Md Shamim Hussain hussam4@rpi.edu Rensselaer Polytechnic Institute Troy, New York, USA Mohammed J. Zaki zaki@cs.rpi.edu Rensselaer Polytechnic Institute Troy, New York, USA Dharmashankar Subramanian dharmash@us.ibm.com IBM T. J. Watson Research Center Yorktown Heights, New York, USA

ABSTRACT

Transformers use the dense self-attention mechanism which gives a lot of flexibility for long-range connectivity. Over multiple layers of a deep transformer, the number of possible connectivity patterns increases exponentially. However, very few of these contribute to the performance of the network, and even fewer are essential. We hypothesize that there are sparsely connected sub-networks within a transformer, called information pathways which can be trained independently. However, the dynamic (i.e., input-dependent) nature of these pathways makes it difficult to prune dense self-attention during training. But the overall distribution of these pathways is often predictable. We take advantage of this fact to propose Stochastically Subsampled self-Attention (SSA) - a general-purpose training strategy for transformers that can reduce both the memory and computational cost of self-attention by 4 to 8 times during training while also serving as a regularization method - improving generalization over dense training. We show that an ensemble of sub-models can be formed from the subsampled pathways within a network, which can achieve better performance than its densely attended counterpart. We perform experiments on a variety of NLP, computer vision and graph learning tasks in both generative and discriminative settings to provide empirical evidence for our claims and show the effectiveness of the proposed method.

CCS CONCEPTS

• Computing methodologies \rightarrow Neural networks; Ensemble methods; Artificial intelligence.

KEYWORDS

Transformer neural networks; Self-attention; Sparse attention; Ensemble methods; Information pathway

ACM Reference Format:

Md Shamim Hussain, Mohammed J. Zaki, and Dharmashankar Subramanian. 2023. The Information Pathways Hypothesis: Transformers are Dynamic Self-Ensembles. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23), August 6–10, 2023, Long Beach, CA, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10. 1145/3580305.3599520

KDD '23, August 6-10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0103-0/23/08...\$15.00 https://doi.org/10.1145/3580305.3599520





1 INTRODUCTION

Transformer neural networks [58] have become ubiquitous in all fields of machine learning including natural language processing (NLP) [14, 50], computer vision [15, 39], and graph learning [26, 64]. The transformer architecture is based on the attention mechanism [4], which allows the model to learn to focus on the most relevant parts of the input. The global self-attention mechanism allows the transformer to update the representation of each element (e.g., token, pixel, node) of the input based on that of all other elements. The relevancy of each element is dictated by the attention weights formed by the network during the update and can be expressed as the self-attention matrix. These weights are dynamically computed by the network for each particular input. This form of flexible weighted aggregation is the key to the success of the transformer. However, the all-to-all nature of the self-attention process incurs a compute and memory cost that increases quadratically with the number of input elements N. Consequently, the self-attention process is the main efficiency bottleneck when the transformer is applied to long inputs. During the self-attention process, if element *i* applies a significant weight to element *j*, information can flow from j to i allowing them to communicate. This way, the self-attention process allows inter-element connections to form arbitrarily within a layer. However, as shown in Fig. 1, in a deep network, this communication may occur indirectly over multiple layers, for example, element k may get updated from element j and then element i may get updated from element k in the next layer, forming a communication channel that spans multiple layers. Over *l* layers, thus there are at least N^{l-1} possible ways for the two elements to communicate. The question that arises is whether all of these exponential numbers of connections contribute to the performance of the network and if not whether some of them can be pruned to save memory and computation costs during training.

Previous works like [38] have shown that the attention matrices of a fully trained transformer are sparse, and a large portion of its elements can be pruned without hurting inference time performance. Despite this sparsity, over multiple layers, connectivities

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

can reach most elements of the input, similar to expander graphs. This inspired some works to pre-define a fixed sparsity pattern to the self-attention matrix [10, 65]. However, this comes at the cost of expressivity since the model is forced to learn the attention weights within the specified *fixed* sparsity pattern. While the underlying connectivity in the self-attention process is sparse, this pattern is also dynamic, i.e., input-dependent and should not be pre-imposed. Also, these connectivities do not work in isolation within a layer but expand over multiple layers to form directed subgraphs of connectivity patterns. We call these dynamically formed sparsely connected subnetworks within the fully connected transformer information pathways. We hypothesize that not only do these pathways use a small portion of the self-attention matrix at each layer to make connections, but there are many such pathways within the network which can work independently. An ensemble of sub-models formed from a subset of pathways can often get performance close to that of the full model. Thus, we hypothesize that the transformer can be viewed as an ensemble of these sub-models, which are internally aggregated by the attention process. We use the term self-ensemble to point out that all of these sub-models use the same set of transformer weights, and vary only in inter-element connectivity. These connectivities are input dependent, and the transformer uses the pathways to perform dynamic inference on each element of the input based on the other elements. We call the information pathways that contribute to the generalization performance of the transformer important pathways, while other pathways can be deemed redundant or may even overfit the training data. To train a transformer, it is enough to ensure that these important pathways get enough training.

Previously, there has been a wealth of research on pruning the learnable weights of a neural network [19, 21, 35, 37] which reduces the cost of inference. The lottery ticket hypothesis by Frankle and Carbin [17] states that such pruning is possible because of the existence of winning tickets - very sparse subnetworks that exist within the dense network, as early as the initialization. When trained in isolation, these winning tickets can match or even exceed the performance of the dense network. Our information pathways hypothesis makes similar statements about the interconnectivity of the input elements and the dynamic weights of the attention matrix. Similar to the learnable weights, at inference time, the selfattention matrix can be dynamically pruned to reduce the inference cost both in terms of memory and compute [9, 49]. However, this is much trickier during training since the weights of the network are updated in each training step and the pruning pattern is harder to predict. In other words, unlike the winning tickets in the lottery ticket hypothesis, the important information pathways are dynamic, changing from one training sample to another. However, the connectivity patterns of the information pathways can often follow a predictable distribution. We can thus perform biased subsampling to increase the probability of covering important pathways during training while reducing the cost of training.

Our contributions are as follows – we propose a novel method for training transformers called **SSA** (Stochastically Subsampled self-Attention) that reduces both the memory and computational requirements of training while also improving generalization. SSA works by randomly subsampling the self-attention process at each training step, which allows the model to learn different connectivity patterns. We can utilize the locality of connectivity (the local inductive bias) to perform a more intelligent subsampling than random subsampling. We show that SSA can also be performed at inference time to build a self-ensemble of sub-models, each containing a subset of pathways, which can further improve generalization. We propose the information pathways hypothesis as an implication of our empirical results, which states the existence of a small number of sparsely connected and dynamic subnetworks within the transformer, the information pathways, that can be trained independently.

2 RELATED WORK

Randomly dropping part of the network such as activations [57], weights [59] or layers [25] have been seen to improve generalization. For transformers, similarly, dropping attention weights [66] and attention heads [67] have led to better generalization. Among these methods, only a few such as [25] lead to a reduction in training costs. Although dropout was originally formulated for the learnable weights of a network, they were directly adopted for the attention weights [66], which empirically improves generalization. We believe that attention dropout also trains an ensemble of pathways through the network. However, unlike attention dropout, we perform subsampling in a structured manner so that we may save training costs. We also apply local inductive bias while doing so.

After training, pruning parts of the transformer can lead to a reduction in the number of parameters and save memory [8, 46], and can potentially improve generalization [42] and/or efficiency [16] during inference. Our method is focused on stochastically dropping parts of the attention mechanism during training to reduce training costs, and can be used alongside the aforementioned methods. Additionally, we show the regularization effect of SSA and better generalization through ensembles of sparsely connected sub-models during inference.

Our method can also facilitate training on longer inputs, due to the reduction in both the memory and compute cost of self-attention. Previously, many works sought to remedy the computational bottleneck of dense self-attention via architectural modifications. This includes the use of sparse or localized self-attention [5, 10, 39, 65], or low-rank/linear/factorized attention [11, 30, 55, 60] or recurrence [12, 51] and other methods [33, 63]. These often make trade-offs in terms of expressivity, performance or generality to gain efficiency. Recently, many specialized architectures have evolved [31, 52, 54]. Despite these innovations, simple dense and local window based attention mechanisms remain relevant and competitive in many applications [62]. Unlike these approaches, we make innovations in training transformers while allowing fall-back to vanilla dense or locally dense attention at inference time.

Many innovations have also been made to reduce the training cost of transformers on long sequences. Shortformer [47] uses a staged training scheme where training is done first on short inputs followed by longer input sequences, which reduces the cost of training. Curriculum learning has also been used to stabilize training and optimize for large batches [36]. However, these approaches have only been effective in causal generative language modeling or non-causal masked language modeling tasks. Our SSA is applicable

to any causal/non-causal generative or discriminative tasks, on any form of data including text, images, and graphs.

Our self-ensembling method is related to the ensemble methods of neural networks [20, 24, 27]. However, unlike these methods, we do *not* train multiple models and average their predictions/weights. Instead, we train a single model with SSA and form an ensemble of sub-models at inference time using different subsampled attention patterns. This approach resembles Monte Carlo dropout [18], which performs dropout at inference time to make multiple predictions for uncertainty estimation. However, while MC dropout randomly drops activations, we subsample the attention mechanism from a specific distribution. Our main focus is improving generalization through self-ensembling, while its potential use for uncertainty estimation is left for future work.

3 METHOD

3.1 Background

The transformer architecture [58] consists of an encoder and a decoder. An encoder-only architecture can be used for tasks like classification [15] and masked language modeling [14], whereas a decoder-only architecture can be used for generative tasks [7, 50]. Both of these only require self-attention. For tasks like machine translation, an encoder-decoder architecture is used which additionally uses cross-attention in the decoder. We only focus on the self-attention mechanism of the transformer in this work. The key innovation of the transformer is the multihead attention mechanism, which can be expressed as:

$$\operatorname{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \operatorname{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^{T}}{\sqrt{d_{k}}}\right)\mathbf{V} = \mathbf{A}\mathbf{V}$$
(1)

where **Q**, **K**, **V** are matrices containing rows of keys, queries and values. In the case of self-attention, all of them are formed by learnable projections of the embeddings. d_k is the dimensionality of the queries and the keys. A is known as the attention matrix. Element (i, j) of this matrix is formed from the scaled dot product of query q_i and the key k_j followed by a softmax over all j. The normalized weights at row i are used to aggregate the values v_j in updating the representation of position i, thus allowing information to flow from j to i. This process is done for multiple sets of queries, keys and values, where each is called an *attention head*.

Several other terms may be added to the scaled dot product of queries and keys. A masking value $m_{ij} = -\infty$ may be added to prevent the model from attending to future positions (i.e., j > i) for generative modeling or to padding tokens; the softmax function drives the attention matrix to zero at these positions. Another term may be added to encode relative positions. Although this may take different forms [12, 44, 48, 53, 56, 61], we will discuss methods where a relative positional bias r_{i-j} is added to the scaled dot-product, e.g., [39, 48, 53]. Our method should apply to other forms of relative positional encodings as well. With the inclusion of masking and relative positional encodings, the attention matrix becomes:

$$\mathbf{A} = \operatorname{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^{T}}{\sqrt{d_{k}}} + \mathbf{M} + \mathbf{R} \right) = \operatorname{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^{T}}{\sqrt{d_{k}}} + \mathbf{B} \right)$$
(2)

Where, \mathbf{M} is the masking matrix and \mathbf{R} is the relative positional bias matrix. We merge both of these into a single bias matrix \mathbf{B} .

KDD '23, August 6-10, 2023, Long Beach, CA, USA



Figure 2: A conceptual demonstration of the information pathways hypothesis. Embeddings are e_i , information pathways are P_i , and communication channels are c_i . (a) is the full model, (b) and (c) are sub-models with only a subset of pathways.

3.2 The Information Pathways Hypothesis

The information pathways hypothesis is conceptually demonstrated in Fig. 2. We define a *communication channel* c_i as a series of selfattention based connections over multiple layers that let one element of the input gather information from another element. Each element may use many such channels to gather information from the context, i.e., other elements. A set of such connections (which may overlap) that can form a proper representation e_i of a given element is called an *information pathway* P_i . Multiple pathways may work together to form an embedding, but they can work independently as well, and can also be trained independently. The attention mechanism ensures that multiple sampled pathways are properly aggregated. If a pathway is sampled partially, it may introduce some noise in the aggregation. However, if the signals from the fully sampled pathways are strong enough, the network can ignore this noise (similar to a few weak models in an ensemble of mostly strong models). We define a *sub-model* as one that uses only a subset of the pathways P_i as in Fig. 2 (b) and (c). A randomly sampled sub-model can be trained instead of the full model, which trains the sampled subset of the pathways. Even if a pathway is not sampled at a given step, it is trained indirectly because it shares weights with the sampled pathways. If a pathway positively contributes to the generalization performance of a transformer we call it an important information pathway. With a proper sampling scheme, over multiple training steps, we can sample sub-models that cover most of the important information pathways. This is the key idea behind the proposed SSA method, which can efficiently sample the important pathways during training. Then, at inference time, we can use the full model to get the best performance, or we can use a set of sub-models to form an ensemble, which we call an attention self-ensemble. This ensemble often produces more robust predictions than the full model, because of the regularization effect of the sampling process.

3.3 Stochastically Subsampled Self-Attention

In the self-attention process, all of the *N* input elements form keys and values, and again all of the *N* input elements form queries, which is responsible for the $N \times N$ shape of the self-attention matrix, and corresponding quadratic computational cost. To efficiently subsample the self-attention matrix we decouple the elements forming keys and values, which we call source elements, and the ones KDD '23, August 6-10, 2023, Long Beach, CA, USA



Figure 3: (a) Unbiased SSA uses unbiased source shuffling with truncation, (b) locally biased SSA uses locally biased source shuffling and windowed attention. Different attention patterns result from shuffling source indices (red and blue).

Algorithm 1 Unbiased SSA

Input: Subsampled length $k \in \mathbb{N}$; embeddings $\mathbf{X} \in \mathbb{R}^{N \times d}$; query projection matrix $\mathbf{W}_{\mathbf{q}} \in \mathbb{R}^{d \times d_k}$; key projection matrix $\mathbf{W}_{\mathbf{k}} \in \mathbb{R}^{d \times d_k}$; value projection matrix $\mathbf{W}_{\mathbf{v}} \in \mathbb{R}^{d \times d_0}$; bias matrix $\mathbf{B} \in \mathbb{R}^{N \times N}$

Output: Attention head $\mathbf{H} \in \mathbb{R}^{N \times d_v}$

| 1: $\mathcal{P} \leftarrow \operatorname{randperm}(N)$ | ▷ Random permutation of indices: N |
|--|--|
| 2: $\tilde{\mathcal{P}} \leftarrow \mathcal{P}[0:k]$ | \triangleright Truncation: k |
| 3: $X_{target} \leftarrow X$ | \triangleright Target elements: $N \times d$ |
| 4: $X_{\text{source}} \leftarrow X[\tilde{\mathcal{P}}, :]$ | \triangleright Unbiased source sampling: $k \times d$ |
| 5: $\mathbf{Q} \leftarrow \mathbf{X}_{target} \mathbf{W}_{\mathbf{q}}$ | \triangleright Query projection: $N \times d_k$ |
| 6: $\mathbf{K} \leftarrow \mathbf{X}_{source} \mathbf{W}_{\mathbf{k}}$ | \triangleright Key projection: $k \times d_k$ |
| 7: $\mathbf{V} \leftarrow \mathbf{X}_{source} \mathbf{W}_{\mathbf{v}}$ | \triangleright Value projection: $k \times d_v$ |
| 8: $\tilde{\mathbf{B}} \leftarrow \mathbf{B}[:, \tilde{\mathcal{P}}]$ | \triangleright Subsample bias matrix: $N \times k$ |
| 9: $\mathbf{H} \leftarrow \operatorname{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^{T}}{\sqrt{d_{k}}} + \tilde{\mathbf{B}}\right)\mathbf{V}$ | $\triangleright \textit{ Self-attention: } N \times d_v$ |
| return H | |

forming queries, which we call target elements. In our subsampling scheme, all of the elements in the input serve as targets, but each target only attends to a random subset of sources. That is, the queries q_i are formed for all *i* but each of them attends to key-value pairs (k_j, v_j) for a random subset of *j*'s. During sampling, the inclusion of a particular source multiple times for a given target is redundant. To avoid this, we ensure the sources are sampled without replacement for each target element. We propose two forms of SSA: i) *Unbiased SSA*, and ii) *Locally biased SSA*.

Unbiased SSA: In the first implementation of SSA shown in Algorithm 1, we simply shuffle the sources in a random (unbiased) order (in line 1: randperm(N)), and truncate to keep only the first k elements, as shown in Fig. 3(a). By subsampling k sources for each target, unbiased SSA reduces the complexity of the self-attention process from $O(N^2)$ to O(Nk).

Locally Biased SSA: Here, we form local windows for both sources and targets, as shown in Algorithm 2. If both the source and target windows contain local patches of elements, then attention is confined within that window. However, if we rearrange the sources in a locally biased random order (in line 1: localrandperm(N, w, σ)), then the targets can attend to elements beyond their own window, possibly from the entire input with a non-zero probability (Fig. 3(b)). By subsampling *w* local windows, locally biased subsampling pairs

Md Shamim Hussain, Mohammed J. Zaki and Dharmashankar Subramanian

| Algorithm 2 Locally Biased SSA | |
|--|--|
| Input: Number of local windows $w \in \sigma \in \mathbb{R}$; embeddings $X \in \mathbb{R}^N$ $\mathbb{R}^{d \times d_k}$; key projection matrix matrix $\mathbf{W}_{\mathbf{v}} \in \mathbb{R}^{d \times d_v}$; bias matrix | $ \begin{array}{l} \displaystyle \overleftarrow{\in} \mathbb{N}; \text{Standard deviation of local bias} \\ \displaystyle \overset{\times d}{}; \text{query projection matrix } \mathbf{W}_{\mathbf{q}} \in \\ \displaystyle \overleftarrow{\in} \mathbf{W}_{\mathbf{k}} \in \mathbb{R}^{d \times d_{k}}; \text{ value projection} \\ \displaystyle \operatorname{trix} \mathbf{B} \in \mathbb{R}^{N \times N} \end{array} $ |
| Output: Attention head $\mathbf{H} \in \mathbb{R}^{N \times d_v}$ | |
| 1: $\mathcal{P} \leftarrow \text{localrandperm}(N, w, \sigma)$ | ▶ Locally biased random ▶ permutation of indices: N |
| 2: $X_{target} \leftarrow X$ | \triangleright Target elements: $N \times d$ |
| 3: $X_{\text{source}} \leftarrow X[\mathcal{P},:]$ | \triangleright Reindex source elements: $N \times d$ |
| 4: $\mathbf{Q} \leftarrow \mathbf{X}_{target} \mathbf{W}_{\mathbf{q}}$ | \triangleright Query projection: $N \times d_k$ |
| 5: $\mathbf{K} \leftarrow \mathbf{X}_{source} \mathbf{W}_{\mathbf{k}}$ | \triangleright Key projection: $N \times d_k$ |
| 6: $\mathbf{V} \leftarrow \mathbf{X}_{source} \mathbf{W}_{\mathbf{v}}$ | \triangleright Value projection: $N \times d_v$ |
| 7: $\tilde{\mathbf{B}} \leftarrow \mathbf{B}[:, \mathcal{P}]$ | \triangleright Reindex bias matrix: $N \times N$ |
| 8: $\mathbf{Q}_{\mathbf{w}} \leftarrow \text{window} (\mathbf{Q}, \mathbf{w})$ | $\blacktriangleright \textit{ Window query: } w \times {^N/_w} \times d_k$ |
| 9: $\mathbf{K}_{\mathbf{w}} \leftarrow \text{window}(\mathbf{K}, \mathbf{w})$ | $\blacktriangleright Window \ key: w \times {}^N/_w \times d_k$ |
| 10: $\mathbf{V}_{\mathbf{w}} \leftarrow \text{window}(\mathbf{V}, \mathbf{w})$ | \triangleright Window value: $w \times {}^N/_w \times d_v$ |
| 11: $\tilde{\mathbf{B}}_{\mathbf{w}} \leftarrow \text{window}_{\text{diag}}\left(\tilde{\mathbf{B}}, w\right)$ | ▶ Window bias along diagonal ▶ blocks: w × ^N/_w × ^N/_w |
| 12: $\mathbf{H}_{\mathbf{w}} \leftarrow \operatorname{softmax}\left(\frac{\mathbf{Q}_{\mathbf{w}}\mathbf{K}_{\mathbf{w}}^{\mathrm{T}}}{\sqrt{d_{k}}} + \tilde{\mathbf{B}}_{\mathbf{w}}\right)\mathbf{V}_{\mathbf{w}}$ | $\triangleright \textit{ Self-attention: } w \times {^N/_w} \times d_v$ |
| 13: $\mathbf{H} \leftarrow \text{flattenwindow}(\mathbf{H}_{\mathbf{w}})$ | ▷ Flatten windowed attention ▷ head: N × d _n |
| return H | |

each target with only N/w sources, reducing the complexity from $O(N^2)$ to $O(N^2/w)$.

Unbiased SSA is very easy to implement, but in our experiments, we found that locally biased SSA works better both in terms of model performance and efficiency. We pair the same set of sources with all targets for unbiased SSA, or within each window for locally biased SSA. This ensured that we can use highly optimized dense tensor multiplications for attention. Also, we use the same set of sources for all attention heads within a layer. This allows us to perform SSA by simply reindexing the embeddings and the bias matrix, followed by an unaltered/windowed multihead attention. We also use the same reindexing within each mini-batch, although, in a distributed data-parallel setting, each worker may have different indices. Both SSA algorithms can be implemented in any modern deep-learning framework in a few lines of code, without use of sparse tensor operations or custom GPU kernels.

We implement locally biased shuffling (localrandperm (N, w, σ)) by generating permutation indices whereby each index can shift around its original position with a gaussian probability distribution. We do this by adding gaussian noise to the indices:

$$\mathcal{P} = \operatorname{argsort} \left(\{ 1 + n_1, 2 + n_2, 3 + n_3, \dots, N + n_N \} \right)$$
(3)

where $n_i \sim \mathcal{N}(0, \sigma^2)$, and the standard deviation σ controls the amount of local bias. A lower value of σ results in more local bias, whereas $\sigma \to \infty$ would lead to no local bias. The resultant subsampling distribution is shown in Fig. 4 (a), where we can see that the sampling probabilities are concentrated towards the diagonal of the self-attention matrix. For generative tasks, we use a causal version of locally biased SSA, where the permutation indices are resampled for each window, and are constrained to be from 0 to the end of the window. The resulting sampling distribution is shown



Figure 4: Sampling probability of the self-attention matrix for different types of locally biased sampling: (a) gaussian, (b) causal gaussian, and (c) causal gaussian for 2D grids with vertical windows.



Figure 5: Windowed attention with (a) no source shuffling vs. (b) locally biased source shuffling – some sources move to other windows forming long-range connections, some of which are shown in red.

in Fig. 4 (b). For 2D grids, such as images we perform shuffling both horizontally and vertically. For image generation, we partition the grid vertically into w windows. The resultant distribution after locally biased shuffling and windowing is shown in Fig. 4 (c). Here we have flattened the grid row-by-row.

In Fig. 5 we show the implications of locally biased SSA on the subsampled connectivity patterns in a deep network. Simply performing windowed attention in each layer would isolate each window as in Fig. 5 (a). This is why local self-attention methods use either overlapping [5, 10] or shifted windows [39] to ensure connectivity across windows. Instead, we rely on the stochasticity of the sampling process for inter-window connectivity. We can see in Fig. 5 (b) that with locally biased SSA, after a few layers, we have long-distance connections across window boundaries with a non-zero probability while maintaining the same level of sparsity. Note that methods like BigBird [65] achieve this by a combination of local and random attention, which is kept fixed during training and inference. In contrast, the sparsity patterns in SSA are resampled at every training step, and we can fall back to dense attention during inference. Also, slowly reducing local bias (increasing σ) in the deeper layers leads to better generalization. We hypothesize that within the important information pathways, local connections are formed predominantly in the shallower layers while long-range connections are formed in deeper layers. For a given sparsity budget, locally biased SSA can sample these pathways with a higher probability than unbiased SSA. This is why locally biased SSA can achieve better performance at a lower training cost.

3.4 Fine-tuning and Inference

After training with SSA, we fall back to dense attention at inference time, which ensures that the network leverages all information pathways to produce the best output. This is analogous to the rescaling/renormalization in dropout [57] at inference time. In our case, the attention process ensures that the contributions of the pathways are properly aggregated via its weighted averaging process, so that no manual rescaling is required. We call this attention-based renormalization. Often, no extra training is required to ensure proper renormalization and good performance at inference time. However, especially when we apply a high sparsity during training, the network may need some extra adjustment to ensure proper renormalization. A small amount of fine-tuning with dense attention at the end of training is sufficient to ensure good performance at inference time. This is done in the last few epochs ($\leq 10\%$ of the total epochs). This method falls within the category of curriculum learning [6] strategies such as [36, 47]. Although training can be significantly slower without SSA, since this is done only for a few epochs, the overall training time is not significantly affected. This fine-tuning step is not required when we use only moderately sparse attention ($\leq 50\%$ sparsity) during training, because the network does not face a drastic distribution shift from the training to the inference time in this case.

3.5 SSA-based Attention Self-Ensembling

Generation of an ensemble of sub-models using SSA is as simple as performing SSA at inference time on the trained model, drawing multiple sample outputs for the same input and taking an aggregation of the predictions. Although this method leverages the same model weights for each sample prediction, SSA draws a random subsampling pattern each time, producing a set of sub-models that only vary in their attention patterns. We use an average of the predicted probabilities of the sub-models for generative or classification tasks, or a mean of the predicted values for regression tasks. Surprisingly, we found that the average predictions of these sub-models can be more robust and generalizable than that of the full model if SSA is performed meticulously (i.e., if the SSA hyperparameters are chosen carefully). This shows that the full model may suffer from over-capacity, and thus overfit the training data. Even at inference time, SSA can uncover lower capacity models which may have more generalizable traits such as prioritizing long-distance dependencies over short-distance ones. Although SSA-based selfensembling works best when the model is trained with SSA, we found that it can work with a model trained with vanilla dense attention as well, often matching or even outperforming the dense model. Also, the fact that an ensemble of sub-models can be as performant as the full model shows that the transformer can be thought of as an ensemble of these sub-models with the attention mechanism aggregating/merging them into a single model. This also gives evidence in favor of the information pathways hypothesis, by showing that sub-models can be formed from a subset of the connectivities, indicating the existence of alternative information pathways in the transformer which can operate independently.

SSA-based attention self-ensembling works best with SSA training, and can often serve as an alternative to fine-tuning or denseattention fallback. In this case SSA is performed both during training and inference. As a result, we have the same distribution of subsampled attention, so the network does not need to readjust to a different distribution at inference time. Also, the SSA inference for

Md Shamim Hussain, Mohammed J. Zaki and Dharmashankar Subramanian

each sub-model can be much less costly and less memory intensive than the full model which uses dense attention. Although we need to draw multiple samples, this process is embarrassingly parallel and can be easily done on separate workers (CPUs/GPUs/nodes) followed by an aggregation step. All sub-models in a self-ensemble share the same set of parameters, so the total number of parameters is the same as that of the full model. There is no added training cost since we train a single model with SSA. This makes it easier to train and deploy the ensemble. As such, attention self-ensemble is a more general concept and can potentially be used with other forms of stochastic subsampling methods (e.g., attention dropout), and also for uncertainty estimation, similar to [18].

4 EXPERIMENTS

We explore the effectiveness of SSA for various tasks involving transformers. We experiment with different types of data and both generative and discriminative tasks, such as generative modeling of text, image generation, image classification and graph regression. Our experiments cover different granularities of input data as well, e.g., for text, we consider both word-level and character-level inputs, for images we consider both pixel-level and patch-level inputs and for graphs we process individual node-level inputs. Also, we explore different scales such as relatively smaller-scale CIFAR-10 [34] image dataset, medium-scale Enwik8 [40] and WikiText-103 [41] text datasets and large scale ImageNet-1K [13] and PCQM4Mv2 [23] molecular graph datasets. We used the PyTorch [45] library for our experiments. The training was done in a distributed manner with mixed-precision computation on up to 4 nodes, each with 8 NVIDIA Tesla V100 GPUs (32GB RAM/GPU), and two 20-core 2.5GHz Intel Xeon CPUs (768GB RAM). More details about the hyperparameters and the training procedure are provided in the Appendix. Our code is available at https://github.com/shamim-hussain/ssa.

4.1 Generative Language Modeling

Our language modeling experiments showcase the application of SSA to generative modeling of text data, and its ability to handle long-range dependencies. We experiment on the WikiText-103 and the Enwik8 datasets. The WikiText-103 [41] dataset contains a diverse collection of English Wikipedia articles with a total of 103 million word-level tokens. This dataset has been extensively used as a long-range language modeling benchmark. The Enwik8 [40] dataset contains the first 100 million bytes of unprocessed Wikipedia text. This dataset has been used as a benchmark for bytelevel text compression. For both these datasets, we used the 16-layer transformer decoder of Press et al. [48] which uses ALiBi relative positional encodings. We used an input length of 3072 tokens for WikiText-103. We made minor changes to the architecture and training procedure (refer to the Appendix), which allow us to train the model much faster on 32 V100 GPUs, within 9 hours, compared to the 48 hours required by Press et al. [48], while still yielding comparable perplexity. We achieve validation and test perplexities of 17.14 and 17.98, with a sliding window inference (overlap length 2048), compared to 16.96 and 17.68 of Press et al. [48] with vanilla dense attention training. We call this S0 (since SSA was used in 0 layers) and use this as a baseline for SSA results. On Enwik8, we get validation and test BPB (bits per byte) of 1.052 and 1.028 with a

| Table 1: Results on language modeling tasks on WikiText-103 and |
|---|
| Enwik8. Red: best model, Violet: good model; C/M/S: normalized |
| Compute/Memory/Speedup; Ppl.: perplexity; BPB: bits per byte; ar- |
| row indicates if higher or lower is better. |

| | WikiText-103 (Gen.) (#Layers=16 #Params=247M) | | Enwik8 (Gen.) (#Layers=16, #Params=202M) | |
|-----------|--|---|---|--|
| Model* | dev/test Ppl.↓ | $C\downarrow$ / $M\downarrow$ / $S\uparrow$ | dev/test BPB↓ | $C \downarrow / M \downarrow / S \uparrow$ |
| S0(Dense) | 17.14 / 17.98 | 1.00 / 1.00 / 1.00 | 1.052 / 1.028 | 1.00 / 1.00 / 1.00 |
| S16-L2 | 17.12 / 17.84 | 0.83 / 0.74 / 1.15 | 1.052 / 1.028 | 0.80 / 0.67 / 1.34 |
| +FT | 16.95 / 17.68 | 0.85 / 0.77 / 1.13 | 1.050 / 1.026 | 0.82 / 0.71 / 1.30 |
| S16-L4 | 17.39 / 18.13 | 0.75 / 0.62 / 1.31 | 1.081 / 1.058 | 0.70 / 0.51 / 1.64 |
| +FT | 16.91 / 17.60 | 0.78 / 0.65 / 1.27 | 1.052 / 1.029 | 0.73 / 0.56 / 1.55 |
| S12-L4 | 17.29 / 17.95 | 0.81 / 0.71 / 1.22 | 1.047 / 1.024 | 0.78 / 0.64 / 1.48 |
| +FT | 17.09 / 17.86 | 0.83 / 0.74 / 1.20 | 1.044 / 1.024 | 0.80 / 0.67 / 1.41 |
| S16-L6 | 17.49 / 18.30 | 0.72 / 0.57 / 1.39 | *S< <i>l</i> >-L< <i>w</i> >: Loca | lly biased SSA |
| +FT | 17.09 / 17.86 | 0.75 / 0.62 / 1.34 | on the last <i>l</i> layers | with w windows |
| S16-L8 | 17.94 / 18.69 | 0.71 / 0.55 / 1.42 | +FT: Finetuned wi | ithout SSA for the |
| +FT | 17.20 / 17.92 | 0.74 / 0.60 / 1.36 | last 10% epochs | |

sliding window inference (overlap length 3072), which we use as the baseline (i.e., S0). We could not find a comparable dense attention implementation; Al-Rfou et al. [1] achieve a test BPB of 1.06 with a very deep 64-layer transformer. A local transformer achieves a test BPB of 1.10, whereas specialized architectures such as [12, 54] use a longer input length to achieve a test BPB of 0.99, which is comparable to ours. We could train only up to an input length of 4096 with dense attention without gradient accumulation/checkpointing, so we experiment with this input length.

Our experiments are designed to show the effectiveness of SSA in reducing training costs and also as a regularization method. We measure training cost in terms of Compute (FLOPs), Memory (GB) and Speed (steps/sec). We normalize these with respect to S0, to better represent comparative gains achieved with SSA (refer to the Appendix for unnormalized values). We primarily show results for locally biased SSA since it produces the best results, and leave the results for unbiased SSA as an ablation study (refer to the Appendix). We use the causal gaussian sampling scheme described in section 3.3. We tune the SSA parameters σ (in Eq. 3) in different layers for the best validation set results. We applied different numbers of windows with locally biased SSA to achieve different levels of sparsity and regularization, both of which increase with the number of windows. For example, with 4 windows we reduce the attention cost 4 times by only sampling 25% of the self-attention matrix. This is denoted with a suffix '-L4' (Locally biased with 4 windows). We mainly apply SSA to all 16 transformer layers (S16), but we found that sometimes better results can be achieved by leaving the first few layers unsampled, at the cost of some efficiency. For example, we use S12 to denote that SSA has been applied only to the last 12 layers. Also, we produced results for the Fine-Tuning (+FT) scheme where we turn off SSA in the last 10% of the training epochs and finetune the model for dense attention, which leads to better results. For additional fine-tuning, we report the total compute, but average speedup and memory consumption over the training epochs.

The results are presented in Table 1. On WikiText-103 we achieve the best result with S16-L4 after fine-tuning. Here, SSA is used in all layers with 4 windows, which corresponds to only 25% of



Figure 6: Validation loss vs training (a) epochs, (b) time and (c) compute for the WikiText-103 experiment, with (red) and without (blue) SSA and with fine-tuning (green) which begins at epoch 100.

attention being sampled during the majority of the training. We achieve a significant improvement over the baseline (S0) due to the regularization effect of SSA, while also achieving a 1.27x speedup in training, 22% reduction in compute and 35% reduction in memory cost. This method also achieves competitive results on Enwik8, but the best result is achieved by S12-L4, where we leave the first 4 layers unsampled. We think this is due to the higher granularity of character-level data, which makes the locally biased SSA algorithm less effective in predicting the attention patterns in the shallower layers. S12-L4 achieves the best result even without fine-tuning and also has 1.48x speedup in training, 22% reduction in compute and 36% reduction in memory cost. Both S16-L2 and S12-L4 achieve good results even without fine-tuning, which shows that the requirement for fine-tuning arises mainly due to highly sparse sampling. We can reduce the training cost further by using sparser subsampling, for example, with S16-L6 or S16-L8 but this comes at the cost of slightly worse results, even after fine-tuning. We believe this is because, at very high sparsity levels, some of the important pathways remain undertrained, which is corrected only slightly by fine-tuning. Also, at this point, other parts of the network become the bottleneck rather than self-attention, which leads to diminishing returns in terms of training cost reduction.

In Fig. 6 we see how training with SSA progresses compared to dense attention training. From Fig. 6(a) we see that the validation loss of S16-L4 closely follows that of S0 for most of the training in terms of the number of steps. This verifies our claim that the information pathways can be trained independently by showing that even when we are sampling a small subset (25%) of the pathways, training progresses naturally. However, in terms of both wall time and compute, the validation loss of S16-L4 drops much faster than S0. The validation loss plateaus at a slightly higher value than that of S0, but with a slight fine-tuning in the end, it falls even below that of S0. Also, even with fine-tuning, training finishes significantly earlier than S0. Thus, compared to dense attention (S0), SSA delivers significant improvements in performance and efficiency.

4.2 Image Generation and Classification

While some previous works only focus on reducing the cost of training only for generative language modeling [36, 47], we show the generality of our method by also applying it to image generation and classification tasks. We target the unconditional sub-pixel level image generation task on CIFAR-10 [34], which contains 60,000 tiny 32x32x3 images from 10 classes. Each image is flattened into a

| Table 2: Image generation results on CIFAR-10 and image classifi- |
|---|
| cation results on ImageNet-1K. BPD: bits per dimension; Acc.: top-1 |
| accuracy. Red: best model, Violet: good model. |

| CIFAR-10 (Gen.) | | ImageNet-1K (Class.) | | | |
|---|--------------|--|--------------|--------------------|--|
| (#Layers=16, #Params=203M) | | (Swin-T, #Layers=12, #Params=28M) | | | |
| Model | BPD↓ | $C \downarrow / M \downarrow / S \uparrow$ | Model* | Acc.↑ | $C {\downarrow} \ / \ M {\downarrow} \ / \ S {\uparrow}$ |
| S0(Dense) | 2.789 | 1.00 / 1.00 / 1.00 | W7-S0(Dense) | 81.19% | 0.90 / 0.70 / 1.14 |
| S16-L4 | 2.796 | 0.75 / 0.53 / 1.25 | W14-S10-L4 | 80.56% | 0.90 / 0.73 / 1.13 |
| +FT | 2.774 | 0.77 / 0.58 / 1.22 | +FT | 81.15% | 0.91 / 0.76 / 1.13 |
| | | W14-S6-L4 | 81.23% | 0.97 / 0.91 / 1.05 | |
| | | +FT | 81.60% | 0.97 / 0.92 / 1.05 | |
| * W < ω >: Window-size of Swin-T = ω | | W14-S0(Dense) | 81.89% | 1.00 / 1.00 / 1.00 | |

sequence of length 3072 and fed to a transformer decoder, which serves as an autoregressive model. We get a validation BPD (bits per dimension) of 2.789 with dense attention training which we denote as the baseline S0. We could not find a comparable dense attention result in the literature, but some specialized architectures such as [10, 54] have reported comparable results. Our results are presented in Table 2 (left). We see that with fine-tuning, SSA achieves a slightly better result than dense training (S0) while achieving 1.22x speedup, saving 23% compute and 42% memory. Without finetuning, it achieves a slightly worse result but almost halves the memory required for training, which is particularly beneficial for high-resolution image generation.

Beyond generative tasks, we also explore the usefulness of SSA for discriminative tasks such as the large-scale image classification task on the ImageNet-1K dataset [13] which contains 1.28 million images from 1000 classes. It is customary to train transformers on image patches for classification. Instead of the vanilla Vision Transformer [15], we use the Swin Transformer [39] because it achieves state-of-the-art results on ImageNet-1K when trained from scratch. Additionally, we aim to investigate SSA's applicability to locally dense attention based architectures such as the Swin Transformer, which uses a shifted window based local attention mechanism enabling efficient handling of smaller patches (e.g., 4x4). We use the Swin-Tiny model with 12 layers and 28 million parameters and an input resolution of 224x224 in our experiments, and report the top-1 accuracy on the validation set. To demonstrate the usefulness of SSA, we use window sizes of 7x7 and 14x14, denoted by W7 and W14 respectively. A larger window uses more attention and achieves better results, but also requires more compute and memory. The results are presented in Table 2 (right). To apply SSA we subdivide each window into 4 sub-windows (L4). With SSA applied to 10 layers (the last two layers have a resolution of 7x7, where further sub-division is not possible), we can train a W14 model with almost the same training cost as a W7 model. However, even with fine-tuning, we cannot achieve better results than W7. Only by excluding the first 4 layers from SSA and fine-tuning, we attain better accuracy than W7. This accuracy is, however, slightly less than that of W14-S0, but we achieve this at a lower training cost. We believe that this is because the shifted window based attention mechanism is inherently more local than global attention, limiting the regularization effect of locally biased SSA. Moreover, attention is no longer the primary bottleneck. Hence, the savings due to SSA are only incremental. However, SSA can still be utilized to trade off

 Table 3: Graph regression results on PCQM4Mv2 dataset. MAE:

 mean absolute error.

| | PCQM4Mv2 (Regr.) (EGT, #Layers=6, #Params=11M) | | | | |
|-----------|---|----------|----------|----------|--|
| | w/c | FT | + | FT | |
| Model* | dev MAE↓ | Compute↓ | dev MAE↓ | Compute↓ | |
| S0(Dense) | 0.0905 | 1.00 | | | |
| S6-U10 | 0.0907 | 0.96 | 0.0895 | 0.97 | * S < <i>l</i> >-U< <i>x</i> >: |
| S6-U20 | 0.0895 | 0.93 | 0.0876 | 0.94 | Unbiased SSA |
| S6-U30 | 0.0904 | 0.89 | 0.0876 | 0.90 | on the last <i>l</i> |
| S6-U40 | 0.0930 | 0.86 | 0.0879 | 0.87 | layers with |
| S6-U50 | 0.0964 | 0.82 | 0.0908 | 0.84 | x% drop |

accuracy for training cost, as evidenced by the 3% compute and 8% memory savings, as well as the 5% speedup over the locally dense model.

4.3 Molecular Graph Regression

We further show the generality of our method by applying SSA to molecular graph data on the PCQM4Mv2 quantum chemical dataset [23]. Also, we wanted to demonstrate its applicability to newly proposed Graph Transformers [26, 44, 64], which use global self-attention. The PCQM4Mv2 dataset contains 3.8 million molecular graphs, and the target task is to predict a continuous valued property, the HOMO-LUMO gap, for each molecule. For this task, we use the Edge-augmented Graph Transformer (EGT) [26]. We experiment with an ablated variant of EGT called EGT-Simple since it approximately achieves the same performance on PCQM4Mv2 while also being simpler to apply SSA to, but for brevity, we will call this model EGT. We experiment on the EGT_{small} model with 11 million parameters and 6 layers, and report the mean absolute error (MAE) on the validation set. We achieve a baseline MAE of 0.0905, as reported in [26] without SSA, which we call S0.

Graphs are fundamentally different from images and text due to their arbitrary topology and do not have a single simplistic notion of locality. To apply locally biased SSA we must partition the graph into equally sized local windows. There are different possible ways of doing it which may also involve the edge features. Further, we need to do locally biased source shuffling on graph nodes. Since this would require substantial further research, we instead show results for unbiased SSA on graphs, which is straightforward to implement as it does not rely on the notion of locality. We apply SSA to all layers (S6) and drop 10%-50% of source nodes randomly during training. For example, we use the suffix '-U20' to denote that 20% of the source nodes are randomly dropped and we sample the remaining 80%. We also report the result after fine-tuning without SSA for the last 10% of the training epochs (+FT). The results are shown in Table 3. We see that the best results (MAE of 0.0876) are achieved for S6-U20 and S6-U30 with fine-tuning which is not only significantly better than the baseline (S0) but also requires around 10% less compute (FLOPs). For this training, we could not tabulate the memory savings and speedup because in our implementation the data-loading of graphs becomes the bottleneck. We believe that the better results achieved by SSA on graphs are due to its regularization effect, which encourages the network to consider

Table 4: Self-ensembling results by locally biased SSA with 4 windows on WikiText-103 and Enwik8, produced with 50 samples for each input segment. Renormalized results are from Table 1.

| | Wikitext-103 dev/test Ppl.↓ | | Enwik8 dev/test BPB↓ | |
|---------------|--------------------------------|--|---|--------------------------------|
| Model | Renorm. | Ensemble | Renorm. | Ensemble |
| S0(Dense) | 17.14 / 17.98 | 16.86 / 17.46 | 1.052 / 1.028 | 1.066 / 1.042 |
| S16-L4 +FT | 17.39 / 18.13 16.91 / 17.60 | 16.75 / 17.42 16.54 / 17.18 | 1.081 / 1.058 1.052 / 1.029 | 1.086 / 1.062 1.058 / 1.035 |
| S12-L4 +FT | 17.29 / 17.95 17.09 / 17.86 | 16.89 / 17.60 16.80 / 17.51 | 1.047 / 1.024 1.044 / 1.024 | 1.050 / 1.029 1.055 / 1.033 |

long-range interactions. However, unlike locally biased SSA, unbiased SSA cannot employ highly sparse attention without incurring a performance penalty, as evident from the results of S6-U50. At 50% sparsity, the important pathways are rarely sampled and remain undertrained. We leave it as a future research direction to explore the use of locally biased SSA on graphs, which we believe will further improve the performance and efficiency of training.

4.4 Self-ensembling Results

Once a transformer has been trained we can apply SSA at inference time, draw multiple sample predictions from the same input and aggregate them. This way the prediction is made by an ensemble of sub-models, sampled by SSA, which we call self-ensembling. The results of an average of 50 prediction samples drawn by locally biased SSA with 4 windows, which samples 25% attention at each prediction instance for language modeling tasks, are shown in Table 4, and they are compared against their full-model counterpart, which we call renormalized results (since the network merges and normalizes the sub-models into a single model). For WikiText-103, we see that the self-ensembling results are significantly better than their renormalized counterparts. This is true even for S0, which was not trained with SSA but with vanilla dense attention. This shows that SSA-based self-ensembling can improve the performance of the model even when it is not trained with SSA. This also shows the existence of sub-models within a dense transformer, trained with dense attention, which is an implication of the information pathway hypothesis. Results are better when the model is trained with SSA and fine-tuning further improves the results. We think the better results are due to the higher generalizability of the constituent submodels which take advantage of the local inductive bias and higher sparsity regularization. For Enwik8, however, the results are close to but not better than the renormalized counterparts. We think this is because it is more difficult to predict important pathways in character-level prediction tasks than in word-level tasks due to the higher granularity of the data. Future work may uncover the important pathways with a higher success rate and thus form better ensembles.

Self-ensembling can be done for unbiased SSA and regression tasks as well. The results of self-ensembling on the PCQM4Mv2 dataset are presented in Table 5. We take an average of 50 sample predictions for each input graph while following the same SSA scheme during inference as during training. We see that the self-ensembling results are better than the renormalized results for all models that have not been fine-tuned. The self-ensembled

Table 5: Self-ensembling results by unbiased SSA on the PCQM4Mv2 dataset, produced with 50 samples for each input graph. Renormalized results are from Table 3.



Figure 7: Self-ensemble performance for language modeling on (a) WikiText-103 and graph regression on (b) PCQM4Mv2 as a function of the number of samples drawn. Dashed lines show the performance of the renormalized model.

results are even better than that of renormalized fine-tuned results. This shows that self-ensembling can serve as an alternative to fine-tuning. We believe that the better results are due to the regularization effect of SSA, sampling sub-models that consider sparse and long-range dependencies. These results degrade with fine-tuning because the pathways within these models become less predictable by unbiased SSA after fine-tuning.

Fig. 7 shows how the self-ensembling performance improves with the number of samples drawn, for the language modeling task on WikiText-103, and the graph regression task on PCQM4Mv2, and how they compare against the renormalized results. We see that the self-ensembling performance improves with the number of samples drawn for both tasks. From Fig. 7 (a) we see that for S0, which was not trained with SSA, we need to draw upwards of 20 samples to improve the results beyond that of renormalization. But for S16-L4 and their fine-tuned counterparts, which were trained with SSA, we need to draw only 2-5 samples to improve the results beyond that of renormalization. Since we are using SSA at inference time, these samples are faster to produce for the sub-models than the full model. This shows that self-ensembling is a practical option for improving the results of a model that was trained with SSA. We believe the important information pathways are more predictably sampled within a model that was trained with SSA, which leads to the result plateauing with fewer samples. However, this rate of improvement also depends on the amount of sparsity applied by SSA. From Fig. 7 (b) we see that for the graph regression task, we also need to draw only 3-5 samples to improve the results beyond that of renormalization, but S6-U10 which applies only 10% attention drop plateaus much faster than S6-U50 which applies 50% drop. This is because variance increases with the amount of sparsity, but this also produces a more diverse set of sub-models, which often leads to better results.

In Fig. 7, we observe that even when we draw only a single random sample, the results are not significantly worse than the renormalized results. It is important to note that if the information pathways were *not* independent, randomly selecting a set of pathways to form a sub-model would lead to a drastic drop in performance. This shows that the information pathways are indeed independent, e.g., the presence/absence of a particular pathway does not negatively affect the performance of another pathway. We hypothesize that for a single random sample, the reduction in performance is only due to the reduced strength of the ensemble due to the missing pathways, which is quickly recovered as we draw more samples by covering most of the important pathways. Also, the fact that a few sub-models drawn from a predefined distribution can be as performant as the full model shows that the distribution of the information pathways is predictable.

5 CONCLUSION AND FUTURE WORK

In this paper, we presented the information pathways hypothesis which states the existence of sparsely connected sub-networks within the transformer called information pathways. A sub-model formed from a random subset of these pathways can be trained at each training step to reduce the cost of training. We introduce an algorithm called SSA which can take advantage of this fact by stochastically sampling only a subset of attention sources and training the important information pathways with a high probability, which not only reduces training cost but also improves generalization. SSA can be applied to any model that uses dense self-attention, and for both generative and discriminative tasks. We showed the effectiveness of SSA for language modeling, image classification, and graph regression tasks. We also showed that SSA can be applied at inference time to form an ensemble of sub-models from the transformer which can further improve the results beyond that of the full model, by making more robust predictions. We used local bias to improve the performance of SSA by sampling the important pathways with a higher probability.

Our SSA algorithm is simple and easy to implement, but its performance can be further improved by using more sophisticated sampling strategies. The information pathways hypothesis calls for more research into the search for sparsely connected sub-networks within the transformer, and how to better sample them, which could further alleviate the training cost of the transformers while helping them to generalize better using strategies such as attention self-ensembling. We also want to explore the prospect of extending SSA to cross-attention, for tasks such as machine translation.

ACKNOWLEDGMENTS

This work was supported by the Rensselaer-IBM AI Research Collaboration, part of the IBM AI Horizons Network.

REFERENCES

 Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. 2019. Character-level language modeling with deeper self-attention. In Proceedings of the AAAI conference on artificial intelligence, Vol. 33. 3159–3166. KDD '23, August 6-10, 2023, Long Beach, CA, USA

Md Shamim Hussain, Mohammed J. Zaki and Dharmashankar Subramanian

- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. arXiv preprint arXiv:1607.06450 (2016).
- [3] Alexei Baevski and Michael Auli. 2018. Adaptive input representations for neural language modeling. arXiv preprint arXiv:1809.10853 (2018).
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014).
- [5] Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The longdocument transformer. arXiv preprint arXiv:2004.05150 (2020).
- [6] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum learning. In Proceedings of the 26th annual international conference on machine learning. 41–48.
- [7] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020. Generative pretraining from pixels. In International Conference on Machine Learning. PMLR, 1691–1703.
- [8] Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020. The lottery ticket hypothesis for pre-trained bert networks. Advances in neural information processing systems 33 (2020), 15834–15846.
- [9] Zhaodong Chen, Yuying Quan, Zheng Qu, Liu Liu, Yufei Ding, and Yuan Xie. 2022. Dynamic N: M Fine-grained Structured Sparse Attention Mechanism. arXiv preprint arXiv:2203.00091 (2022).
- [10] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. arXiv preprint arXiv:1904.10509 (2019).
- [11] Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, and Others. 2020. Rethinking attention with performers. arXiv preprint arXiv:2009.14794 (2020).
- [12] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. 2019. Transformer-xl: Attentive language models beyond a fixedlength context. arXiv preprint arXiv:1901.02860 (2019).
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition. Ieee, 248–255.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018).
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020).
- [16] Angela Fan, Edouard Grave, and Armand Joulin. 2019. Reducing transformer depth on demand with structured dropout. arXiv preprint arXiv:1909.11556 (2019).
- [17] Jonathan Frankle and Michael Carbin. 2018. The lottery ticket hypothesis: Finding sparse, trainable neural networks. arXiv preprint arXiv:1803.03635 (2018).
- [18] Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*. PMLR, 1050–1059.
- [19] Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. Advances in neural information processing systems 28 (2015).
- [20] Lars Kai Hansen and Peter Salamon. 1990. Neural network ensembles. IEEE transactions on pattern analysis and machine intelligence 12, 10 (1990), 993–1001.
- [21] Babak Hassibi and David Stork. 1992. Second order derivatives for network pruning: Optimal brain surgeon. Advances in neural information processing systems 5 (1992).
- [22] Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016).
- [23] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. 2021. Ogb-lsc: A large-scale challenge for machine learning on graphs. arXiv preprint arXiv:2103.09430 (2021).
- [24] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. 2017. Snapshot ensembles: Train 1, get m for free. arXiv preprint arXiv:1704.00109 (2017).
- [25] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. 2016. Deep networks with stochastic depth. In *European conference on computer vision*. Springer, 646–661.
- [26] Md Shamim Hussain, Mohammed J Zaki, and Dharmashankar Subramanian. 2022. Global self-attention as a replacement for graph convolution. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 655–665.
- [27] Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2018. Averaging weights leads to wider optima and better generalization. arXiv preprint arXiv:1803.05407 (2018).
- [28] Armand Joulin, Moustapha Cissé, David Grangier, Hervé Jégou, et al. 2017. Efficient softmax approximation for GPUs. In *International conference on machine learning*. PMLR, 1302–1310.

- [29] Heewoo Jun, Rewon Child, Mark Chen, John Schulman, Aditya Ramesh, Alec Radford, and Ilya Sutskever. 2020. Distribution augmentation for generative modeling. In *International Conference on Machine Learning*. PMLR, 5006–5019.
- [30] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In International Conference on Machine Learning. PMLR, 5156–5165.
- [31] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. arXiv preprint arXiv:1911.00172 (2019).
- [32] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014).
- [33] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. arXiv preprint arXiv:2001.04451 (2020).
- [34] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [35] Yann LeCun, John Denker, and Sara Solla. 1989. Optimal brain damage. Advances in neural information processing systems 2 (1989).
- [36] Conglong Li, Minjia Zhang, and Yuxiong He. 2021. Curriculum learning: A regularization method for efficient and stable billion-scale gpt model pre-training. arXiv preprint arXiv:2108.06084 (2021).
- [37] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. 2016. Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710 (2016).
- [38] Liu Liu, Zheng Qu, Zhaodong Chen, Yufei Ding, and Yuan Xie. 2021. Transformer Acceleration with Dynamic Sparse Attention. arXiv preprint arXiv:2110.11299 (2021).
- [39] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 10012–10022.
- [40] Matt Mahoney. 2011. Large text compression benchmark.
- [41] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. arXiv preprint arXiv:1609.07843 (2016).
- [42] Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? Advances in neural information processing systems 32 (2019).
- [43] 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. arXiv preprint arXiv:1904.01038 (2019).
- [44] Wonpyo Park, Woong-Gi Chang, Donggeon Lee, Juntae Kim, et al. 2022. GRPE: Relative Positional Encoding for Graph Transformer. In *ICLR2022 Machine Learning for Drug Discovery*.
- [45] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems 32 (2019).
- [46] Sai Prasanna, Anna Rogers, and Anna Rumshisky. 2020. When bert plays the lottery, all tickets are winning. arXiv preprint arXiv:2005.00561 (2020).
- [47] Ofir Press, Noah A Smith, and Mike Lewis. 2020. Shortformer: Better language modeling using shorter inputs. arXiv preprint arXiv:2012.15832 (2020).
- [48] Ofir Press, Noah A Smith, and Mike Lewis. 2021. Train short, test long: Attention with linear biases enables input length extrapolation. arXiv preprint arXiv:2108.12409 (2021).
- [49] Zheng Qu, Liu Liu, Fengbin Tu, Zhaodong Chen, Yufei Ding, and Yuan Xie. 2022. DOTA: detect and omit weak attentions for scalable transformer acceleration. In Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems. 14–26.
- [50] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- [51] Jack W Rae, Anna Potapenko, Siddhant M Jayakumar, and Timothy P Lillicrap. 2019. Compressive transformers for long-range sequence modelling. arXiv preprint arXiv:1911.05507 (2019).
- [52] Hongyu Ren, Hanjun Dai, Zihang Dai, Mengjiao Yang, Jure Leskovec, Dale Schuurmans, and Bo Dai. 2021. Combiner: Full attention transformer with sparse computation cost. Advances in Neural Information Processing Systems 34 (2021), 22470–22482.
- [53] Adam Roberts, Colin Raffel, Katherine Lee, Michael Matena, Noam Shazeer, Peter J Liu, Sharan Narang, Wei Li, and Yanqi Zhou. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. (2019).
- [54] Aurko Roy, Mohammad Saffar, Ashish Vaswani, and David Grangier. 2021. Efficient content-based sparse attention with routing transformers. *Transactions of* the Association for Computational Linguistics 9 (2021), 53-68.
- [55] Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. 2021. Linear transformers are secretly fast weight memory systems. arXiv preprint arXiv:2102.11174 (2021).
- [56] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. arXiv preprint arXiv:1803.02155 (2018).
- [57] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.

- [58] 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. 5998–6008.
- [59] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. 2013. Regularization of neural networks using dropconnect. In *International conference on machine learning*. PMLR, 1058–1066.
- [60] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. arXiv preprint arXiv:2006.04768 (2020).
- [61] Kan Wu, Houwen Peng, Minghao Chen, Jianlong Fu, and Hongyang Chao. 2021. Rethinking and improving relative position encoding for vision transformer. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 10033– 10041.
- [62] Wenhan Xiong, Barlas Oğuz, Anchit Gupta, Xilun Chen, Diana Liskovich, Omer Levy, Wen-tau Yih, and Yashar Mehdad. 2021. Simple Local Attentions Remain Competitive for Long-Context Tasks. arXiv preprint arXiv:2112.07210 (2021).
- [63] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. 2021. Nyströmformer: A nyström-based algorithm for approximating self-attention. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 35. 14138–14148.
- [64] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. 2021. Do Transformers Really Perform Bad for Graph Representation? arXiv preprint arXiv:2106.05234 (2021).
- [65] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. Advances in Neural Information Processing Systems 33 (2020), 17283–17297.
- [66] Lin Zehui, Pengfei Liu, Luyao Huang, Junkun Chen, Xipeng Qiu, and Xuanjing Huang. 2019. DropAttention: a regularization method for fully-connected selfattention networks. arXiv preprint arXiv:1907.11065 (2019).
- [67] Wangchunshu Zhou, Tao Ge, Ke Xu, Furu Wei, and Ming Zhou. 2020. Scheduled drophead: A regularization method for transformer models. arXiv preprint arXiv:2004.13342 (2020).

A DATA AND CODE AVALABILITY

Data: All datasets used in this work are publicly available. The dataset sources are listed in Table 6.

Code: The code is available at https://github.com/shamim-hussain/ssa.

| Dataset | Source |
|--------------|---|
| WikiText-103 | https://huggingface.co/datasets/wikitext |
| Enwik8 | http://mattmahoney.net/dc/textdata.html |
| CIFAR-10 | https://www.cs.toronto.edu/~kriz/cifar.html |
| ImageNet-1k | https://image-net.org/ |
| PCQM4Mv2 | https://ogb.stanford.edu |

Table 6: Dataset sources.

B HYPERPARAMETERS AND TRAINING DETAILS

B.1 Generative Language Modeling

For language modeling on both WikiText-103 and Enwik8 datasets, we used the 16-layer transformer decoder of Press et al. [48] which uses ALiBi relative positional encodings. We used the fairseq toolkit [43] to perform these experiments. We used an input length of 3072 tokens for WikiText-103. Adaptive input embeddings [3] and adaptive softmax [28] output were used to handle a large vocabulary of size around 260K. For Enwik8, we used a simple vector embedding and a vanilla softmax output layer. We used the same architecture and hyperparameters as Press et al. [48], except that we changed the activation function from ReLU to GELU [22]. For Enwik8, we also add a final Layer Normalization [2] layer before the softmax

layer. On WikiText-103 we trained for 64,000 steps (16,000 linear learning rate warmup steps, followed by 48,000 steps of cosine decay) with the Adam [32] optimizer, a maximum learning rate of 0.001 and an increased batch size of 64. For Enwik8, we trained for 10,000 steps (4,000 warmup steps followed by linear decay) with a maximum learning rate of 0.001 and a minimum learning rate of 0.005. Again we used a batch size of 64 and the Adam optimizer.

We tune the SSA parameter σ (in Eq. 3) in different layers for the best validation set results. We express its value as a fraction of the input length so that it is independent of the input length. For WikiText-103 we use a value of $\sigma = 0.2$ in the first layer and linearly increase it to $\sigma = 0.35$ in the deepest layer. For Enwik8 we start with $\sigma = 0.1$ and linearly increase it to $\sigma = 0.225$ in the deepest layer.

B.2 Image Generation and Classification

For image generation on CIFAR-10, we use a 16-layer transformer with similar architectural and hyperparameter settings as in the previous section, but we use the 2D relative position bias for positional encoding, similar to [39]. Instead of using dropout regularization, we use the augmentation techniques described in [29] to achieve better generalization. We use the Adam optimizer with a maximum learning rate of 0.002 and a batch size of 128. We train for 100,000 steps in total, with an initial learning rate warmup of over 4,000 steps, followed by cosine decay. To perform locally biased SSA on this 2D data we divide the image vertically into 4 windows of 8 rows. Locally biased source shuffling is performed in both the horizontal and vertical directions while preserving causality at the window level. We set the SSA parameter σ to 0.25 times dimensions (length/width), in both the horizontal and vertical directions, and all layers.

For ImageNet-1K classification, we use the Swin-Tiny model with 12 layers and 28 million parameters and an input resolution of 224x224. We use the same architecture, hyperparameters, augmentation, and training scheme as in [39]. To apply locally biased SSA within the 14x14 windows, we further subdivided the windows into 4, 7x7 sub-windows. Locally biased source shuffling was performed in both horizontal and vertical directions, with the value of σ as 0.75 times the window size (i.e., 14), but we further ensured that sources were constrained within their own 14x14 (shifted) windows after shuffling (yet they can move beyond the smaller 7x7 sub-windows).

B.3 Molecular Graph Regression

For graph regression on the PCQM4Mv2 dataset, we use the Edgeaugmented Graph Transformer (EGT) described in [26]. EGT uses additional channels to represent and update edge embeddings, which makes it slightly different from the standard transformer. However, we experiment with an ablated variant of EGT called EGT-Simple which reduces the edge representations to relative positional encodings. However, for brevity, we call this model EGT. We experiment on the EGT_{small} model with 11 million parameters and 6 layers. We use the same hyperparameters and training scheme as in [26], except, we do not use attention dropout in these models because SSA works as a regularization method. KDD '23, August 6-10, 2023, Long Beach, CA, USA

Table 7: Baseline (S0) training cost for different datasets.

| Dataset | Model | Compute (Exa FLOP) | Memory/GPU (GB) | Time/step (ms) |
|--------------|----------------------|-----------------------|--------------------|-------------------|
| WikiText-103 | Transfo. Decoder | 7.6 | 21.3 | 492 |
| Enwik8 | Transfo. Decoder | 1.7 | 30.5 | 628 |
| CIFAR-10 | Transfo. Decoder | 23.8 | 29.1 | 1059 |
| ImageNet-1k | Swin-Tiny | 3.2 | 4.0 | 146 |
| PCQM4Mv2 | EGT _{small} | 0.5 | - | - |

C BASELINE TRAINING COST

In our experiments, we normalized the training costs with respect to the baseline model S0, the dense attention model without SSA. However, we also report the baseline training costs in terms of absolute values in Table 7 for completeness. Note that for the PCQM4Mv2 dataset, we could not faithfully compute the memory consumption and the training time due to the data loading bottleneck. However, we can still compare the cost of the baseline models with SSA models in terms of compute.

D LOCALLY BIASED VS UNBIASED SSA

In the results presented in our experiments, we claimed that local bias was an important ingredient in improving the performance of SSA. Here, we directly compare the results for locally biased SSA and unbiased SSA for the same level of sparsity and when they are applied to the same subset of layers.

Table 8: Locally biased vs unbiased SSA results for language modeling tasks on WikiText-103 and Enwik8. Locally biased SSA results are from Table 1.

| % Attention | Wikitext-103 | | Enwi | k8 |
|---------------|----------------|---------------|----------------|---------------|
| | dev/test Ppl.↓ | | dev/test | BPB↓ |
| Sampled | Locally Biased | Unbiased | Locally biased | Unbiased |
| 50% attention | 17.12 / 17.84 | 17.45 / 18.18 | 1.052 / 1.028 | 1.114 / 1.087 |
| +FT | 16.95 / 17.68 | 16.96 / 17.69 | 1.050 / 1.026 | 1.063 / 1.042 |
| 25% attention | 17.39 / 18.13 | 19.33 / 20.15 | 1.081 / 1.058 | 1.329 / 1.287 |
| +FT | 16.91 / 17.60 | 17.89 / 18.69 | 1.052 / 1.029 | 1.100 / 1.075 |

The results for language modeling are presented in Table 8 where the same level of sparsity is applied to all layers for both types of SSA (using 2 or 4 window attention for 50% or 25% sampling, respectively). We see that unbiased SSA performs slightly worse than locally biased SSA for 50% sampling of attention. This gap can be reduced with fine-tuning. However, when we only sample 25% of attention, training is significantly hampered for unbiased SSA, and the results cannot be made comparable to locally biased SSA, even with fine-tuning. This is because unbiased SSA cannot sample the important pathways with a high enough probability for the training to progress gracefully. This shows the necessity of local bias for sampling at high sparsity levels.

The results for image classification are presented in Table 9 where the same level of sparsity (25% sampled, 75% dropped) is applied to 10 layers (using 4 sub-windows) or the first 4 layers are excluded. In all cases, we see that locally biased SSA performs better than unbiased SSA, but we do get good results with unbiased SSA when we exclude the first 4 layers. This shows that local bias is important for SSA to work well, but it is less important in deeper layers than in the shallower layers. In deeper layers, the model tends to form long-distance dependencies, which are more predictable by unbiased SSA. This is why we see that unbiased SSA performs better when we exclude the first 4 layers.

Table 9: Locally biased vs unbiased SSA results on the ImageNet-1K image classification task. Locally biased SSA results are presented from Table 2.

| # Layers | dev Acc. ↑ | | |
|-----------|----------------|----------|--|
| Sampled | Locally Biased | Unbiased | |
| 10 layers | 80.56% | 80.21% | |
| +FT | 81.15% | 80.80% | |
| 6 layers | 81.23% | 81.21% | |
| +FT | 81.60% | 81.36% | |

Table 10: Additional self-ensembling results for language modeling tasks on WikiText-103, produced with 50 samples per input segment. Renormalized results are from Table 1.

| | dev/test Ppl.↓ | |
|--------|----------------|---------------|
| Model | Renorm. | Ensemble |
| S16-L6 | 17.49 / 18.30 | 17.01 / 17.80 |
| +FT | 17.09 / 17.86 | 16.83 / 17.53 |
| S12-L8 | 17.94 / 18.69 | 17.41 / 18.10 |
| +FT | 17.20 / 17.92 | 17.04 / 17.69 |

E ADDITIONAL SELF-ENSEMBLING RESULTS

We present additional self-ensembling results for language modeling on WikiText-103 in Table 10, for higher levels of sparsity – with 6 and 8 windows we sample as little as 16.7% and 12.5% attention, respectively. We see similar results as presented in the main section with self-ensembling significantly improving over their renormalization counterparts. This shows that self-ensembling can improve performance even with a high level of sparsity.