

# SwiftPruner: Reinforced Evolutionary Pruning for Efficient Ad Relevance

Li Lyna Zhang  
Microsoft Research  
lzhani@microsoft.com

Youkow Homma  
Microsoft  
youkow.homma@microsoft.com

Yujing Wang  
Microsoft  
yujwang@microsoft.com

Min Wu  
Microsoft  
albw@microsoft.com

Mao Yang  
Microsoft Research  
maoyang@microsoft.com

Ruofei Zhang  
Microsoft  
bzhang@microsoft.com

Ting Cao  
Microsoft Research  
ting.cao@microsoft.com

Wei Shen\*  
Microsoft  
sashen@microsoft.com

## ABSTRACT

Ad relevance modeling plays a critical role in online advertising systems including Microsoft Bing. To leverage powerful transformers like BERT in this low-latency setting, many existing approaches perform ad-side computations offline. While efficient, these approaches are unable to serve cold start ads, resulting in poor relevance predictions for such ads. This work aims to design a new, low-latency BERT via structured pruning to empower real-time online inference for cold start ads relevance on a CPU platform. Our challenge is that previous methods typically prune all layers of the transformer to a high, uniform sparsity, thereby producing models which cannot achieve satisfactory inference speed with an acceptable accuracy.

In this paper, we propose SwiftPruner - an efficient framework that leverages evolution-based search to automatically find the best-performing *layer-wise sparse* BERT model under the desired latency constraint. Different from existing evolution algorithms that conduct random mutations, we propose a reinforced mutator with a latency-aware multi-objective reward to conduct better mutations for efficiently searching the large space of layer-wise sparse models. Extensive experiments demonstrate that our method consistently achieves higher ROC AUC and lower latency than the uniform sparse baseline and state-of-the-art search methods. Remarkably, under our latency requirement of 1900us on CPU, SwiftPruner achieves a **0.86%** higher AUC than the state-of-the-art uniform sparse baseline for BERT-Mini on a large scale real-world dataset. Online A/B testing shows that our model also achieves a significant 11.7% cut in the ratio of defective cold start ads with satisfactory real-time serving latency.

\*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CIKM '22, October 17–21, 2022, Atlanta, GA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9236-5/22/10...\$15.00

<https://doi.org/10.1145/3511808.3557139>

## CCS CONCEPTS

• Information systems → Information retrieval; • Computing methodologies → Artificial intelligence.

## KEYWORDS

Transformer, Pruning, Neural Architecture Search, Ad Relevance

### ACM Reference Format:

Li Lyna Zhang, Youkow Homma, Yujing Wang, Min Wu, Mao Yang, Ruofei Zhang, Ting Cao, and Wei Shen. 2022. SwiftPruner: Reinforced Evolutionary Pruning for Efficient Ad Relevance. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management (CIKM '22)*, October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3511808.3557139>

## 1 INTRODUCTION

In sponsored search, ad relevance measures the semantic similarity between a user's search query and an ad. Given a user query, the system calculates its relevance with hundreds of potential ads to prevent irrelevant ads recommendation and to assist ranking and pricing of the most relevant ads. As such, ad relevance directly affects user and advertiser satisfaction with the product and serves as one of the most important tasks in online advertising.

At Microsoft Bing, we previously released the TwinBERT [21] and AutoADR [7] models which leverage the latest advances in pre-trained language models (e.g., BERT) [8, 34] and neural architecture search (NAS) to effectively predict ad relevance. Due to the expensive inference cost of BERT, we decoupled the inferences for ad and query via a two-tower structure which consists of online query-side and offline ad-side sub-models. While this design avoids the time-infeasible online ad-side inference for hundreds of ads per query, the offline ad-side inference presents its own challenges. In particular, we encounter **coverage issues** due to *cold start ads*, or *cold ads* for short<sup>1</sup>: for a new ad, there is an inevitable delay between the time when it is available for online serving and when its offline-computed ad-side features are published online.

In this work, we aim to design a new semantic query-ad relevance model for cold ads which can be computed fully online

<sup>1</sup>When the offline-computed ad-side data is not yet available during ad serving time online, we call such ads as cold ads.

on CPU under the real-time latency constraints. We therefore discard the two-tower structure and instead focus on **single-tower** BERT [8] models where the query and ad content are concatenated as the model input. Considering the real-time inference constraints, we select the state-of-the-art tiny transformer, BERT-Mini [33] as our model structure for cold ads, which achieves comparable accuracy with our previous models (*i.e.*, TwinBERT and AutoADR) via knowledge distillation. Although optimized with both knowledge distillation [28, 33] and quantization [3, 15, 44], the inference latency of BERT-Mini is still too high to be served online on our CPU-based platform. To afford the real-time inference for cold ads relevance, we call for new model compression methods to effectively reduce the latency of tiny transformers.

Structured pruning methods have shown to be effective in reducing inference latency on standard hardware while mitigating accuracy loss. Previous works [16, 17, 23, 40, 42] have demonstrated the large redundancy in transformers, but they focus on large transformers (e.g., BERT-Base and BERT-Large) and simply set a uniform sparsity ratio for all layers. However, we observe that individual attention and feed-forward (FFN) layers have different redundancy in tiny transformers (*i.e.* BERT-Mini). Pruning some critical encoder layers can lead to a significant accuracy drop while pruning other redundant layers can reduce latency without affecting the model accuracy. Moreover, attention and FFN layers contribute differently to the model latency. Therefore, we study layer-wise sparsity to achieve better trade-offs between accuracy and latency.

Layer-wise sparsity leads to a large design space to compromise between model accuracy and latency. While it is challenging to determine the sparsity for each layer from such a large space, NAS approaches such as reinforcement learning (RL) [26, 45] and evolution search (EA) [11, 27] have demonstrated effectiveness in automatically finding good architectures. Inspired by this, the questions naturally arise: *Can we leverage NAS to search the space of layer-wise sparse models to find the best performing model under a given latency constraint? What kind of NAS algorithm is the most effective for our ad relevance scenario?*

To this end, we propose SwiftPruner, which automatically searches the best-performing layer-wise sparsity for structured pruning under a given latency constraint. Inspired by NAS, we formulate the pruning problem as a multi-objective optimization problem that considers both accuracy and inference latency of pruned transformers. First, we design a large search space that factorizes the BERT-Mini structure into layer-wise sparse configurations. Then, an efficient search algorithm is employed to identify layer-wise sparse settings that can achieve both low latency and high accuracy. Finally, we select the pruned model with maximum validation Receiver Operating Characteristic (ROC) AUC under a certain latency constraint as the final model, named SwiftBERT.

Despite the success of EA and RL in NAS, both of them face specific challenges in our scenario. EA-based search shows simplicity and stability in many tasks [11, 27]. However, its evolution process is highly reliant on the mutation actions, and the search efficiency has no guarantee due to the uncontrollable random mutation. RL is effective but more time consuming than EA. Additionally, the stability of RL depends on careful hyper-parameter selection. Thus, it's difficult to tune for jointly optimizing accuracy and latency in our scenario.

To address the above limitations, we propose a reinforced evolution algorithm to search for the optimal layer-wise sparsity with a latency-aware multi-objective reward for transformers. Unlike traditional EA methods that randomly select a layer for sparsity mutation, our method introduces a reinforced mutator to learn (*i*) which layer to mutate and (*ii*) what sparsity value to assign to the mutated layer. Our search method leverages the advantages of both EA and RL: we not only accelerate the search via reinforced mutations but also maintain the stability of EA. Furthermore, to reduce the search cost caused by latency measurement, we build a latency predictor that can accurately predict the inference latency for sparse transformers.

We conduct extensive experiments with a real-world dataset for ad relevance under many latency constraints. For all latency constraints, SwiftPruner consistently achieves higher AUC and lower latency compared to the uniform sparsity baseline and state-of-the-art layer-wise search methods. Remarkably, SwiftPruner can achieve a maximum of 43.46% latency reduction on CPU for the BERT-Mini with a minimal AUC loss of 0.32%. Compared to other search methods, SwiftPruner demonstrates its superiority through higher AUC (up to +0.48%) under all latency constraints, better search efficiency, as well as robustness to hyper-parameter selection. Finally, online A/B testing shows that SwiftBERT achieves a significant 11.7% cut in the ratio of defective cold ads and 2.37% increase in click-through-rate. This new model has been shipped into the Microsoft Bing ad relevance production model.

We summarize the main contributions of our work as follows:

- We propose SwiftPruner, an efficient framework that leverages NAS to search for the best-performing layer-wise sparse BERT under a desired latency constraint.
- We introduce a reinforced evolutionary search algorithm which efficiently predicts a layer and sparsity value for mutation. We demonstrate its superiority over EA and RL-based search methods in our scenario.
- We apply SwiftPruner to ad relevance prediction and produce SwiftBERT models with various latency levels. Compared to uniform sparsity and other layer-wise search methods, SwiftBERT shows much higher AUC and lower inference latency.
- After adding SwiftBERT to the ad relevance production pipeline, the ratio of defective cold ads is significantly reduced according to A/B testing results. To the best of our knowledge, SwiftBERT is the first single-tower, query-document crossing BERT architecture at Microsoft Bing that serves fully online within a real-time latency on CPU. Our method is also applicable to other scenarios such as dynamically generated ads [13, 38] and search relevance [43].

## 2 RELATED WORKS

**Knowledge distillation.** Knowledge distillation [14, 28, 32] is a powerful method which leverages scores, structures, or weights from a large, teacher BERT model to train a smaller, student BERT model. While students that are 2-3× smaller than the teacher can be trained with minimal accuracy loss, Turc et al. [33] show that there is a steep drop-off in accuracy when distilling to tiny transformer students like those required in our scenario.

**Quantization.** Quantization [15, 44] of transformer models from FP32 to Int8 can produce models up to 4× as small with negligible accuracy loss regardless of the starting model size. While the model compression ratio from quantization is impressive, inference speedup often relies on specialized hardware [5, 44]. On our CPU deployment platform which is a standard Intel CPU, we find the benefit from model quantization is limited due to the lack of hardware support [3].

**Pruning** can be categorized into: 1) unstructured and 2) structured pruning. Although unstructured methods [10, 29, 36] reach a relatively high sparsity ratio without significant accuracy drop, they yield very little actual latency benefits. Many common deployment platforms on modern hardware, particularly on CPU, often keep the original dense shapes of the uncompressed model as it is difficult to leverage irregular sparse patterns for acceleration.

Structured pruning, on the other hand, allows for the removal of coherent groups of weights and can reduce latency without special hardware support. Early works [24, 35] showed that some attention heads in BERT can be removed without significant accuracy loss. Later, structured pruning was also expanded to act on both the attention and FFN layers of transformers [16, 22, 23, 42]. Recently, Lagunas et al. [17] introduced a hybrid pruning named *nn\_pruning* which prunes structured blocks in attention heads and paired rows/columns in FFN intermediate layers. When combined with movement pruning [29], which learns the importance score for weights during fine-tuning, *nn\_pruning* is able to achieve effective compression of BERT models. However, these methods focus on maximum parameter removal by ranking the importance scores. They pre-define a target sparsity (compression ratio) and assign a uniform sparsity for all layers. In our work, we find our production BERT models have different redundancy per layer and hence require layer-wise sparsity ratios.

**Layer-wise sparsity.** In computer vision tasks, it has been demonstrated that different CNN layers have different redundancy, and rule- or heuristic-based unstructured pruning [9, 18, 31] may be used to find the appropriate layer-wise sparsities to reduce their redundancy. AMC [12] is the first work that leveraged reinforcement learning (RL) to automatically search the optimal layer-wise sparsity for CNNs. To accelerate exploration, AMC conducts weight magnitude pruning. However, this method is less effective for fine-tuned BERT models on downstream tasks [29]. Unlike these previous works, we focus on pruning BERT models for NLP tasks. In our work, we propose reinforced evolutionary search, which is more efficient than RL. During the search, we conduct movement-based structured pruning rather than magnitude pruning. Therefore, we can efficiently find the optimal layer-wise sparsity for a target latency constraint with minimal accuracy loss.

**AutoML.** Hyper-parameter optimization (HPO) and neural architecture search (NAS) are both prevalent approaches for conducting efficient optimization within a large search space. The Tree Parzen Estimator (TPE) [1] algorithm is one of the most popular methods for hyper-parameter optimization which utilizes Bayesian optimization to search high-dimensional conditional spaces. However, TPE relies on the quality of initial samples and has limitations in robustness [30]. On the other hand, reinforcement learning (RL) and evolution search (EA) have achieved state-of-the-art performance

in NAS tasks. Works focusing on RL-based search [26, 45] train an RNN controller with reinforcement learning to determine a sequence of operators and connections that specify the architecture of a neural network. One limitation of RL-based search is the need for careful hyper-parameter selection to guarantee stability. In contrast, EA is less sensitive to hyper-parameter selection. Real et al. [27] introduces aging evolution, which samples a best parent candidate at each iteration and then conducts random mutations to create children candidates. While EA addresses RL’s shortcomings, EA’s search efficiency cannot be guaranteed due to its heavy reliance on random mutation.

In our work, we propose a reinforced RNN controller to conduct predictable mutation which greatly accelerates EA’s search efficiency. The idea of reinforced mutations on EA has been previously investigated in RENAS [6], but they focus on architecture search for CV tasks without considering inference latency. To the best of our knowledge, our method is the first to propose and apply this idea to prune NLP models to a desired latency.

## 3 AD RELEVANCE SCENARIO

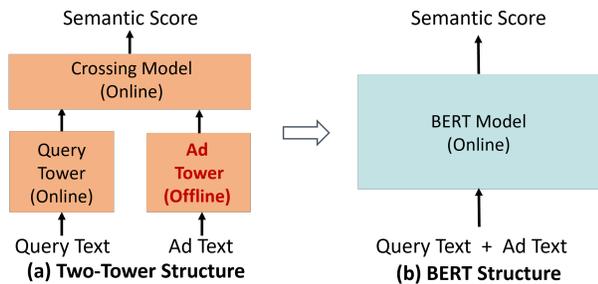
### 3.1 Background

In sponsored search systems, the user submits a *query* to the search engine, and the engine delivers *ads* along with the search results. The ads shown to the user are chosen based on a number of parameters including *relevance* between the query and ad, pricing, and user behavior. When ads are more relevant to the user’s query, the user is more satisfied with the search product and the advertiser receives increased value from their bid. Thus, ad relevance serves as a cornerstone to a healthy sponsored search ecosystem.

The Microsoft Bing ad relevance system consists of a cascade of models which feed into one another to make many ad relevance-based decisions in the ad stack. These decisions include filtering out the most irrelevant ads, placing relevant ads higher in the search results, and pricing ads based on their relevance. One of the primary goals of the ad relevance system is to minimize the number of *defective* ad impressions. An ad is determined to be defective for a query if their semantic similarity is low as determined by a set of pre-determined guidelines. As such, some of the most important models in the ad relevance system are the query-ad **semantic relevance models**. These semantic models take the query and ad text as input and produce a prediction of defectiveness, called the **semantic score**, which is then used by other models.

Currently, our system contains a few semantic relevance models with the most powerful being TwinBERT [21] and AutoADR [7] which leverage the latest advances in transformers. These two models share a similar two-tower structure as shown in Fig. 1(a) which decouples the query and ad inferences for efficient computation. The query-side model serves online and is computed just once per user query. On the other hand, the ad-side model is inferred offline for the billions of servable ads, and the embedding outputs are then published to the online service. Given a candidate query-ad pair, the semantic score is then computed online by running a small crossing model on the query and ad tower outputs.

The two-tower structure allows us to use large, powerful models in our production system; however, it also introduces a challenge. When advertisers create new, servable ads, there is a delay between



**Figure 1: Comparison of the hybrid offline-online two-tower structures of TwinBERT and AutoADR (a) vs the fully online BERT structure of SwiftBERT (b).**

Model	Size	AUC (%)	Latency (us)
BERT Teacher Model	2B	88.98	-
BERT-Mini (4L-256H) [33]	11.3M	87.15	3274.24
2L-256H [33]	9.7M	85.16	1727.67
4L-128H [33]	4.8M	84.23	1841.29

**Table 1: AUC and latency of distilled models from [33].**

when the ad can be shown to users and when the ad-side tower computation is published to the online system. This delay results in ads that are missing their offline-computed ad-side data at the online serving stage. We call such ads, **cold ads**. Conversely, we call ads that have their ad-side data online, hot ads. For cold ads, our ad relevance system currently relies on weak semantic relevance models that do not leverage transformers. As a result, millions of defective cold ads pollute the system daily and degrade the overall quality of ads shown to users. To address these defective cold ads, we call for a new, powerful transformer-based semantic relevance model which can be used to specifically serve cold ads.

### 3.2 Serving Cold Ad Relevance

In order to compute a semantic score for cold ads, we cannot use offline components in the model. We therefore discard the two-tower structure and instead opt for a fully online, single-tower BERT where the input is a concatenation of the query and ad text as shown in Fig. 1(b). While this single-tower structure addresses the coverage issue, a new challenge arises. For each query-hot ad pair, the two-tower structure allows us to run just a small crossing model online for each query-ad pair to get their semantic score. This crossing model can be run on hundreds of ads per query with minimal impact on system latency. On the other hand, for every query-cold ad pair, we must now run an entire BERT model online. Inferencing even a very small BERT model on tens of ads per query has a measurable impact on overall system latency. Therefore, our new cold ads model must be designed to run within a strict latency budget. Specifically, we find that the new single tower model must run in less than 1900us per query-cold ad pair on our platform.

To start, we distill some of the smallest, state-of-the-art tiny transformers from Turc et al. [33] which meet our system latency budget of 1900us. The chosen models are a 2-layer, 256-hidden dimensional BERT (2L-256H) and a 4-layer, 128-hidden dimensional BERT (4L-128H). After distilling the two models using our powerful

Layer	Base config.	Sparsity space	Prune pattern
MSA ( $a$ )	4	(0, 0.25, 0.5, 0.75)	head
FFN ( $f$ )	1024	(0, 0.01, 0.02, ...0.99)	dimension

**Table 2: Candidate sparsity for each encoder layer. The total search space contains  $\sim 2 \times 10^{10}$  settings for pruning.**

teacher model, we find there is a large gap in ROC AUC between these two students and the teacher as shown in Table 1.

To strike a better balance between prediction performance and latency, we examine the next largest model checkpoint from [33]: BERT-Mini, which is a 4-layer, 256-hidden dimensional BERT model with 4 attention heads and 1024 FFN dimensions. While BERT-Mini achieves significantly better performance as shown in Table 1, its inference time is over 70% above the latency limit. In order to retain its good performance while decreasing the latency, we explore different model compression methods. We first apply Int8 quantization [15, 44] on our CPU deployment platform, but the inference latency is still too high beyond our requirement. We next turn to pruning to close the remaining latency gap. In the following sections, we study how our SwiftPruner method can be applied to BERT-Mini to generate a new, low-latency semantic relevance model for cold ads, which we call SwiftBERT.

## 4 METHODOLOGY

In this section, we present the SwiftPruner for layer-wise pruning of tiny transformers (i.e., BERT-Mini) such that the pruned model meets the given latency constraints.

We formulate our problem as:

$$\begin{aligned} (a_1, f_1, \dots, a_l, f_l)^* &= \arg \min_{a_1, f_1, \dots, a_l, f_l} L(A(a_1, f_1, \dots, a_l, f_l)) \\ \text{s.t. } LAT(A(a_1, f_1, \dots, a_l, f_l)^*) &\leq T, \end{aligned} \quad (1)$$

where  $A(\cdot)$  is the transformer model with a specific sparsity setting, and  $a_i$  and  $f_i$  are the sparsity ratios for the attention and FFN subnetworks in the  $i^{th}$  encoder layer, respectively. SwiftPruner aims to find the optimal layer-wise sparsity  $(a_1, f_1, \dots, a_l, f_l)^*$  from the  $1^{st}$  to  $l^{th}$  encoder layer that has the minimum cross-entropy loss,  $L$ , (corresponding to the maximum accuracy) while the latency (denoted by  $LAT$ ) also meets the constraint,  $T$ .

To achieve this, we first design a fine-grained sparsity search space which contains a large amount of models that can meet various latency constraints after pruning. Then we introduce a reinforced evolutionary search algorithm equipped with a latency predictor. Instead of random mutation, the algorithm learns to mutate a layer with predicted sparsity towards the best accuracy-latency trade-off. Fig. 2 illustrates the overall pipeline.

### 4.1 Search Space Design

Unlike previous NAS works [6, 41] that search the optimal operator for each layer, our goal is to search for the optimal sparsity setting for each layer of the transformer model. Since our deployment platform has no support for accelerating unstructured pruned models, we design a large transformer search space for structured pruning, as detailed in Table 2. Following previous works [17, 23, 24, 40] which show that the majority of attention heads and FFN intermediate dimensions can be pruned with minimal accuracy loss, we

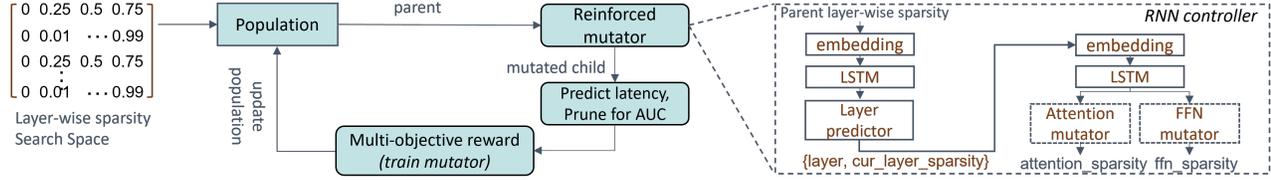


Figure 2: Overview of SwiftPruner for searching the optimal layer-wise sparse transformers.

also prune the attention layers via head pruning and FFN layers via dimension pruning. This hybrid pruning granularity setting has demonstrated a better accuracy-efficiency trade-off [17, 40]. Concretely, we use the entire head as the smallest pruning granularity in attention. The FFN layer contains two paired linear layers, so we prune the rows in the first layer (i.e., the intermediate layer output) and the corresponding columns in the second linear layer (i.e., the second layer input).

**Head sparsity.** For an attention layer with  $H$  heads, we allow pruning of  $\{0, 1, \dots, H-1\}$  heads (at least one head is retained to avoid pruning the entire layer). Since our base model, BERT-Mini [33], has 4 attention heads in one encoder layer, we have 4 candidate settings of sparsity ratio: 0, 0.25, 0.5 and 0.75, which correspond to the pruning of 0, 1, 2 and 3 heads, respectively.

**FFN sparsity.** The original FFN intermediate dimension size  $F$  of BERT-Mini is 1024. A full search space would contain all possible dimension sizes,  $F$ , but would also lead to search space explosion for larger dimensions. Fortunately, we observe that there are negligible latency differences between  $F$  and  $F+1$  on our platform. Therefore, we reduce the search space of dimensions to 100 candidates, where different sparsity leads to noticeable latency change. As shown in Table 2, we allow FFN layer sparsity ratios to range from 0 to 0.99 with a step size of 0.01 (i.e., a step of  $\sim 10$  in dimension size).

In total, for our 4 encoder-layer SwiftBERT model, the search space contains  $\sim 2 \times 10^{10}$  candidate settings, which is extremely large and poses practical challenges for traditional NAS search algorithms. If a candidate model could be evaluated in a few seconds via weight sharing and performance prediction [4, 25, 37], a naive evolution algorithm could quickly find good performing models by evaluating thousands of candidates. However, evaluating the AUC for layer-wise settings requires considerable training cost on our large-scale real-world dataset. Unlike in NAS tasks, introducing weight sharing in pruning poses difficulty in accurate AUC evaluation. Therefore, we require a higher efficiency search algorithm that can find the optimal layer-wise sparse setting in a few hundred search trials. In the following sub-section, we introduce our proposed search method to tackle this problem.

## 4.2 Reinforced Aging Evolution

To efficiently search the large space of candidates, we propose the reinforced aging evolution as shown in Fig. 2. Algorithm 1 formulates the overall procedure of SwiftPruner. Following the original aging evolution, we first randomly initialize a population of  $P$  models from the search space, where each model is encoded with a specific layer-wise sparsity setting. After this, evolution improves the initial population in mutation iterations (line 3). At each iteration, we

### Algorithm 1 Reinforced Aging Evolution Algorithm

**Input:** population size  $P$ , total number of models to explore  $N$ , sample size  $S$ , latency\_auc tradeoff  $\alpha$ , target latency  $T$

- 1:  $population^{(0)} \leftarrow \text{initialize}(P)$
- 2:  $history \leftarrow population^{(0)}$
- 3: **for**  $i=1: (N-P)$  **do**
- 4:    $parent \leftarrow \text{sample\_with\_maximum\_reward}(population^{(i-1)}, S)$
- 5:    $child \leftarrow \text{reinforced\_mutate}(parent)$
- 6:    $reward \leftarrow \text{get\_reward}(child, T, \alpha)$  (see Equation 2)
- 7:   update mutator with maximizing  $reward$  by REINFORCE [39]
- 8:    $population^{(i)} \leftarrow \text{add } child \text{ to right of } population^{(i-1)}$ , remove the oldest model from left of  $population^{(i-1)}$
- 9:   add  $child$  to  $history$
- 10: **end for**
- 11: return maximum-AUC model under the  $T$  from the  $history$

sample  $S$  random models from the population and select the model with maximum reward as the *parent*. We then run the reinforced mutator to mutate a layer sparsity on the *parent* and construct a new model, called the *child*. Once the *child* is constructed, we compute its reward by Equation 2, which considers both accuracy and inference latency. We train the reinforced mutator by maximizing the reward with REINFORCE [39]. At the end of each iteration, we remove the oldest model in the current population and add the *child* to the population for the next iteration. After the evolution process finishes, we collect all the explored models that are under the target latency constraint  $T$  and select the model with maximum validation AUC as the final model.

**Reinforced mutator.** We now introduce the key component in reinforced aging evolution, namely the reinforced mutator, which conducts sparsity mutations on the sampled *parent* model. In the original aging evolution algorithm, the mutator randomly samples one layer and performs a random modification of the sparsity ratio for this layer. However, as described in Section 2, uncontrollable mutation cannot guarantee the efficiency of optimizing our reward.

Instead of random mutations, the reinforced mutator is designed to 1) predict a layer to mutate, and 2) predict a sparsity ratio for the selected layer, so that the mutated *child* has a better chance to achieve the optimal AUC-latency trade-off. Fig. 2 (right) illustrates the architecture of our reinforced mutator. Guided by the two prediction targets, we design an RNN controller which sequentially selects a new layer to mutate and the corresponding new sparsity ratio. The controller is implemented as a recurrent neural network consisting of two sub-networks.

In the first sub-network, we take the layer-wise sparsity of the *parent* as input. The sub-network starts with an embedding layer with a bidirectional LSTM to learn the effect of each layer-wise

sparsity ratio of the *parent*. Based on the output of the LSTM, we apply a fully connected layer with *Softmax* as the layer predictor. The predictor estimates the mutation probability of each layer, and we sample a layer  $l_i$  for mutation according to the probability.

After determining the layer  $l_i$  for mutation, we run the second sub-network to predict a new sparsity ratio for this layer. We take the predicted layer index  $l_i$  and its current sparsity ratio  $f_i$  of the *parent* as the input to this model. As shown in Fig. 2 (right), another LSTM layer is leveraged to learn the effect of changing the sparsity ratio. Since attention and FFN layers have different sparsity ratios (shown in Table 2), we design two different fully-connected classification layers for attention and FFN layers, called the *Attention mutator* and *FFN mutator*, respectively. The corresponding classification branch is selected according to the type of layer  $l_i$ . For instance, if layer  $l_i$  is an attention layer, we forward the input to the *Attention mutator* branch to predict a new sparsity ratio,  $f_i^*$ .

**Latency-aware reward.** The objective of our reinforced mutator is to maximize the latency-aware reward as follows:

$$\text{reward} = \text{AUC}(m) \times \left[ \frac{\text{LAT}(m)}{T} \right]^w \quad (2)$$

where  $\text{AUC}(m)$  denotes the AUC of model  $m$ ,  $\text{LAT}(m)$  denotes the inference latency on the target platform, and  $T$  is the target latency.  $w$  is the weight factor defined as:

$$w = \begin{cases} 0, & \text{if } \text{LAT}(m) \leq T \\ \alpha, & \text{otherwise} \end{cases} \quad (3)$$

In our scenario, we have a hard limit on the pruned model latency. Therefore, we set  $w = 0$  if the latency of the searched model  $m$  is less than  $T$  to find the layer-wise sparsity with the highest AUC under constraint,  $T$ . Otherwise, we set  $w = \alpha$  to penalize the objective value and discourage mutations that violate the latency constraint.  $\alpha$  is a negative number, and in our experiment, we empirically set  $\alpha = -1$ .

The computation of AUC and latency in the reward function is time-consuming. To eliminate the cost of latency measurement, we build a latency predictor that can accurately predict the latency for layer-wise sparse models. To reduce AUC computation cost, we conduct fast structured pruning on a small training set. We detail these optimizations below.

**Latency prediction.** Measuring the latency for a model requires deploying the model on the target platform and then running the model many times on the same input. As such, it would be time-consuming to measure the candidate model latency at each iteration of the reinforced aging evolution algorithm. In lieu of this, we build a latency predictor which can accurately predict the latency for each layer-wise sparse model on our deployment platform. First, we randomly sample layer-wise sparse models from our search space and measure the latency on the target platform. We then train a random forest regression model [19] to predict the latency values from the sampled models' layer-wise sparsity ratios.

**Structured pruning.** We conduct structured pruning to obtain the AUC in the reward function. It's non-trivial to select a structured pruning algorithm for our scenario. In our application, the base model BERT-Mini is first initialized with pretraining weights [33]

and then fine-tuned on our real-world training set. Since movement pruning [29] has demonstrated the advances in removing unimportant neurons for fine-tuning pretrained models, we leverage Hugging Face's `nn_pruning` library [17], which is an improved version of movement pruning to support structured pruning.

The original `nn_pruning` conducts block pruning in the attention layer. To determine which blocks are pruned, it divides weight matrices  $W_q$  (Query  $Q$ ),  $W_k$  (Key  $K$ ),  $W_v$  (Value  $V$ ), and  $W_o$  (output weight) into individual blocks. For instance, when the block size equals the head size, the Key matrix  $W_k$  in the 4-head BERT-Mini will be divided into  $W_{k1}$ ,  $W_{k2}$ ,  $W_{k3}$ ,  $W_{k4}$ . Each block is assigned with a score, which is used to measure the importance for block weights to final model accuracy. However, since the importance scores for each head's  $W_{ki}$ ,  $W_{qi}$ ,  $W_{vi}$  are learned independently, there is no guarantee to prune an entire head. Therefore, we cannot directly apply block pruning in our scenario. For example, when the sparsity is set to prune 1 head, block pruning removes one block in each  $W_q$ ,  $W_k$ ,  $W_v$ . Based on the importance scores, we might prune the first head by  $W_{q1}$ , the second head by  $W_{k2}$  and the third head by  $W_{v3}$ . In this case, no head would get removed and there would be no latency reduction.

To solve this problem, we apply a simple but effective method. For each head  $i$ , we constrain its corresponding four matrices  $W_{qi}$ ,  $W_{ki}$ ,  $W_{vi}$  and  $W_{oi}$  with one shared importance score, so that all matrices in one head will be pruned at the same time. We evaluate the two different implementations on pruning the same model (i.e., the output model with 80% sparsity using the original `nn_pruning` in Fig. 3), and they achieve comparable AUC scores.

Additionally, to further reduce the search cost, we randomly sample a small subset of the training data which is sufficiently large enough to evaluate model performance. As a result, we can quickly get the approximate AUC for ranking a candidate model in 4 minutes on a 4-V100 GPU node.

## 5 EVALUATION

### 5.1 Dataset and Setting

**Dataset.** The training set for our study is a large-scale sample of real impressed query-ad pairs from Microsoft Bing's search log. Since it would be laborious to manually label each query-ad pair in the training set as defective/non-defective, we follow the practice in AutoADR [7] of generating the defective probability for each sample using a powerful BERT teacher model. The teacher model itself is trained on millions of human-classified query-ad pairs. All of our pruned models are trained on the binary cross entropy loss between the teacher score and model score. The test set in our experiments contains 100k query-ad pairs which are classified as defective/non-defective by human judges, and we report the ROC AUC of the model scores on this set. The teacher has an AUC of 88.98% on the test set.

**Latency measurement.** To build the latency predictor, we randomly sample 5000 models from our sparsity search space and measure their inference latency on a Intel(R) Core(TM) i7-7700 CPU device. The inference framework is ONNX runtime 1.8.2. To make full utilization of ONNX optimizations, we apply quantization [3] from FP32 to Int8 as well as ONNX transformer optimizations [2] for all models. We report the average latency of 1000 runs with an

Model	Size	AUC (%)	Latency
BERT-Mini [33]	11.3M	87.15	3274.24 us
uniform pruning [17] (85% sparsity)	8.86M	85.97	1876.44 us
SwiftBERT	8.99M	<b>86.83</b>	<b>1851.16 us</b>

**Table 3: End-to-end performance under the required 1900us latency constraint. SwiftBERT achieves +0.86% better AUC than uniform pruned model with smaller latency.**

input sequence length of 38 and batch size of 1. The sampled layer-wise sparsity and measured latency pairs are split 8:2 between the training and validation sets for the latency prediction model. The RMSE is only 70us and RMSPE is 3.88% on the validation data. This suggests that the latency prediction model can be used to replace the expensive latency measurement with little error introduced.

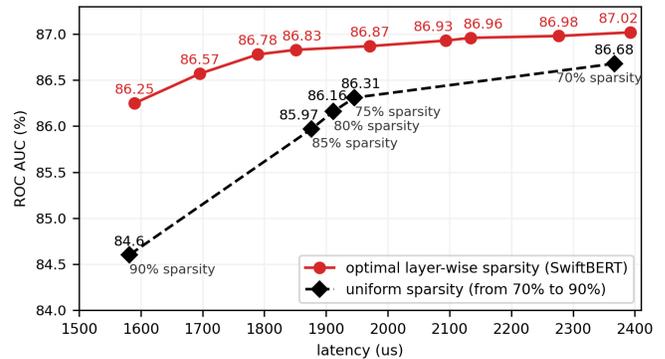
**Search algorithm details.** In our experiment, we explore  $N=500$  models in total. To set the population size  $P$  and sample size  $S$ , we consider 20, 50 and 100 as in [6, 27]. The results show there is no obvious difference after training to converge. We finally set the  $P$  and  $S$  to 50. The initial populations are randomly sampled from our search space with a relaxed latency constraint (i.e.,  $1.15\times$  of the target constraint). For the mutation controller, the embedding size is 104 (4 candidates in Attention and 100 candidates in FFN), and the hidden state size of the 2-layer LSTM is 100. We use the Adam optimizer with a learning rate of 0.001.

At each mutation step, the controller samples a model, and we calculate the latency-aware multi-objective reward for this mutation action. As introduced in Section 4.2, we predict the latency for the sampled model using our latency predictor and conduct layer-wise pruning to get the ROC AUC score. To reduce the search cost incurred during pruning, we prune on a subset of 500 mini-batches, which is randomly sampled from our full training set. The AUC is evaluated on a small validation set for the search algorithm. When the mutation process finishes, we collect all models that fulfill the latency constraint, and select the model with maximum validation AUC. In our experiment, the search process can finish within 24 hours on a 4-V100 GPU node.

**Final layer-wise pruning and evaluation.** For the final model setting selected by our search algorithm, we conduct layer-wise pruning using the full large-scale training set. Concretely, we run 3 epochs of pruning and 1 epoch of fine-tuning with a batch size of 8192 and initial learning rate of  $3e-5$ . The other settings and hyperparameters are kept the same as `nn_pruning`<sup>2</sup> [17]. We report the final ROC AUC on our test set.

## 5.2 SwiftPruner Under Various Latency Constraints

**Setup.** We now report our reinforced evolutionary pruning performance and compare with the state-of-the-art uniform sparsity method, `nn_pruning` [17] by Hugging Face. Since `nn_pruning` does not support latency-aware pruning, we set multiple sparsity ratios for comparison. Following the original paper, we apply hybrid-filled pruning for the best AUC and efficiency trade-off and use square block ( $32\times 32$ ) pruning on the attention layer and dimension pruning on the FFN layer. A fine-tuning is conducted after the pruning



**Figure 3: Compared to uniform sparsity pruning, SwiftPruner consistently outputs the best model settings under various latency constraints. Moreover, SwiftPruner is able to construct models for a wide range of latency constraints whereas uniform sparsity model latencies are fragmented and difficult to tune to a specific latency requirement.**

to re-fill the 0-valued neurons in the pruned model. For fair comparison, we run 3 epochs of pruning and 1 epoch of fine-tuning under the same random seed. To demonstrate the effectiveness of SwiftPruner, we set a wide range of latency constraints from 1600us to 2400us.

**Results and Analysis.** Fig. 3 gives a full comparison of models with different compression methods and latency constraints. Our results show that by searching the optimal layer-wise sparsity, SwiftPruner yields both significant latency and AUC improvement compared to the uniform sparsity baseline. More specifically, we observe the following: (i) layer-wise sparsity demonstrates the superiority in removing redundancy in tiny transformers compared to uniform sparsity. Under the same-level latency, SwiftPruner consistently outperforms uniform sparsity pruning. Specifically, SwiftBERT (1851.16us) achieves +0.86% higher AUC than the 85% uniform sparsity model (1876.44us). The advantages of layer-wise sparsity becomes more competitive under extremely low latency constraints. SwiftBERT (1590.00us) outperforms the 90% uniform sparsity model (1581.05us) by a significant +1.65% higher AUC. (ii) Our method demonstrates the effectiveness in pruning tiny transformer to various latency constraints, which is crucial in practical deployment. In Fig. 3, we can see that the final latency is not linearly reduced by the increase in uniform sparsity. In particular, it’s hard to tune a sparsity to obtain a model with around 1700us latency via `nn_pruning`. Meanwhile, SwiftPruner can prune a model while retaining the maximum AUC under all latency constraints.

In our online system, the latency of SwiftBERT is constrained to 1900us. Table 3 suggests that we can even achieve much lower latency than the 1900us constraint with a satisfactory ROC AUC. Remarkably, our searched layer-wise sparsity model (1851.16us) can significantly reduce the original BERT-Mini’s latency by 43.46% with an acceptable AUC loss of just 0.32%. The uniform pruned model, on the other hand, has a 1.18% AUC drop under the 1900us constraint, which demonstrates the efficiency of layer-wise sparsity in this extremely low-latency regime.

<sup>2</sup>The code of `nn_pruning` is at: [https://github.com/huggingface/nn\\_pruning](https://github.com/huggingface/nn_pruning)

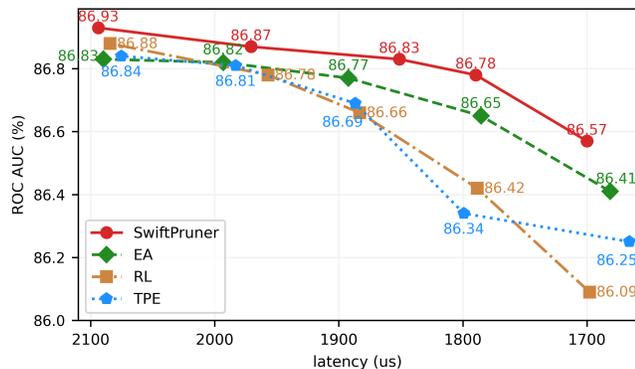


Figure 4: Comparison of search methods under different latency constraints,  $\alpha = -1$ . The X-axis is reversed.

### 5.3 Comparison of Search Methods

The previous section demonstrated that our searched layer-wise sparsity configuration consistently outperforms uniform sparsity. In this section, we compare SwiftPruner with other methods that are also applicable to search layer-wise sparse models.

**Comparison baselines.** We implement three other baselines, which are representative of the state-of-the-art in model architecture and hyper-parameter search.

- Tree Parzen Estimator (TPE). We apply the TPE algorithm [1] by formulating the problem as a hyper-parameter search where the hyper-parameters are the sparsity values for the attention heads and FFN of the 4 encoder layers.
- Reinforcement Learning (RL). Our implementation follows the well-known ENAS [26]. We design and train an RNN controller to predict the sparsity for each layer. The controller consists of 4 LSTM sub-networks, where each sub-network predicts the attention and FFN sparsity for a corresponding encoder layer.
- Aging Evolution (EA). Our implementation follows the original algorithm in [27]. EA is conducted with the same setting to SwiftPruner, except that the mutation action is made by randomly sampling a layer and its new sparsity setting.

For fair comparison, we use one training recipe and hyper-parameter setting for all methods. Specifically, all three baselines search for the same number of 500 models. For EA, the population size  $P$  and sample size  $S$  are the same with ours (i.e., set to 50). All search methods optimize the same reward function in Equation 2.

**Search results comparison.** We first compare the final AUC for the searched models. Fig. 4 reports comparisons between SwiftPruner and the other search methods under five latency constraints (1700us - 2100us). Compared to the three baselines, SwiftPruner consistently reaches the highest AUC under all latency constraints. Specifically, we improve the AUC by 0.1% (EA), 0.09% (TPE), and 0.05% (RL) under the maximum 2100us constraint. We also observe that the AUC improvements by SwiftPruner are even higher under lower latency constraints. Significantly, under the smallest 1700us constraint, we outperform EA, TPE, and RL by 0.16%, 0.32%, and 0.48% higher AUC, respectively.

**Search efficiency comparison.** For a deeper understanding of various search methods, we not only compare the AUC scores of the best searched models, but also compare the search efficiency. Since

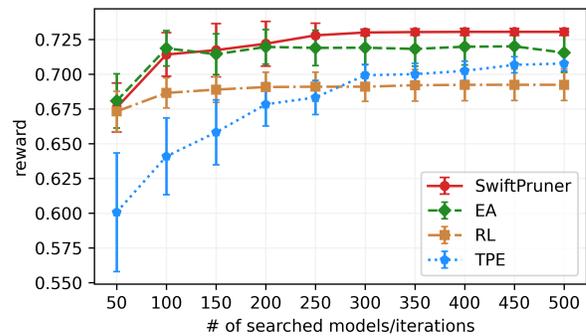


Figure 5: Search efficiency comparison under the same search iteration ( $\alpha = -1$ ).

$\alpha$	TPE	RL	EA	SwiftPruner
-0.3	86.71	86.65	86.75	<b>86.82</b>
-0.7	86.79	86.30	86.76	<b>86.84</b>
-1	86.69	86.66	86.77	<b>86.83</b>

Table 4: AUC (%) achieved by different latency-AUC tradeoff  $\alpha$ . We set the same 1900us constraint for all methods.

all the search methods optimize both latency and AUC, we compare the achieved reward (computed by Equation 2) with trade-off ratio  $\alpha = -1$ . Fig. 5 compares the search efficiency of SwiftPruner to EA, RL and TPE under the 1900us constraint. For each method, we keep track of the searched models for every 50 iterations. Since EA and SwiftPruner are evolution-based methods, we evaluate them by the reward mean and variance of models in the population (i.e., 50 models). For TPE and RL, we select the top 50 models with the best reward over time.

As shown in Fig. 5, SwiftPruner achieves better efficiency than EA, RL and TPE. Compared to TPE and RL, SwiftPruner consistently achieves higher rewards all the time. Compared to EA, SwiftPruner degrades a little over the first few iterations of reinforced mutation, but it quickly surpasses EA after 150 iterations.

**The impact of AUC-latency trade-off  $\alpha$ .**  $\alpha$  controls the AUC-latency trade-off during the search process. A smaller  $\alpha$  indicates a larger penalty on the latency. To study the impact of different  $\alpha$  on search methods, we set three  $\alpha$ : -0.3, -0.7 and -1. We run SwiftPruner and our baselines with different  $\alpha$  under the same 1900us constraint.

Table 4 summarizes the searched AUC under different  $\alpha$ . From the results, we observe that: (i) for each  $\alpha$ , SwiftPruner reaches the highest AUC; (ii) while TPE and RL AUC vary significantly across the three  $\alpha$  settings, SwiftPruner and EA have stable performance with little variance. This suggests that evolution-based algorithms are more robust to hyper-parameter selection.

In summary, we demonstrate that SwiftPruner outperforms other search methods with higher AUC models, better search efficiency and greater robustness to hyper-parameter selection.

### 5.4 Pruned Model Visualization

We now visualize the SwiftBERT structures and identify some heuristics for designing more efficient pruned BERT architectures. By doing so, we hope to motivate and inspire further study of layer-wise pruning and architecture design of transformer models.

Encoder Layers	BERT-Mini (AUC: 87.15%, 3274us)		nn_pruning (AUC:86.16%, 1911us)		SwiftBERT (AUC:86.57%, 1695us)		SwiftBERT (AUC: 86.83%, 1851us)		SwiftBERT (AUC: 86.92%, 2093us)	
	Head Count	FFN Size	Head Count	FFN Size	Head Count	FFN Size	Head Count	FFN Size	Head Count	FFN Size
No.1	4	1024	2	204	2	40	2	276	2	204
No.2	4	1024	2	204	4	102	4	256	4	328
No.3	4	1024	2	204	1	40	1	204	2	512
No.4	4	1024	2	204	1	30	1	133	1	614

**Figure 6: Visualization of pruned model structure. Insights for ad relevance: first two encoders prefer larger attention head count; FFN has more redundancy than attention.**

Metric	Cold ads	Cold+Hot ads
Defect Rate	-11.7%	-3.3%
RPM	+7.94%	+0.27%
CTR	+2.37%	+0.65%
QBR	-4.43%	-0.16%
Avg. Latency	-	+0.10%
99.9th % Latency	-	+1.50%

**Table 5: Online KPI impact of adding SwiftBERT 1900 to serve as a new cold ads semantic relevance model. The online test shows significant positive KPI’s in both cold ads segment and on all (cold+hot) ads with minimal latency impact.**

Fig. 6 shows the pruned structures of both nn\_pruning and SwiftPruner. We observe that deeper layers have more redundancy than lower layers in BERT-Mini, which is consistent with the observations in CNN models [20]. Compared to uniform sparsity, our layer-wise sparsity method prunes less in crucial layers and prunes more in less-crucial layers. For instance, while nn\_pruning prunes 2 heads in each encoder layer, our method automatically learns to keep all attention heads in the 2nd encoder layer and prunes 3 heads in the 3rd and 4th layers. As a result, SwiftBERT (1851us) has a comparable latency with nn\_pruning but a much higher AUC (+0.67%). Moreover, we observe that SwiftPruner learns to keep more attention heads and prune more FFN intermediate layers under tighter latency constraints. For example, SwiftBERT (1695us) retains the same attention structure as SwiftBERT (1851us), but keeps only 0.03%-0.1% of the original FFN intermediate dimensions.

## 5.5 Impact on Ad Production

In this section, we study the impact of our offline AUC gain by integrating the SwiftBERT model with latency constraint 1900us (SwiftBERT 1900) into the Microsoft Bing ad relevance system. To do so, we add SwiftBERT 1900 as a new query-ad semantic relevance model for cold ads whose semantic score is then used to inform decisions later in the ad stack as described in Section 3.

For the task of filtering defective cold ads, we find SwiftBERT 1900’s semantic score can lift the relevance system’s defective cold ad classifier performance by a significant +0.96% in ROC AUC.

To measure the other effects of adding SwiftBERT 1900 to the relevance system, we conduct an online A/B test and study the impact on both the cold ads and overall (i.e., hot+cold ads) segments. The KPI results from our online A/B test is outlined in Table 5 and a description of each metric is provided below.

**Defect Rate:** Ratio of defective ads which are shown to users. One major goal of the ad relevance system is to minimize this rate.

**Revenue per Mille (RPM):** This metric is the revenue per 1000 queries. Increasing this value while maintaining or improving quality of ads represents an improvement in the overall ad system.

**Click-through-rate (CTR):** This is the average number of clicks per ad impression. CTR tends to be correlated with the defect rate and positive trends signal an improvement in the quality of ads shown to users.

**Quick back rate (QBR):** This is the percent of clicks that result in users “quickly” returning to the search page, indicating they are unsatisfied with the clicked ads. QBR tends to correlate with the defect rate and reduced QBR suggests ad quality improvements.

**Avg. Latency:** Average latency of the system to process all ads per query. We would like our models to not impact the system latency significantly.

**99.9th% Latency:** Latency of the system to process all ads for the slowest 0.1-percentile of queries. We would like to keep this number small such that even the slowest queries with the largest ad loads do not time out.

In Table 5, we see a strict improvement in every KPI in both the cold and overall segments. In particular, we observe a great 11.7% cut in cold ad Defect Rate which translates to a 3.3% cut in overall ad Defect Rate. Furthermore, we see strong positive metrics in user behavior with a 2.37% increase in CTR and a substantial 4.43% cut in QBR for cold ads delivered. Finally, the online A/B test shows a greater revenue after adding SwiftBERT 1900 with a 7.94% RPM increase among cold ads and a 0.27% RPM increase among hot+cold ads. Since we add a new semantic model which runs at real time for all cold ads at online serving stage, the latency of the system increases, as expected. However, due to our efficient pruning methodology, the latency impact of running SwiftBERT 1900 for all cold ads is reasonably small with just a 1.50% increase in 99.9th% latency. This means even after adding SwiftBERT 1900, our system can continue to serve user queries with high ad load without timing out.

In summary, our SwiftBERT 1900 model delivers significant performance improvements in ads relevance, user behavior metrics, and revenue while adding minimal latency to the system.

## 6 CONCLUSION AND FUTURE WORKS

In this paper, we introduced SwiftPruner, a novel end-to-end framework that searches the best-performing layer-wise sparse BERT models under a desired latency constraint. SwiftPruner provides an efficient reinforced evolutionary search algorithm, which learns to conduct effective mutations rather than random ones. Extensive experiments demonstrate its effectiveness and efficiency compared to other search methods. Real-world online testing shows that our pruned model, SwiftBERT achieves a significant defective cold ads reduction and enables real-time, transformer-based online serving for cold ads on CPU at Microsoft Bing.

Finally, SwiftPruner’s fast online inference provides benefits beyond just the cold ads scenario. We have also successfully deployed SwiftPruner to serve relevance in dynamic ads [13, 38] where offline ad-side computation is infeasible because ad text is dynamically generated online. Additionally, we plan to apply SwiftPruner to search relevance [43] and vertical ads in the future.

## REFERENCES

- [1] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. *Advances in Neural Information Processing Systems* (2011), 2546–2554.
- [2] Microsoft blog. 2021. Microsoft open sources breakthrough optimizations for transformer inference on GPU and CPU. <https://cloudblogs.microsoft.com/opensource/2020/01/21/microsoft-onnx-opensource-optimizations-transformer-inference-gpu-cpu/>
- [3] Microsoft blog. 2021. Optimizing BERT model for Intel CPU Cores using ONNX runtime default execution provider. <https://cloudblogs.microsoft.com/opensource/2021/03/01/optimizing-bert-model-for-intel-cpu-cores-using-onnx-runtime-default-execution-provider/>
- [4] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. 2020. Once-for-All: Train One Network and Specialize it for Efficient Deployment. In *ICLR*.
- [5] Wentao Chen, Hailong Qiu, Jian Zhuang, Chutong Zhang, Yu Hu, Qing Lu, Tianchen Wang, Yiyu Shi, Meiping Huang, and Xiaowe Xu. 2021. Quantization of Deep Neural Networks for Accurate Edge Computing. arXiv:2104.12046 [cs.CV]
- [6] Yukang Chen, Gaofeng Meng, Qian Zhang, Shiming Xiang, Chang Huang, Lisen Mu, and Xinggang Wang. 2019. Renas: Reinforced evolutionary neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4787–4796.
- [7] Yiren Chen, Yaming Yang, Hong Sun, Yujing Wang, Yu Xu, Wei Shen, Rong Zhou, Yunhai Tong, Jing Bai, and Ruofei Zhang. 2020. AutoADR: Automatic Model Design for Ad Relevance. In *CIKM*.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.
- [9] Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. 2021. Rigging the Lottery: Making All Tickets Winners.
- [10] Mitchell A. Gordon, Kevin Duh, and Nicholas Andrews. 2020. Compressing BERT: Studying the Effects of Weight Pruning on Transfer Learning. arXiv:2002.08307 [cs.CL]
- [11] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. 2020. Single Path One-Shot Neural Architecture Search with Uniform Sampling.
- [12] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. 2018. AMC: AutoML for Model Compression and Acceleration on Mobile Devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- [13] J. Weston Hughes, Keng-hao Chang, and Ruofei Zhang. 2019. Generating Better Search Engine Text Advertisements with Deep Reinforcement Learning. In *KDD*.
- [14] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. TinyBERT: Distilling BERT for Natural Language Understanding.
- [15] Sehoon Kim, Amir Gholami, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. 2021. I-bert: Integer-only bert quantization. *arXiv preprint arXiv:2101.01321* (2021).
- [16] Young Jin Kim and Hany Hassan Awadalla. 2020. Fastformers: Highly efficient transformer models for natural language understanding. *arXiv preprint arXiv:2010.13382* (2020).
- [17] Francois Lagunas, Ella Charlaix, Victor Sanh, and Alexander M. Rush. 2021. Block Pruning For Faster Transformers. In *EMNLP*.
- [18] Jaeho Lee, Sejun Park, Sangwoo Mo, Sungsoo Ahn, and Jinwoo Shin. 2021. Layer-adaptive Sparsity for the Magnitude-based Pruning. In *International Conference on Learning Representations*.
- [19] Breiman Leo. 2001. Random Forests. In *Machine Learning*, 5–32.
- [20] Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Tim Kwang-Ting Cheng, and Jian Sun. 2019. MetaPruning: Meta Learning for Automatic Neural Network Channel Pruning. In *ICCV*.
- [21] Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. TwinBERT: Distilling Knowledge to Twin-Structured Compressed BERT Models for Large-Scale Retrieval. In *CIKM*.
- [22] Jiachen Mao, Huanrui Yang, Ang Li, Hai Li, and Yiran Chen. 2021. TPrune: Efficient Transformer Pruning for Mobile Devices. *ACM Trans. Cyber-Phys. Syst.* 5, 3, Article 26 (apr 2021), 22 pages. <https://doi.org/10.1145/3446640>
- [23] JS McCarley, Rishav Chakravarti, and Avirup Sil. 2019. Structured pruning of a bert-based question answering model. *arXiv preprint arXiv:1910.06360* (2019).
- [24] Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one?. In *NeurIPS*.
- [25] Bert Moons, Parham Noorzad, Andrii Skliar, Giovanni Mariani, Dushyant Mehta, Chris Lott, and Tijmen Blankevoort. 2021. Distilling Optimal Neural Networks: Rapid Search in Diverse Spaces. In *ICCV*.
- [26] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. 2018. Efficient Neural Architecture Search via Parameter Sharing.
- [27] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized Evolution for Image Classifier Architecture Search. In *AAAI*.
- [28] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv:1910.01108 [cs.CL]
- [29] Victor Sanh, Thomas Wolf, and Alexander M Rush. 2020. Movement pruning: Adaptive sparsity by fine-tuning. In *NeurIPS*.
- [30] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical Bayesian Optimization of Machine Learning Algorithms. In *Advances in Neural Information Processing Systems*.
- [31] Wei Sun, Aojun Zhou, Sander Stuijk, Rob Wijnhoven, Andrew O Nelson, Henk Corporaal, et al. 2021. DominoSearch: Find layer-wise fine-grained N: M sparse schemes from dense neural networks. *Advances in Neural Information Processing Systems* 34 (2021).
- [32] Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. MobileBERT: a Compact Task-Agnostic BERT for Resource-Limited Devices. In *ACL*.
- [33] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-Read Students Learn Better: The Impact of Student Initialization on Knowledge Distillation. *CoRR abs/1908.08962*. arXiv:1908.08962
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, and Lukasz Kaiser. 2017. Attention is All You Need. In *NIPS*.
- [35] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*.
- [36] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2020. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems.
- [37] Dilin Wang, Meng Li, Chengyue Gong, and Vikas Chandra. 2021. AttentiveNAS: Improving Neural Architecture Search via Attentive Sampling. In *Conference on Computer Vision and Pattern Recognition*.
- [38] Xiting Wang, Xinwei Gu, Jie Cao, Zhihua Zhao, Yulan Yan, Bhuvan Middha, and Xing Xie. 2021. Reinforcing Pretrained Models for Generating Attractive Text Advertisements. In *KDD*.
- [39] Ronald J. Williams. 1992. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. (1992).
- [40] Mengzhou Xia, Zexuan Zhong, and Danqi Chen. 2022. Structured Pruning Learns Compact and Accurate Models. In *ACL*.
- [41] Jin Xu, Xu Tan, Renqian Luo, Kaitao Song, Jian Li, Tao Qin, and Tie-Yan Liu. 2021. NAS-BERT: Task-Agnostic and Adaptive-Size BERT Compression with Neural Architecture Search. In *kdd*.
- [42] Zhewei Yao, Linjian Ma, Sheng Shen, Kurt Keutzer, and Michael W. Mahoney. 2021. MLPPruning: A Multilevel Structured Pruning Framework for Transformer-based Models. arXiv:2105.14636 [cs.CL]
- [43] Dawei Yin, Yuening Hu, Jiliang Tang, Tim Daly, Mianwei Zhou, Hua Ouyang, Jianhui Chen, Changsung Kang, Hongbo Deng, Chikashi Nobata, Jean-Marc Langlois, and Yi Chang. 2016. Ranking Relevance in Yahoo Search (*KDD '16*). Association for Computing Machinery, 323–332.
- [44] Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8bert: Quantized 8bit bert. *arXiv preprint arXiv:1910.06188* (2019).
- [45] Barret Zoph and Quoc V. Le. 2017. Neural Architecture Search with Reinforcement Learning.