
Efficient Representation Learning via Adaptive Context Pooling

Chen Huang¹ Walter Talbott¹ Navdeep Jaitly¹ Josh Susskind¹

Abstract

Self-attention mechanisms model long-range context by using pairwise attention between all input tokens. In doing so, they assume a fixed attention granularity defined by the individual tokens (*e.g.*, text characters or image pixels), which may not be optimal for modeling complex dependencies at higher levels. In this paper, we propose *ContextPool* to address this problem by adapting the attention granularity for each token. Inspired by the success of ConvNets that are combined with pooling to capture long-range dependencies, we learn to pool neighboring features for each token before computing attention in a given attention layer. The pooling weights and support size are adaptively determined, allowing the pooled features to encode meaningful context with varying scale. We show that *ContextPool* makes attention models more expressive, achieving strong performance often with fewer layers and thus significantly reduced cost. Experiments validate that our *ContextPool* module, when plugged into transformer models, matches or surpasses state-of-the-art performance using less compute on several language and image benchmarks, outperforms recent works with learned context sizes or sparse attention patterns, and is also applicable to ConvNets for efficient feature learning.

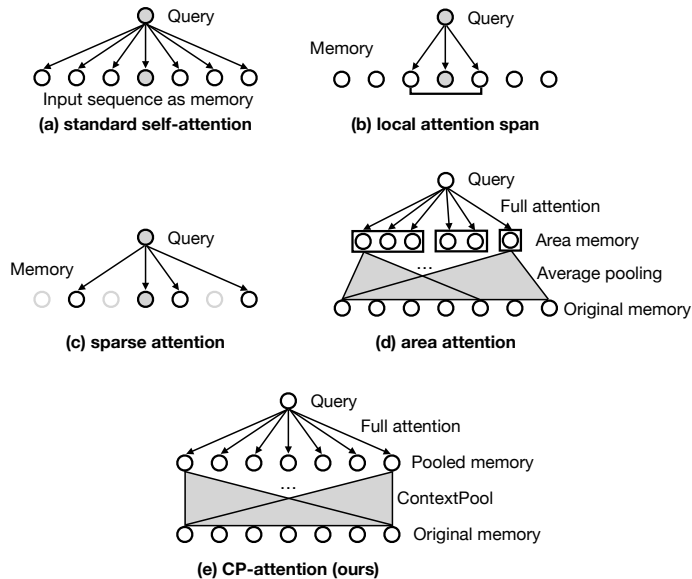


Figure 1. Comparing transformers with (a) standard self-attention (Vaswani et al., 2017), (b-c) efficient attention mechanisms with localized (Yang et al., 2018) or other sparsity patterns (Li et al., 2019a) that lose the full-attention capacity, and (d) area attention (Li et al., 2019b) that maintains an extra memory formed by average pooling with a predefined set of pool sizes. (e) Our *ContextPool* learns to pool with adaptive weighting and support size for each token in-place, before computing full attention.

1. Introduction

Transformers (Vaswani et al., 2017) have achieved great success in the domains of natural language processing (NLP) (Devlin et al., 2019) and computer vision (Dosovitskiy et al., 2021). These models benefit from the self-attention mechanism, which computes correlations between all pairs of tokens in an input sequence. Self-attention enables transformers to capture long-range context, which is

¹Apple Inc., Cupertino, United States. Correspondence to: Chen Huang <chen-huang@apple.com>.

important in both language and vision tasks.

However, each attention layer uses pairwise relationships between individual tokens (*e.g.*, text characters and image pixels), which implies a fixed granularity for attention. This ignores the context around each token, which can vary substantially in scale in the vision and language domains, *e.g.*, from characters to words and from phrases to sentence. Therefore, self-attention with fixed granularity can be fundamentally limited for modeling complex distribution of contextual dependencies, and several layers of self-attention might be needed to make up for this fixed granularity.

Recent vision transformers such as Swin transformer (Liu et al., 2021) and PVT (Wang et al., 2021) adopt a hierarchical architecture to compute self-attention at various scales. However, such attention scale or granularity is pre-

determined rather than learned. Similarly, Li *et al.* (2019b) proposed to use a predefined set of pooling sizes to form a multi-scale ‘area memory’, which accounts for varying context range but in fixed architecture. In BP-Transformer (Ye *et al.*, 2019), a fine-to-coarse attention is computed from multi-scale attention spans via automatic binary partitioning, but the resulting local span sequences might still hurt the capacity of *full* attention.

In this paper we propose *ContextPool*, a drop-in and low-cost module for both the transformer and convolutional networks (ConvNets) to enhance their capacity to model long-range context with dynamic scales, and hence to facilitate efficient representation learning. The idea behind ContextPool is in general inspired by ConvNets, which have local receptive fields and pooling operations. Here we similarly learn to pool neighboring features for each token at every attention layer before computing full attention in transformers. Importantly, the pooling weights and support size are input-adaptive. This allows the pooled features to encode meaningful context with dynamic scale. As a result, self-attention among pooled features can explicitly capture high-level dependencies between contexts.

We show our simple ContextPool makes attention models more expressive, achieving strong performance often with fewer layers. This leads to significantly reduced cost without much sacrifice in accuracy. On the other hand, when we can maintain the same level of compute cost, ContextPool consistently improves performance as it can model longer range of context. When compared to recent transformer models that reduce cost by sparsifying the attention matrix (Yang *et al.*, 2018; Sukhbaatar *et al.*, 2019; Child *et al.*, 2019; Ainslie *et al.*, 2020), our ContextPool method preserves the full attention capability (see comparison in Fig. 1) and can be considered orthogonal to those efficient techniques.

Experiments show that our ContextPool module significantly improves transformer models in terms of performance-cost trade-off, matching or surpassing state-of-the-art performance with less compute on several language and image benchmarks. ContextPool also outperforms recent works with adaptive context size or sparse attention, and is applicable to ConvNets for efficient representation learning. To summarize, our main **contributions** are:

- We introduce ContextPool to encode varying-sized context for each token in an attention layer, giving rise to self-attention with adaptive granularity to model high-level dependencies.
- We show ContextPool-based transformers achieve competitive performance with much less compute on several language and image benchmarks, and outperform prior works with adaptive context size or sparse attention patterns.

- ContextPool is applicable to ConvNets with strong image recognition performance, showing its promise to be a generic module for efficient representation learning.

2. Related Work

Context in Transformers is captured by the attention mechanism between all pairs of tokens from the entire input sequence. When the network goes deeper, high-level contextual dependencies emerge. However, full attention scales quadratically with the sequence length as existing attention models are trained to attend to individual tokens with a fixed granularity, *e.g.*, text characters and image pixels. Hence the vanilla Transformer (Vaswani *et al.*, 2017) is prohibitive for learning long sequences such as long documents or high-resolution images (modeled as long sequences of image patches).

Recent works build on the hierarchical architecture to improve the capability of long-range context modeling. In the vision domain for example, hierarchical transformers such as Swin transformer (Liu *et al.*, 2021), PVT (Wang *et al.*, 2021) and ViL (Zhang *et al.*, 2021) rely on predefined image pyramids to compute self-attention at multiple scales, and can thus model long sequences of image patches at a much higher resolution. However, both the scaling scheme and effective attention granularities remain fixed in these methods. In a similar spirit, ‘area attention’ (Li *et al.*, 2019b) computes multi-scale attention which is generic for both language and vision tasks. Specifically, attention is computed against a multi-scale memory formed by pooling the original memory with predetermined pool sizes. This not only requires larger memory but also does not adapt the context range based on content. Finally, the BP-Transformer (Ye *et al.*, 2019) computes attention using multi-scale attention spans that encode fine-to-coarse contexts, but it imposes a sparsity prior on the attention mechanism, which is adaptive, but which might hurt its capacity.

Efficient Transformers mostly use sparsity or low-rank assumptions on the attention matrix to reduce cost. For sparse attention methods, one can sparsify the attention matrix with predefined patterns like local window (Yang *et al.*, 2018; Sukhbaatar *et al.*, 2019; Child *et al.*, 2019), blockwise (Qiu *et al.*, 2020), log-sparse (Li *et al.*, 2019a) or axial (Ho *et al.*, 2019) patterns and their combinations (Beltagy *et al.*, 2020; Ainslie *et al.*, 2020; Zaheer *et al.*, 2020). The sparsity patterns can also be learned as in (Kitaev *et al.*, 2020; Roy *et al.*, 2022; Tay *et al.*, 2020). These sparse attention methods, despite their sub-quadratic cost, often have a reduced model capacity because each token can only attend to a subset of tokens. Generally, sparse attention needs more layers to model full contextual dependencies in a long sequence (Child *et al.*, 2019). Another family of efficient transformers approximate the attention matrix using low-rank projections (Wang *et al.*,

2020) or feature maps of particular kernels (Katharopoulos et al., 2020). Such low-rank methods preserve the full attention capability with low computational cost, but suffer from the lossy approximations of potentially full-rank attention. We depart from the above-mentioned methods, aiming for efficient, full attention without sparse or low-rank approximations. Nevertheless, our ContextPool module can be embedded within the internals of several of these models.

There are some recent attempts to accelerate transformers by directly reducing the number of tokens to process in attention layers. Ryoo *et al.* (2021) proposed to ‘tokenize’ the input images by aggregating their feature maps into a few tokens, while DynamicViT (Rao et al., 2021) relies on an extra neural network to prune tokens for a fully trained ViT (Dosovitskiy et al., 2021). We provide a novel perspective for parameter-efficient self-attention given any amount of tokens. By learning to pool the token features with adaptive weighting and pool size, we obtain more expressive tokens from fewer layers.

Context in ConvNets is efficiently captured by convolutions, which summarize local neighborhoods with shared weights and when combined with pooling, can model long-term dependencies. Recent works indicate that ConvNets benefit from using different kernel sizes at different convolutional layers (Pintea et al., 2021). Therefore, many methods choose to learn adaptive kernel size to account for data-dependent context or receptive field. Concretely, they scale kernels by dilation and learn dilation factors over shifted Delta-dirac functions (Dai et al., 2017), scalable Gaussian functions (Shelhamer et al., 2019) or Gaussian derivative filters (Pintea et al., 2021; Tomen et al., 2021). Another method of receptive field learning in ConvNets is based on learning pooling functions with adaptive pooling regions (Coates & Ng, 2011; Jia et al., 2012). Our ContextPool method is also applicable to ConvNets. By learning dynamic pooling weights and support size, it is shown to be competitive with existing methods while maintaining low computational cost.

3. ContextPool for Transformers

3.1. Standard Transformers

A standard transformer model (Vaswani et al., 2017) is a chain of self-attention modules (self-attention plus feed-forward layers). The input of each self-attention layer is a feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ from the preceding layer. \mathbf{X} is a sequence of n tokens $=\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ each of dimension d . The attention layer operates on all the token features in \mathbf{X} . Specifically, each token \mathbf{x}_i is first transformed to the query $\mathbf{q}_i = \mathbf{W}^q \mathbf{x}_i$, key $\mathbf{k}_i = \mathbf{W}^k \mathbf{x}_i$ and value $\mathbf{v}_i = \mathbf{W}^v \mathbf{x}_i$ with learned projection matrices $\{\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v\} \in \mathbb{R}^{d \times d}$. Then the attention score of one query \mathbf{q} attending to all the

keys $\{\mathbf{k}_i\}$ stored in a memory is given by:

$$a_i = \frac{\exp(\mathbf{q}^T \mathbf{k}_i)}{\sum_{j=1}^n \exp(\mathbf{q}^T \mathbf{k}_j)}. \quad (1)$$

The final output \mathbf{o}_q from querying the memory with \mathbf{q} is obtained by taking a weighted average of all the values $\{\mathbf{v}_i\}$ in memory:

$$\mathbf{o}_q = \sum_{i=1}^n a_i \mathbf{v}_i. \quad (2)$$

In practice, multi-head self-attention is used in transformers, where multiple projections are learned to compute attention within different heads. The outputs are then concatenated and projected into refined token features.

Drawback The above self-attention mechanism assumes a fixed granularity over which to construct the query and key vectors for individual tokens. However, such a fixed granularity may be sub-optimal for modeling context with different scales. Consider neural machine translation with word-based tokens – translating numerals from one language to another requires little context, but translating ambiguous pronouns (*e.g.*, “it”) requires long range-contextual cues from neighboring tokens. One might argue that this difficulty can be resolved by using deeper models, where self-attention in deeper layers can capture the interactions between single tokens and bake them into deep feature representations. This can progressively correct for the fixed attention granularity at lower layers, but requires more computation that may be avoidable with an adaptive strategy.

3.2. Adaptive Context Pooling

Motivation We motivate our method using Fig. 2(a). In language modelling, if we can piece words together to form phrases, we can gradually capture the phrasal patterns and useful context information. This helps to disambiguate the pronoun “it” by linking it to the phrase “The couch”. Similarly, in image understanding, pooling similar image patches can enable the model to learn semantics of a bird’s body parts. To account for the special role of context in obtaining adaptive attention granularity, we introduce an explicit way of learning context-aware token features. We do this by learning to pool neighboring features for each token (*ContextPool*). Self-attention between such pooled features can thus be context-aware and model high-level dependencies, without requiring multiple self-attention layers.

Therefore, our ContextPool method needs an input-adaptive pooling function. Below, we describe how to learn that with adaptive pooling weights and pooling size. Specifically, given the input token feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, we pool for each token $\mathbf{x}_i \in \mathbf{X}$ with learned weights $\mathbf{w} \in \mathbb{R}^{n \times 1}$ and a Gaussian mask $\mathbf{g}^i \in \mathbb{R}^{n \times 1}$ (acting as a soft, local pooling window), generating a contextual feature matrix $\mathbf{Y} \in \mathbb{R}^{n \times d}$ of the same size of \mathbf{X} (see Fig. 2(b)).

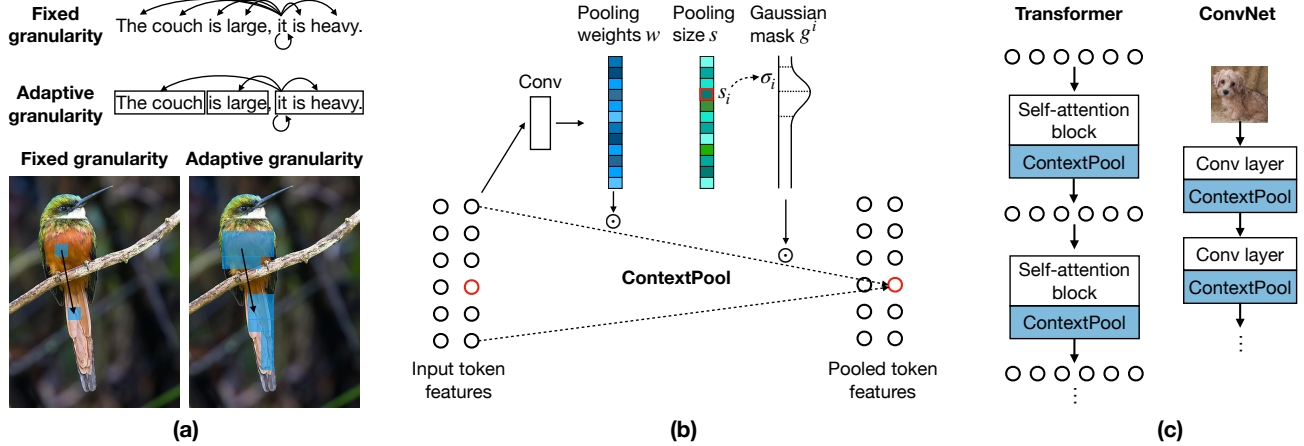


Figure 2. (a) Motivation: the proposed ContextPool seeks to achieve adaptive attention granularity through adaptive context pooling around each token and then computing context-wise attention. This helps to capture high-level dependencies and is useful to model ambiguous pronoun “it” by associating with neighboring phrases rather than single words, or to model interactions between varying-sized object parts. (b) For adaptive ContextPool, we learn the pooling weights and support size dynamically for each token. (c) Our ContextPool module is applicable to both transformers and ConvNets for efficient feature learning. For transformers, the ContextPool module is placed after each attention block, whose output token features are pooled to the same number of features for use in the next attention block. While for ConvNets, ContextPool replaces the conventional pooling function (please refer to supplementary materials for details).

Adaptive pooling weights differ from the uniform ones in the popular average pooling function. We found it helpful to reweight the neighboring token features $\{\mathbf{x}_j\}$ during pooling based on their contextual support to \mathbf{x}_i . One widely used approach of measuring such support is based on nonlocal feature similarity as in (Wang et al., 2018):

$$w_j = \frac{\exp(\theta(\mathbf{x}_i)^T \phi(\mathbf{x}_j))}{\sum_{j=1}^n \exp(\theta(\mathbf{x}_i)^T \phi(\mathbf{x}_j))}, \quad (3)$$

where $\theta(\mathbf{x}_i) = \mathbf{W}^\theta \mathbf{x}_i$ and $\phi(\mathbf{x}_j) = \mathbf{W}^\phi \mathbf{x}_j$ are embeddings with learnable projections $\{\mathbf{W}^\theta, \mathbf{W}^\phi\} \in \mathbb{R}^{d \times d}$.

We dub such learned pooling weights w as nonlocal weights (**NL weights**). The intuition behind NL weights is that similar features in the context are likely to correspond to semantically related entities. Therefore, nonlocal similarity pooling in form of $\sum_{i=1}^n w_i \mathbf{x}_i$ can provide contextual information to increase (or decrease) the probability of a semantic region or segment. Note we only introduce NL weights as a comparing baseline.

One limitation of NL weights is that each weight w_j in Eq. (3) only depends on a feature pair $(\mathbf{x}_i, \mathbf{x}_j)$, overlooking the potential contributions from other features to \mathbf{x}_i . Here we turn to learning w_j by a mapping function $m(\cdot)$ conditioned on all the token features $\{\mathbf{x}_i\}$ in \mathbf{X} . In fact, we predict the pooling weights $w = m(\mathbf{X})$ all at once, where $m(\cdot)$ is implemented as two convolutional layers. Hence the prediction of w is collaborative and more efficient than NL weights prediction.

Adaptive pooling size Pooling with adaptive weights, how-

ever, does not take into account the location relationships between tokens. Here we introduce a locality prior to bias pooling towards the local context around considered token. Note that learning the pooling weights alone might also be able to find local patterns in the learned weights. However, the locality prior can simplify learning by allowing factorized and independent predictions of pooling weights and scope. Our experiments support this hypothesis with favorable results. The locality prior also shares a similar high-level idea with the effective receptive field (Luo et al., 2016), which is shown to have a Gaussian distribution.

We learn a Gaussian mask for each token to implement soft, localized pooling with *adaptive* pooling size rather than a hand-picked one. Specifically, we learn the mapping function $m(\cdot)$ to predict both the pooling weights $w \in \mathbb{R}^{n \times 1}$ and sizes $s \in \mathbb{R}^{n \times 1}$ for n input tokens conditioned on their features \mathbf{X} , i.e., $\{w, s\} = m(\mathbf{X})$. We implement $m(\cdot)$ again by two convolutional layers, but with the channel size set to 2 now. This enables generating the vectors of w and s altogether, which are normalized by a softmax function for ease of training. Given the normalized pooling size $s_i \in [0, 1]$, we then transform it to the standard deviation $\sigma_i = rn \cdot s_i$ of a Gaussian mask $g^i \sim \mathcal{N}(i, \sigma_i^2)$. Here r is a scalar empirically set as 0.1.

By multiplying the learned pooling weights w with the Gaussian mask g^i for token \mathbf{x}_i , we arrive at our final ContextPool function:

$$\mathbf{y}_i = f_{ave}(\mathbf{X} \odot \gamma(w) \odot \gamma(g^i)) = \sum_{j=1}^n \mathbf{x}_j \cdot w_j \cdot g_j^i, \quad (4)$$

where $y_i \in Y$ denotes the ContextPooled features, f_{ave} denotes average pooling function, $\gamma(\cdot)$ is a broadcasting function for element-wise multiplication \odot . We set the normalization factor as $C(\mathbf{X}) = \sum_j w_j \cdot g_j^i$.

As shown in Fig. 2(c), our ContextPool module can be placed after different attention blocks. After each attention block, we take its outputs as input token features and pool them to the same number of features for use in the next attention block. During training, we jointly learn the main model and ContextPool parameters. We also show the applicability of our ContextPool method to ConvNets in supplementary materials.

4. Results

We evaluate the proposed ContextPool (dubbed ‘‘CP’’ as a prefix) module mainly in the transformer architecture to show how strengthened context modeling can benefit self-attention in a parameter-efficient way. We validate such benefits on both language and vision tasks that require a good context modeling capability. Supplementary materials also show that our ContextPool can be seamlessly integrated into ConvNets in place of the conventional pooling function. ContextPool leads to strong results on standard image classification benchmarks, being competitive or even better than those ConvNets with adaptive kernel size or receptive field. This comes at low computational overhead, showing the potential of ContextPool to be a generic module for efficient representation learning.

4.1. Tasks, Datasets and Implementation

Neural Machine Translation For language tasks, we first experiment on the token-level Neural Machine Translation (NMT) task. We use both the WMT 2014 English-to-German (EN-DE) dataset with about 4.5 million English-German sentence pairs, and the and English-French (EN-FR) dataset with about 36 million English-French sentence pairs. A token is a byte pair or a word piece as in (Vaswani et al., 2017). We compare with different methods all using three transformer architectures as defined in (Li et al., 2019b): Small (2 layers), Base and Big (6 layers) models. For our method, we insert ContextPool after every attention layer. Following (Li et al., 2019b), we train for 250k iterations for Small and Base models, and for 600k iterations with a smaller batch size for the Big model due to the memory constraint. We use Adam optimizer with the same learning rate schedule in (Vaswani et al., 2017).

Autoregressive Language Modeling We also evaluate ContextPool on the autoregressive language modeling task at character level. Compared to the token-level task, character-level task is harder due to much longer sequences, which would hypothetically benefit more from stronger con-

text modeling. We use enwik8 and text8 datasets, each with 100M characters and 90M/5M/5M for train/dev/test as in (Mahoney, 2009). For testing, we follow (Beltagy et al., 2020) to split the dataset into overlapping sequences of length 32k with step size 512, and then calculate the Bits Per Character (BPC) of predicting 512 characters from previous 32k.

We use the same 12-layer model architecture with Longformer (Beltagy et al., 2020). We train our models in 3 stages with increasing sequence lengths (2048, 4096, 8192) and different batch sizes (32, 32, 16). All models are trained for a total of 530k steps with linear learning rate warmup. We also use dropout rates 0.2 and weight decays 0.01.

Image Classification We benchmark different transformer models on the widely used ImageNet-1K classification dataset (Deng et al., 2009). There are 1.28M training and 50k validation images from 1k classes. The top-1 accuracy on a single crop is reported. We consider the regular training setting in (Touvron et al., 2021) where no external training data are used. The input image resolution is 224^2 by default. For higher resolutions like 384^2 , we fine-tune the 224^2 trained models. We train for 300 epochs with the AdamW optimizer, using a cosine decay learning rate scheduler and linear warm-up (20 epochs). When fine-tuning on higher resolution images, we tune for 30 epochs with a similar training recipe as in (Liu et al., 2021). We have batch size 1024, initial learning rate 0.001, weight decay 0.05, and the max norm of gradient clipping 1. Stronger data augmentation is found to benefit our ContextPool method. Therefore we use a larger degree of augmentation with the augmentation techniques in (Touvron et al., 2021) such as RandAugment (Cubuk et al., 2020), making our pooled token features more robust.

4.2. Ablations and Comparisons

Ablation on adaptive pooling weights and size We start with ablation studies on these two core components of our ContextPool method and compare against their alternatives. For this purpose, both the NMT and image classification tasks are considered for a comprehensive comparison. For NMT, we choose the English-German (EN-DE) translation task using the Base model. While for image classification, the ViT-B/16 model (Dosovitskiy et al., 2021) (the ‘‘Base’’ variant with 16×16 input patch size) is used.

Tables 1 and 2 summarize the results. We observe that our ContextPool method can consistently improve the baseline transformers at only marginal overhead (in memory, FLOPs and speed), due to the efficiency of adaptive pooling functions implemented by convolutions. For the learning of pooling weights, we first compare with those unnormalized weights without using softmax (middle cell). We obtained slightly worse results for both tasks using

Table 1. Ablations on token-level translation (EN-DE task) using the Base model. Speed (steps / s) is measured on a V100 GPU. CP denotes the use of our full ContextPool module ($w \odot g^i$). The middle and bottom cells compare with alternative weightings and locality priors respectively for context pooling.

Method	Memory (G)	Speed	BLEU \uparrow
Base	17.2	1.20	28.16
CP-Base ($w \odot g^i$)	17.6	1.12	28.91
Unnormalized weights $\odot g^i$	17.6	1.13	28.79
Uniform weights $\odot g^i$	17.4	1.16	28.52
NL weights $\odot g^i$	21.3	0.84	28.66
No locality prior $\odot w$	17.4	1.15	28.31
Fixed window $\odot w$	17.4	1.15	28.55
Adaptive window $\odot w$	17.6	1.12	28.74
Random sparse $\odot w$	17.4	1.15	28.14

un-normalized weights, which confirms the need of normalization for effective weighting (note the pooling size predictions were always softmax normalized). One straightforward alternative to our learned weighting is the use of uniform weights, *i.e.*, to perform average pooling. By doing so, we save the learning cost for the weights but suffer from apparent performance loss. We can also choose to learn NL weights as in Eq. (3), which is equivalent to learning extra, single-head self-attention weights in transformers and is thus much more costly than our lightweight convolutional method. Further, NL weights are found to be less competitive than ours due to the lack of feature interactions in pairwise weights computation.

Tables 1 and 2 (bottom cell) compare several baselines to replace our learned Gaussian mask that imposes a soft locality prior for pooling. When we remove the locality prior entirely, we save compute again but observe a big drop in performance for both tasks. This suggests that context pooling indeed benefits from a local “receptive field” (similar to the findings in (Luo et al., 2016)). It also suggests the difficulty of disentangling the local prior from the pooling weights by learning the latter alone in an unfactorized way. The “Fixed window” baseline is one simple remedy to this issue by associating a fixed local window to the pooling function, where the window size is hand-picked on validation data. We see immediate help from this baseline (relative to “no locality prior”). On the other hand, we find pooling at random sparse locations will slightly hurt performance. Finally, learning adaptive local windows performs close to our method with adaptive soft Gaussian masks, but the benefits of the latter still hold with consistent gains.

Ablation on the design choice of ContextPool module

Recall that our default ContextPool module is implemented as a convolutional mapping function $m(\mathbf{X})$, which maps the input feature matrix \mathbf{X} into arrays of pooling weights

Table 2. Ablations on ImageNet-1K classification. Top1 is top-1 accuracy. Throughput (images / s) is measured on a V100 GPU. CP denotes the use of our full ContextPool module ($w \odot g^i$). The middle and bottom cells compare with alternative weightings and locality priors respectively for context pooling.

Method	FLOPs (G)	Throughput	Top1 \uparrow
ViT-B/16	55.4	85.9	77.9
CP-ViT-B/16 ($w \odot g^i$)	56.7	84.1	79.9
Unnormalized weights $\odot g^i$	56.6	84.2	79.7
Uniform weights $\odot g^i$	56.1	84.8	78.9
NL weights $\odot g^i$	68.8	69.2	79.4
No locality prior $\odot w$	56.0	85.1	78.3
Fixed window $\odot w$	56.0	85.1	78.9
Adaptive window $\odot w$	56.7	84.1	79.6
Random sparse $\odot w$	56.0	85.1	78.1

and size. Table 3 compares such a CNN-based design choice against alternatives like fully-connected MLP and self-attention layers. Here we conduct the comparing experiments on both transformers and ConvNets (detailed in supplementary materials) for a more comprehensive ablation. Note for transformers, we still benchmark on the same tasks as in Tables 1 and 2, with identical task settings and baseline models.

The MLP-based ContextPool module in Table 3 can be considered as the simplest form of $m(\cdot)$, which maps each feature vector $x_i \in \mathbf{X}$ to its corresponding pooling weights. We can see that MLP is more compute-efficient than our convolutional module but worse in performance. The reason is that such MLP module operates individually for x_i without considering feature interactions when predicting their pooling weights, while convolutional layers leverage neighboring features to do so. Note we can use a giant MLP that predicts for all $\{x_i\}$ together, which becomes collaborative but at a much higher cost.

Alternatively, we can implement $m(\cdot)$ using a (single-head) self-attention layer as in Eq. (3). However, as mentioned before, such mapping function $m(\cdot)$ is not only costly with quadratic computation, but also limited in modeling feature interactions. As shown in Table 3, the inferiority also translates to the ConvNet framework. Note we can improve by modeling richer feature interactions with more than one attention layers, but this will further increase the cost.

4.3. Visualizations and Analysis

Now we visualize what have been learned in our pooling weights and pooling sizes (in the form of soft Gaussian mask). Since visualization is easier on images with spatial grids, we take the ViT-B/16 model and visualize the predictions from our ContextPool module after the second attention layer.

Table 3. Ablation study on the design choice of ContextPool (CP) module. For transformers, we choose the same NMT and image classification tasks as in Tables 1 and 2, with identical task settings and baseline models. We also include the ConvNet experiments (details in supplementary materials) for more comprehensive ablations.

Method	Transformer (EN-DE translation)			Transformer (ImageNet classification)			ConvNet (CIFAR-10 classification)		
	Memory (G)	Speed	BLEU \uparrow	FLOPs (G)	Throughput	Top1 \uparrow	FLOPs (G)	Size	Accuracy \uparrow
Baseline	17.2	1.20	28.16	55.4	85.9	77.9	3.7	0.66M	92.9
+ CNN-based CP (default)	17.6	1.12	28.91	56.7	84.1	79.9	3.9	0.68M	93.4
+ MLP-based CP	17.3	1.17	28.33	56.5	85.2	78.7	3.8	0.67M	93.1
+ Self-attention-based CP	21.3	0.84	28.66	68.8	69.2	79.4	4.4	0.67M	93.2

Table 4. The BLEU scores for token-level translation on the WMT 2014 EN-DE and EN-FR datasets. We compare our CP-attention with standard attention (Vaswani et al., 2017), local attention (Yang et al., 2018) and area attention (Li et al., 2019b).

Model	Standard attention		Local attention		Area attention		CP-attention (ours)	
	EN-DE	EN-FR	EN-DE	EN-FR	EN-DE	EN-FR	EN-DE	EN-FR
Small	22.55	31.93	22.71	32.48	23.20	32.93	23.67	33.24
Base	28.16	38.97	28.32	39.04	28.52	39.19	28.91	39.36
Big	29.26	41.00	29.31	41.17	29.77	41.46	30.11	41.59

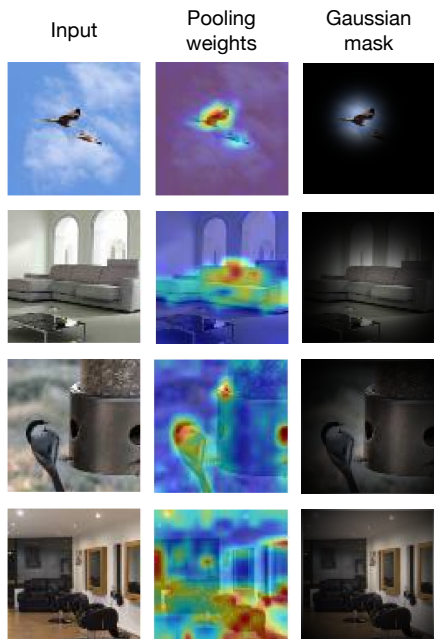


Figure 3. Visualizations of the pooling weights and size (in the form of soft Gaussian mask) predicted by our ContextPool module on example ImageNet images. We observe that the pooling weights are learned to aggregate diverse information from different locations or object parts, while the pooling size is learned to capture either local or global image context depending on the input.

We are able to observe from Fig. 3 that: 1) The pooling weights are learned to aggregate diverse information, and seem to go beyond feature similarity (the main intuition of NL weights). The last image gives one example where the pooling weights highlight some dissimilar regions around

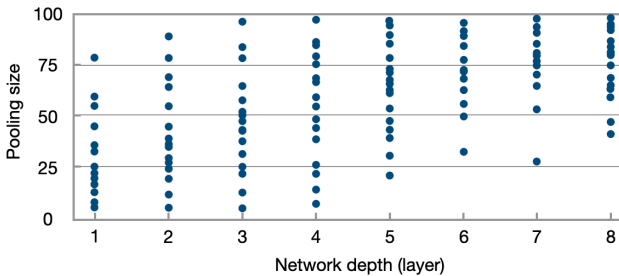


Figure 4. Distributions of the predicted pooling size by our ContextPool module in different attention layers (ViT-B/16).

the window and ceiling, which can instead accumulate evidence for the target class of “room”. 2) The learned pooling size is indeed input dependent, capturing the local or global context adaptively. Fig. 4 further shows the distributions of pooling size in different layers. Interestingly, the predicted pooling size remains diverse within each layer, but in general tends to increase at higher layers to capture long-range dependencies.

4.4. Comparing to SOTAs on Language Tasks

Table 4 evaluates our ContextPool-based attention model on the token-level NMT task using both EN-DE and EN-FR datasets. Comparison is made against standard attention and other variants that model context differently. Three transformer architectures are adopted as in (Li et al., 2019b).

It is observed that local attention only achieves marginal gains over standard attention, mainly because the locality is added to the attention mechanism which hurts the full attention capacity. Area attention preserves full attention by

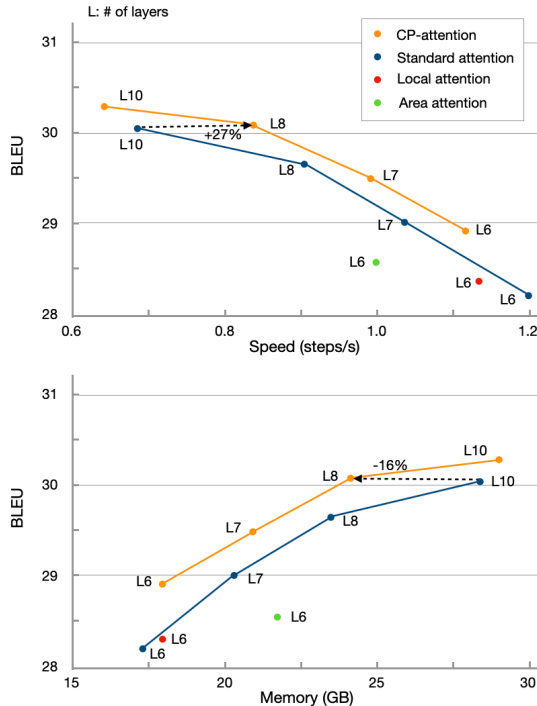


Figure 5. Performance-cost comparisons on token-level translation (EN-DE task) using the Base model. The number of layers ranges from $L = 6$ to 10.

allowing queries to attend to the whole memory. The memory is a multi-scale one to encode context of varying scales. Despite the strong BLEU scores from area attention, it is not flexible enough to model content-dependent context due to the use of fixed set of pooling sizes when constructing the multi-scale memory. Our ContextPool is able to meaningfully outperform area attention across datasets and model sizes, thanks to its adaptiveness during context pooling.

Fig. 5 further compares the above methods in terms of computation and memory complexities. Given the default number of layers $L = 6$, our CP-attention not only outperforms others at the same L , but also strikes a better trade-off between performance and cost. For instance, our CP-attention ($L = 6$) obtains a higher BLEU score 28.91 at a noticeably faster speed and lower memory than area attention, since the latter needs to maintain a multi-scale memory online. More importantly, we are able to utilize our saved compute in the form of additional layers (increasing L to 7). This way, we further improve the model capacity and BLEU score, but our speed and memory remain comparable to those of area attention with only $L = 6$ layers. When we continue to train a deeper model with CP, we found significantly boosted parameter efficiency over the one without CP. Interestingly, our CP-attention with $L = 8$ layers obtains a similar BLEU score with the 10-layer vanilla attention model (without CP), leading to 27% faster speed and 16%

Table 5. BPC (\downarrow) and model size on enwik8 and text8 for autoregressive language modeling. The number of layers is included in the parenthesis. CP denotes the use of our ContextPool module.

Model	#Param	text8		enwik8	
		Dev	Test	Dev	Test
T12	44M	-	1.18	-	1.11
Transformer-XL	41M	-	-	-	1.06
Adaptive local	38M	1.05	1.11	1.04	1.02
BP-Transformer	38M	-	1.11	-	1.02
Longformer	41M	1.04	1.10	1.02	1.00
Reformer	-	-	-	-	1.05
CP-Transformer (12)	39M	1.04	1.09	1.02	0.99
CP-Transformer (14)	44M	1.02	1.07	1.01	0.97
CP-Transformer (11)	36M	1.05	1.11	1.03	1.01
CP-Adaptive local	39M	1.05	1.10	1.03	1.01
CP-Longformer	43M	1.03	1.09	1.02	0.99

less memory used. On the other hand, when we train shallower models, the performance gap (with vs. without CP) becomes larger, *e.g.*, $\Delta\text{BLEU}=1.21$ when $L=4$. This again demonstrates our improved model expressiveness.

Finally, we evaluate on the challenging task of character-level autoregressive language modeling (see Table 5). BPC results are reported on the Dev/Test sets of enwik8 and text8 datasets. We compare with the baseline models of T12 (Al-Rfou et al., 2019) and Transformer-XL (Dai et al., 2019), as well as four representative methods of sparse attention. Among them, Adaptive local (Sukhbaatar et al., 2019) and BP-Transformer (Ye et al., 2019) use the local window as a sparsity prior, but with a learned window size and multi-scale windows respectively. Longformer (Beltagy et al., 2020) uses a combined sparsity pattern (global+local window), while Reformer (Kitaev et al., 2020) chooses to learn the patterns.

The above sparse attention methods differ from our ContextPool method in their loss of full attention capacity despite the improved efficiency. Our method on the other hand, computes full attention over ContextPooled token features. Note our context pooling function does have a locality prior, similar to existing sparsity priors based on local window. But the critical difference is that our locality prior is only applied to feature pooling, not to the following full attention process.

Table 5 confirms the benefits of full attention models. Our CP module when applied to the standard 12-layer transformer, makes a strong baseline CP-Transformer (12) that has a small model size (39M parameters) but consistently outperforms the compared sparse attention methods. We are able to further lower the model size to 36M when inserting CP to a 11-layer model without sacrificing the performance much, due to the boosted model expressiveness. When the

Table 6. ImageNet-1K top1 classification accuracy. Throughput (images/s) is measured on a V100 GPU. CP denotes the use of our ContextPool module. The number of attention layers is included in the parenthesis.

Method	image	#Param	FLOPS	image/s	Top1
ViT-B/16	384 ²	86M	55.4G	85.9	77.9
ViT-L/16	384 ²	307M	190.7G	27.3	76.5
DeiT-S	224 ²	22M	4.6G	940.4	79.8
DeiT-B	224 ²	86M	17.5G	292.3	81.8
DeiT-B	384 ²	86M	55.4G	85.9	83.1
Swin-S	224 ²	50M	8.7G	436.9	83.0
Swin-B	224 ²	88M	15.4G	278.1	83.5
Swin-B	384 ²	88M	47.0G	84.7	84.5
CP-ViT-B/16 (12)	384 ²	88M	57.2G	85.1	79.2
CP-ViT-B/16 (10)	384 ²	75M	48.7G	96.1	76.8
CP-Swin-B	224 ²	89M	16.8G	272.3	84.3
CP-Swin-B	384 ²	89M	48.9G	81.4	85.6

saved parameters are re-invested in constructing a deeper model (14 layers) that has comparable model size of 44M, we attain new state-of-the-art performance on both enwik8 and text8.

The bottom cell of Table 5 examines if our ContextPool is complementary to sparse attention. The answer is positive given our consistent gains over two sparse attention baselines. Intuitively, sparse attention would benefit more from our expressive token features that are context-aware.

4.5. Comparing to SOTAs on Image Classification

Table 6 evaluates our CP method on ImageNet classification and compares with the state-of-the-art methods ViT (Dosovitskiy et al., 2021), DeiT (Touvron et al., 2021) and Swin-T (Liu et al., 2021). It is shown that when CP is simply applied to the 12-layer ViT-B/16 model, performance gains are achieved at low overhead. When we plug CP into a smaller CP-ViT-B/16 model with 10 layers, this model can perform even comparably to ViT-L/16 despite being much more efficient. We further show CP is applicable to the Swin transformer that computes multi-scale attention with an image pyramid. Our CP method proves helpful for the two Swin-B models using different input image resolutions, and achieves a strong top1 accuracy of 85.6%.

5. Conclusions and Future Work

In this paper we have shown how adaptive pooling of features for a location based on context can improve the results for a transformer model, both by reducing the number of layers needed to achieve similar accuracy and by improving accuracy of models with the same number of layers. For future work we hope to apply this technique more broadly to

other domains, such as speech recognition that have multi-level contextual dependencies that span different, dynamic extents. It is our hope that adaptive pooling can benefit other such domains. In addition, a common dynamic pooling mechanism across Convnets and transformers can help to simplify hybrid architectures that adapt to context, opening up new efficient design choices.

Acknowledgements

The authors want to thank Shih-Yu Sun, Hesam Najafi Shoushtari, Kelsey Ho and many others at Apple for helpful discussions during the course of this project. We also thank the ICML reviewers for providing useful feedback.

References

- Ainslie, J., Ontaño, S., Alberti, C., Cvicek, V., Fisher, Z., Pham, P., Ravula, A., Sanghai, S., Wang, Q., and Yang, L. ETC: Encoding long and structured data in transformers. In *EMNLP*, 2020.
- Al-Rfou, R., Choe, D., Constant, N., Guo, M., and Jones, L. Character-level language modeling with deeper self-attention. In *AAAI*, 2019.
- Beltagy, I., Peters, M. E., and Cohan, A. Longformer: The long-document transformer. *arXiv:2004.05150*, 2020.
- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv:1904.10509*, 2019.
- Coates, A. and Ng, A. Selecting receptive fields in deep networks. In *NeurIPS*, 2011.
- Cubuk, E. D., Zoph, B., Shlens, J., and Le, Q. Randaugment: Practical automated data augmentation with a reduced search space. In *NeurIPS*, 2020.
- Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., and Wei, Y. Deformable convolutional networks. In *ICCV*, 2017.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., and Salakhutdinov, R. Transformer-XL: Attentive language models beyond a fixed-length context. In *ACL*, 2019.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer,

- M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlisby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- Ho, J., Kalchbrenner, N., Weissenborn, D., and Salimans, T. Axial attention in multidimensional transformers. *arXiv:1912.12180*, 2019.
- Jia, Y., Huang, C., and Darrell, T. Beyond spatial pyramids: Receptive field learning for pooled image features. In *CVPR*, 2012.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *ICML*, 2020.
- Kitaev, N., Kaiser, L., and Levskaya, A. Reformer: The efficient transformer. In *ICLR*, 2020.
- Li, S., Jin, X., Xuan, Y., Zhou, X., Chen, W., Wang, Y., and Yan, X. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. In *NeurIPS*, 2019a.
- Li, Y., Kaiser, L., Bengio, S., and Si, S. Area attention. In *ICML*, 2019b.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- Luo, W., Li, Y., Urtasun, R., and Zemel, R. Understanding the effective receptive field in deep convolutional neural networks. In *NeurIPS*, 2016.
- Mahoney, M. Large text compression benchmark. <http://mattmahoney.net/dc/textdata>, 2009.
- Pintea, S., Tömen, N., Goes, S., Loog, M., and van Gemert, J. Resolution learning in deep convolutional networks using scale-space theory. *IEEE Transactions on Image Processing*, 30:8342 – 8353, 2021.
- Qiu, J., Ma, H., Levy, O., Yih, W.-t., Wang, S., and Tang, J. Blockwise self-attention for long document understanding. In *EMNLP*, 2020.
- Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., and Hsieh, C.-J. DynamicViT: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, 2021.
- Roy, A., Saffar, M., Vaswani, A., and Grangier, D. Efficient content-based sparse attention with routing transformers. *Transactions of the Association for Computational Linguistics*, 9(0):53–68, 2022.
- Ryoo, M. S., Piergiovanni, A., Arnab, A., Deghani, M., and Angelova, A. Tokenlearner: Adaptive space-time tokenization for videos. In *NeurIPS*, 2021.
- Shelhamer, E., Wang, D., and Darrell, T. Blurring the line between structure and learning to optimize and adapt receptive fields. *arXiv:1904.11487*, 2019.
- Sukhbaatar, S., Grave, E., Bojanowski, P., and Joulin, A. Adaptive attention span in transformers. In *ACL*, 2019.
- Tay, Y., Bahri, D., Yang, L., Metzler, D., and Juan, D.-C. Sparse sinkhorn attention. In *ICML*, 2020.
- Tomen, N., Pintea, S.-L., and Van Gemert, J. Deep continuous networks. In *ICML*, 2021.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., and Jegou, H. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *NeurIPS*, 2017.
- Wang, S., Li, B., Khabsa, M., Fang, H., and Ma, H. Linformer: Self-attention with linear complexity. *arXiv:2006.04768*, 2020.
- Wang, W., Xie, E., Li, X., Fan, D.-P., Song, K., Liang, D., Lu, T., Luo, P., and Shao, L. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021.
- Wang, X., Girshick, R., Gupta, A., and He, K. Non-local neural networks. In *CVPR*, 2018.
- Yang, B., Tu, Z., Wong, D. F., Meng, F., Chao, L. S., and Zhang, T. Modeling localness for self-attention networks. In *EMNLP*, 2018.
- Ye, Z., Guo, Q., Gan, Q., Qiu, X., and Zhang, Z. Bp-transformer: Modelling long-range context via binary partitioning. *arXiv:1911.04070*, 2019.
- Zaheer, M., Guruganesh, G., Dubey, K. A., Ainslie, J., Alberti, C., Ontanon, S., Pham, P., Ravula, A., Wang, Q., Yang, L., et al. Big bird: Transformers for longer sequences. In *NeurIPS*, 2020.
- Zhang, P., Dai, X., Yang, J., Xiao, B., Yuan, L., Zhang, L., and Gao, J. Multi-scale vision longformer: A new vision transformer for high-resolution image encoding. In *ICCV*, 2021.