

---

# Hysteresis Activation Function for Efficient Inference

---

**Moshe Kimhi\***

Technion - Israel Institute of Technology  
Department of Computer Science  
moshekimhi@cs.technion.ac.il

**Idan Kashani\***

Technion - Israel Institute of Technology  
Department of Computer Science  
idan-kashani@cs.technion.ac.il

**Avi Mendelson**

Technion - Israel Institute of Technology  
Department of Computer Science

**Chaim Baskin**

Ben-Gurion University of the Negev  
School of Electrical and Computer Engineering

## Abstract

The widely used ReLU is favored for its hardware efficiency, as the implementation at inference is a one bit sign case, yet suffers from issues such as the “dying ReLU” problem, where during training, neurons fail to activate and constantly remain at zero, as highlighted by Lu et al. [16]. Traditional approaches to mitigate this issue often introduce more complex and less hardware-friendly activation functions. In this work, we propose a Hysteresis Rectified Linear Unit (HeLU), an efficient activation function designed to address the “dying ReLU” problem with minimal complexity. Unlike traditional activation functions with fixed thresholds for training and inference, HeLU employs a variable threshold that refines the backpropagation. This refined mechanism allows simpler activation functions to achieve competitive performance comparable to their more complex counterparts without introducing unnecessary complexity or requiring inductive biases. Empirical evaluations demonstrate that HeLU enhances model generalization across diverse datasets, offering a promising solution for efficient and effective inference suitable for a wide range of neural network architectures.


## 1 Introduction

The “dying ReLU” problem, as observed with the ReLU activation function, occurs because ReLU outputs zero for negative inputs. This can lead to neurons becoming inactive during training, thereby halting their learning capability [16]. This phenomenon significantly impacts the network’s ability to adapt and generalize effectively. Various alternatives to ReLU have been proposed [9, 10, 19, 25, 27], which allow small negative values to persist, and thus avoid “dying ReLU” during training. However, these modifications often come at the cost of less efficient inference time.

This inspired us to explore the use of robust mechanisms, with zero computational cost, by embracing Hysteresis. Hysteresis, derived from the Greek word meaning “lagging behind”, describes a fundamental property observed in various physical and engineering systems, including computer architecture. In computational contexts, hysteresis refers to the persistence of a state or behavior even after the input conditions that triggered it have changed or been removed. In computer architecture, hysteresis is prominently observed in memory elements such as flip-flops and in signal processing circuits. It plays a crucial role in ensuring proper timing and synchronization, preventing unintended state changes due to noise or transient signals. Engineers utilize hysteresis to design robust and dependable circuits capable of maintaining stable outputs amidst varying input conditions, thereby enhancing overall system performance and reliability.

---

\*Equal contribution.



(a) The threshold transition of bit voltage between 0 and 1 (forward) and 1 and 0 (backward) inspired the ResNet 40-4. In the image classification task with concept of using different thresholds for activation over parametric network, we see that majority of the functions in the forward and backward passes of the network. (b) Weight distribution of all pre-activations of Wide and 1 (forward) and 1 and 0 (backward) inspired the ResNet 40-4. In the image classification task with concept of using different thresholds for activation over parametric network, we see that majority of the functions in the forward and backward passes of the network. In red we mark the pre-activations that we allow regret with HeLU.

By introducing hysteresis, which adjusts the activation threshold, used during backpropagation, the proposed HeLU mitigates the “dying ReLU” problem without incurring additional computational costs during inference and only negligible costs during training. This approach enables neurons to remain active even for inputs that would otherwise result in zero activation in traditional ReLU networks. By maintaining non-zero gradients and promoting continuous learning in potentially dormant neurons, Hysteresis ReLU offers a promising direction for achieving attractive inference properties in lean networks with comparable performance. For instance, we observed a performance gain of +2.96 on CIFAR10, +2.19 on CIFAR100 and +1.23 over Imagenette in the computer vision domain, and +0.51 on the GLUE benchmark across 8 datasets, all compared to the simple ReLU.

## 2 Related Work

To increase throughput and reduce energy consumption in growing DNN models for deployment on edge devices, common methods include pruning and quantization. Pruning techniques remove less significant weights or features, requiring sophisticated re-routing in hardware to achieve actual benefits [6, 22]. Quantization, another prominent approach, reduces the precision of weights and activations to fixed precisions such as INT8 or INT4, or even binary representation [2, 26], and mixed-precision [17, 24, 12]. This reduces the memory footprint and enables faster matrix multiplications, which are the most common operations in DNNs.

We argue that for accelerated inference, focusing on the most fundamental building blocks, such as activation functions, can reveal further improvements in throughput. These improvements are particularly significant for more efficient models.

Traditional methods may overlook the computational efficiency of simpler components like activation functions, which, when optimized, can lead to substantial performance gains. Specifically, the ReLU activation function, due to its simplicity and computational efficiency, demonstrates the potential for enhanced throughput without the need for additional multiplications. Here we explore some alternative activation functions that does not benefit from this property.

### 2.1 Activation Functions

A deep neural network (DNN) is a parametric approximation of a function, constructed by repeatedly applying a linear transformation to the input followed by a non-linearity. These non-linear functions, known as activation functions, are crucial for enabling the network to learn complex mappings. Over time, the selection of activation functions has been influenced by both engineering and scientific considerations [7]. Enhancing the power of activation functions can improve the performance of models like transformers, often by incorporating additional operations such as products and linear projections, as seen in Gated Linear Units (GLUs) [27].

**Sigmoid** ( $f(x) = \sigma(x) = \frac{1}{1+e^{-x}}$ ), which was designed to mimic the firing rate of biological neurons. However, the Sigmoid function is susceptible to the “vanishing gradients” problem, where repeated

multiplications during backpropagation result in gradients that approach zero, hindering the training process [29].

**ReLU** ( $f(x) = \max\{0, x\}$ ) [20] activation function is the most prominent one in DNNs for many years, due to its empirical success. It is simple to implement in hardware, as signed representations simply look at the sign bit (MSB) and pass forward the number itself, or zero if negative. This means it does not require any intermediate calculations or additional parameters. In other words, the Arithmetic logic unit (ALU) of a processor does not need to execute any mathematical unit, unlike other cased activation functions. ReLU does not suffer from vanishing gradient, but from “dying ReLU” phenomenon, where a neuron might be deactivated during training producing a constant output of 0, and gradient updates would not change that.

In the recent years, **GELU** ( $f(x) = x\Phi(x)$ ) [10] where  $\Phi(x)$  is the Gaussian Error function, has proven great success in many transformer models and has become a standard to use. Typically used with approximation  $\frac{x}{2}[1 + \text{Tanh}(\sqrt{\frac{2}{\pi}}(x + 0.044715x^3))]$ , that even if faster, still require computation of the Tanh function. Other similar functions such as **Swish** ( $f(x) = x\sigma(x)$ ) [25], have become common in vision models. GELU is inspired by stochastic regularization and was formulated with a probabilistic approach of gating. It views the activation of a neuron and the following dropout as a single component, that is produced by the neuron. One thing to note about GELU is its curve, where its derivative is negative. This is in contrary to the former known activation functions, that are monotonic, and serves as a mean of regularization over the gradients, encouraging optimization to be more robust to sub-optimums, as local minima and saddle points. We refer the reader to LLMCompas [34], that analyze the hardware complexity and utilization of LLMs, showing that for GPT3, the GELU operations cause 400-500 G elements per second on NVIDIA A100. Moreover, the analysis shows that specifically for GPU usage, GELU can cause around 6% of the latency of GPT3, where quantized models become more popular, increasing this number significantly.

A recent trend of trainable activation functions gains popularity, with the notable Kolmogorov-Arnold Network (**KAN**) [15]. However empirical evident shown learnable activation functions are extremely inefficient, as they can take up to 10 times longer to train, and do not have any inference utilization grantees [7], they left out of the scope of this work.

Despite the advancements in activation functions, many of these alternatives do not benefit from the computational simplicity of ReLU, which merely requires a single-bit condition check and no multiplications. This makes ReLU particularly attractive for applications requiring fast and efficient inference, highlighting the trade-off between the advanced capabilities of newer functions and the hardware efficiency of ReLU, such as this work.

### 3 HeLU

Our formulation leverage the robustness of hysteresis, that preforms regularization by setting different threshold for the forward and backward steps on neurons of network, during training, as seen in Figure 2. HeLU is defined exactly like ReLU, for the forward pass, but for the backpropagation, it is shifted backwards by a hyperparameter,  $\alpha \in \mathbb{R}$ .

We define the Hysteresis Rectifier Linear Unit (HeLU) as:

**Definition 1 (HeLU)**  $HeLU_{\alpha}(x) = ReLU(x) = \max\{0, x\}$

Whereas, we override the autograd derivative that is used during backpropagation, and use instead the following term, lightly abusing the notation:

**Definition 2 (HeLU Modified Derivative)**  $\frac{d}{dx} HeLU_{\alpha}(x) =$

$$\begin{cases} 0 & x \leq -\alpha \\ 1 & -\alpha < x \end{cases}$$

Shifting back the derivative, essentially, refines the trigger of “dying ReLU”, by demanding a greater step to turn off the neuron. For convenience, we added a simple implementation of HeLU in Torch-like pseudo code in Algorithm 1.




Figure 2: ReLU, GELU and HeLU functions and their derivatives

---

**Algorithm 1** HeLU Function in PyTorch Style

---

```

1: class HeLUFunction(torch.autograd.Function)
2: @staticmethod
3: def forward(ctx, z, alpha)
4:   ReLU  $\leftarrow$  torch.where(z > 0, z, 0)
5:   ctx.save_for_backward(z)
6:   ctx.alpha  $\leftarrow$  alpha
7:   return ReLU
8: @staticmethod
9: def backward(ctx, grad_output)
10:  z,  $\leftarrow$  ctx.saved_tensors
11:  alpha  $\leftarrow$  ctx.alpha
12:  grad_positive  $\leftarrow$  torch.ones_like(z)
13:  grad_HeLU  $\leftarrow$  torch.where(z > -alpha, grad_positive, 0)
14:  return grad_HeLU * grad_output

```

---

## 4 Experiments

We evaluate the performance of HeLU along with common counterpart activations (GELU, ELU, and ReLU) across two prominent domains: computer vision and natural language processing, specifically for discriminative tasks. Our results indicate that HeLU performs comparably to GELU and outperforms ReLU while retaining the computational efficiency of ReLU. Consequently, our experiments emphasize the accelerated paradigm of the aforementioned tasks, as described in this section.

### 4.1 Image Classification

We conducted experiments on CIFAR10 and CIFAR100 datasets [13], as well as on a subset of ImageNet containing 10 common classes, known as Imagenette [11]. Consistent with [10], all image classification experiments employed the 40-4 Wide ResNet architecture [33]. For the CIFAR datasets, we trained for 100 epochs with a learning rate of 0.01, utilizing the momentum SGD optimizer with standard horizontal flip and crop augmentations. For Imagenette, we trained for 50 epochs with the same optimization algorithm but without augmentations. As shown in Table 1, HeLU demonstrates superior performance compared to the non-shifted threshold ReLU function, outperforming other methods or being second only to GELU in certain cases, such as CIFAR.

### 4.2 General Language Understanding Evaluation

#### Model Architecture and Training

Our experiments were conducted using *Cramming* [8], a highly optimized training framework of *BERT-base-uncased* [5] model, including many techniques to boost the training and the inference of

Table 1: Image classification Accuracy of Wide ResNet 40-4 with various activations.

Activation Function	CIFAR10	CIFAR100	Imagenette
ReLU	92.84 ± 0.34	75.31 ± 0.94	76.76 ± 1.50
ELU	93.59 ± 0.75	71.85 ± 1.77	76.67 ± 2.00
GELU	<b>96.11</b> ± 0.51	<b>79.26</b> ± 0.95	<b>78.45</b> ± 2.51
HeLU ( $\alpha = 0.001$ )	95.80 ± 0.63	77.51 ± 1.20	77.99 ± 1.21
HeLU ( $\alpha = 0.05$ )	95.31 ± 0.73	75.62 ± 1.47	76.67 ± 2.35
HeLU ( $\alpha = 0.01$ )	95.35 ± 0.94	75.84 ± 1.95	75.73 ± 2.85
HeLU ( $\alpha = 0.1$ )	95.10 ± 0.77	68.75 ± 1.45	72.32 ± 4.52

Table 2: Comparison of different activation functions on GLUE tasks (Larger is better).

Name	MNLI Accuracy (m)	MNLI Accuracy (mm)	QQP F1	QNLI Acc	SST-2 Acc	STS-B Spearman	MRPC F1	RTE Acc	WNLI Acc	Avg
GELU	81.9	82.52	<b>86.6</b>	<b>89.3</b>	<b>90.8</b>	<b>82.4</b>	82.0	53.1	56.3	<b>78.32</b>
ReLU	<b>82.7</b>	<b>82.9</b>	85.7	87.8	90.3	74.3	<b>83.2</b>	55.2	47.9	76.67
HeLU $\alpha = 0.05$	82.0	82.59	86.0	<b>89.3</b>	90.6	76.9	81.6	<b>56.3</b>	<b>57.7</b>	78.02
HeLU $\alpha = 1$	81.4	81.62	84.7	88.0	86.5	18.2	81.3	54.2	<b>57.7</b>	70.4
HeLU $\alpha = 2$	32.7	32.95	70.0	78.0	50.9	13.2	81.2	54.5	53.5	51.88

models, like automatic mixed precision (AMP) [18], flash attention [4], torch compile [1]. All the added optimizations are orthogonal to HeLU, and we believe one might like to use them in parallel with HeLU, to speed up inference. The training objective was also slightly modified, compared to the original *BERT*, including only masking loss, and not next sentence prediction, as shown in *RoBERTA* [14]. The pretraining phase took 3327000 steps, out of them, 1% for warmup.

Hyperparameters for BERT pretraining were selected by ablation and are specified in Table 4. At first, we performed a hyperparameter search on the GELU baseline, mainly to find the appropriate learning rate. After that, we used the selected hyperparameters to train another 3 HeLU models, where all GELU activations were replaced with HeLU with the following alphas: 0.05, 1, 2, and one ReLU model with ReLU activations. We used wikipedia-bookcorpus as **Pretraining Data**, similar to the original paper of *BERT*.

### Evaluation Benchmarks

After pretraining, we evaluated the models on 8 out of the 9 natural language understanding tasks included in GLUE [30]. Every model was fine-tuned for each of the tasks, using the *Cramming* code base. The results reflected from Table 2 shows an improvement using HeLU  $\alpha = 0.05$  over ReLU. Inline with the Vision results, GELU still achieves more accurate results on most datasets. Note that HeLU aim to bridge the gap between ReLU and GELU without memory and compute additional cost.

Another experiment conducted on QBERT [28], a post-training quantization of the BERT model without any post-training calibration, as shown in Table 3, highlights the gap in performance between QBERT and BERT. The results indicate that the tested activation functions suffer a small degradation in performance when quantized to 8-bit, whereas HeLU shows a mild improvement.

Table 3: Comparison of different activation functions on GLUE tasks, showing the difference between QBERT-INT8 and BERT performance (Larger is better).

Name	MNLI Accuracy (m)	MNLI Accuracy (mm)	QQP F1	QNLI Acc	SST-2 Acc	STS-B Spearman	MRPC F1	RTE Acc	WNLI Acc	Avg
ReLU	-0.4589	-0.2238	<b>0.0726</b>	<b>0.0915</b>	<b>0.4587</b>	0.0284	<b>0.1316</b>	0.3610	-2.8169	-0.2617
GELU	-1.4978	-2.0952	-0.1234	0.0732	0.3440	-0.0123	0.0535	<b>1.0830</b>	0.0000	-0.2417
HeLU $\alpha = 0.05$	<b>0.1822</b>	<b>1.6273</b>	-0.1400	-0.0549	0.3440	<b>0.2296</b>	0.1251	0.7220	<b>4.2254</b>	<b>0.8067</b>

### Training Environment

We used  $3 \times$  NVIDIA GEFORCE RTX 2080 Ti. Each GPU has memory capacity of close to 12GB. Each pretraining took us approximately 1 week. The seed for the pretraining, fine-tuning and evaluation of all tested models is 42.

The pretraining of the GELU models was using CUDAGraphs [21], an optimization mechanism that allows a series of CUDA kernels to be executed as a single unit, reducing the launching overhead and achieving speedup, provided by the training framework. However, we had to omit the optimization when pretraining HeLU and ReLU models due to a lack of memory required for torch compilation. Thus, we consider a further optimized version of the GELU model in our comparison.




Figure 3: Comparing HeLU and GELU in BERT. Top figures are Inference time by dataset, using half precision and QINT8 versions. Bottom figures are Throughput using half precision and QINT8.

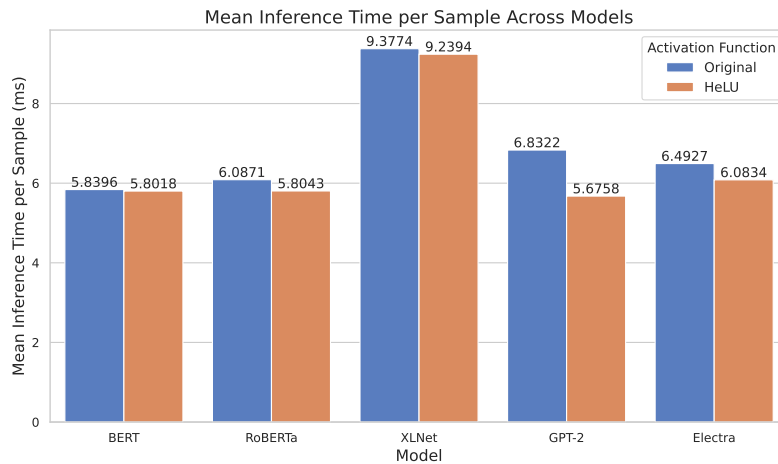


Figure 4: Analysis of various transformer architectures inference time per sample, using their original activation and HeLU.

### 4.3 Computational Analysis

We assess the inference time and throughput of BERT with and without QINT8 quantization (using only AMP for the latter) by running quantized models on a CPU and non-quantized models on an NVIDIA L40 GPU. The evaluation iterates over batches of 256 drawn from a subset of up to 1024 randomly selected samples from the validation set of each GLUE benchmark dataset, with measurements repeated 10 times. The results of this experiment are illustrated in Figure 3. Our findings reveal a notable trend: HeLU achieves faster inference times, particularly in the non-quantized model, and higher throughput, especially in the quantized model, where memory constraints are more pronounced. The reported standard deviation is calculated based on batch-wise processing rather than individual samples, reflecting the variability introduced by the random selection of samples from the validation set. Consequently, datasets with considerable variation in sequence length may exhibit high variance, though this pattern remains consistent across different non-linearity functions.

Another experiment presented in Figure 4, shows inference time per sample using various transformer architectures, namely BERT [5], RoBERTa [14], XLNet [32], GPT-2 [23] and Electra [3] on single NVIDIA T4 GPU. We denote the original implementation from [31] as Original, since some models like GPT-2 used a modified version of GELU. In all models we observed a reduction in inference time, where the experiment averaged over 100,000 examples. The total timing of the examples benefit range from 121.63 seconds (GPT-2) to only 4.735 seconds (BERT). While BERT seem insignificant, we show in Figure 3 that in some cases, like the CoLA or STSB datasets it can be very significant.

Hyperparameter	Value
Batch Size	288
Gradient Clipping	0.5
Optimizer	AdamW
Learning Rate	0.0001
Betas	[0.9, 0.98]
Epsilon	1E-12
Weight Decay	0.01
Total Steps	3,327,000
Warmup Steps	33,270
Scheduler	Budget-One-Cycle

Table 4: Pretraining Hyperparameters of BERT

## 5 Conclusion

In this study, we advocate for using shifted activation threshold during training, HeLU, a novel hysteresis activation function designed to enhance the robustness of activations against the “dying ReLU” phenomenon. By leveraging the gradients of a shifted activation function, our approach maintains simplicity while improving performance across various domains. Our empirical results indicate that a relatively straightforward modification can yield significant performance gains, highlighting the potential of HeLU for universal application. This study paves the way for exploring efficient activation functions that are both computationally and environmentally sustainable. While extremely efficient, one drawback of HeLU is the search for the optimal hyper-parameter; Future research should further investigate the pre-activation statistics and find an optimal value for  $\alpha$ . It is important to note that HeLU is intended to enhance efficiency rather than replace activation functions in systems where accuracy is paramount.

## References

[1] J. Ansel, E. Yang, H. He, N. Gimelshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski, G. Chauhan, A. Chourdia, W. Constable, A. Desmaison, Z. DeVito, E. Ellison, W. Feng, J. Gong, M. Gschwind, B. Hirsh, S. Huang, K. Kalambarkar, L. Kirsch, M. Lazos, M. Lezcano, Y. Liang, J. Liang, Y. Lu, C. K. Luk, B. Maher, Y. Pan, C. Puhersch, M. Reso, M.-A. Saroufim, M. Y. Siraichi, H. Suk, S. Zhang, M. Suo, P. Tillet, X. Zhao, E. Wang, K. Zhou,

- R. Zou, X. Wang, A. Mathews, W. Wen, G. Chanan, P. Wu, and S. Chintala. Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 2024.
- [2] R. Banner, Y. Nahshan, E. Hoffer, and D. Soudry. Post-training 4-bit quantization of convolution networks for rapid-deployment, 2019.
- [3] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning. Electra: Pre-training text encoders as discriminators rather than generators, 2020.
- [4] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness, 2022.
- [5] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.
- [6] E. Diao, G. Wang, J. Zhan, Y. Yang, J. Ding, and V. Tarokh. Pruning deep neural networks from a sparsity perspective, 2023.
- [7] S. R. Dubey, S. K. Singh, and B. B. Chaudhuri. Activation functions in deep learning: A comprehensive survey and benchmark, 2022.
- [8] J. Geiping and T. Goldstein. Cramming: Training a language model on a single gpu in one day, 2022.
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, 2015.
- [10] D. Hendrycks and K. Gimpel. Gaussian error linear units (gelus), 2023.
- [11] J. Howard. Imagemang, 2019.
- [12] M. Kimhi, T. Rozen, A. Mendelson, and C. Baskin. Amed: Automatic mixed-precision quantization for edge devices. *Mathematics*, 12(12), 2024.
- [13] A. Krizhevsky. Learning multiple layers of features from tiny images. Master’s thesis, University of Toronto, 2009.
- [14] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019.
- [15] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark. Kan: Kolmogorov-arnold networks, 2024.
- [16] L. Lu, Y. Su, and G. E. Karniadakis. Collapse of deep and narrow neural nets, 2018.
- [17] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu. Mixed precision training, 2018.
- [18] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu. Mixed precision training, 2018.
- [19] D. Misra. Mish: A self regularized non-monotonic activation function, 2019.
- [20] V. Nair and G. Hinton. Rectified linear units improve restricted boltzmann machines vinod nair, 06 2010.
- [21] V. Nguyen, M. Carilli, S. B. Eryilmaz, V. Singh, M. Lin, N. Gimelshein, A. Desmaison, and E. Yang. Cudagraphs, 2024.
- [22] M. M. Pasandi, M. Hajabdollahi, N. Karimi, and S. Samavi. Modeling of pruning techniques for deep neural networks simplification, 2020.
- [23] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.



- [24] M. Rakka, M. E. Fouda, P. Khargonekar, and F. Kurdahi. Mixed-precision neural networks: A survey, 2022.
- [25] P. Ramachandran, B. Zoph, and Q. V. Le. Searching for activation functions, 2017.
- [26] T. Rozen, M. Kimhi, B. Chmiel, A. Mendelson, and C. Baskin. Bimodal-distributed binarized neural networks. *Mathematics*, 10(21), 2022.
- [27] N. Shazeer. Glu variants improve transformer, 2020.
- [28] S. Shen, Z. Dong, J. Ye, L. Ma, Z. Yao, A. Gholami, M. W. Mahoney, and K. Keutzer. Q-bert: Hessian based ultra low precision quantization of bert, 2019.
- [29] L. Ven and J. Lederer. Regularization and reparameterization avoid vanishing gradients in sigmoid-type networks, 2021.
- [30] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding, 2019.
- [31] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. 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, Oct. 2020. Association for Computational Linguistics.
- [32] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. Salakhutdinov, and Q. V. Le. Xlnet: Generalized autoregressive pretraining for language understanding, 2019.
- [33] S. Zagoruyko and N. Komodakis. Wide residual networks, 2016.
- [34] H. Zhang, A. Ning, R. Prabhakar, and D. Wentzlaff. A hardware evaluation framework for large language model inference, 2023.