

NEUROPRUNE: A Neuro-inspired Topological Sparse Training Algorithm for Large Language Models

Amit Dhurandhar*, Tejaswini Pedapati*, Ronny Luss*,
Soham Dan, Aurelie Lozano, Payel Das and Georgios Kollias

IBM Research, Yorktown Heights, NY

Abstract

Transformer-based Language Models have become ubiquitous in Natural Language Processing (NLP) due to their impressive performance on various tasks. However, expensive training as well as inference remains a significant impediment to their widespread applicability. While enforcing sparsity at various levels of the model architecture has found promise in addressing scaling and efficiency issues, there remains a disconnect between how sparsity affects network topology. Inspired by brain neuronal networks, we explore sparsity approaches through the lens of network topology. Specifically, we exploit mechanisms seen in biological networks, such as preferential attachment and redundant synapse pruning, and show that principled, model-agnostic sparsity approaches are performant and efficient across diverse NLP tasks, spanning both classification (such as natural language inference) and generation (summarization, machine translation), despite our sole objective not being optimizing performance. NEUROPRUNE is competitive with (or sometimes superior to) baselines on performance and can be up to 10x faster in terms of training time for a given level of sparsity, simultaneously exhibiting measurable improvements in inference time in many cases.

1 Introduction

In the past decade, transformer-based models (Vaswani et al., 2017) leveraging *attention* mechanisms have led to state-of-the-art performance on NLP tasks and other multimodal applications, in both classification and generation settings. Despite the performance improvements, the computational overhead required for training, and inference, hinders progress. The models are large and are typically parameterized by many dense matrices, which also begs the question as to whether this complexity is necessary for better performance.

Sparsity in general neural networks has been considered using sparse regularizations on weights (Liu et al., 2024) and weight thresholding/masking (Liu et al., 2020). Specifically for transformers, various attention masking patterns have been studied (Shi et al., 2020). Another direction for inducing sparsity is to remove entire attention heads altogether (Michel et al., 2019). Previous sparse methods, however, give little emphasis to the topology of the networks being trained (Xia et al., 2022).

In this paper, we study how certain network topologies can be exploited in transformer-based large language models (LLMs) to offer sparser models (in terms of fewer parameters and fewer attention heads overall) while maintaining performance. Our framework, which we call NEUROPRUNE, is model-agnostic as well as task-agnostic, and is a dynamic sparse training method inspired by biological neuronal networks present in the brain. For example, (Yuan et al., 2019) discusses two stages by which connections (synapses) in a neuronal network evolve in the brain. First, an overabundance of synapses is created, which is similar to the pretraining an LLM. In the second stage, synapses are judiciously removed until stability in the network is achieved, which is akin to fine-tuning an LLM for a particular task by inducing sparsity, or at a higher level, by removing attention heads. Our framework relates sparsity within the Multi-Layer Perceptron (MLP) layers and the attention heads, as well as sparsity at the level of attention heads, to two distinct processes that take place during that second stage of neuronal network development: *preferential attachment* (i.e. rich-get-richer) (Lynn et al., 2024) and the *elimination of redundancy* (Wang et al., 2011).

Preferential attachment inspired regularization
Within MLP layers and attention heads, NEUROPRUNE, is motivated by a well-known network concept, called *preferential attachment*, which was found to be highly relevant in neuronal networks

*Equal contribution




Figure 1: Resulting sparsity patterns ($\approx 50\%$, $\approx 90\%$) as determined by NEUROPRUNE in an intermediate Transformer layer of a BERT(-base) model learned on the SST2 GLUE benchmark dataset (Wang et al., 2019). NEUROPRUNE sparsifies according to a preferential attachment topology as entire rows/columns of the attention and MLP matrices are zeroed out. Quantitatively, the standard deviation (sd) between the connectivity of neurons in the MLP layers increases up to two orders of magnitude (50%: 12.12, 90%: 4.16) compared with standard fine-tuning (0.13) as seen in Figure 2. This increase in sd is indicative of preferential attachment, similar to what is seen in biological neurons (Lynn et al., 2024), where a minimal *rich-get-richer* mechanism is used to produce sparse and heavy-tailed networks. The pattern is qualitatively similar for other layers too, as can be seen in the Appendix.

in the brain (Vértes et al., 2012; Lynn et al., 2024) over the last two decades. The general notion is that over time, neurons with more connections build even more connections, while those with fewer connections are removed. Similarly, our framework induces weighted l_1 sparsity in MLP layers (weight inversely \propto connectivity/degree) and group sparsity within attention heads, so that influential neurons (measured by attention parameters) are maintained while those with little influence are pruned. Modeling the removal of weak synapses is an established approach (Chechik et al., 1999) to understanding the refinement process of neurons in the brain. For LLMs, this effort is illustrated in Figure 1 where, NEUROPRUNE, sparsifies the parameter matrices of transformers by zeroing out entire rows in attention and MLP layers. Quantitatively, this non-uniformity in connectivity is verified in Figure 2.

Redundancy-based pruning While structured sparsity aims at preferential attachment, such prun-




Figure 2: Left is the MLP degree distribution for the SST2 dataset using a BERT model indicative of preferential attachment for NEUROPRUNE as sparsity increases (echoing the degree distribution in brain functional networks (Vértes et al., 2012)). Standard fine-tuning creates a dense network (black vertical line). Right we see the non-uniformity in connectivity at different sparsity %s across GLUE tasks using NEUROPRUNE, indicative of this preferential attachment across tasks.




Figure 3: Left is the fraction of times a head is removed using NEUROPRUNE when fine-tuning on SST2 with a BERT model. The overall numbers (blue curve) are averaged across layers and runs (\pm sd), where at least 10 heads are removed. We also show individual layer numbers averaged across runs for the top three layers where most pruning of heads happens (Figure 7(right)). We see there is a significant bias towards keeping the last head in each layer leading to a more modular structure and also showcasing preferential attachment, as neurons from the previous and next layers will connect only to these heads. The middle head (head 5) is also maintained more than most other heads, as it replaces many of the earlier heads. Results averaged across GLUE tasks on the right are similar.

ing (by zeroing out weights) cannot determine which connections are redundant. Elimination of redundant connections is an important aspect of the refinement process (Wang et al., 2011; Hashimoto and Kano, 2013) that takes place after the brain develops very dense networks of connections. In the case of LLMs, we hypothesize that such redundancy can be measured by similarity between attention heads, whereupon similar attention heads can then be merged, resulting in reduced complexity while maintaining functionality (i.e. performance is maintained on downstream tasks). Such removal of redundancy is conjectured in Lichtman and Colman (2020) to be unique to the neuron development in the central nervous system of vertebrates. Figure 3 illustrates how often heads are found to be redundant using NEUROPRUNE. Generally, the last head is found to be the least redundant, while the middle head also exhibits limited redundancy. In

Figure 7(right), we see head removal as a function of layers and find that the last three layers have the highest number of redundant heads.

Contributions In this paper, we propose a neuro-inspired topological sparse training algorithm with custom attention and MLP (structured) sparsity regularizations based on preferential attachment, and a novel redundancy-based head pruning scheme, which we map to the dominating set problem (Allan and Laskar, 1978) in theoretical computer science.

Our approach has the following benefits:

- It is task agnostic.
- It is easily adaptable to different transformer-based LLM architectures as it does not add additional (mask) variables to do the pruning. We apply it to BERT (encoder), T5 (encoder-decoder) and OPT (decoder) models.
- It learns sparsity patterns exhibiting principled topological structure.
- It results in LLMs with a competitive and even sometimes superior performance on different benchmarks and tasks (GLUE, summarization, machine translation), although our proposal is more neuroscience-motivated than solely trying to maximize performance.
- It is generally much faster to train than the competing baselines, with time per epoch being similar to standard fine-tuning. It also exhibits inference speedups as the topological constraints encourage $N:M$ -type sparsity.

2 Notation

A Transformer consists of multiple identical units. Each unit in turn is comprised of a Multi-Head Attention (MHA) Layer and a Feed Forward (FFN) or MLP Layer (used interchangeably). Each attention layer is partitioned into multiple heads H_i composed of Query Q_{H_i} , Key K_{H_i} and Value V_{H_i} parameter matrices. If d denotes the embedding dimension of each token in an input matrix X then,

$$H_i = \text{softmax} \left(\frac{XQ_{H_i}K_{H_i}^T X^T}{\sqrt{d}} \right) X V_{H_i}$$

A MHA layer with k heads computes the attention of all heads in parallel and concatenates them. $MHA = \text{Concat}(H_1, \dots, H_k)W^O$, where W^O is an output dense matrix.

The FFN layer in turn has two linear layers, one to expand the dimensions $L_{in} : \mathcal{R}^{d_e} \times \mathcal{R}^d$ and the other to project it back to the original dimension

$L_{out} : \mathcal{R}^d \times \mathcal{R}^{d_e}$. Typically $d_e \gg d$ (e.g. in BERT $d_e = 3072$ and $d = 768$) with L denoting the concatenation of all the MLP layers.

If Q, K, V are the Query, Key, and Value matrices of all the heads in a MHA layer concatenated together, then $Q_{H_i}, K_{H_i}, V_{H_i}$ corresponds to the $(i-1)\frac{d}{k} + 1$ to $i\frac{d}{k}$ columns in each such matrix respectively. Let L_{in, \mathcal{H}_i} denote the corresponding columns of the MLP and let $A_{\mathcal{H}_i} = [Q_{\mathcal{H}_i}, K_{\mathcal{H}_i}, V_{\mathcal{H}_i}]$.

The use of the superscript in $A^{(l)}$ denotes (attention) layer l of the transformer. We use \vec{x} to signify that x is a vector.

3 Related Work

Quantization (Ahmadian et al., 2023), Knowledge Distillation to a smaller model (Gu et al., 2023), and Model Pruning are some ways to alleviate the extensive computational cost required by LLMs. Here we review prior work on model pruning, which is most relevant to us.

3.1 Unstructured Pruning

Unstructured pruning (Frantar and Alistarh, 2023; Sun et al., 2023; Han et al., 2016; Tanaka et al., 2020; Lee et al., 2019) removes the less salient parameters from the model, thereby achieving sparsity. Based on the lottery ticket hypothesis, Frankle and Carbin (2019) performs iterative magnitude pruning. Prasanna et al. (2020); Chen et al. (2021) apply the lottery ticket hypothesis to the BERT model. This class of pruning algorithms attain high sparsity while largely maintaining accuracy, but are mostly post hoc. Moreover, the resulting pruned models do not provide much inference speedup.

3.2 Structured Pruning

A neural network can be divided into blocks or components. For instance, channels, kernels, and layers for a convolution neural network, an attention head, and a fully connected layer for a transformer. Structured Pruning (Anwar et al., 2017; Fan et al., 2020) involves removing an entire component, thus eliminating some of the multiply-and-accumulate computations, thereby accelerating inference.

Head Pruning Michel et al. (2019) were the first to examine if all the heads are necessary for a BERT model. They defined the *importance of a head* by the drop in performance of the model upon removing the head. Voita et al. (2019) apply gates to each head and learn these gates using l_0 regularization.

Li et al. (2021) also learn gates and identifies a subset of heads for each layer of the BERT model such that the drop in the model’s performance is minimal. They sample the top-k heads based on their importance score and use the Gumbel-softmax trick to make the top-k formulation differentiable. This method was shown to have superior results to other competitors on BERT, and we thus compare it with NEUROPRUNE’s head pruning strategy.

Block and Layer Pruning Fan et al. (2020) and Sajjad et al. (2020) experimented with dropping different transformer layers, such as every alternate layer, or top-k layers, or middle layers, and found inference speedups. Lagunas et al. (2021) divide the MHA and FFN layers into several blocks, and apply masks to each of the blocks to prune them.

CoFI (Xia et al., 2022) also prunes a transformer by applying gates to each of the heads m_h , one mask to the entire MHA layer, and finally, one to the MLP layer in the block. The model is then trained using l_0 regularization to learn these gates. The sparsity constraints are imposed using Lagrangian multipliers. To further boost the performance of the pruned model, CoFI jointly prunes and performs layer-wise distillation.

As NEUROPRUNE also prunes the attention matrices, the feed-forward layer, and attention heads, this is the closest baseline when varying the percentage of sparsity. Note that in contrast to CoFI, we do not require any additional mask variables where an architecture has to be modified, and hence, our approach is easily transferable across model architectures. We demonstrate this via experiments on BERT (an encoder-only model), T5 (an encoder-decoder model), and OPT (a decoder-only model).

4 Method

We propose (topological) sparsification of a Transformer block at three levels: i) The two (expand and contract) Multi-Layer Perceptron (MLP) layers, ii) the attention layers, and iii) head pruning at the level of attention layers. Our method, NEUROPRUNE, is detailed in Algorithm 1 with two sub-procedures given in Algorithms 2 and 3. The three sparsifications are described next.

4.1 MLP Sparsification

Preferential sparsification of the MLP layers is conceptually the simplest component of NEUROPRUNE. For each L_{in} and L_{out} matrix in each Transformer layer, a weighted l_1 penalty is added

Algorithm 1: NeuroPrune

Input: Model $M(\cdot; A_0, L_0)$ with N_L layers, initial Q, K, V concatenated matrices $A_0^{(l)}$ and MLP weight matrices $L_0^{(l)}$ for layer l , sparsity parameters $\alpha, \beta \geq 0$, # epochs n_e , loss function λ , redundancy threshold θ

for $i = 1$ to n_e **do**

$(A_i, L_i) \leftarrow$

$SGD\left(\lambda(M(\cdot; A_{i-1}, L_{i-1})) +$

$\sum_{l=1}^{N_L} \alpha R_{attn}^{(l)}(A_{i-1}^{(l)}) + \beta R_{mlp}^{(l)}(L_{i-1}^{(l)})\right)$

 # See Eq (1-2) for $R_{mlp}^{(l)}(\cdot), R_{attn}^{(l)}(\cdot)$

for $l = 1$ to N_L **do**

 # See Algorithm 2

$(A_i^{(l)}, L_i^{(l)}) \leftarrow$

$ELIM_REDUNDANT(A_i^{(l)}, L_i^{(l)}, \theta)$

Output: (A_{n_e}, L_{n_e})

to the training objective, where the weights for each row of entries in the matrix are inversely proportional to the (fractional) connectivity of that neuron. Specifically, let $n_{in,i}$ be the number of entries in the i^{th} row of L_{in} with absolute values less than some small $\epsilon > 0$ (with $n_{out,i}$ similarly defined for L_{out}). The MLP regularizer added to the training loss for layer l is as follows:

$$R_{mlp}^{(l)}(L^{(l)}) = \frac{1}{d} [n_{in,1}^{(l)}, \dots, n_{in,d}^{(l)}] \cdot |L_{in}^{(l)}| \cdot \vec{1}_d \quad (1)$$

$$+ \frac{1}{d_e} [n_{out,1}^{(l)}, \dots, n_{out,d}^{(l)}] \cdot |L_{out}^{(l)}| \cdot \vec{1}_{d_e}$$

where, $|\cdot|$ denotes an element-wise absolute value and $\vec{1}_d$ is a d -dimensional vector of 1s. In essence, Equation (1) penalizes neurons with less connectivity more than the densely connected ones. This explicitly encourages preferential attachment, yielding a training process where sparsely connected neurons are likely to be weeded out.

4.2 Attention Sparsification

It is not obvious what topological sparsity based on preferential attachment would entail for attention. We conceive of a novel way of inducing such sparsity by leveraging the rich literature on group sparsity (Bach et al., 2012; Hu et al., 2017).

Considering the connectivity of an input embedding neuron to the output neurons of an attention layer, it is evident that the i^{th} embedding dimension only interacts with the i^{th} row of the Q, K and

Algorithm 2: ELIM_REDUNDANT

Input: Attention layer $A = [K, Q, V]$,
 MLP layer L , k attention head index
 subsets $\mathcal{H}_1, \dots, \mathcal{H}_k$, distance threshold θ
Initialize: $S = I_k$ # k dim. identity matrix
Find similar heads
 $S[i, j] = 1\{l_\infty(A_{\mathcal{H}_i}, A_{\mathcal{H}_j}) \leq \theta\} \quad \forall i, j$
 $no_similar = 1\{\sum_{i \neq j} S[i, j] > 0\}$
if $no_similar = 0$ **then**
 | **Output:** (A, L)
Find redundant heads using similarity
See Algorithm 3
 $D \leftarrow \text{FIND_DOMINATING}(S)$
Adjust (A, L) for redundant heads
for $i = 1$ **to** k **do**
 | $\mathcal{J} = \{\mathcal{H}_j | D[i, j] = 1, i \neq j\}$
 | **for** $\mathbf{J} \in \mathcal{J}$ **do**
 | | $L_{in, \mathbf{J}} = L_{in, \mathbf{J}} + L_{in, \mathcal{H}_i}$
 | | Prune($A_{\mathcal{H}_i}, L_{\mathcal{H}_i}$) from (A, L) if
 | | $\mathcal{J} \neq \phi$
Output: (A, L)

V matrices. These interactions can be visualized as connections to the output neurons. However, even one non-zero entry in the i^{th} row of Q, K, V leads to the i^{th} input neuron being connected to all output neurons. Hence, to remove the effect of this neuron on the output neurons, one needs to zero out the i^{th} row in all three matrices. In other words, a group sparsity penalty, where each group is a row of the concatenated $A = [Q, K, V]$ matrix, is desired. Such a penalty encourages sparse rows to become more sparse as it tries to eliminate those rows by making them (almost) zero, again showcasing preferential behavior.

Rather than adding extra masking variables to implement preferential behavior, we leverage group sparsity and apply an l_p^q norm penalty on the rows of $[Q, K, V]$, where $p = 1$ and $q = 0.5$. The l_1^5 penalty was seen to be more robust to other choices in (Hu et al., 2017) as it leads to a sharp reduction in the parameter values belonging to a group. As such we add the following regularization, corresponding to the attention matrix at layer l , to the training loss, where $A^{(l)}$ is the concatenated $[Q, K, V]$ matrix:

$$R_{\text{attn}}^{(l)}(A^{(l)}) = \sum_{i=1}^d \sqrt{\sum_{j=1}^{3d} |A_{ij}^{(l)}|} \quad (2)$$

Note that the above constraint is applied across heads in the attention layer as it considers the en-

Algorithm 3: FIND_DOMINATING

Input: Similarity matrix S for k points
Initialize: $D = I_k$ # $D[i, j] = 1$ indicates
 point i can be replaced by point j
Find redundant points using similarity
for $i = 1$ **to** k **do**
 | $j^* = i$
 | **for** $j = 1$ **to** k **do**
 | | $\vec{s} = S[j^*, :] - S[j, :]$
 | | **if** $(1 \notin \vec{s} \text{ and } -1 \in \vec{s})$ **or**
 | | $(\vec{s} = \vec{0} \text{ and } j^* < j)$ **then**
 | | | $j^* = j$
 | | $D[i, j^*] = 1$
Output: D

tire Q, K, V matrices (hence the inner summation over $3d$ entries). Additionally, while the standard l_2 group penalty induces weights within a group to be similar, this l_1^5 group penalty allows sparsity patterns to be learned within a group. For example, in Figures 1 and 8, while entire rows are often removed, we also observe that certain rows only exhibit sparsity in Q and K while leaving corresponding rows of V dense, which is still valuable as it indicates that attention may not be required for those neurons/dimensions.

4.3 Head Pruning

Unlike the attention and MLP sparsifications mentioned above, head pruning is done after each epoch as seen in Algorithm 1 (NEUROPRUNE). The main idea here is to remove heads in a layer that are similar to other heads and are hence deemed redundant. We want to remove as many heads as possible in order to get maximum sparsification. NEUROPRUNE accomplishes this by determining which heads are similar to many other heads, and then maintaining such heads while removing others. Note that similarity is not transitive, and thus removal of heads is not trivial. Algorithm 2 details these steps.

NEUROPRUNE removes heads that are *dominated* by other heads, i.e., the dominated head is similar to only a subset of heads that the dominating head is similar to. The problem of keeping a minimum number of heads based on similarity can be mapped to the *dominating set problem* (Allan and Laskar, 1978), where each head is a vertex and each edge indicates being similar. We want to find the minimum number of vertices such that they, along with their adjacent vertices, account for all the vertices in the graph. This problem is

NP-Hard and our approach (detailed in Algorithm 3) is a quadratic-time approximation to solve this, where it biases towards keeping latter heads in a layer. Since our algorithm also biases towards keeping vertices (heads) with high degrees, our head pruning scheme also elicits preferential behavior.

An important note to make is that, unlike previous methods (Michel et al., 2019; Li et al., 2021), NEUROPRUNE does not prune according to head importance, but rather *head redundancy*, and hence even important heads can get pruned. The experiments indeed show that the average head importance is quite high across the heads we eliminate. This can lead to more aggressive pruning and faster train times as witnessed in our experiments.

4.4 NEUROPRUNE

Algorithm 1 puts together the above regularizations and head pruning. Our fine-tuning procedure is very similar to the vanilla Stochastic Gradient Descent (SGD) methods (Bottou et al., 2018) that are typically used for training LLMs. In each epoch, the *SGD(-)* term refers to running a single epoch of any SGD algorithm over a batched dataset. The key additions made in NEUROPRUNE are the MLP and attention regularizations, which appear in the objective being passed to *SGD(-)*. Head pruning is done after each epoch of stochastic gradient descent in the inner **for** loop in Algorithm 1.

Table 1: NeuroPrune (NP) vs l_1 pruning on the CNNDaily summarization dataset using T5-base. FT stands for standard fine tuning. As can be seen we are most of the time better on rouge metrics and as well as inference time. The train times are similar. Best values for each sparsity % (s) are bolded.

s	Meth.	\uparrow Rouge ₁	\uparrow Rouge ₂	\uparrow Rouge _L	\uparrow Rouge _{Lsum}	\downarrow Inf. Time(s)	\downarrow Train Time(s)
0	FT	43.18	20.47	30.77	40.41	0.455	24603
25	NP	43.07	20.34	30.7	40.31	0.451	24620
	l_1	42.19	20.12	29.29	39.33	0.454	24621
50	NP	41.96	19.52	29.73	39.2	0.442	24605
	l_1	42.18	19.02	29.29	38.65	0.451	24601
70	NP	41.6	18.45	28.56	37.93	0.431	24623
	l_1	40.1	18.02	27.51	36.63	0.441	24628
80	NP	36.92	16.78	26.29	34.35	0.427	24614
	l_1	34.27	14.95	25.11	32.79	0.437	24610
90	NP	33.92	13.78	24.29	31.35	0.415	24602
	l_1	31.88	11.94	23.18	29.22	0.422	24608
95	NP	32.17	13.72	23.66	30.97	0.406	24610
	l_1	30.25	11.21	21.42	28.16	0.417	24611

5 Experiments

We now test our method in two different settings: i) varying sparsity and ii) varying number of heads. In each setting, we run our method on the GLUE (Wang et al., 2019) tasks, where the dev set is used

for testing¹. For i) we also test our method on the CNN/Daily Mail (Nallapati et al., 2016) summarization task. For ii) we also run our method on a machine translation task for German to English on the IWSLT (Cettolo et al., 2014) dataset.

Baselines and Models: When varying sparsity we compare against CoFI (Xia et al., 2022), which is a state-of-the-art (SOTA) method for inducing structured sparsity in LLMs. Since the author-shared code is for BERT (encoder) we compare with CoFI on BERT(-base) for GLUE tasks. We additionally implemented NEUROPRUNE for T5 (encoder-decoder) (Raffel et al., 2020) and OPT (decoder) (Zhang et al., 2022) models. Since a CoFI implementation was not available for T5 and OPT, and would require modifying the architecture, we apply l_1 sparsity to the attention and MLP blocks of the transformer as a baseline. For summarization, we use a T5(-base) model and again l_1 as a baseline.

When varying the number of heads, CoFI does not provide an easy way to control for this number, and hence we compare against a specialized head pruning method called Differential Subset Pruning (DSP) (Li et al., 2021), another SOTA head pruning method. Here too, code is available for BERT, but not for T5 or OPT, and hence we compare NeuroPrune on BERT(-base) for the GLUE tasks. We do not show head removal results for the other models as there are no natural baselines like we had for sparsity (l_1). For machine translation, we use an 18-layer encoder-decoder model with 6 heads per layer as done in (Li et al., 2021).

Metrics: For performance we report accuracy (Acc) for the GLUE datasets except COLA where Mathew’s Correlation (MCorr) is a standard metric, Rouge for summarization and BLEU scores for machine translation. We also report (average) inference and train times. Please note that by train time we mean the total time taken to train the model using dynamic sparsity and inference time is the time taken by the model to process a test example based on the sparse model that has been learned.

Hardware: All experiments were conducted on an NVIDIA A100 GPU with 40 GB memory.

Setup Details: NeuroPrune results were obtained by varying α and β from 10^{-7} to 0.1 in multiples of 10. The l_1 penalty parameter also took these values. θ took values in $\{0.15, 0.2, 0.25\}$ for the head removal experiments where the default was set to 0.15 for the GLUE and summarization tasks.

¹For MNLI we report the matched dev set accuracies.

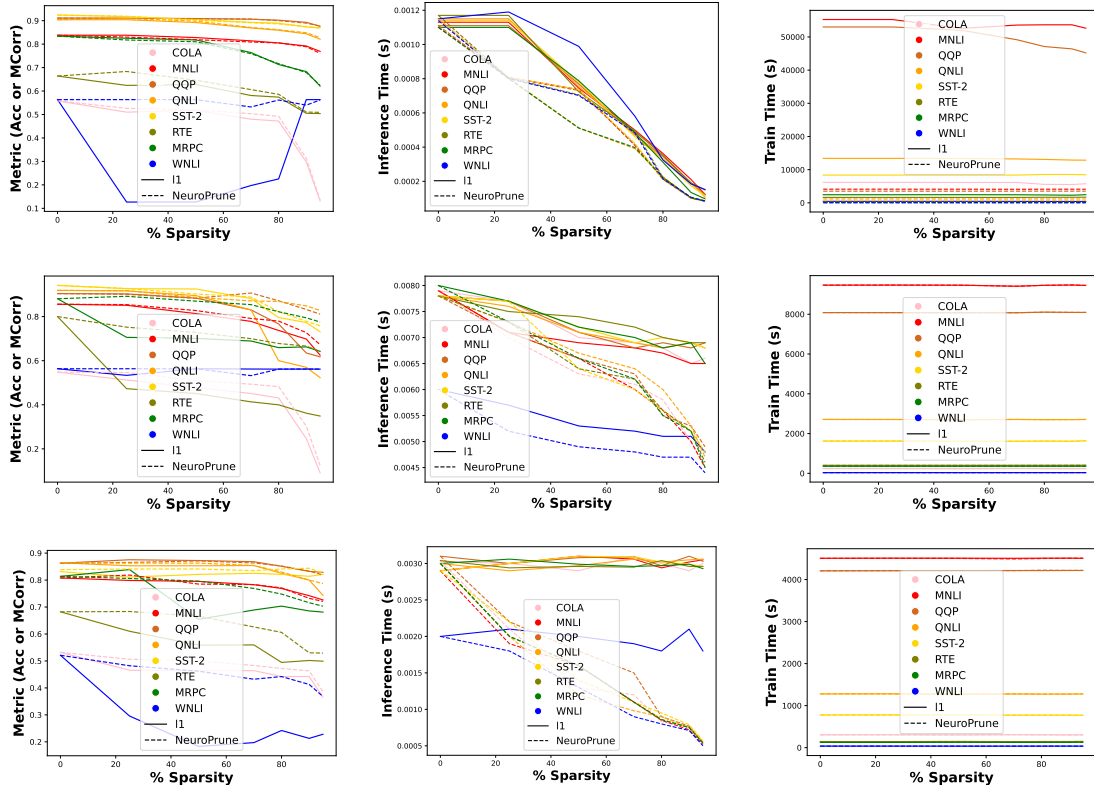


Figure 4: Performance (1st column), inference time (2nd column) and train time (3rd column) for NEUROPRUNE and CoFI/ l_1 on GLUE tasks at different sparsity percentages. The 1st, 2nd and 3rd rows correspond to BERT-base, T5-base and OPT-125m models respectively. In the 1st row we see that NEUROPRUNE outperforms CoFI on the smaller GLUE datasets and is competitive on larger ones, with consistently better inference and train times. In the next two rows, we see that NEUROPRUNE is largely better than l_1 sparsity, especially at intermediate sparsities (25-80%), with notable inference time gains and comparable train time. Qualitatively, similar results are obtained for T5-large and OPT-1.3b as seen in Figure 11 in the appendix.

It was sufficient to run NeuroPrune for a single epoch for the larger GLUE datasets (viz. MNLI, QQP, QNLI and SST2) and the summarization task, while for the other smaller GLUE datasets we ran it for 3 to 5 epochs. ϵ for MLP sparsification was set to 10^{-4} . CoFI finetunes the model before it starts pruning. For smaller GLUE datasets, finetuning before pruning epochs is 4 and total epochs is 100, while for larger datasets, these numbers are 2 and 20 respectively. All the other parameters were unchanged. To make the comparison fair, we turned off the distillation option. DSP runs 3 epochs of joint (finetuning and mask learning) training. For machine translation, NeuroPrune results were obtained with $\alpha \in \{0.005, 0.01, 0.05, 0.1, 0.25, 0.5\}$, $\beta \in \{0.05, 0.1\}$ and θ varying from 0.1 to 0.4 in increments of 0.02. ϵ was set as in GLUE. Since there was no pretrained model, 15 epochs of pretraining were done followed by 15 epochs of NeuroPrune finetuning for a total of 30 epochs. DSP joint training was also done with a total of 30 epochs.

5.1 Varying sparsity percentage

We report results for sparsity percentages of 25, 50, 70, 80, 90 and 95 for GLUE and summarization.

GLUE: In Figure 4 we see the behavior of our method on GLUE datasets for BERT, T5, and OPT models. We see that our method is always competitive with CoFI, outperforming it on the smaller datasets. This is possible because we do not add extra variables to the model, which when coupled with the topological constraints results in stabler performance. The structured sparsity also gives improved inference times and the train time is much lower, since not only do we need to run our method only for a few epochs, but the time per epoch is the same as standard fine-tuning. CoFI on the other hand, as recommended in (Xia et al., 2022), requires many epochs of running to reach specified sparsity levels as it first fine tunes and then prunes.

For T5 and OPT we compare with l_1 sparsification. As can be seen, NEUROPRUNE is better than l_1 in most cases w.r.t. performance, especially for in-between sparsities. We believe this happens be-




Figure 5: Performance (left), inference time (center) and train time (right) for NEUROPRUNE and DSP on GLUE tasks, where different number of heads are present in a BERT-base model are shown above. NEUROPRUNE is better or similar (rarely worse) in performance to DSP on most datasets, where it is notably more efficient to train. Inference time is (slightly) improved when many heads are removed, however, the DSP code (simply) masks heads rather than explicitly pruning them like ours does and hence if these masked heads are removed the inference time of DSP might also improve as shown in their paper.

cause NEUROPRUNE can choose between attention or MLP (α , β parameters), on which to sparsify more when optimizing performance even though the constraints are structured. We also see benefits in inference time, again possibly because of the structured sparsification. Train times are similar as we run l_1 for the same number of epochs as NEUROPRUNE and its per-epoch time is similar to that of standard fine-tuning. Performance with varying sparsity on T5-large and OPT-1.3b are seen in Figure 11 in the appendix.

Summarization: For summarization we see similar qualitative behavior of NEUROPRUNE vs l_1 sparsity for the T5 model. NEUROPRUNE is best on the Rouge metrics and its inference time is also slightly better. The train times are again similar.




Figure 6: Performance (left) and Efficient Frontier (right) for NEUROPRUNE and DSP on a German to English translation task. NEUROPRUNE consistently outperforms DSP across models with varying numbers of attention heads present and typically with much high levels of sparsity. The efficient frontier shows that for similar levels of performance and same number of heads, NEUROPRUNE finds much sparser models.

5.2 Varying number of heads

Beyond sparsity, we now observe the behavior of NEUROPRUNE w.r.t. the number of heads pruned. We roughly keep 10, 25, 50, 75, 100, 125 and 144 (i.e. all) heads. DSP is a strong competitor here.

GLUE: As can be seen we generally perform better or similar to DSP, but rarely worse. We believe this

is because we have an additional knob of parameter sparsification, which can make heads similar as sparsity increases and given our novel redundancy-based head pruning algorithm we can effectively replace these heads with others keeping the overall performance of the model largely intact. The training time for our method is also significantly better since we can achieve the necessary head pruning in a few epochs, in addition to the fact that each epoch takes similar time as standard fine-tuning. The inference time is better, but that improvement may be reduced if the DSP code explicitly removed heads rather than just masking them².




Figure 7: Left, we see the relative importance of removed heads for NeuroPrune and DSP averaged across runs, where at least 10 heads are removed. We see that NeuroPrune removes more important heads than DSP does because of its redundancy based head elimination algorithm. Right, we see the (average) # of heads removed per layer using NeuroPrune when fine tuning on the GLUE Benchmark datasets with a BERT(-base) model for cases where at least 10 heads are removed. As can be seen more heads are pruned from the later layers and the first one as compared to the middle layers.

Machine translation: We see in Figure 6 (left) that NEUROPRUNE outperforms DSP on a machine translation task and with higher levels of sparsity. Figure 6 (right) illustrates efficient frontiers for both NEUROPRUNE and DSP; each curve offers a variety of models (varying sparsity and number of heads present) that achieve roughly similar perfor-

²We use HuggingFace’s `prune_heads()` function for this.

mance. NEUROPRUNE clearly dominates DSP in terms of sparsity. DSP was run for varying numbers of heads present compared with varying sparsity parameters for NEUROPRUNE which generated more models. Unlike with GLUE tasks, we implemented NEUROPRUNE regularizations and head pruning within the DSP codebase which uses the fairseq toolkit (Ott et al., 2019) and used head masking for pruning. Since attention heads are still attached to the model, inference and train times between the methods on this task were similar; however note that the sparsity led to faster convergence, i.e., significantly better performance, for NEUROPRUNE when the number of heads present was greater than 70. NEUROPRUNE also has the advantage that it can be adapted to any architecture, whereas DSP must modify an architecture to include gates. Note that Li et al. (2021) show results from Michel et al. (2019) (which significantly underperforms) and Voita et al. (2019) (which has similar performance but again requires direct architecture modification).

5.3 Other Insights

Topological sparsity: As seen in Figure 1 (and Figures 8, 9, 10 in the appendix) we see that our constraints try to eliminate neurons both in the MLP layers as well as the attention layers. This is observed as a row/column (sparsity) pattern in these matrices. Interestingly, sometimes only the Q and K entries in a row are sparsified, although the group sparsity constraint is applied to QKV jointly. The similarity in sparsification patterns for matrices Q and K directly reflects the symmetry of the operations they are jointly involved in, when computing the attention coefficients as inner products of *each row* of Q with *each row* of K , compactly as QK^T , unlike V which is simply a linear projection of the token embeddings. The MLP portions, as seen in Figure 1, exhibit preferential attachment with increasing sparsity which is consistent with brain functional networks.

Importance of removed heads: In Figure 7(left), we see the *relative importance of removed heads*, defined as the ratio between the importance of removed heads in a layer to the ones that remain averaged across all the layers. The individual head importances are computed as the sum of the absolute output dense weights that emerge from a head. We see that NEUROPRUNE can prune more important heads than DSP as it finds redundant heads, thereby sparsifying more aggressively, which is ev-

idenced by the faster training time on GLUE tasks for similar levels of head removal.

Head removal: As seen in Figure 5 there is a high bias to keep the last head in each layer. This is because NEUROPRUNE prefers keeping later heads in a layer if it is similar to a set of heads that another head may be similar to, thus encouraging more modularity in a layer. In Figure 7(right) in the appendix, we see that later layers, and often the first layer, lose more heads than the intermediate ones, which is consistent with (Li et al., 2021).

6 Discussion

This work shows that NEUROPRUNE, inspired by sparsity, modularity, and preferential attachment topology in brain functional networks, is competitive and sometimes even superior to other SOTA dynamic sparse training methods, even though it does not solely try to optimize performance. It is also more efficient than them in train time with speed-ups also seen in inference. Not to mention it is easily transferable across transformer architectures (i.e. encoder, encoder-decoder, decoder) as it does not require adding additional gating variables to the architecture or modifying the architecture in any way. Also note that since dynamic sparse training and post-pruning are complimentary to each other one can potentially apply both, that is, the latter after the former. Moreover, since dynamic sparse training approaches employ sparsity during training the benefits of sparsity in terms of storage and also sometimes computation are leveraged both during training and inference. In contrast, post-pruning methods apply pruning after (dense) network training, benefiting from sparsity at inference alone.

There are multiple avenues for future research. First, it would be interesting to combine our redundancy-based and head-importance pruning methods to produce even more aggressive and efficient pruning of heads. Second, the structured strategies we used for fine-tuning could also be tried during pre-training. Third, one could test the generalizability of models trained using NEUROPRUNE on related, but different tasks and measure if similar gains can be secured. We believe our work will spur further progress on efficient dynamic sparse training methods that are also performant.

Limitations

All the datasets we considered were in English. Results may vary for other languages. We applied our method to three LLMs, but more architectures could be tested in the future not to mention more tasks. Our method although easy to adapt to different architectures while also being efficient, it does not allow the user to specify the exact number of heads one wants to prune and architectures are limited to Transformer based LLMs. This is implicitly a function of the threshold θ and the similarity of the pre-trained/fine tuned heads as well as the sparsity. We also have three hyper-parameters (α , β and θ) that need to be specified for each run.

Ethics Statement

Our work could be used to dynamically sparsify other LLMs and models, while fine tuning or pre-training. The sparsification may result in reduced alignment if the LLM is aligned to certain values and especially if those values are not encapsulated by the loss function that is used to fine tune using our method. So although one may use the method to create smaller models one has to be cognizant of what aspects may have been lost in the process. The method is easy to adapt to transformer based models and hence could possibly be widely used but improvements such as being able to specify number of heads etc. might be beneficial in future versions.

References

- Arash Ahmadian, Saurabh Dash, Hongyu Chen, Bharat Venkitesh, Stephen Gou, Phil Blunsom, Ahmet Üstün, and Sara Hooker. 2023. Intriguing properties of quantization at scale. *arXiv preprint arXiv:2305.19268*.
- Robert B. Allan and Renu Laskar. 1978. On domination and independent domination numbers of a graph. *Discrete Mathematics*, 23(2):73–76.
- Sajid Anwar, Kyuyeon Hwang, and Wonyong Sung. 2017. Structured pruning of deep convolutional neural networks. *ACM J. Emerg. Technol. Comput. Syst.*, 13(3):32:1–32:18.
- Francis Bach, Rodolphe Jenatton, Julien Mairal, and Guillaume Obozinski. 2012. Structured Sparsity through Convex Optimization. *Statistical Science*, 27(4):450 – 468.
- Léon Bottou, Frank E. Curtis, and Jorge Nocedal. 2018. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. 2014. Report on the 11th IWSLT evaluation campaign. In *Proceedings of the 11th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 2–17.
- Gal Chechik, Isaac Meilijson, and Eytan Ruppin. 1999. Neuronal regulation: A biologically plausible mechanism for efficient synaptic pruning in development. *Neurocomputing*, 26-27:633–639.
- Xiaohan Chen, Yu Cheng, Shuohang Wang, Zhe Gan, Zhangyang Wang, and Jingjing Liu. 2021. [Earlybert: Efficient BERT training via early-bird lottery tickets](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2195–2207. Association for Computational Linguistics.
- Angela Fan, Edouard Grave, and Armand Joulin. 2020. [Reducing transformer depth on demand with structured dropout](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Jonathan Frankle and Michael Carbin. 2019. [The lottery ticket hypothesis: Finding sparse, trainable neural networks](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Elias Frantar and Dan Alistarh. 2023. [Sparsegpt: Massive language models can be accurately pruned in one-shot](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10323–10337. PMLR.
- Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2023. Knowledge distillation of large language models. *arXiv preprint arXiv:2306.08543*.
- Song Han, Huizi Mao, and William J. Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*.
- Kouichi Hashimoto and Masanobu Kano. 2013. Synapse elimination in the developing cerebellum. *Cellular and Molecular Life Sciences*, 70:4667–4680.
- Yaohua Hu, Chong Li, Kaiwen Meng, Jing Qin, and Xiaoqi Yang. 2017. Group sparse optimization via lp,q regularization. *Journal of Machine Learning Research*, 18(30):1–52.
- François Lagunas, Ella Charlaix, Victor Sanh, and Alexander M. Rush. 2021. [Block pruning for faster transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages

- 10619–10629. Association for Computational Linguistics.
- Namhoon Lee, Thalaisyasingam Ajanthan, and Philip H. S. Torr. 2019. [Snip: single-shot network pruning based on connection sensitivity](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jiaoda Li, Ryan Cotterell, and Mrinmaya Sachan. 2021. [Differentiable subset pruning of transformer heads](#). *Trans. Assoc. Comput. Linguistics*, 9:1442–1459.
- Jeff W. Lichtman and Howard Colman. 2020. Synapse elimination and indelible memory author links open overlay panel. *Neuron*, 25(2):269–278.
- J. Liu, Z. Xu, R. Shi, R. C. C. Cheung, and H. K. H. So. 2020. Dynamic sparse training: Find efficient sparse network from scratch with trainable masked layers. In *ICLR*.
- Ziming Liu, Eric Gan, and Max Tegmark. 2024. Seeing is believing: Brain-inspired modular training for mechanistic interpretability. *Entropy*, 26(1).
- Christopher W Lynn, Caroline M Holmes, and Stephanie E Palmer. 2024. Heavy-tailed neuronal connectivity arises from hebbian self-organization. *Nature Physics*, pages 1–8.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *Advances in Neural Information Processing Systems*, volume 32.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Caglar Gulcehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53.
- Sai Prasanna, Anna Rogers, and Anna Rumshisky. 2020. [When BERT plays the lottery, all tickets are winning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 3208–3229. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Hassan Sajjad, Fahim Dalvi, Nadir Durrani, and Preslav Nakov. 2020. [Poor man’s BERT: smaller and faster transformer models](#). *CoRR*, abs/2004.03844.
- H. Shi, J. Gao, X. Ren, H. Xu, X. Liang, Z. Li, and J. T. Kwok. 2020. Sparsebert: Rethinking the importance analysis in self-attention. In *ICML*.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J. Zico Kolter. 2023. [A simple and effective pruning approach for large language models](#). *CoRR*, abs/2306.11695.
- Hidenori Tanaka, Daniel Kunin, Daniel L. K. Yamins, and Surya Ganguli. 2020. [Pruning neural networks without any data by iteratively conserving synaptic flow](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*.
- Petra E Vértes, Aaron F Alexander-Bloch, Nitin Gogtay, Jay N Giedd, Judith L Rapoport, and Edward T Bullmore. 2012. Simple models of human brain functional networks. *Proceedings of the National Academy of Sciences*, 109(15):5868–5873.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5797–5808. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 24th International Conference on Learning Representations*.
- Hao Wang, Hong Liu, and Zhong wei Zhang. 2011. Elimination of redundant synaptic inputs in the absence of synaptic strengthening. *Journal of Neuroscience*, 31(46):16675–16684.
- Mengzhou Xia, Zexuan Zhong, and Danqi Chen. 2022. [Structured pruning learns compact and accurate models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1513–1528. Association for Computational Linguistics.
- Ye Yuan, Jian Liu, Peng Zhao, Fu Xing, Hong Huo, and Tao Fang. 2019. Structural insights into the dynamic evolution of neuronal networks as synaptic density decreases. *Frontiers in Neuroscience*, 13:892.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv:2205.01068*.

A Additional Figures

In Figures 8, 9 and 10 we sparsity patterns induced by NeuroPrune in a BERT-base model at different sparsity percentages on the SST2 dataset.

Performance with varying sparsity on T5-large and OPT-1.3b are seen in Figure 11.

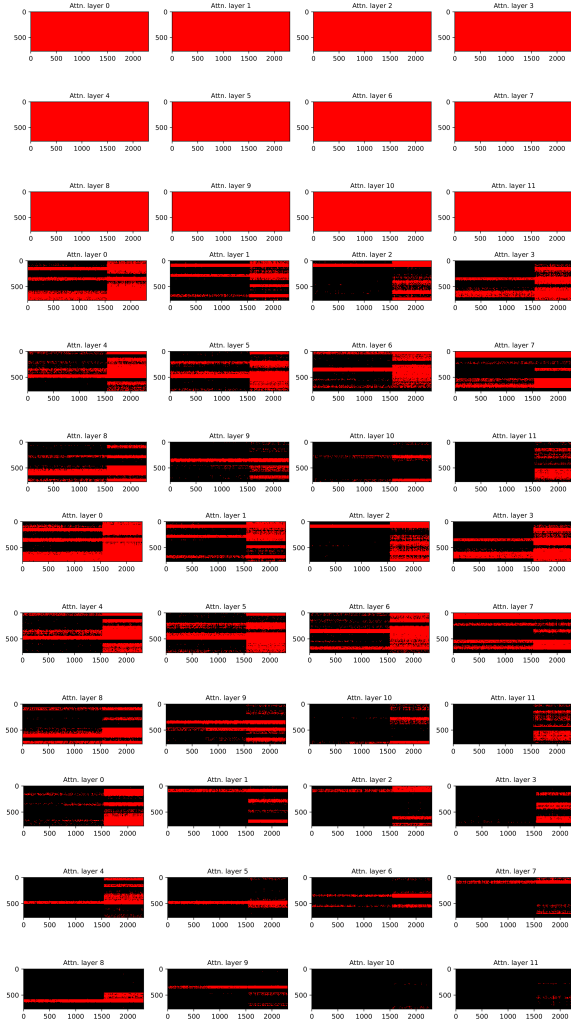


Figure 8: Attention layers for BERT model where, top three rows correspond to standard fine tuning, the next sets of three correspond to $\approx 25\%$, $\approx 50\%$ and $\approx 90\%$ sparsity using NeuroPrune. As can be seen NeuroPrune encourages preferential attachment topology.




Figure 9: MLP_in layers for BERT model where, top three rows correspond to standard fine tuning, the next sets of three correspond to $\approx 25\%$, $\approx 50\%$ and $\approx 90\%$ sparsity using NeuroPrune. As can be seen NeuroPrune encourages preferential attachment topology.



Figure 10: MLP_out layers for BERT model where, top three rows correspond to standard fine tuning, the next sets of three correspond to $\approx 25\%$, $\approx 50\%$ and $\approx 90\%$ sparsity using NeuroPrune. As can be seen NeuroPrune encourages preferential attachment topology.

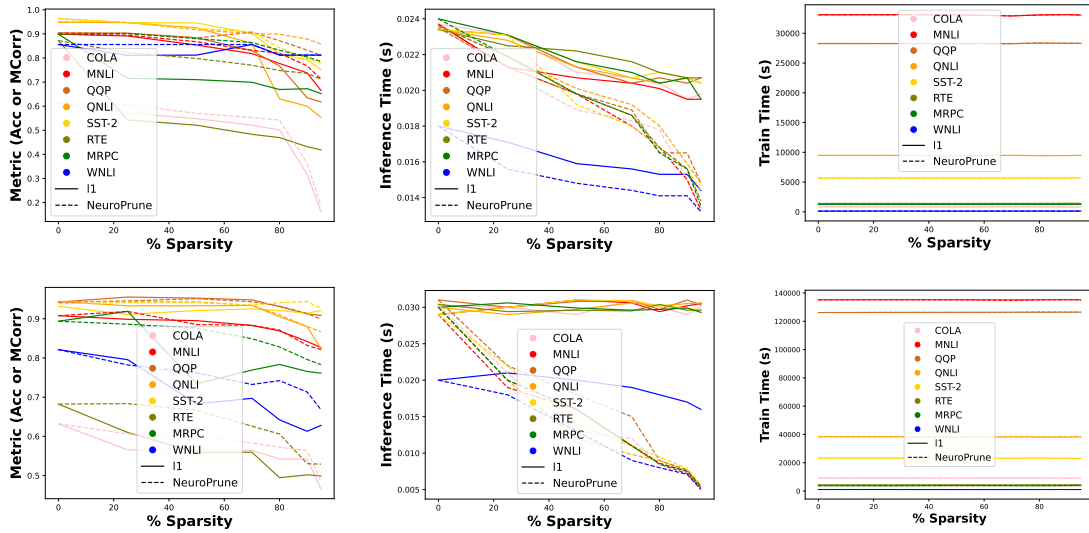


Figure 11: Performance (1st column), inference time (2nd column) and train time (3rd column) for NEUROPRUNE and l_1 on GLUE tasks at different sparsity percentages. The 1st and 2nd rows correspond to T5-large and OPT-1.3b models respectively. In the two rows, we see that NEUROPRUNE is largely better than l_1 sparsity, especially at intermediate sparsities (25-80%), with notable inference time gains and comparable train time. The results are qualitatively similar to those seen with smaller versions of these models in the main paper.