	16.82	1.64%	0.52m+ 1.67m	23.67	2.37%	0.73m+ 3.34m	74.50	3.80%	4.16m+ 16.77m
DIMES†	RL+MCTS	16.84	1.77%	0.97m+ 1.67m	23.68	2.44%	2.08m+ 3.34m	74.10	3.23%	4.65m+ 16.77m
UTSP†	UL+MCTS	17.11	3.41%	1.37m+ 1.67m	24.14	4.40%	3.35m+ 3.34m	—	—	—
Ours	SoftDist+MCTS	**16.78**	**1.44%**	**0.00m+ 1.67m**	**23.63**	**2.20%**	**0.00m+ 3.34m**	**74.03**	**3.13%**	**0.00m+ 16.78m**

Experiment 1

- Heatmaps generated by ML underperform those generated by simpler methods

Method	Type	TSP-500 Length ↓	TSP-500 Gap ↓	TSP-500 Time ↓	TSP-1000 Length ↓	TSP-1000 Gap ↓	TSP-1000 Time ↓	TSP-10000 Length ↓	TSP-10000 Gap ↓	TSP-10000 Time ↓
Concorde	OR(exact)	16.55*	—	37.66m	23.12*	—	6.65h	N/A	N/A	N/A
Gurobi	OR(exact)	16.55	0.00%	45.63h	N/A	N/A	N/A	N/A	N/A	N/A
LKH-3 (default)	OR	16.55	0.00%	46.28m	23.12	0.00%	2.57h	71.78*	—	8.8h
LKH-3 (less trails)	OR	16.55	0.00%	3.03m	23.12	0.00%	7.73m	71.79	—	51.27m
Nearest Insertion	OR	20.62	24.59%	0s	28.96	25.26%	0s	90.51	26.11%	6s
Random Insertion	OR	18.57	12.21%	0s	26.12	12.98%	0s	81.85	14.04%	4s
Farthest Insertion	OR	18.30	10.57%	0s	25.72	11.25%	0s	80.59	12.29%	6s
EAN	RL+S	28.63	73.03%	20.18m	50.30	117.59%	37.07m	N/A	N/A	N/A
EAN	RL+S+2-opt	23.75	43.57%	57.76m	47.73	106.46%	5.39h	N/A	N/A	N/A
AM	RL+S	22.64	36.84%	15.64m	42.80	85.15%	63.97m	431.58	501.27%	12.63m
AM	RL+G	20.02	20.99%	1.51m	31.15	34.75%	3.18m	141.68	97.39%	5.99m
AM	RL+BS	19.53	18.03%	21.99m	29.90	29.23%	1.64h	129.40	80.28%	1.81h
GCN	SL+G	29.72	79.61%	6.67m	48.62	110.29%	28.52m	N/A	N/A	N/A
GCN	SL+BS	30.37	83.55%	38.02m	51.26	121.73%	51.67m	N/A	N/A	N/A
POMO+EAS-Emb	RL+AS	19.24	16.25%	12.80h	N/A	N/A	N/A	N/A	N/A	N/A
POMO+EAS-Lay	RL+AS	19.35	16.92%	16.19h	N/A	N/A	N/A	N/A	N/A	N/A
POMO+EAS-Tab	RL+AS	24.54	48.22%	11.61h	49.56	114.36%	63.45h	N/A	N/A	N/A
DIFUSCO#	SL+MCTS	16.63	0.51%	3.61m+ 1.67m	23.39	1.18%	11.86m+ 3.34m	73.76	2.77%	28.51m+ 16.87m
Att-GCN†	SL+MCTS	16.82	1.64%	0.52m+ 1.67m	23.67	2.37%	0.73m+ 3.34m	74.50	3.80%	4.16m+ 16.77m
DIMES†	RL+MCTS	16.84	1.77%	0.97m+ 1.67m	23.68	2.44%	2.08m+ 3.34m	74.10	3.23%	4.65m+ 16.77m
UTSP†	UL+MCTS	17.11	3.41%	1.37m+ 1.67m	24.14	4.40%	3.35m+ 3.34m	—	—	—
Ours	SoftDist+MCTS	**16.78**	**1.44%**	**0.00m+ 1.67m**	**23.63**	**2.20%**	**0.00m+ 3.34m**	**74.03**	**3.13%**	**0.00m+ 16.78m**

Experiment 2

- MCTS-based paradigm underperforms LKH-3, even though it consumes more resources

Method	Supervision	Hardware	Score ↑ TSP-500	TSP-1000	TSP-10000
ATT-GCN†	✓	GTX 1080 Ti GPU	0.74%	3.87%	24.66%
DIMES†	✗	NVIDIA P100 GPU	0.68%	3.75%	28.99%
UTSP†	✗	NVIDIA V100 GPU	0.35%	2.08%	—
DIFUSCO#	✓	8×NVIDIA V100 GPUs	2.39%	7.78%	33.82%
SoftDist	✗	NVIDIA A100 GPU	0.84%	4.17%	29.88%

Experiment 3

- In both time-sensitive and non-time-sensitive scenarios, LKH-3 is a more practical choice

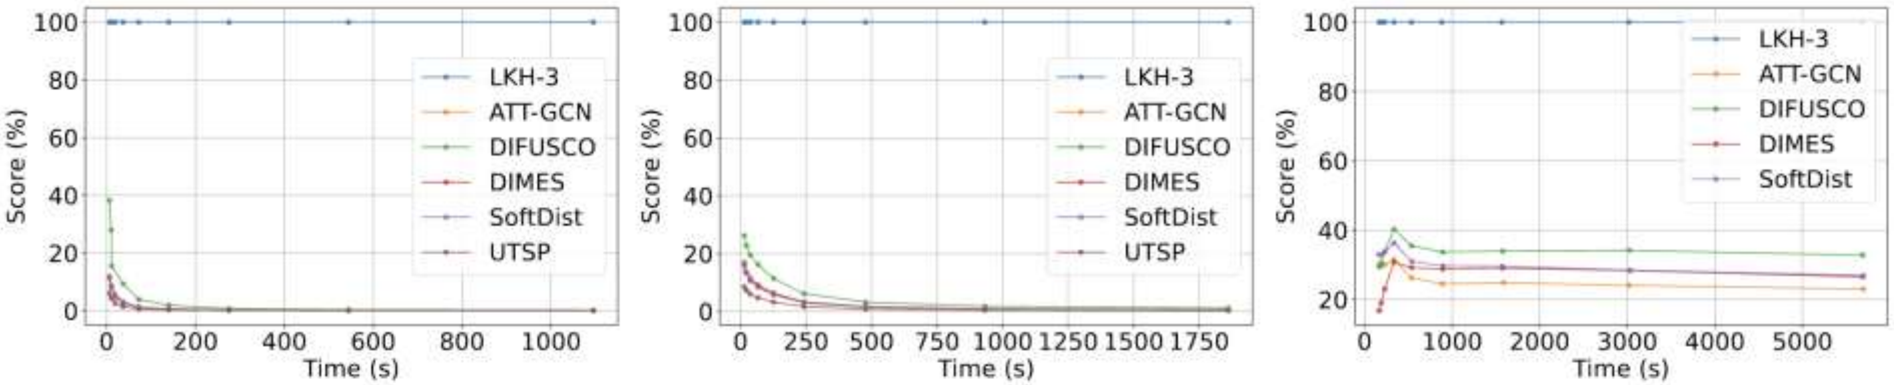

(a) TSP-500 with default MCTS settings. (b) TSP-1000 with default MCTS settings. (c) TSP-10000 with default MCTS settings.

Experiment 4

- By fine-tuning MCTS parameters, the influence of ML-generated heatmaps diminishes; even zero-input heatmaps yield similar performance

Method	TSP-500				TSP-1000			
	Length ↓	Gap ↓	Time ↓	Score ↑	Length ↓	Gap ↓	Time ↓	Score ↑
LKH-3 (default)	16.55*	0.00%	46.28m	—	23.12*	0.00%	2.57h	—
ATT-GCN†	16.72	1.02%	0.52m+0.67m	5.38%	23.48	1.58%	0.73m+1.43m	13.77%
DIMES†	16.75	1.26%	0.97m+0.68m	4.35%	23.61	2.11%	2.08m+1.45m	10.29%
DIFUSCO#	16.69	0.90%	3.61m+0.68m	6.12%	23.48	1.56%	11.86m+1.43m	13.93%
UTSP†	16.73	1.09%	1.37m+0.68m	5.05%	23.50	1.65%	3.35m+1.45m	13.18%
SoftDist	16.72	1.03%	0.00m+0.68m	5.32%	23.52	1.73%	0.00m+1.44m	12.56%
Zeros	16.72	1.06%	0.00m+0.68m	5.20%	23.55	1.85%	0.00m+1.44m	11.72%

Experiment 4

- By fine-tuning MCTS parameters, the influence of ML-generated heatmaps diminishes; even zero-input heatmaps yield similar performance

Method	TSP-500				TSP-1000			
	Length ↓	Gap ↓	Time ↓	Score ↑	Length ↓	Gap ↓	Time ↓	Score ↑
LKH-3 (default)	16.55*	0.00%	46.28m	—	23.12*	0.00%	2.57h	—
ATT-GCN†	16.72	1.02%	0.52m+0.67m	5.38%	23.48	1.58%	0.73m+1.43m	13.77%
DIMES†	16.75	1.26%	0.97m+0.68m	4.35%	23.61	2.11%	2.08m+1.45m	10.29%
DIFUSCO#	16.69	0.90%	3.61m+0.68m	6.12%	23.48	1.56%	11.86m+1.43m	13.93%
UTSP†	16.73	1.09%	1.37m+0.68m	5.05%	23.50	1.65%	3.35m+1.45m	13.18%
SoftDist	16.72	1.03%	0.00m+0.68m	5.32%	23.52	1.73%	0.00m+1.44m	12.56%
Zeros	16.72	1.06%	0.00m+0.68m	5.20%	23.55	1.85%	0.00m+1.44m	11.72%

Discussion

- Key Issue:
  - Reliance on surrgate loss functions lacking rigorous theoretical foundation
  - Surrogate loss function do not directly optimize original TSP loss, resulting in uncertaion test performance

Discussion

- Key Issue:
  - Reliance on surrgate loss functions lacking rigorous theoretical foundation
  - Surrogate loss function do not directly optimize original TSP loss, resulting in uncertaion test performance
- Continued Dependence on Post-Hoc Search Methods:
  - Contradicts goal of using ML in OR to develop generalizable, autonomous algorithms

Discussion

- Key Issue:
  - Reliance on surrgate loss functions lacking rigorous theoretical foundation
  - Surrogate loss function do not directly optimize original TSP loss, resulting in uncertaion test performance
- Continued Dependence on Post-Hoc Search Methods:
  - Contradicts goal of using ML in OR to develop generalizable, autonomous algorithms
- Future Direction of ML4CO:
  - Focus on developing more effective heatmap generation methods with a theoretical basis
  - Explore end-to-end solution generation methods

Thanks!

- Poster: Today 11:30 a.m. – 1 p.m.
  Hall C 4-9 #900
- Contact: yfxia@smail.nju.edu.cn
- Paper:

arXiv:2406.03503