

Subformer: Exploring Weight Sharing for Parameter Efficiency in Generative Transformers

Machel Reid[†], Edison Marrese-Taylor^{‡,†}, Yutaka Matsuo[†]

[†]The University of Tokyo

[‡]National Institute of Advanced Industrial Science and Technology (AIST)

{machelreid, emarrese, matsuo}@weblab.t.u-tokyo.ac.jp

Abstract

Transformers have shown improved performance when compared to previous architectures for sequence processing such as RNNs. Despite their sizeable performance gains, as recently suggested, the model is computationally expensive to train and with a high parameter budget. In light of this, we explore parameter-sharing methods in Transformers with a specific focus on generative models. We perform an analysis of different parameter sharing/reduction methods and develop the Subformer. Our model combines sandwich-style parameter sharing, which overcomes naive cross-layer parameter sharing in generative models, and self-attentive embedding factorization (SAFE). Experiments on machine translation, abstractive summarization and language modeling show that the Subformer can outperform the Transformer even when using significantly fewer parameters.¹

1 Introduction

Recent improvements in NLP tasks can be attributed to the Transformer (Vaswani et al., 2017) model. The model has led to better deeply contextualized representations (Devlin et al., 2019), better machine translation systems (Vaswani et al., 2017), and language models (Baevski and Auli, 2019; Dai et al., 2019). Despite their success, one main drawback of training these models is their computational cost, being a greatly limiting factor for many, with training times and memory usage ballooning as model sizes increase to attain better performance.

With this in mind, there has been recent interest in making the Transformer more parameter-efficient to reap its performance benefits while making the model more computationally efficient and able to scale better. Many approaches have focused on

automating model design with neural architecture search that aim at finding more efficient Transformer variations using gradient descent (Wu et al., 2020; So et al., 2019; Mehta et al., 2020a). As such, these techniques are expensive, requiring a significant amount of GPU hours to find good designs. In contrast to these approaches, the model by (Lan et al., 2020) has directly tackled model parameter reduction in the context of deeply contextualized word representations, still attaining similar (or better) performance. In this paper, we take a similar approach and look to explore whether these ideas can be applied to sequence-to-sequence models in a simple manner by designing the Subformer.

The Subformer incorporates two novel techniques: (1) SAFE (Self-Attentive Factorized Embeddings), in which we use a small self-attention layer to reduce embedding parameter count, and (2) Sandwich-style Parameter Sharing, in which we develop a simple and intuitive technique for parameter sharing to be effective in Transformer models.

We evaluate the Subformer on machine translation, abstractive summarization, and language modeling, showing that our model can achieve similar or better performance compared with a base/big Transformer with a $\sim 40\%$ parameter reduction and minimal modification to the original architecture, reinforcing the existing over-parameterization claims (Fan et al., 2020; Mehta et al., 2020a; Lan et al., 2020). On WMT’14 EN-DE we achieve a BLEU score of 29.3, compared to Transformer-big’s 28.6 with 13M fewer parameters, and we also outperform the Transformer-XL model with a significant 3.6 lower perplexity and 37% fewer parameters.

2 The Subformer

Let us start by defining the notation to be used throughout the paper. We refer to the model dimension as d_m , feed-forward projection dimension as \vec{d}_m , the vocabulary size as V , and the number of

¹<https://github.com/machelreid/subformer>

layers as L . Note that, unlike standard Transformer models, in which the embedding dimension is kept the same as d_m , we disentangle to embedding dimension to reduce parameter count. For this reason, we denote the embedding dimension to be d_e .

2.1 SAFE: Self-Attentive Factorized Embeddings

We propose to reduce the number of parameters in our embedding layers, which can take up to 25% of the total parameter count (in the case of Transformer-base, Vaswani et al., 2017), using a small self-attention layer. Specifically, we look to reduce the embedding size by disentangling the model dimension from the embedding dimension, reducing the embedding dimension d_e , and then projecting this to the model dimension d_m using a small self-attention sub-layer followed by a feed-forward module.

Given vocabulary size V , the usage of a standard embedding layer would result in $V \times d_m$ parameters. However, considering that the power of Transformers lies in their ability to learn contextual representations with attention, using a smaller d_e for non-contextual embeddings and then projecting to d_m is intuitively an effective method for parameter reduction (Lan et al., 2020). This results in a significant parameter reduction for values of $d_e \ll d_m$.

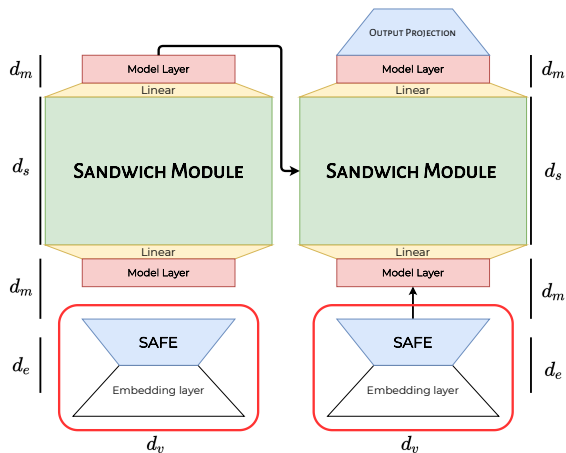


Figure 1: **The Subformer** with its four main components: (1) *SAFE embeddings* (colored blue) and output projection layers, (2) *model layers* which are placed at the top and bottom of the model (colored red), (3) *Sandwich Module*, in which a wider shared layer composes the central part of our encoder/decoder, (4) *projection layers*, which allow for the interaction between the model layers and Sandwich Module despite their different dimensions (colored yellow).

MODEL	Param.	BLEU
DeFINE (Mehta et al., 2020b)	52M	27.0
$d_e = 128$, Linear	48M	26.0
$d_e = 256$, Linear	53M	27.1
$d_e = 256$, 2-Layer Linear	54M	27.2
$d_e = 128$, SAFE	48M	26.6
$d_e = 256$, SAFE	54M	27.7
Vaswani et al. (2017)	65M	27.3
TRANSFORMER-BASE (reimpl.)	61M	27.7

Table 1: Experiments on the impact on SAFE vs a regular linear projection using TRANSFORMER-BASE on the WMT’14 EN-DE machine translation benchmark

When using SAFE, the embedding layer is composed of a regular token \rightarrow vector embedding matrix $\mathbf{E} \in \mathbb{R}^{V \times d_e}$. This is followed by projecting the embeddings (summed with the positional encodings (Vaswani et al., 2017), denoted by PE) to the model dimension d_m using SAFE. Once we have our SAFE embeddings, we feed them through the first model layer—the base of the sandwich. The output of this first layer is then projected to the sandwich dimension d_s . Once fed through the shared sandwich layers, we then project the output back to the model dimension using an MLP. The output of the projection is then fed through the final model layer to produce the output vectors.

Current models (Baevski and Auli, 2019; Dai et al., 2019; Lan et al., 2020) often use a single linear projection, i.e. $V \times d_e + d_e \times d_m$. In contrast, we empirically show that simply contextualizing this projection with a small self-attention layer results in stronger performance with a minimal addition of parameters—especially in the encoder-decoder case, where the input embedding layer and output projection are often tied (Press and Wolf, 2017) (Table 1).

2.2 Sandwich-style Parameter Sharing

Weight sharing techniques, despite being surprisingly effective, have been relatively unexplored in the context of generative Transformers. However, this has been shown to be powerful for leveraging models with large capacity and less memory usage/computation (Wu et al., 2019; Lan et al., 2020).

Given that the output of each Transformer layer depends directly on its two sub-layers—multiheaded attention and the feedforward module—when discussing alternatives for parameter sharing across

MODEL	Param.	BLEU
<i>All-Shared</i>	24M	14.3
<i>All-Shared</i> , $d_m = 768$	41M	22.0
<i>All-Shared (Independent FFN)</i>	27M	22.4
<i>All-Shared (except last)</i>	31M	23.2
<i>Every 2 layers shared</i>	38M	27.2
SANDWICH	38M	27.3
SANDWICH, $L = 8$	38M	27.7
Vaswani et al. (2017)	65M	27.3
TRANSFORMER-BASE (Our reimpl.)	61M	27.7

Table 2: Experiments performed on WMT’14 EN-DE using different parameter sharing techniques.

transformer layers there are several options. As we aim to leverage the aforementioned properties of weight sharing, we performed preliminary experiments, investigating the capabilities of weight sharing in the following five settings: (1) **All-shared** Naively sharing all encoder and all decoder layers—that is including both of their sub-layers, following Lan et al. (2020); Dehghani et al. (2019). (2) **All-Shared (Independent FFN)** Naively sharing all encoder and all decoder layers but allowing each layer $l \in [2, \dots, L]$ to have an independent feed-forward sub-layer. (3) **All-Shared (except last)** Sharing weights across layers $l \in [1, \dots, L - 1]$ such that layer L remains independent. (4) **Every 2 layers shared** Sharing every two layers, i.e. $[1, 2], [3, 4], [5, 6]$ in the case of a 6-layer transformer. (5) **Sandwich** Finally, we only share the middle or central layers (i.e. $2 \leq l \leq L - 1$), leaving layers 1 and L to have independent sets of parameters.

Table 2 summarizes the results of our exploratory study. As can be seen, naive parameter sharing/tying approaches do not offer any advantages, hurting performance significantly ($\sim 50\%$) when compared to the regular Transformer. However, our results also show that when combined properly, using Sandwich-style parameter sharing, we can attain a good balance of parameter reduction and performance. When compared to tasks such as pre-training deep contextualized word representations, generative tasks such as machine translation require informative token-level representations for each input token to be accurately translated. In this context, we surmise that the success of Sandwich-style parameter sharing on this sequence-to-sequence task is a result of it being able to have the input and output layers (arguably, the most important layers) be trained independently, allowing them to learn

different operations than the sandwich layers.

Combined Architecture Having introduced our proposed parameter reduction techniques, we will now explain the Subformer architecture. The Subformer is composed of four main components, for both the encoder and decoder: the embedding layer, the model layers, the sandwich module and the projection layers. We disentangle the sandwiched layer dimension from that of the model layer, allowing the sandwich layer width to be larger than the rest of the model. For this reason, we denote the dimension of the sandwiched layer to be d_s and its corresponding feed-forward dimension to be \vec{d}_s .

3 Experimental Setup

We apply our method to a variety of sequence modeling tasks: neural machine translation, summarization, and language modeling. We compare Transformer-base and big (Vaswani et al., 2017) with the Subformer trained in the same setting. We also include simple sandwich-style parameter sharing (denoted as Sandwich- $\{\text{base}, \text{big}\}$) and the usage of *only* SAFE as an ablation of what these techniques do in their naive forms when decoupled. Additional implementation and training details with hyperparameter settings are in the appendix.

Machine Translation We evaluate our model on two standard translation benchmarks: WMT’14 English-German (EN-DE) and WMT’16 English-Romanian (EN-RO). Following previous work, we evaluate all models using tokenized BLEU (Papineni et al., 2002). In order to better contextualize our results, we consider parameter-efficient models DeLight (Mehta et al., 2020a) (contemporaneous work), and the Evolved Transformer (So et al., 2019).

Abstractive Summarization We test the model’s ability to process long documents on the CNN-DailyMail summarization benchmark (Hermann et al., 2015; Nallapati et al., 2016), comprising over 280K news articles paired with multi-sentence summaries. We also compare effects of BART (Lewis et al., 2020) pretraining (details in Appendix). For this task we contextualize our results with specialized architectures such as Pointer-Generator Networks (See et al., 2017), and methods leveraging pretraining: RobertaShare (Rothe et al., 2020), BertExtAbs (Liu and Lapata, 2019), and BART (Lewis et al., 2020). We evaluate using ROUGE 1,2,L (Lin, 2004).

BASE MODELS	WMT'14 EN-DE		WMT'16 EN-RO	
	Param.	BLEU	Param.	BLEU
DeLighT	37M	27.6	22M	34.3
Evolved Transformer	48M	27.7	—	—
DeLighT	54M	28.0	52M	34.7
Evolved Transformer	64M	28.2	—	—
Transformer-base (orig)	65M	27.3	62M	34.2 [†]
Transformer-base (ours)	61M	27.7	62M	34.1
Sandwich-base	38M	27.3	—	—
Only SAFE, $d_e = 256$	54M	27.7	—	—
SUBFORMER-SMALL	38M	27.7	20M	34.1
SUBFORMER-BASE	52M	28.1	48M	34.7
SUBFORMER-MID	63M	28.5	—	—

Table 3: Results on WMT'14 EN-DE and WMT'16 EN-RO task, for our base models. The [†] superscript indicates results from Kasai et al. (2020).

BIG MODELS	Param.	BLEU
Evolved Transformer	222M	29.0
Transformer-big (orig)	213M	28.4
Transformer-big (ours)	210M	28.6
Sandwich-big	122M	28.6
SUBFORMER-LARGE	197M	29.3

Table 4: Results on WMT'14 EN-DE for large models.

Language Modeling We evaluate on the large-scale Wikitext-103 dataset (Merity et al., 2017). Models are evaluated in terms of perplexity on the test portion. In order to better contextualize our results, we consider the QRNN (Merity et al., 2018), Transformer-XL (Dai et al., 2019) and Deep Equilibrium Model (DEQ) (Bai et al., 2019), which also employs parameter sharing.

4 Results

Machine Translation Tables 3 and 4² summarize our results on machine translation. Firstly, we note that our Transformer baselines outperform Vaswani et al. (2017) (base: 27.3 \rightarrow 27.7, big: 28.4 \rightarrow 28.6). We surmise that this is due to training for longer and with a larger batch size.

The Subformer outperforms our Transformer baselines when trained in the same setting, with similar or fewer parameters. SUBFORMER-SMALL reduces parameters by 40%, matching performance with our Transformer baselines. SUBFORMER-BASE and MID, outperform our model significantly (0.4, 0.8 BLEU) when using a fewer/similar

²In all tables, results from other work used to contextualize our results are placed above the double bar.

number of parameters. Furthermore, we note that SUBFORMER-MID performs similarly to the Transformer-big model (210M params.) in Table 4, despite a 70% parameter reduction.

For our large set of models (Table 4), Sandwich-big achieves the same performance as our Transformer-big reimplementation, but with 40% fewer parameters. We believe this to be an indication of the capability of Sandwich-style parameter sharing as the encoder/decoder layers get wider, while also providing further evidence for the overparameterization of the Transformer. Subformer-large, with 197M parameters achieves a significant 0.7 BLEU score gain over Transformer-big, despite using 13M fewer parameters.

Language Modeling Results for language modeling can be seen in Table 5. The Subformer uses adaptive input embeddings (Baevski and Auli, 2019) instead of SAFE embeddings, following common practice. We also train two Transformer baselines with the same setup—one with the same amount of parameters and another with a similar parameter count to Transformer-XL—to provide better context for comparison. Task-specific techniques that can be applied during training, such as caching (Dai et al., 2019) or other methods applied during inference time (Khandelwal et al., 2020; Krause et al., 2018) can further improve all models so we do not focus on these.

MODEL	Param.	CL	PPL
QRNN	151M	—	33.00
DeLighT	99M	480	24.14
Transformer-XL	151M	640	24.03
DEQ	110M	—	23.20
Adaptive Inputs	247M	480	19.03
Adaptive Inputs (4 Layer)	96M	480	26.42
Adaptive Inputs (8 Layer)	146M	480	22.32
SUBFORMER	83M	480	20.88
	96M	480	20.39
	122M	480	19.90

Table 5: Results on the Wikitext-103 benchmark. CL stands for Context Length.

As seen in Table 5, the Subformer outperforms the baselines by a significant margin (between 1.4 and 6.5 perplexity), with a significant reduction in parameters.

Abstractive Summarization For the CNN/Daily Mail summarization task we use Subformer-base. We also pretrain a Transformer

MODEL	Params.	R1	R2	RL
Ptr-Gen+Cov	—	39.5	17.3	36.4
RobertaShare	152M	40.3	18.9	37.6
BertExtAbs	220M	41.7	19.4	38.8
BART	406M	44.2	21.3	40.9
Transformer (3 Layer)	57M	40.0	17.5	36.7
Transformer	77M	40.1	17.6	36.8
Transformer-BART	77M	41.2	18.7	37.6
SUBFORMER-BASE	57M	40.9	18.3	37.7
SUBFORMER-BART	57M	41.6	19.2	38.4

Table 6: Results on CNN-Daily Mail.

Model	Param.	Iterations	Dev. PPL
Adaptive Inputs	146M	272K	22.31
SUBFORMER	83M	97K	20.84

Table 7: Iterations to convergence on WIKITEXT-103

and Subformer model with the same architecture on Wikipedia (details in the Appendix). As can be seen in Table 6, the Subformer outperforms two Transformer baselines with both the same parameter count and its respective Transformer-base configuration in both settings, demonstrating the Subformer’s performance on a variety of tasks and with longer sequences.

Discussion on Speed and Convergence We found training time to consistently speed up by 10-30%, and inference speed to either be faster by 10-20% (keeping $d_s = d_m$) to be slightly slower by 10-30% (when $d_s \gg d_m$) (due to more operations, similar to Lan et al. (2020)). The Subformer converges faster most likely due to fewer parameters to optimize. Given the fewer number of parameters, it can be expected for the models to converge with fewer iterations. We test this on the task of language modeling, where we found that the Subformer converged 65% faster than its Transformer counterpart, as shown in Table 7. We also measure inference speed on our machine translation models (Table 8).

5 Conclusion

In this paper we have presented the Subformer, a parameter-efficient Transformer-based generative model primarily based on two simple parameter factorization/sharing techniques. Our empirical results on three sequence modeling tasks show that the Subformer can achieve similar or better performance compared with a base/big Transformer with

Model	Param.	Speed \uparrow	BLEU
DeLighT	37M	0.30x	27.6
Transformer	61M	1.00x	27.7
SAFE, $d_e = 256$	54M	1.17x	27.7
SANDWICH-BASE	38M	1.26x	27.3
SUBFORMER-BASE	52M	0.75x	28.1

Table 8: Inference speed for our models measured on a single V100 GPU on WMT’14 En-De (batch size: 384, 1.00x = 5135 tokens)

a $\sim 40\%$ parameter reduction. Furthermore, the simplicity of our approach allows the Subformer to be used in conjunction with other parameter reduction techniques in the literature, for even smaller but performant models. We hope our work incites interest in parameter sharing techniques for an even wider range of Transformer models, ultimately helping reduce their computational cost in general.

Ethical Considerations

This work has impact in the field of natural language processing, and develops a more efficient approach for learning performant generative models. As with much of language technology has the potential to be both used for good and used maliciously. We also experiment with pretraining, learning representations in an unsupervised way, which is likely to capture and amplify biases found in the data. However, our approach has a potential positive impact given the lower cost and energy expenditure needed to train our proposed model.

References

- Alexei Baeovski and Michael Auli. 2019. [Adaptive input representations for neural language modeling](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. 2019. [Deep equilibrium models](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 688–699.
- Robin Cheong and Robel Daniel. 2019. Transformers. zip: Compressing transformers with pruning and quantization. Technical report, Stanford University, Stanford, California.
- 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.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Lukasz Kaiser. 2019. **Universal transformers**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. **Exploiting deep representations for neural machine translation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4253–4262, Brussels, Belgium. Association for Computational Linguistics.
- Sergey Edunov, Alexei Baevski, and Michael Auli. 2019. **Pre-trained language model representations for language generation**. 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 4052–4059, Minneapolis, Minnesota. Association for Computational Linguistics.
- Angela Fan, Edouard Grave, and Armand Joulin. 2020. **Reducing transformer depth on demand with structured dropout**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Prakhar Ganesh, Yao Chen, Xin Lou, Mohammad Ali Khan, Yin Yang, Deming Chen, Marianne Winslett, Hassan Sajjad, and Preslav Nakov. 2020. **Compressing Large-Scale Transformer-Based Models: A Case Study on BERT**. *arXiv:2002.11985 [cs, stat]*.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. **Mask-predict: Parallel decoding of conditional masked language models**. 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 6112–6121, Hong Kong, China. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. **Teaching machines to read and comprehend**. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- Jungo Kasai, James Cross, Marjan Ghazvininejad, and Jiatao Gu. 2020. **Non-autoregressive machine translation with disentangled context transformer**. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 5144–5155. PMLR.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. **Generalization through memorization: Nearest neighbor language models**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. **Reformer: The efficient transformer**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. 2018. **Dynamic evaluation of neural sequence models**. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2766–2775, Stockholmsmassan, Stockholm Sweden. PMLR.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. **ALBERT: A lite BERT for self-supervised learning of language representations**. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. **Deterministic non-autoregressive neural sequence modeling by iterative refinement**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Brussels, Belgium. Association for Computational Linguistics.
- Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. 2019. **Snip: single-shot network pruning based on connection sensitivity**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.

- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Sachin Mehta, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2020a. [Delight: Very deep and light-weight transformer](#).
- Sachin Mehta, Rik Koncel-Kedziorski, Mohammad Rastegari, and Hannaneh Hajishirzi. 2020b. [Define: Deep factorized input token embeddings for neural sequence modeling](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [An analysis of neural language modeling at multiple scales](#).
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Damos, Erich Elsen, David García, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. [Mixed precision training](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- 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 the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). 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 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Gabriele Prato, Ella Charlaix, and M. Rezagholizadeh. 2019. Fully Quantized Transformer for Improved Translation. *ArXiv*.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Sascha Rothe, Shashi Narayan, and Aliaksei Severyn. 2020. [Leveraging pre-trained checkpoints for sequence generation tasks](#). *Transactions of the Association for Computational Linguistics*, 8:264–280.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- 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.
- David R. So, Quoc V. Le, and Chen Liang. 2019. [The evolved transformer](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5877–5886. PMLR.
- Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. 2013. [On the importance of initialization and momentum in deep learning](#). In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pages 1139–1147. JMLR.org.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob

Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Felix Wu, Angela Fan, Alexei Baevski, Yann N. Dauphin, and Michael Auli. 2019. [Pay less attention with lightweight and dynamic convolutions](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

Zhanghao Wu, Zhijian Liu, Ji Lin, Yujun Lin, and Song Han. 2020. [Lite transformer with long-short range attention](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. [Q8BERT: Quantized 8Bit BERT](#). *arXiv:1910.06188 [cs]*.

A Designing the Subformer

Naming The Subformer is a play on words, referencing its small size - i.e. *sub-*, as well as the *Sandwich*-style parameter sharing technique, referencing the type of sandwich.

Architecture When using SAFE, our parameter count would result in $V \times d_e + 5d_e^2 + d_e \times d_m$ parameters, where the first term represents the embedding layer, the value $5d_e^2$ groups the query, key, and value projections and 2 output feed-forward layers, and $d_e \times d_m$ represents the linear projection from the embedding dimension to the model dimension. As mentioned in the paper, this results in a significant parameter reduction for values of $d_e \ll d_m$.

As we tie the decoder’s output projection layer (returning a distribution over the vocabulary) with the decoder’s input embedding matrix, we project the decoder’s last hidden state (with dimension d_m) to d_e using a two layer MLP. Also, when we perform encoder attention in the decoder’s Sandwich Module, we simply linearly project the query from the decoder from d_s to d_m and then project it back to d_s once the attention operation is complete.

Memory Footprint Table 9 summarizes the memory footprint of our proposed techniques. In this table, the benefits of Sandwich-style parameter sharing can be seen as the number of independent layers is controlled to be $L \leq 3$, however, Transformers generally need to be deeper (with a standard of $L = 6$) to learn more meaningful representations with the parameter count scaling linearly with respect to the layer count. Similarly, the benefits of disentangling the model dimension from the embedding dimension can be seen as well. Due to the parameter reduction attained by these techniques, the models can be trained in memory-constrained scenarios with a larger batch size.

B Data and Training Details

Training was done on 8 GPUs on a single DGX-1 Machine, while training on 16 GPUs was done using multiple compute nodes of a compute cluster. We train all base/small models on 8 NVIDIA Tesla V100 GPUs. For all big/large models, we train on 16 NVIDIA Tesla V100 GPUs. All models were trained with mixed precision (Micikevicius et al., 2018) and are implemented in PyTorch (Paszke et al., 2019) using our modification of fairseq

(Ott et al., 2019).

Machine Translation We train using 8192 tokens per GPU an update frequency of 2, for small, base models. For large models, we train on 16 GPUs with 4096 tokens per GPU with an update frequency of 2. We follow the training setup of Ghazvininejad et al. (2019): we use the same weight initialization scheme as BERT (Devlin et al., 2019), sampling weights from $\mathcal{N}(0, 0.02)$, initializing biases to zero and setting layer normalization parameters β and γ to be 0 and 1, respectively. For regularization we use the best of [0.1, 0.2, 0.3] dropout, weight decay of 0.01, while using label-smoothed cross entropy loss with $\epsilon = 0.1$. We train using an effective batch size of 128K tokens. The models are trained using Adam (Kingma and Ba, 2015), with hyper-parameters $\beta = (0.9, 0.999)$ and $\epsilon = 10^{-6}$. We warm up the learning rate to a peak of 5×10^{-4} within 10K iterations and then decay the learning rate with the inverse square root schedule. When creating the final model, we use the checkpoint with the lowest loss on the development set and generate using a beam size of 5 (Vaswani et al., 2017), tuning the length penalty of $\alpha \in [0.0, 0.2, \dots, 2.0]$ in the validation set. We perform early stopping, training for a maximum of 250K iterations.

We use the following settings for our models: (1) Subformer-small has $d_m = 512$, $d_s = 768$, $d_e = 256$ and $L = 8$, (2) Subformer-base has $d_m = 512$, $d_s = 1024$, $\vec{d}_s = 3072$, $d_e = 320$, (3) Subformer-mid has $d_m = 768$, $d_s = 768$, $d_e = 350$ and (4) Subformer-large has $d_m = 1024$, $d_s = 2048$ and $d_e = 512$. For WMT’16 EN-RO, our small model has $d_m = 320$, $d_s = 512$ and $d_e = 192$ and our base model has $d_m = 512$, $d_s = 640$, and $d_e = 384$.

In terms of datasets, we make use of the same pre-processed data used by Ghazvininejad et al. (2019) for WMT’14 EN-DE with a 32K BPE (Sennrich et al., 2016) vocabulary and during evaluation we perform de-hyphenation (Vaswani et al., 2017). We use the same data as Lee et al. (2018) for WMT’16 EN-RO with a 35K BPE vocabulary.

Abstractive Summarization We follow Edunov et al. (2019) and use the official ROUGE-1.5.5.pl script with parameters `-m -a -n 2`. As mentioned in the paper, our model configuration is the same as Subformer-base,

Method	Embedding Memory Usage	Model Memory Usage
Transformer	$d_m \times V$	$L(4d_m^2 + 2(\vec{d}_m \times d_m))$
Sandwich (naive)	—	$3(4d_m^2 + 2(4\vec{d}_m \times d_m))$
Subformer	$d_e \times V + 5d_e^2 + d_e \times d_m$	$2(4d_m^2 + 2(\vec{d}_m \times d_m)) + (4d_s^2 + 2(\vec{d}_s \times d_s + 2(d_s \times d_m)))$

Table 9: Memory space required by each method given a stack of encoder layers. *Sandwich (naive)* refers to simply performing Sandwich style parameter sharing with no other modifications to the architecture.

but we set $d_e = 256$. Articles are truncated to 400 tokens (See et al., 2017) and we use a BPE vocabulary of 32K types (Edunov et al., 2019). We follow the training schedule of Edunov et al. (2019). During inference, we tune generation length in the range of {40, 50, 60} and use tri-gram blocking, following standard practice. When pretraining, we pretrain on Wikipedia (14GB) for 100K iterations, using a batch size of 512K tokens. We use a learning rate of $7e-4$, warmed up over 10K iterations.

Language Modeling When training our language models, we use 8 GPUs with 3072 tokens per GPU and an update frequency of 3, following Baevski and Auli (2019). Models are trained using Nesterov’s accelerated gradient optimizer (Sutskever et al., 2013), warming up the learning rate to 1.0 for 16K iterations, and then annealing for 270K iterations using a cosine annealing schedule. We use three configurations: (1) $d_m = 768, \vec{d}_m = 3072, d_s = 1536, \vec{d}_s = 6144$, (2) $d_m = 768, \vec{d}_m = 4096$ and $d_s = 2048, \vec{d}_s = 6144$ and (3) $d_m = 1024, \vec{d}_m = 4096$ and $d_s = 2048, \vec{d}_s = 6144$. All models use $L = 12$. Our considered dataset, Wikitext-103, contains 103M tokens and has a vocabulary of nearly 270K.

C Extended Related Work

Improving Transformers Given the effectiveness of the Transformer, improving the architecture has been of much interest to the NLP community. Within this domain, one branch of research concerns the reduction of the quadratic complexity (w.r.t. sequence length) of the Transformer’s core self-attention mechanism (Wu et al., 2019; Kitaev et al., 2020), pushing it down to linear or log-linear complexity. The second branch of research regards improving the expressiveness of Transformer models, by using more layers (Dou et al., 2018), or by improving the architecture (Wu et al., 2019; So et al., 2019). The third branch of research regards improving the parameter efficiency of Transform-

ers. Approaches towards this goal include neural architecture search approaches (So et al., 2019; Wu et al., 2020), where new Transformer-based architectures are learned using gradient descent, more manually crafted approaches (Dehghani et al., 2019; Mehta et al., 2020a), as well as weight-sharing approaches (Lan et al., 2020; Wu et al., 2019). The work most similar to ours is ALBERT (Lan et al., 2020) in which complete weight sharing is used to pre-train deep contextualized word representations (Peters et al., 2018; Devlin et al., 2019). Different from this work, we focus on common NLP generative/sequence-to-sequence tasks versus large-scale pre-training and develop an approach to increase model capacity while reducing parameter footprint tailored to this setting.

Compressing Transformers We find prior work on pruning and quantizing Transformer models to reduce their size with a focus on sequence-to-sequence settings like machine translation (Prato et al., 2019), on encoder-based methods like BERT (Zafrir et al., 2019; Ganesh et al., 2020) or with a more generic scope in mind (Cheong and Daniel, 2019; Lee et al., 2019). Our approach is orthogonal to these since we directly aim at reducing the number of parameters of Transformer models by proposing architecture modifications and weight sharing techniques.

Reducing Embedding Dimensionality in Sequence Models As embeddings can substantially increase the parameter count as the vocabulary size increases, especially in sequence modeling scenarios, embedding reduction techniques have been proposed, including using a linear projection to project to a lower dimension (Baevski and Auli, 2019; Dai et al., 2019) or using combinations of block sparse transformations (Mehta et al., 2020b,a). We propose a self-attention based projection layer, SAFE, which we empirically show to outperform the aforementioned linear projection methods with a similar parameter count.