



End-to-End Supermask Pruning: Learning to Prune Image Captioning Models

Jia Huei Tan^a, Chee Seng Chan^{a,*}, Joon Huang Chuah^b

^a Center of Image and Signal Processing (CISIP), Department of Artificial Intelligence, Universiti Malaya, Kuala Lumpur 50603, Malaysia

^b Department of Electrical Engineering, Universiti Malaya, Kuala Lumpur 50603, Malaysia

ARTICLE INFO

Article history:

Received 13 April 2021

Revised 23 August 2021

Accepted 4 October 2021

Available online 5 October 2021

Keywords:

Image captioning

Deep network compression

Deep learning

ABSTRACT

With the advancement of deep models, research work on image captioning has led to a remarkable gain in raw performance over the last decade, along with increasing model complexity and computational cost. However, surprisingly works on compression of deep networks for image captioning task has received little to no attention. For the first time in image captioning research, we provide an extensive comparison of various unstructured weight pruning methods on three different popular image captioning architectures, namely *Soft-Attention*, *Up-Down* and *Object Relation Transformer*. Following this, we propose a novel end-to-end weight pruning method that performs gradual sparsification based on weight sensitivity to the training loss. The pruning schemes are then extended with encoder pruning, where we show that conducting both decoder pruning and training simultaneously prior to the encoder pruning provides good overall performance. Empirically, we show that an 80% to 95% sparse network (up to 75% reduction in model size) can either match or outperform its dense counterpart. The code and pre-trained models for *Up-Down* and *Object Relation Transformer* that are capable of achieving CIDEr scores >120 on the MS-COCO dataset but with only 8.7 MB and 14.5 MB in model size (size reduction of 96% and 94% respectively against dense versions) are publicly available at <https://github.com/jiahuei/sparse-image-captioning>.

© 2021 Elsevier Ltd. All rights reserved.

1. Introduction

Over the past decade, continuous research on image captioning using deep neural networks (DNNs) has led to a steady improvement in the overall model performance. For instance, CIDEr¹ scores [1] of state-of-the-art (SOTA) models have doubled from 66 points [2] to 130 points and beyond [3,4] recently on the MS-COCO dataset [5]. However, such gains are usually achieved at the expense of model size using heavily parameterised models, where the decoder size had quadrupled from 12 million [6] to 55 million [3] parameters (see Table 3 for details).

In an effort to reduce model size, various pruning techniques have been proposed to remove unimportant weights from a network. Generally, there are multitudes of benefits to be gained from weight pruning: it provides opportunities for improvements in terms of i) *speed*, ii) *storage*, and iii) *energy consumption*, es-

pecially during the deployment stage. For *speed*, highly-sparse models with significantly fewer non-zero parameters can enjoy faster run-times when combined with efficient SpMM kernels [7,8]. This is particularly true for Recurrent Neural Network (RNN) and Transformer whose matrix-multiplication computations are bottlenecked by bandwidth [9]. For *storage*, compressed models are easier to be deployed onto mobile devices. Moreover, compressing SOTA model checkpoints into tens of MB can potentially accelerate the dissemination of research findings, result reproduction and experimentation. Finally, for *energy consumption*, small RNN kernels produced via pruning can be stored in on-chip SRAM cache with lower energy requirements rather than DRAM memory [10], reducing carbon footprint.

While there is no shortage of pruning methods for image classification and translation tasks [11–14] (see Sec. 2 for more), their applicability to multimodal contexts such as image captioning is still under-explored. To the best of our knowledge, there is only one prior work that involved pruning an image captioning model, and that is by Dai et al. [15]. We hypothesise that this lack of progress is due to several difficulties. Firstly, weights are shared and reused across time steps, complicating variational pruning methods proposed for feed-forward networks [12]. Secondly, naively performing structured pruning on Long-Short Term

* Corresponding author.

E-mail addresses: tanjiahuei@siswa.um.edu.my (J.H. Tan), cs.chan@um.edu.my (C.S. Chan), jhchuah@um.edu.my (J.H. Chuah).

¹ Consensus-based Image Description Evaluation (CIDEr) is a widely-used metric for caption quality by measuring the level of consensus between generated captions and ground-truth captions.

Table 1
Caption statistics on MS-COCO test set.

Approaches	Sparsity (%)	Caption stats.	
		Unique (%)	Av. len.
SA LSTM Dense	-	42.1	9.09
SMP	80.0	42.5	9.11
	97.5	44.4	8.99
SA GRU Dense	-	42.4	9.15
SMP	80.0	43.1	9.13
	97.5	42.0	8.94
UD Dense	-	53.0	9.46
SMP	95.0	58.6	9.46
	99.1	61.7	9.30
ORT Dense	-	61.2	9.52
SMP	95.0	62.4	9.46
	99.1	61.1	9.28

Memory (LSTM) kernels can lead to invalid units due to inconsistent dimensions [16]. Thirdly, whereas small CIFAR datasets allow quick experimentation and iteration for Convolutional Neural Network (CNN) pruning [13,17], there is a lack of an equivalent dataset in image captioning. Finally, image captioning is an inherently complex multimodal task; thus any proposed method must be able to perform well on both image and language domains.

To this end, this paper attempts to answer the following questions:

- 1) Which weight pruning method produces the best results on image captioning models?
- 2) Is there an ideal sparsity where a sparse model can match or even outperform its dense counterpart?
- 3) What is the ideal prune-finetune sequence for pruning both the pre-trained encoder and decoder?
- 4) Can a sparse captioning model outperform a smaller but dense model?

with an extensive comparison of various unstructured weight pruning methods on three different SOTA image captioning architectures, namely *Soft-Attention* (SA) [6], *Up-Down* (UD) [18] and *Object Relation Transformer* (ORT) [3]. Following this, we propose a novel end-to-end weight pruning method that performs gradual sparsification while maintaining the overall model performance. The pruning schemes are then extended with encoder pruning, where several prune-finetune sequences are explored. Empirically, we show that conducting both decoder pruning and training simultaneously prior to the encoder pruning-and-finetuning provides better raw performance. Also, we show that for a given performance level, a large-sparse LSTM captioning model is better than a small-dense one in terms of model costs.

As a summary, the core contributions of this paper are three-fold. Firstly, this is the first extensive attempt at exploring unstructured model pruning for image captioning task. Empirically, we show that 80% to 95% sparse networks can either match or even slightly outperform their dense counterparts (Sec. 5.2). In addition, we propose a pruning method – Supermask Pruning (SMP) that performs continuous and gradual sparsification during training stage based on parameter sensitivity in an end-to-end fashion. Secondly, we investigate an ideal way to combine pruning with fine-tuning of pre-trained CNN, and show that both decoder pruning and training should be done before pruning the encoder (Sec. 5.4). Finally, we release the pre-trained sparse models for UD and ORT that are capable of achieving CIDEr scores >120 on the MS-COCO dataset; yet are only 8.7 MB (reduction of 96% compared to dense UD) and 14.5 MB (reduction of 94% compared to dense ORT) in

model size (Fig. 1 and Sec. 5.3). Our code and pre-trained models are publicly available².

2. Related Works

2.1. Image Captioning

Since the advent of deep neural networks, research on image captioning can be characterised by numerous architectural innovations in pursuit of raw performance (see [19] for a complete survey). The first major innovation came in the form of an end-to-end captioning network that directly generates a caption given an image [2,20]. Next came visual attention, in which one or more CNN feature maps were used to guide and condition the caption generation process [6,21]. There are also numerous works that used attributes as a way to directly inject salient information into the decoder [22–24]. Following that, [18] employed an object detector to generate image features as a form of hard-attention; which along with Transformer, became a popular captioning paradigm [3,4,25]. Concurrently, substantial effort has been put into reinforcement learning which allowed non-differentiable caption metrics to be used as optimisation objectives [26,27]. While these methods have been successful in advancing the SOTA performance, minimal effort has been made on reducing model cost [28,29], which is the main motivation of this work.

2.2. Structured or Channel Pruning

Structured pruning is a coarse-grain pruning technique whereby entire rows, columns or channels of fully connected or convolutional weights are removed. There are extensive prior work in this direction targeted at feed-forward CNNs, including [30–34] just to name a few. At the same time, structured pruning of RNNs is also widely explored [16,35,36].

Since structured pruning reduces model dimensions, the resulting network is more amenable to run-time speed-up. However, this advantage comes with several costs: **(a) Architectural constraints:** For gated RNNs such as LSTM, structured pruning requires that the pruned rows and columns of the recurrent weight kernels be aligned with each other; otherwise it may lead to invalid units [16]. The same is true for attention kernels, which is extensively used in modern captioning architectures. **(b) Lower sparsity:** Structured pruning usually provides lower sparsity for a given performance loss [37,38], often in the range of 40% (1.7×) to 90% (10×). In contrast, we demonstrate that unstructured pruning can prune an order of magnitude more at 99% (100×) while maintaining performance (see Fig. 4).

2.3. Unstructured Pruning

Recently, unstructured pruning has enjoyed emerging support, including Fast SpMM kernels [7–9] and block-sparsity support by NVIDIA Ampere GPU³ and HuggingFace Transformers library⁴. While there exist numerous unstructured pruning methods [12,13,39], we focus on methods applied to RNN and NLP models with the following characteristics: **(a) Straightforward:** Reasonably simple to implement and integrate into a standard deep network training workflow. **(b) Effective:** Able to prune at least 80% of parameters without compromising performance. **(c) Efficient:** Does not require expensive iterative pruning and retraining cycles. Thus, we arrive at the following pruning methods as a solid starting point for exploring image captioning model pruning:

² <https://github.com/jiahuei/sparse-image-captioning>

³ <https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth>

⁴ https://github.com/huggingface/pytorch_block_sparse

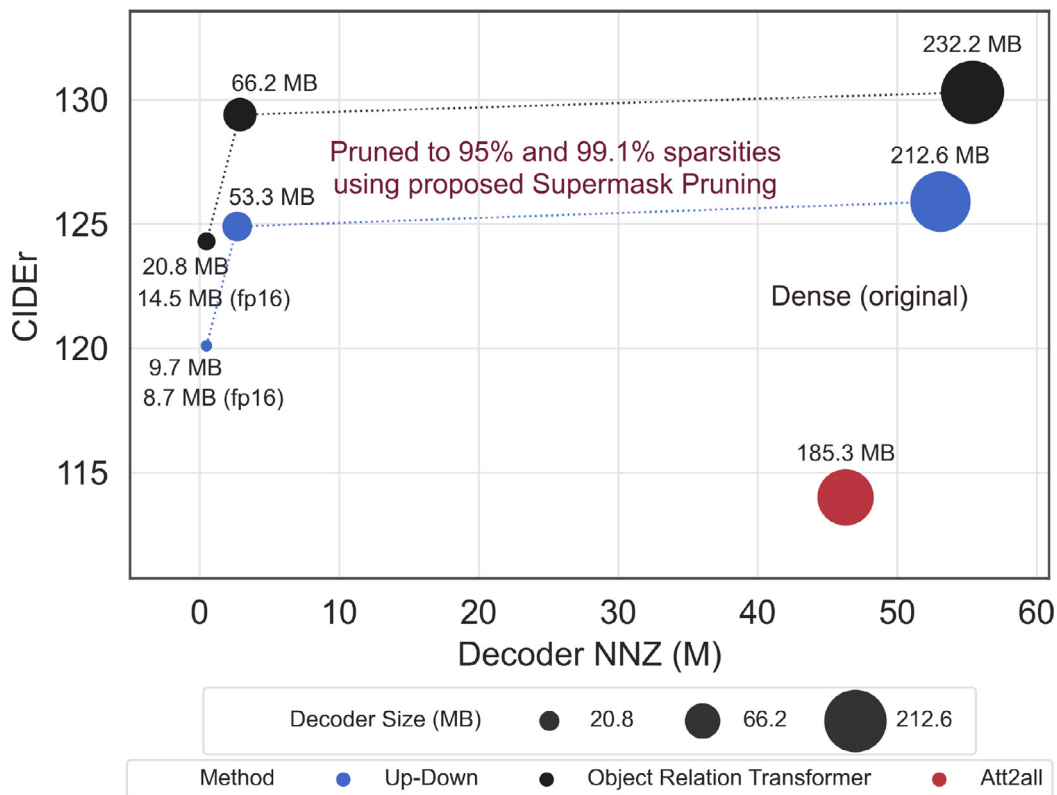


Fig. 1. Our proposed Supermask Pruning (SMP) method can produce 99.1% sparse networks that are capable of achieving CIDEr scores of 120 and above on the MS-COCO dataset, see Table 3.

(1) **Hard / one-shot magnitude-based pruning:** [40] first investigated three magnitude-based schemes for translation model with multi-layer LSTM, namely class-blind, class-uniform and class-distribution. *Class-blind* removes parameters with the smallest absolute value regardless of weight class. In contrast, *class-uniform* prunes every layer to the same sparsity level. *Class-distribution* [10] prunes parameters smaller than a global factor of the class standard deviation. Experiments found that class-blind produced the best results.

(2) **Gradual magnitude-based pruning:** First introduced by [41] to prune parameters gradually over the course of training, it was extended by [11] via a simplified pruning curve with reduced hyperparameters. The simplified slope has a single phase, governed by a cubic function that determines the sparsity level at each training step. Their method is tested on deep CNN, stacked LSTM and GNMT models on classification, language modelling and translation.

(3) **SNIP:** [14] proposed a saliency criterion for identifying structurally important connections. The criterion is computed as the absolute magnitude of the derivative of training loss with respect to a set of multiplicative pruning masks. Guided by the saliency criterion, single-shot pruning of CNN and RNN were performed at initialisation, prior to training. It is evaluated on the task of image classification.

(4) **Lottery ticket (LT):** It is a seminal work by [17] which put forth “The Lottery Ticket Hypothesis”. It states that there exists a subnetwork in a randomly initialised dense neural network, such that when trained in isolation can match the test accuracy of the original network. By iteratively pruning and resetting networks to their original initialisation values, the authors found sparse networks that can reach the original dense accuracy within equal or shorter training iterations. It is tested on CNNs for image classification.

Table 2

Comparison with H-LSTM [15,44]. Both encoder and decoder of the SMP models are pruned.

Approaches	RNN (K)		MS-COCO test scores			
	NNZ	FLOP	B-1	B-4	C	S
SMP (80%)	562	1128	73.9	33.4	102.5	18.8
H-LSTM + GP	394	670	71.9	-	95.4	-
SMP (90%)	270	545	72.9	32.8	99.0	18.3
H-LSTM + GP	163	277	71.4	-	93.3	-
SMP (95%)	116	236	72.0	31.7	94.7	17.5

(5) **Supermask:** The work by [42] explored various aspects of Lottery Tickets in order to determine the reason behind its success. In the process, the authors discovered that binary pruning masks can be learned in an end-to-end fashion. However as formulated, only the masks were optimised, and there is no straightforward way to control network sparsity. Another work by [43] optimised both masks and weights of CNNs, yet similarly, its final sparsity is influenced indirectly via a set of regularisation hyperparameters. In this work, we extend *Supermask* with a **novel sparsity loss** (Eq. (6)) to directly control the final sparsity of the network.

Among the prior works on model pruning, only the work by Dai et al. [15] involved image captioning. However, there exists several differences with our work: (a) only the H-LSTM cell is pruned; (b) CNN weight pruning is not investigated; (c) grow-and-prune (GP) method [44] used requires expensive and time-consuming “grow” and “prune-retrain” cycles. In contrast, our approach prunes both encoder and decoder in-parallel with regular training. Nevertheless, we provide a performance comparison with H-LSTM in Table 2.

Table 3
Single-model comparison with captioning SOTA. NNZ and model size calculations exclude CNN.

Approaches	NNZ (M)	Model size(MB)	MS-COCO test scores				
			B-1	B-4	M	C	S
DeepVS [2]	-	-	62.5	23.0	19.5	66.0	-
SA [6]	11.9 1 ^a	70.1 ^a	70.7	24.3	23.9	-	-
ALT-ALTM [48]	-	-	75.1	35.5	27.4	110.7	20.3
ARL [49]	-	-	75.9	35.8	27.8	111.3	-
Att2all [26]	46.3 ^a	185.3 ^a	-	34.2	26.7	114.0	-
UD [18]	53.2 ^a	212.6 ^a	79.8	36.3	27.7	120.1	21.4
ORT [3]	55.4 ^a	232.2 ^a	80.5	38.6	28.7	128.3	22.6
M2 [4]	-	-	80.8	39.1	29.2	131.2	22.6
UD (95.0%)	2.7	53.3	79.7	38.5	27.9	124.9	20.9
(99.1%, float16)	0.5	8.7 ^b	78.9	37.2	27.3	120.1	20.0
ORT (95.0%)	2.9	66.2	80.5	39.1	28.5	129.4	21.6
(99.1%, float16)	0.6	14.5 ^b	79.4	37.6	27.8	124.3	20.9

^a Based on reimplementation, size in float32.
^b Size in float32: 9.7 MB (UD), 20.8 MB (ORT).

3. Supermask Revisited

Supermask [42] is a network training method proposed by Zhou et al. as part of their work on studying the Lottery Tickets phenomenon [17]. Their work aimed to uncover the critical elements that contributed towards the good performance of “winning tickets”: sparse networks that emerged from iterative prune-reset cycles. In the process, the authors discovered that an untrained, randomly initialised network could attain test performance that is significantly better than chance. This is achieved by applying a set of well-chosen masks to the network weights, effectively pruning it. These masks were hence named “Supermask”, in that they are able to boost performance even without training of the underlying weights.

3.1. Learning Supermasks

Supermasks are learned in an end-to-end fashion via stochastic gradient descent (SGD). For every weight matrix W to be pruned, a gating matrix G with the same shape as W is created. This gating matrix G operates as a masking mechanism that determines which of the parameter $w \in W$ will be involved in both the forward-execution and back-propagation of the graph. For a model with R layers, we now have two sets of parameters: gating parameters $\phi = \{G_{1:R}\}$ and network parameters $\theta = \{W_{1:R}, B_{1:R}\}$. To this end, the effective weight tensor W' is computed following Eq. (1):

$$W' = W \odot G^b \tag{1}$$

where $W, G \in \mathbb{R}^D$ are the original weight and gating matrices with shape D ; and superscript $(\cdot)^b$ indicates binary variables. \odot is element-wise multiplication.

In order to achieve the desired masking effect, G^b must contain only “hard” binary values, i.e. $G^b \in \{0, 1\}^D$. Therefore, matrix G containing continuous values is transformed into binary matrix G^b using a composite function $z(\sigma(\cdot))$. Here, $\sigma(\cdot)$ is a point-wise function that squeezes continuous values into the interval $(0, 1)$; whereas $z(\cdot)$ is another point-wise function that samples from the output of $\sigma(\cdot)$. This is shown in Eq. (2):

$$G^b = z(\sigma(G)) \tag{2}$$

Sampling from $\sigma(G)$ is done by treating $\sigma(G)$ as Bernoulli random variables, and then performing an “unbiased draw”. Unbiased draw is the sampling process where each gating value $g \in [0, 1]$ is binarised to 1 with probability g and 0 otherwise, i.e. $z(\cdot) = \text{Bern}(\cdot)$. Sigmoid function is employed as $\sigma(\cdot)$. Finally, the effective weight W' can be computed as follows by modifying Eq. (1):

$$W' = W \odot \text{Bern}(\sigma(G)) \tag{3}$$

Before training, all the gating variables $\phi = \{G_{1:R}\}$ are initialised with the same constant value m , whereas the weights of the network are initialised randomly. The authors found that the utilisation of $\text{Bern}(\cdot)$ helped to mitigate the bias arising from the constant value initialisation by injecting stochasticity into the training process.

Although Supermask is an effective pruning technique, the formulation as presented does not allow for easy control of final network sparsity. Instead, the pruning ratios were indirectly controlled via the pruning mask initialisation magnitude. In order to address this limitation, we proposed Supermask Pruning (SMP) method that is explained next.

4. Supermask Pruning (SMP)

In this paper, we propose a simple yet effective method to directly control the final weight sparsity of models pruned based on the Supermask framework. To achieve this, a *novel sparsity loss* L_s is formulated which allows one to drive the sparsity level of gating variables ϕ to a user-specified level s_{target} . We name our method Supermask Pruning (SMP), and an overview is illustrated in Fig. 2. The complete algorithm is given in Algorithm 1.

4.1. Sparsity Loss

Technically, a straightforward way to influence the sparsity and pruning rate of Supermask is to introduce an L_1 regularisation term as follows:

$$L_s = |*| s_{target} - \left(1 - \frac{p_{nnz}}{p_{total}}\right) \tag{4}$$

where p_{nnz} is the number of non-zero (NNZ) gating parameters; p_{total} is the total number of gating parameters; n and n_{max} are the current and final training step respectively. Such a regularisation term as formulated in Eq. (4) would apply a downward pressure on the magnitude of the gating parameters G over the course of training, so that by the end of training, most of the gating parameters would have magnitudes smaller than zero. Ideally, these negative-valued gating parameters would represent weights that are least important, and can thus be removed without significant performance impact. At the same time, smaller gating magnitudes will cause more weights to be dropped more frequently, which in turn would allow the network to learn to depend on fewer parameters.

However, while naively applying the regularisation term (Eq. (4)) can produce networks with the desired sparsities, it does not achieve optimal performance. Our preliminary experiments found that constant application of L_s causes weights to be

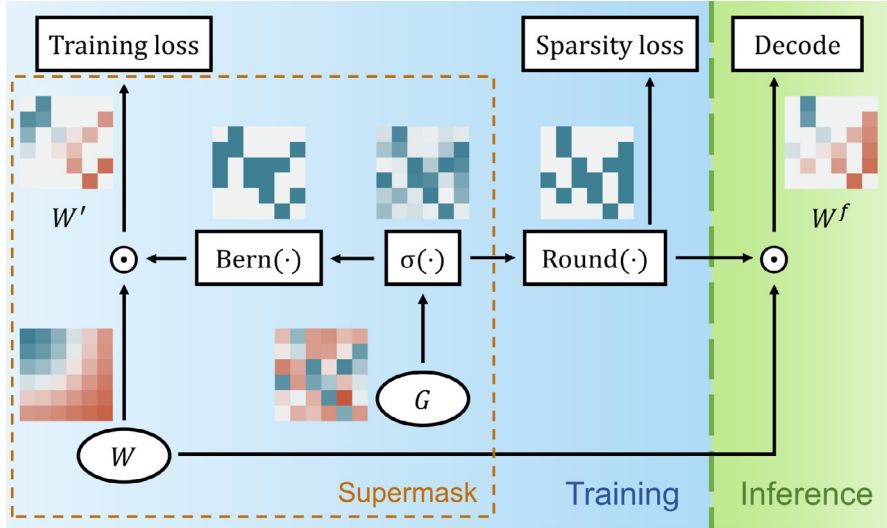


Fig. 2. An overview of our proposed Supermask Pruning (SMP). During the training stage, each weight is probabilistically retained via Bernoulli sampling on the gating parameters which are sparsified via a sparsity loss. Upon training completion, all weights are sparsified by multiplying with binarised gating parameters.

Algorithm 1: Supermask Pruning (SMP)

Require: Model parameters $\theta = \{W_{1:R}, B_{1:R}\}$, gating parameters $\phi = \{G_{1:R}\}$, sparsity target s_{target} , maximum training step n_{max} , optimizer η , training data D , loss function L

Output: Final model parameters $\theta^f = \{W^f, B^f\}$

$\theta \leftarrow \text{ModelSpecificInitializer}$; \Rightarrow Initialise model parameters
 $\phi \leftarrow m$; \Rightarrow Initialise gating parameters with constant m

for $n \in \{0, 1, \dots, n_{max}\}$ **do**
 $(I_n, C_n) \sim D$; \Rightarrow Sample a mini-batch of training data
 $W' \leftarrow W \odot \text{Bern}(\sigma(G))$; \Rightarrow Sample effective weights, refer Eq. (3)
 $\alpha \leftarrow 1 - \frac{1}{2} \left(1 + \cos\left(\frac{n\pi}{n_{max}}\right)\right)$; \Rightarrow Sparsity loss annealing, refer Eq. (5)
 $p_{nnz} \leftarrow \|\text{Round}(\sigma(\phi))\|^1$; \Rightarrow Compute NNZ, refer Eq. (7)
 $L_s \leftarrow \alpha * |s_{target} - \left(1 - \frac{p_{nnz}}{p_{total}}\right)|$; \Rightarrow Sparsity loss, refer Eq. (6)
 $L \leftarrow L_c(I_n, C_n; W', B) + \lambda_s L_s$; \Rightarrow Final loss, refer Eq. (9)
 $\theta \leftarrow \theta - \eta \left(\frac{\partial}{\partial \theta}(L)\right)$; \Rightarrow Update model parameters
 $\phi \leftarrow \phi - \eta \left(\frac{\partial}{\partial \phi}(L)\right)$; \Rightarrow Update gating parameters

end
 $W^f \leftarrow W \odot \text{Round}(\sigma(G))$; \Rightarrow Compute final weights, refer Eq. (10)
 $B^f \leftarrow B$;
 Discard ϕ ; \Rightarrow Gating parameters can be discarded

our final sparsity loss L_s is given as:

$$\alpha = 1 - \frac{1}{2} \left(1 + \cos\left(\frac{n\pi}{n_{max}}\right)\right) \tag{5}$$

$$L_s = \alpha * |s_{target} - \left(1 - \frac{p_{nnz}}{p_{total}}\right)| \tag{6}$$

Note that to compute p_{nnz} , it is necessary to sample from the gating parameters. However, instead of using $\text{Bern}(\cdot)$ as the sampling function, we perform a “maximum-likelihood (ML) draw” [43] to sample from $\sigma(G)$ in order to ensure determinism when calculating the sparsity. ML draw involves thresholding each value g at 0.5, i.e. $z(\cdot) = \text{Round}(\cdot)$. For a model with R layers and gating variables $\phi = \{G_{1:R}\}$, this p_{nnz} computation takes the following form:

$$p_{nnz} = \|\text{Round}(\sigma(\phi))\|^1 \tag{7}$$

As both the sampling functions $\text{Bern}(\cdot)$ and $\text{Round}(\cdot)$ are non-differentiable, gradient back-prop has to be performed via an estimator. On this front, [45] had explored several gradient estimators for stochastic neurons, and the *straight-through estimator* (STE) is found to be simple yet performant. Hence, back-prop is calculated by treating both sampling functions as identity functions, such that the gradients are estimated as $\delta z(\sigma(g))/\delta \sigma(g) = 1$ such that:

$$\frac{\delta L}{\delta \sigma(g)} = \frac{\delta L}{\delta z(\sigma(g))} \tag{8}$$

Finally given an image I and caption C , the overall cost function for training the image captioning model θ with gating variables ϕ is a weighted combination of captioning loss L_c and sparsity loss L_s :

$$L(I, C, s_{target}) = L_c + \lambda_s L_s \tag{9}$$

dropped too early in the training process. In other words, it leads to an over-aggressive pruning schedule. To mitigate this, we propose to perform loss annealing by adding a variable weight α to the L_s term in the cost function.

Our idea is at the beginning of training, the value of α is set to zero to allow network learning to progress without any pruning being done. As training progresses, the value of α is gradually increased, forcing the model towards a sparse solution. Specifically, this loss annealing is done using an inverted cosine curve. Our experiments found that such gradual weight pruning produces better results, which is consistent with the observations in [11,41]. Thus,

4.2. Inference

After model training is completed, all the weight matrices W are transformed into sparse matrices via element-wise multiplication with G^b . This can be done by sampling from G using $\text{Round}(\cdot)$, after which G can be discarded. In other words, the final weights W^f are calculated as:

$$W^f = W \odot \text{Round}(\sigma(G)) \quad (10)$$

The final sparse network can then be stored in appropriate sparse matrix formats such as Coordinate List (COO) or Compressed Sparse Row (CSR) in order to realise compression in terms of storage size. This can be done easily using PyTorch, as it supports parameter saving in COO format (as of release 1.6). Alternatively, regular compression algorithms such as *gzip* can be used to compress the model.

5. Experiments

In this section, we first present the setup of our experiments, followed by the results obtained from over 6,000 GPU hours using 2 Titan X GPUs.

5.1. Experiment Setup

Architectures: Three different popular image captioning architectures are used in this work: *Soft-Attention* (SA) [6], *Up-Down* (UD) [18] and *Object Relation Transformer* (ORT) [3]. SA consists of Inception-V1 [46], and a single layer LSTM or GRU with single-head attention function. Other details such as context size, attention size and image augmentation follow [29]. For UD and ORT, we reuse the public implementations^{5 6}.

Hyperparameters: For all training, we utilise Adam [47] as the optimiser, with an epsilon of 1×10^{-2} for SA and UD. The SA models were trained for 30 epochs, whereas UD and ORT models were trained for 15 epochs. Cosine LR schedule was used for SA and UD, whereas ORT follows [3]. Following [41] and [10], lower dropout rates are used for sparse networks to account for their reduced capacity. The rest follows [29].

For Supermask Pruning (SMP), training of the gating variables ϕ is done with a higher constant learning rate (LR) of 100 without annealing. This requirement of a higher LR is also noted in [42]. All ϕ are initialised to a constant $m = 5.0$.

The other pruning methods are trained as follows. *Hard:* Pruning is applied after decoder training is completed. It is then retrained for 10 epochs. *Gradual:* Pruning begins after the first epoch is completed and ends at half of the total epochs, following the heuristics outlined in [41]. Pruning frequency is 1000. We use the standard scheme where each layer is uniformly pruned. *SNIP:* Pruning is done at initialisation using one batch of data. Implementation is based on the authors' code⁷. *Lottery Ticket:* Winning tickets are produced using hard-blind, hard-uniform and gradual pruning. For a fair comparison with other single-shot pruning methods, we follow the one-shot protocol instead of the iterative protocol.

Inference is performed using beam search without length normalisation.

Datasets: Experiments are performed on MS-COCO [5] which is a public English captioning dataset. Following prior captioning works, we utilise the "Karpathy" split [2], which assigns 5,000 images for validation, 5,000 for testing and the rest for training. Pre-processing of captions is done following [29].

Metrics: Evaluation scores are obtained using the publicly available MS-COCO evaluation toolkit⁸, which computes BLEU, METEOR, ROUGE-L, CIDEr and SPICE (B, M, R, C, S).

5.2. Pruning Image Captioning Models

In this section, we attempt to answer Questions (1) and (2) in Sec. 1 via extensive performance comparisons of the pruning methods at multiple sparsity levels. We first present the pruning results on SA in Fig. 3, followed by UD and ORT in Fig. 4. Pruning is applied to all learnable parameters except for normalisation layers and biases. All the results herein were obtained using teacher-forcing with cross-entropy loss.

Which pruning method produces the best results? Our proposed end-to-end Supermask Pruning (SMP) method provides a good performance relative to the dense baselines. This observation is valid even at high pruning ratios of 95% and above. In particular, the relative drops in CIDEr scores for UD and ORT are only marginal (−3.1% to −4.7%) even at a 111× pruning rate. This is in contrast with competing methods whose performance drops are either double or even triple compared to ours, especially on SA and UD. To further support this observation, we compute the uniqueness and length of captions produced by our sparse SMP models. Results in Table 1 shows that they are largely unaffected by the pruning rate.

Among the competing methods, gradual pruning generally outperforms hard pruning, especially at higher sparsity levels when NNZ falls to 0.6 M and below. On the other hand, the results of LTs indicates that model resetting in a one-shot scenario does not outperform direct application of the underlying pruning method. We note that better results have been reported using iterative prune-reset-train cycles, however that would lead to excessively long training times and unfair comparisons with other pruning methods.

Another notable result is the relatively poor performance of SNIP when applied to image captioning. We can observe in Fig. 3 that the performance of SNIP is acceptable at 80% sparsity only. Any higher sparsity levels than this quickly led to a collapse in caption quality, as indicated by the metric scores. We tried accumulating the saliency criterion across 100 batches in an attempt to improve the result, but the improvement is limited with a huge gap from the baseline⁹. All in all, these results reflect the difficulty of pruning generative models, as well as the importance of testing on larger datasets.

Is there an ideal sparsity? A broad trend that emerged from Fig. 3 and 4 is that the model performance is more dependent on the remaining NNZ parameters after pruning, rather than the sparsity level. Both the UD and ORT models, which are about 4× larger than the SA model, can achieve substantially higher sparsity. On the extreme end, we were able to prune 99.1% of parameters from the networks, while suffering only −3.5 CIDEr points for UD and −5.4 CIDEr points for ORT.

In addition, there are indeed ideal sparsity levels where sparse models can either match or outperform their dense counterparts. This occurs at an 80% sparsity for SA, and at a 95% sparsity for both UD and ORT. We did not further investigate the performance of these models at lower sparsities, as although it is reasonable to expect better performance, the model sizes also increase substantially.

All in all, these results showcase the strength of SMP across pruning ratios from 80% to 99.1%, while managing good performance relative to the dense baselines and other pruning methods.

⁸ <https://github.com/salaniz/pycocoevalcap>

⁹ To ensure there are no critical errors on our implementation, we had successfully reproduced the results for *LSTM-b* on MNIST with a lower error rate of 1.281% averaged across 20 runs.

⁵ <https://github.com/ruotianluo/self-critical.pytorch/tree/3.2>

⁶ https://github.com/yahoo/object_relation_transformer

⁷ <https://github.com/namhoonlee/snip-public>

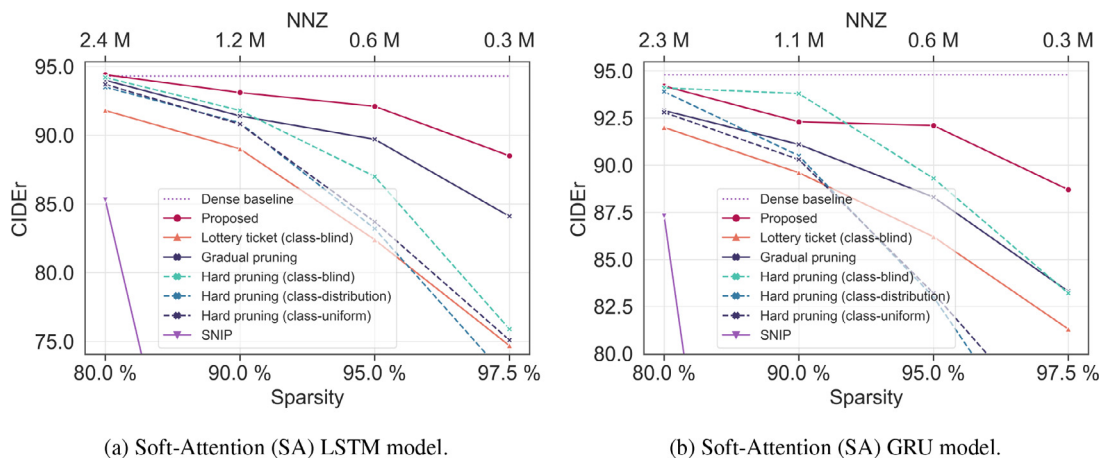


Fig. 3. Pruning performance on MS-COCO. λ_s is set following $\lambda_s = \max(5, 0.5/(1 - s_{target}))$. Sparsity, compression and NNZ figures exclude normalisation and bias parameters.

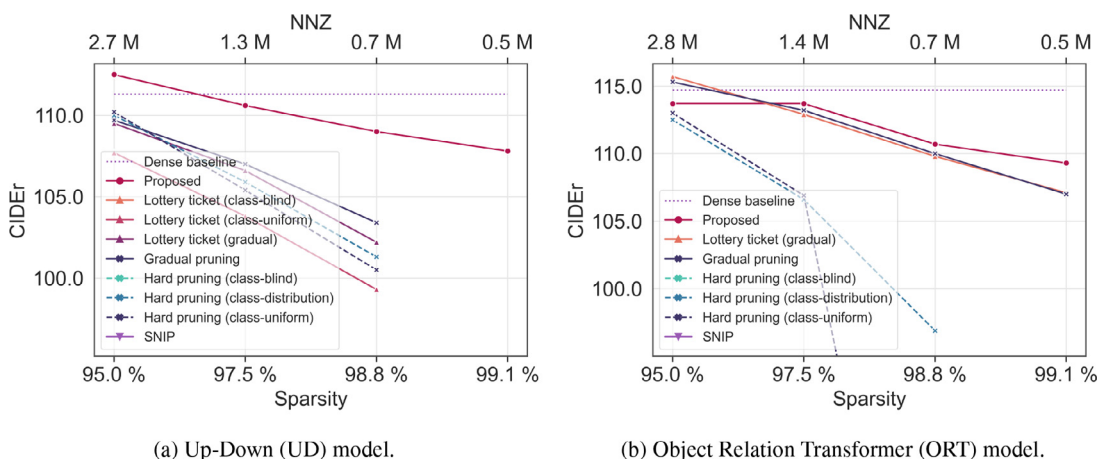


Fig. 4. Pruning performance on MS-COCO. Both SNIP and hard-blind methods failed to converge well. In order from 80% to 99.1% sparsity, λ_s is set to 80, 80, 80, 120 for UD and 120, 120, 80, 120 for ORT. Sparsity, compression and NNZ figures exclude normalisation and bias parameters.

5.3. SOTA Comparison

In this section, we compare models pruned using our proposed SMP against both H-LSTM by Dai et al. [15,44] and standard captioning SOTA approaches.

H-LSTM comparison: In Table 2, we provide the compression rate and model performance comparisons with [15]. Our SMP models are SA models trained and fine-tuned using teacher-forcing. As it can be seen, both of the SMP models at 90% and 95% sparsities with smaller RNN sizes outperform H-LSTM on both BLEU-4 and CIDEr. Furthermore, SMP does not require the expensive and time-consuming process of “grow-prune-retrain” cycles as required by [44].

SOTA comparison: To demonstrate that sparse SMP models are competitive with standard SOTA works, we compare UD and ORT models pruned using SMP against several SOTA approaches in Table 3. We optimised our models for BLEU-4 and CIDEr using SCST [26], but with the mean of rewards as baseline following [4,50]. Sparse models are saved in PyTorch COO format. For float16 models, weights are converted back to single-precision prior to computation.

From the results, it is evident that our pruned models are still capable of obtaining good captioning performance. In fact, our 95% sparse UD and ORT models managed to outperform their original dense counterparts. This is consistent with the findings in Section 5.2, which found that 95% sparsity is ideal. Finally, despite having a relatively small model size of 10 MB and 21 MB, our 99.1%

sparse models provided good results as well. The 99.1% sparse UD model, in particular, is able to match the dense UD model on CIDEr while outperforming it on BLEU-4.

5.4. Pruning Sequence for Encoder

In this section, we attempt to answer Question (3), which asks: *what is the ideal prune-finetune sequence for the encoder?* To answer this, we devised three prune-finetune schemes for the SA model as follows:

Scheme A: Start from scratch: Train the decoder while pruning both the encoder and decoder. Then, fine-tune both with gating ϕ frozen (i.e. not updated).

Scheme B: Start from a trained decoder: Fine-tune and prune both the encoder and decoder.

Scheme C: Start from a trained and pruned decoder: Fine-tune both the encoder and decoder, but only prune the encoder. Decoder ϕ are left frozen.

We paired each of the schemes with three pruning methods from the previous section, namely i) class-blind hard pruning, ii) gradual pruning and iii) SMP. All learnable parameters were pruned except for normalisation layers and biases. For schemes where gating parameters ϕ are frozen, we still apply $Bern(\cdot)$ to sample from $\sigma(\phi)$. However, we also found that there is minimal difference in the final performance when $Round(\cdot)$ is used instead. Scheme A is not evaluated for hard-blind as it requires a trained model prior to pruning.

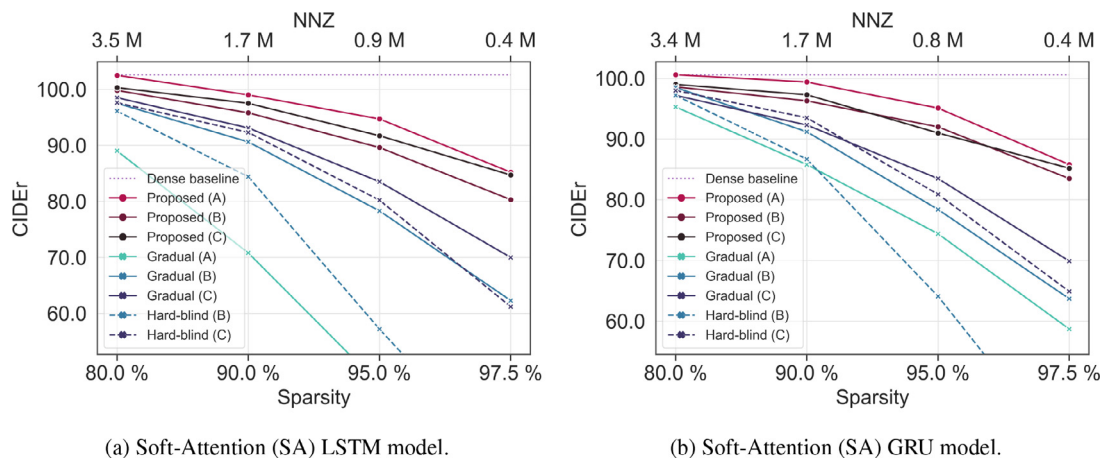


Fig. 5. Pruning performance on MS-COCO when both encoder and decoder are pruned. Sparsity, compression and NNZ figures exclude normalisation and bias parameters.

Table 4 Large-sparse versus small-dense models. Both encoder and decoder are pruned and fine-tuned.

Approaches	Cost (M)		MS-COCO test scores				
	NNZ	FLOP	B-1	B-4	M	C	S
Dense-L	15.33	4376	73.2	32.6	25.0	98.0	18.1
Sparse (80.0%)	3.08	901	72.1	31.4	24.9	95.9	17.7
Dense-M	3.37	731	70.3	29.3	23.3	86.9	16.4
Sparse (90.0%)	1.56	533	72.2	31.4	24.7	94.9	17.7
Dense-S	2.67	340	67.1	26.7	21.7	76.6	14.7
Sparse (95.0%)	0.80	307	70.0	29.2	23.3	86.0	16.1
Sparse (97.5%)	0.42	195	66.6	25.8	21.3	73.1	14.3

From Fig. 5, it is evident that Scheme A produces polarised results. Specifically, it is the best when paired with SMP, yet is the worst with gradual and hard-blind. On the other hand, Scheme C is consistently favoured over Scheme B for all three pruning methods. This shows that better performance can be attained when pruning and training for the decoder are done in-parallel rather than separately.

Comparing the three different pruning methods, we can see that the trends are consistent with the results obtained for decoder pruning in the previous section. Across different sparsity levels, our SMP method produces the best performance. At 80% sparsity, there is barely any performance loss relative to the baselines, with a mere -0.1 in CIDEr score for LSTM and no difference for GRU. At the other extreme with 2.5% of parameters, we managed CIDEr scores of 85.2 for LSTM and 85.8 for GRU, while gradual and hard-blind scored 70.0 and below.

5.5. Large-Sparse vs Small-Dense

Can a sparse model outperform a smaller dense model? Towards that end, empirical results are given in Table 4. All models are based on SA in Sec. 5.2 but with different CNN and LSTM sizes. The CNNs are MobileNet-V1: Dense-L and Sparse have a width multiplier of 1.0; Dense-M and -S have a width of 0.5 and 0.25 respectively. Moreover, Dense-M and -S have a word embedding size of 88, with an attention and LSTM size of 128. The FLOP counts are for generating a 9-word caption from a 224×224 image using a beam size of 3 (average caption length is 9). Sparse models are pruned using SMP.

Comparing models with similar metric scores, large-sparse models often have smaller NNZ and FLOP counts than their dense counterparts. Notably, a 95% sparse model can provide comparable performance as Dense-M that is larger and heavier ($4.3 \times$ NNZ and

$2.37 \times$ FLOP). This further showcases the strength of model pruning and solidifies the observations made in works on RNN pruning [11,41].

5.6. Qualitative Results and Visualisations

In this section, we present examples of captions generated, as well as visualisations of training progression, final layer-wise sparsities and weight distribution of sparse SMP models.

Qualitative results: Figure 6 shows the captions produced by our sparse UD and ORT models from Table 3. From the samples, we can see that the overall caption quality is satisfactory with sufficient details, such as umbrellas, living room, fence and school bus. Object counts are largely correct except for 5th image in which a bird is confused for two. The last image shows captions with bad endings, which is a side-effect of SCST optimisation.

Training progression: Meanwhile in Fig. 7, we can observe the effects of cosine annealing α from Eq. (5) and the sparsity regularisation weightage λ_s from Eq. (9) on the final weighted sparsity loss term. Loss annealing allows the model to focus on learning useful representations to solve the captioning task during the early stages of training, and then move towards a sparse solution during the middle to late stages when the training has stabilised. Note that whereas both figures show that sparsity levels only start to increase at around 25% of total training steps, the pruning process actually began much earlier. The average value of gating variables ϕ began to decrease around 10% into training, and continued to drop towards -8.0 throughout later stages of the training process. We can also observe that the training loss (XE loss) remained relatively stable throughout the training and pruning process even for a 97.5%-sparse model.

Layer-wise sparsities: For Inception-V1 encoder (Fig.) pruned using SMP or hard-blind pruning, we can see that earlier convolution layers with fewer parameters are pruned less heavily than later layers. This behaviour is consistent with findings in [7]. We can also see that the 3×3 convolution kernel of the second branch of each Inception module is pruned the most compared to the rest.

For LSTM decoder (Fig.), SMP and hard-blind pruning consistently prune “QK” layer (the second layer of the 2-layer attention MLP) the least, whereas “Key” and “Query” layers were pruned most heavily. Finally, “Embedding” consistently receives more pruning than “Output” despite having fewer parameters. This may indicate that there exists substantial information redundancy in the word embeddings matrix as noted in [29,51].

For MobileNet-V1 encoder (Fig.), SMP consistently prunes point-wise (1×1) convolution kernels significantly more than



Fig. 6. MS-COCO captions generated by 99.1% sparse UD (dashed blue box) and ORT (solid red box) models.




Fig. 7. Training progression of the 97.5% sparse SA model using SMP on MS-COCO. “Gating average” is the average value of gating variables ϕ ; “XE loss” refers to L_c in Eq. (9); “Weighted sparsity loss” refers to $\lambda_s L_s$ in Eq. (9) ($\lambda_s = 20$); “Sparsity loss” refers to L_s without cosine annealing in Eq. (6).

depth-wise kernels. This is a desirable outcome as point-wise operations overwhelmingly dominate the computation budget of separable convolutions both in terms of FLOP count and parameters [7,52].

Weight distribution: Lastly, distributions of non-zero weights are visualised via kernel density estimation plots in Fig. 9. We can see that the remaining weights have a bi-modal distribution centred around zero. Notably, there are still considerable amounts of small-magnitude weights after fine-tuning, even for extremely-sparse models.

5.7. Ablation Studies

This section investigates the impact of hyperparameters on the performance of SMP.

Table 5 shows the effect of different gating initialisation values m . From the results, we can establish that the best overall performance is achieved when $m = 5.0$. This can be attributed to the fact that initialisation value of 5.0 allows model parameters θ to be retained with high probability at early stages of training, leading to better convergence. This observation is also consistent with the works of [11,41], where it is found that gradual pruning can lead to better model performance. Thus, we recommend setting $m = 5.0$.

Table 5

Effects of varying initialisation value m of gating parameters ϕ on a 80% sparse SA model.

m	MS-COCO test scores					
	B-1	B-4	M	R	C	S
5.0	71.6	31.4	24.6	52.8	94.4	17.5
2.5	71.3	31.1	24.4	52.5	93.1	17.4
0.0	71.3	30.6	24.4	52.6	92.4	17.3
- 2.5	70.8	30.1	24.1	52.1	91.1	17.0
- 5.0	70.5	29.5	23.6	51.8	88.0	16.5

Table 6

Effects of varying sparsity loss weightage λ_s with $s_{target} = 0.9$ on SA model.

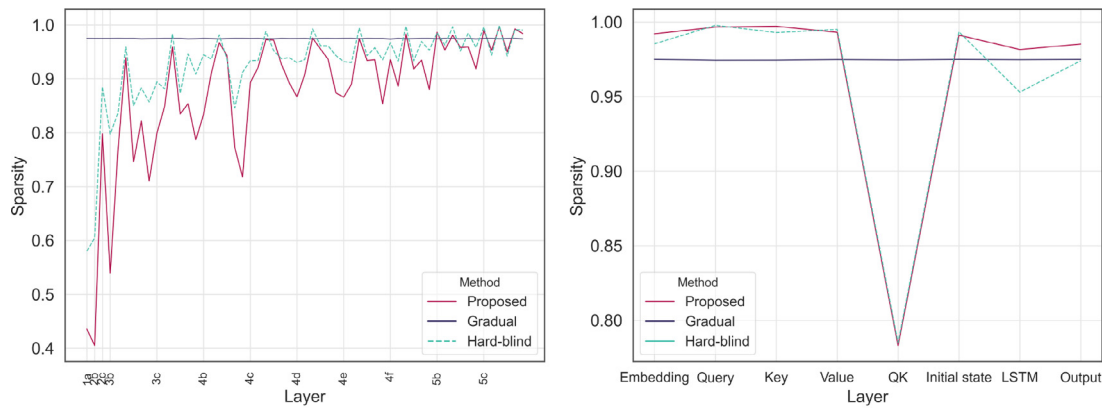
λ_s	Sparsity (%)	MS-COCO test scores					
		B-1	B-4	M	R	C	S
1.0	66.2	71.4	31.0	24.7	52.7	94.1	17.4
2.0	83.0	71.6	31.0	24.6	52.7	93.9	17.5
5.0	90.0	71.4	30.8	24.4	52.4	93.1	17.3
10.0	90.0	71.1	30.6	24.4	52.5	92.7	17.3

Table 6 shows the effect of sparsity regularisation weightage λ_s . This is an important hyperparameter that could affect the final sparsity level at convergence, with higher sparsity target s_{target} requiring larger λ_s . From the results, we can see that low values lead to insufficient sparsity. At the same time, we found that setting λ_s to a large value does not necessarily degrade its final performance, as λ_s of 80 and 120 were used for UD and ORT models in Fig. 4.

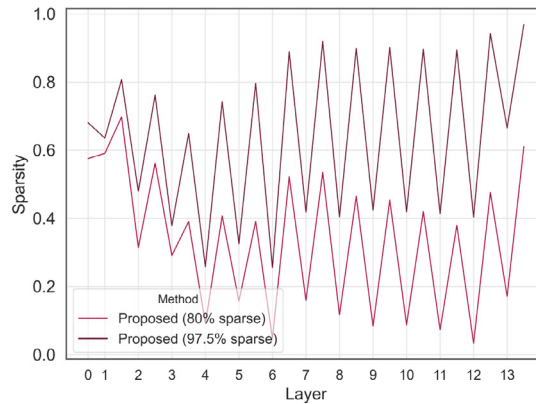
6. Discussion

In the formulation of SMP, the sparsity loss L_s is annealed using an inverted cosine curve α defined in Eq. (5). This annealing schedule is inspired by works on gradual pruning as well as works on Variational Recurrent Auto-Encoder (VRAE) for text generation. In particular, [41] has found that gradual pruning is 7% to 9% better than hard pruning. Our experiments that compared gradual-uniform pruning and hard-uniform have found this to be generally true, especially at high sparsity levels (see Sec. 5.2 and 5.4). Meanwhile, loss annealing is also used to train VRAE for text generation in the form of Kullback-Leibler (KL) annealing [53]. Specifically, the KL regularisation term is gradually introduced during training in order to shift the model from a vanilla RAE to a VRAE. In the same spirit, SMP gradually transitions the model from dense to sparse, as shown in Fig. 7.

Another perspective on the effectiveness of gradual pruning or sparsity annealing can be found in the notion of “Information Plas-




(a) Inception-V1. (b) LSTM decoder.



(c) MobileNet-V1.

Fig. 8. Layer-wise final sparsity levels of SA model. (a) and (b) are 97.5% sparse.



(a) SA model, 97.5% sparse. (b) UD model, 99.1% sparse. (c) ORT model, 99.1% sparse.

Fig. 9. Distribution of non-zero weights in models pruned using SMP.

ticity” introduced by [54]. In the work, it is found that DNN optimisation exhibits two distinct learning phases: a critical “memorisation phase” during which information stored in the weights as measured by Fisher Information rapidly increase, followed by a “forgetting phase” where the amount of information contained gradually decrease and the network is less adaptable to change. This suggests that an ideal pruning schedule should impose sparsity constraints while the network has passed its critical learning phase, and at the same time still plastic enough to adapt to such changes.

Beyond this, sensitivity-based pruning is another equally important aspect of SMP. In SNIP [14], weights are pruned based on the absolute magnitude of the derivative of training loss with respect to the multiplicative pruning masks. In contrast, SMP achieves this by updating the gating parameters according to their gradients.

This crucial difference meant that whereas SNIP removes weights with the least influence on training loss regardless of sign, SMP will also remove weights with negative influence (i.e. increase training loss).

Moreover, SNIP computes sensitivity at initialisation using one or more batches of training data. This implies that SNIP can be sensitive to the choice of weight initialisation scheme, as stated in the paper. In contrast, SMP performs continuous and gradual sparsification throughout the training process, making it less sensitive to weight initialisation. In fact, Section 5.2 shows that SMP can be used on a variety of architectures, each with its own set of initialisation schemes.

By combining these insights, we are able to realise several benefits. Firstly, SMP achieves good performance across sparsity levels from 80% to 99.1% (111× reduction in NNZ parameters). This is in

contrast with competing methods [11,40] where there is a significant performance drop-off starting from sparsity level of 90% (see Sec. 5.2). Secondly, our SMP sparsity loss allows explicit control of the overall pruning ratio and compression desired by simply specifying the target sparsity s_{target} . The pruning ratio for each layer is also automatically determined (see Fig. 8). In contrast, works like [13,43] control sparsity levels indirectly via a set of regularisation hyperparameters. Last but not least, SMP can be easily implemented on top of any model, and be integrated seamlessly into a typical training process. Only 2 main hyperparameters needed to be tuned (gating learning rate and λ_s), instead of up to 4 as in [11,41]. Since pruning is performed in-parallel with training, we can avoid the complexities and costs associated with iterative train-and-prune [15] or reinforcement learning techniques [55]. Complexities associated with variational pruning [12] such as the local reparameterisation trick can also be avoided.

7. Conclusion and Future Work

This paper presented empirical results on the effectiveness of unstructured weight pruning methods on various image captioning architectures, including RNN and Transformer architectures. In addition, we presented an effective end-to-end weight pruning method – Supermask Pruning – that performs continuous and gradual sparsification based on parameter sensitivity. Subsequently, the pruning schemes are extended by adding encoder pruning, where we showed that conducting decoder pruning and training simultaneously provides good performance. We also demonstrated that using appropriate pruning methods, ideal sparsity levels can be found in the range of 80% to 95%. These sparse networks can match or outperform their dense counterparts. Finally, we show that for a given performance level, a large-sparse LSTM captioning model outperforms a small-dense one in terms of model costs. In short, this is the first extensive attempt at exploring unstructured model pruning for image captioning. We hope that this work can spur new research interest in this direction and subsequently serve as benchmark for future image captioning pruning works.

We believe that this work opens up a sea of directions for future works. Firstly, optimised sparse matrix multiplication kernels and block-sparsity patterns can be implemented in order to realise speed-up at inference time. Finally, there are many other pruning methods that are yet to be tested, including variational pruning and saliency-based methods.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] R. Vedantam, C. Lawrence Zitnick, D. Parikh, CIDEr: Consensus-based image description evaluation, in: IEEE CVPR, 2015, pp. 4566–4575.
- [2] A. Karpathy, L. Fei-Fei, Deep visual-semantic alignments for generating image descriptions, in: IEEE CVPR, 2015, pp. 3128–3137.
- [3] S. Herdade, A. Kappeler, K. Boakey, J. Soares, Image captioning: Transforming objects into words, in: NeurIPS, 2019, pp. 11137–11147.
- [4] M. Cornia, M. Stefanini, L. Baraldi, R. Cucchiara, Meshed-Memory Transformer for Image Captioning, in: IEEE/CVF CVPR, 2020, pp. 10578–10587.
- [5] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C.L. Zitnick, Microsoft COCO: Common objects in context, in: ECCV, 2014, pp. 740–755.
- [6] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, Y. Bengio, Show, attend and tell: Neural image caption generation with visual attention, in: ICML, 2015, pp. 2048–2057.
- [7] E. Elsen, M. Dukhan, T. Gale, K. Simonyan, Fast sparse ConvNets, in: IEEE/CVF CVPR, 2020, pp. 14629–14638.
- [8] Z. Wang, SparseRT: Accelerating Unstructured Sparsity on GPUs for Deep Learning Inference, in: Proceedings of the ACM International Conference on PACT, 2020, pp. 31–42.

- [9] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. Oord, S. Dieleman, K. Kavukcuoglu, Efficient Neural Audio Synthesis, in: ICML, 80, 2018, pp. 2415–2424.
- [10] S. Han, J. Pool, J. Tran, W.J. Dally, Learning both weights and connections for efficient neural networks, in: NIPS, 2015, pp. 1135–1143.
- [11] M. Zhu, S. Gupta, To prune, or not to prune: exploring the efficacy of pruning for model compression, ICLR, Workshop Track Proceedings (2018).
- [12] N. Chirkova, E. Lobacheva, D. Vetrov, Bayesian Compression for Natural Language Processing, in: Proceedings of EMNLP, 2018, pp. 2910–2915.
- [13] C. Louizos, M. Welling, D.P. Kingma, Learning Sparse Neural Networks through L_0 Regularization, ICLR (2018).
- [14] N. Lee, T. Ajanthan, P.H. Torr, SNIP: Single-shot network pruning based on connection sensitivity, ICLR (2019).
- [15] X. Dai, H. Yin, N. Jha, Grow and Prune Compact, Fast, and Accurate LSTMs, IEEE Transactions on Computers 69 (3) (2020) 441–452.
- [16] W. Wen, Y. He, S. Rajbhandari, M. Zhang, W. Wang, F. Liu, B. Hu, Y. Chen, H. Li, Learning Intrinsic Sparse Structures within Long Short-Term Memory, ICLR (2018).
- [17] J. Frankle, M. Carbin, The lottery ticket hypothesis: Finding sparse, trainable neural networks, ICLR (2019).
- [18] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, L. Zhang, Bottom-up and top-down attention for image captioning and visual question answering, in: IEEE CVPR, 2018, pp. 6077–6086.
- [19] M.Z. Hossain, F. Sohel, M.F. Shiratuddin, H. Laga, A comprehensive survey of deep learning for image captioning, ACM CSUR 51 (6) (2019) 1–36.
- [20] Y.H. Tan, C.S. Chan, Phrase-based image caption generator with hierarchical LSTM network, Neurocomputing 333 (2019) 86–100.
- [21] K. Fu, J. Li, J. Jin, C. Zhang, Image-text surgery: Efficient concept learning in image captioning by generating pseudopairs, IEEE TNNLS 29 (12) (2018) 5910–5921.
- [22] H. Chen, G. Ding, Z. Lin, S. Zhao, J. Han, Show, Observe and Tell: Attribute-driven Attention Model for Image Captioning, in: IJCAI, 2018, pp. 606–612.
- [23] G. Ding, M. Chen, S. Zhao, H. Chen, J. Han, Q. Liu, Neural image caption generation with weighted training and reference, Cognitive Computation 11 (6) (2019) 763–777.
- [24] J. Ji, Z. Du, X. Zhang, Divergent-convergent attention for image captioning, Pattern Recognition 115 (2021) 107928.
- [25] Q. Wang, W. Huang, X. Zhang, X. Li, Word-sentence framework for remote sensing image captioning, IEEE Transactions on Geoscience and Remote Sensing (2020).
- [26] S.J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, V. Goel, Self-Critical Sequence Training for Image Captioning, in: IEEE CVPR, 2017, pp. 1179–1195.
- [27] H. Chen, G. Ding, S. Zhao, J. Han, Temporal-difference learning with sampling baseline for image captioning, in: AAAI, 2018, pp. 6706–6713.
- [28] S.N. Parameswaran, Exploring Memory and Time Efficient Neural Networks for Image Captioning, in: NCVPRIPG, Springer, 2017, pp. 338–347.
- [29] J.H. Tan, C.S. Chan, J.H. Chuah, COMIC: Toward A Compact Image Captioning Model With Attention, IEEE TMM 21 (10) (2019) 2686–2696.
- [30] J.-H. Luo, J. Wu, Autopruner: An end-to-end trainable filter pruning method for efficient deep model inference, Pattern Recognition 107 (2020) 107461.
- [31] Z. Zhuang, M. Tan, B. Zhuang, J. Liu, Y. Guo, Q. Wu, J. Huang, J. Zhu, Discrimination-aware channel pruning for deep neural networks, in: NeurIPS, 2018, pp. 875–886.
- [32] M. Lin, R. Ji, Y. Wang, Y. Zhang, B. Zhang, Y. Tian, L. Shao, HRank: Filter Pruning using High-Rank Feature Map, in: IEEE/CVF CVPR, 2020, pp. 1529–1538.
- [33] B. Li, B. Wu, J. Su, G. Wang, L. Lin, EagleEye: Fast Sub-net Evaluation for Efficient Neural Network Pruning, in: ECCV, 2020, pp. 639–654.
- [34] M. Lin, R. Ji, Y. Zhang, B. Zhang, Y. Wu, Y. Tian, Channel Pruning via Automatic Structure Search, in: IJCAI, 2020, pp. 673–679.
- [35] N. Yu, C. Weber, X. Hu, Learning Sparse Hidden States in Long Short-Term Memory, in: ICANN, Springer, 2019, pp. 288–298.
- [36] L. Wen, X. Zhang, H. Bai, Z. Xu, Structured pruning of recurrent neural networks through neuron selection, Neural Networks 123 (2020) 134–141.
- [37] E.J. Crowley, J. Turner, A. Storkey, A. O’Boyle, Pruning neural networks: Is it time to nip it in the bud? in: Workshop on Compact DNNs with Industrial Applications, NIPS, 2018, pp. 1–10.
- [38] Z. Liu, M. Sun, T. Zhou, G. Huang, T. Darrell, Rethinking the value of network pruning, ICLR (2018).
- [39] C. Wang, R. Grosse, S. Fidler, G. Zhang, EigenDamage: Structured Pruning in the Kronecker-Factored Eigenbasis, in: ICML, 2019, pp. 6566–6575.
- [40] A. See, M.-T. Luong, C.D. Manning, Compression of Neural Machine Translation Models via Pruning, in: Proceedings of The SIGLL CoNLL, ACL, 2016, pp. 291–301.
- [41] S. Narang, E. Elsen, G. Diamos, S. Sengupta, Exploring Sparsity in Recurrent Neural Networks, ICLR (2017).
- [42] H. Zhou, J. Lan, R. Liu, J. Yosinski, Deconstructing lottery tickets: Zeros, signs, and the supermask, in: NeurIPS, 2019, pp. 3597–3607.
- [43] S. Srinivas, A. Subramanya, R. Venkatesh Babu, Training sparse neural networks, in: IEEE CVPR Workshops, 2017, pp. 138–145.
- [44] X. Dai, H. Yin, N. Jha, NeST: A neural network synthesis tool based on a grow-and-prune paradigm, IEEE Transactions on Computers 68 (10) (2019) 1487–1497.
- [45] Y. Bengio, N. Léonard, A. Courville, Estimating or propagating gradients through stochastic neurons for conditional computation, arXiv preprint arXiv:1308.3432 (2013).

- [46] S. Ioffe, C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *ICML*, 37, 2015, pp. 448–456.
- [47] D.P. Kingma, J. Ba, Adam: A Method for Stochastic Optimization, *ICLR* (2014).
- [48] S. Ye, J. Han, N. Liu, Attentive Linear Transformation for Image Captioning, *IEEE TIP* 27 (11) (2018) 5514–5524.
- [49] J. Wang, W. Wang, L. Wang, Z. Wang, D.D. Feng, T. Tan, Learning visual relationship and context-aware attention for image captioning, *Pattern Recognition* 98 (2020) 107075.
- [50] R. Luo, A Better Variant of Self-Critical Sequence Training, *arXiv preprint arXiv:2003.09971* (2020).
- [51] K. Shi, K. Yu, Structured Word Embedding for Low Memory Neural Network Language Model, in: *Interspeech*, 2018, pp. 1254–1258.
- [52] A.G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, H. Adam, MobileNets: Efficient convolutional neural networks for mobile vision applications, *arXiv preprint arXiv:1704.04861* (2017).
- [53] S. Bowman, L. Vilnis, O. Vinyals, A. Dai, R. Jozefowicz, S. Bengio, Generating Sentences from a Continuous Space, in: *Proceedings of the SIGNLL CoNLL, ACL*, 2016, pp. 10–21.
- [54] A. Achille, M. Rovere, S. Soatto, Critical learning periods in deep networks, *ICLR* (2018).
- [55] Y. He, J. Lin, Z. Liu, H. Wang, L.-J. Li, S. Han, AMC: AutoML for model compression and acceleration on mobile devices, in: *ECCV*, 2018, pp. 784–800.

Jia Huei Tan received his bachelor's degree in Telecommunication Engineering from the Universiti Malaya, Malaysia in 2016. He is currently pursuing a Ph.D. degree at the Faculty of Computer Science and Information Technology, Universiti Malaya, Malaysia. His research interests include deep learning for computer vision and natural language understanding.

Chee Seng Chan is currently an Associate Professor with the Faculty of Computer Science and Information Technology, Universiti Malaya, Kuala Lumpur, Malaysia. From 2020 till 2022, he is seconded to Ministry of Science and Technology and Innovation (MOSTI, Malaysia) as an undersecretary. His research interests include computer vision and machine learning with focus on scene understanding. He is also interested in the interplay between language and vision: generating sentential descriptions about complex scenes. Dr. Chan was the Founding Chair for IEEE Computational Intelligence Society (Malaysia Chapter), Organising Chair for the Asian Conference on Pattern Recognition (2015), General Chair for the IEEE Workshop on Multimedia Signal Processing (2019), and IEEE Visual Communications and Image Processing (2013). He is a Chartered Engineer registered under the Engineering Council, U.K. He was the recipient of several notable awards, such as Young Scientist Network-Academy of Sciences Malaysia in 2015 and the Hitachi Research Fellowship in 2013.

Joon Huang Chuah received B.Eng. (Hons.) degree from the Universiti Teknologi Malaysia, the M.Eng. degree from the National University of Singapore, and the M.Phil. and Ph.D. degrees from the University of Cambridge. He is currently Head of VIP Research Group at the Department of Electrical Engineering, Faculty of Engineering, Universiti Malaya. He is a Chartered Engineer registered under the Engineering Council, U.K., and also a Professional Engineer registered under the Board of Engineers, Malaysia. He was the Honorary Treasurer of IEEE Computational Intelligence Society (CIS) Malaysia Chapter and the Honorary Secretary of IEEE Council on RFID Malaysia Chapter. He is the Vice Chairman of the Institution of Engineering and Technology (IET) Malaysia Network. In addition, he is a Fellow and the Honorary Secretary of the Institution of Engineers, Malaysia (IEM). His main research interests include image processing, computational intelligence, IC design, and scanning electron microscopy.