

Training Deeper Neural Machine Translation Models with Transparent Attention

Ankur Bapna* Mia Xu Chen* Orhan Firat* Yuan Cao* Yonghui Wu
ankurbpn, miachen, orhanf, yuancao@google.com
Google AI

Abstract

While current state-of-the-art NMT models, such as RNN seq2seq and Transformers, possess a large number of parameters, they are still shallow in comparison to convolutional models used for both text and vision applications. In this work we attempt to train significantly (2-3x) deeper Transformer and Bi-RNN encoders for machine translation. We propose a simple modification to the attention mechanism that eases the optimization of deeper models, and results in consistent gains of 0.7-1.1 BLEU on the benchmark WMT'14 English-German and WMT'15 Czech-English tasks for both architectures.

1 Introduction

The past few years have seen significant advances in the quality of machine translation systems, owing to the advent of neural sequence to sequence models. While current state of the art models come in different flavours, including Transformers (Vaswani et al., 2017), convolutional seq2seq models (Gehring et al., 2017) and LSTMs (Chen et al., 2018), all of these models follow the seq2seq with attention (Bahdanau et al., 2015) paradigm.

While revolutionary new architectures have contributed significantly to these quality improvements, the importance of larger model capacities cannot be downplayed. The first major improvement in NMT quality since the switch to neural models, amongst other factors, was brought about by a huge scale up in model capacity (Zhou et al., 2016; Wu et al., 2016). While there are multiple approaches to increase capacity, deeper models have been shown to extract more expressive features (Mhaskar et al., 2016; Telgarsky, 2016; Eldan and Shamir, 2015), and have resulted in significant gains for vision tasks over the past few years (He et al., 2015; Srivastava et al., 2015).

Despite this being an obvious avenue for improvement, research in deeper models is often restricted by computational constraints. Additionally, deep models are often plagued by trainability concerns like vanishing or exploding gradients (Bengio et al., 1994). These issues have been studied in the context of capturing long range dependencies in recurrent architectures (Pascanu et al., 2012; Hochreiter et al., 2001), but resolving these deficiencies in Transformers or LSTM seq2seq models deeper than 8 layers is unfortunately under-explored (Wang et al., 2017; Barone et al., 2017; Devlin, 2017).

In this study we take the first step towards training extremely deep models for translation, by training deep encoders for Transformer and LSTM based models. As we increase the encoder depth the vanilla Transformer models completely fail to train. We also observe sub-optimal performance for LSTM models, which we believe is associated with trainability issues. To ease optimization we propose an enhancement to the attention mechanism, which allows us to train deeper models and results in consistent gains on the WMT'14 En→De and WMT'15 Cs→En tasks.

2 Transparent Attention

While the effect of attention on the forward pass is exalted with visualizations and linguistic interpretations, its influence on the gradient flow is often forgotten. Consider the original seq2seq model *without* attention (Sutskever et al., 2014). To propagate the error signal from the last layer of the decoder to the first layer of the encoder, it has to pass through multiple time-steps in the decoder, survive the encoder-decoder bottleneck, and pass through multiple time-steps in the encoder, before reaching the parameter to be updated. There is some loss of information at every step, especially in

* Equal contribution.




Figure 1: Grad-norm ratio (r_t) vs training step (t) comparison for a 6 layer (blue) and 20 layer (red) Transformer trained on WMT 14 En→De.

the early stages of training. Attention (Bahdanau et al., 2015) creates a direct path from the decoder to the topmost layer of the encoder, ensuring its efficient dispersal over time. This increase in inter-connectivity significantly shortens the credit-assignment path (Britz et al., 2017), making the network less susceptible to optimization pathologies like vanishing gradients.

For deeper networks the error signal also needs to traverse along the depth of the encoder. We propose an extension to the attention mechanism that behaves akin to creating weighted residual connections along the encoder depth, allowing the dispersal of error signal simultaneously over encoder depth and time. Using trainable weights, this ‘transparent’ attention allows the model the flexibility to adjust the gradient flow to different layers in the encoder depending on its training phase.

2.1 Experimental Setup

We train our models on the standard WMT’14 En→De dataset. Each sentence is tokenized with the Moses tokenizer before breaking into subword units similar to (Sennrich et al., 2016). We use a shared vocabulary of 32k units for each language pair. We report all our results on newstest 2014, and use a combination of newstest 2012 and newstest 2013 for validation. To verify our results, we also evaluate our models on WMT’15 Cs→En. Here we use newstest 2013 for validation and newstest 2015 as the test set. To evaluate the models we compute BLEU on the tokenized, true-case output. We report the mean post-convergence score over a window of 21 checkpoints, obtained using dev performance, following (Chen et al., 2018).




Figure 2: Grad-norm ratio (r_t) vs training step (t) comparison for a 6 layer (blue) and 20 layer (red) RNMT+ model trained on WMT 14 En→De.

2.2 Baseline Experiments

We base our study on two architectures: Transformer (Vaswani et al., 2017) and RNMT+ (Chen et al., 2018). We choose a smaller version of each model to fit deep encoders with up to 20 layers on a single GPU. All our models are trained on eight P100 GPUs with synchronous training, and optimized using Adam (Kingma and Ba, 2014). For both architectures we train four models, with 6, 12, 16 and 20 encoder layers. We use 6 and 8 decoder layers for all our transformers and RNMT+ experiments respectively. We also report performance for the standard Transformer Big and RNMT+ setups, as described in (Chen et al., 2018), for comparison against higher capacity models.

Transformer: We use the latest version of the Transformer base model, using the implementation from (Chen et al., 2018). We modify the learning rate schedule to use a learning rate of 3.0 and 40,000 warmup steps.

RNMT+: We implemented a smaller version of the En→De RNMT+ model based on the description in (Chen et al., 2018), with 512 LSTM nodes in both encoder and decoder.

2.3 Analysis

From Tables 1 and 2, we notice that the deeper Transformer encoders completely fail to train. To understand what goes wrong we keep track of the grad norm ratio $r_t = \left(\frac{\|\nabla_{h_1} L^{(t)}\|}{\|\nabla_{h_N} L^{(t)}\|} \right)$, $t = 1 \dots T$, where $L^{(t)}$ is the loss at time step t , N is the number of layers in the encoder, h_1 is the output of the first encoder layer, h_N is the output of the N -th encoder layer, and T is the total number of train-




Figure 3: Grad-norm ratio (r_t) vs training step for 20 layer Transformer with transparent attention.

ing steps. We use r_t as a diagnostic measure for two reasons: First, it indicates if training is suffering from exploding or vanishing gradients. Second, when a network is properly trained the lowest layers usually converge quickly, whereas the top-most layers take longer (Raghu et al., 2017). We therefore expect that, for a healthy training process, r_t is relatively large during the early stages of training when updates to lower layers are larger than upper layers. We observe this in most successful Transformer and RNMT+ training runs.

Figure 1 illustrates the r_t curves for the 6-layer and 20-layer Transformers. As expected, the shallow model has a high r_t value during early stages of training. For the deep model, however, r_t remains flat at a much smaller value throughout training. We also observe that r_t remains below 1.0 for both models, although the problem seems much less severe for the shallow model.

From Tables 3 and 4, we also observe that the performance of deep RNMT+ encoders is not significantly impacted, reaching the level of the 6 layer model. This is supported by the RNMT+ r_t curves in Figure 2, which indicate few differences in the learning dynamics of the shallow and deep models. This contrasts with the Transformer experiments, where increasing the depth leads to an unstable training process.

To gain further insights into the stability of the two architectures we completely remove the residual connections from their encoders. Residual connections have been shown, in theory and practice, to improve training stability and performance of deeper networks (see (He et al., 2015; Philipp et al., 2017; Hardt and Ma, 2017; Orhan, 2017)). Removing residual connections leads to disastrous results for the Transformer, where the training pro-




Figure 4: Plot illustrating the variations in the learned attention weights $s_{i,6}$ for the 20 layer Transformer encoder over the training process.

cess either does not converge or results in significantly worse results. On the other hand, the 6 layer RNMT+ converges with only a slight degradation in quality. Deeper versions of RNMT+ fail to train in the absence of residual connections.

2.4 Regulating Deep Encoder Gradients with Transparent Attention

Our baseline experiments reveal that mechanisms to regulate gradient flow can be critical to improving the optimization of deeper encoders. Since the only difference between our shallow and deep models is the number of layers in the encoder, the trainability issues are likely to be associated with gradient flow through the encoder.

To improve gradient flow we let the decoder attend weighted combinations of all encoder layer outputs, instead of just the top encoder layer. Similar approaches have been found to be useful in deep convolutional networks, for example (Shen and Zeng, 2016; Huang et al., 2016a; Srivastava et al., 2015; Huang et al., 2016b), but this remains un-investigated in sequence-to-sequence models. We formulate our proposal below.

Assume the model has N encoder layers and M encoder-decoder attention modules. For Transformer models each decoder layer attends the encoder, so M is equivalent to the number of decoder layers ($M = 6$). For RNMT+, attention is only applied in the first decoder layer, thus $M = 1$. Let the activations from the i -th encoder layer be $\{h_t^i | t = 1 \dots T\}$, and embeddings be layer 0. Then the traditional attention module attends to $\{h_t^N | t = 1 \dots T\}$. In transparent attention we evaluate M weighted combinations of the encoder outputs, one corresponding to each attention mod-

En→De WMT 14	Transformer (Base)				(Big)
Encoder layers	6	12	16	20	6
Num. Parameters	94M	120M	137M	154M	375M
Baseline	27.26	*	*	*	27.94
Baseline - residuals	*	6.00	*	*	N/A
Transparent	27.52	27.79	28.04	27.96	N/A

Table 1: BLEU scores on En→De newstest 2014 with Transformers. * indicates that a model failed to train.

Cs→En WMT 15	Transformer (Base)				(Big)
Encoder layers	6	12	16	20	6
Num. Parameters	94M	120M	137M	154M	375M
Baseline	27.20	*	*	*	27.76
Baseline - residuals	25.83	*	*	*	N/A
Transparent	27.41	27.69	27.93	27.80	N/A

Table 2: BLEU scores Cs→En newstest 2015 with Transformers. * indicates that a model failed to train.

ule. We define a $(N + 1) \times M$ weight vector W , which is learned during training.¹ We apply dropout to W since we empirically found it helpful to stabilize training. We then compute softmax s to normalize the weights.

$$s_{i,j} = \frac{e^{W_{i,j}}}{\sum_{k=0}^N e^{W_{k,j}}}, \quad j = 1 \dots M \quad (1)$$

We now define

$$z_t^j = \sum_{i=1}^{N+1} s_{i,j} h_t^i, \quad t = 1 \dots T, \quad j = 1 \dots M \quad (2)$$

Now attention module j attends to $\{z_t^j \mid t = 1 \dots T\}$. Since in RNMT+ a projection is applied to the encoder final layer output, we apply a projection to the weighted combination of encoder outputs before the attention module.

3 Results and Analysis

Our results, from tables 1 and 2, indicate that adding transparent attention improves the performance of most of our transformer experiments, but the gains are most pronounced for deeper models. While the baseline transformer fails to train with 12 layers or deeper encoders, transparent attention allows us to train encoders with up to 20 layers, improving by more than 0.7 BLEU points on both datasets. Relative to Transformer Big, deeper models seem to result in better or comparable performance with less than half the model capacity.

¹Here +1 is for the embedding layer.

We also observe gains of 0.7 and 1.0 BLEU for RNMT+ models, on En→De and Cs→En respectively, as indicated by Tables 3 and 4. However, experiments comparing wide models against deeper ones are inconclusive. While deeper models perform slightly better than a wide model with double their capacity on Cs-En, they are clearly out-performed by the larger model on En-De.

The r_t plot in Figure 3, also indicates that the learning dynamics now resemble what we expect to see with stable training. We also notice that the scale of r_t now resembles that of the RNMT+ model, although the lower layers converge more slowly for the Transformer, possibly because it uses a much smaller learning rate.

A plot of the weights $s_{i,j}$, in Figure 4, also seems to support our findings. The scalar weights for the lowest embeddings layer grow rapidly in the early stages of training, but once these layers converge the weights for layers 16 and 20 become much larger. The weights for the top few layers remain comparable at convergence, suggesting that the observed gains in performance might also be partially associated with an ensembling effect of the encoder features, similar to the effect observed in (Peters et al., 2018).

En→De WMT 14	RNMT+ (512)				(1024)
Encoder layers	6	12	16	20	6
Num. Parameters	128M	165M	191M	216M	379M
Baseline	26.63	26.32	26.49	26.33	28.49
Baseline - residuals	26.37	*	*	*	N/A
Transparent	26.61	26.87	27.07	27.33	N/A

Table 3: BLEU scores on En→De newstest 2014 with RNMT+. * indicates that a model failed to train.

Cs→En WMT 15	RNMT+ (512)				(1024)
Encoder layers	6	12	16	20	6
Num. Parameters	128M	165M	191M	216M	379M
Baseline	25.77	25.86	26.02	25.75	26.66
Baseline - residuals	25.43	*	*	*	N/A
Transparent	26.69	26.74	26.79	26.72	N/A

Table 4: BLEU scores Cs→En newstest 2015 with RNMT+. * indicates that a model failed to train.

4 Conclusions and Future Work

In this work we explore deeper encoders for Transformer and RNMT+ based machine translation models. We observe that Transformer models are extremely difficult to train when encoder depth is increased beyond 12 layers. While RNMT+ models train with deeper encoders, we did not observe any big performance improvements.

We associated the difficulty in training deeper encoders with hindered gradient flow, and resolved it by proposing the transparent attention mechanism. This enabled us to successfully train deeper Transformer and RNMT+ models, resulting in consistent gains in translation quality on both WMT’ 14 En→De and WMT’ 15 Cs→En.

Our results show that there is potential for improvement in translation quality by training deeper architectures, even though they pose optimization challenges. While this study explores training deeper encoders for narrow models, we plan to further study extremely deep and wide models to utilize the full strength of these architectures.

5 Acknowledgments

We would like to thank the Google Brain and Google Translate teams for their foundational contributions to this project.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly

learning to align and translate. In *International Conference on Learning Representations*.

Antonio Valerio Miceli Barone, Jindřich Helcl, Rico Sennrich, Barry Haddow, and Alexandra Birch. 2017. Deep architectures for neural machine translation. *arXiv preprint arXiv:1707.07631*.

Yoshua Bengio, Patrice Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166.

Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. Massive exploration of neural machine translation architectures. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1442–1451, Copenhagen, Denmark. Association for Computational Linguistics.

Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Niki Parmar, Mike Schuster, Zhifeng Chen, et al. 2018. The best of both worlds: Combining recent advances in neural machine translation. *arXiv preprint arXiv:1804.09849*.

Jacob Devlin. 2017. Sharp models on dull hardware: Fast and accurate neural machine translation decoding on the cpu. *arXiv preprint arXiv:1705.01991*.

Ronen Eldan and Ohad Shamir. 2015. The power of depth for feedforward neural networks. *CoRR*, abs/1512.03965.

Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. *CoRR*, abs/1705.03122.

- Moritz Hardt and Tengyu Ma. 2017. Identity matters in deep learning. *ICLR*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *CoRR*, abs/1512.03385.
- Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, Jürgen Schmidhuber, et al. 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies.
- Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. 2016a. Densely connected convolutional networks. *CoRR*, abs/1608.06993.
- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. 2016b. Deep networks with stochastic depth. *CoRR*, abs/1603.09382.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Hrushikesh Mhaskar, Qianli Liao, and Tomaso Poggio. 2016. Learning functions: when is deep better than shallow. *arXiv preprint arXiv:1603.00988*.
- A. Emin Orhan. 2017. Skip connections as effective symmetry-breaking. *CoRR*, abs/1701.09175.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2012. Understanding the exploding gradient problem. *CoRR*, abs/1211.5063.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *CoRR*, abs/1802.05365.
- George Philipp, Dawn Song, and Jaime G. Carbonell. 2017. Gradients explode - deep networks are shallow - resnet explained. *CoRR*, abs/1712.05577.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. Svcca: Singular vector canonical correlation analysis for deep understanding and improvement. *arXiv preprint arXiv:1706.05806*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Falong Shen and Gang Zeng. 2016. Weighted residuals for very deep networks. *CoRR*, abs/1605.08831.
- 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.
- Matus Telgarsky. 2016. Benefits of depth in neural networks. *CoRR*, abs/1602.04485.
- 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.
- Mingxuan Wang, Zhengdong Lu, Jie Zhou, and Qun Liu. 2017. Deep neural machine translation with linear associative unit. *arXiv preprint arXiv:1705.00861*.
- 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.
- Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. Deep recurrent models with fast-forward connections for neural machine translation. *arXiv preprint arXiv:1606.04199*.