

MicroNet for Efficient Language Modeling

Zhongxia Yan

Massachusetts Institute of Technology

ZXYAN@MIT.EDU

Hanrui Wang

Massachusetts Institute of Technology

HANRUI@MIT.EDU

Demi Guo

Harvard University

DGUO@COLLEGE.HARVARD.EDU

Song Han

Massachusetts Institute of Technology

SONGHAN@MIT.EDU

Abstract

It is important to design compact language models for efficient deployment. We improve upon recent advances in both the language modeling domain and the model-compression domain to construct parameter and computation efficient language models. We use an efficient transformer-based architecture with adaptive embedding and softmax, differentiable non-parametric cache, Hebbian softmax, knowledge distillation, network pruning, and low-bit quantization. In this paper, we provide the winning solution to the NeurIPS 2019 MicroNet Challenge in the language modeling track. Compared to the baseline language model provided by the MicroNet Challenge, our model is 90 times more parameter-efficient and 36 times more computation-efficient while achieving the required test perplexity of 35 on the Wikitext-103 dataset. We hope that this work will aid future research into efficient language models, and we have released our full source code on [GitHub](#).

Keywords: Language model, model compression, Transformer

1. Introduction

Language modeling has been one of the most commonly studied sequence modeling tasks and one of the most studied tasks in the natural language processing community. Within the past few years, numerous works have improved the state-of-the-art language modeling results with deep neural network (DNN)-based sequence models. The Long Short-Term Memory (LSTM) network (Hochreiter and Schmidhuber, 1997) was designed to model dependencies in sequential data, and has successfully been applied to language modeling in a series of works, including the AWD-LSTM (Merity et al., 2017). Nevertheless, the LSTM architecture suffers from the vanishing gradient effect, limiting its ability to model long-term dependencies. On the other hand, the development of the Transformer attention architecture (Vaswani et al., 2017) inspired new Transformer-based language models, such as Transformer-XL (Dai et al., 2019), which achieved new state-of-the-art in language modeling benchmarks. In addition, several works have developed architecture-agnostic enhancements such as adaptive embedding and adaptive softmax (Baeovski and Auli, 2018), non-parametric cache (Grave et al., 2016), Hebbian softmax (Rae et al., 2018), and dynamic evaluation (Krause et al., 2017). While these works focused on improving the predictive

accuracy of language models, relatively few works have optimized for parameter and computational efficiency, which are critical for tasks with hardware constraints.

On the other hand, there have been numerous advances in parameter and computation efficient neural network architectures and model-compression techniques for deep neural networks. SqueezeNet (Iandola et al., 2016) and MobileNet (Howard et al., 2017) are efficient convolutional neural networks that take advantage of 1x1 convolutions and depthwise separable convolutions, respectively. Ma et al. (2019) applies Block-term Tensor Decomposition (Lathauwer, 2008) to compress transformer-based language models. Architecture agnostic model compression techniques such as knowledge distillation (Hinton et al., 2015), network pruning (Han et al., 2016), and trained quantization (Han et al., 2016) are commonly used techniques to either increase the predictivity of neural networks or decrease the model size. There are also recent efforts focusing on automatically designing efficient models (He et al., 2018; So et al., 2019; Wang et al., 2020b) and designing specialized accelerators to process the compressed models (Parashar et al., 2017; Zhang et al., 2020). Nevertheless, these techniques are mostly developed on convolutional neural networks and computer vision tasks, and only a handful of recent works like DistilBERT (Sanh et al., 2019) and HAT (Wang et al., 2020a) applied them to natural language tasks.

In this work, we integrate advances in both the language modeling domain and the model-compression domain to construct parameter- and computation-efficient language models. Specifically, we evaluated our model with the criteria of the NeurIPS 2019 MicroNet Challenge (Gale et al., 2019). Compared to the baseline language model provided by the MicroNet Challenge, our model is 90 times more parameter-efficient and 36 times more computation-efficient while maintaining good performance. Our entry into the MicroNet Challenge achieved the top performance in parameter- and computation-efficiency in the language modeling track.

2. NeurIPS 2019 MicroNet Challenge Language Modeling Track

The NeurIPS 2019 MicroNet Challenge asks contestants to build efficient yet still performant models (Gale et al., 2019). In particular, the language modeling track of the competition asks participants to train efficient word-level language models on the Wikitext-103 Dataset (Merity et al., 2016), a standard benchmark dataset composed of 103 million words in the training set, 217 thousand words in the validation set, and 245 thousand words in the test set, with a total of 267735 word tokens in the vocabulary. The challenge requires that entries achieve a word-level perplexity of below 35 on the test set, but otherwise does not factor perplexity into the score.

2.1. Scoring

Assuming an entry achieved the prerequisite perplexity threshold of 35, it is scored based on two criteria:

1. Parameter storage. This is the total number of 32-bit model parameters required to be loaded from disk to perform inference. Reduce-precision model parameters with less than 32-bits are counted as a fractional number. For example, an 8-bit parameter counts as $\frac{1}{4}$ th of a parameter.

2. Math operations. This is the mean number of 32-bit arithmetic operations required to perform inference on each token in the tokenized version of the test set. Multiplies and additions count separately. Tokens are assumed to be fed sequentially to the model, and the model must predict the next token before receiving it. Conversion to and from reduced-precision format (e.g., int8) do not count as operations. Reduced-precision operations are counted analogously to the parameter storage.

The total score of an entry is normalized by the size of the LSTM language model in [Rae et al. \(2018\)](#), which has 159M 32-bit parameters and 318M 32-bit math operations

$$\text{Score} = \frac{\text{Parameter Storage}}{159\text{M}} + \frac{\text{Math Operations}}{318\text{M}}.$$

3. Core Language Model

For a corpus of tokens $\mathbf{x} = (x_1, \dots, x_N)$, we take the standard approach to modeling the joint probability $P(\mathbf{x})$ by factorizing it into a product of conditional probabilities

$$P(\mathbf{x}) = P(x_1) \prod_{n=2}^N P(x_n | x_{1:n-1}).$$

We model each conditional probability $P(x_n | x_{1:n-1})$ term with a transformer-based neural network with a categorical distribution output. Here discuss the set of language modeling enhancements that are applicable on top of the vanilla transformer ([Vaswani et al., 2017](#)).

3.1. Transformer-XL Model

The Transformer-XL model extends the vanilla Transformer decoder model ([Vaswani et al., 2017](#)) by adding segment-level recurrence with state reuse and relative positional embeddings ([Dai et al., 2019](#)). These modifications improve the inference-time computational efficiency greatly when comparing with the vanilla Transformer. To maintain a fixed context size for predicting each token, the vanilla transformer requires full forward-computation of the entire context due to its absolute positional embedding, while the Transformer-XL merely needs to compute the forward-computation for one token due to its relative positional embeddings. We note this computational advantage of the Transformer-XL and incorporate it into our model.

3.2. Joint Optimization of Groups of Short Contexts

While Transformer-XL focuses on modeling long-term dependencies with per-layer context size C_{xl} on the order of thousands of tokens, we experiment with using *short* per-layer context size $C \ll C_{xl}$ in the MicroNet challenge; this approach allows us to train *multiple* contexts jointly and reduces the test-time computation cost.

Transformer-XL and our model both consist of L layers, allowing the size of the context to grow linearly with the number of layers. However, due to limitations of GPU memory, Transformer-XL only backpropagates to the previous $C_{xl} - 1$ dependencies at training-time—dropping the gradients to further dependencies—despite using $L(C_{xl} - 1) + 1$ previous

dependencies at test-time. In addition, Dai et al. (2019) reported diminishing decreases in perplexity as C_{xl} increases.

In contrast, we choose a relatively short per-layer context size C on the order of a few hundred tokens but jointly optimize over an *extended context* of C_e tokens. Each additional layer allows us to backpropagate to $C - 1$ more dependencies at training-time, so in total we backpropagate to dependencies up to $\min(C_e, L(C - 1) + 1)$ tokens away across L layers. In practice we choose $C_e \geq L(C - 1) + 1$ so that we can backpropagate to up to $L(C - 1) + 1$ dependencies at training-time. At test-time, each prediction has an efficient per-layer context of C and a total context of $L(C - 1) + 1$. In Appendix C, we verify that our choice of $C \ll C_{xl}$ is ideal in parameter- and computation-efficient settings. We illustrate joint training across short contexts in Figure 1.

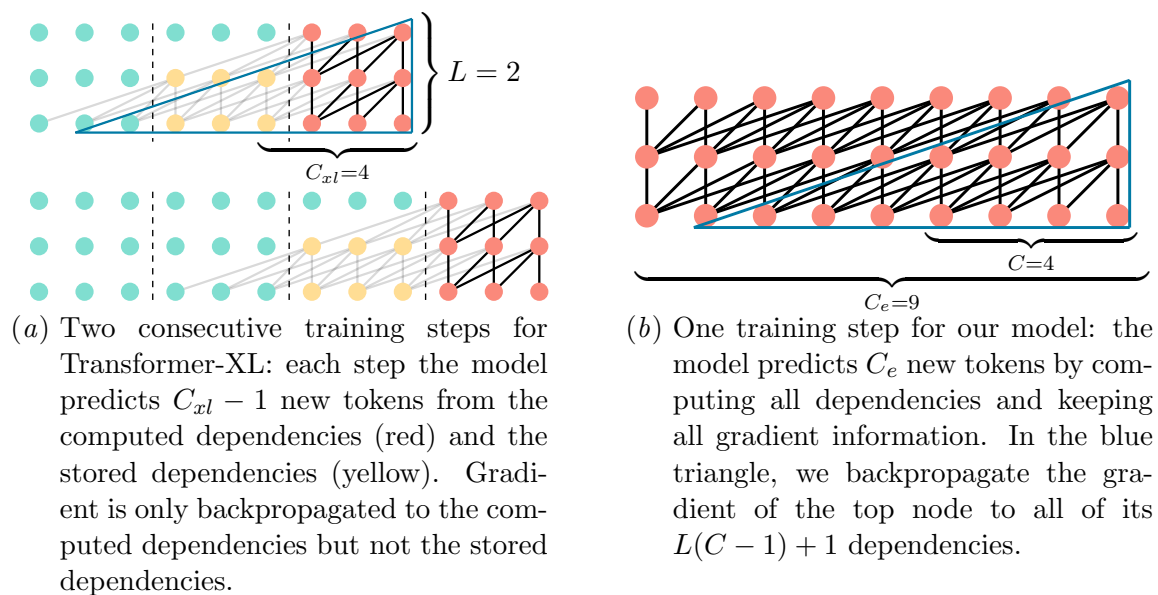


Figure 1: Transformer-XL vs. our model. Black lines represent forward- and backward-propagation while gray lines represent forward-propagation only. Each blue triangle represents all the dependencies of the top node in the triangle.

3.3. Adaptive Embedding and Softmax

We use the adaptive embedding and adaptive softmax as presented in Baevski and Auli (2018). Naive word embedding and softmax require parameter, memory, and computation costs proportional to the vocabulary size, which significantly reduces the batch size available for tasks like Wikitext-103, which has a vocabulary size of 267735. Intuitively, adaptive embedding and softmax allow us to allocate more representational power to the embedding space of tokens that occur more frequently in the training set; we sort the tokens in the vocabulary by frequency and assign smaller embedding vectors to less frequent vocabulary.

3.4. Hebbian Updates

We experiment with using Hebbian softmax (Rae et al., 2018), which updates the output embedding by combining traditional gradient descent and interpolation with the last hidden activations. Rae et al. (2018) suggests that non-parametric Hebbian updates help the embedding memorize the representations of infrequent tokens.

3.5. Differentiable Non-parametric Cache

We explore the use of a non-parametric cache with a cache size on the order of thousands of activations. As described by Grave et al. (2016), the non-parametric cache stores the most recent activations of the network and their corresponding labels, which improves prediction for infrequent tokens that appeared in the recent context. We define cache size n_{cache} and model activation h . To calculate cross entropy loss for token x_n , we use predicted token probability $P(x_n)$ given by

$$P(x_n|x_{1:n-1}) = (1 - \lambda_{\text{cache}})P_{\text{softmax}}(x_n|x_{1:n-1}) + \lambda_{\text{cache}}P_{\text{cache}}(x_n|x_{1:n-1})$$

$$P_{\text{cache}}(x_n|x_{1:n-1}) \propto \sum_{i=1}^{n_{\text{cache}}} e^{\theta_{\text{cache}} h_{n-i}^T h_n} \mathbb{1}_{x_{n-i}=x_n},$$

where P_{softmax} is the output probability of the language model and P_{cache} is the binned softmax of cache similarities.

While Grave et al. (2016) and Rae et al. (2018) perform hyperparameter grid-search over θ_{cache} and λ_{cache} on the validation set, we make an enhancement to jointly optimize the cache hyperparameters with gradient descent during training. We make an additional observation that batches at training-time are randomly sampled with no relation to the previous batch, while batches at test-time are iterated in order. To exploit this structure, we explore the effect of local search on the cache parameters θ_{cache} and λ_{cache} on the validation set *after* training. Since we may use a different cache size for local search, we define n_{cache}^t and n_{cache}^s to denote the cache size at training-time and local search-time, respectively.

4. Compression Techniques

We experiment with three compression techniques—knowledge distillation, pruning, and quantization—on top of our core language model.

4.1. Knowledge Distillation

The knowledge distillation technique (Hinton et al., 2015) first learns a larger, more predictive teacher model P_{teacher} then transfers its behavior into a more compact student model. For efficiency, we compute the largest 30 “soft” labels $S_{30} \subset \{P_{\text{teacher}}(x|x_{1:n-1}) : x \in \text{Vocab}\}$ for token x_n in the training set. We train the student model with distillation loss $\mathcal{L}_{\text{soft}}$ in addition to the standard crossentropy loss $\mathcal{L}_{\text{hard}}$, which uses x_n as the ground truth label. We apply teacher annealing (Clark et al., 2019) to linearly reduce the weight of the $\mathcal{L}_{\text{soft}}$

from λ_{\max} to λ_{\min} in T training steps. The total loss for token x_n at step t is

$$\begin{aligned}\mathcal{L}_t(x_n) &= (1 - \lambda_{\text{soft}}(t))\mathcal{L}_{\text{hard}}(x_n) + \lambda_{\text{soft}}(t)\mathcal{L}_{\text{soft}}(P_{\text{teacher}}) \\ \mathcal{L}_{\text{soft}}(P_{\text{teacher}}) &= - \sum_{x \in S_{30}} P_{\text{teacher}}(x|x_{1:n-1}) \log P(x|x_{1:n-1}) \\ \lambda_{\text{soft}}(t) &= \lambda_{\max} - \frac{t}{T}(\lambda_{\max} - \lambda_{\min})\end{aligned}$$

4.2. Pruning

To prune a given model, we perform sensitivity analysis (Li et al., 2016) on the model’s parameter matrices $\Phi = \{\phi_1, \dots, \phi_p\}$ to analyze how the performance of the model degrades with sparsity ρ_ϕ of each parameter. We empirically obtain the $\xi_\phi = f_\phi(\rho_\phi)$, which maps the sparsity of ϕ to the model perplexity ξ_ϕ . Next, we choose a target model sparsity ρ as a hyperparameter, which is a weighted combination of the parameter sparsities ρ_ϕ .

$$\rho = \frac{\sum_{\phi \in \Phi} |\phi| \rho_\phi}{\sum_{\phi \in \Phi} |\phi|} = \frac{\sum_{\phi \in \Phi} |\phi| f_\phi^{-1}(\xi^*)}{\sum_{\phi \in \Phi} |\phi|}$$

In practice, we solve for the perplexity threshold ξ^* by doing a binary search on values of ξ^* until we achieve a model sparsity of ρ . We then use Automatic Gradual Pruning (Zhu and Gupta, 2017) to simultaneously prune parameters with target sparsities $\rho_\phi = f_\phi(\xi^*)$, which specifies different levels of aggressiveness for different parameters. We explore the effect on model size and performance for several values of ρ .

4.3. Quantization

To quantize a given model, we choose a bit-width w to quantize our model to, then perform one step of quantization-aware training with linear-range symmetric fake-quantization (Jacob et al., 2018) on the model parameters and activations. To ensure that our model can be expressed in reduced-precision integer representation, we make sure that the last w bits of the mantissa of the inverse scale factors are 0. We perform quantization only on the non-embedding weights and on activations with no normalization operations; we do not quantize the output of layer normalization and softmax layers. Due to competition rules, we assume that all addition operations take place as 32-bit operations.

5. Experiments

We discuss experiments on the effect of the training cache size n_{cache}^t and the effects of compression techniques. Due to space constraint, we defer the detailed list of model configurations, comparison with state-of-the-art language models, analysis of per-layer context size, ablation study on the training cache sizes, ablation study on Hebbian softmax, and analysis of search cache sizes to Appendices A, B, C, D, E, and F, respectively. We depict our overall pipeline in Figure 2 with performances results and processing time estimates.

		Core LM	
Base	74M 149M	Teacher	53.9M 111M
	30hr		23.0 48hr
+ Adaptive Embed	8.3M 17M	Core LM	8.3M 18M
	39.5		33.6 30hr
+ Training Cache	18M	+ Distill	32.4
	33.7	Prune	5.6M 13M
+ Hebbian Softmax	33.6		34.6 15hr
Cache Search	33.0 3min	Cache Search	33.8 3min
		Quantize	1.8M 8.8M
			34.1 2min

Figure 2: Our pipeline without (left) and with (right) compression techniques, using the hyperparameters in Appendix A. From top to bottom, each stack displays the progression of techniques. Each row displays associated metrics: parameters (top left), operations (top right), validation perplexity (bottom left), and estimated processing time (bottom right). Metrics are displayed when changed from the previous row, with green for desirable change and red for undesirable. Red rows represent Core LM techniques, blue rows represent compression techniques, and gray rows represent cache search; joined rows represent joint training.

5.1. Overall Results

We report our performance with the hyperparameters detailed in Appendix A. Without compression techniques, our core language model (Core LM) achieves a validation perplexity of 33.6 with 8.3M parameters and 18M math operations. Adding the compression techniques, our model achieves a validation perplexity of 34.1 and test perplexity of 35.0, using 1.8M 32-bit parameters and 8.8M 32-bit math operations for a final MicroNet score of 0.0387. This is equivalent to a 90-fold reduction in parameter size and a 36-fold reduction in math operations compared to the MicroNet baseline. We compare our results to those of other MicroNet challenge language modeling track participants in Table 1.

Participant	MicroNet Score
Ours*	0.0387
MIT-HAN-Lab (Ours [†])	0.0475
Clova AI / Kyoto University	0.1657
JAIST / ISM	0.8232

Table 1: Performance of teams in the MicroNet Challenge Language Modeling track (Gale et al., 2019). Ours[†] is our original submission to the challenge, which contains an evaluation error. Ours* is the work presented in this paper with the error resolved. Clova AI / Kyoto University does not report methods. JAIST / ISM uses a QRNN base-model (Bradbury et al., 2016) rather than a Transformer base-model.

5.2. Cache Effect Analysis

We analyze the token-level effect of the non-parametric cache on different token bins within the validation loss. Grave et al. (2016) states that the cache increase prediction accuracy on rare labels within the cache range. We analyze three models with training cache sizes $n_{\text{cache}}^t \in \{\text{no-cache}, 1000, 2000\}$. In Figure 3(a), we bin the loss by token index sorted in decreasing order of token frequency and compare loss incurred in $n_{\text{cache}}^t \in \{\text{no-cache}, 1000\}$ to loss incurred in $n_{\text{cache}}^t = 2000$. We observe that most of the additional loss in $n_{\text{cache}}^t \in \{\text{no-cache}, 1000\}$ is attributed to rare tokens, with token indices $> 10^3$ in the vocabulary.

In Figure 3(b), we define the token gap to be the distance between a token and the previous occurrence of that token in the validation set. Similarly to above, we bin the loss by token gap and identify which token gap incur additional loss. We observe that $n_{\text{cache}}^t = 2000$ significantly outperforms no-cache on tokens with gaps between 100 and 2000 and slightly underperforms no-cache on tokens with gaps ≥ 2000 . We also observe that $n_{\text{cache}}^t = 2000$ outperforms $n_{\text{cache}}^t = 1000$ on tokens with gaps between 1000 and 2000. In both cases, the non-parametric cache improves prediction accuracy within the cache range.

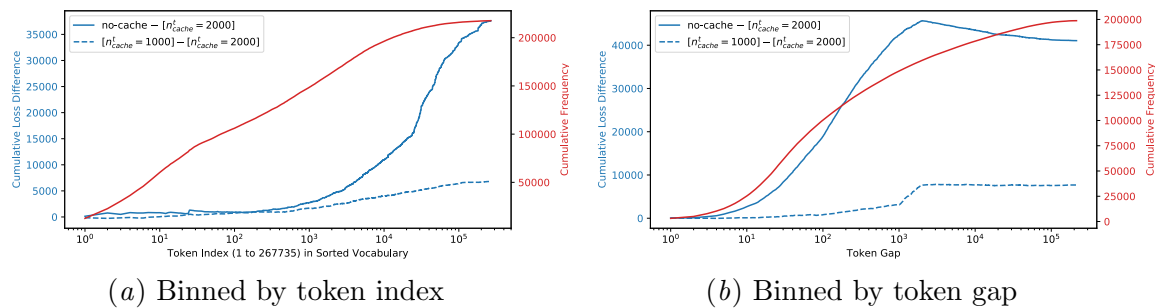


Figure 3: Blue: cumulative sum over the binned validation loss-difference between $n_{\text{cache}}^t \in \{\text{no-cache}, 1000\}$ and $n_{\text{cache}}^t = 2000$. Red: cumulative bin size for comparison. If the loss-difference were equally distributed across all tokens, the cumulative loss-difference would exactly match the cumulative bin size. All experiments are performed with $C = 97$, Hebbian softmax, and no compression techniques.

5.3. Effect of Distillation, Pruning, and Quantization

We take per-layer context $C \in \{97, 129, 257\}$ and run an ablation study on the compression techniques, comparing the performance at each stage of compression to the performance of the Core LM. We perform cache search *before* quantization because we empirically found that quantization immobilizes the cache parameters, and cache search after quantization fails to yield better parameters. Consistent with Appendix C, we observe that $C \in \{129, 257\}$ results in significantly lower perplexity than $C = 97$ for much of the pipeline before quantization. Surprisingly, the perplexities of $C \in \{129, 257\}$ increases significantly after quantization, while the perplexity of $C = 97$ only increases marginally. We

suspect that this phenomenon is due to the fact that $C = 97$ is more reliant on the cache; $C = 97$ has cache weight $\lambda_{\text{cache}} = 0.15$, larger than 0.13 for $C \in \{129, 257\}$. Overall, we find $[C = 97, \rho = 0.358, n_{\text{cache}}^s = 3000, \text{bits} = 9]$ to be the best hyperparameters for the MicroNet challenge.

Core LM		+Distill	+Prune		+Search		+Quantize						
C	Val	Val	ρ	Val	n_{cache}^s	Val	bits	Params	Ops	Val	Test		
97	33.6	32.9	0.239	33.6	2000	33.2	8	1.8M	9.5M	34.7	35.6		
							9	2.0M	9.7M	33.5	34.4		
			0.358	34.6	2000	34.1	3000	33.8	9	1.8M	8.5M	34.5	35.3
									10	1.9M	8.7M	34.2	35.1
									8	1.6M	8.6M	35.2	36.1
									9	1.8M	8.9M	34.1	35.0
			4000	33.6	8	1.6M	8.9M	35.1	36.0				
					9	1.8M	9.1M	34.0	34.8				
			0.477	36.2	2000	35.7	9	1.6M	7.5M	36.1	36.8		
					4000	35.2	10	1.7M	8.4M	35.3	36.0		
129	33.3	32.7	0.358	34.0	2000	33.5	9	1.8M	8.6M	35.0	36.0		
							10	1.9M	8.8M	33.8	34.8		
							3000	33.2	9	1.8M	8.9M	34.7	35.7
257	32.2	31.6	0.358	33.3	2000	32.9	9	1.8M	9.1M	34.3	35.2		

Table 2: Performance of different choices of per-layer context C , model sparsity ρ , local search cache size n_{cache}^s , and quantization bit. We bold the overall best result for the MicroNet challenge. “Val” and “Test” denote respective perplexities.

6. Conclusions

In our work, we combined transformer-based methods and cache-based language modeling methods to significantly reduce the total amount of parameters and computation while maintaining good perplexity. Furthermore, we showed that model-compression techniques achieve large reductions in parameter size and computation. We hope that this work will aid future research into efficient language models, and we have released our full source code on [GitHub](#).

Acknowledgments

We thank Facebook Faculty Award, AWS Machine Learning Award, and AMD for sponsoring this research. We are grateful to Phillip Isola for helpful discussions.

References

- Alexei Baevski and Michael Auli. Adaptive input representations for neural language modeling. *CoRR*, abs/1809.10853, 2018. URL <http://arxiv.org/abs/1809.10853>.
- Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models, 2019.
- James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. Quasi-recurrent neural networks. *CoRR*, abs/1611.01576, 2016. URL <http://arxiv.org/abs/1611.01576>.
- Kevin Clark, Minh-Thang Luong, Urvashi Khandelwal, Christopher D. Manning, and Quoc V. Le. Bam! born-again multi-task networks for natural language understanding. *CoRR*, abs/1907.04829, 2019. URL <http://arxiv.org/abs/1907.04829>.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *CoRR*, abs/1901.02860, 2019. URL <http://arxiv.org/abs/1901.02860>.
- Trevor Gale, Erich Elsen, Sara Hooker, Olivier Temam, Scott Gray, Jongsoo Park, Cliff Young, Utku Evci, Niki Parmar, and Ashish Vaswani. Micronet challenge, 2019. URL <https://micronet-challenge.github.io/>.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. Improving neural language models with a continuous cache. *CoRR*, abs/1612.04426, 2016. URL <http://arxiv.org/abs/1612.04426>.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016. URL <https://arxiv.org/abs/1510.00149>.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018. URL <https://arxiv.org/pdf/1802.03494.pdf>.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. URL <http://arxiv.org/abs/1704.04861>.
- Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. *CoRR*, abs/1602.07360, 2016. URL <http://arxiv.org/abs/1602.07360>.

- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014.
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. Dynamic evaluation of neural sequence models. *CoRR*, abs/1709.07432, 2017. URL <http://arxiv.org/abs/1709.07432>.
- Lieven Lathauwer. Decompositions of a higher-order tensor in block terms—part ii: Definitions and uniqueness. *SIAM J. Matrix Analysis Applications*, 30:1033–1066, 01 2008. doi: 10.1137/070690729.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *CoRR*, abs/1608.08710, 2016. URL <http://arxiv.org/abs/1608.08710>.
- Xindian Ma, Peng Zhang, Shuai Zhang, Nan Duan, Yuexian Hou, Dawei Song, and Ming Zhou. A tensorized transformer for language modeling. *CoRR*, abs/1906.09777, 2019. URL <http://arxiv.org/abs/1906.09777>.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *CoRR*, abs/1609.07843, 2016. URL <http://arxiv.org/abs/1609.07843>.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing LSTM language models. *CoRR*, abs/1708.02182, 2017. URL <http://arxiv.org/abs/1708.02182>.
- Angshuman Parashar, Minsoo Rhu, Anurag Mukkara, Antonio Puglielli, Rangharajan Venkatesan, Brucek Khailany, Joel Emer, Stephen W Keckler, and William J Dally. Scnn: An accelerator for compressed-sparse convolutional neural networks. *ACM SIGARCH Computer Architecture News*, 45(2):27–40, 2017. URL <https://arxiv.org/abs/1708.04485>.
- Jack W. Rae, Chris Dyer, Peter Dayan, and Timothy P. Lillicrap. Fast parametric learning with activation memorization. *CoRR*, abs/1803.10049, 2018. URL <http://arxiv.org/abs/1803.10049>.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2019.
- David R So, Chen Liang, and Quoc V Le. The evolved transformer. *arXiv preprint arXiv:1901.11117*, 2019. URL <https://arxiv.org/abs/1901.11117>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.

Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. Hat: Hardware-aware transformers for efficient natural language processing. In *Annual Conference of the Association for Computational Linguistics*, 2020a.

Tianzhe Wang, Kuan Wang, Han Cai, Ji Lin, Zhijian Liu, Hanrui Wang, Yujun Lin, and Song Han. Apq: Joint search for network architecture, pruning and quantization policy. In *Conference on Computer Vision and Pattern Recognition*, 2020b.

Zhekai Zhang, Hanrui Wang, Song Han, and William J. Dally. Sparch: Efficient architecture for sparse matrix multiplication. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2020. URL <https://arxiv.org/abs/2002.08947>.

Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression, 2017.

Neta Zmora, Guy Jacob, Lev Zlotnik, Bar Elharar, and Gal Novik. Neural network distiller, June 2018. URL <https://doi.org/10.5281/zenodo.1297430>.

Appendices

A. Configurations and Hyperparameters

Our best model use the following hyperparameters. Unless otherwise stated, experiments discussed in this paper also use these hyperparameters. The full implementation can be found in the code release.

1. Adaptive embedding and softmax: after sorting in order of increasing frequency, we divide the vocabulary into three bins: $[1, 3500]$, $[3501, 25000]$, $[25001, 267735]$. For the tokens in each bin, we use embedding vectors of sizes $[256, 64, 4]$, respectively.
2. Base transformer: we use a base transformer with context size $C = 97$, extended context $C_e = 1152$, $L = 8$ layers, input and output dimension $d_{\text{model}} = 256$, $h = 8$ attention heads, key and value dimensions $d_k = d_v = 24$, and an “inner” fully connected layer dimension $d_{\text{ff}} = 768$.
3. Cache training: we use a cache of size $n_{\text{cache}}^t = 2000$ with initial parameters $\theta_{\text{cache},0} = 0.016$ and $\lambda_{\text{cache},0} = 0.07$.
4. Hebbian softmax: we use a minimum discount factor $\gamma_{\text{hebbian}} = 0.01$ and a smoothing limit $T_{\text{hebbian}} = 500$ classes.
5. Training: we train for $T = 200000$ steps with the Adam Optimizer (Kingma and Ba, 2014), learning rate 0.0001, and cosine learning rate decay with 1000 linear-warmup steps. We choose the largest batch size that fits in memory.
6. Distillation: we distill from a teacher model trained with $L = 16$, embedding vectors of size $[512, 256, 16]$, $d_{\text{model}} = 512$, $d_k = d_v = 64$, $d_{\text{ff}} = 1536$, learning rate 0.0005, dropout 0.1, no cache, and otherwise the same hyperparameters as the student model. The teacher model has 53.8M parameters and a validation perplexity of 23.0. We train the student model with $\lambda_{\text{max}} = 0.5$ and $\lambda_{\text{min}} = 0.05$.
7. Pruning: we choose a target sparsity $\rho = 0.358$. We initialize the model with the best model from the distillation step, prune this model to an initial sparsity $\rho_0 = 0.16$, then train for 175000 steps, pruning every 1000 steps. We use the Distiller (Zmora et al., 2018) implementation of the Adaptive Gradual Pruning algorithm.
8. Cache local search: we perform local search on the cache parameters θ_{cache} and λ_{cache} with a search cache size $n_{\text{cache}}^s = 3000$.
9. Quantization: we quantize the model to 9 bits using the Distiller (Zmora et al., 2018) implementation of quantization-aware training and symmetric linear-range quantization.

We train the large teacher model on eight GeForce RTX 2080 Ti GPUs and all other models on a single GeForce GTX Titan X GPU, RTX 2080 Ti, or V100 GPU.

B. Comparisons to SOTA

In Table 3, we compare the performance of our models against performances of state-of-the-art methods reported by Bai et al. (2019). These methods use full word embeddings with embedding dimension 512, while we use embedding dimension 256 to reduce parameters. We compare a base-model with full word embedding and also compare our best small adaptive embedding models (with and without cache). We find that our base-model performance is comparable to Transformer-XL despite requiring half as many embedding parameters. Adding adaptive embedding and cache allows us to stay competitive with state-of-the-art methods while using much fewer parameters.

Model	Params	Embedding Params	Model Params	n_{cache}^t	Test PPL
Transformer-XL Small	139M	134M	4.9M	0	35.8
DEQ-Transformer Small	138M	134M	4.5M	0	32.4
Ours, Base	73.6M	68.5M	5.0M	0	36.5
Ours, Adaptive	8.3M	3.3M	5.0M	0	41.3
Ours, Adaptive, Cache	8.3M	3.3M	5.0M	2000	34.9

Table 3: Our models are competitive with SOTA transformer-based models while using much less embedding parameters. Here we do not use Hebbian softmax. We note that the DEQ Transformer Bai et al. (2019) uses root-finding methods to optimize model parameters and does not provide an estimate on computations required.

C. Context Length Experiment

We perform a sweep over per-layer context length C as detailed in Table 4 and analyze model performance and math operations. We see that using $C = 257$ or $C = 129$ achieves a good trade-off between the number of operations and the validation perplexity. Interestingly, $C \in \{513, 1025\}$ results in both increasing computation cost and worse perplexity; we suspect that the latter may be due to the small model size that we use.

D. Training Cache Size Ablation Study

We conduct an ablation study on the effect of the training cache size on perplexity. In Table 5, we see that the non-parametric cache is crucial to good performance; the 8.3M parameter model equipped with $n_{\text{cache}}^t = 2000$ performs similarly to a model with between 11M and 15M parameters with no cache.

E. Hebbian Softmax Ablation Study

We perform an ablation study on the Hebbian softmax technique presented in Rae et al. (2018). Unlike in their work, we do not see a significant decrease in perplexity when using

Context C	Params	Operations	Val PPL
1025	8.5M	23.4M	38.4
513	8.4M	20.1M	37.7
257	8.3M	18.5M	37.5
129	8.3M	17.6M	38.6
97	8.3M	17.4M	39.5
65	8.3M	17.2M	40.5

Table 4: Performance of different per-layer context C with only adaptive softmax; we do not use cache, Hebbian softmax, or compression techniques.

n_{cache}^t	Params	Operations	Val PPL
0	8.3M	17.4M	39.2
0	11.0M	23.3M	37.6
0	15.2M	31.6M	32.4
1000	8.3M	17.9M	34.8
2000	8.3M	18.4M	33.6

Table 5: Non-parametric cache improves the performance of the 8.3M parameter model greatly at little extra cost. We use $C = 97$ and Hebbian softmax for all models, but do not apply local cache search or compression techniques.

the Hebbian softmax technique. A potential explanation is that they use an LSTM base-model and regular softmax, whereas we use a transformer base-model and adaptive softmax.

		No Cache	With Trained Cache
Rae et al. (2018)	No Hebbian	36.0	34.5
	With Hebbian	34.1	29.7
Ours	No Hebbian	39.5	33.7
	With Hebbian	39.2	33.6

Table 6: Comparison of validation perplexity with and without Hebbian softmax in our model vs in Rae et al. (2018). We use context $C = 97$ for all experiments here. For our experiments with trained cache, we use $n_{\text{cache}}^t = 2000$.

F. Cache Training and Cache Search Experiment

As mentioned in Subsection 3.5, we sample random extended contexts C_e at training-time but sequentially iterate over contexts in the validation and test sets; the latter setup allows for stronger temporal locality. We therefore expect that even if we train cache weight λ_{cache} at training-time, increasing this hyperparameter yields better performance at test-times. We empirically observe that the validation perplexity is a smooth function of θ_{cache} and λ_{cache} with a clear global minimum, so we perform local search on θ_{cache} and λ_{cache} to minimize the validation perplexity, instead of using grid-search. In Table 7, we compare the performance of combinations of training cache size and search cache size, which we denote as $(n_{\text{cache}}^t, n_{\text{cache}}^s)$. Interestingly, (2000, no-search) significantly outperforms (no-cache, 2000), which suggests that jointly training of the cache hyperparameters is preferable to solely performing hyperparameter search over the validation set. We suspect that joint training learns more cache-friendly final layer activations h . We also observe that performing local search *after training* still improves performance, with bigger improvements for larger n_{cache}^s . Correspondingly, we observe that the cache weight λ_{cache} increases with n_{cache}^s .

Training n_{cache}^t	Search n_{cache}^s	Val PPL	λ_{cache}	θ_{cache}
no-cache	2000	34.8	0.0922	0.0197
no-cache	4000	34.7	0.0964	0.0213
no-cache	6000	34.8	0.0963	0.0221
2000	no-search	33.6	0.0689	0.0242
2000	2000	33.0	0.145	0.0266
2000	4000	32.6	0.156	0.0276
2000	6000	32.4	0.158	0.0281
2000	8000	32.4	0.160	0.0284

Table 7: Performance of models with different combinations of training and local search cache sizes. Note that (2000, 2000), (2000, 4000), (2000, 6000), and (2000, 8000) merely performs local search on the trained model from (2000, no-search). We use $C = 97$ and Hebbian softmax for all models.