Auto Graph Encoder-Decoder for Model Compression and Network Acceleration

Sixing Yu
Department of Computer Science
Iowa State University
yusx@iastate.edu

Arya Mazaheri
Department of Computer Science
Technische Universität Darmstadt
mazaheri@cs.tu-darmstadt.de

Ali Jannesari
Department of Computer Science
Iowa State University
jannesar@iastate.edu

***Abstract*—Model compression aims to deploy deep neural networks (DNN) to mobile devices with limited computing power and storage resource. However, most of the existing model compression methods rely on manually defined rules, which requires domain expertise. In this paper, we propose an Auto Graph encoder-decoder Model Compression (AGMC) method combined with graph neural networks (GNN) and reinforcement learning (RL) to find the best compression policy. We model the target DNN as a graph and use GNN to learn the embeddings of the DNN automatically. In our experiments, we first compared our method with rule-based DNN embedding methods to show the graph auto encoder-decoder's effectiveness. Our learning-based DNN embedding achieved better performance and a higher compression ratio with fewer search steps. Moreover, we evaluated the AGMC on CIFAR-10 and ILSVRC-2012 datasets and compared handcrafted and learning-based model compression approaches. Our method outperformed handcrafted and learning-based methods on ResNet-56 with 3.6% and 1.8% higher accuracy, respectively. Furthermore, we achieved a higher compression ratio than state-of-the-art methods on MobileNet-V2 with just 0.93% accuracy loss.**

***Index Terms*—Model Compression, Network Pruning, RL**

I. Introduction

With the increasing demand to deploy DNN models on edge devices (e.g., mobile phones, robots, self-driving cars, etc.), which usually have limited storage and computing power, model compression techniques become an essential part of efficient DNN deployment. Network pruning [6], [7], [25], factorization [29], [35], knowledge distillation [14], [26], [27], and parameter quantization [6], [16], [38] are among the most well-known compression techniques. However, these methods heavily rely on manually defined rules by experts, which requires an extensive amount of time and might not necessarily lead to a fully compressed model. Recently, automatic model compression [11], [24], [38] is gaining momentum. Wang et al. [38] proposed a Bayesian automatic model compression method trained in a one-shot manner to find reasonable quantization policies. He et al. [11] proposed an automatic model compression method based on reinforcement learning (RL). However, when representing the DNN, they manually defined the embedding vector for each hidden layer, which ignores the rich structural information between the neural network's hidden layers.

To overcome the shortcomings of the existing methods, we propose a graph-based auto encoder-decoder model compression method AGCM combines GNN [18], [40], [42] and reinforcement learning [21], [32], [34] to learn the compression strategy of DNNs without expert knowledge. GNN is a powerful technique to learn the graph's structural information. Thus, we use GNN to model the DNN as a graph and learn the hidden layer's representation. We perform model compression by predicting each hidden layer of DNN's pruning ratio and evaluate the compressed model's performance using RL.

Neural networks can be easily represented as computational graphs, which contain wealthy structural information. However, it is unrealistic to directly transform a DNN into a computational graph because a DNN may involve billions of calculations [8], and the resulting computation graph could be huge. Thus, we apply the idea of Motif-Graph [23], as shown in Figure 1, to construct a hierarchical graph for DNNs.

We modeled a DNN as a hierarchical computation graph and introduced a graph encoder based on Differentiable Graph Pooling (DIFFPOOL) GNN [41] to learn the embedding. DIFFPOOL GNN focuses on graph classification and can learn the embeddings of graphs very well. Then, we introduced a decoder to decode the representation of DNN and learn the features of each hidden layer within the DNN. We use each hidden layer's features as the state vector of the environment, let the RL agent look for each hidden layer's compression ratios, and generate the corresponding candidate model. The RL reward function is based on the performance of the candidate model.

In essence, this paper makes the following contributions:

- To the best of our knowledge, this work is the first work to model the DNN as a hierarchical graph and use GNN to embed the DNN.
- We introduce a learning-based autoencoder to perform embedding the DNNs without human effort.
- Under the AGMC framework, we introduced an innovative scheme to perform structured pruning and unstructured pruning together in a single-shot manner.

II. Related Work

Many previous works focus on model compression, such as knowledge distillation [14], [26], [27], parameter quantization [6], [16], [38], factorization [29], [35], and network pruning [6], [7], [25]. The network pruning techniques are among the

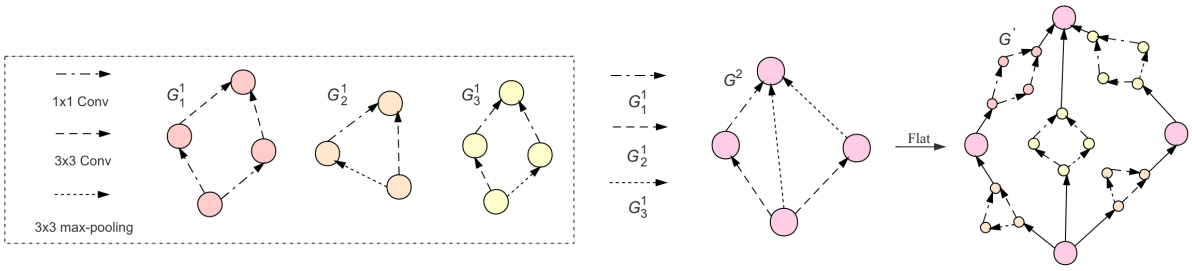

Fig. 1. A two-layer hierarchical computation graph. The nodes on the graph denotes the feature maps of input data, and the edges corresponding to primitive operations. In layer 1 (Left),we have three edge types associated with three primitive operations. And each graph in layer 1 corresponding to a compound operation composed of the primitive operations. The computation graphs $\{G_1^1$,G_2^1,$G_3^1\}$are becomes the primitive operation in the layer 2. And we can flat the G^2 by replacing the edges with computation graphs in layer 1.

most widely used methods and can dramatically shrink the model size. In this paper, we compress the DNN by network pruning.

Network pruning includes two major categories: structured pruning and unstructured pruning, and aims to evaluate the parameters' importance in the DNN and remove the parameters with a lower rank. The unstructured pruning [4], [43] prunes individual unimportant elements in weight tensors and can achieve a high compression ratio. Although the unstructured pruning can accelerate DNN with specialized hardware [5], [17], it fails in parallel implementations like GPUs. The structured pruning [13], [44] overcomes the limitation of unstructured pruning. For example, the filter pruning [10], [12] on CNNs removes the redundant channels from feature maps. As neural networks are typically over-parameterized, network pruning has achieved outstanding results and can remove even 90% of the parameters in specific models [1]. However, conventional network pruning methods rely primarily on handcraft and rule-based policies, which require human efforts and domain expertise and might not necessarily lead to a fully compressed model.

Recently, many learning-based network pruning methods [11], [24], [39] have been proposed. Liu et al. [24] proposed an ADMM-based [2] structured weight pruning method and an innovative additional purification step for further weight reduction. He et al. [11] proposed AutoML for network pruning, which leverage reinforcement learning to predict each hidden layer's compression policies. However, they manually defined DNN's embedding and ignored the neural network's enormous structural information. In our work, we introduced an AutoGraph encoder-decoder, which automatically learns the embeddings from the DNN's rich topology information.

III. Methodology

A. Deep Neural Network to Graph

We noticed that neural networks in deep learning frameworks, such as TensorFlow and PyTorch, are represented as computational graphs, containing rich structural information. However, it is unrealistic to transform a DNN into a computational graph without simplification, as a DNN may involve

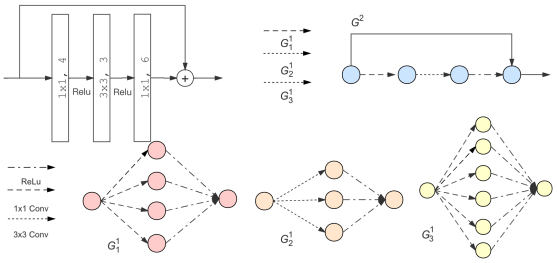

Fig. 2. An example of two-layer hierarchical computation graph for a simple ResNet [8] Block. In layer 1, for example, the G_1^1 denotes the 1x1 convolution with 4 channels and activate the feature maps with ReLu. And then the G_1^1 becomes a primitive operation in layer 2, and corresponding to an edge type in G^2. The G^2 denotes the computations of the ResNet Block, and the nodes on the G^2 are the feature maps on the ResNet Block.

billions of calculations [8]. Thus, we employ the idea of the Motif graph [23] to build a hierarchical computational graph for DNNs.

We model our DNN as an $l-$layer single-source and single-sink hierarchical computation graph $G^l = (V^l, E^l, \mathcal{G}^{l-1})$, where $\mathcal{G}^{l-1} = \{G_0^{l-1}, G_1^{l-1}, ...\}$ is the primitive operation set at layer l. For a hierarchical computation graph in layer t, each node of the graph corresponds to a hidden state of the data, and each directed edge with a specific edge type associates with a primitive operation at layer t. The primitive operations at layer t is compound operations composed of primitive operations at layer $t-1$. Figure 1 shows the idea of a 2-layer hierarchical computation graph. For example, in layer 1, we choose three primitive operations $\mathcal{G}^0 = \{1\times1$ conv, 3×3 conv, 3×3 max-pooling$\}$ corresponding to three edge types, and the computation graph G_1^1 denotes a compound operation composed of the primitive operations in $\mathcal{G}^0$, each directed edge in G_1^1 corresponding to a primitive operation in $\mathcal{G}^0$. The operation G_1^1 denotes is

$$y = assemble(conv3(conv1(x)), conv3(conv1(x))) \quad (1)$$

In layer 2, $\mathcal{G}^1 = \{G_1^1,\ G_2^1,\ G_3^1\}$, each edge type in G^2 corresponds to a computation in $\mathcal{G}^1$, and we can flatten the

G^2 to $G^{'}$ by replacing the edges to their corresponding lower-level computation graph. Figure 2 shows an example for constructing a 2-layer hierarchical computation graph for a ResNet [8] block.

The hierarchical computation graph's size depends on the primitive operation we choose in G^1. As long as we reasonably choose primitive operation for G^1, we can construct a reasonable computation graph $G^{'}$ for DNN. In our experiments, we choose the commonly used operations in machine learning as primitive operations (e.g., convolution, pooling, etc.).

B. Model Compression with GNN and RL

Figure 3 shows an overview of the AGMC. We introduced an auto model compression method, combined GNN [18] with RL [21], to automatically find the best compression strategy for each hidden layer of DNN. We model the DNN as a graph and learn the DNN's representation g through GCN based encoder. And decode the representation g to the features of each hidden layer $s_i \in S, i = 1, 2, ..T$, where T is the number of hidden layer. Then we take the S as the environment state vectors to the RL agent and search for the hidden layer's compression policy $a_i \in A, i = 1, 2, ..T$. The compressed DNN's performance is feedback to the RL agent as a reward to find the best compression policy.

C. Auto graph encoder-decoder

We introduced an auto graph encoder-decoder to learn the features of DNN's hidden layers automatically. The graph encoder aims to embed the graph. We use GCN to embed the nodes on the graph and then pool it to learn the representation g. In this paper, we introduced two graph encoders for different DNNs by applying different pooling strategies, Mean Pool and Differentiable Pool. Moreover, we introduced the LSTM [36] based decoder, which takes the previous layer's embedding and compression policy from the RL agent as input to predict DNN's hidden layers' embeddings.

1) Graph Mean Pool: Mean Pool GCN learns graph representation $g \in \mathbb{R}^{1\times d}$ by averaging node features across the node dimension:

$$X^{(l)} = GCN_l(X^{(l-1)}) \in \mathbb{R}^{N\times d} \tag{2}$$

$$g = \frac{1}{N}\sum_{n=1}^{N} {x^l}_n \tag{3}$$

Where X^l is the node embedding matrix at layer l, ${x^l}_i$ is the embedding of node n at layer l, N is the number of nodes on the graph, and d is the embedding size. Mean Pool GCN first learns nodes embedding $X^{(l)}$ through a l-layer GCN, and then mean pool the graph, and it works well when graph size is small. However, when the hierarchical computation graph's size is large, Mean Pool GCN is inherently flat and unable to infer and hierarchically aggregate the information [41].

2) Differentiable Pool: Differentiable Pool (DIFFPOOL) GNN focuses on graph classification and is a state-of-the-art [41] method to learn graph representation. DIFFPOOL GNN learns a differentiable soft assignment M at each layer of a deep GNN, mapping nodes to a set of clusters based on their learned embeddings, and generate a smaller graph for the next GNN layer based on the clusters. And each node in the graph corresponding to a cluster of nodes at the previous layer. In this paper, we introduced the DIFFPOOL GCN encoder. We take the input adjacency matrix $A^{(l)}$ and denote the node embedding matrix at layer l as Z^l. The DIFFPOOL coarsens the input graph, generating a new coarsened adjacency matrix $A^{(l+1)}$ and a new matrix of embeddings $X^{(l+1)}$ for each of the nodes/clusters in this coarsened graph:

$$Z^l = GCN^l_{embed}(A^l, X^l) \in \mathbb{R}^{n_l\times d} \tag{4}$$

$$M^l = softmax(GCN^l_{pool}(A^l, X^l)) \in \mathbb{R}^{n_l\times n_{l+1}} \tag{5}$$

$$X^{(l+1)} = M^{(l)T} Z^{(l)} \in \mathbb{R}^{n_{l+1}\ \times\ d} \tag{6}$$

$$A^{(l+1)} = M^{(l)T} A^{(l)} M^{(l)} \ \in \mathbb{R}^{n_{l+1}\ \times\ n_{l+1}} \tag{7}$$

where the d is the embedding size and the n_l denotes the nodes at layer l and clusters at layer $l-1$. We set the final layer with 1 cluster and we can get the final out put $Z \in \mathbb{R}^{1\ \times\ d}$ as the representation g of the graph.

3) Decoder: The decoder aims to learn the embedding of each DNN's hidden layer, and we take the hidden layer's features as the state vectors for RL agent. Because the state vectors in the RL environment are determined by the previous state and the action, the decoder takes the previous layer's feature vector and compression policy as input:

$$s_t = Decoder_{lstm}\left(s_{t-1}, a_{t-1}\right) \tag{8}$$

For the $t-$th hidden layer, we use the feature s_{t-1} of the previous hidden layer and the compression policy a_{t-1} (action selected by the RL agent) to calculate the environment states.

D. DDPG Agent

1) The State an Action Space: The traditional RL state space is determined by the environment, which is typically fixed. The RL agent finds the best strategy by learning the rewards given by the environment. In AGMC, RL state space is determined by the auto graph encoder-decoder, which means that our environment is also learnable. We use the performance of the compressed model as RL agent's feedback. We choose a continuous action space $a \in (0, 1]$ to determine each hidden layer's compression policy.

2) Deep Deterministic Policy Gradient: To find the best strategy among continuous action space, we adopt DDPG [21], [32]. DDPG agent receives the state s from the environment and predicts the next action by the feedback of rewards and punishments. Similar to AMC [11], we also use truncated normal distribution for the exploration noise process:

$$\mu'\left(st\right) \ \sim\ TN\left(\mu\left(st \mid \theta^{\mu}_t\right), \sigma^2, 0, 1\right) \tag{9}$$

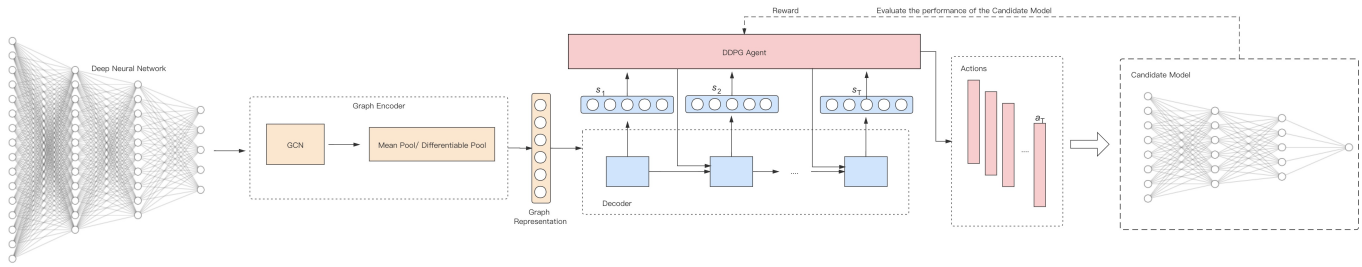

Fig. 3. Overview of the AGMC. The AGMC first embedding the DNN through a GCN based graph encoder. In the graph encoder, we introduced two pooling strategies for different DNNs, and learn the representation of the DNN. The decoder decodes the DNN representation and learn the layer embeddings. And then, the DDPG agent takes the layer embeddings as the environment's state vectors to predict the compression ratio for each hidden layer. Finally, after compressed all the hidden layers, DDPG take the performance of the compressed candidate model as the reward to update the AGMC.

And we update the agent with the loss function:

$$Loss = \frac{1}{N}\sum_i \left(y_i - Q\left(s_i, a_i \middle| \theta^Q\right)\right)^2 \tag{10}$$

$$y_i = r_i - b + \gamma Q\left(s_{i+1}, \mu\left(s_{i+1}\right) \middle| \theta^Q\right) \tag{11}$$

3) Reward: We use the performance of the candidate model as the reward. Specifically, we evaluate the candidate model by the error and the FLOPs of the candidate model. In different search protocols, DDPG apply different reward function [11]. For resource-constrained compression, we take the reward:

$$R_{err} = -Error \tag{12}$$

For accuracy-guaranteed compression, we use the reward:

$$R_{FLOPs} = -Error \times \log(FLOPs) \tag{13}$$

$$R_{Param} = -Error \times \log(\#Param) \tag{14}$$

4) Algorithm for the desired model reduction: The reward we use offers small or no incentive for model size and FLOPs reduction. Without constraint, the RL agent tends to search for a tiny compression ratio for each layer. To get a desired model size reduction, we apply Algorithm 1 to constraint the action a. According to the original scale, the AGMC achieves the desired model size or FLOPs reduction by re-scaling all the action.

Algorithm 1: Re-scaling the actions for the desired model size-reduction

Input: The actions $a = \{a_0, ..., a_T\}$, the upper bound of actions a_{max}, the model size (#FLOPs/#Parameters etc.) of each hidden layer $W = \{W_0, ..., W_T\}$, and the desired model size reduction d

Output: The actions a' after re-scaling

$W_{all} = \sum_t W_t$
$W_{reduced} = \sum_t W_t a_t$
if $W_{reduced} < d$ **then**
  $d_{rest} = d - W_{reduced}$
  for $i = 1, 2, ..., T$ **do**
    $a_i += (d_{rest} * (a_i / \sum_t a_t)) / W_i$
    $a'_i = min(a_{max}, a_i)$
return a'

IV. Experiment

In this section, we first compare the AGMC with AMC [11], which manually defines DNN's layer embeddings, and random search, which do not have layer embeddings. This comparison shows the effectiveness of our embedding technique. Then, we evaluate AGMC on CIFAR-10 [19] and ILSVRC-2012 [28] dataset with popular CNN models(VGG-16 [33], MobileNet [15], MobileNet-V2 [31], ResNet-20/50/56 [8]). We perform the structured channel pruning and fine-grained pruning for FLOPs-constrained compression and accuracy-guaranteed compression on the CIFAR-10 [19] dataset by predicting the pruning ratio of hidden layers. For FLOPs-constrained compression, we focus on the channel pruning (filter pruning) on the convolutional layers, which are the most computationally intensive. For accuracy-guaranteed compression, we apply fine-grained pruning to prune individual unimportant elements in weight tensors. We further perform structured pruning and unstructured pruning together in a single-shot manner on ILSVRC-2012 [28] dataset.

We construct 2-layer hierarchical computation graphs for DNNs. The primitive operations we choose in level 1 are commonly used operations in machine learning $\mathcal{G}^0 = \{conv1/3/7, Relu, BatchNorm, (Max/Average)Pooling2/3, Padding, Splitting\}$. Before we embedding the DNN, we flate the hierarchical computation graph as shown in Figure1. For ResNet-20/50/56, we apply Mean-Pool GCN as a graph encoder. For VGG-16 and MobileNet-V2, we apply Differentiable Pool GCN as a graph encoder. In our DDPG agent, our actor-network μ and critic network Q have two hidden layers, each with 300 units. The μ's output layer applies the sigmoid function to bound the actions within $(0, 1)$. We use $\tau = 0.01$ for the soft target updates. In the first 100 episodes, our agent searches with random action. Then exploits 300 episodes with exponentially decayed noise and trains the network with 64 batch size and 2000 as replay buffer size. In the CIFAR-10 dataset, we sample $15K$ images from the training set to fast fine-tuning our candidate model and $5K$ images from the test

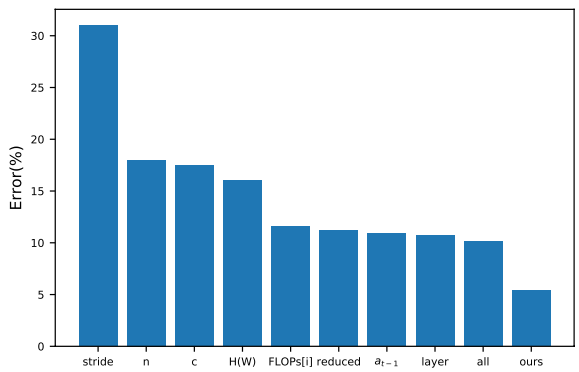

Fig. 4. Effectiveness of Layer Embeddings, evaluated with spatial decomposition for ResNet-20 2×. With the only Stride as the layer embeddings, we get an error of 31%; it is difficult to distinguish different layers. Combined the Stride and number of filter n, the performance is better with an error of 18%. And combining all the 11 features defined in AMC [11], we get a 10.2% error. By comparison, our learning-based layer embeddings with an error of 5.38% outperform the manually defined layer embeddings.

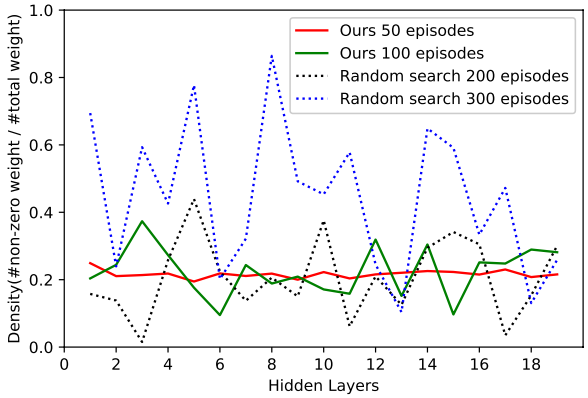

Fig. 5. Comparison with Random search on ResNet-20. Random search 200 and 300 episodes, the DDPG agent gets the compressed models with 71% and 88.41% accuracy. The AGMC search for 50 episodes and 100 episodes with an accuracy of 93.8% and 94.6%. Moreover, our learning-based achieved a higher compression ratio with fewer episodes. (Accuracy computes on the part of CIFAR-10 trainset with $5K$ images)

set to calculate reward. In ILSVRC-2012, we split $20K$ images from the test-set to fast fine-tuning our candidate model and $5K$ images to calculate reward.

A. Effectiveness of DNN Embeddings

Layer embeddings are essential for the DDPG agent to find a compression policy. However, many auto model compression methods [11], [24] relies on manually defined rules for DNN embedding. We introduced an Auto Graph Encoder-Decoder to embed the DNN without human efforts and domain expertise.

We compare AGMC with AMC [11] to show the effectiveness of our embeddings for the DNN. In AMC, for each layer t, they manually defined 11 features that characterize the state s_t:

$$(t, n, c, h, w, stride, k, FLOPs(t), reduced, rest, a_{t-1})$$

where t is the layer id, the dimension of the kernel is $n \times c \times k \times k$, and the input is $c \times h \times w$. $FLOPs(t)$ is the FLOPs of layer t. Reduced is the total number of reduced FLOPs in previous layers. Rest is the number of remaining FLOPs in the following layers. The above layer embedding rules may miss important information, such as the number of parameters in each hidden layer, which are only applicable to convolutional layers. In AGMC, we learn the layer embeddings from graph encoder-decoder, which does not require expert knowledge and applicable for all kinds of hidden layers.

Figure 4 shows the spatial decomposition evaluation for AMC's layer embeddings [9] under ResNet-20 [8] 2×. Straightforwardly, combining all the 11 features, the performance is better than combining part of them. However, even combining all the features, our learning-based embeddings are outperformed the manually defined embeddings.

Moreover, we compare the AGMC with a random search without layer embeddings. Using the R_{err} as a reward, we apply the DDPG reinforcement learning agent to a random search compression policy for ResNet-20 and ResNet-56 on CIFAR-10 and compress the DNNs by pruning the filters for convolutional layers. Moreover, we evaluate the candidate model on a subset of the training set ($5K$ images). The DDPG agent random searches 300 episodes and saves the best performance candidate model and fine-tuning the candidate model. As shown in Figure 5 for ResNet-20, compared to the random search, the AGMC get a better result. We find the candidate model with fewer episodes and higher accuracy and more massive FLOPs reduction. Figure 6 shows the results in ResNet-56. The ResNet-56 is much deeper and more challenging. With fewer searching episodes, AGMC gets an accuracy of 95.64% with 60% FLOPs reduction, which outperforms the random search by a large margin. Furthermore, the random search without layer embedding is pruning irregularly on ResNet-56. In contrast, The Model can learn that $3 \times 3Conv$ has more redundancy and prunes more on them.

We further evaluate the effectiveness of our embeddings by clustering the nodes on the graph. In the hierarchical computation graph, each node corresponding to a feature map in DNN. Neurons in the same hidden layer often have similar structural information. We label 20% nodes on the graph and perform semi-supervised classification. The graph encoder can successfully classify nodes according to their layers with a 98% accuracy.

B. CIFAR-10

We conducted FLOPs-Constrained compression and accuracy-guaranteed compression and analyzed the effectiveness of AGMC on CIFAR-10 [19].

Channel Pruning on CONV Layers We conducted channel pruning on convolutional Layers for FLOPs-constrained compression. Moreover, we compared our approach with three handcraft empirical policies [13], [20]: uniform, shallow,

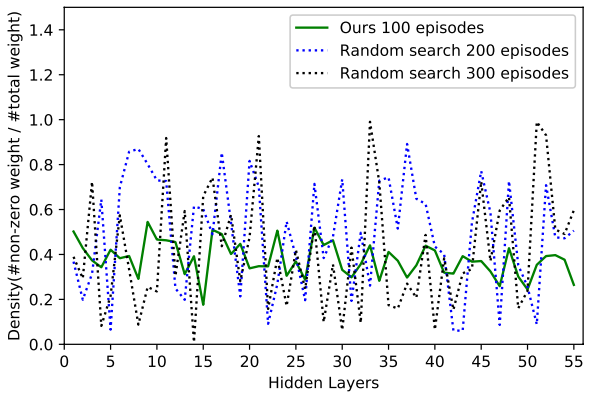

Fig. 6. Comparison with random search on ResNet-56. Random search 200 and 300 episodes, the DDPG agent gets the compressed models with 87.43% and 90.93% accuracy. The AGMC search for 100 episodes with an accuracy of 95.64%. (Accuracy computes on the part of CIFAR-10 trainset with $5K$ images)

Model	Policy	Ratio	Test $Acc\%$
ResNet-20	deep	50% FLOPs	79.6
	shallow		83.2
	uniform		84
	AMC (R_{Err})		86.4
	AGMC (R_{Err})		**88.42**
ResNet-56	uniform	50% FLOPs	87.5
	deep		88.4
	AMC (R_{Err})		90.2
	AGMC(R_{Err})		**92.0**
ResNet-56	Random search	$> 50\%$ Params	90.93
	AGMC(R_{Param})		**95.64**
ResNet-50	AMC (R_{Param})	$> 60\%$ Params	93.64
	AGMC(R_{Param})		**94.62**

TABLE I
COMPARISON RESULTS ON RESNET-20/50/56 USING THE CIFAR-10 DATASET. WE PERFORM FLOPS-CONSTRAINED COMPRESSION ON RESNET-20/56 WITH CHANNEL PRUNING AND USE THE R_{Err} AS A REWARD, WHILE PERFORM ACCURACY GUARANTEED COMPRESSION WITH FINE-GRAINED PRUNING BY USING R_{Param} AS A REWARD.

and deep. The uniform policy sets the compression ratio uniformly, and the shallow and deep policies aggressively prune shallow and deep layers, respectively. We further compared with an AutoML model compression method AMC [11].

As shown in Table I, we prune 50% FLOPs for ResNet-20 and ResNet-56 by using R_{err} to find the pruning policy. After searching, we fine-tune the best performance candidate model on the training dataset. Our method outperforms the handcraft empirical policies by a large margin with 88.42% accuracy on ResNet20 and 92.00% on ResNet-56. These handcraft empirical policies heavily rely on experts' manually defined rules and might not lead to an optimal compression policy. And compare to the learning-based method AMC [11], which uses manually defined layer embeddings, AGMC with 2.02% higher accuracy on ResNet20 and 1.8% on ResNet-56 than the AMC.

Model	Policy	FLOPs	Params	$\Delta Acc.\%$
VGG-16	FP [20]	20%	–	-14.6
	RNP [22]		–	-3.58
	SPP [37]		–	-2.3
	AMC [11]		–	**-1.4**
	AGMC		70%	-1.7
MobileNet	uniform [3]	56%	–	−2.5
	uniform [3]	41%	–	-3.7
	AMC [11]	40%	–	**-1.7**
	AGMC	40%	70%	-2.42
MobileNetV2	uniform [30]	70%	–	-2
	AMC [11]		–	-1
	AGMC		70%	**-0.93**

TABLE II
COMPARISON RESULTS ON VGG-16, MOBILENET AND MOBILENET-V2 USING THE ILSVRC-2012 DATASET. THE COLUMN FLOPS DENOTES THE RATIO BETWEEN THE FLOPS OF THE COMPRESSED MODEL AND THE ORIGINAL MODEL ON CONVOLUTION LAYERS. AND WE DO NOT CONSIDER THE DENSE LAYERS' FLOPS HERE. THE COLUMN PARAMS DENOTES THE RATIO BETWEEN THE PARAMS OF THE COMPRESSED AND THE ORIGINAL MODEL ON DENSE LAYERS.

Fine-Grained Pruning We aim to conduct accuracy-guaranteed compression with fine-grained pruning. We do not set the desired model size reduction with accuracy-guaranteed compression and let the DDPG agent find the compression policy with the best accuracy freely. Fine-grained pruning prunes individual unimportant elements in weight tensors, which can achieve a higher compression ratio and be accelerated with specialized hardware [5], [17].

Using the R_{Param} as a reward, we compress the ResNet-50 and ResNet-56 without fine-tuning, and compare the AGMC with random search and AMC [11]. Table I shows the results, with 50% parameter reduction on ResNet-56, AGMC outperforms random search with 4.71% higher accuracy. And with 60% parameter reduction on ResNet-50, AGMC outperforms AMC with 0.98% higher accuracy

C. ILSVRC-2012

This sub-section evaluates the AGMC on the ILSVRC-2012 [28] dataset, which is much more challenging than the CIFAR-10 [19]. We compared the AGMC with state-of-art handcraft methods: SPP [37], FP [20] and RNP [22], and auto model compression method AMC [11]. The SPP prunes DNNs by analyzing each layer and measures the reconstruction error to determine the pruning ratios. FP evaluates the performance of single-layer-pruning and estimates the sensitivity of each layer. Layers with lower sensitivity are pruned more aggressively. RNP introduced an RL-based method and groups all convolutional channels into four sets for training.

The AGMC learns DNN's embeddings from its rich topology structure information automatically and can learn the embeddings with all hidden layers, not only the convolutional layers. So the AGMC is not limited to pruning the convolutional layers. In this experiment, the AGMC prunes convolution layers and linear layers together, thus performing filter pruning on convolutional layers and fine-grained pruning

on linear layers, respectively. Our baselines only prunes filter on convolutional layers.

We prune the VGG-16 [33], MobileNet and MobileNet-V2 [31]. Table II shows the results. The AGMC prunes convolutional and dense layers together, where the baselines only prune filter on convolutional layers. On MobileNet-V2, which are already highly compact, the AGMC prunes 30% FLOPs on convolutional layers and 30% parameters on dense layers and outperforms all the baselines with higher accuracy and lower density. The AGMC only fast fine-tuning on $20K$ images split from the test set of ILSVRC-2012. However, all the baselines are fine-tuning on the whole training set. We believe that if we are fine-tuning the compressed model on the whole training set, we could achieve better performance and a higher compression ratio.

Moreover, the AGMC performs structured pruning and unstructured pruning together in a single-shot manner, which is more challenging, obtaining the benefits from structured and unstructured pruning. We can not only get a high compression ratio but also speed up on parallel devices like GPUs.

V. Conclusion

In this paper, we proposed AGMC, which combines GNN and RL to explore the deep network compression policies automatically. To the best of our knowledge, we are the first to model DNNs as hierarchical computational graphs for model compression. We introduced a novel auto graph encoder-decoder to learn the embedding of DNNs. We performed comprehensive experiments on CIFAR-10 and ILSVRC-2012 datasets. Together with the efficient DNN embedding techniques, our model outperforms the handcrafted and learning-based methods by a large margin. On ResNet-20/50/56, we outperform all the baselines with a 1.8% higher accuracy. Moreover, under the AGMC framework, we introduced a novel pruning scheme that performs structured and unstructured pruning together in a single-shot manner, thus benefiting from structured and unstructured pruning. Using this method, we achieved a higher compression ratio with only 0.93% accuracy loss on MobileNet-V2.

References

[1] C. Blakeney, Y. Yan, and Z. Zong. Is pruning compression?: Investigating pruning via network layer similarity. *WACV*, pages 903–911, 2020.

[2] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, Jan. 2011.

[3] H. Gao, Z. Wang, L. Cai, and S. Ji. Channelnets: Compact and efficient convolutional neural networks via channel-wise convolutions. *IEEE transactions on pattern analysis and machine intelligence*, pages 1–1, 2020.

[4] Y. Guo, A. Yao, and Y. Chen. Dynamic network surgery for efficient dnns. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, pages 1379–1387. Curran Associates, Inc., 2016.

[5] S. Han, J. Kang, H. Mao, Y. Hu, X. Li, Y. Li, D. Xie, H. Luo, S. Yao, Y. Wang, H. Yang, and B. J. W. Dally. ESE: Efficient speech recognition engine with sparse LSTM on FPGA. *Proc. of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 75–84, 2017.

[6] S. Han, H. Mao, and J. W. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *International conference on learning representations*, 2015.

[7] B. Hassibi and G. D. Stork. Second order derivatives for network pruning: Optimal brain surgeon. *NIPS*, pages 164–171, 1992.

[8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. pages 770–778, 2016.

[9] Y. He and S. Han. Adc: Automated deep compression and acceleration with reinforcement learning. *arXiv preprint arXiv:1802.03494*, 2, 2018.

[10] Y. He, G. Kang, X. Dong, Y. Fu, and Y. Yang. Soft filter pruning for accelerating deep convolutional neural networks. In *Proc. of the 27th International Joint Conference on Artificial Intelligence*, pages 2234–2240, 2018.

[11] Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, and S. Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proc. of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018.

[12] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019.

[13] Y. He, X. Zhang, and J. Sun. Channel pruning for accelerating very deep neural networks. In *Proc. of the IEEE International Conference on Computer Vision*, pages 1389–1397, 2017.

[14] E. G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *CoRR*, 2015.

[15] G. A. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv: Computer Vision and Pattern Recognition*, 2017.

[16] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2704–2713, 2018.

[17] H. Ji, L. Song, L. Jiang, H. H. Li, and Y. Chen. Recom: An efficient resistive accelerator for compressed deep neural networks. In *Proc of. Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 237–240. IEEE, 2018.

[18] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[19] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.

[20] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

[21] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[22] J. Lin, Y. Rao, J. Lu, and J. Zhou. Runtime neural pruning. In *Advances in Neural Information Processing Systems*, pages 2181–2191, 2017.

[23] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu. Hierarchical representations for efficient architecture search. *International Conference on Learning Representations (ICLR)*, 2018.

[24] N. Liu, X. Ma, Z. Xu, Y. Wang, J. Tang, and J. Ye. Autocompress: An automatic dnn structured pruning framework for ultra-high compression rates. In *Proc. of Artificial Intelligence Conference (AAAI)*, pages 4876–4883, 2020.

[25] P. Molchanov, A. Mallya, S. Tyree, I. Frosio, and J. Kautz. Importance estimation for neural network pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11264–11272, 2019.

[26] W. Park, D. Kim, Y. Lu, and M. Cho. Relational knowledge distillation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019.

[27] A. Polino, R. Pascanu, and D. Alistarh. Model compression via distillation and quantization. *ICLR*, 2018.

[28] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.

[29] T. N. Sainath, B. Kingsbury, V. Sindhwani, E. Arisoy, and B. Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6655–6659, 2013.

[30] M. Sandler, G. A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. *arXiv: Computer Vision and Pattern*

Recognition, 2018.

[31] M. Sandler, G. A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *CVPR*, pages 4510–4520, 2018.

[32] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller. Deterministic policy gradient algorithms. volume 1, 06 2014.

[33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *international conference on learning representations*, 2015.

[34] S. R. Sutton, A. D. Mcallester, P. S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. *Neural Information Processing Systems*, pages 1057–1063, 1999.

[35] S. Swaminathan, D. Garg, R. Kannan, and F. Andres. Sparse low rank factorization for deep neural network compression. *Neurocomputing*, pages 185–196, 2020.

[36] S. K. Tai, R. Socher, and D. C. Manning. Improved semantic representations from tree-structured long short-term memory networks. *International Workshop on the ACL2 Theorem Prover and Its Applications*, 2015.

[37] H. Wang, Q. Zhang, Y. Wang, and R. Hu. Structured probabilistic pruning for deep convolutional neural network acceleration. *British Machine Vision Conference*, 2017.

[38] J. Wang, H. Bai, J. Wu, and J. Cheng. Bayesian automatic model compression. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):727–736, 2020.

[39] X. Xiao, Z. Wang, and S. Rajasekaran. Autoprun: Automatic network pruning by regularizing auxiliary parameters. *Advances in Neural Information Processing Systems (NIPS)*, pages 13681–13691, 2019.

[40] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *Proc. of International Conference on Learning Representations (ICLR)*, 2019.

[41] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, and J. Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in neural information processing systems*, pages 4800–4810, 2018.

[42] M. Zhang and Y. Chen. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, pages 5165–5175, 2018.

[43] T. Zhang, S. Ye, K. Zhang, J. Tang, W. Wen, M. Fardad, and Y. Wang. A systematic DNN weight pruning framework using alternating direction method of multipliers. *ECCV*, 2018.

[44] Z. Zhuang, M. Tan, B. Zhuang, J. Liu, Y. Guo, Q. Wu, J. Huang, and J.-H. Zhu. Discrimination-aware channel pruning for deep neural networks. *Advances in Neural Information Processing Systems (NIPS)*, pages 875–886, 2018.