

DCLP: Neural Architecture Predictor with Curriculum Contrastive Learning

Shenghe Zheng, Hongzhi Wang*, Tianyu Mu

Massive Data Computing Lab, Harbin Institute of Technology
shenghez.zheng@gmail.com, {wangzh, mutianyu}@hit.edu.cn

Abstract

Neural predictors have shown great potential in the evaluation process of neural architecture search (NAS). However, current predictor-based approaches overlook the fact that training a predictor necessitates a considerable number of trained neural networks as the labeled training set, which is costly to obtain. Therefore, the critical issue in utilizing predictors for NAS is to train a high-performance predictor using as few trained neural networks as possible. Although some methods attempt to address this problem through unsupervised learning, they often result in inaccurate predictions. We argue that the unsupervised tasks intended for the common graph data are too challenging for neural networks, causing unsupervised training to be susceptible to performance crashes in NAS. To address this issue, we propose a Curriculum-guided Contrastive Learning framework for neural Predictor (DCLP). Our method simplifies the contrastive task by designing a novel curriculum to enhance the stability of unlabeled training data distribution during contrastive training. Specifically, we propose a scheduler that ranks the training data according to the contrastive difficulty of each data and then inputs them to the contrastive learner in order. This approach concentrates the training data distribution and makes contrastive training more efficient. By using our method, the contrastive learner incrementally learns feature representations via unsupervised data on a smooth learning curve, avoiding performance crashes that may occur with excessively variable training data distributions. We experimentally demonstrate that DCLP has high accuracy and efficiency compared with existing predictors, and shows promising potential to discover superior architectures in various search spaces when combined with search strategies. Our code is available at: <https://github.com/Zhengsh123/DCLP>.

1 Introduction

To reduce the design cost of Deep Neural Networks (DNNs), NAS has emerged as a promising technique to automatically design neural networks for specific tasks (Zoph and Le 2017; Qin et al. 2022). NAS has been widely adopted for various tasks such as image classification (Li et al. 2023) and semantic segmentation (Zhang et al. 2021). In typical NAS, searched architectures are evaluated to guide search (He and

Zhao 2021). However, the estimation is often expensive as it usually requires training DNNs from scratch to assess their performance. For instance, NASNet (Zoph et al. 2018) needs around 40K GPU hours for estimation on CIFAR-10, making it unaffordable for many applications. Hence, predictor-based NAS (Ning et al. 2020) has gained popularity due to its close-to-zero estimation costs and satisfactory results.

For predictor-based NAS, a key challenge is the requirement for trained neural networks to train the predictor. Training a neural network to converge may take several days, and then obtaining a substantial amount of training data is costly. Two solutions have been proposed to overcome this challenge. The first is to optimize the predictor to better utilize limited labeled data (Chen et al. 2021b), while the second is to optimize the training method by incorporating unsupervised learning to utilize unlabeled data obtained at a low cost (Tang et al. 2020). However, optimizing the predictor disregards the upper limit to the predictive ability learned from limited labeled data. Then, they still require a great number of training data (Ning et al. 2020). On the other hand, although the introduction of unsupervised learning reduces the need for labeled data, it does not enable predictors to achieve desired performance (Jing and Xu 2022).

Unsupervised neural predictors fail because the pretext tasks designed for regular graph data are too complex for neural networks (Jing and Xu 2022). For example, contrastive learning, a popular unsupervised approach, requires the model to differentiate data similarity and provide similar embeddings for similar data (He et al. 2020). However, determining the similarity between neural networks is challenging even for experienced experts. Moreover, randomly selecting samples in vanilla contrastive learning creates excessive data distribution differences during training, which further complicates the learning process. Therefore, employing unsupervised learning meant for other data to neural predictors results in poor performance (Tang et al. 2020).

To simplify the contrastive task for better adaptation to neural predictors, the issue of high variability in the distributions of randomly selected positive data for contrastive training, which may hinder the learning process for weak learners, must be addressed. This variability in data distribution arises from differences in the difficulty of positive samples during contrastive training. Specifically, when positive data are more similar to each other, the learner can classify

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

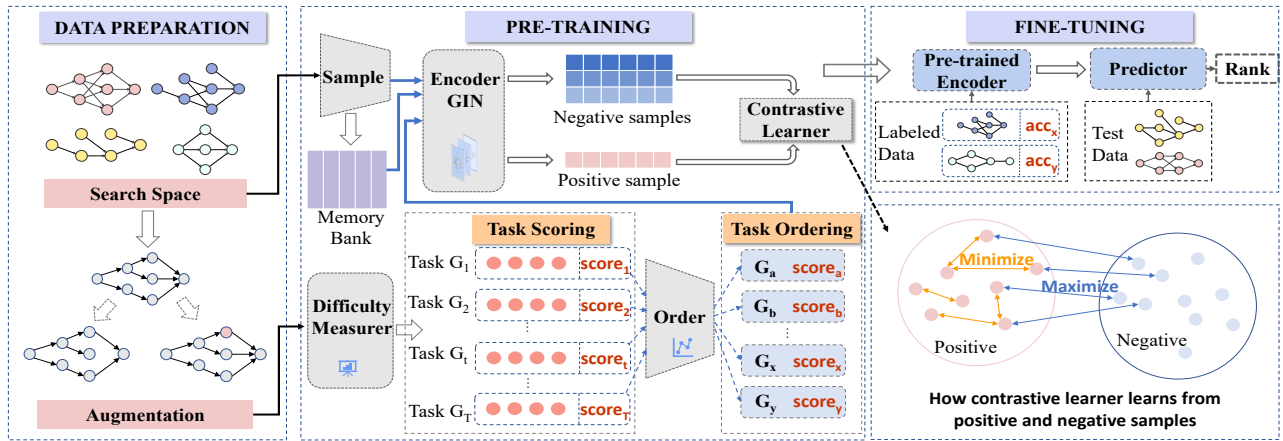


Figure 1: The overall framework of DCLP, comprises data preparation, pre-training, and fine-tuning. Left: The data preparation, uses edge perturbation and attribute masking to augment graphs in the search space to obtain the positive training samples. Middle: The pre-training. The bottom is the curriculum learning section, which schedules the training order of each positive data. The upper part is the contrastive learning process. A good encoder for neural networks is obtained by pre-training. Right: The upper part is the fine-tuning process, using a small amount of labeled data to fine-tune the encoder and obtain the predictor.

them into the same class more easily, making such samples less challenging for contrastive learning. Then, our goal is to construct a training method with relatively stable data difficulty variation, allowing the contrastive learner to converge more consistently. To achieve this, inspired by curriculum learning (Bengio et al. 2009), we incorporate human knowledge to design a smooth learning curve, which enables the encoder to learn efficiently from unlabeled data. We propose a curriculum-guided contrastive learning method for neural predictors called DCLP as shown in Figure 1.

In DCLP, a novel metric is developed to measure the difficulty of positive data based on the distance between positive and original graphs. Next, a scheduler is employed to control the training order of data, prioritizing difficulty as the criterion as shown in the middle part of Figure 1. As training progresses, the scheduler gradually selects more difficult data, enabling the encoder to learn from harder data. However, unlike vanilla curriculum where data is learned from easy to hard, in DCLP, there are intervals of decreasing difficulty within the overall increasing trend. The decreasing intervals are arranged because as the data difficulty increases, the noise increases, and the model may overfit the noise, leading to poor performance. The uninterrupted use of simple data corrects the bias caused by noise to make the training more efficient. The scheduler infuses human knowledge into training, leading to efficient convergence of contrastive encoders, minimizing the empirical risk of training, and improving generalization ability (Wang, Chen, and Zhu 2021).

The key insight of DCLP is the difficulty of training varies across data, and mixing multiple difficulties in the same training round exacerbates the learning difficulty. Thus, the stepwise curriculum aims to improve the performance of predictors by creating a smooth training schedule.

After contrastive training, we use limited labeled data for fine-tuning. However, the typical regression loss is too stringent for NAS, which only requires the accurate predicted

ranking (Wang et al. 2021). Then, we propose using ranking as the optimization target. Combining DCLP with various search strategies for NAS, our pre-training method employs designed learning curves to enhance contrastive tasks, enabling smooth extraction of valuable information from unlabeled data for prediction. Additionally, we demonstrate the good generalization ability of DCLP to improve NAS performance while reducing consumption under multiple search strategies. Our contributions are summarized as follows:

- We utilize contrastive learning in neural predictors to leverage unlabeled data. This approach reduces the requirement for labeled training data of the predictor and improves its generalization ability.
- We propose a novel curriculum method to guide the contrastive task, which makes the predictor converge faster and perform better in NAS.
- We conduct comprehensive experiments, showing that DCLP outperforms the popular predictors on various benchmarks and search spaces that can be used in actual scenarios.

2 Related Work

Neural Architecture Search (NAS). For NAS, the focus is primarily on search methods such as reinforcement learning (Zoph and Le 2017), evolutionary search (Huang et al. 2022), and random search (Yang and Esperança 2020). However, a challenge is the costly evaluation. Although the one-shot method utilizes super-networks to address the issue, it remains expensive (Liu, Simonyan, and Yang 2019). Our focus is on the neural predictor, an efficient method that utilizes predictive models to decrease estimation costs.

Architecture Performance Predictors. The predictor is a method for reducing the evaluation cost by directly estimating neural networks. PNAS (Cai, Zhu, and Han 2019) introduces LSTM-based predictors. Other methods include random forests (Liu, Tang, and Sun 2021) and GNNs (Jing and

Tian 2020). However, they only use labeled data, while using unlabeled data can reduce the cost of getting data. For instance, Semi-Supervised Assessor (Tang et al. 2020) utilizes both labeled and unlabeled data, but the prediction is not ideal. Instead, DCLP makes unsupervised training more stable through curriculums. GMAE (Jing and Xu 2022) uses reconstruction as the pretext task. The difference is that DCLP uses contrastive training, achieving better results efficiently. Specifically, CTNAS (Chen et al. 2021a) predicts the relationship between two neural networks. Although it uses contrastive learning and curriculum, it differs greatly from DCLP. While contrastive learning in DCLP is unsupervised, CTNAS refers to its supervised training as contrastive. Additionally, CTNAS uses the curriculum to update results, whereas our curriculum schedules stable training. Thus, there are fundamental differences between CTNAS and DCLP. More details are available in Appendix G.

Contrastive Learning. Contrastive learning is an unsupervised method that has achieved success in various domains such as CV (He et al. 2020) and NLP (Devlin et al. 2019). It aims to train encoders to obtain embeddings for data by minimizing the distance between positive items (similar data) while maximizing the distance between negative data.

Curriculum Learning. Curriculum learning is a training method that orders data to simulate human learning (Bengio et al. 2009), and it has been shown to enhance generalization ability in various domains, such as CV (Pentina, Sharman-ska, and Lampert 2015) and NLP (Platanios et al. 2019). In this study, we integrate it into neural predictors and combine it with contrastive learning to leverage its benefits.

3 Methodology

In this section, we present DCLP, a contrastive neural predictor with curriculum guidance. Figure 1 provides an overview of DCLP, including the contrastive learning in Sec. 3.1, and the curriculum used for training in Sec. 3.2. Additionally, we discuss fine-tuning in Sec. 3.3, and its potential use for NAS in Sec. 3.4.

Before delving into the specifics, it is necessary to define the search space. We use cells as the search unit in line with most NAS works (Zoph and Le 2017). Each architecture is composed of L-stacked cells. Each cell can be represented as a DAG G . Cells have two formats: operation on nodes (OON) and operation on edges (OOE). OON is used as the default. In an OON cell, each node represents an operation, and the edge $e_{i \rightarrow j}$ indicates the information flow from node i to j . Details are illustrated in Appendix B.

3.1 Contrastive Learning Framework

The purpose of contrastive learning is to train an encoder for feature extraction from unlabeled neural networks. This subsection introduces data augmentation and the encoder, followed by the contrastive training approach.

Data Augmentation. Augmentation is employed to generate useful data to train a robust encoder. Given a graph G , we create an augmented graph \hat{G} using a predefined distribution $p_{\theta, \mu}(G)$, where θ represents the augmentation method and μ is the change ratio. To ensure appropriate augmentation for

graphs, we use edge perturbation and attribute masking (Wu et al. 2021). Edge perturbation randomly edits edges, while attribute masking modifies node properties. In DCLP, positive data are generated by augmentation, while negative data are obtained by random sampling.

Encoder Method. Encoder is used to convert neural networks into embeddings. In contrastive learning, similar data should have similar embeddings. Therefore, we need a model that can distinguish between graphs, so we employ Graph Isomorphism Network (GIN) (Xu et al. 2019), as it is highly capable of identifying differences in graphs. The propagation function of the k -th layer of a GIN is as follows:

$$h_n^k = MLP^k(h_n^{k-1} + \sum_{n' \in N(n)} h_{n'}^{k-1}) \quad (1)$$

where h_n^k is the encoding of node v_n in the k -th layer, and $N(n)$ is the set of node adjacent to v_n . Since a graph feature is required, the features can be integrated as follows:

$$h_G = \text{concat}(\text{readout}(\{h_n^k | n \in G\}) | k = 0, \dots, K) \quad (2)$$

Contrastive Training. Contrastive learning trains the encoder by letting it determine data similarity and ensure close embeddings for positive data. Then, the similarity between networks G_i and G_j is measured using Radial Basis Function (RBF) with encoder representations denoted as $z(G)$:

$$s(G_i, G_j) = \exp\left(-\frac{d(z(G_i), z(G_j))}{2\sigma^2}\right) \quad (3)$$

where d is distance measurement such as Euclidean distance. RBF maps items into a high-dimensional space to separate data that are hard to separate in the original space.

The contrastive task requires similar data to yield a larger s . InfoNCE is used as the optimized target. It is applied to positive pairs $\{q, k_+\}$ and negative data $\{k_i\}_{i=1}^N$ as follows:

$$L = -\log \frac{\exp(s(q \cdot k_+)/\tau)}{s(q \cdot k_+)/\tau + \sum_{i=1}^N \exp(s(q \cdot k_i)/\tau)} \quad (4)$$

where τ is a temperature parameter. Minimizing Eq.4 allows the encoder to extract features from unlabeled data by ensuring higher score of positive data than negative. Then, the encoder can reasonably encode networks for prediction.

3.2 Curriculum Learning

Training predictors using contrastive learning pose challenges when determining the similarity of neural networks compared to regular data. Random data ordering in existing contrastive methods may lead to convergence issues due to varying data distributions (Chu et al. 2021a). To address this, we propose a novel curriculum that schedules data to create a smooth learning curve, simplifying the contrastive task and enabling stable learning without training crashes. The curriculum involves two steps: difficulty measurement and training scheduler, which we will discuss in detail.

Difficulty Measurement. The difficulty function evaluates the difficulty of positive items in the contrastive task that aims to classify similar samples into the same class. Difficult positive data are those less similar to the original data, potentially causing misclassification by the contrastive learner.

Algorithm 1: Training procedure of DCLP

Input: Training set \mathcal{D} , the number of iterations \mathcal{T} , the augmentation distribution p , GIN encoder f

- 1: **pre-training**
- 2: Initialize parameters α of f with kaiming initialization
- 3: **for** $t = 1$ to \mathcal{T} **do**
- 4: $G \leftarrow \text{RandomSample}(\mathcal{D})$
- 5: Initialize $\hat{G} = \emptyset$
- 6: **for** $i = 1$ to N **do**
- 7: $\hat{G}_i \sim p_\theta(\hat{G}|G)$
- 8: $\hat{G} \leftarrow \hat{G} \cup \hat{G}_i$
- 9: **end for**
- 10: $G_t \leftarrow \text{argmax}(\mathcal{P}_t(\hat{G}_i))_{i=1}^N$
- 11: $z \leftarrow f(G), z' \leftarrow f(G_t)$
- 12: $f_\alpha \leftarrow f_\alpha - \nabla_{f_\alpha} L_{NCE}$
- 13: **end for**
- 14: **fine-tuning**
- 15: $S_F \leftarrow \text{LimitedRandomSample}(D)$
- 16: $L \leftarrow \text{SelectedLossFunction}(S_F)$
- 17: **for** $t = 1$ to \mathcal{T} **do**
- 18: $f_\alpha \leftarrow f_\alpha - \nabla_{f_\alpha} L$
- 19: **end for**

Output: the trained predictor f_α

Therefore, the difficulty of augmented data is calculated based on their similarity to the origin. The higher the similarity, the lower the difficulty. To calculate the difficulty of a positive graph \hat{G} related to the original graph G , we use the Wasserstein Distance(WD). WD measures the variability between the two graphs, which is a reasonable measure of difficulty. For practicality, the augmentation distribution can be considered uniform. Additionally, we approximate WD using Levenshtein Distance (LD) as follows:

$$\begin{aligned}
 L(\hat{G}, G) &= \left(\int |p_{\theta, \mu}^{-1}(\hat{G}|G) - p^{-1}(G) dt \right)^{1/p} \\
 &\approx \int_{\mathbb{R}} LD(\hat{G}_t, G_t, t) dt \\
 &= cLD(\hat{G}, G)
 \end{aligned} \tag{5}$$

The reasonableness of the approximation lies in that we need to calculate the distance between two uniform distributions denoted by μ . Each point's distance can be viewed as an approximation of the LD between the graphs generated by related parameters at that point. By integrating these approximations, we obtain an approximation of the final LD. Since the difficulty value in Eq.5 can be approximated as a multiple of LD, and c is a constant, the WD between the two graphs can be directly approximated by LD. This expresses the similarity between networks in terms of their structural similarity. More details can be found in Appendix F.1. Then the difficulty of \hat{G} can be defined as:

$$L(G, \hat{G}) = \frac{LD(G, \hat{G})}{|G| \cdot |\hat{G}|} \tag{6}$$

where $|G|$ is the size of G measuring by the number of edges.

Training Scheduler. After defining difficulty, we design the training order based on difficulty, controlled by a parameter that determines the training preference for easy or hard data. The probability of sampling \hat{G}_i at step t is then defined as:

$$\mathcal{P}_t(\hat{G}_i) = \frac{\exp(\tau_t L(\hat{G}_i, G_i))}{\sum_{j=1}^n \exp(\tau_t L(\hat{G}_j, G_j))} \tag{7}$$

where

$$\begin{aligned}
 \tau_t &= \frac{\tau_T - \tau_m}{\sigma} g(\sigma t) + \tau_m, \tau_m = \frac{\tau_1 + \tau_T}{2} \\
 \sigma_t &= g^{-1}\left(\frac{t}{T}\sigma + \frac{\sigma}{k} \sin \frac{n\pi t}{T}\right)
 \end{aligned} \tag{8}$$

where t denotes steps ($t \in \{1, 2, \dots, T\}$), g is tanh, $\sigma = 0.9$, and $n > 1, k > 1$ controls the frequency and range of difficulty decrease respectively, the decreasing will be discussed in the next paragraph. τ_t determines the preference for easy/hard data. For $\tau_t < 0$, larger difficulty corresponds to a smaller selection probability. Conversely, when $\tau_t > 0$, the curriculum favors hard data. Our strategy gradually adjusts τ_t to achieve robust learning. Initially, we set a negative τ_1 to emphasize learning from easy data at the start without being affected by noise and to quickly reach a relatively good ability. Then we gradually increase τ_t to encourage deeper learning from hard data. Finally, we use a positive τ_T , smoothly transitioning from easy to hard data.

However, unlike the conventional easy-to-hard curriculum, we introduce a fluctuating term that causes the curriculum to fluctuate in difficulty controlled by n and k . This non-monotonic increase in difficulty includes intervals of decreasing difficulty. Therefore, the model can revisit simpler data periodically during training. It is necessary since the harder positive data includes more noise, i.e., negative samples that are mistakenly labeled as positive. Training with increasing difficulty often causes the model to focus excessively on the noise, resulting in loss of learned knowledge in the correct data. Intermittent training with simple data can counteract the learned misinformation from the noise, improving training stability, and being more robust.

The curriculum facilitates the encoder to acquire a stronger ability by increasing data difficulty. Additionally, incorporating easy data prevents overfitting that may occur as difficulty increases, leading to efficient convergence.

As shown in Algorithm 1, we use the difficulty function to calculate the difficulty of each positive data and sort them according to the scheduler (Line 10). Next, we feed the positive samples into the contrastive learner for training.

3.3 Fine Tuning

After pre-training, to transfer the pre-trained model to the prediction task, fine-tuning is in demand. However, using mean squared error (MSE) between predicted and actual results as loss function is too strict, as NAS focuses on the best networks. To alleviate the issue of overfitting, we propose two functions that take ranking into account.

MSE+Normalization. As we prioritize ranking, normalization can widen the performance gap between two networks. Inaccurate rank judgments lead to greater MSE, facilitating

Algorithm 2: Predictor-based Random Search Strategy

Input: Search space \mathcal{A} , predictor f , iteration steps \mathcal{T} , the number of networks sampled in t -th iteration \mathcal{N}^t , the number of networks selected in each iteration \mathcal{K} .

- 1: Initialize $t = 0$, an architecture set $S = \emptyset$
- 2: **for** $t = 1$ to \mathcal{T} **do**
- 3: Randomly sample \mathcal{N}^t architectures $S^t = \{a_j^t\}_{j=1}^{\mathcal{N}^t}$ from the search space \mathcal{A} without repeat.
- 4: Evaluate items in S^t by trained predictor f , get $\hat{S}^t = \{(a_j^t, y_j^t)\}_{j=1}^{\mathcal{N}^t}$, where y is the predicted performance.
- 5: Select top- k items from \hat{S}^t according to y , get \tilde{S}^t .
- 6: $S \leftarrow S \cup \tilde{S}^t$.
- 7: **end for**
- 8: Evaluate each architecture in S by training to obtain the corresponding ground-truth \mathcal{Y}

Output: the architecture a^* with the best true performance.

the learning of their relative ranking relationship. Further details can be found in the supporting material F.2.

Ranking Loss. We utilize ListMLE (Burgess 2010) to optimize predicted ranking. We use o to represent the actual sorting result and $S = \{s_1 \dots s_n\}$ to denote the predicted accuracy of n architectures. Then, S is sorted according to o to obtain $\hat{S} = \{s_{o_1}, s_{o_2}, \dots, s_{o_n}\}$. Our objective is to maximize the probability of this \hat{S} to make the predicted ranking similar to the actual one. The loss function is as follows:

$$L = - \sum_{i=1}^n \log \frac{\exp(s_{o_i})}{\sum_{j=i}^n \exp(s_{o_j})} \quad (9)$$

These two optimization targets align the optimization goal with the prediction task, improving the model’s fitness to NAS while reducing the required labeled data.

3.4 Search Method

After fine-tuning, we use DCLP with various search strategies, such as random search (RS), evolution search (EA), and reinforcement learning (RL), to discover the optimal network. The fast estimation provided by DCLP aids in the efficient exploration of search space. A larger exploration space increases the likelihood of finding a better result.

Random search is outlined in Algorithm 2, while other strategies are in Appendix E. Each iteration involves randomly sampling N networks, and the top- k based on predicted performance are added to the candidate pool (Lines 4-6). After T iterations, only a few networks ($k \cdot T$) are trained to obtain actual performance, from which the optimal network is selected, making it an efficient approach.

4 Experiments

We evaluate DCLP through experiments for image classification and language tasks. We introduce experimental settings first and validate the superiority in NAS benchmarks, along with ablation experiments for further analysis.

Image Classification. We conduct experiments on two NAS benchmarks and a large search space. Details are as follows.

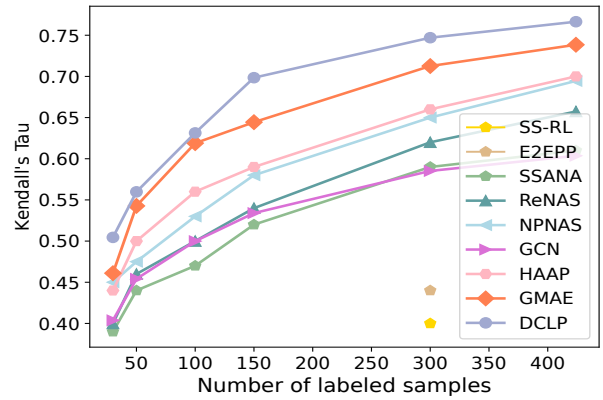


Figure 2: The comparison between DCLP and the state-of-the-art predictors on NAS-Bench-101.

- **NAS-Bench-101.** It is an OON set of 423k CNN (Ying et al. 2019). Each node can be a 1x1 convolution, etc. Test accuracy on CIFAR-10 is available within the search space.
- **NAS-Bench-201.** This is a NAS benchmark in OOE space (Dong and Yang 2020) with 15,625 structures, providing performance on CIFAR-10, CIFAR-100, and ImageNet.
- **DARTS Space.** It is an OOE search space (Liu, Simonyan, and Yang 2019), consisting of 10^{18} architectures. It contains cells with each edge offering an operation choice.

Language Model. We demonstrate the efficiency of DCLP on the language model task. It typically uses RNN instead of CNN in image tasks. The success of DCLP on both types of neural networks shows its generalization capability.

- **DARTS RNN Space:** The recurrent cell is an OOE space. The operations available are ReLU, Tanh, etc.

Evaluation Metrics. We evaluate the predictor using Kendall’s Tau (Sen 1968) for ranking accuracy, with a score closer to 1 indicating a better predictor. Additionally, the image classification accuracy and perplexity for language models are used to measure the performance of the NAS method.

4.1 Predictor Evaluation

In this section, we evaluate the predictive ability using Kendall’s Tau on NAS-Bench-101. Predictors with high Kendall’s Tau are more likely to discover exceptional networks. We compare DCLP with popular predictors (Wei et al. 2021; Xu et al. 2021; Jing and Xu 2022; Wen et al. 2020; Liu, Tang, and Sun 2021; Kipf 2017) to demonstrate its superiority in requiring less labeled data while maintaining high performance. Our goal is to develop efficient predictors for NAS with minimal labeled data, and DCLP shows outstanding performance with small labeled sets.

As shown in Figure 2, DCLP outperforms various popular predictors, particularly with smaller training sets. This is consistent with our focus on achieving high performance on small training data as mentioned above.

The results demonstrate DCLP’s superior generalization with small training sets compared to solely supervised learning approaches (Xu et al. 2021; Wen et al. 2020; Kipf

Method	CIFAR-10		CIFAR-100		ImageNet-16-120		Query	Search cost (GPU h)	Type
	Acc(%)	R(%)	Acc(%)	R(%)	Acc(%)	R(%)			
NASWOT (2021)	93.15±0.08	7.0	69.98±1.22	5.01	44.44±2.10	3.01	-	0.06	zero+RS
TENAS (2021)	93.18±0.39	6.6	70.24±1.05	3.68	43.55±8.24	6.6	-	0.43	zero
HAAP (2021)	93.75±0.17	1.16	71.08±0.19	0.97	45.22±0.26	0.94	150	0.011	pred+EA
GMAE (2022)	94.03±0.11	0.21	72.56 ±0.16	0.13	46.09±0.27	0.19	150	0.009	pred+RS
HOP (2021b)	94.1±0.12	0.14	72.64± 0.11	0.10	46.29±0.19	0.1	150	0.015	pred+RS
FreeREA (2023)	94.20±0.02	0.08	73.30±0.31	0.01	46.34±0.00	0.09	100	0.012	EA
DCNAS (2022)	94.29±0.07	0.14	73.02±0.16	0.02	46.41±0.14	0.08	-	3.9	supernet
DCLP+RL	94.29±0.05	0.14	72.83±0.15	0.096	46.49±0.25	0.03	50	0.007	pred+RL
DCLP+RS	94.34±0.03	0.025	73.5±0.3	0.0001	46.54±0.03	0.02	50	0.004	pred+RS

Table 1: Comparison with the SOTA methods on NAS-Bench-201. Query indicates the amount of labeled data needed for training the predictor, the less the better. R means ranking, pred means predictor, and zero means zero-cost estimation proxy.

Method	Accuracy(%)	Query	Type
HAAP (2021)	93.69±0.22	300	pred+EA
FreeREA (2023)	93.80±0.02	500	EA
ReNAS (2021)	93.90±0.21	423	pred+EA
CTNAS (2021a)	93.92±0.18	423	pred+RS
GMAE (2022)	93.98±0.15	300	pred+EA
HOP (2021b)	94.09±0.11	300	pred+RS
DCLP+RL	94.14±0.08	300	pred+RL
DCLP+RS	94.17±0.06	300	pred+RS

Table 2: The comparison of NAS on NAS-Bench-101. Query indicates the amount of labeled training data, the less the better, and pred indicates predictor.

2017). This is consistent with our intention of reducing dependence on labeled data by using curriculum-guided contrastive learning to train the encoder. While methods relying on supervised learning focus on improving the predictor structure to extract more knowledge from limited labeled data, the amount of information remains limited, leading to poor performance. Hence, DCLP’s advantage over these methods reinforces the effectiveness of contrastive pre-training in constructing robust neural predictors.

On the other hand, DCLP outperforms unsupervised predictors (Wei et al. 2021; Jing and Xu 2022; Liu, Tang, and Sun 2021), validating the effectiveness of our curriculum-guided learning. The moderate learning and ranking of training data enhance the contrastive learning process, aligning with our motivation. Since our key point is the pre-training for neural predictors, extending this approach to other methods is natural. This extension is critical in constructing effective predictors that rely on limited trained architectures.

4.2 Search Strategy Performance Evaluation

In this section, we compare DCLP with popular methods to demonstrate its effectiveness with search methods. We report our setups in Appendix D. As mentioned above, we aim to achieve the best results with the smallest labeled sets in the shortest time. Therefore, we introduce two evaluation metrics: *Query* and *Cost*. *Query* represents the number of trained networks required for predictor, and smaller values indicate more efficient training. *Cost* encompasses computa-

Method	Test error (%)	Cost(GPU h)		Type
		train	search	
DARTS- (2021b)	2.59±0.08	-	9.6	supernet
CTNAS (2021a)	2.59±0.04	120	2.4	pred+RL
CATE (2021)	2.56±0.08	-	75	pred+RS
HOP (2021b)	2.52±0.04	-	60	pred+RS
GMAE (2022)	2.50±0.03	40	0.2	pred+BO
DCNAS (2022)	2.50±0.01	-	480	supernet
DCLP-RL	2.50±0.02	4.0	0.08	pred+RL
DCLP-RS	2.48±0.02	4.0	0.05	pred+RS

Table 3: The comparison of NAS in DARTS space with CIFAR-10 as the dataset. The pred means predictor.

tional power consumed in training and search. For predictor-based methods, it includes the cost of training neural networks to obtain training data, along with the time for training the predictor and search process. The unit of cost is GPU hour (i.e. GPU h), which represents one hour of calculation on a GPU. We use a single RTX3090 as the platform.

Tables 1 and 2 present experiments on NAS-Bench-201 and NAS-Bench-101, respectively. DCLP achieves the highest precision, with minimal training data and cost. This is attributed to leveraging knowledge from unlabeled data effectively. Additionally, DCLP achieves better results through curriculum compared to the predictor which employs self-supervised methods. Moreover, it offers nearly zero-cost estimation, accelerating evaluation while maintaining superior accuracy compared to agent-dependent methods. Interestingly, combining DCLP with RS yields the best results as RS is a strong method (Yang and Esperança 2020; Bergstra and Bengio 2012). DCLP enables RS to explore a wide range effectively, enhancing its performance.

Table 3 and 4 present experimental results on DARTS space for image classification and language model tasks. More results on ImageNet are available in the Appendix. DCLP outperforms predictor-based and other evaluation strategies, showcasing its superiority in large search spaces with reduced training and search costs. Notably, DCLP achieves better results in less time compared to other methods, further validating its effectiveness. Additionally, DCLP’s success in the large CNN and RNN search spaces suggests its potential for broader application to various neural structures, reinforcing the effectiveness of our




Figure 3: The ablation of contrastive learning, curriculum learning, and fine-tuning targets on NAS-Bench-101.

Method	Test PPL	Cost(GPU h)		Type
		train	search	
DARTS- (2021b)	59.21	-	4.2	supernet
DCNAS (2022)	58.21	-	3.2	supernet
β -DARTS (2022)	58.17	-	3.1	supernet
GAME (2022)	57.41	3.5	0.15	pred+BO
HOP (2021b)	57.34	-	4.0	pred+RS
CL-V2 (2022)	56.94	-	4.8	supernet
DCLP-RS	56.19	2.0	0.04	pred+RS

Table 4: The comparison of NAS in DARTS space, and we use PTB as the dataset for the language model task. Pred means predictor, and zero means zero-cost estimation proxy.

curriculum-based contrastive training for neural predictors.

4.3 Ablation Study

With/Without Contrastive Learning. In Figure 3(a), we demonstrate that using contrastive learning for pre-training is superior to training the predictor from scratch. This supports the objective of using contrastive learning to acquire information about the features of neural networks through the contrastive task on unlabeled data. By leveraging this information, the labeled data needed for the prediction task’s training process is substantially decreased, leading to resource savings while obtaining a better predictor.

Choice of Curriculum Learning Methods. As shown in Figure 3(b), we compare results without curriculum and using different curriculum methods, revealing a significant impact. Surprisingly, learning from hard to easy is worse than no curriculum. This is because determining the similarity of neural networks is challenging, even when simplified as graph isomorphism (NP-hard). This leads to an unclear boundary between difficult positive and negative samples so the difficult data will contain noise, i.e., negative samples are incorrectly labeled as positive. Then, if the training is too difficult at the beginning, the encoder tends to focus more on the noise that often appears in hard data, which degrades performance. On the other hand, learning without a curriculum is ineffective as the learner struggles to converge with randomly selected data of varying difficulties. Tasks progressing from simple to hard yield excellent results, confirming

Contrastive Learning	Curriculum Learning	Ranking Loss	Test error (%)
✗	✗	✗	6.42±2.00
✗	✗	✓	5.80±1.75
✓	✗	✓	5.36±0.95
✓	easy to hard	✓	2.66±0.17
✓	ours	✓	2.48±0.02

Table 5: The ablation experiment in DARTS space on CIFAR-10. If without pre-training, we use the time for pre-training to obtain more labeled data for supervised training.

the correctness of the overall trend. However, using complex data introduces noise, requiring some easy data for parameter optimization. Our method cycles between hard and easy tasks, avoiding noise exploration and ensuring the ability to learn complex data for improved performance.

Choice of Fine-Tuning Targets. Figure 3(c) shows the impact of different fine-tuning functions. The choice of target function significantly affects the predictor’s performance. Notably, ranking loss and MSE+norm outperform vanilla MSE. This result aligns with our concept. The effectiveness of using ranking as the optimization target in this task is demonstrated. As ranking loss emphasizes the rankings, it shows particular advantages as training sets grow in size.

Table 5 compares the effects of various techniques in DCLP using DARTS space. When pre-training is not used, we allocate the time to obtain more labeled training data. The results further support the significance of curriculum-guided contrastive learning for high-performance predictors.

5 Conclusion

In this paper, we propose DCLP, a curriculum-guided contrastive pre-training approach that improves neural predictor performance in NAS. By pre-training the encoder with difficulty-scheduled positive data, the training process becomes stable and efficient. Fine-tuning with a suitable target further improves the predictor. DCLP reduces reliance on labeled data, enabling efficient and accurate predictor-based NAS. It offers a promising approach for neural predictors in NAS and future work will focus on designing self-supervised pretext tasks to boost performance further.

Acknowledgments

This work is supported by the National Key Research and Development Program of China (2021YFB3300502); National Natural Science Foundation of China (NSFC) (62232005, 62202126, 92267203).

References

- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48.
- Bergstra, J.; and Bengio, Y. 2012. Random search for hyperparameter optimization. *Journal of machine learning research*, 13(2).
- Burges, C. J. 2010. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11(23-581): 81.
- Cai, H.; Zhu, L.; and Han, S. 2019. ProxylessNAS: Direct Neural Architecture Search on Target Task and Hardware. In *International Conference on Learning Representations*.
- Cavagnero, N.; Robbiano, L.; Caputo, B.; and Averta, G. 2023. FreeREA: Training-Free Evolution-based Architecture Search. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1493–1502.
- Chen, W.; Gong, X.; and Wang, Z. 2021. Neural Architecture Search on ImageNet in Four GPU Hours: A Theoretically Inspired Perspective. In *International Conference on Learning Representations*.
- Chen, Y.; Guo, Y.; Chen, Q.; Li, M.; Zeng, W.; Wang, Y.; and Tan, M. 2021a. Contrastive neural architecture search with neural architecture comparators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9502–9511.
- Chen, Z.; Zhan, Y.; Yu, B.; Gong, M.; and Du, B. 2021b. Not All Operations Contribute Equally: Hierarchical Operation-adaptive Predictor for Neural Architecture Search. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10508–10517.
- Chu, G.; Wang, X.; Shi, C.; and Jiang, X. 2021a. CuCo: Graph Representation with Curriculum Contrastive Learning. In *IJCAI*, 2300–2306.
- Chu, X.; Wang, X.; Zhang, B.; Lu, S.; Wei, X.; and Yan, J. 2021b. {DARTS}-: Robustly Stepping out of Performance Collapse Without Indicators. In *International Conference on Learning Representations*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Dong, X.; and Yang, Y. 2020. NAS-Bench-201: Extending the Scope of Reproducible Neural Architecture Search. In *International Conference on Learning Representations*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- He, X.; and Zhao, K. 2021. AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212: 106622.
- Huang, T.; You, S.; Wang, F.; Qian, C.; Zhang, C.; Wang, X.; and Xu, C. 2022. Greedynasv2: Greedier search with a greedy path filter. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11902–11911.
- Jing, K.; and Xu, J. 2022. Graph Masked Autoencoder Enhanced Predictor for Neural Architecture Search. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 3114–3120. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Jing, L.; and Tian, Y. 2020. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11): 4037–4058.
- Kipf, T. N. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations*.
- Li, G.; Yang, Y.; Bhardwaj, K.; and Marculescu, R. 2023. ZiCo: Zero-shot NAS via inverse Coefficient of Variation on Gradients. In *The Eleventh International Conference on Learning Representations*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2019. DARTS: Differentiable Architecture Search. In *International Conference on Learning Representations*.
- Liu, Y.; Tang, Y.; and Sun, Y. 2021. Homogeneous architecture augmentation for neural predictor. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12249–12258.
- Mellor, J.; Turner, J.; Storkey, A.; and Crowley, E. J. 2021. Neural architecture search without training. In *International Conference on Machine Learning*, 7588–7598. PMLR.
- Ning, X.; Zheng, Y.; Zhao, T.; Wang, Y.; and Yang, H. 2020. A generic graph-based neural architecture encoding scheme for predictor-based nas. In *European Conference on Computer Vision*, 189–204. Springer.
- Pan, J.; Sun, C.; Zhou, Y.; Zhang, Y.; and Li, C. 2022. Distribution Consistent Neural Architecture Search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10884–10893.
- Pentina, A.; Sharmanska, V.; and Lampert, C. H. 2015. Curriculum learning of multiple tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5492–5500.
- Platanios, E. A.; Stretcu, O.; Neubig, G.; Poczos, B.; and Mitchell, T. 2019. Competence-based Curriculum Learning for Neural Machine Translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language*

- Technologies, Volume 1 (Long and Short Papers)*, 1162–1172. Minneapolis, Minnesota: Association for Computational Linguistics.
- Qin, Y.; Zhang, Z.; Wang, X.; Zhang, Z.; and Zhu, W. 2022. NAS-Bench-Graph: Benchmarking Graph Neural Architecture Search. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Sen, P. K. 1968. Estimates of the regression coefficient based on Kendall’s tau. *Journal of the American statistical association*, 63(324): 1379–1389.
- Tang, Y.; Wang, Y.; Xu, Y.; Chen, H.; Shi, B.; Xu, C.; Xu, C.; Tian, Q.; and Xu, C. 2020. A semi-supervised assessor of neural architectures. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1810–1819.
- Wang, R.; Chen, X.; Cheng, M.; Tang, X.; and Hsieh, C.-J. 2021. RANK-NOSH: Efficient Predictor-Based Architecture Search via Non-Uniform Successive Halving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10377–10386.
- Wang, X.; Chen, Y.; and Zhu, W. 2021. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wei, C.; Tang, Y.; Niu, C. N. C.; Hu, H.; Wang, Y.; and Liang, J. 2021. Self-supervised representation learning for evolutionary neural architecture search. *IEEE Computational Intelligence Magazine*, 16(3): 33–49.
- Wen, W.; Liu, H.; Chen, Y.; Li, H.; Bender, G.; and Kindermans, P.-J. 2020. Neural predictor for neural architecture search. In *European Conference on Computer Vision*, 660–676. Springer.
- Wu, L.; Lin, H.; Tan, C.; Gao, Z.; and Li, S. Z. 2021. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE Transactions on Knowledge and Data Engineering*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- Xu, Y.; Wang, Y.; Han, K.; Tang, Y.; Jui, S.; Xu, C.; and Xu, C. 2021. Renas: Relativistic evaluation of neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4411–4420.
- Yan, S.; Song, K.; Liu, F.; and Zhang, M. 2021. Cate: Computation-aware neural architecture encoding with transformers. In *International Conference on Machine Learning*, 11670–11681. PMLR.
- Yang, A.; and Esperança, P. M. 2020. NAS evaluation is frustratingly hard. In *International Conference on Learning Representations*.
- Ye, P.; Li, B.; Li, Y.; Chen, T.; Fan, J.; and Ouyang, W. 2022. b-darts: Beta-decay regularization for differentiable architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10874–10883.
- Ying, C.; Klein, A.; Christiansen, E.; Real, E.; Murphy, K.; and Hutter, F. 2019. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, 7105–7114. PMLR.
- Zhang, X.; Xu, H.; Mo, H.; Tan, J.; Yang, C.; Wang, L.; and Ren, W. 2021. Dcnas: Densely connected neural architecture search for semantic image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13956–13967.
- Zhou, Y.; Wang, X.; Chen, H.; Duan, X.; Guan, C.; and Zhu, W. 2022. Curriculum-nas: Curriculum weight-sharing neural architecture search. In *Proceedings of the 30th ACM International Conference on Multimedia*, 6792–6801.
- Zoph, B.; and Le, Q. 2017. Neural Architecture Search with Reinforcement Learning. In *International Conference on Learning Representations*.
- Zoph, B.; Vasudevan, V.; Shlens, J.; and Le, Q. V. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8697–8710.