

When to Use Efficient Self Attention? Profiling Text, Speech and Image Transformer Variants

Anuj Diwan, Eunsol Choi, David Harwath

Department of Computer Science

The University of Texas at Austin

{anuj.diwan, eunsol, harwath}@utexas.edu

Abstract

We present the first unified study of the efficiency of self-attention-based Transformer variants spanning text, speech and vision. We identify input length thresholds (*tipping points*) at which efficient Transformer variants become more efficient than vanilla models, using a variety of efficiency metrics (latency, throughput, and memory). To conduct this analysis for speech, we introduce L-HuBERT, a novel local-attention variant of a self-supervised speech model. We observe that these thresholds are (a) much higher than typical dataset sequence lengths and (b) dependent on the metric and modality, showing that choosing the right model depends on modality, task type (long-form vs. typical context) and resource constraints (time vs. memory). By visualising the breakdown of the computational costs for transformer components, we also show that non-self-attention components exhibit significant computational costs. We release our profiling toolkit at <https://github.com/ajd12342/profiling-transformers>.

1 Introduction and Related Work

Transformers (Vaswani et al., 2017) are widely adopted across NLP (Devlin et al., 2019; Brown et al., 2020), Speech Processing (Mohamed et al., 2022) and Computer Vision (Dosovitskiy et al., 2021). Studies have shown that scaling models up improves performance (Chowdhery et al., 2022), making efficiency an important research topic. Many Transformer variants focus on improving the efficiency of self-attention, motivated by its asymptotic quadratic time/space complexity with respect to the input sequence length.¹ While these Transformer variants are designed to be asymptotically faster, in practice they may actually be slower, especially given modest input lengths that are typical of many tasks.

¹We refer the readers to Tay et al. (2022) for a comprehensive overview of efficient Transformers.

Our paper presents two main analyses. First, we visualize the *layerwise* efficiency of such models to locate bottlenecks and attempt to answer the question “*is self-attention the true bottleneck?*” We find that in the non-asymptotic case, non-self-attention layers contribute significantly to the overall cost, especially for speech architectures due to the input waveform tokenizer in models like HuBERT (Hsu et al., 2021). Second, *when* should we use self-attention-based efficient Transformers? Comparing efficient variants with vanilla models at different input lengths, we find that this *tipping point* where efficient variants outperform vanilla architectures is much higher than typical input lengths of existing benchmarks across all modalities, calling into question the efficacy of using such efficient Transformers and requiring new benchmarks. We introduce a local-attention variant of a speech Transformer, HuBERT, to conduct this analysis. Together, our analyses suggest that current approaches that focus on improving self-attention might not be the most effective for improving efficiency.

2 Efficiency Metrics

Model efficiency is an umbrella term for a suite of *efficiency metrics*, which do not always correlate with, and sometimes contradict, each other (Dehghani et al., 2022). Further, different metrics are relevant to different end use-cases. To cover most use-cases, we evaluate a set of four metrics; two for computational time and two for memory usage: **Throughput**: Number of examples processed per sec, given inputs of a given sequence length, using the maximum possible batch size for a given GPU. **Latency-Inference**: Time (in ms) to run inference for 1 unbatched input of a given sequence length. **Max-Memory**: The allocated GPU memory (MiB) for processing 1 input of a given sequence length. **Parameter Count**: Number of model parameters.

We profile models in all modalities in *training* mode and *inference* mode. For training, while

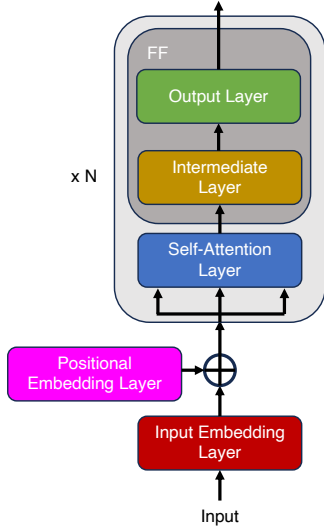


Figure 1: Transformer layer types profiled in our layerwise efficiency profiling experiments.

Transformer architectures often use prediction heads with a larger output space (e.g., for text generation), we choose a lightweight binary classification head for profiling.

Layerwise Efficiency Metrics We also profile some metrics and models in a **layerwise** fashion to locate their efficiency bottlenecks. Our goal is twofold: a) provide an empirical approach to efficient model design, as an alternative to theoretical analyses or mental models (e.g. self-attention is $O(n^2)$) and b) empirically answer the question "to what degree is self-attention the bottleneck?"

We identify important layer types (Self-Attention, Feedforward, etc.) and profile the Latency-Inference and Parameter Count metrics per-layer-type to obtain a fine-grained understanding of which layer types and indices (layer 0 vs 11) contribute the most to model efficiency costs. We use param counts as a proxy for memory (profiling real layerwise memory usage is non-trivial due to Pytorch memory allocation intricacies). We profile the layers depicted in Figure 1; more details in Appendix E.

3 Local-Attention Speech Model

Efficient transformers (Xiong et al., 2021; Ma et al., 2021) have not received as much attention in Speech as they have in NLP and CV, perhaps due to two reasons. First, there is a relative lack of long-context speech benchmarks as compared to those in NLP (LRA (Tay et al., 2021) and QuALITY (Pang et al., 2022)). Second, when performing speech

Model	WER ↓	WER (w/ FT) ↓
HuBERT Base	7.09	3.4
L-HuBERT (32 100)	21.06 14.48	8.52 7.39

Table 1: WERs on the SUPERB ASR task.

tasks like automatic speech recognition (ASR), it is typical to segment a long speech signal into small individual utterances and perform ASR independently on each. For example, most Librispeech examples are less than 5 seconds. Many popular speech models like HuBERT (Hsu et al., 2021) tokenize the waveform at 50 tokens per second, implying that a typical utterance has only several hundred tokens; far below the number of tokens in long-context NLP tasks. Nevertheless, textless speech models (Lakhotia et al., 2021) are becoming more feasible, motivating the modelling of long speech utterances.

Local HuBERT Model To investigate the efficiency of the self-attention layer in speech models, we introduce the *Local HuBERT* model which replaces HuBERT’s self-attention with the Longformer (Beltagy et al., 2020) sliding-window self-attention. In this attention mechanism, every token attends to tokens within a local window context, rather than the full token sequence. Our model is similar to the temporally windowed-attention Transformer acoustic model proposed by Alastruey et al. (2021) for speech translation; our approach differs by using the self-supervised HuBERT model as our basis, and we evaluate on ASR. Choosing the appropriate window size for the local attention context is key; we explore 32 and 100 token contexts, corresponding to 640 ms and 2 s, inspired by phone recognition models that typically incorporate similar context sizes (Peddinti et al., 2015; feng Yeh et al., 2019).

ASR Results We initialize the L-HuBERT model with pretrained HuBERT Base weights (pretrained with full self-attention), and then replace self-attention with sliding-window self-attention; due to limited compute, we did not pretrain L-HuBERT from scratch using sliding-window attention. We then evaluate L-HuBERT on Librispeech (Panayotov et al., 2015) ASR via the SUPERB (Yang et al., 2021) benchmark under two settings; a) **Freeze**: freezing the model and only training projection weights and b) **Finetune**: fully finetune the model. We use the default S3PRL² hyperparams; but we

²<https://github.com/s3prl/s3prl>

Model	Emb	Pos	SA	Interm	Output	Others
BERT	23.8M	-	29M	28.3M	28.3M	0.6M
HuBERT	4.2M	5.1M	29M	28.3M	28.3M	0.2M
ViT	0.6M	-	29M	28.3M	28.3M	0.6M

Table 2: Layer-wise parameter counts. Emb: Input Embedding, Pos: Positional Emb. SA: Self-Attention, Interm: Intermediate.

train for 200k steps for Freeze and 104k steps for Finetune. Both models converge by 104k steps; we train Freeze for longer to eke out as much performance as possible, while we stop training Finetune due to limited compute.

We report Word Error Rate (WER) on Librispeech test-clean in Table 1; lower is better. In the frozen setting (middle column), we see a large WER increase over HuBERT; we hypothesize that this is due to the attention layer mismatch since we initialize L-HuBERT with HuBERT weights that were pretrained with full self attention, rather than pretraining L-HuBERT from scratch. However, in the finetuning setting, the gap between HuBERT Base and L-HuBERT narrows considerably and using a larger window size achieves better performance. As our L-HuBERT model is a reasonable architecture capable of moderate ASR performance, we can continue to study its computational efficiency (we profile the window-100 variant).

4 Methods and Implementation

We analyze the Base versions of the BERT (Devlin et al., 2019), Longformer (Beltagy et al., 2020) and Nyströmformer (Xiong et al., 2021) models for text; the HuBERT (Hsu et al., 2021) and L-HuBERT (Section 3) models for speech; and Vision Transformer (Dosovitskiy et al., 2021) and Swin Transformer (Liu et al., 2021) models for vision; BERT, HuBERT and ViT are standard Transformer encoder architectures. Longformer, L-HuBERT and Swin use fixed-pattern self-attention while Nyströmformer uses approximate self-attention.

4.1 Sequence Length Ranges

We profile our models on a wide range of input sequence lengths to cover both avg. sequence lengths of commonly used contemporary datasets (Table 3) and typical sequence lengths of long-context tasks. Details about how we compute sequence lengths in Table 3 can be found in Appendix B. Most image datasets use images resized to 224 or 512 pixels. Below, $\text{range}(a, b, c)$ means a range from a to b

in steps of c . Since there is no difference between synthetic and real inputs from a computational complexity standpoint, we use synthetic inputs to more easily control for their sequence lengths.

Text Modality The input is ‘*This is a sentence.*’ repeated n times, $n \in \text{range}(10, 560, 10)$ i.e. $\text{range}(62, 3362, 60)$ tokens for all tokenizers.

Speech Modality The inputs have durations in $\text{range}(1, 50, 0.5)$ sec i.e. $\text{range}(50, 2500, 25)$ tokens for all featurizers (CNNs with 20 ms framerate). Our sampling strategy is in Appendix A.

Image Modality We use square inputs of dimension in $\text{range}(32, 1024, 32)$ pixels by rescaling a fixed image. The # tokens depend on featurizer patch size, which is different for different models.

4.2 Implementational Details

We profile time-based metrics (latency/throughput) using Pytorch CUDA Events³ by executing 20 iterations sequentially. The first few iterations serve as GPU warm-start; thus, we report the average of the last 10. We record Max-Memory with `torch.cuda.max_memory_allocated()` and param counts with `torchinfo` (Yep, 2020).

To profile throughput, we *approximate* the max batch size that fits on a single GPU using a linear estimator; more details in Appendix C. Finally, we profile the layerwise Latency-Inference metric using `torchprof` (Wong, 2020). We attach profiling hooks to modules of interest (e.g. Self-Attention, Embedding), giving us execution times of their `forward()` functions (other modules/functions are not profiled). We use the Huggingface (Wolf et al., 2020) implementations of text and image models and `fairseq` (Ott et al., 2019) implementations for speech models; more details in Appendix D.

5 Profiling Results

5.1 Layerwise Profiling Results

Figure 2 shows the layerwise Latency-Inference for all 3 vanilla architectures in each modality. Figures for efficient models are in Appendix F. Color darkness represents the layer index (layer 0 is darkest). Table 2 shows the layerwise param count.

Asymptotically, self-attention dominates the computation. However, since the average seq length for most text and speech tasks is less than 1000 tokens and most image datasets are used at

³<https://pytorch.org/docs/stable/generated/torch.cuda.Event.html>

	Text							Speech					
Dataset	SST	MNLI	SQ	ON	CNN	HPQA	TQA	TEDL	LJS	VoxC	Libri	S-SQuAD	Spotify
# of tokens	23	36	177	506	863	1,316	6,589	301	328	390	615	3080	101400

Table 3: Average token sequence lengths. Left to right: Stanford Sentiment Treebank, MultiNLI, SQuAD2.0, OntoNotes, CNN-DailyMail, HotpotQA, TriviaQA, TEDLIUM, LJSpeech, VoxCeleb Speaker Recognition, Librispeech, Spoken SQuAD, Spotify Podcasts.




Figure 2: Layerwise latency of different vanilla Transformer architectures in inference mode.

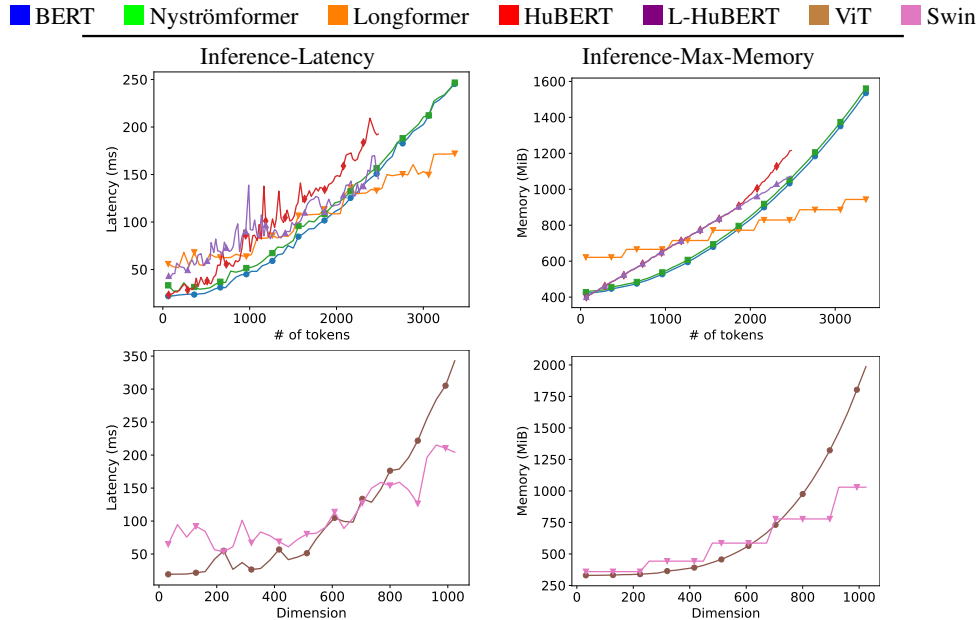


Figure 3: Overall Inference-time Profiling Results. Text and speech models in first row, vision models in second.

a max dimension of 512, at these points, non-self-attention components take up 35%, 58.8% and 43.75% latency for NLP, speech and images. Additionally, parameter counts of SA are also comparable to Intern/Output layers. This shows that it is also important to direct efficiency efforts for other model components.

While the latency associated with embedding layers is minimal for BERT, they are sizable for HuBERT. HuBERT uses a CNN feature extractor with different strides and kernel sizes and consumes more time in the earlier CNN layers as opposed to later ones, as is visible in Figure 2, which shows

darker shades i.e. earlier layers dominating the computation. Optimal efficiency strategies can thus differ across modalities, e.g. Wu et al. (2022) slims down this CNN feature extractor embedding layer. On the other hand, embedding layers take up a lot of *parameters* in BERT; thus, it may be helpful to shrink the BERT embedding layer for memory purposes (as opposed to *latency* for HuBERT). Finally, analyzing Transformer variants (Appendix F), we see that self-attention in Longformer, Swin and L-HuBERT encouragingly scales latency linearly, but with large overhead for smaller inputs.

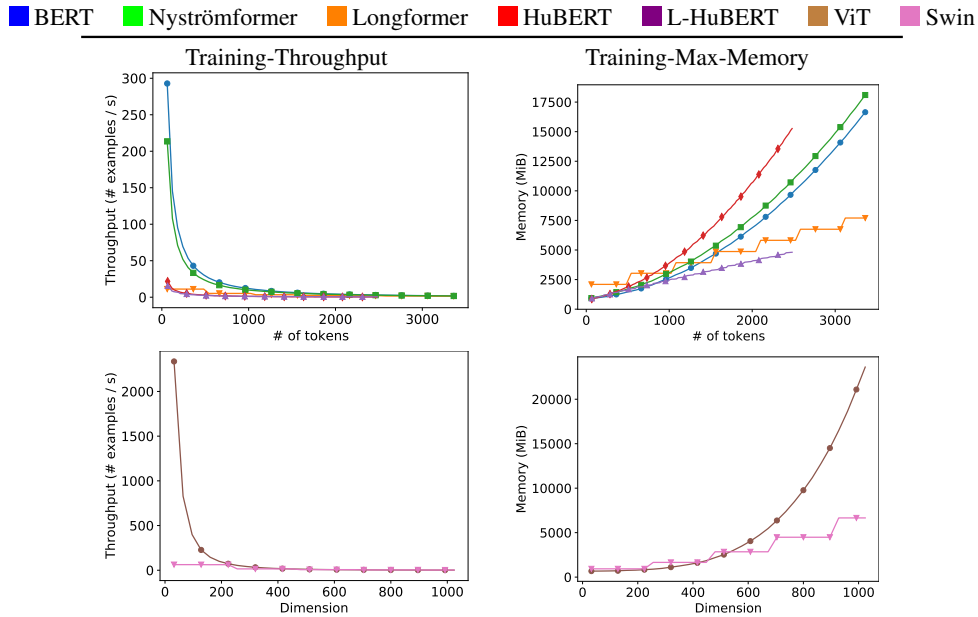


Figure 4: Overall Training-time Profiling Results. Text and speech models in first row, vision models in second.

5.2 Overall Profiling Results

Our profiling results are in Figures 3 and 4. Inference Throughput is in the Appendix at Figure 6, exhibiting similar trends as training Throughput.

Tipping Point Analysis We see that most variants are slower and more memory hungry than vanilla models for input lengths of typical-context tasks. We define the *tipping point* for each modality: the input length at which the variant becomes more efficient than the vanilla model. For text and speech, it is 1750 – 2000 tokens for inference latency and max-memory, greater than typical input lengths (Table 3). However, while the tipping point for training max-memory is ≈ 1500 tokens for text (still a large number), it is $\approx 0 - 250$ for speech, an encouraging result. For images, it is 500 – 700 pixels for all metrics apart from throughput. This is less reasonable for 224 pixel datasets but good for high resolution image datasets (512/1024). All variants are either worse or comparable than vanilla models across modalities for throughput.

We hypothesize that some efficient models suffer from additional overheads; while vanilla attention benefits from highly optimized matrix multiplication, windowed attention requires complex reshaping and preprocessing.

Choosing the Right Model Depends on Resource Constraints Our results show that the choice of the right model depends on resource constraints. Suppose that one is training models under a time constraint; then, throughput is the bottleneck and

efficient models would not be a good fit. On the other hand, efficient models are useful for long context memory-constrained inference.

Local Attention and Excessive Padding The Longformer pads input lengths to be a multiple of 512 and Swin requires input dimension to be a multiple of 224. This slows shorter inputs down and results in extremely low performance (measured by all 3 metrics) as compared to vanilla models.

Comparing Parameter Counts The Longformer uses more parameters compared to vanilla BERT (148M vs. 109M) because it uses two sets of Q,K,V projection matrices for its global and local attention operations; sharing these may decrease its memory usage. For other modalities, efficient models do not incur more parameters.

6 Conclusion

We present an empirical efficiency analysis of vanilla Transformers and their self-attention-based efficient variants across modalities, metrics and input context sizes. We find substantial differences across modalities and metrics when analyzing the tipping point for efficient variants. Finally, the layerwise analysis finds that self-attention is not the only bottleneck. We recommend that all efficient model papers should report such cross-modal, layerwise profiling results on multiple efficiency metrics covering a variety of use-cases to provide a full picture of the benefits of the model.

Limitations

We focus primarily on comparing model efficiencies using a variety of efficiency metrics and do not consider model performance; one can perform a more elaborate analysis of performance-efficiency tradeoffs, which we did not do here.

We only profile a total of seven models across three modalities while there are more efficient variants and vanilla Transformers proposed in the literature. While we choose our models to be as representative of each modality and efficiency technique as possible, we cannot extrapolate results to other model variants and other modalities. In particular, modalities like video and genomics and efficiency approaches like quantization would be interesting to profile, which we did not do.

Acknowledgements

We thank the reviewers and the meta-reviewer of the ACL community for helpful feedback on the draft. This work was partially funded by a grant from UT Machine Learning Lab.

References

- Belen Alastruey, Gerard I. Gállego, and Marta R. Costajussà. 2021. [Efficient Transformer for Direct Speech Translation](#).
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The Long-Document Transformer](#). *ArXiv preprint*, abs/2004.05150.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [PaLM: Scaling Language Modeling with Pathways](#).
- Ann Clifton, Sravana Reddy, Yongze Yu, Aasish Pappu, Rezvaneh Rezapour, Hamed Bonab, Maria Eskevich, Gareth Jones, Jussi Karlgren, Ben Carterette, and Rosie Jones. 2020. [100,000 Podcasts: A Spoken English Document Corpus](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5903–5917, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Mostafa Dehghani, Yi Tay, Anurag Arnab, Lucas Beyer, and Ashish Vaswani. 2022. [The Efficiency Mismatcher](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. 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.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Ching feng Yeh, Jay Mahadeokar, Kaustubh Kalgaonkar, Yongqiang Wang, Duc Le, Mahaveer Jain, Kjell Schuberth, Christian Fuegen, and Michael L. Seltzer. 2019. [Transformer-transducer: End-to-end speech recognition with self-attention](#). *ArXiv*, abs/1910.12977.
- 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.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. **Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation.** In *Speech and Computer*, pages 198–208. Springer International Publishing.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. **HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units.** *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Keith Ito and Linda Johnson. 2017. **The LJ Speech Dataset.** <https://keithito.com/LJ-Speech-Dataset/>.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. **TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension.** In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Kushal Lakhotia, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, and Emmanuel Dupoux. 2021. **On generative spoken language modeling from raw audio.** *Transactions of the Association for Computational Linguistics*, 9:1336–1354.
- Chia-Hsuan Li, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee. 2018. **Spoken SQuAD: A Study of Mitigating the Impact of Speech Recognition Errors on Listening Comprehension.** In *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*, pages 3459–3463. ISCA.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. **Swin Transformer: Hierarchical Vision Transformer using Shifted Windows.** In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 9992–10002. IEEE.
- Xuezhe Ma, Xiang Kong, Sinong Wang, Chunting Zhou, Jonathan May, Hao Ma, and Luke Zettlemoyer. 2021. **Luna: Linear Unified Nested Attention.** In *NeurIPS*.
- Abdelrahman Mohamed, Hung yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaloe, Tara N. Sainath, and Shinji Watanabe. 2022. **Self-Supervised Speech Representation Learning: A Review.** *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1179–1210.
- Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. **VoxCeleb: A Large-Scale Speaker Identification Dataset.** In *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, pages 2616–2620. ISCA.
- 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.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. **Librispeech: An ASR corpus based on public domain audio books.** In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pages 5206–5210. IEEE.
- Richard Yuanzhe Pang, Alicia Parrish, Nitish Joshi, Nikita Nangia, Jason Phang, Angelica Chen, Vishakh Padmakumar, Johnny Ma, Jana Thompson, He He, and Samuel Bowman. 2022. **Quality: Question answering with long input texts, yes!** In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5336–5358, Seattle, United States. Association for Computational Linguistics.
- Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. **A time delay neural network architecture for efficient modeling of long temporal contexts.** In *Proc. Interspeech 2015*, pages 3214–3218.
- Sameer S. Pradhan and Nianwen Xue. 2009. **OntoNotes: The 90% solution.** In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Tutorial Abstracts*, pages 11–12, Boulder, Colorado. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. **Know what you don’t know: Unanswerable questions for SQuAD.** In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. **Recursive deep models for semantic compositionality over a sentiment treebank.** In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.

- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. 2021. [Long Range Arena : A Benchmark for Efficient Transformers](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2022. [Efficient Transformers: A Survey](#). In *ACM Comput. Surv.*, volume 55, New York, NY, USA. Association for Computing Machinery.
- 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.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). 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 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Alexander William Wong. 2020. torchprof. <https://github.com/awwong1/torchprof>.
- Felix Wu, Kwangyoung Kim, Jing Pan, Kyu J. Han, Kilian Q. Weinberger, and Yoav Artzi. 2022. [Performance-Efficiency Trade-Offs in Unsupervised Pre-Training for Speech Recognition](#). In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7667–7671.
- Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. 2021. [Nyströmformer: A Nyström-based Algorithm for Approximating Self-Attention](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 14138–14148. AAAI Press.
- Shu-Wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y. Lin, Andy T. Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, Tzu-Hsien Huang, Wei-Cheng Tseng, Kottik Lee, Da-Rong Liu, Zili Huang, Shuyan Dong, Shang-Wen Li, Shinji Watanabe, Abdelrahman Mohamed, and Hung-yi Lee. 2021. [SUPERB: Speech Processing Universal PERFORMANCE Benchmark](#). In *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association, Brno, Czechia, 30 August - 3 September 2021*, pages 1194–1198. ISCA.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Tyler Yep. 2020. torchinfo. <https://github.com/TylerYep/torchinfo>.

A Sampling Speech Utterances for Profiling

To obtain speech inputs of length i seconds to $i + 0.5$ seconds for all i less than 12 seconds, we sample 5 speech utterances from the training set of the Librispeech dataset (Panayotov et al., 2015) whose lengths fall within this range and compute aggregate metrics over these 5 utterances. Since the Librispeech dataset does not contain extremely long speech utterances, for i of length greater than 12 seconds, we adopt a different approach to generate inputs. To generate such an input utterance of length between i and $i + 0.5$ seconds, we first sample 5 speech utterances from the Librispeech training set of input length ranging from $\frac{i}{5}$ to $\frac{i+0.5}{5}$ and concatenate them to obtain utterances of length ranging from i to $i + 0.5$ as desired. We do this 5 times to get 5 different utterances and compute aggregate metrics over these 5 utterances.

B Computing Token Lengths for NLP and Speech Datasets

We compute average sequence token lengths for 7 NLP datasets and 6 speech datasets. For all speech datasets, we compute mean utterance durations and multiply durations by 50 to obtain number of tokens (model framerates are 20 ms i.e. $\times 50$). For TEDLIUM (Hernandez et al.), LJSpeech (Ito and Johnson, 2017), VoxCeleb Speaker Recognition Dataset (Nagrani et al., 2017) and Librispeech (Panayotov et al., 2015), we compute

mean validation-set *utterance* durations; for Spoken SQuAD (Li et al., 2018), we report mean validation-set *paragraph* duration and for the Spotify English Podcasts dataset (Clifton et al., 2020), we report mean *podcast* duration directly obtained from Clifton et al. (2020).

SST (Socher et al., 2013). We use test-set sentences. We use the HuggingFace BERTTokenizer. **MNLI** (Williams et al., 2018). We use validation-matched-set examples by concatenating the premise and the hypothesis. We use the HuggingFace BERTTokenizer.

SQuAD2.0 (Rajpurkar et al., 2018). We use validation-set examples by concatenating the context and the question. We use the HuggingFace BERTTokenizer.

OntoNotes (Pradhan and Xue, 2009). We obtain this number from the Longformer (Beltagy et al., 2020) paper.

CNN-Dailymail (Hermann et al., 2015). We use the 3.0.0 version of the dataset and use test-set articles. We use the HuggingFace BERTTokenizer.

HotpotQA (Yang et al., 2018). We obtain this number from the Longformer (Beltagy et al., 2020) paper.

TriviaQA (Joshi et al., 2017). We obtain this number from the Longformer (Beltagy et al., 2020) paper.

C Implementing Throughput Profiling

To profile Throughput, we need to compute the max batch size that can fit on a single GPU. We *approximately* predict this using a linear estimator as follows. We first record the memory B reserved on the GPU after just loading the model. Next, we independently run batches of sizes 1 and 2 and record memory usages M_1 and M_2 . We use an NVIDIA Quadro RTX 8000 GPU with a maximum memory of 45000 MiB. Thus, assuming a linear relationship between batch size and memory consumption, we predict a maximum batch size of $bsz = \frac{45000-B}{M_2-M_1}$. In practice, this is an overestimate; we keep decreasing the batch size by a factor of 0.9 until no OOM errors occur and this is our final estimate.

D Implementational Details for Models

We use the following HuggingFace configurations: bert-base-uncased for BERT, allenai/longformer-base-4096 for Longformer, uw-madison/nystromformer-4096 for Nyströmformer,

google/vit-base-patch16-224 for ViT and microsoft/swin-base-patch4-window7-224 for Swin. The BERT model natively supports a maximum of 512 tokens as input because it has 512 positional embeddings; we modify the positional embedding computation to allow an arbitrarily long input to be provided. The Longformer internally pads all input lengths to a multiple of 512. For Swin, we pad images to have an input dimension that is a multiple of 224; this is necessary due to the windowed attention mechanism in Swin. In fact, the Swin model natively supports only a 224×224 resolution; we make a small modification in order to support resolutions that are multiples of 224. We use the HuBERT Base model for both HuBERT and L-HuBERT.

E Transformer Layer Types

Input Embedding Layer. (■/red) Maps the input sequence into fixed-dimensional embeddings. This is a linear layer for text and a CNN featurizer for image/speech.

Positional Embedding Layer. (■/fuchsia) For text and image models this is part of the input embedding layer. For speech models, this is a very wide convolution layer.

Self Attention Layer.(■/blue) The multi-head self attention block, which computes self-attention outputs and maps the result to the model dimension.

Intermediate Layer.(■/yellow) Linear layer of the feedforward block that maps the output from the Self Attention block into the ‘feedforward dimension’ (typically 4x the model dimension).

Output Layer.(■/green) Second linear layer of the feedforward block, which maps the output from Intermediate layer back to the model dimension.

Other Layers.(■/black) Other modules (activations, layer normalizations, other linear layers, etc.) not covered by the above components.

F Additional Profiling Analyses

We report layerwise profiling runs for efficient self-attention variants and inference-time throughput profiling runs for all variants in this section at Figures 5 and 6.

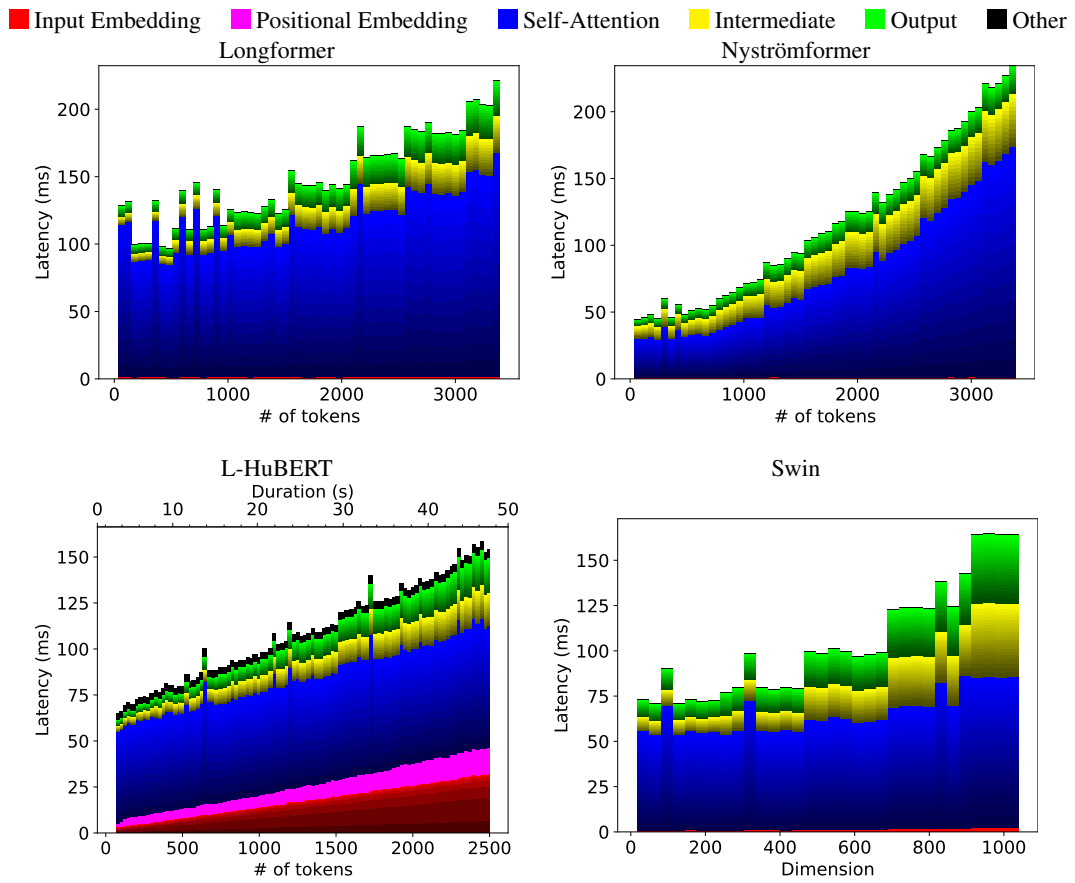


Figure 5: Layerwise latency of different Transformer variants in inference mode.




Figure 6: **Throughput** Profiling Results in inference mode.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
The Limitations section
- A2. Did you discuss any potential risks of your work?
The Limitations section
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3, 4

- B1. Did you cite the creators of artifacts you used?
Section 3, 4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not explicitly, since we use publicly available Huggingface and Fairseq models that are intended for research use
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
We use publicly available Huggingface and Fairseq models that are intended for research use
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 4 and 4.2, Appendices B,D
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
We only use datasets to profile models over different sequence lengths, but don’t use the content of the dataset itself. Thus we report the relevant statistic i.e. dataset sequence length.

C Did you run computational experiments?

Section 3, 4.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4.1, 4.2, Appendix C.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 4.2

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4.2

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 3, 4

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.