

# The Best of Both Worlds: Combining Recent Advances in Neural Machine Translation

Mia Xu Chen \*

Orhan Firat \*

Ankur Bapna \*

Melvin Johnson

Wolfgang Macherey

George Foster

Llion Jones

Niki Parmar

Noam Shazeer

Ashish Vaswani

Jakob Uszkoreit

Lukasz Kaiser

Mike Schuster Zhifeng Chen Yonghui Wu Macduff Hughes

miachen, orhanf, ankurbpn, yonghui@google.com

Google AI

## Abstract

The past year has witnessed rapid advances in sequence-to-sequence (seq2seq) modeling for Machine Translation (MT). The classic RNN-based approaches to MT were first out-performed by the convolutional seq2seq model, which was then out-performed by the more recent Transformer model. Each of these new approaches consists of a fundamental architecture accompanied by a set of modeling and training techniques that are in principle applicable to other seq2seq architectures. In this paper, we tease apart the new architectures and their accompanying techniques in two ways. First, we identify several key modeling and training techniques, and apply them to the RNN architecture, yielding a new RNMT+ model that outperforms all of the three fundamental architectures on the benchmark WMT'14 English→French and English→German tasks. Second, we analyze the properties of each fundamental seq2seq architecture and devise new hybrid architectures intended to combine their strengths. Our hybrid models obtain further improvements, outperforming the RNMT+ model on both benchmark datasets.

## 1 Introduction

In recent years, the emergence of seq2seq models (Kalchbrenner and Blunsom, 2013; Sutskever et al., 2014; Cho et al., 2014) has revolutionized the field of MT by replacing traditional phrase-based approaches with neural machine translation (NMT) systems based on the encoder-decoder paradigm. In the first architectures that surpassed

the quality of phrase-based MT, both the encoder and decoder were implemented as Recurrent Neural Networks (RNNs), interacting via a soft-attention mechanism (Bahdanau et al., 2015). The RNN-based NMT approach, or RNMT, was quickly established as the de-facto standard for NMT, and gained rapid adoption into large-scale systems in industry, e.g. Baidu (Zhou et al., 2016), Google (Wu et al., 2016), and Systran (Crego et al., 2016).

Following RNMT, convolutional neural network based approaches (LeCun and Bengio, 1998) to NMT have recently drawn research attention due to their ability to fully parallelize training to take advantage of modern fast computing devices, such as GPUs and Tensor Processing Units (TPUs) (Jouppi et al., 2017). Well known examples are ByteNet (Kalchbrenner et al., 2016) and ConvS2S (Gehring et al., 2017). The ConvS2S model was shown to outperform the original RNMT architecture in terms of quality, while also providing greater training speed.

Most recently, the Transformer model (Vaswani et al., 2017), which is based solely on a self-attention mechanism (Parikh et al., 2016) and feed-forward connections, has further advanced the field of NMT, both in terms of translation quality and speed of convergence.

In many instances, new architectures are accompanied by a novel set of techniques for performing training and inference that have been carefully optimized to work in concert. This ‘bag of tricks’ can be crucial to the performance of a proposed architecture, yet it is typically under-documented and left for the enterprising researcher to discover in publicly released code (if any) or through anecdotal evidence. This is not simply a problem for reproducibility; it obscures the central scientific question of how much of the observed gains come from the new architecture

\* Equal contribution.

and how much can be attributed to the associated training and inference techniques. In some cases, these new techniques may be broadly applicable to other architectures and thus constitute a major, though implicit, contribution of an architecture paper. Clearly, they need to be considered in order to ensure a fair comparison across different model architectures.

In this paper, we therefore take a step back and look at which techniques and methods contribute significantly to the success of recent architectures, namely ConvS2S and Transformer, and explore applying these methods to other architectures, including RNMT models. In doing so, we come up with an enhanced version of RNMT, referred to as RNMT+, that significantly outperforms all individual architectures in our setup. We further introduce new architectures built with different components borrowed from RNMT+, ConvS2S and Transformer. In order to ensure a fair setting for comparison, all architectures were implemented in the same framework, use the same pre-processed data and apply no further post-processing as this may confound bare model performance.

Our contributions are three-fold:

1. In ablation studies, we quantify the effect of several modeling improvements (including multi-head attention and layer normalization) as well as optimization techniques (such as synchronous replica training and label-smoothing), which are used in recent architectures. We demonstrate that these techniques are applicable across different model architectures.
2. Combining these improvements with the RNMT model, we propose the new RNMT+ model, which significantly outperforms all fundamental architectures on the widely-used WMT'14 En→Fr and En→De benchmark datasets. We provide a detailed model analysis and comparison of RNMT+, ConvS2S and Transformer in terms of model quality, model size, and training and inference speed.
3. Inspired by our understanding of the relative strengths and weaknesses of individual model architectures, we propose new model architectures that combine components from the RNMT+ and the Transformer model, and achieve better results than both individual architectures.

We quickly note two prior works that provided empirical solutions to the difficulty of training NMT architectures (specifically RNMT). In (Britz et al., 2017) the authors systematically explore which elements of NMT architectures have a significant impact on translation quality. In (Denkowski and Neubig, 2017) the authors recommend three specific techniques for strengthening NMT systems and empirically demonstrated how incorporating those techniques improves the reliability of the experimental results.

## 2 Background

In this section, we briefly discuss the commonly used NMT architectures.

### 2.1 RNN-based NMT Models - RNMT

RNMT models are composed of an encoder RNN and a decoder RNN, coupled with an attention network. The encoder summarizes the input sequence into a set of vectors while the decoder conditions on the encoded input sequence through an attention mechanism, and generates the output sequence one token at a time.

The most successful RNMT models consist of stacked RNN encoders with one or more bidirectional RNNs (Schuster and Paliwal, 1997; Graves and Schmidhuber, 2005), and stacked decoders with unidirectional RNNs. Both encoder and decoder RNNs consist of either LSTM (Hochreiter and Schmidhuber, 1997; Gers et al., 2000) or GRU units (Cho et al., 2014), and make extensive use of residual (He et al., 2015) or highway (Srivastava et al., 2015) connections.

In Google-NMT (GNMT) (Wu et al., 2016), the best performing RNMT model on the datasets we consider, the encoder network consists of one bi-directional LSTM layer, followed by 7 uni-directional LSTM layers. The decoder is equipped with a single attention network and 8 uni-directional LSTM layers. Both the encoder and the decoder use residual skip connections between consecutive layers.

In this paper, we adopt GNMT as the starting point for our proposed RNMT+ architecture.

### 2.2 Convolutional NMT Models - ConvS2S

In the most successful convolutional sequence-to-sequence model (Gehring et al., 2017), both the encoder and decoder are constructed by stacking multiple convolutional layers, where each layer

contains 1-dimensional convolutions followed by a gated linear units (GLU) (Dauphin et al., 2016). Each decoder layer computes a separate dot-product attention by using the current decoder layer output and the final encoder layer outputs. Positional embeddings are used to provide explicit positional information to the model. Following the practice in (Gehring et al., 2017), we scale the gradients of the encoder layers to stabilize training. We also use residual connections across each convolutional layer and apply weight normalization (Salimans and Kingma, 2016) to speed up convergence. We follow the public ConvS2S codebase<sup>1</sup> in our experiments.

### 2.3 Conditional Transformation-based NMT Models - Transformer

The Transformer model (Vaswani et al., 2017) is motivated by two major design choices that aim to address deficiencies in the former two model families: (1) Unlike RNMT, but similar to the ConvS2S, the Transformer model avoids any sequential dependencies in both the encoder and decoder networks to maximally parallelize training. (2) To address the limited context problem (limited receptive field) present in ConvS2S, the Transformer model makes pervasive use of self-attention networks (Parikh et al., 2016) so that each position in the current layer has access to information from all other positions in the previous layer.

The Transformer model still follows the encoder-decoder paradigm. Encoder transformer layers are built with two sub-modules: (1) a self-attention network and (2) a feed-forward network. Decoder transformer layers have an additional cross-attention layer sandwiched between the self-attention and feed-forward layers to attend to the encoder outputs.

There are two details which we found very important to the model’s performance: (1) Each sub-layer in the transformer (i.e. self-attention, cross-attention, and the feed-forward sub-layer) follows a strict computation sequence: *normalize* → *transform* → *dropout* → *residual-add*. (2) In addition to per-layer normalization, the final encoder output is again normalized to prevent a blow up after consecutive residual additions.

In this paper, we follow the latest version of the

<sup>1</sup><https://github.com/facebookresearch/fairseq-py>

Transformer model in the Tensor2Tensor<sup>2</sup> codebase.

### 2.4 A Theory-Based Characterization of NMT Architectures

From a theoretical point of view, RNNs belong to the most expressive members of the neural network family (Siegelmann and Sontag, 1995)<sup>3</sup>. Possessing an infinite Markovian structure (and thus an infinite receptive fields) equips them to model sequential data (Elman, 1990), especially natural language (Grefenstette et al., 2015) effectively. In practice, RNNs are notoriously hard to train (Hochreiter, 1991; Bengio et al., 1994; Hochreiter et al., 2001), confirming the well known dilemma of trainability versus expressivity.

Convolutional layers are adept at capturing local context and local correlations by design. A fixed and narrow receptive field for each convolutional layer limits their capacity when the architecture is shallow. In practice, this weakness is mitigated by stacking more convolutional layers (e.g. 15 layers as in the ConvS2S model), which makes the model harder to train and demands meticulous initialization schemes and carefully designed regularization techniques.

The transformer network is capable of approximating arbitrary squashing functions (Hornik et al., 1989), and can be considered a strong feature extractor with extended receptive fields capable of linking salient features from the entire sequence. On the other hand, lacking a memory component (as present in the RNN models) prevents the network from modeling a state space, reducing its theoretical strength as a sequence model, thus it requires additional positional information (e.g. sinusoidal positional encodings).

Above theoretical characterizations will drive our explorations in the following sections.

## 3 Experiment Setup

We train our models on the standard WMT’14 En→Fr and En→De datasets that comprise 36.3M and 4.5M sentence pairs, respectively. Each sentence was encoded into a sequence of sub-word units obtained by first tokenizing the sentence with the Moses tokenizer, then splitting tokens into sub-word units (also known as “wordpieces”) using the approach described in (Schuster and Nakajima, 2012).

<sup>2</sup><https://github.com/tensorflow/tensor2tensor>

<sup>3</sup>Assuming that data complexity is satisfied.

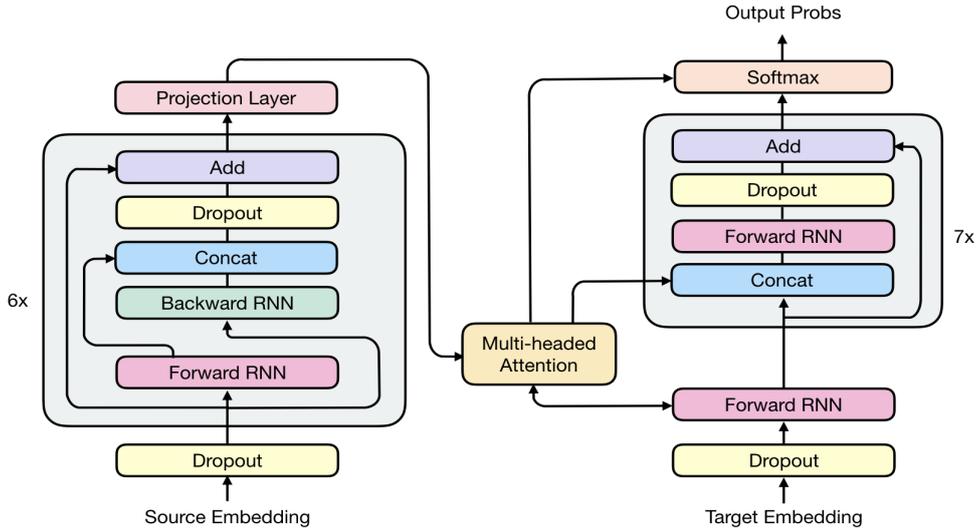


Figure 1: Model architecture of RNMT+. On the left side, the encoder network has 6 bidirectional LSTM layers. At the end of each bidirectional layer, the outputs of the forward layer and the backward layer are concatenated. On the right side, the decoder network has 8 unidirectional LSTM layers, with the first layer used for obtaining the attention context vector through multi-head additive attention. The attention context vector is then fed directly into the rest of the decoder layers as well as the softmax layer.

We use a shared vocabulary of 32K sub-word units for each source-target language pair. No further manual or rule-based post processing of the output was performed beyond combining the sub-word units to generate the targets. We report all our results on newstest 2014, which serves as the test set. A combination of newstest 2012 and newstest 2013 is used for validation.

To evaluate the models, we compute the BLEU metric on tokenized, true-case output.<sup>4</sup> For each training run, we evaluate the model every 30 minutes on the dev set. Once the model converges, we determine the best window based on the average dev-set BLEU score over 21 consecutive evaluations. We report the mean test score and standard deviation over the selected window. This allows us to compare model architectures based on their mean performance after convergence rather than individual checkpoint evaluations, as the latter can be quite noisy for some models.

To enable a fair comparison of architectures, we use the same pre-processing and evaluation methodology for all our experiments. We refrain from using checkpoint averaging (exponential moving averages of parameters) (Junczys-Dowmunt et al., 2016) or checkpoint ensembles (Jean et al., 2015; Chen et al., 2017) to focus on

<sup>4</sup>This procedure is used in the literature to which we compare (Gehring et al., 2017; Wu et al., 2016).

evaluating the performance of individual models.

## 4 RNMT+

### 4.1 Model Architecture of RNMT+

The newly proposed RNMT+ model architecture is shown in Figure 1. Here we highlight the key architectural choices that are different between the RNMT+ model and the GNMT model. There are 6 bidirectional LSTM layers in the encoder instead of 1 bidirectional LSTM layer followed by 7 unidirectional layers as in GNMT. For each bidirectional layer, the outputs of the forward layer and the backward layer are concatenated before being fed into the next layer. The decoder network consists of 8 unidirectional LSTM layers similar to the GNMT model. Residual connections are added to the third layer and above for both the encoder and decoder. Inspired by the Transformer model, per-gate layer normalization (Ba et al., 2016) is applied within each LSTM cell. Our empirical results show that layer normalization greatly stabilizes training. No non-linearity is applied to the LSTM output. A projection layer is added to the encoder final output.<sup>5</sup> Multi-head additive attention is used instead of the single-head attention in the GNMT model. Similar to GNMT, we use the

<sup>5</sup>Additional projection aims to reduce the dimensionality of the encoder output representations to match the decoder stack dimension.

bottom decoder layer and the final encoder layer output after projection for obtaining the recurrent attention context. In addition to feeding the attention context to all decoder LSTM layers, we also feed it to the softmax by concatenating it with the layer input. This is important for both the quality of the models with multi-head attention and the stability of the training process.

Since the encoder network in RNMT+ consists solely of bi-directional LSTM layers, model parallelism is not used during training. We compensate for the resulting longer per-step time with increased data parallelism (more model replicas), so that the overall time to reach convergence of the RNMT+ model is still comparable to that of GNMT.

We apply the following regularization techniques during training.

- **Dropout:** We apply dropout to both embedding layers and each LSTM layer output before it is added to the next layer’s input. Attention dropout is also applied.
- **Label Smoothing:** We use uniform label smoothing with an uncertainty=0.1 (Szegedy et al., 2015). Label smoothing was shown to have a positive impact on both Transformer and RNMT+ models, especially in the case of RNMT+ with multi-head attention. Similar to the observations in (Chorowski and Jaitly, 2016), we found it beneficial to use a larger beam size (e.g. 16, 20, etc.) during decoding when models are trained with label smoothing.
- **Weight Decay:** For the WMT’14 En→De task, we apply L2 regularization to the weights with  $\lambda = 10^{-5}$ . Weight decay is only applied to the En→De task as the corpus is smaller and thus more regularization is required.

We use the Adam optimizer (Kingma and Ba, 2014) with  $\beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 10^{-6}$  and vary the learning rate according to this schedule:

$$lr = 10^{-4} \cdot \min \left( 1 + \frac{t \cdot (n - 1)}{np}, n, n \cdot (2n)^{\frac{s-nt}{e-s}} \right) \quad (1)$$

Here,  $t$  is the current step,  $n$  is the number of concurrent model replicas used in training,  $p$  is the number of warmup steps,  $s$  is the start step of the exponential decay, and  $e$  is the end step of the decay. Specifically, we first increase the learning rate linearly during the number of warmup steps, keep

it a constant until the decay start step  $s$ , then exponentially decay until the decay end step  $e$ , and keep it at  $5 \cdot 10^{-5}$  after the decay ends. This learning rate schedule is motivated by a similar schedule that was successfully applied in training the Resnet-50 model with a very large batch size (Goyal et al., 2017).

In contrast to the asynchronous training used for GNMT (Dean et al., 2012), we train RNMT+ models with synchronous training (Chen et al., 2016). Our empirical results suggest that when hyper-parameters are tuned properly, synchronous training often leads to improved convergence speed and superior model quality.

To further stabilize training, we also use adaptive gradient clipping. We discard a training step completely if an anomaly in the gradient norm value is detected, which is usually an indication of an imminent gradient explosion. More specifically, we keep track of a moving average and a moving standard deviation of the log of the gradient norm values, and we abort a step if the norm of the gradient exceeds four standard deviations of the moving average.

## 4.2 Model Analysis and Comparison

In this section, we compare the results of RNMT+ with ConvS2S and Transformer.

All models were trained with synchronous training. RNMT+ and ConvS2S were trained with 32 NVIDIA P100 GPUs while the Transformer Base and Big models were trained using 16 GPUs.

For RNMT+, we use sentence-level cross-entropy loss. Each training batch contained 4096 sentence pairs (4096 source sequences and 4096 target sequences). For ConvS2S and Transformer models, we use token-level cross-entropy loss. Each training batch contained 65536 source tokens and 65536 target tokens. For the GNMT baselines on both tasks, we cite the largest BLEU score reported in (Wu et al., 2016) without reinforcement learning.

Table 1 shows our results on the WMT’14 En→Fr task. Both the Transformer Big model and RNMT+ outperform GNMT and ConvS2S by about 2 BLEU points. RNMT+ is slightly better than the Transformer Big model in terms of its mean BLEU score. RNMT+ also yields a much lower standard deviation, and hence we observed much less fluctuation in the training curve. It takes approximately 3 days for the Transformer

Base model to converge, while both RNMT+ and the Transformer Big model require about 5 days to converge. Although the batching schemes are quite different between the Transformer Big and the RNMT+ model, they have processed about the same amount of training samples upon convergence.

Model	Test BLEU	Epochs	Training Time
GNMT	38.95	-	-
ConvS2S <sup>7</sup>	39.49 ± 0.11	62.2	438h
Trans. Base	39.43 ± 0.17	20.7	90h
Trans. Big <sup>8</sup>	40.73 ± 0.19	8.3	120h
RNMT+	41.00 ± 0.05	8.5	120h

Table 1: Results on WMT14 En→Fr. The numbers before and after ‘±’ are the mean and standard deviation of test BLEU score over an evaluation window. Note that Transformer models are trained using 16 GPUs, while ConvS2S and RNMT+ are trained using 32 GPUs.

Table 2 shows our results on the WMT’14 En→De task. The Transformer Base model improves over GNMT and ConvS2S by more than 2 BLEU points while the Big model improves by over 3 BLEU points. RNMT+ further outperforms the Transformer Big model and establishes a new state of the art with an averaged value of 28.49. In this case, RNMT+ converged slightly faster than the Transformer Big model and maintained much more stable performance after convergence with a very small standard deviation, which is similar to what we observed on the En-Fr task.

Table 3 summarizes training performance and model statistics. The Transformer Base model

<sup>6</sup>Since the ConvS2S model convergence is very slow we did not explore further tuning on En→Fr, and validated our implementation on En→De.

<sup>7</sup>The BLEU scores for Transformer model are slightly lower than those reported in (Vaswani et al., 2017) due to four differences:

1) We report the mean test BLEU score using the strategy described in section 3.

2) We did not perform checkpoint averaging since it would be inconsistent with our evaluation for other models.

3) We avoided any manual post-processing, like unicode normalization using Moses replace-unicode-punctuation.perl or output tokenization using Moses tokenizer.perl, to rule out its effect on the evaluation. We observed a significant BLEU increase (about 0.6) on applying these post processing techniques.

4) In (Vaswani et al., 2017), reported BLEU scores are calculated using mteval-v13a.pl from Moses, which re-tokenizes its input.

Model	Test BLEU	Epochs	Training Time
GNMT	24.67	-	-
ConvS2S	25.01 ± 0.17	38	20h
Trans. Base	27.26 ± 0.15	38	17h
Trans. Big	27.94 ± 0.18	26.9	48h
RNMT+	28.49 ± 0.05	24.6	40h

Table 2: Results on WMT14 En→De. Note that Transformer models are trained using 16 GPUs, while ConvS2S and RNMT+ are trained using 32 GPUs.

is the fastest model in terms of training speed. RNMT+ is slower to train than the Transformer Big model on a per-GPU basis. However, since the RNMT+ model is quite stable, we were able to offset the lower per-GPU throughput with higher concurrency by increasing the number of model replicas, and hence the overall time to convergence was not slowed down much. We also computed the number of floating point operations (FLOPs) in the model’s forward path as well as the number of total parameters for all architectures (cf. Table 3). RNMT+ requires fewer FLOPs than the Transformer Big model, even though both models have a comparable number of parameters.

Model	Examples/s	FLOPs	Params
ConvS2S	80	15.7B	263.4M
Trans. Base	160	6.2B	93.3M
Trans. Big	50	31.2B	375.4M
RNMT+	30	28.1B	378.9M

Table 3: Performance comparison. Examples/s are normalized by the number of GPUs used in the training job. FLOPs are computed assuming that source and target sequence length are both 50.

## 5 Ablation Experiments

In this section, we evaluate the importance of four main techniques for both the RNMT+ and the Transformer Big models. We believe that these techniques are universally applicable across different model architectures, and should always be employed by NMT practitioners for best performance.

We take our best RNMT+ and Transformer Big models and remove each one of these techniques independently. By doing this we hope to learn two things about each technique: (1) How much does

it affect the model performance? (2) How useful is it for stable training of other techniques and hence the final model?

Model	RNMT+	Trans. Big
Baseline	41.00	40.73
- Label Smoothing	40.33	40.49
- Multi-head Attention	40.44	39.83
- Layer Norm.	*	*
- Sync. Training	39.68	*

Table 4: Ablation results of RNMT+ and the Transformer Big model on WMT’14 En → Fr. We report average BLEU scores on the test set. An asterisk ‘\*’ indicates an unstable training run (training halts due to non-finite elements).

From Table 4 we draw the following conclusions about the four techniques:

- **Label Smoothing** We observed that label smoothing improves both models, leading to an average increase of 0.7 BLEU for RNMT+ and 0.2 BLEU for Transformer Big models.
- **Multi-head Attention** Multi-head attention contributes significantly to the quality of both models, resulting in an average increase of 0.6 BLEU for RNMT+ and 0.9 BLEU for Transformer Big models.
- **Layer Normalization** Layer normalization is most critical to stabilize the training process of either model, especially when multi-head attention is used. Removing layer normalization results in unstable training runs for both models. Since by design, we remove one technique at a time in our ablation experiments, we were unable to quantify how much layer normalization helped in either case. To be able to successfully train a model without layer normalization, we would have to adjust other parts of the model and retune its hyper-parameters.
- **Synchronous training** Removing synchronous training has different effects on RNMT+ and Transformer. For RNMT+, it results in a significant quality drop, while for the Transformer Big model, it causes the model to become unstable. We also notice that synchronous training is only successful when coupled with a tailored learning rate schedule that has a warmup stage at the beginning (cf. Eq. 1 for RNMT+ and

Eq. 2 for Transformer). For RNMT+, removing this warmup stage during synchronous training causes the model to become unstable.

## 6 Hybrid NMT Models

In this section, we explore hybrid architectures that shed some light on the salient behavior of each model family. These hybrid models outperform the individual architectures on both benchmark datasets and provide a better understanding of the capabilities and limitations of each model family.

### 6.1 Assessing Individual Encoders and Decoders

In an encoder-decoder architecture, a natural assumption is that the role of an encoder is to build feature representations that can best encode the meaning of the source sequence, while a decoder should be able to process and interpret the representations from the encoder and, at the same time, track the current target history. Decoding is inherently auto-regressive, and keeping track of the state information should therefore be intuitively beneficial for conditional generation.

We set out to study which family of encoders is more suitable to extract rich representations from a given input sequence, and which family of decoders can make the best of such rich representations. We start by combining the encoder and decoder from different model families. Since it takes a significant amount of time for a ConvS2S model to converge, and because the final translation quality was not on par with the other models, we focus on two types of hybrids only: Transformer encoder with RNMT+ decoder and RNMT+ encoder with Transformer decoder.

Encoder	Decoder	En→Fr Test BLEU
Trans. Big	Trans. Big	40.73 ± 0.19
RNMT+	RNMT+	41.00 ± 0.05
Trans. Big	RNMT+	<b>41.12 ± 0.16</b>
RNMT+	Trans. Big	39.92 ± 0.21

Table 5: Results for encoder-decoder hybrids.

From Table 5, it is clear that the Transformer encoder is better at encoding or feature extraction than the RNMT+ encoder, whereas RNMT+ is better at decoding or conditional language modeling, confirming our intuition that a stateful de-

coder is beneficial for conditional language generation.

## 6.2 Assessing Encoder Combinations

Next, we explore how the features extracted by an encoder can be further enhanced by incorporating additional information. Specifically, we investigate the combination of transformer layers with RNMT+ layers in the same encoder block to build even richer feature representations. We exclusively use RNMT+ decoders in the following architectures since stateful decoders show better performance according to Table 5.

We study two mixing schemes in the encoder (see Fig. 2):

(1) *Cascaded Encoder*: The cascaded encoder aims at combining the representational power of RNNs and self-attention. The idea is to enrich a set of stateful representations by cascading a feature extractor with a focus on vertical mapping, similar to (Pascanu et al., 2013; Devlin, 2017). Our best performing cascaded encoder involves fine tuning transformer layers stacked on top of a pre-trained frozen RNMT+ encoder. Using a pre-trained encoder avoids optimization difficulties while significantly enhancing encoder capacity. As shown in Table 6, the cascaded encoder improves over the Transformer encoder by more than 0.5 BLEU points on the WMT’14 En→Fr task. This suggests that the Transformer encoder is able to extract richer representations if the input is augmented with sequential context.

(2) *Multi-Column Encoder*: As illustrated in Fig. 2b, a multi-column encoder merges the outputs of several independent encoders into a single combined representation. Unlike a cascaded encoder, the multi-column encoder enables us to investigate whether an RNMT+ decoder can distinguish information received from two different channels and benefit from its combination. A crucial operation in a multi-column encoder is therefore how different sources of information are merged into a unified representation. Our best multi-column encoder performs a simple concatenation of individual column outputs.

The model details and hyperparameters of the above two encoders are described in Appendix A.5 and A.6. As shown in Table 6, the multi-column encoder followed by an RNMT+ decoder achieves better results than the Transformer and the RNMT model on both WMT’14 benchmark tasks.

Model	En→Fr BLEU	En→De BLEU
Trans. Big	40.73 ± 0.19	27.94 ± 0.18
RNMT+	41.00 ± 0.05	28.49 ± 0.05
Cascaded	<b>41.67 ± 0.11</b>	28.62 ± 0.06
MultiCol	41.66 ± 0.11	<b>28.84 ± 0.06</b>

Table 6: Results for hybrids with cascaded encoder and multi-column encoder.

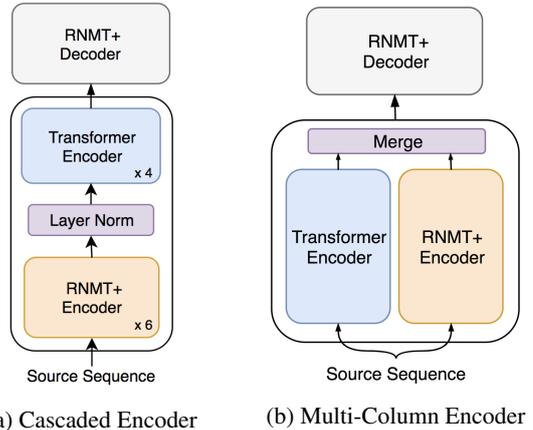


Figure 2: Vertical and horizontal mixing of Transformer and RNMT+ components in an encoder.

## 7 Conclusion

In this work we explored the efficacy of several architectural and training techniques proposed in recent studies on seq2seq models for NMT. We demonstrated that many of these techniques are broadly applicable to multiple model architectures. Applying these new techniques to RNMT models yields RNMT+, an enhanced RNMT model that significantly outperforms the three fundamental architectures on WMT’14 En→Fr and En→De tasks. We further presented several hybrid models developed by combining encoders and decoders from the Transformer and RNMT+ models, and empirically demonstrated the superiority of the Transformer encoder and the RNMT+ decoder in comparison with their counterparts. We then enhanced the encoder architecture by horizontally and vertically mixing components borrowed from these architectures, leading to hybrid architectures that obtain further improvements over RNMT+.

We hope that our work will motivate NMT researchers to further investigate generally applicable training and optimization techniques, and that our exploration of hybrid architectures will open paths for new architecture search efforts for NMT.

Our focus on a standard single-language-pair translation task leaves important open questions to be answered: How do our new architectures compare in multilingual settings, i.e., modeling an *interlingua*? Which architecture is more efficient and powerful in processing finer grained inputs and outputs, e.g., characters or bytes? How transferable are the representations learned by the different architectures to other tasks? And what are the characteristic errors that each architecture makes, e.g., linguistic plausibility?

## Acknowledgments

We would like to thank the entire Google Brain Team and Google Translate Team for their foundational contributions to this project. We would also like to thank the entire Tensor2Tensor development team for their useful inputs and discussions.

## References

- Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](https://arxiv.org/abs/1607.06450). *CoRR* abs/1607.06450. [http://arxiv.org/abs/1607.06450](https://arxiv.org/abs/1607.06450).
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](https://arxiv.org/abs/1409.0473). In *International Conference on Learning Representations*. [http://arxiv.org/abs/1409.0473](https://arxiv.org/abs/1409.0473).
- Y. Bengio, P. Simard, and P. Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *Trans. Neur. Netw.* 5(2):157–166.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. [Massive exploration of neural machine translation architectures](https://www.aclweb.org/anthology/D17-1151). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1442–1451. <https://www.aclweb.org/anthology/D17-1151>.
- Hugh Chen, Scott Lundberg, and Su-In Lee. 2017. [Checkpoint ensembles: Ensemble methods from a single training process](https://arxiv.org/abs/1710.03282). *CoRR* abs/1710.03282. [http://arxiv.org/abs/1710.03282](https://arxiv.org/abs/1710.03282).
- Jianmin Chen, Rajat Monga, Samy Bengio, and Rafal Józefowicz. 2016. [Revisiting distributed synchronous SGD](https://arxiv.org/abs/1604.00981). *CoRR* abs/1604.00981. [http://arxiv.org/abs/1604.00981](https://arxiv.org/abs/1604.00981).
- Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder-decoder for statistical machine translation](https://arxiv.org/abs/1406.1078). In *Conference on Empirical Methods in Natural Language Processing*. [http://arxiv.org/abs/1406.1078](https://arxiv.org/abs/1406.1078).
- Jan Chorowski and Navdeep Jaitly. 2016. [Towards better decoding and language model integration in sequence to sequence models](https://arxiv.org/abs/1612.02695). *CoRR* abs/1612.02695. [http://arxiv.org/abs/1612.02695](https://arxiv.org/abs/1612.02695).
- Josep Maria Crego, Jungi Kim, Guillaume Klein, Anabel Rebollo, Kathy Yang, Jean Senellart, Egor Akhanov, Patrice Brunelle, Aurelien Coquard, Yongchao Deng, Satoshi Enoue, Chiyo Geiss, Joshua Johanson, Ardas Khalsa, Raoum Khiari, Byeongil Ko, Catherine Kobus, Jean Lorieux, Leidiana Martins, Dang-Chuan Nguyen, Alexandra Priori, Thomas Riccardi, Natalia Segal, Christophe Serivan, Cyril Tiquet, Bo Wang, Jin Yang, Dakun Zhang, Jing Zhou, and Peter Zoldan. 2016. [Systran’s pure neural machine translation systems](https://arxiv.org/abs/1610.05540). *CoRR* abs/1610.05540. [http://arxiv.org/abs/1610.05540](https://arxiv.org/abs/1610.05540).
- Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2016. [Language modeling with gated convolutional networks](https://arxiv.org/abs/1612.08083). *CoRR* abs/1612.08083. [http://arxiv.org/abs/1612.08083](https://arxiv.org/abs/1612.08083).
- Jeffrey Dean, Greg S. Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V. Le, Mark Z. Mao, MarcAurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y. Ng. 2012. Large scale distributed deep networks. In *NIPS*.
- Michael Denkowski and Graham Neubig. 2017. [Stronger baselines for trustable results in neural machine translation](https://www.aclweb.org/anthology/W17-3203). In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, Vancouver, pages 18–27. [http://www.aclweb.org/anthology/W17-3203](https://www.aclweb.org/anthology/W17-3203).
- Jacob Devlin. 2017. [Sharp models on dull hardware: Fast and accurate neural machine translation decoding on the cpu](https://aclweb.org/anthology/D17-1300). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 2820–2825. [http://aclweb.org/anthology/D17-1300](https://aclweb.org/anthology/D17-1300).
- Jeffrey L. Elman. 1990. Finding structure in time. *Cognitive Science* 14(2):179 – 211.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](https://arxiv.org/abs/1705.03122). *CoRR* abs/1705.03122. [http://arxiv.org/abs/1705.03122](https://arxiv.org/abs/1705.03122).
- Felix A Gers, Jürgen Schmidhuber, and Fred Cummins. 2000. Learning to forget: Continual prediction with LSTM. *Neural computation* 12(10):2451–2471.
- Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He.

2017. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR* abs/1706.02677. <http://arxiv.org/abs/1706.02677>.
- Alex Graves and Jürgen Schmidhuber. 2005. Frame-wise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks* 18(5):602 – 610. IJCNN 2005.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. 2015. Learning to transduce with unbounded memory. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*. MIT Press, Cambridge, MA, USA, NIPS’15, pages 1828–1836. <http://dl.acm.org/citation.cfm?id=2969442.2969444>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *CoRR* abs/1512.03385. <http://arxiv.org/abs/1512.03385>.
- Sepp Hochreiter. 1991. Untersuchungen zu dynamischen neuronalen netzen. *Diploma, Technische Universität München* 91:1.
- Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2(5):359 – 366.
- Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*.
- Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, and et al. 2017. In-datacenter performance analysis of a tensor processing unit. *CoRR* abs/1704.04760. <http://arxiv.org/abs/1704.04760>.
- Marcin Junczys-Dowmunt, Tomasz Dwojak, and Rico Sennrich. 2016. The amu-uedin submission to the wmt16 news translation task: Attention-based nmt models as feature functions in phrase-based smt. *arXiv preprint arXiv:1605.04809*.
- Nal Kalchbrenner and Phil Blunsom. 2013. Recurrent continuous translation models. In *Conference on Empirical Methods in Natural Language Processing*.
- Nal Kalchbrenner, Lasse Espeholt, Karen Simonyan, Aaron van den Oord, Alex Graves, and Koray Kavukcuoglu. 2016. Neural machine translation in linear time. *CoRR* abs/1610.10099. <http://arxiv.org/abs/1610.10099>.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980. <http://arxiv.org/abs/1412.6980>.
- Yann LeCun and Yoshua Bengio. 1998. The handbook of brain theory and neural networks. MIT Press, Cambridge, MA, USA, chapter Convolutional Networks for Images, Speech, and Time Series, pages 255–258. <http://dl.acm.org/citation.cfm?id=303568.303704>.
- Ankur P. Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *EMNLP*.
- Razvan Pascanu, Çağlar Gülçehre, Kyunghyun Cho, and Yoshua Bengio. 2013. How to construct deep recurrent neural networks. *CoRR* abs/1312.6026. <http://arxiv.org/abs/1312.6026>.
- Tim Salimans and Diederik P. Kingma. 2016. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *CoRR* abs/1602.07868. <http://arxiv.org/abs/1602.07868>.
- M. Schuster and K.K. Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45(11):2673–2681.
- Mike Schuster and Kaisuke Nakajima. 2012. Japanese and Korean voice search. *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*.
- H.T. Siegelmann and E.D. Sontag. 1995. On the computational power of neural nets. *Journal of Computer and System Sciences* 50(1):132 – 150.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. *CoRR* abs/1505.00387.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*. pages 3104–3112.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. Rethinking the inception architecture for computer vision. *CoRR* abs/1512.00567. <http://arxiv.org/abs/1512.00567>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *CoRR* abs/1706.03762. <http://arxiv.org/abs/1706.03762>.

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, Lukasz 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. <http://arxiv.org/abs/1609.08144>.

Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. [Deep recurrent models with fast-forward connections for neural machine translation](#). *CoRR* abs/1606.04199. <http://arxiv.org/abs/1606.04199>.