

A Additional Related Work

A.1 Per-Instance Search

Once the neural network is trained over a collection of problem instances, per-instance fine-tuning can be used to improve the quality of solutions via local search. For DRL solvers, Bello et al. [7] fine-tuned the policy network on each test graph, which is referred as *active search*. Hottung et al. [28] proposed three active search strategies for efficient updating of parameter subsets during search. Zheng et al. [77] tried a combination of traditional reinforcement learning with Lin-Kernighan-Helsgaun (LKH) Algorithm [49, 24]. Hottung et al. [27] performed per-instance search in a differentiable continuous space encoded by a conditional variational auto-encoder [39]. With a heatmap indicating the promising parts of the search space, discrete solutions can be found via beam search [31], sampling [42], guided tree-search [48], dynamic programming [44], and Monte Carlo Tree Search (MCTS) [19]. In this paper, we mainly adopt greedy, sampling, and MCTS as the per-instance search techniques.

B Per-instance Search

In this section, we describe the decoding strategies used in our paper. Given a fine-tuned (i.e., after active search) continuous parameterization $\theta_s^{(T)}$ of the solution space, the per-instance search decoding aims to search for a feasible solution that minimizes the cost function c_s .

Greedy Decoding generates the solution through a sequential decoding process similar to the auxiliary distribution designed for each combinatorial optimization problem, where at each step, the variable k with the highest score θ_k is chosen to extend the partial solution. For TSP, the first node in the permutation is picked at random.

Sampling Inspired by Kool et al. [42], we propose to parallelly sample multiple solutions according to the auxiliary distribution and report the best one. The continuous parameterization is divided by a temperature parameter τ . The parallel sampling of solutions in DIMES is very efficient due to the fact that it only relies on the final parameterization $\theta_s^{(T)}/\tau$ but not on neural networks.

Monte Carlo Tree Search Inspired by [19], for the TSP task, we also leverage a more advanced reinforcement learning-based searching approach, i.e., Monte Carlo tree search (MCTS), to find high-quality solutions. In MCTS, k -opt transformation actions are sampled guided by the continuous parameterization $\theta_s^{(T)}$ to improve the current solutions. The MCTS iterates over the simulation, selection, and back-propagation steps, until no improving actions exists among the sampling pool. For more details, please refer to [19].

C Implementation Details

C.1 Neural Architecture for TSP

Anisotropic Graph Neural Networks We follow Joshi et al. [33] on the choice of neural architectures. The backbone of the graph neural network is an anisotropic GNN with an edge gating mechanism [9]. Let \mathbf{h}_i^ℓ and \mathbf{e}_{ij}^ℓ denote the node and edge features at layer ℓ associated with node i and edge ij , respectively. The features at the next layer is propagated with an anisotropic message passing scheme:

$$\mathbf{h}_i^{\ell+1} = \mathbf{h}_i^\ell + \alpha(\text{BN}(\mathbf{U}^\ell \mathbf{h}_i^\ell + \mathcal{A}_{j \in \mathcal{N}_i}(\sigma(\mathbf{e}_{ij}^\ell) \odot \mathbf{V}^\ell \mathbf{h}_j^\ell))), \quad (18)$$

$$\mathbf{e}_{ij}^{\ell+1} = \mathbf{e}_{ij}^\ell + \alpha(\text{BN}(\mathbf{P}^\ell \mathbf{e}_{ij}^\ell + \mathbf{Q}^\ell \mathbf{h}_i^\ell + \mathbf{R}^\ell \mathbf{h}_j^\ell)). \quad (19)$$

where $\mathbf{U}^\ell, \mathbf{V}^\ell, \mathbf{P}^\ell, \mathbf{Q}^\ell, \mathbf{R}^\ell \in \mathbb{R}^{d \times d}$ are the learnable parameters of layer ℓ , α denotes the activation function (we use SiLU [15] in this paper), BN denotes the Batch Normalization operator [30], \mathcal{A} denotes the aggregation function (we use mean pooling in this paper), σ is the sigmoid function, \odot is the Hadamard product, and \mathcal{N}_i denotes the outlinks (neighborhood) of node i . We use a 12-layer GNN with width 32.

The node and edge features at the first layer \mathbf{h}_i^0 and e_{ij}^0 are initialized with the absolute position of the nodes and absolute length of the edges, respectively. After the anisotropic GNN backbone, a Multi-Layer Perceptron (MLP) is appended and generates the final continuous parameterization θ for all the edges. We use a 3-layer MLP with width 32.

Graph Sparsification As described, we focus on developing a neural TSP solver for graphs with tens of thousands of nodes. Because the number of edges in the graph grows quadratically to the number of nodes, a densely connected graph is intractable for an anisotropic GNN when it is applied to large graphs. Therefore, we use a simple heuristic to sparsify the original graph. Specifically, we prune the outlinks of each node such that it is only connected to k nearest neighbors. The continuous parameterization θ is also pruned accordingly. As a result, the computation complexity of our method is reduced from $O(n^2)$ to $O(nk)$, where n is the number of nodes in the graph.

C.2 Neural Architecture for MIS

Graph Convolutional Networks We follow Li et al. [48] on the choice of neural architecture, i.e., using Graph Convolutional Network (GCN) [40], since θ is merely scores for each node. Specifically, the GCN backbone consists of multiple layers $\{\mathbf{h}^l\}$ where $\mathbf{h}^l \in \mathbb{R}^{N \times C^l}$ is the feature layer in the l -th layer and C^l is the number of feature channels in the l -th layer. We initialize the input layer \mathbf{h}^0 with all ones and \mathbf{h}^{l+1} is computed from the previous layer \mathbf{h}^l with layer-wise convolutions:

$$\mathbf{h}^{l+1} = \sigma(\mathbf{h}^l \mathbf{U}_0^l + \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \mathbf{h}^l \mathbf{U}_1^l), \quad (20)$$

where $\mathbf{U}_0^l \in \mathbb{R}^{C^l \times C^{l+1}}$ and $\mathbf{U}_1^l \in \mathbb{R}^{C^l \times C^{l+1}}$ are trainable weights in the convolutions of the network, \mathbf{D} is the degree matrix of \mathbf{A} with its diagonal entry $\mathbf{D}(i, i) = \sum_j \mathbf{A}(j, i)$, and $\sigma(\cdot)$ is the ReLU [55] activation function. After the GCN backbone, a 10-layer Multi-Layer Perceptron (MLP) with residual connections [22] is appended and generates the final continuous parameterization θ for all the nodes.

D Experimental Details

D.1 TSP

Training For TSP-500, we train our model for 120 meta-gradient descent steps (1.5 h in total) with $T = 15$. For TSP-1000, we train our model for 120 meta-gradient descent steps (1.7 h in total) with $T = 14$. For TSP-10000, we train our model for 50 meta-gradient descent steps (10 h in total) with $T = 12$. We generate 3 instances per meta-gradient descent step. We use the AdamW optimizer [50] with learning rate 0.005 and weight decay 0.0005 for meta-gradient descent steps, and with learning rate 0.05 for REINFORCE gradient descent steps. For other learning-based baseline methods, we download and rerun the source codes published by their original authors based on their pre-trained models.

Hardware We follow the hardware environment suggested by Fu et al. [19]. For the three traditional algorithms, since their source codes do not support running on GPUs, they run on Intel Xeon Gold 5118 CPU @ 2.30GHz. To ensure fair comparison, learning-based methods run on GTX 1080 Ti GPU during the testing phase. MCTS runs on Intel Xeon Gold 6230 80-core CPU @ 2.10GHz, where we use 64 threads for TSP-500 and TSP-1000, and 16 threads for TSP-10000. For the training phase, we train our model on NVIDIA Tesla P100 16GB GPU.

Reproduction We implement DIMES for TSP based on PyTorch Geometric [17] in LibTorch and PyTorch [62]. Our code for TSP is publicly available.⁷ The test instances are provided by Fu et al. [19].⁸

D.2 MIS

Training For SAT, we train our model for 50k meta-gradient steps with $T = 1$. For ER-[700-800], we train our model for 150k meta-gradient steps with $T = 1$. For ER-[9000-11000], we initialize

⁷<https://github.com/DIMESTeam/DIMES> (MIT license)

⁸<https://github.com/Spider-scnu/TSP> (MIT license)

our model from the checkpoint of ER-[700-800], and further train it for 200 meta-gradient steps. We use a batch size of 8 on all datasets and Adam optimizer [38] with learning rate 0.001 for the meta-gradient descent step, and with learning rate 0.0002 for REINFORCE gradient descent steps. For other learning-based baseline methods, we mainly use an integrated implementation⁹ provided by Böther et al. [8].

Hardware All the methods are trained and evaluated on a single NVIDIA Ampere A100 40 GB GPU, with AMD EPYC 7713 64-Core CPUs.

Reproduction Our code for MIS is publicly available.¹⁰ Following Böther et al. [8], for SAT, we use the “Random-3-SAT Instances with Controlled Backbone Size” dataset¹¹ and randomly split it into 39500 training instances and 500 test instances. For the Erdős-Rényi graphs, both training and test instances are randomly generated.

E Proofs

In this section, we follow the notation introduced in Section 3.

E.1 Convergence of Solution Distributions

The following propositions show that p_θ and q_θ converge to the *same* solution. They imply that we can optimize q_θ instead of p_θ .

Proposition 1 (TSP version). *Let $0 < \delta \ll 1$ be a sufficiently small number. If $q_\theta^{\text{TSP}}(f) \geq 1 - \delta$ for a solution $f \in \mathcal{F}$, then we also have $p_\theta(f) \geq 1 - O(\delta)$.*

Proposition 2 (MIS version). *Suppose that θ is normalized (i.e., $\sum_i \exp(\theta_i) = 1$) and uniformly bounded w.r.t. a solution $f \in \mathcal{F}$ (i.e., $\sum_i f_i \exp(\theta_i) / \exp(\sum_i f_i \theta_i) \leq L$ for a constant $L > 0$). Let $0 < \delta \ll 1$ be a sufficiently small number. If $q_\theta^{\text{MIS}}(f) \geq 1 - \delta$, then we also have $p_\theta(f) \geq 1 - O(\delta)$.*

Remark. Propositions 1 & 2 imply that if q_θ converges to f ($\delta \rightarrow 0_+$), then p_θ also converges to f .

Proof for TSP. Using the bound of $q_\theta^{\text{TSP}}(f)$, we have for any node j :

$$q_{\text{TSP}}(\pi_f \mid \pi_f(0) = j) = n q_\theta^{\text{TSP}}(f) - \sum_{i \neq j} q_{\text{TSP}}(\pi_f \mid \pi_f(0) = i) \quad (21)$$

$$\geq n q_\theta^{\text{TSP}}(f) - (n - 1) \quad (22)$$

$$\geq n(1 - \delta) - (n - 1) = 1 - O(\delta). \quad (23)$$

Thus, for any edge (i, j) in the tour π_f and any edge $(i, k) \neq (i, j)$,

$$\theta_{i,j} - \theta_{i,k} = \log \frac{\exp(\theta_{i,j})}{\exp(\theta_{i,k})} \quad (24)$$

$$\geq \log \frac{q_\theta(\pi_f(1) = j \mid \pi_f(0) = i)}{1 - q_\theta(\pi_f(1) = j \mid \pi_f(0) = i)} \quad (25)$$

$$\geq \log \frac{q_\theta(\pi_f \mid \pi_f(0) = i)}{1 - q_\theta(\pi_f \mid \pi_f(0) = i)} \quad (26)$$

$$\geq \log \frac{1 - O(\delta)}{O(\delta)}. \quad (27)$$

Note that for any edge (i, j) in the tour f (denoted by $(i, j) \in \pi_f$) and any solution $g \in \mathcal{F} \setminus \{f\}$, there exist a unique k_i^g such that edge (i, k_i^g) is in the tour π_g , and $(i, k_i^g) \neq (i, j)$ for at least one

⁹<https://github.com/MaxiBoether/mis-benchmark-framework> (No license)

¹⁰<https://github.com/DIMESTeam/DIMES> (MIT license)

¹¹https://www.cs.ubc.ca/~hoos/SATLIB/Benchmarks/SAT/CBS/descr_CBS.html

edge $(i, j) \in \pi_f$. Then,

$$p_{\theta}(f) = \frac{1}{1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \exp(-\sum_{(i,j) \in f} (\theta_{i,j} - \theta_{i,k_i^g}))} \quad (28)$$

$$= \frac{1}{1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \exp(-\sum_{(i,j) \in f \setminus g} (\theta_{i,j} - \theta_{i,k_i^g}))} \quad (29)$$

$$\geq \frac{1}{1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \exp(-\sum_{(i,j) \in f \setminus g} \log \frac{1-O(\delta)}{O(\delta)})} \quad (30)$$

$$= 1 - O(\delta). \quad (31)$$

□

Proof for MIS. Let $|g|$ denote the size of a solution $g \in \mathcal{F}$, i.e., $|g| = \sum_i g_i$. With a little abuse of notation, let $g \in \mathcal{F}$ also denote the corresponding independent set. Note that

$$\frac{\max_{i \notin f} \exp(\theta_i)}{\max_{i \notin f} \exp(\theta_i) + \sum_{i \in f} \exp(\theta_i)} \quad (32)$$

$$\leq \frac{\sum_{i \notin f} \exp(\theta_i)}{\sum_{i \notin f} \exp(\theta_i) + \sum_{i \in f} \exp(\theta_i)} \quad (33)$$

$$= \frac{\sum_{i \notin f} \exp(\theta_i)}{\sum_i \exp(\theta_i)} = \sum_{i \notin f} q_{\text{MIS}}(a_1 = i) \quad (34)$$

$$= q_{\text{MIS}}(a_1 \notin f) \leq 1 - q_{\theta}^{\text{MIS}}(f) \leq \delta. \quad (35)$$

This implies

$$\max_{i \notin f} \exp(\theta_i) \leq \frac{\delta}{1-\delta} \sum_{i \in f} \exp(\theta_i). \quad (36)$$

Recall that we have assumed in Section 3.2.2 that each $f' \in \mathcal{F}$ is not a proper subset of any other $f'' \in \mathcal{F}$. Thus for any $f, g \in \mathcal{F}$, we have $f \setminus g \neq \emptyset$, and $g \setminus f \neq \emptyset$. Note also that $\exp(\theta_i) \leq \sum_j \exp(\theta_j) = 1$ for all nodes i . Hence,

$$p_{\theta}(f) = \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\exp(\sum_i g_i \theta_i)}{\exp(\sum_i f_i \theta_i)}\right)^{-1} \quad (37)$$

$$= \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\prod_{i \in g \setminus f} \exp(\theta_i)}{\prod_{i \in f \setminus g} \exp(\theta_i)}\right)^{-1} \quad (38)$$

$$\geq \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\max_{i \in g \setminus f} \exp(\theta_i)}{\prod_{i \in f \setminus g} \exp(\theta_i)}\right)^{-1} \quad (39)$$

$$\geq \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\max_{i \notin f} \exp(\theta_i)}{\prod_{i \in f} \exp(\theta_i)}\right)^{-1} \quad (40)$$

$$\geq \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\frac{\delta}{1-\delta} \sum_{i \in f} \exp(\theta_i)}{\prod_{i \in f} \exp(\theta_i)}\right)^{-1} \quad (41)$$

$$\geq \left(1 + \sum_{g \in \mathcal{F} \setminus \{f\}} \frac{\delta}{1-\delta} \cdot L\right)^{-1} \quad (42)$$

$$= 1 - O(\delta). \quad (43)$$

□

E.2 First-Order Approximation of Meta-Gradient

The following proposition gives a first-order approximation formula of the meta-gradient.

Proposition 3. *Let $F_{\Phi}(\kappa_s, A_s)$ be a GNN F with parameter Φ and input (κ_s, A_s) , $\mathcal{L}(\Phi | \{s\})$ be a loss function, and $\alpha > 0$ be a learning rate. Suppose $\Phi_s^{(0)} = \Phi$, and $\Phi_s^{(t)} = \Phi_s^{(t-1)} - \alpha \nabla_{\Phi_s^{(t-1)}} \mathcal{L}(\Phi_s^{(t-1)} | \{s\})$ for $1 \leq t \leq T$, and $\theta_s^{(T)} = F_{\Phi_s^{(T)}}(\kappa_s, A_s)$. Then,*

$$\nabla_{\Phi} \theta_s^{(T)} = \nabla_{\Phi_s^{(T)}} F_{\Phi_s^{(T)}}(\kappa_s, A_s) + O(\alpha).$$

Proof. The proof resembles [58]. By chain rule,

$$\nabla_{\Phi_s^{(0)}} \Phi_s^{(T)} = \prod_{t=1}^T \nabla_{\Phi_s^{(t-1)}} \Phi_s^{(t)} \quad (44)$$

$$= \prod_{t=1}^T \nabla_{\Phi_s^{(t-1)}} (\Phi_s^{(t-1)} - \alpha \nabla_{\Phi_s^{(t-1)}} \mathcal{L}(\Phi_s^{(t-1)} | \{s\})) \quad (45)$$

$$= \prod_{t=1}^T (I - \alpha \nabla_{\Phi_s^{(t-1)}}^2 \mathcal{L}(\Phi_s^{(t-1)} | \{s\})) \quad (46)$$

$$= I + \sum_{k=1}^T (-\alpha)^k \sum_{1 \leq t_1 < \dots < t_k \leq T} \prod_{i=1}^k \nabla_{\Phi_s^{(t_i-1)}}^2 \mathcal{L}(\Phi_s^{(t_i-1)} | \{s\}) \quad (47)$$

$$= I + O(\alpha). \quad (48)$$

Hence,

$$\nabla_{\Phi} \theta_s^{(T)} = \nabla_{\Phi_s^{(0)}} \Phi_s^{(T)} \nabla_{\Phi_s^{(T)}} F_{\Phi_s^{(T)}}(\kappa_s, A_s) \quad (49)$$

$$= (I + O(\alpha)) \nabla_{\Phi_s^{(T)}} F_{\Phi_s^{(T)}}(\kappa_s, A_s) \quad (50)$$

$$= \nabla_{\Phi_s^{(T)}} F_{\Phi_s^{(T)}}(\kappa_s, A_s) + O(\alpha). \quad (51)$$

□

F Additional Experiments for TSP

F.1 Performance on TSP-100

We trained DIMES on TSP-100 and evaluate it on TSP-100 with $T = 10$ and 0 (i.e., with and without meta-learning). Since MCTS is the best per-instance search scheme for DIMES (see Table 1), we also use MCTS here. When using AS, we fine-tune DIMES on each instance for 100 steps. We compare DIMES with learning-based methods listed in Section 4.1.2. Results of baselines are taken from Fu et al. [19]. The results are presented in Table 4.

As is shown in the table, DIMES outperforms all learning-based methods, and its results are very close to optimal lengths given by exact solvers. The results suggest that DIMES achieves the best in-distribution performance among learning-based methods. Notably, with meta-learning ($T = 10$), even when DIMES does not fine-tune (i.e., no active search) for each problem instance in evaluation, it still outperforms all other learning-based methods. This again demonstrates the efficacy of meta-learning to DIMES.

F.2 Extrapolation Performance

We evaluate the extrapolation performance of DIMES (i.e., trained on smaller graphs and tested on larger graphs). We train the model on TSP-100 and test it on TSP-500/1000/10000. For testing, we use RL+S ($\tau = 0.01$) without active search. The results are reported in Table 5 in comparison with corresponding results trained on larger graphs (TSP- n).

Table 4: Results on TSP-100. * indicates the baseline for computing the performance drop.

Method	Type	Length ↓	Drop ↓
Concorde	OR (exact)	7.7609*	—
Gurobi	OR (exact)	7.7609*	—
LKH-3	OR	7.7611	0.0026%
EAN	RL+S	8.8372	13.8679%
EAN	RL+S+2-OPT	8.2449	6.2365%
AM	RL+S	7.9735	2.7391%
AM	RL+G	8.1008	4.3791%
AM	RL+BS	7.9536	2.4829%
GCN	SL+G	8.4128	8.3995%
GCN	SL+BS	7.8763	1.4828%
Att-GCN	SL+MCTS	7.7638	0.0370%
DIMES ($T = 0$)	RL+MCTS	7.7647	0.0490%
DIMES ($T = 0$)	RL+AS+MCTS	7.7618	0.0116%
DIMES ($T = 10$)	RL+MCTS	7.7620	0.0142%
DIMES ($T = 10$)	RL+AS+MCTS	7.7617	0.0103%

Table 5: Results of DIMES (RL+S). “Trained on TSP-100” indicates extrapolation performance.

Setting	TSP-500		TSP-1000		TSP-10000	
	Length ↓	Drop ↓	Length ↓	Drop ↓	Length ↓	Drop ↓
Trained on TSP- n	18.84	13.84%	26.36	14.01%	85.75	19.48%
Trained on TSP-100	19.21	16.07%	27.21	17.69%	86.24	20.16%

From the table we can observe that the performance of DIMES does not drop much, which demonstrates the nice extrapolation performance of DIMES. One of our hypotheses is that graph sparsification in our neural network (see Appendix C.1) avoids the explosion of activation values in the graph neural network. Another hypothesis is that meta learning tends to not generate too extreme values in (see point 9 of our previous response) and hence improve the generalization capability.

F.3 Stability of Training

We compare the training settings of AM [42], POMO [45], and DIMES in Table 6. The training costs of AM and POMO are obtained from their papers¹². A training step means a gradient descent step of the GNN. That is, for AM/POMO, a training step means a gradient descent step over a batch; for DIMES, a training step means a meta-gradient descent step.

The table shows that DIMES is much more sample-efficient than AM/POMO. Notably, DIMES achieves stable training using only 3 instances per meta-gradient descent step. Hence, its total training time is accordingly much shorter, even though its per-step time is longer. Moreover, the stability of training enables us to use a larger learning rate, which also accelerates training.

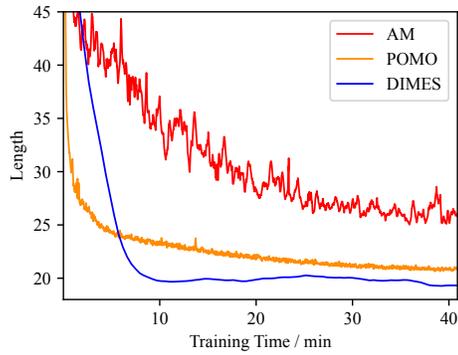
To further illustrate the fast stable training of DIMES, we compare the dynamics of training among AM, POMO, and DIMES in Figure 1. We closely follow the training settings of their papers, i.e., we train AM/POMO on TSP-100 and DIMES on TSP-500. For AM/POMO, we train their models on our hardware by re-running their public source code. The performance is evaluated using TSP-500 test instances. For DIMES, we use RL+S in evaluation.

From Figure 1a, we can observe that DIMES stably converges to a better performance within fewer time, while the dynamics of training AM/POMO are slower and less stable. From Figure 1b, we can observe that DIMES converges at much fewer training steps. The results again demonstrate that the training of DIMES is fast and stable.

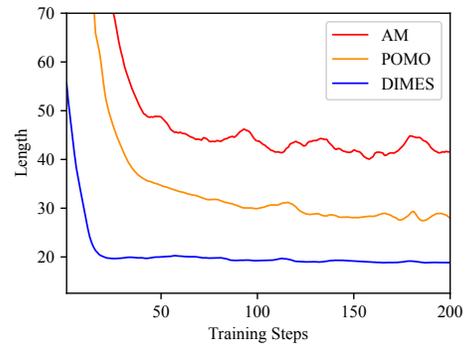
¹²For AM/POMO, per-step training time is estimated by total training time divided by total training steps.

Table 6: Comparison of training settings for TSP-500/1000/10000.

Setting	AM	POMO	DIMES
Training problem scale	TSP-100	TSP-100	TSP-500/1000/10000
Training descent steps	250,000	312,600	120/120/50
Per-step training instances	512	64	3
Total training instances	128,000,000	20,000,000	360/360/150
Per-step training time	0.66 s	0.28 s	45 s/51 s/12 m
Total training time	2 d	1 d	1.5 h/1.7 h/10 h
Training GPUs	2	1	1



(a) Performance vs training time.



(b) Performance vs training steps.

Figure 1: Evaluation performance vs training cost.