

Input Combination Strategies for Multi-Source Transformer Decoder

Jindřich Libovický and Jindřich Helcl and David Mareček

Charles University, Faculty of Mathematics and Physics

Institute of Formal and Applied Linguistics

Malostranské náměstí 25, 118 00 Prague, Czech Republic

{libovicky, helcl, marecek}@ufal.mff.cuni.cz

Abstract

In multi-source sequence-to-sequence tasks, the attention mechanism can be modeled in several ways. This topic has been thoroughly studied on recurrent architectures. In this paper, we extend the previous work to the encoder-decoder attention in the Transformer architecture. We propose four different input combination strategies for the encoder-decoder attention: serial, parallel, flat, and hierarchical. We evaluate our methods on tasks of multimodal translation and translation with multiple source languages. The experiments show that the models are able to use multiple sources and improve over single source baselines.

1 Introduction

The Transformer model (Vaswani et al., 2017) recently demonstrated superior performance in neural machine translation (NMT) and other sequence generation tasks such as text summarization or image captioning (Kaiser et al., 2017). However, all of these setups consider only a single input to the decoder part of the model.

In the Transformer architecture, the representation of the source sequence is supplied to the decoder through the encoder-decoder attention. This attention sub-layer is applied between the self-attention and feed-forward sub-layers in each Transformer layer. Such arrangement leaves many options for the incorporation of multiple encoders.

So far, attention in sequence-to-sequence learning with multiple source sequences was mostly studied in the context of recurrent neural networks (RNNs). Libovický and Helcl (2017) explicitly capture the distribution over multiple inputs by projecting the input representations to a shared vector space and either computing the attention over all hidden states at once, or hierarchically, using another level of attention applied on the con-

text vectors. Zoph and Knight (2016) employ a gating mechanism for combining the context vectors. Voita et al. (2018) adapted the gating mechanism for use within the Transformer model for context-aware MT. The other approaches are however not directly usable in the Transformer model.

We propose a number of strategies of combining the different sources in the Transformer model. Some of the strategies described in this work are an adaptation of the strategies previously used with recurrent neural networks (Libovický and Helcl, 2017), whereas the rest of them is a novel contribution devised for the Transformer architecture. We test these strategies on multimodal machine translation (MMT) and multi-source machine translation (MSMT) tasks.

This paper is organized as follows. In Section 2, we briefly describe the decoder part of the Transformer model. We propose a number of input combination strategies for the multi-source Transformer model in Section 3. Section 4 describes the experiments we performed, and Section 5 shows the results of quantitative evaluation. An overview of the related work is given in Section 6. We discuss the results and conclude in Section 7.

2 Transformer Decoder

The Transformer architecture is based on the use of attention. Attention, as conceptualized by Vaswani et al. (2017), can be viewed as a soft-lookup function operating on an associative memory. For each query vector in query set Q , the attention computes a set of weighted sums of values V associated with a set of keys K , based on their similarity to the query.

The variant of the attention function used in the Transformer architecture is called *multi-head scaled dot-product* attention. Scaled dot-product

of queries and keys is used as the similarity measure. Given the dimension of the input vectors d , the attention is computed as follows:

$$\mathcal{A}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V. \quad (1)$$

In the multi-head variant, the vectors that represent the queries, keys, and values are linearly transformed to a number of projections (usually with smaller dimension), called *attention heads*. The attention is computed in each head independently and the outputs are concatenated and projected back to the original dimension:

$$\mathcal{A}^h(Q, K, V) = \sum_{i=1}^h C_i W_i^O \quad (2)$$

where $W_i^O \in \mathbb{R}^{d_h \times d}$ are trainable parameter matrices used as projections of the attention head outputs of dimension d_h to the model dimension d , and

$$C_i = \mathcal{A}(QW_i^Q, KW_i^K, VW_i^V) \quad (3)$$

where W^Q , W^K , and $W^V \in \mathbb{R}^{d \times d_h}$, are trainable projection matrices used to project the attention inputs to the attention heads.

The model itself consists of a number of layers, each of which is divided in three sub-layers: self-attention, encoder-decoder (or cross) attention, and a feed-forward layer. Both of the attention types use identical sets for keys and values. The states of the previous layer are used as the query set. The self-attention sub-layer attends to the previous decoder layer (i.e. the sets of queries and keys are identical). Since the decoder works autoregressively from left to right, during training, the self-attention is masked to prevent attending to the future positions in the sequence. The encoder-decoder attention sub-layer attends to the final layer of the encoder. The feed-forward sub-layer consists of a single non-linear projection (usually to a space with larger dimension), followed by a linear projection back to the vector space with the original dimension. The input of each sub-layer is summed with the output, creating a residual connection chain throughout the whole layer stack.

3 Proposed Strategies

We propose four input combination strategies for multi-source variant of the Transformer network,




Figure 1: Schemes of computational steps for the serial, parallel, flat, and hierarchical attention combination in a single layer of the decoder.

as illustrated in Figure 1. Two of them, serial and parallel, model the encoder-decoder attentions independently and are a natural extension of the sub-layer scheme in the transformer decoder. The other two versions, flat and hierarchical, are inspired by approaches proposed for RNNs by Libovický and Helcl (2017) and model joint distributions over the inputs.

Serial. The serial strategy (Figure 1a) computes the encoder-decoder attention one by one for each input encoder. The query set of the first cross-attention is the set of the context vectors computed by the preceding self-attention. The query set of each subsequent cross-attention is the output of the preceding sub-layer. All of these sub-layers are interconnected with residual connections.

Parallel. In the parallel combination strategy (Figure 1b), the model attends to each encoder independently and then sums up the context vectors. Each encoder is attended using the same set of queries, i.e. the output of the self-attention sub-layer. Residual connection link is used between the queries and the summed context vectors from the parallel attention.

$$\mathcal{A}_{para}^h(Q, K_{1:n}, V_{1:n}) = \sum_{i=1}^n \mathcal{A}^h(Q, K_i, V_i) \quad (4)$$

Flat. The encoder-decoder attention in the flat combination strategy (Figure 1c) uses all the states of all input encoders as a single set of keys and values. Thus, the attention models a joint distribution over a flattened set of all encoder states. Unlike the approach taken in the recurrent setup (Libovický and Helcl, 2017), where the flat combination strategy requires an explicit projection of the encoder states to a shared vector space, in the Transformer models, the vector spaces of all layers are tied with residual connections. Therefore, the intermediate projection of the states of each encoder is not necessary.

$$K_{flat} = V_{flat} = \text{concat}_i(K_i) \quad (5)$$

$$\mathcal{A}_{flat}^h(Q, K_{1:n}, V_{1:n}) = \mathcal{A}^h(Q, K_{flat}, V_{flat}) \quad (6)$$

Hierarchical. In the hierarchical combination (Figure 1d), we first compute the attention independently over each input. The resulting contexts are then treated as states of another input and the

attention is computed once again over these states.

$$K_{hier} = V_{hier} = \text{concat}_i(\mathcal{A}^h(Q, K_i, V_i)) \quad (7)$$

$$\mathcal{A}_{hier}^h(Q, K_{1:n}, V_{1:n}) = \mathcal{A}^h(Q, K_{hier}, V_{hier}) \quad (8)$$

4 Experiments

We conduct our experiments on two different tasks: multimodal translation and multi-source machine translation. We use Neural Monkey (Helcl and Libovický, 2017)¹ for design, training, and evaluation of the experiments.

In all experiments, the encoder part of the network follows the Transformer architecture as described by Vaswani et al. (2017).

We optimize the model parameters using Adam optimizer (Kingma and Ba, 2014) with initial learning rate 0.2, and Noam learning rate decay (Vaswani et al., 2017) with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$, and 4,000 warm-up steps. The size of a mini-batch size of 32 for MMT, and 24 for multi-source MT experiments.

During decoding, we use beam search of width 10 and length normalization of 1.0 (Wu et al., 2016).

4.1 Multimodal Translation

The goal of MMT (Specia et al., 2016) is translating image captions from one language into another given both the source and image as the input. We use Multi30k dataset (Elliott et al., 2016) containing triplets of images, English captions and their English translations into German, French and Czech. The dataset contains 29k triplets for training, 1,014 for validation and a test set of 1,000. We experiment with all language pairs available in this dataset.

We extract image feature using the last convolutional layer of the ResNet network (He et al., 2016) trained for ImageNet classification. We apply a linear projection into 512 dimensions on the image representation, so it has the same dimension as the rest of the model. For each language pair, we create a shared wordpiece-based vocabulary of approximately 40k subwords. We share the embedding matrices across the languages and we use the transposed embedding matrix as the output projection matrix as proposed by Press and Wolf (2017).

We use 6 layers in the textual encoder and decoder, and set the model dimension to 512. We

¹<http://github.com/ufal/neuralmonkey>

set the dimension of the hidden layers in the feed-forward sub-layers to 4096. We use 16 heads in the attention layers.

During the evaluation, we follow the preprocessing used in WMT Multimodal Translation Shared Task (Specia et al., 2016).

Conclusions of previous work show (Elliott and Kádár, 2017) that the improved performance of the multimodal models compared to textual models can come from improving the input representation. In order to test whether it is also the case with our models or the models explicitly use the visual input, we perform an adversarial evaluation similar to Elliott (2018). We evaluate the model while providing a random image and observe how it affects the score and observe whether their quality drops.

4.2 Multi-Source MT

In this set of experiment, we attempt to generate a sentence in a target language, given equivalent sentences in multiple source languages.

We use the Europarl corpus (Tiedemann, 2012) for training and testing the MSMT. We use Spanish, French, German, and English as source languages and Czech as a target language. We selected an intersection of the bilingual sub-corpora using English as a pivot language. Our dataset contains 511k 5-tuples of sentences for training, 1k for validation and another 1k for testing.

Due of the memory demands of having four encoders, we use a smaller model than in the previous experiment. The encoders only have 4 layers and the decoder has 6 layers with embeddings size 256, feed-forward layers dimension 2048, and 8 attention heads. We use a shared word-piece vocabulary of 48k subwords. As in the MMT experiments, the transposition of the embedding matrix is reused as the parameters of the output projection layer (Press and Wolf, 2017).

We use bilingual English-to-Czech translation as a single source baseline. The baseline uses vocabulary of 42k subwords from Czech and English only.

Similarly to the MMT, we also perform adversarial evaluation. To evaluate the importance of the source languages for the translation quality, when randomizing one of the source languages.

5 Results

We evaluate the results using BLEU (Papineni et al., 2002) and METEOR (Denkowski and Lavie, 2011) as implemented in MultEval.² The results of the MMT task are tabulated in Table 1. The results of the multi-source MT are shown in Table 2.

In MMT, the input combination significantly surpassed the text-only baseline in English-to-French translation. The performance in other target languages is only slightly better than the textual baseline.

The only worse score was achieved by the flat combination strategy. We hypothesize this might be because the optimization failed to find a common representation of the input modalities that could be used to compute the joint distribution.

The adversarial evaluation with randomly selected input images shows that all our models rely on both inputs while generating the target sentence and that providing incorrect visual input harms the model performance. The modality gating in the hierarchical attention combination seems to make the models more robust to noisy visual input.

In the multi-source translation task, all the proposed strategies perform better than single-source translation from English to Czech. Among the combination strategies, the best-scoring is the serial stacking of the attentions. In multimodal translation, the flat combination has shown to be the best-performing strategy.

Analysis of the attention distribution shows that the serial strategy use information from all source languages. The parallel strategy almost does not use the Spanish source and the flat strategy prefers the English source. The hierarchical strategy uses information from all source languages, however the attentions are sometimes more fuzzy than in the previous strategies. Figure 2 shows what source languages were attended on different layers of the encoder. Other examples of the attention visualization are shown in Appendix A.

The adversarial evaluation shows all the models used English as a primary source. Providing incorrect English source harms. Introducing noise into other languages affects the score in much smaller scale.

	MMT: en→de			MMT: en→fr			MMT: en→cs		
	BLEU	METEOR	adv.BLEU	BLEU	METEOR	adv.BLEU	BLEU	METEOR	adv.BLEU
baseline	38.3 ± .8	56.7 ± .7	—	59.6 ± .9	72.7 ± .7	—	30.9 ± .8	29.5 ± .4	—
serial	38.7 ± .9	57.2 ± .6	37.3 ± .6	60.8 ± .9	75.1 ± .6	58.9 ± .9	31.0 ± .8	29.9 ± .4	29.7 ± .8
parallel	38.6 ± .9	57.4 ± .7	38.2 ± .8	60.2 ± .9	74.9 ± .6	58.9 ± .9	31.1 ± .9	30.0 ± .4	30.4 ± .8
flat	37.1 ± .8	56.5 ± .6	35.7 ± .8	58.0 ± .9	73.3 ± .7	57.0 ± .9	29.9 ± .8	29.0 ± .4	28.2 ± .8
hierarchical	38.5 ± .8	56.5 ± .6	38.1 ± .8	60.8 ± .9	75.1 ± .6	60.2 ± .9	31.3 ± .9	30.0 ± .4	31.0 ± .8

Table 1: Quantitative results of the MMT experiments on the 2016 test set. Column ‘adv. BLEU’ is an adversarial evaluation with randomized image input.

	MSMT		Adversarial evaluation (BLEU)			
	BLEU	METEOR	en	de	fr	es
baseline	16.5 ± .5	20.5 ± .3	—	—	—	—
serial	20.5 ± .6	23.5 ± .5	8.1 ± .4	19.7 ± .5	19.5 ± .6	18.4 ± .5
parallel	20.5 ± .6	23.3 ± .3	1.4 ± .2	18.7 ± .5	17.9 ± .5	20.3 ± .5
flat	20.4 ± .6	23.3 ± .3	0.2 ± .1	19.9 ± .6	20.0 ± .6	19.6 ± .5
hierarchical	19.4 ± .5	22.7 ± .3	4.2 ± .3	18.3 ± .5	18.3 ± .5	15.3 ± .5

Table 2: Quantitative results of the MMT experiment. The adversarial evaluation shows the BLEU score when one input language was changed randomly.




Figure 2: Attention over contexts in the hierarchical strategy over the decoder layers.

6 Related Work

MMT was so far solved only within the RNN-based architectures. Elliott et al. (2015) report significant improvements with a non-attentive model. With attentive models (Bahdanau et al., 2014), the additional visual information usually did not improve the models significantly (Caglayan et al., 2016; Helcl and Libovický, 2017) in terms of BLEU score. Our models slightly outperform these models in the single model setup.

²<https://github.com/jhclark/multeval>

Except for using the image features direct input to the model, they can be used as an auxiliary objective (Elliott and Kádár, 2017). In this setup, the visually grounded representation, improves the MMT significantly, achieving similar results that our models achieved using only the Multi30k dataset.

To our knowledge, multi-source MT has also been studied only using the RNN-based models. Dabre et al. (2017) use simple concatenation of source sentences in various languages and process them with a single multilingual encoder.

Zoph and Knight (2016) try context concatenation and hierarchical gating method for combining context vectors in attention models with multiple inputs encoded by separate encoders. In all of their experiments, the multi-source methods significantly surpass the single-source baseline. Nishimura et al. (2018) extend the former approach for situations when of the source languages is missing, so that the translation system does not overly rely on a single source language like some of the models presented in this work.

7 Conclusions

We proposed several input combination strategies for multi-source sequence-to-sequence learning using the Transformer model (Vaswani et al., 2017). Two of the strategies are a straightforward extension of cross-attention in the Trans-

former model: the cross-attentions are combined either serially interleaved by residual connections or in parallel. The two remaining strategies are an adaptation of the flat and the hierarchical attention combination strategies introduced by Libovický and Helcl (2017) in context of recurrent sequence-to-sequence models.

The results on the MMT task show similar properties as in RNN-based models (Caglayan et al., 2017; Libovický and Helcl, 2017). Adding visual features significantly improves translation into French and brings minor improvements on other language pairs. All the attention combinations perform similarly with the exception of the flat strategy which probably struggles with learning a shared representation of the input tokens and the image representation.

Evaluation on multi-source MT shows significant improvements over the single-source baseline. However, the adversarial evaluation suggests that the model relies heavily on the English input and only uses the additional source languages for minor modifications of the output. All attention combinations performed similarly.

Acknowledgments

This research received support from the grant No. 18-02196S and No. P103/12/G084 of the Grant Agency of the Czech Republic, and the grant No. 976518 of the Grant Agency of the Charles University. This research was partially supported by SVV project number 260 453.

References


- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Ozan Caglayan, Walid Aransa, Adrien Bardet, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, Marc Masana, Luis Herranz, and Joost van de Weijer. 2017. Lium-cvc submissions for wmt17 multimodal translation task. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 432–439, Copenhagen, Denmark. Association for Computational Linguistics.
- Ozan Caglayan, Walid Aransa, Yaxing Wang, Marc Masana, Mercedes García-Martínez, Fethi Bougares, Loïc Barrault, and Joost van de Weijer. 2016. Does multimodality help human and machine for translation and image captioning? In *Proceedings of the First Conference on Machine Translation*, pages 627–633, Berlin, Germany. Association for Computational Linguistics.
- Raj Dabre, Fabien Cromières, and Sadao Kurohashi. 2017. Enabling multi-source neural machine translation by concatenating source sentences in multiple languages. *CoRR*, abs/1702.06135.
- Michael Denkowski and Alon Lavie. 2011. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 85–91, Edinburgh, United Kingdom. Association for Computational Linguistics.
- Desmond Elliott. 2018. Adversarial evaluation of multimodal machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium. Association for Computational Linguistics.
- Desmond Elliott, Stella Frank, and Eva Hasler. 2015. Multi-language image description with neural sequence models. *CoRR*, abs/1510.04709.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. Multi30k: Multilingual english-german image descriptions. In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Desmond Elliott and Ákos Kádár. 2017. Imagination improves multimodal translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 130–141, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778. IEEE Computer Society.
- Jindřich Helcl and Jindřich Libovický. 2017. Neural Monkey: An open-source tool for sequence learning. *The Prague Bulletin of Mathematical Linguistics*, 107:5–17.
- Jindřich Helcl and Jindřich Libovický. 2017. CUNI system for the WMT17 multimodal translation task. In *Proceedings of the Second Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 450–457, Copenhagen, Denmark. Association for Computational Linguistics.
- Lukasz Kaiser, Aidan N. Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017. One model to learn them all. *CoRR*, abs/1706.05137.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.

- Jindřich Libovický and Jindřich Helcl. 2017. Attention strategies for multi-source sequence-to-sequence learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 196–202, Vancouver, Canada. Association for Computational Linguistics.
- Yuta Nishimura, Katsuhito Sudoh, Graham Neubig, and Satoshi Nakamura. 2018. Multi-source neural machine translation with missing data. In *The Second Workshop on Neural Machine Translation and Generation (WNMT)*, Melbourne, Australia.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163. Association for Computational Linguistics.
- Lucia Specia, Stella Frank, Khalil Sima'an, and Desmond Elliott. 2016. A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation*, pages 543–553, Berlin, Germany. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6000–6010. Curran Associates, Inc.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 30–34, San Diego, California. Association for Computational Linguistics.


A Attention Visualizations

We show cross-attention visualizations for the four proposed combination strategies on Multi-source MT. The Czech target wordpieces are in rows, the source Spanish, French, German, and English wordpieces are concatenated and shown in columns. These attentions were taken from the


decoder's fourth layer and were averaged across the individual heads. For serial and parallel strategy the cross-attention weights sum to one for each language separately, the flat strategy has only one common cross-attention, and for the hierarchical strategy visualization the cross-attention weights for individual languages are multiplied by the weights of the attention over contexts.




a) serial



b) parallel



c) flat



d) hierarchical