

Query-Key Normalization for Transformers

Alex Henry, Prudhvi Raj Dachapally, Shubham Pawar, Yuxuan Chen

Cyndx Technologies

{alex.henry, prudhvi.dachapally, shubham.pawar, ethan.chen}@cyndx.com

Abstract

Low-resource language translation is a challenging but socially valuable NLP task. Building on recent work adapting the Transformer’s normalization to this setting, we propose QKNORM, a normalization technique that modifies the attention mechanism to make the softmax function less prone to arbitrary saturation without sacrificing expressivity. Specifically, we apply ℓ_2 normalization along the head dimension of each query and key matrix prior to multiplying them and then scale up by a learnable parameter instead of dividing by the square root of the embedding dimension. We show improvements averaging 0.928 BLEU over state-of-the-art bilingual benchmarks for 5 low-resource translation pairs from the TED Talks corpus and IWSLT’15.¹

1 Introduction

The Transformer (Vaswani et al., 2017) remains the architecture of choice for machine translation. Since its introduction, various architectural and functional modifications have been made to improve its performance on NMT datasets (Ahmed et al., 2017; Zhang et al., 2018; Wang et al., 2019; Dai et al., 2019; Zhao et al., 2019). Translating low-resource languages presents special challenges. Recent strategies for adapting Transformers to this socially valuable task include exploiting transfer learning with many-to-many multilingual models (Aharoni et al., 2019), reducing model depth (van Biljon et al., 2020), and adding a regularization penalty for diverging from the predictions of a monolingual language model pretrained on the target language (Baziotis et al., 2020). This paper builds on recent work on layer normalization for

low-resource language pairs, introducing a normalization technique that tries to keep the input to softmax attention within an appropriate range.

Layer normalization. For Transformers and other NLP models, layer normalization (Ba et al., 2016) yields significantly better performance than batch normalization (Ioffe and Szegedy, 2015), in part because NLP models tend to exhibit greater variance in batch statistics during training, for example compared to computer vision (Shen et al., 2020). Layer normalization boosts performance in deeper networks chiefly by controlling their gradients (Xu et al., 2019). It re-scales and re-centers activation distributions (though re-centering may be unnecessary, see Zhang and Sennrich 2019). The type of normalization used and the placement of that normalization within the Transformer are both crucial to Transformer performance (Nguyen and Salazar, 2019).

Softmax attention. Given a matrix X embedding a sequence of tokens, attention transforms each embedding into a mixture of itself and other elements of the sequence according to the importance of their connections for the modeling task at hand. In the case of multihead self-attention, the vectors of X are projected linearly into Query, Key and Value matrices. The operation

$$\text{softmax}(QK^T) \quad (1)$$

defines a distribution for each token over all the others in its sequence that sums to 1. Multiplying by V then yields a new matrix where the embedding of each token is a weighted average of the vectors in V .

Richter and Wattenhofer (2020) propose replacing the softmax function in attention because it constrains attention’s output to the convex hull spanned by the vectors in V , limiting model flexibility. For

¹Code to reproduce our experiments is available at <https://github.com/CyndxAI/QKNorm>

the softmax over the vocabulary in next word prediction, Demeter et al. (2020) find that the norms of word embeddings drown out their angular displacements, with the consequence that words with smaller norms are systematically less likely to be predicted.

In this work, we replace the dot product inside of softmax attention with cosine similarity scaled up by a learnable parameter. This technique yields improved performance in low-resource bilingual translation, which we conjecture is because it binds QK^T to a narrower range in a way that makes it easier to learn more diffuse attention patterns wherever these prove valuable.

2 Background

Nguyen and Salazar (2019) achieve state-of-the-art bilingual performance on 5 low-resource translation pairs from the TED Talks (Qi et al., 2018) and IWSLT’15 (Cettolo et al., 2015) corpora. This work builds directly on theirs, applying our technique to the same 5 benchmarks. Their model combines three normalization techniques that we describe below: FIXNORM (Nguyen and Chiang, 2018), PRENORM (Klein et al., 2017; Domhan, 2018; Vaswani et al., 2018; Chen et al., 2018), and SCALENORM, which they introduce as a replacement for layer normalization. They report that each technique contributes about 0.3 BLEU for an average improvement of 1.1 BLEU across the test sets for their 5 language pairs.

FIXNORM sets word embeddings to unit length, which aids rare word translation (Nguyen and Chiang, 2018). PRENORM simply changes the location of layer normalization within the Transformer architecture, applying it to the input to each sublayer instead of after the residual connection. Moving layer normalization ahead of the residual connection enhances stability because the residual path is allowed to stay an identity map, instead of contributing terms to the gradient that could cause it to explode or vanish (Wang et al., 2019; Nguyen and Salazar, 2019). Interestingly, Nguyen and Salazar (2019) find PRENORM to be superior in low-resource but not high-resource translation settings.

Lastly, SCALENORM replaces layer normalization with ℓ_2 normalization along the embedding dimension, multiplied by a learnable scalar parameter initialized with $\frac{1}{\sqrt{d}}$ (where d is the embedding dimension; the same term is used in scaled dot

product attention (Vaswani et al., 2017)).

In other words, SCALENORM applies ℓ_2 normalization along the embedding dimension of Q , K and V , and it does so *before* the input to multihead attention gets split into heads.

Building on their work, we combine FIXNORM, PRENORM, and vanilla layer normalization (LAYERNORM) with a new technique we call query-key normalization (QKNORM), surpassing their model’s performance on each of the same 5 translation pairs by an average of 0.928 test BLEU.

QKNORM applies ℓ_2 normalization to Q and K *only*, and it does so along the *head* dimension (which is the same dimension as the embedding dimension, but *after* multihead attention has split its input into separate heads). Q and K thus become \hat{Q} and \hat{K} , where the i th row vector of \hat{Q} (the i th embedding in the sequence) is given by:

$$\hat{q}_i = \frac{q_i}{\|q_i\|} \quad (2)$$

The effect is to make each element of QK^T the cosine similarity of the corresponding pair of contextual token representations instead of their dot product. This is similar to Luo et al. (2018), who propose replacing the dot product in fully-connected networks between layer weights and previous layer outputs with cosine similarity.

Like SCALENORM, we also multiply by a learnable parameter that we initialize according to a rule of thumb we describe below. Unlike SCALENORM, QKNORM complements LAYERNORM rather than replacing it.

3 Dot Products and the Softmax Function

Softmax attends only to the differences between values. For example,

$$\begin{aligned} & \text{softmax}([760, 752, 750]) \\ &= \text{softmax}([12, 4, 2]) \\ &= [0.99962, 0.00034, 0.00005]. \end{aligned}$$

Since the dot product is unbounded, differences between elements that may be insignificantly small on a relative basis can silence all other signals in the attention weights applied to V . We conjecture that this limits the complexity of the patterns that attention heads can learn.

The impact is more obvious in less sophisticated Transformer implementations (perhaps in part because subsequent advances have mitigated the same

Encoder Layer 4




Figure 1: Scaled Dot Product Attention. Self-attention heatmaps for 4 heads from one encoder layer displaying more “concentrated” attention, consistent with the conjecture that unnormalized dot products in QK^T saturate the softmax and limit the attention patterns that can be learned.

Encoder Layer 4




Figure 2: Query-Key Normalized Attention. Self-attention heatmaps of the same 4 heads in Figure 1. QKNORM enables more diffuse attention patterns.

issue in different ways). Figures 1 and 2 show a heatmap comparison of encoder weights trained using the code for The Annotated Transformer², the first with scaled dot product attention and the second with QKNORM.

The models containing these encoders were trained for 10 epochs on IWSLT 2016 *de*→*en* (Cettolo et al., 2016) using the Annotated Transformer implementation, with the baseline model scoring 19.4 BLEU and the QKNORM model scoring 24.33 BLEU on the test set, computed with the SacreBLEU Python package (Post, 2018).

Though this heatmap comparison is obviously not systematic, we think the visual at least provides a plausible intuition for the incremental gain this technique achieves, with scaled dot product attention exhibiting the kind of “winner-take-all” behavior we would expect from a softmax near saturation.

In comparison to dot products, cosine similarities are bounded by $[-1, 1]$ which creates the opposite problem as input to softmax – the differences

²<https://nlp.seas.harvard.edu/2018/04/03/attention.html>

between values are too small for softmax to let the model effectively ignore connections between words it should not attend to. Instead of dividing by \sqrt{d} as in scaled dot product attention we scale up using a learnable parameter that we initialize with a value that depends on the length of the sequences in the training data (and hence on the number of elements in QK^T):

$$g_0 = \log_2(L^2 - L) \quad (3)$$

where L is the 97.5th percentile sequence length across all training data sequences for source and target.

The attention operation thus changes from

$$\text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (4)$$

to

$$\text{softmax}(g * \hat{Q}\hat{K}^T)V \quad (5)$$

where \hat{Q} and \hat{K} are Q and K with ℓ_2 -normalization applied along their head dimensions and g is a learnable scalar parameter initialized with g_0 as computed in (3).

	Examples	Source + Target Tokens	Number of Parameters	Training Time (in hours)	Development BLEU	GPU	L
gl → en	10k	0.37M	31,051,880	6	23.45	T4	79
sk → en	61k	2.32M	48,356,907	11	31.34	T4	75
en → vi	133k	5.99M	48,431,538	19	28.77	T4	72
en → he	212k	7.88M	48,401,538	38	31.16	T4	72
ar → en	214k	8.09M	48,499,512	26	37.94	P100	75

Table 1: Summary of data and model training information. Number of examples and number of tokens taken directly from Nguyen and Salazar (2019). L is the 97.5th percentile sequence length across all training data sequences.

	en→vi	ar→en	en→he	gl→en	sk→en
Nguyen and Salazar (2019)	32.79	36.09	28.28	22.01	32.58
QKNORM + LAYERNORM	33.24	36.75	28.96	24.21	33.23

Table 2: Comparison of test BLEU (Papineni et al., 2002), scored using the Moses toolkit scripts provided in the repo for Nguyen and Salazar (2019). $p < 0.01$ using bootstrap resampling (Koehn, 2004). Both architectures use PRENORM and FIXNORM. The Nguyen and Salazar (2019) architecture uses SCALENORM where we instead use vanilla layer normalization (Ba et al., 2016), and scaled dot product attention where we use QKNORM.

4 Experiments and Results

We follow the implementation in the repository for Nguyen and Salazar (2019), both in replicating their performance and as a starting point for our version (and also for computing BLEU as reported in Table 2).³ We train on the same 5 low-resource translation pairs as Nguyen and Salazar (2019): 4 from the TED Talks corpus (Qi et al., 2018)⁴ – Arabic, Slovak, and Galician translated to English, and English translated to Hebrew – and 1 from the IWSLT’15 corpus (Cettolo et al., 2015), English to Vietnamese. The repository for Nguyen and Salazar (2019) provides the tokenized text they used for English to Vietnamese.

Tokenization and BLEU. Apart from BPE (Sennrich et al., 2016), their repository does not include the code they used for tokenization, so for the other 4 language pairs we used the tokenization script from the repository for Qi et al. (2018).⁵

The repository for Nguyen and Salazar (2019) includes two Moses⁶ scripts for scoring BLEU, `multi-bleu.perl` and `multi-bleu-detok.perl`. We can’t use `multi-bleu.perl` for the 4 TED Talks pairs without being able to replicate their tokenization because scores from that script are not comparable

³https://github.com/tnq177/Transformers_without_tears

⁴http://phontron.com/data/ted_talks.tar.gz

⁵https://github.com/neulab/word-embeddings-for-nmt/blob/master/ted_reader.py

⁶<https://github.com/moses-smt/mosesdecoder>

when there are differences in tokenization, unlike `multi-bleu-detok.perl` (Post, 2018). We use `multi-bleu.perl` to score `en→vi` (since we have their preprocessed text for this pair) and `multi-bleu-detok.perl` to score the 4 TED Talks pairs.

For additional confirmation, we also score all models using SacreBLEU (Post, 2018) after detokenizing with NLTK’s `TrebankWordDetokenizer` (Bird and Loper, 2004). These scores are reported in Table 3. All the detokenized BLEU scores from Table 2 are basically unchanged in Table 3, with the exception of `en→vi`. The best scores for the baseline model we could get on `en→vi` were 32.48 for Moses `multi-bleu.perl` and 32.41 for SacreBLEU, though in Table 2 we report the `multi-bleu.perl` score from Nguyen and Salazar (2019), 32.79. Our model’s score for the same pair comes in 0.06 BLEU lower as well.

Following the Nguyen and Salazar (2019) repository, we perform BPE using `fastBPE`⁷. We also use the same Moses code for bootstrap resampling (Koehn, 2004).

Model hyperparameters. Although PRENORM has been shown to make warmup less important for Transformers using scaled dot product attention (Nguyen and Salazar, 2019; Xiong et al., 2020), we obtained our best results using 8,000 steps of linear warmup. How much linear warmup matters for QKNORM and why it matters are both subjects for further investigation. We used the same validation-

⁷<https://github.com/glample/fastBPE>

	en→vi	ar→en	en→he	gl→en	sk→en
Nguyen and Salazar (2019)	32.41	36.09	28.28	22.01	32.58
QKNORM + LAYERNORM	33.18	36.75	28.96	24.21	33.22

Table 3: Comparison of test BLEU (Papineni et al., 2002), scored using SACREBLEU (Post, 2018).

based decay scheme as Nguyen and Salazar (2019) and allowed models to train until they had reached the minimum learning rate. For all other model hyperparameters and preprocessing settings we followed Nguyen and Salazar (2019) and the code in the lead author’s GitHub repository. As in their repository, we calculate test BLEU on the translation from the epoch with the highest validation BLEU.

Results. Incorporating QKNORM and using layer normalization instead of SCALENORM boosted performance by an average of 0.928 BLEU across the test sets for the 5 translation pairs. On IWSLT’15 *en→vi*, our SacreBLEU test score of 33.18 is only 0.09 BLEU lower than Provilkov et al. (2020), who use BPE-dropout to increase BLEU 1.49 over the same model with vanilla BPE.

5 Conclusion

In this paper, we introduced a normalization technique that modifies the attention mechanism in Transformers and demonstrated its utility for low-resource bilingual translation by building it into an existing Transformer implementation with state-of-the-art performance on 5 low-resource language pairs. QKNORM improves performance for each of the 5 pairs, with an average test BLEU increase of 0.928. We pointed to possible explanations for its effectiveness but identifying exactly where it helps and why requires further research. First, we plan to combine our approach with the fairseq Transformer implementation (Ott et al., 2019) and apply it to the FLORES dataset (Guzmán et al., 2019), investigating the effect of QKNORM on the optimal depth, number of attention heads, and warmup schedule for low-resource translation, in combination with recent advances like BPE-dropout (Provilkov et al., 2020). Next, we plan to look at high-resource settings to see whether the benefits of query-key normalization dissipate with access to more training data. Lastly, we intend to study how QKNORM impacts what attention heads actually learn, adapting methods from BERT attention studies such as Clark et al. (2019).

Acknowledgments

The authors would like to thank the reviewers for their valuable and insightful comments, and Toan Q. Nguyen for helpful clarifications and suggestions along the way.

References

- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Karim Ahmed, Nitish Shirish Keskar, and Richard Socher. 2017. [Weighted transformer network for machine translation](#).
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#).
- Christos Baziotis, Barry Haddow, and Alexandra Birch. 2020. [Language model prior for low-resource neural machine translation](#).
- Elan van Biljon, Arnu Pretorius, and Julia Kreutzer. 2020. [On optimal transformer depth for low-resource language translation](#).
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.
- M. Cettolo, Nihues Jan, Stüker Sebastian, L. Bentivogli, R. Cattoni, and M. Federico. 2016. The iwslt 2016 evaluation campaign.
- M. Cettolo, J. Nihues, S. Stüker, L. Bentivogli, R. Cattoni, and M. Federico. 2015. The iwslt 2015 evaluation campaign.
- Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Mike Schuster, Noam Shazeer, Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Zhifeng Chen, Yonghui Wu, and Macduff Hughes. 2018. [The best of both worlds: Combining recent advances in neural machine translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 76–86, Melbourne, Australia. Association for Computational Linguistics.

- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- David Demeter, Gregory Kimmel, and Doug Downey. 2020. [Stolen probability: A structural weakness of neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2191–2197, Online. Association for Computational Linguistics.
- Tobias Domhan. 2018. [How much attention do you need? a granular analysis of neural machine translation architectures](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1799–1808, Melbourne, Australia. Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- Sergey Ioffe and Christian Szegedy. 2015. [Batch normalization: Accelerating deep network training by reducing internal covariate shift](#). volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France. PMLR.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Chunjie Luo, Jianfeng Zhan, Xiaohe Xue, Lei Wang, and Rui Ren. 2018. [Cosine Normalization: Using Cosine Similarity Instead of Dot Product in Neural Networks: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part I](#), pages 382–391.
- Toan Nguyen and David Chiang. 2018. [Improving lexical choice in neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 334–343, New Orleans, Louisiana. Association for Computational Linguistics.
- Toan Q. Nguyen and Julian Salazar. 2019. [Transformers without tears: Improving the normalization of self-attention](#). In *Proc. Workshop on Spoken Language Translation*. To appear.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. [BPE-dropout: Simple and effective subword regularization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892, Online. Association for Computational Linguistics.
- Ye Qi, Devendra Sachan, Matthieu Felix, Sarguna Padmanabhan, and Graham Neubig. 2018. [When and why are pre-trained word embeddings useful for neural machine translation?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 529–535, New Orleans, Louisiana. Association for Computational Linguistics.
- Oliver Richter and Roger Wattenhofer. 2020. [Normalized attention without probability cage](#).
- 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 (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Sheng Shen, Zhewei Yao, Amir Gholami, Michael Mahoney, and Kurt Keutzer. 2020. Powernorm: Rethinking batch normalization in transformers. In *ICML*.

Ashish Vaswani, Samy Bengio, Eugene Brevdo, François Chollet, Aidan Gomez, Stephan Gouws, Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. [Tensor2Tensor for neural machine translation](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 193–199, Boston, MA. Association for Machine Translation in the Americas.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz 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 5998–6008. Curran Associates, Inc.

Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. [Learning deep transformer models for machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1810–1822, Florence, Italy. Association for Computational Linguistics.

Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. [On layer normalization in the transformer architecture](#).

Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. 2019. [Understanding and improving layer normalization](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 4381–4391. Curran Associates, Inc.

Biao Zhang and Rico Sennrich. 2019. [Root mean square layer normalization](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 12381–12392. Curran Associates, Inc.

Biao Zhang, Deyi Xiong, and Jinsong Su. 2018. [Accelerating neural transformer via an average attention network](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Melbourne, Australia. Association for Computational Linguistics.

Guangxiang Zhao, Junyang Lin, Zhiyuan Zhang, Xuancheng Ren, Qi Su, and Xu Sun. 2019. [Explicit sparse transformer: Concentrated attention through explicit selection](#).

Number of Heads	Test BLEU
2	32.40
4	33.16
8	33.24
16	32.42
32	32.30

Table 4: IWSLT’15 $en \rightarrow vi$ test BLEU for QKNORM varying the number of attention heads.

Percentile	Test BLEU
75th	32.58
90th	32.89
92.5th	32.64
95th	33.13
97.5th	33.24
99th	32.64
Maximum Word Count	33.10

Table 5: IWSLT’15 $en \rightarrow vi$ test BLEU for QKNORM varying the training set word count percentile used to initialize the learnable scaling factor g .

Appendix

A Varying the Number of Heads

In Table 4, we show the performance of QKNORM on the $en \rightarrow vi$ test set varying the number of heads. Even when the number of heads is 32 (with head dimension 16), the performance remains stable.

B Equation 3

Intuitively, longer sequences require more scaling to make it at least possible for the maximum values in QK^T to softmax to 1. We arrived at Equation 3 empirically by applying softmax to similarity matrices of word vectors scaled up with various heuristics. Like \sqrt{d} in scaled dot product attention (Vaswani et al., 2017), Equation 3 is a rule of thumb but it initializes a learnable parameter.

We determined the best value of L in Equation 3 by running the $en \rightarrow vi$ translation task with different percentile values. Table 5 shows the results from those experiments.

C Ablation Experiments

Table 6 shares test performance on $en \rightarrow vi$ when we ablate specific components of QKNORM. The biggest performance drop in these experiments comes from omitting g , the learnable scaling factor. This is unsurprising because if we don’t scale up

Experiment	Test BLEU
Without g	24.53
Without LAYERNORM	31.56
Without FIXNORM	32.63
Without FIXNORM or PRENORM	32.20
ℓ_2 -normalizing V along with Q and K	32.34

Table 6: Ablation Experiments.

$\hat{Q}\hat{K}^T$ its values are all within $[-1, 1]$ and softmax is a function of the differences between values.