
Area Attention

Yang Li¹ Lukasz Kaiser¹ Samy Bengio¹ Si Si¹

Abstract

Existing attention mechanisms are trained to attend to individual items in a collection (the memory) with a predefined, fixed granularity, e.g., a word token or an image grid. We propose area attention: a way to attend to areas in the memory, where each area contains a group of items that are structurally adjacent, e.g., spatially for a 2D memory such as images, or temporally for a 1D memory such as natural language sentences. Importantly, the shape and the size of an area are dynamically determined via learning, which enables a model to attend to information with varying granularity. Area attention can easily work with existing model architectures such as multi-head attention for simultaneously attending to multiple areas in the memory. We evaluate area attention on two tasks: neural machine translation (both character and token-level) and image captioning, and improve upon strong (state-of-the-art) baselines in all the cases. These improvements are obtainable with a basic form of area attention that is parameter free.

1. Introduction

Attentional mechanisms have significantly boosted the accuracy on a variety of deep learning tasks (Bahdanau et al., 2014; Luong et al., 2015; Xu et al., 2015). They allow the model to selectively focus on specific pieces of information, which can be a word in a sentence for neural machine translation (Bahdanau et al., 2014; Luong et al., 2015) or a region of pixels in image captioning (Xu et al., 2015; Sharma et al., 2018).

An attentional mechanism typically follows a *memory-query* paradigm, where the memory M contains a collection of items of information from a source modality such as the

embeddings of an image (Xu et al., 2015) or the hidden states of encoding an input sentence (Bahdanau et al., 2014; Luong et al., 2015), and the query q comes from a target modality such as the hidden state of a decoder model. In recent architectures such as Transformer (Vaswani et al., 2017), self-attention involves queries and memory from the same modality for either encoder or decoder. Each item in the memory has a key-value pair, (k_i, v_i) , where the key is used to compute the probability a_i regarding how well the query matches the item (see Equation 1).

$$a_i = \frac{\exp(f_{att}(q, k_i))}{\sum_{j=1}^{|M|} \exp(f_{att}(q, k_j))} \quad (1)$$

The typical choices for f_{att} include dot products qk_i (Luong et al., 2015) and a multilayer perceptron (Bahdanau et al., 2014). The output O_q^M from querying the memory M with q is then calculated as the sum of all the values in the memory weighted by their probabilities (see Equation 2), which can be fed to other parts of the model for further calculation. During training, the model learns to attend to specific pieces of information given a query. For example, it can associate a word in the target sentence with a word in the source sentence for translation tasks.

$$O_q^M = \sum_{i=1}^{|M|} a_i v_i \quad (2)$$

Attention mechanisms are typically designed to focus on individual items in the entire memory, where each item defines the granularity of what the model can attend to. For example, it can be a character for a character-level translation model, a word for a word-level model, a grid cell for an image-based model or a hidden state in a latent space. Such a construction of attention granularity is predetermined rather than learned. While this kind of item-based attention has been helpful for many tasks, it can be fundamentally limited for modeling complex attention distribution that might be involved in a task.

In this paper, we propose *area attention*, as a general mechanism for the model to attend to a group of items in the memory that are structurally adjacent. In area attention, each unit for attention calculation is an area that can contain one or more than one item. Because each of these areas

¹Google Research, Mountain View, CA, USA. Correspondence to: Yang Li <liy@google.com>.

can aggregate a varying number of items, the granularity of attention is learned from the data rather than predetermined. Note that area attention subsumes item-based attention because when an area contains a single item, it is equivalent to regular attention mechanisms. Area attention can be used along multi-head attention (Vaswani et al., 2017). With each head using area attention, multi-head area attention allows the model to attend to multiple areas in the memory. As we show in the experiments, the combination of both achieved the best results.

Extensive experiments with area attention indicate that area attention outperforms regular attention on a number of recent models for two popular tasks: machine translation (both token and character-level translation on WMT'14 EN-DE and EN-FR), and image captioning (trained on COCO and tested for both in-domain with COCO40 and out-of-domain captioning with Flickr 1K). These models involve several distinct architectures, such as the canonical LSTM seq2seq with attention (Luong et al., 2015) and the encoder-decoder Transformer (Vaswani et al., 2017; Sharma et al., 2018).

2. Related Work

The issue of item-grouping such as ranges or segments of a sentence, beyond individual tokens, has been investigated for problems such as dependency parsing or constituency parsing in natural language processing. Recent works (Wang & Chang, 2016; Stern et al., 2017; Kitaev & Klein, 2018) represent a sentence segment by subtracting the encoding of the first token from that of the last token in the segment, assuming the encoder captures contextual dependency of tokens. The popular choices of the encoder are LSTM (Wang & Chang, 2016; Stern et al., 2017) or Transformer (Kitaev & Klein, 2018). In contrast, the representation of an area (or a segment) in area attention, for its basic form, is defined as the mean of all the vectors in the segment where each vector does not need to carry contextual dependency. We calculate the mean of each area of vectors using subtraction operation over a summed area table (Viola & Jones, 2001) that is fundamentally different from the subtraction applied in these previous works.

Lee et al. proposed a rich representation for a segment in coreference resolution tasks (Lee et al., 2017), where each span (segment) in a document is represented as a concatenation of the encodings of the first and last words in the span, the size of the span and an attention-weighted sum of the word embeddings within the span. Again, this approach operates on encodings that have already captured contextual dependency between tokens, while area attention we propose does not require each item to carry contextual or dependency information. In addition, the concept of range, segment or span in all the above works is proposed in a specific context and addresses a unique language-related

task, rather than aiming for improving general attentional mechanisms that can be applied to any problems.

Instead of using softmax as attention activation function, sigmoid has been used to allow multiple items to be attended (Shen & Lee, 2016; Rei & Søgaard, 2018). An important distinction is that using sigmoid activation alone does not enforce the constraint for attended items to be structurally adjacent while area attention does.

Previous works have proposed several methods for capturing structures in attention calculation. For example, Kim et al. used a conditional random field to directly model the dependency between items, which allows multiple "cliques" of items to be attended to at the same time (Kim et al., 2017). Niculae and Blondel approached the problem, from a different angle, by using regularizers to encourage attention to be placed onto contiguous segments (Niculae & Blondel, 2017). In image captioning tasks, previous work showed that it is beneficial to attend to semantic regions or concepts on an image (Pedersoli et al., 2016; Zheng et al., 2017; Anderson et al., 2017; Lu et al., 2018; You et al., 2016). They often train a dedicated sub-network such as Fast R-CNN (Girshick, 2015) to extract region or object proposals.

Compared to these previous works, area attention we propose here does not require to train a special network or sub-network, or use an additional loss (regularizer) to capture structures, and can be entirely parameter free. It allows a model to attend to information at a varying granularity, which can be at the input layer where each item might lack contextual information, or in the latent space. While region proposal-based methods can probably extract better-quality regions as they are often pre-trained with labeled image regions, area attention is more lightweight and generally applicable. It is easy to apply area attention to existing single or multi-head attention mechanisms. By enhancing Transformer, an attention-based architecture, (Vaswani et al., 2017) with area attention, we achieved state-of-art results on a number of tasks.

3. Area-Based Attention Mechanisms

An area is a group of structurally adjacent items in the memory. When the memory consists of a sequence of items, a 1-dimensional structure, an area is a range of items that are sequentially (or temporally) adjacent and the number of items in the area can be one or multiple. Many language-related tasks are categorized in the 1-dimensional case, e.g., machine translation or sequence prediction tasks. In Figure 1, the original memory is a 4-item sequence. By combining the adjacent items in the sequence, we form area memory where each item is a combination of multiple adjacent items in the original memory. We can limit the maximum area size to consider for a task, e.g., 3 in Figure 1.




Figure 1: An illustration of area attention for the 1-dimensional case. In this example, the memory is a 4-item sequence and the maximum size of an area allowed is 3.

When the memory contains a grid of items, a 2-dimensional structure, an area can be any rectangular region in the grid (see Figure 2). This resembles many image-related tasks, e.g., image captioning. Again, we can limit the maximum size allowed for an area. For a 2-dimensional area, we can set the maximum height and width for each area. In this example, the original memory is a 3x3 grid of items and the maximum height and width allowed for each area is 2.




Figure 2: An illustration of area attention for the 2-dimensional case. In this example, the memory is a 3x3 grid and the dimension allowed for an area is 2x2.

As we can see, many areas can be generated by combining adjacent items. For the 1-dimensional case, the number of areas that can be generated is $|R| = (L - S)S + (S + 1)S/2$ where S is the maximum size of an area and L is the length of the sequence. For the 2-dimensional case, there are a quadratic number of areas can be generated from the original memory: $|R| = |R_v||R_h|$. $|R_v| = (L_v - H)H + (H + 1)H/2$ and $|R_h| = (L_h - W)W + (W + 1)W/2$ where L_v and L_h are the height and width of the memory grid and H and W are the maximum height and width allowed for a rectangular area.

To be able to attend to each area, we need to define the key and value for each area that contains one or multiple items in the original memory. As the first step to explore area attention, we define the key of an area, μ_i , simply as the mean vector of the key of each item in the area.

$$\mu_i = \frac{1}{|r_i|} \sum_{j=1}^{|r_i|} k_{i,j} \quad (3)$$

where $|r_i|$ is the size of the area r_i . For the value of an area, we define it as the the sum of all value vectors in the area.

$$v_i^{r_i} = \sum_{j=1}^{|r_i|} v_{i,j} \quad (4)$$

With the keys and values defined, we can use the standard way for calculating attention as discussed in Equation 1 and Equation 2. Note that this basic form of area attention (Eq.3 and Eq.4) is parameter-free—it does not introduce any parameters to be learned. Essentially, Equation 3 and 4 use average and sum pooling over an area of vectors. It is possible to use other pooling methods to compute the key and value vector for each area such as max pooling, which we will discuss later.

3.1. Combining Area Features

Alternatively, we can derive a richer representation of each area by using features other than the mean of the key vectors of the area. For example, we can consider the standard deviation of the key vectors within each area.

$$\sigma_i = \sqrt{\frac{1}{|r_i|} \sum_{l=1}^{|r_i|} (k_{i,l} - \mu_i)^2} \quad (5)$$

We can also consider the height and width of each area, $h_i, 1 \leq h_i \leq H$ and $w_i, 1 \leq w_i \leq W$, as the features of the area. To combine these features, we use a multi-layer perceptron. To do so, we treat h_i and w_i as discrete values and project them onto a vector space using embedding (see Equation 6 and 7).

$$e_i^h = \mathbb{1}(h_i)E^h \quad (6)$$

$$e_i^w = \mathbb{1}(w_i)E^w \quad (7)$$

where $\mathbb{1}(h_i)$ and $\mathbb{1}(w_i)$ are the one-hot encoding of h_i and w_i , and $E^h \in R^{H \times S}$ and $E^w \in R^{W \times S}$ are the embedding matrices. S is the depth of the embedding. We concatenate them to form the representation of the shape of an area.

$$e_i = [e_i^h, e_i^w] \quad (8)$$

We then combine them using a single-layer perceptron followed by a linear transformation (see Equation 9).

$$k_i^T = \phi(\mu_i W_\mu + \sigma_i W_\sigma + e_i W_e) W_d \quad (9)$$

where ϕ is a nonlinear transformation such as ReLU, and $W_\mu \in \mathbb{R}^{D \times D}$, $W_\sigma \in \mathbb{R}^{D \times D}$, $W_e \in \mathbb{R}^{2S \times D}$ and $W_d \in \mathbb{R}^{D \times D}$. W_μ , W_σ , W_e and W_d are trainable parameters.

3.2. Fast Computation Using Summed Area Table

If we naively compute μ_i , σ_i and v_i^T , the time complexity for computing attention will be $O(|M|A^2)$ where $|M|$ is the size of the memory that is L for a 1-dimensional sequence or $L_v L_h$ for a 2-dimensional memory. A is the maximum size of an area, which is S in the one dimensional case and WH in the 2-dimensional case. This is computationally expensive in comparison to the attention computed on the original memory, which is $O(|M|)$. To address the issue, we use summed area table, an optimization technique that has been used in computer vision for computing features on image areas (Viola & Jones, 2001). It allows constant time to calculate a summation-based feature in each rectangular area, which allows us to bring down the time complexity to $O(|M|A)$ —We will report on the actual time cost in our experimental section.

Summed area table is based on an integral image (Szeliski, 2010), I , which can be efficiently computed in a single pass of the memory. With the integral image, we can calculate the key and value of each area in constant time. We present the Pseudo code for performing Eq. 3, 4 and 5 as well as the shape size of each area in Algorithm 1 and 2. These Pseudo code are designed based on highly efficient Tensor operations¹.

4. Experiments

We experimented with area attention on two important tasks: neural machine translation (including both token and character-level translation) and image captioning, where attention mechanisms have been a common component in model architectures for these tasks. The architectures we investigate involves several popular encoder and decoder choices, such as LSTM (Hochreiter & Schmidhuber, 1997) and Transformer (Vaswani et al., 2017). The attention mechanisms in these tasks include both self attention and encoder-decoder attention. Note that area attention does not change the size of queries and only expands the size of keys and values. The basic form of area attention is completely parameter free. As a result, all the models in Table 1-5 using (Eq.3 & 4) have the same number of parameters as the corresponding baseline models, which allows a fair comparison.

¹See TensorFlow implementation of Area Attention as well as its integration with Transformer and LSTM in <https://github.com/tensorflow/tensor2tensor>.

Algorithm 1: Compute the vector sum and the size of each area, for all the qualified rectangular areas on a given grid.

Input: A tensor G in shape of $[H, W, D]$ that represents a grid with height H and width W where each item is a vector of depth D .

Output: Sum of vectors of each area, U , and height and width of each area, S_h and S_w .

Hyperparameter: maximum allowed area width W_a and height H_a .

```

1 Compute horizontal integral image  $I_h$  by cumulative sum
2   along horizontal direction over  $G$ ;
3 Compute integral image  $I_{hv}$  by cumulative sum along
4   vertical directions over  $I_h$ ;
5 Compute  $I$  by padding all-zero vectors to the left and
6   top of  $I_{hv}$ ;
7 for  $h = 1, \dots, H_a$  do
8   for  $w = 1, \dots, W_a$  do
9      $I_1 \leftarrow I[h + 1 :, w + 1 :, :]$ ;
10     $I_2 \leftarrow I[: -h - 1, : -w - 1, :]$ ;
11     $I_3 \leftarrow I[h + 1 :, : -w - 1, :]$ ;
12     $I_4 \leftarrow I[: -h - 1, w + 1 :, :]$ ;
13     $\bar{U} = I_1 + I_2 - I_3 - I_4$ ;
14     $\bar{S}_h \leftarrow [h]^{(H-h) \times (W-w)}$ ; Fill tensor with value  $h$ 
      for the height of each area;
15     $\bar{S}_w \leftarrow [w]^{(H-h) \times (W-w)}$ ; Fill tensor with value  $w$ 
      for the width of each area;
16     $S_h \leftarrow [S_h \ \bar{S}_h]$ , reshape  $\bar{S}_h$  to  $[-1, 1]$  and
      concatenate on the first dimension;
17     $S_w \leftarrow [S_w \ \bar{S}_w]$ , reshape  $\bar{S}_w$  to  $[-1, 1]$  and
      concatenate on the first dimension;
18     $U \leftarrow [U \ \bar{U}]$ , reshape  $\bar{U}$  to  $[-1, D]$  and
      concatenate on the first dimension;
19 return  $U$ ,  $S_h$  and  $S_w$ .
```

4.1. Token-Level Neural Machine Translation

Transformer has recently (Vaswani et al., 2017) established the state of art performance on WMT 2014 English-to-German and English-to-French tasks, while LSTM with encoder-decoder attention has been a popular choice for neural machine translation tasks. We use the same dataset as the one used in (Vaswani et al., 2017) in which the WMT 2014 English-German (EN-DE) dataset contains about 4.5 million English-German sentence pairs, and the English-French (EN-FR) dataset has about 36 million English-French sentence pairs (Wu et al., 2016). A token is either a byte pair (Britz et al., 2017) or a word piece (Wu et al., 2016) as in the original Transformer experiments. We performed three runs for each configuration and report the average of these runs. * stands for statistical significance ($p < 0.05$) for comparison with regular attention and ** for statistical significance when comparing with all the other model conditions.

Algorithm 2: Compute the vector mean, standard deviation, and sum as well as the size of each area, for all the qualified rectangular areas on a grid.

Input: A tensor G in shape of $[H, W, D]$ that represents a grid with height H and width W where each item is a vector of depth D .

Output: Vector mean μ , standard deviation σ and sum U as well as height S_h and width S_w of each area.

- 1 Acquire U , S_h and S_w using Algorithm 1 with input G ;
 - 2 Acquire U' using Algorithm 1 with input $G \odot G$ where \odot is for element-wise multiplication;
 - 3 $\mu \leftarrow U \oslash S$ where \oslash is for element-wise division;
 - 4 $\mu' \leftarrow U' \oslash S$;
 - 5 $\sigma \leftarrow \sqrt{\mu' - \mu \odot \mu}$;
 - 6 **return** μ , σ , U , as well as S_h and S_w .
-

4.1.1. TRANSFORMER TOKEN-LEVEL MT EXPERIMENTS

Transformer heavily uses attentional mechanisms, including both self-attention in the encoder and the decoder, and attention from the decoder to the encoder. We vary the configuration of Transformer to investigate how area attention impacts the model. In particular, we investigated the following variations of Transformer: *Tiny* (#hidden layers=2, hidden size=128, filter size=512, #attention heads=4), *Small* (#hidden layers=2, hidden size=256, filter size=1024, #attention heads=4), *Base* (#hidden layers=6, hidden size=512, filter size=2048, #attention heads=8) and *Big* (#hidden layers=6, hidden size=1024, filter size=4096 for EN-DE and 8192 for EN-FR, #attention heads=16).

During training, sentence pairs were batched together based on their approximate sequence lengths. All the model variations except *Big* uses a training batch contained a set of sentence pairs that amount to approximately 32,000 source and target tokens and were trained on one machine with 8 NVIDIA P100 GPUs for a total of 250,000 steps. Given the batch size, each training step for the Transformer Base model, on 8 NVIDIA P100 GPUs, took 0.4 seconds for Regular Attention, 0.5 seconds for the basic form of Area Attention (Eq.3 and Eq.4), 0.8 seconds for Area Attention using multiple features (Eq.9 and Eq.4).

For *Big*, due to the memory constraint, we had to use a smaller batch size that amounts to roughly 16,000 source and target tokens and trained the model for 600,000 steps. Each training step took 0.5 seconds for Regular Attention, 0.6 seconds for the basic form of Area Attention (Eq.3 and 4), 1.0 seconds for Area Attention using multiple features (Eq.9 and 4). Similar to previous work, we used the Adam optimizer with a varying learning rate over the course of training—see (Vaswani et al., 2017) for details.

We applied area attention to each of the Transformer variation to both encoder and decoder self-attention, and the encoder-decoder attention in the first two layers. We found area attention consistently improved Transformer on all the model variations (see Table 1), even with the basic form of area attention where no additional parameters are used (Eq.3 and Eq.4). For Transformer Base, area attention achieved BLEU scores (EN-DE: 28.52 and EN-FR: 39.27) that surpassed the previous results for both EN-DE and EN-FR.

For EN-FR, the performance of Transformer Big with regular attention—a baseline—does not match what was reported in the Transformer paper (Vaswani et al., 2017), largely due to a different batch size and the different number of training steps used, although area attention still outperformed the baseline consistently. On the other hand, area attention with Transformer Big achieved BLEU 29.77 on EN-DE that improved upon the state-of-art result of 28.4 reported in (Vaswani et al., 2017) with a significant margin.

4.1.2. LSTM TOKEN-LEVEL MT EXPERIMENTS

We used a 2-layer LSTM for both encoder and decoder. The encoder-decoder attention is based on multiplicative attention where the alignment of a query and a memory key is computed as their dot product (Luong et al., 2015). We vary the size of LSTM and the number of attention heads to investigate how area attention can improve LSTM with varying capacity on translation tasks. The purpose is to observe the impact of area attention on each LSTM configuration, rather than for a comparison with Transformer.

Because LSTM requires sequential computation along a sequence, it trains rather slow compared to Transformer. To improve GPU utilization we increased data parallelism by using a much larger batch size than training Transformer. We trained each LSTM model on one machine with 8 NVIDIA P100. For a model has 256 or 512 LSTM cells, we trained it for 50,000 steps using a batch size that amounts to approximately 164,000 source and target tokens. When the number of cells is 1024, we had to use a smaller batch size with roughly 131,000 tokens, due to the memory constraint, and trained the model for 625,000 steps.

In these experiments, we used areas of maximum size 2 and attention was computed from the output of the decoder’s top layer to that of the encoder. Similar to what we observed with Transformer, area attention consistently improves LSTM architectures in all cases (see Table 2).

4.2. Character-Level Neural Machine Translation

Compared to token-level translation, character-level translation requires the model to address significantly longer sequences, which are a more difficult task and often less studied. We speculate that the ability to combine adjacent

Table 1: The BLEU scores on token-level translation tasks for the variations of the Transformer-based architecture.

Model	Regular Attention		Area Attention (Eq.3 and 4)		Area Attention (Eq.9 and 4)	
	EN-DE	EN-FR	EN-DE	EN-FR	EN-DE	EN-FR
Tiny	18.58	27.03	19.13*	27.4*	19.27*	27.91**
Small	22.55	31.93	22.84	32.31*	23.2**	32.93**
Base	28.16	38.97	28.47	39.27*	28.52*	39.19
Big	29.26	41.0	29.49	41.18	29.77*	41.46*

Table 2: The BLEU scores on token-level translation tasks for the LSTM-based architecture with varying model capacities.

#Cells	#Heads	Regular Attention		Area Attention (Eq.3,4)		Area Attention (Eq.9,4)	
		EN-DE	EN-FR	EN-DE	EN-FR	EN-DE	EN-FR
256	1	16.58	22.77	19.26*	29.35*	19.46*	29.79**
256	4	16.73	28.1	20.25*	30.49*	20.74**	30.2*
512	1	18.65	30.32	21.82*	32.80*	21.80*	32.73*
512	4	19.16	30.55	23.09*	33.75*	23.41*	34.09**
1024	1	19.4	31.99	23.69*	34.65*	23.48*	34.76*
1024	4	20.21	32.21	24.55*	35.95*	24.85**	35.97*

characters, as enabled by area attention, is likely useful to improve a regular attentional mechanisms. Likewise, we experimented with the same set of Transformer and LSTM-based architectures for this task (see the appendix for experimental details).

Transformer has not been used for character-level translation tasks. We found area attention consistently improved Transformer across all the model configurations. The best result we found in the literature is $BLEU = 23.75$ in (Kalchbrenner et al., 2017) and next $BLEU = 22.62$ reported by (Wu et al., 2016). We achieved $BLEU = 26.57$ for the English-to-German character-level translation task and $BLEU = 35.16$ on the English-to-French character-level translation task. Note that these accuracy gains are based on the basic form of area attention (see Eq.3 and Eq.4), which does not add any trainable parameters to the model.

Similarly, we tested LSTM architectures on the character-level translation tasks. We found area attention outperformed the baselines in all the conditions (see Table 4).

4.3. Image Captioning

Image captioning is the task to generate natural language description of an image that reflects the visual content of an image. This task has been addressed previously using a deep architecture that features an image encoder and a language decoder (Xu et al., 2015; Sharma et al., 2018). The image encoder typically employs a convolutional net such as

Table 3: The BLEU scores on character-level translation tasks for the Transformer-based architecture with varying model capacities.

Model	Regular		Area (Eq.3, 4)	
	EN-DE	EN-FR	EN-DE	EN-FR
Tiny	6.9	9.58	8.2	11.69*
Small	12.16	18.43	13.46*	21.05*
Base	24.88	33.26	24.96	33.68
Big	25.96	33.82	26.57	35.16

Table 4: The BLEU scores on character-level translation tasks for the LSTM-based architecture with varying model capacities.

Cell,Head	Regular		Area (Eq.3, 4)	
	EN-DE	EN-FR	EN-DE	EN-FR
256, 1	9.9	16.82	11.23	17.27
256, 4	11.99	17.36	12.51	19.85*
512, 1	16.37	23.7	16.86	24.13
512, 4	16.43	24.71	17.61	25.36
1024, 1	20.26	27.83	21.05*	29.48
1024, 4	21.28	28.53	21.71	30.68*

Table 5: Test accuracy of image captioning models on COCO40 (in-domain) and Flickr 1K (out-of-domain) tasks.

Model	COCO40		Flickr 1K	
	CIDEr	ROUGE-L	CIDEr	ROUGE-L
Benchmark (Sharma et al., 2018)	1.032	0.700	0.359	0.416
Benchmark Replicate	1.034	0.701	0.355	0.409
2×2 Eq.3 & 4	1.060	0.704	0.364	0.420
3×3 Eq.3 & 4	1.060	0.706	0.377	0.419
3×3 Eq.9 & 4	1.045	0.707	0.372	0.420

ResNet (He et al., 2015) to embed the images and then uses a recurrent net such as LSTM or Transformer (Sharma et al., 2018) to encode the image based on these embeddings. For the decoder, either LSTM (Xu et al., 2015) or Transformer (Sharma et al., 2018) has been used for generating natural language descriptions. In many of these designs, attention mechanisms have been an important component that allows the decoder to selectively focus on a specific part of the image at each step of decoding, which often leads to better captioning quality.

In this experiment, we follow a champion condition in the experimental setup of (Sharma et al., 2018) that achieved state-of-the-art results as our benchmark model. It uses a pre-trained Inception-ResNet to generate 8×8 image embeddings, a 6-layer Transformer for image encoding and a 6-layer Transformer for decoding. The benchmark model has a hidden size of 512 and uses 8-head regular attention. To investigate how area attention improves the captioning accuracy, particularly regarding self-attention and encoder-decoder attention computed off the image, we add area attention with different maximum area sizes to the first 2 layers of the image encoder self-attention and encoder-decoder (caption-to-image) attention (see Table 5), which both resemble a 2-dimensional area attention case. 2×2 stands for the maximum area size 2 by 2 and 3×3 for 3 by 3. For the 2×2 case, an area can be 1 by 1, 2 by 1, 1 by 2, and 2 by 2 as illustrated in Figure 2. 3×3 allows more area shapes.

Similar to (Sharma et al., 2018), we trained each model based on the training & development sets provided by the COCO dataset (Lin et al., 2014), which has 82K images for training and 40K for validation. Each of these images have at least 5 groundtruth captions. The training was conducted on a distributed learning infrastructure (Dean et al., 2012) with 10 GPU cores where updates are applied asynchronously across multiple replicas. We then tested each model on the COCO40 (Lin et al., 2014) and the Flickr 1K (Young et al., 2014) test sets. Flickr 1K is out-of-domain for the trained model. For each experiment, we report CIDEr (Vedantam et al., 2014) and ROUGE-L (Lin & Och, 2004)

metrics. For both metrics, higher number means better captioning accuracy—the closer distances between the predicted and the groundtruth captions. Similar to the previous work (Sharma et al., 2018), we report the numerical values returned by the COCO online evaluation server² for the COCO C40 test set. Previous work (Sharma et al., 2018) has revealed that human evaluation would give a more complete examination about the model accuracy, which we leave out in this work as our focus is on area attention as a general mechanism.

We found models with area attention outperformed the benchmark on both CIDEr and ROUGE-L metrics with a large margin (see Table 5). The models with 2×2 Eq.3 and 3×3 Eq.3 do not use any additional parameters beyond the benchmark model. 3×3 achieved the best results overall. 3×3 Eq. 9 adds a small fraction of the number of parameters to the benchmark model and did not seem to improve on the parameter-free version of area attention, although it still outperformed the benchmark.

5. Discussions

In this paper, we focus on mean (Equation 3) and sum pooling (Equation 4) as a way to compute keys and values for each area. As we can see from the experimental results, these simple parameter-free area representations can bring accuracy gain to a range of tasks. As mentioned earlier, it is possible to use other methods such as max pooling for this purpose. We experimented with max pooling on the Transformer model for both machine translation and image captioning tasks. For token-level translation, max pooling with Transformer Base achieved BLEU 28.48 for EN-DE and 39.21 for EN-FR. For character-level translation, max pooling with Transformer Base achieved 24.92 and 33.84, respectively. Similarly, for image captioning tasks, max pooling with a 2×2 maximum area size achieved CIDEr 1.055 and ROUGE-L 0.706 for COCO40 official tests, and CIDEr 0.365 and ROUGE-L 0.416 for Flickr 1K tests. While max pooling seems to offer comparable re-

²<http://mscoco.org/dataset/#captions-eval>

sults on some of the tasks, the exact solution cannot be efficiently computed. Without using summed area table, iterating over each area is significantly slower. Alternatively, we can compute and pool over all the areas in parallel. But it requires significantly more memory that limits the use of a large batch size or the handling of a long sequence such as character-level translation. It is possible to calculate approximate max pooling based on summed area table (Van Vliet, 2004), although it can incur numerical problems such as underflow. To handle very long sequences, we can apply area attention to a neighborhood of a sequence given a query instead of the entire sequence. These ideas deserve further investigation.

We found the benefit of area attention is more pronounced when a model is relatively small (see Table 1-4). It also appears that the improvement for LSTM-based architectures is quite substantial particularly on token-level translation tasks, e.g., Table 2. For character-level translation tasks, although the improvement from area attention is quite consistent across model conditions, there is not as much statistical significance as we acquired for token-level translation. One reason is that there is a larger variance in the results from which we speculate that more training iterations are needed for better convergence.

In addition to offering better accuracy, we want to better understand how area attention works, particularly regarding if area attention is able to capture structural or semantic coherence in the data. To do so, we analyze the learned multi-head area self-attention in Transformer encoder for the image captioning task (see examples in Figure 3 and additional examples in the appendix). From these examples, we can see area attention often appropriately captures the image areas that are relevant to the query grid. In particular, many of the top-attended areas (shown in bold) include more than one grid with a variety of shapes, depending on the scene.

Similarly, we have analyzed the self-attention for the character-level machine translation tasks (see examples in the appendix). The analysis reveals that area attention enables multi-head attention in Transformer to attend to the whole word that the query character belongs to as well as other relevant words in the sentences. This shows that area attention allows the model to attend to appropriate granularity of information that is more consistent with the structural and semantics coherence in the data.

6. Conclusions


In this paper, we present a novel attentional mechanism by allowing the model to attend to areas as a whole. An area contains one or a group of items in the memory to be attended. The items in the area are either spatially adja-



(a) a young boy holding a frisbee in his hand



(b) a colorful bird perched on a branch of a tree



(c) a man on a beach throwing a frisbee

Figure 3: The images on the left show the query grid and the images on the right are attended areas. The top attended area is highlighted with the color of the corresponding attention head.

cent when the memory has 2-dimensional structure, such as images, or temporally adjacent for 1-dimensional memory, such as natural language sentences. Importantly, the size of an area, i.e., the number of items in an area or the level of aggregation, can vary depending on the learned coherence of the adjacent items, which gives the model the ability to attend to information at varying granularity. Area attention contrasts with the existing attentional mechanisms that are item-based. We evaluated area attention on two tasks: neural machine translation and image captioning, based on model architectures such as Transformer and LSTM. Area attention is able to offer further improvement on accuracy consistently across a variety of tasks over these strong baselines.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful feedback that substantially improved the paper. We also want to thank the readers of the early versions of the paper for their constructive comments.

References

- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., and Zhang, L. Bottom-up and top-down attention for image captioning and VQA. *CoRR*, abs/1707.07998, 2017. URL <http://arxiv.org/abs/1707.07998>.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014. URL <http://arxiv.org/abs/1409.0473>.
- Britz, D., Goldie, A., Luong, M., and Le, Q. V. Massive exploration of neural machine translation architectures. *CoRR*, abs/1703.03906, 2017. URL <http://arxiv.org/abs/1703.03906>.
- Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Le, Q. V., Mao, M. Z., Ranzato, M., Senior, A., Tucker, P., Yang, K., and Ng, A. Y. Large scale distributed deep networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems, NIPS'12*, pp. 1223–1231, USA, 2012. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2999134.2999271>.
- Girshick, R. B. Fast R-CNN. *CoRR*, abs/1504.08083, 2015. URL <http://arxiv.org/abs/1504.08083>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Kalchbrenner, N., Espeholt, L., Simonyan, K., van den Oord, A., Graves, A., and Kavukcuoglu, K. Neural machine translation in linear time. *arXiv preprint arXiv:1610.10099v2*, 2017.
- Kim, Y., Denton, C., Hoang, L., and Rush, A. M. Structured attention networks. *CoRR*, abs/1702.00887, 2017. URL <http://arxiv.org/abs/1702.00887>.
- Kitaev, N. and Klein, D. Constituency parsing with a self-attentive encoder. *CoRR*, abs/1805.01052, 2018. URL <http://arxiv.org/abs/1805.01052>.
- Lee, K., He, L., Lewis, M., and Zettlemoyer, L. End-to-end neural coreference resolution. *CoRR*, abs/1707.07045, 2017. URL <http://arxiv.org/abs/1707.07045>.
- Lin, C.-Y. and Och, F. J. Orange: A method for evaluating automatic evaluation metrics for machine translation. In *Proceedings of the 20th International Conference on Computational Linguistics, COLING '04*, Stroudsburg, PA, USA, 2004. Association for Computational Linguistics. doi: 10.3115/1220355.1220427. URL <https://doi.org/10.3115/1220355.1220427>.
- Lin, T., Maire, M., Belongie, S. J., Bourdev, L. D., Girshick, R. B., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. URL <http://arxiv.org/abs/1405.0312>.
- Lu, J., Yang, J., Batra, D., and Parikh, D. Neural baby talk. *CoRR*, abs/1803.09845, 2018. URL <http://arxiv.org/abs/1803.09845>.
- Luong, M., Pham, H., and Manning, C. D. Effective approaches to attention-based neural machine translation. *CoRR*, abs/1508.04025, 2015. URL <http://arxiv.org/abs/1508.04025>.
- Niculae, V. and Blondel, M. A regularized framework for sparse and structured neural attention. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3338–3348. Curran Associates, Inc., 2017.
- Pedersoli, M., Lucas, T., Schmid, C., and Verbeek, J. Areas of attention for image captioning. *CoRR*, abs/1612.01033, 2016. URL <http://arxiv.org/abs/1612.01033>.
- Rei, M. and Søgaard, A. Zero-shot sequence labeling: Transferring knowledge from sentences to tokens. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 293–302, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1027. URL <https://www.aclweb.org/anthology/N18-1027>.
- Sharma, P., Ding, N., Goodman, S., and Soricut, R. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pp. 2556–2565, 2018. URL <https://aclanthology.info/papers/P18-1238/p18-1238>.

- Shen, S. and Lee, H. Neural attention models for sequence classification: Analysis and application to key term extraction and dialogue act detection. *CoRR*, abs/1604.00077, 2016. URL <http://arxiv.org/abs/1604.00077>.
- Stern, M., Andreas, J., and Klein, D. A minimal span-based neural constituency parser. *CoRR*, abs/1705.03919, 2017. URL <http://arxiv.org/abs/1705.03919>.
- Szeliski, R. *Computer Vision: Algorithms and Applications*. Springer-Verlag, Berlin, Heidelberg, 1st edition, 2010. ISBN 1848829345, 9781848829343.
- Van Vliet, L. Robust local max-min filters by normalized power-weighted filtering. volume 1, pp. 696–699, 01 2004. doi: 10.1109/ICPR.2004.1334273.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- Vedantam, R., Zitnick, C. L., and Parikh, D. Cider: Consensus-based image description evaluation. *CoRR*, abs/1411.5726, 2014. URL <http://arxiv.org/abs/1411.5726>.
- Viola, P. and Jones, M. Rapid object detection using a boosted cascade of simple features. pp. 511–518, 2001.
- Wang, W. and Chang, B. Graph-based dependency parsing with bidirectional LSTM. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*, 2016. URL <http://aclweb.org/anthology/P/P16/P16-1218.pdf>.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144, 2016. URL <http://arxiv.org/abs/1609.08144>.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A. C., Salakhutdinov, R., Zemel, R. S., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. *CoRR*, abs/1502.03044, 2015. URL <http://arxiv.org/abs/1502.03044>.
- You, Q., Jin, H., Wang, Z., Fang, C., and Luo, J. Image captioning with semantic attention. *CoRR*, abs/1603.03925, 2016. URL <http://arxiv.org/abs/1603.03925>.
- Young, P., Lai, A., Hodosh, M., and Hockenmaier, J. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. 2: 67–78, 01 2014.
- Zheng, H., Fu, J., , and Luo, J. Learning multi-attention convolutional neural network for fine-grained image recognition (iccv 2017 oral). October 2017.