003 004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

032

040

INVERTED ACTIVATIONS: REDUCING MEMORY FOOT-PRINT IN NEURAL NETWORK TRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

The scaling of neural networks with increasing data and model sizes necessitates the development of more efficient deep learning algorithms. A significant challenge in neural network training is the memory footprint associated with activation tensors, particularly in pointwise nonlinearity layers that traditionally save the entire input tensor for the backward pass, leading to substantial memory consumption. In this paper, we propose a modification to the handling of activation tensors in pointwise nonlinearity layers. Our method involves saving the output tensor instead of the input tensor during the forward pass. Since the subsequent layer typically also saves its input tensor, this approach reduces the total memory required by storing only one tensor between layers instead of two. This optimization is especially beneficial for transformer-based architectures like GPT, BERT, Mistral, and Llama. To enable this approach, we utilize the inverse function of the nonlinearity during the backward pass. As the inverse cannot be computed analytically for most nonlinearities, we construct accurate approximations using simpler functions. Experimental results demonstrate that our method significantly reduces memory usage without affecting training accuracy or computational performance. Our implementation is provided as a drop-in replacement for standard nonlinearity layers in the PyTorch framework, facilitating easy adoption without requiring architectural modifications. The code is available at https://github.com/removed/for/anonimity.

1 INTRODUCTION

The remarkable scaling of neural networks with
increasing data and model sizes has created a
continuous demand for training and deploying
larger models more efficiently. This necessity
drives the ongoing development of deep learning algorithms that are both computation and
memory-efficient.

A significant portion of memory consumption 041 during neural network training arises from ac-042 tivation tensors, alongside model parameters 043 and optimizer statistics. Activation tensors 044 are stored during the forward pass to facilitate gradient computations in the backward pass. 046 Even simple and computationally fast opera-047 tions, such as pointwise nonlinearities, con-048 tribute substantially to this memory footprint by storing entire input tensors. In most popular 050 deep learning frameworks, pointwise nonlinearities like GELU Hendrycks & Gimpel (2016) 051 and SiLU Elfwing et al. (2018) save the whole 052

Listing 1: Pseudocode of hypothetical use of our drop-in replacement for GELU that will use less memory without any speed or quality degradation. Exact memory gains for several popular huggingface¹ models are presented in the following table. Exact model setups are given in appendix A.1.

Model	Memory Saving
BERT Audio Spectral Transformer ViT CLIP	$-22.9\% \\ -24.0\% \\ -23.8\% \\ -23.4\%$

model = create_model_as_usual()
for layer in model.layers:
 layer.act_fn = InvActGELU()
train_as_usual(model)

⁰⁵³

¹https://huggingface.co/models

input tensor for the backward pass. This practice leads to considerable memory usage, especially in transformer-based architectures where such activations are prevalent. 056

In this work, we propose a modification to the forward and backward passes of element-wise nonlin-057 earity layers. Specifically, we suggest saving the output tensor instead of the input tensor during 058 the forward pass. Since the subsequent layer in the computation graph typically also saves its input 059 tensor, our approach ensures that only one tensor is stored between layers rather than two, effectively 060 reducing the memory footprint by nearly 25% in practice. 061

- For a novel approach in implementing nonlinearity functions to gain widespread adoption, it should 062 ideally meet three key conditions: 063
 - 1. The method should not introduce any additional computational error, ensuring that model performance remains unaffected.
 - 2. It should not slow down the model, maintaining or improving the speed of forward and backward computations.
 - 3. The approach should be user-friendly, meaning it can serve as a straightforward, drop-in replacement for existing nonlinearity layers without requiring modifications to the architecture or specific tuning of adjacent layers.

072 For example, highly optimized building blocks like Triton NVIDIA Corporation or CUDA-based 073 GeGLU implementations for models such as Llama and Mistral, or efficiently fused MLP blocks for 074 transformer architectures, can offer significant speed and/or memory advantages, but are applicable to a very narrow range of use cases. In contrast, for general research purposes and R&D, truly 075 versatile layers are highly desirable. In this paper, we demonstrate that our proposed method meets 076 all three of these criteria, making it a promising candidate for broader adoption in deep learning 077 systems. 078

079 We implemented this approach as a drop-in replacement for standard nonlinearity layers within the 080 PyTorch Ansel et al. (2024) framework, ensuring seamless integration for researchers and practitioners without the need for architectural modifications. In Listing 1, we illustrate the straightforward 081 integration of our nonlinearity layers into existing pipelines, which requires only the substitution of standard nonlinearities with our InvAct (Inverted Activation) layers, without the need for additional 083 modifications to the architecture or hyperparameters. 084

METHOD 2

Background Knowledge. The computation graph during neural network training can be expressed as a series of operations f_i , each of which transforms an input tensor $\mathbf{X}^{(i)}$ into an output tensor $\mathbf{X}^{(i+1)}$. This output tensor subsequently serves as the input for the next operation in the computation graph:

 $\ldots \to \mathbf{X}^{(i)} \xrightarrow{f_i} \mathbf{X}^{(i+1)} \to \ldots \to L.$

Here, L is the final loss that is optimized using gradient descent. Each operation stores some addi-094 tional information $\mathbf{S}^{(i)}$, referred to as activation tensors, for the backward pass. During the backward 095 pass, these activation tensors are used to compute the gradient of the loss L with respect to all inter-096 mediate tensors of the computation:

$$\ldots \leftarrow \frac{\partial L}{\partial \mathbf{X}^{(i)}} \xleftarrow{\text{gradient } f_i, \mathbf{S}^{(i)}}{\frac{\partial L}{\partial \mathbf{X}^{(i+1)}}} \leftarrow \frac{\partial L}{\partial L}.$$

100 For pointwise nonlinearity layers, the forward computation f_i is simply the element-wise application of a function $f : \mathbb{R} \to \mathbb{R}$ to the input tensor $\mathbf{X}^{(i)}$: 102

$$\mathbf{X}_{j_{1},\dots,j_{k}}^{(i+1)} = f\left(\mathbf{X}_{j_{1},\dots,j_{k}}^{(i)}\right),\tag{1}$$

while the backward computation is the Hadamard product of the gradient with respect to the output 105 tensor and the derivative of f evaluated at the input tensor: 106

107

064

065

066

067

068

069

070

071

085

086

087 880

089

090

091

092

098 099

101

$$\frac{\partial L}{\partial \mathbf{X}^{(i)}} = \frac{\partial L}{\partial \mathbf{X}^{(i+1)}} \odot f'\left(\mathbf{X}^{(i)}\right).$$
(2)

108 109 110 1 def forward(X): 2 S = X < T111 Y = f(X)112 save_for_backward(Y, S) 4 113 return Y 5 114 6 115 7 def backward(dY): 8 Y, S = saved_for_backward() 116 9 return dY * f'($f^{-1}(Y, S)$) 117

Listing 2: Pseudocode of inverted activation approach for forward and backward passes of pointwise nonlinearity layer. Boolean array S should be saved in an optimal way (1 bit per element) to achieve memory footprint reduction.



Figure 1: Two monotonous parts of the GELU or SiLU nonlinearity function, each being an invertible function.

Here, $f'(\mathbf{X}^{(i)})$ is computed element-wise, similarly to Eq. (1).

Since both Eq. (1) and Eq. (2) are applied independently to each element of the corresponding tensors, we can simplify the notation by letting $x \in \mathbb{R}$ represent an input element of the nonlinearity and $y \in \mathbb{R}$ represent the corresponding output:

$$y = f(x), \quad \frac{\partial L}{\partial x} = \frac{\partial L}{\partial y} f'(x).$$
 (3)

134 135 It is evident that to perform the backward computation in Eq. (2) efficiently, one can store the entire 136 input tensor $\mathbf{X}^{(i)}$ as an activation tensor $\mathbf{S}^{(i)}$. This approach is adopted by most popular deep 137 learning frameworks, such as PyTorch and Jax, for nonlinearity layers like GELU and SiLU, which 138 are widely used in transformer-based neural networks.

Inverted Activation. In this work, we propose modifying the forward and backward computations to save the output tensor $\mathbf{X}^{(i+1)}$ instead of the input tensor $\mathbf{X}^{(i)}$. If the subsequent layer after the pointwise nonlinearity also saves its input tensor, the two layers together will only store one tensor instead of two, thereby significantly reducing the overall memory footprint.

143 The backward computation in Eq. (3) can be rewritten as:

$$\frac{\partial L}{\partial x} = \frac{\partial L}{\partial y} f'(x) = \frac{\partial L}{\partial y} f'\left(f^{-1}(y)\right).$$

Here, f^{-1} is the inverse function of the nonlinearity f, which gives our method its name: **Inverted** Activations (or **InvAct** for short).

Space-efficient Boolean Indicator. Both GELU and SiLU are non-invertible functions, but they consist of two monotonic parts, each of which is individually invertible. Thus, we can save additional information during the forward pass – a Boolean indicator that specifies whether each element of the input tensor belongs to the first monotonic part or the second part, which helps disambiguate f^{-1} :

144 145 146

123

124 125 126

127 128

129

130 131 132

133

$$= \begin{cases} 1 & \text{if } x < T \\ 0 & \text{otherwise} \end{cases}$$
(4)

Here, T is the point of separation between the two monotonic halves (see Fig. 1 for a graphical explanation). It is possible to store the tensor **S** in a space-efficient Boolean array, where each element requires only 1 bit of memory. In our implementation, we achieve this by storing the Boolean array **S** in a tensor **S**_{compressed} with the data type 'uint8', in which **S**[*i*] corresponds to the (*i* mod 8)-th bit of the (*i* ÷ 8)-th element:

s

$$\mathbf{S}[i] = \mathbf{S}_{\text{compressed}} [i \div 8] \gg (i \mod 8)$$

180

181

182

183 184

211

162

163



Figure 2: Derivative of the nonlinearity and its approximation in two coordinate systems: plotted as a function of the output variable y (left two columns) and the input variable x (right column). The y-based representation highlights how equations Eqs. (5) to (8) approximate the respective regions of $f'(f^{-1}(y))$. The x-based representation, in turn, helps assess the approximation quality of the layer based on how it is used in a neural network context.

Here, \gg denotes the bitwise right shift, and & denotes the bitwise AND operation. Pseudocode for the resulting pointwise nonlinearity layer can be found in Section 2.1.

Approximation. Both GELU and SiLU are complex functions, and to the best of our knowledge, neither their inverses f^{-1} nor $f' \circ f^{-1}$ can be analytically derived. Therefore, we propose to use approximations based on simple base functions, similar to the fast GELU approximation Hendrycks & Gimpel (2016).

There are many decisions to be made when constructing the approximation for $f' \circ f^{-1}$: whether to approximate only f^{-1} and evaluate f' in the standard way, or to approximate $f' \circ f^{-1}$ as a whole; which primitive functions to use; what form the approximation should take; and how to balance complexity and accuracy. In this work, we chose to approximate $f' \circ f^{-1}$ as a whole. The resulting approximation is highly accurate – so much so that the training process using our approximation is indistinguishable from training with the original nonlinearity layer, as we will demonstrate in our experiments in Section 3.

For GELU we use the following approximations:

$$x < T \Rightarrow f'\left(f^{-1}(y)\right) \approx q^{\text{left}}(y) = c_0\sqrt{y+c_1} * (2y+c_2\sqrt{-y})(|c_3y^2+|c_4y+c_5|+c_6|+c_7),$$
(5)

$$x \ge T \Rightarrow f'\left(f^{-1}(y)\right) \approx q^{\operatorname{right}}(y) = 1 + \left(c_0 + c_1\sqrt{\tilde{y}} + c_2\tilde{y}\right)e^{c_3(c_4 - \tilde{y})^3}, \text{ where } \tilde{y} = y - f(T).$$
(6)

For SiLU we use the following approximations:

$$x < T \Rightarrow f'\left(f^{-1}(y)\right) \approx q^{\text{left}}(y) = \left(c_0 + c_1\sqrt{\tilde{y}} + c_2\tilde{y} + c_3\tilde{y}^2\right) * (1-y) + y,\tag{7}$$

$$x \ge T \Rightarrow f'(f^{-1}(y)) \approx q^{\text{right}}(y) = \left(1 + \left(c_0 + c_1\sqrt{\tilde{y}} + c_2\tilde{y}\right)e^{c_3(c_4 - \tilde{y})^3}\right) * (1 - y) + y, \quad (8)$$

where
$$\tilde{y} = y - f(T)$$
 (9)

The coefficients c_i are precomputed and fine-tuned specifically for each function (GELU or SiLU) and each approximation region (x < T or $x \ge T$). All coefficient values used in this work are provided in appendix A.2. In Fig. 2, we illustrate the functions $f'(f^{-1}(y))$ for x < T and $x \ge T$, along with their respective approximations q^{left} and q^{right} . The plots also show the approximation errors, defined as $f'(f^{-1}(y)) - q^{\text{left}}(y)$ for x < T and $f'(f^{-1}(y)) - q^{\text{right}}(y)$ for $x \ge T$. It is important to note that the presented approximations are not necessarily the most optimal or minimal representations of $f'(f^{-1}(y))$. As we are not experts in constructing such approximations, it is possible that simpler, more accurate, and computationally efficient alternatives exist, which could further enhance the performance of InvAct. Nonetheless, our current approximations are wellsuited for the specific task at hand, as demonstrated by their effectiveness, accuracy and speed in the experiments detailed in Section 3.

222

224

230

2.1 Alternatives to Bit-Compressed Boolean Indicator

An essential component for inverting nonlinearities such as GELU or SiLU is an indicator function that identifies which monotonic part the original x belongs to. Currently, we store this information in a space-efficient, bit-compressed Boolean tensor, where each Boolean value occupies only 1 bit of memory. In this section, we will discuss two alternative methods for storing this information without requiring even 1 bit per element.

Sign-bit Inverted Activation. Both GELU and SiLU functions are bounded, which means there exists some constant C (see Fig. 1) such that for all x, $f(x) \ge C$, where f represents either the GELU or SiLU function. This implies that the value f(x) - C is always greater than zero. Instead of storing f(x) for the backward pass, we can store f(x) - C with its sign bit modified according to the indicator function s (see Eq. (4)).

For most popular floating-point formats used in deep learning, such as float32, float16, and bfloat16, the sign bit is the last bit of a floating-point number. This approach allows us to eliminate the need for additional memory to store the 1 bit per element for the indicator function, making the entire nonlinearity layer memory-free in terms of footprint. The modified value y = f(x) does not pose an issue for subsequent computations, as C is a predefined global constant for a given activation function. For example, the computation for an activation function followed by a fully-connected linear layer can be expressed as follows:

```
1 N # number of bits in floating point data type
243
     2 C # minimum of activation function f
244
     3 def forward(X, W_linear):
245
         S = X < T
     4
246
         Y = f(X) - C
     5
247
     6
        Y = (Y \& (~(1 << N))) | (S << N) \# sets sign bit of tensor Y
                                              # to indicator function S
248
     7
         save_for_backward(Y)
     8
249
     9
250
         return (Y + C) @ W_linear
    10
251
    11
    12 def backward(dY, W_linear):
        Y = saved_for_backward()
253
    13
         S = Y \& (1 << N) # extracts indicator function from last bit of tensor Y
    14
254
    15
        Y = Y^{A}S
                          # reverts sign bit of tensor Y back to zero
       dW = (Y + C) \cdot T @ dY
    16
256
         dX = dY * f'(f^{-1}(Y + C, S))
    17
257
       return dY, dW
    18
```

258 259

Here, the forward pass takes the input tensor X and the matrix W_{linear} of the fully-connected linear layer and returns $f(X) \cdot W_{\text{linear}}$.

Unfortunately, to the best of our knowledge, it is not possible to implement such a layer in PyTorch that works independently of what comes after the nonlinearity function, meaning that you would have to write a custom implementation for each new architectural block. As a result, the sign-bit approach does not satisfy condition 3 outlined in the Introduction (see Section 1), and thus we do not consider it a viable candidate for broad adoption.

Precision Bit Inverted Activation. The second approach to store the indicator function s (see Eq. (4)) is to use the lowest precision bit of a floating-point number (for most popular floating-point data types, this is the first bit in a base-2 representation). During the forward pass, we compute y = f(x) and modify this bit of y to be equal to s (see Eq. (4)).

270 Compared to the previous approach, this method does not violate condition 3, as it allows for a conve-271 nient drop-in replacement for classical activation functions on all modern deep learning frameworks. 272 However, it is more challenging to evaluate the accuracy of this approach. Both the standard inverted 273 activation and the sign-bit inverted activation perform the forward pass exactly and introduce only a 274 minimal approximation error during the backward pass. This method, on the other hand, sacrifices one precision bit for the pointwise nonlinearity computation during the forward pass. Moreover, 275 with the ongoing trend of reducing floating-point precision, loss of accuracy may become larger in 276 the future. For this reason, we do not consider it a viable candidate for broad adoption at this time, as it requires extensive experimentation to validate its accuracy. However, we still regard it as an 278 interesting approach that may find applications in the future. 279

This is a pseudocode for the forward and backward passes of this approach:

289

290

291

294 295

296

297

298

299

300 301

302

303

304 305

306 307

310

311

312

313

314

315

316

317 318

319

320

321

280

1

2

3

4

5

6

7

8

def forward(X): S = X < TY = f(X)Y = (Y & (~1)) | S save_for_backward(Y) return Y 9 def backward(dY): Y = saved_for_backward() 10 $S = Y \& 1 \ \# \ \text{extracts}$ indicator fucntion from last bit of tensor Y 11 12 return dY \star f'(f⁻¹(Y, S))

292 293

3 **EXPERIMENTS**

The validation of our method's efficiency is two-fold:

- 1. We compare the computational efficiency of the forward and backward passes of our inverted activation nonlinearity layer in various settings.
- 2. We demonstrate that the approximation of the nonlinearity gradient does not affect the quality of gradient descent during training.

With these results, we can confidently assert that inverted nonlinearity layers can be safely used in practice to reduce the memory required for training neural networks without any speed or accuracy degradation. All tests and measurements were performed on a single NVIDIA A100 GPU.

3.1 COMPUTATIONAL EFFICIENCY

In this section, we compare the computational efficiency of PyTorch Ansel et al. (2024) nonlinearity 308 layers with our Triton-based NVIDIA Corporation implementations of inverted activation nonlinearities. We conduct three groups of tests (with detailed setups presented in appendix A.3): 309

- 1. Application of nonlinearity to a tensor.
 - 2. Application of several layers, one of which is a nonlinearity layer. For this purpose, we selected several popular building blocks of modern neural networks:
 - Linear layer + nonlinearity
 - Transformer MLP block: Linear layer + nonlinearity + Linear layer
 - GeGLU Shazeer (2020): Hadamard product of a linear layer + nonlinearity with another linear layer
 - 3. Full neural network training iterations for the following popular network architectures:
 - BERT
 - Llama v3.1 8B

The corresponding measurements are presented in table 1. The results indicate that the difference 322 in computational efficiency between our inverted activation nonlinearity layer and the standard non-323 linearity layer within various architectural blocks (such as MLP, Linear + activation, or GeGLU) is

Precision Bit InvAct (ms) Torch (ms) InvAct (ms) Plain GELU 0.820.86 (+5.75%)0.86(+4.96%)Linear + GELU 11.78 11.80 (+0.13%)11.79 (+0.06%)Tranformer MLP Block 75.4175.46 (+0.06%)75.41(-0.00%)GeGLU 94.08 94.13 (+0.05%)94.09(+0.00%)BERT 1112.12 1112.62 (+0.05%)1112.25 (+0.01%)Llama3.1-8b 182.84 182.94 (+0.06%)182.80(-0.02%)

less than 1%. Furthermore, for full-scale neural networks like BERT and Llama, this difference is
 even more negligible, being less than 0.1%.

Table 1: Comparison of computation time (in milliseconds) for different activation function implementatoins and neural network architectures. We compare following architecture blocks: plain nonlinearity layer, fully-connected layer with nonlinearity, Transformer MLP block, GeGLU Shazeer (2020), BERT Devlin et al. (2019) model, Llama3.1-8b model. For each architecture block we compare following activation functions: torch GELU, our inverted activation and precision-bit inverted activation. We show difference in percents w.r.t. torch nonlinearity. Our layer is lower then 1% slower compared to torch nonlinearity inside some computation block, and lower then 0.1% slower inside BERT and Llama.

3.2 EFFECT OF THE APPROXIMATION ERROR



Figure 3: Comparison of BERT model fine-tuning with original and inverted GELU activations. The first row (except the rightmost column) compares two runs with Torch GELU and inverted GELU using identical random seeds. The second row shows losses averaged across 16 runs for each activation function. The leftmost column displays the training loss during the first 100 iterations. The middle column shows the training loss during the last 100 iterations out of a total of 2000. The rightmost column compares the validation losses. The observed losses, both training and validation, are nearly identical, with any discrepancies attributable to minor approximation inaccuracies that are absolutely insignificant and negligible for practical purposes.

367

327

328

330

331

332

333

334

335 336

337

338

339

340

341

342

343 344 345

346 347

348

349

350

351

352

353

354

355

356

357

358

In this section, we investigate the impact of our approximation method on the training process and final model performance. To evaluate this, we conducted experiments using two popular neural network architectures: BERT and Llama. We trained these models on standard benchmark tasks using both the original PyTorch GELU activation function and our approximated inverted GELU function for BERT, as well as the PyTorch SiLU activation function and our approximated inverted SiLU function for Llama.

For BERT, we fine-tuned the model on the Yelp Review dataset Yelp (2020) for sentiment classification. We compared the training process of models using both the original GELU activation function and our inverted GELU activation. Comparisons of the training and validation losses are presented in Fig. 3. We observe that the approximation inaccuracy of the inverted activation does not affect BERT's training quality, as both training and validation losses are nearly indistinguishable over all



Figure 4: Same as Fig. 3, but for the Llama v3.1 8B model with SiLU nonlinearity. The first row (except the rightmost column) compares two runs with Torch SiLU and inverted SiLU using identical random seeds. The second row shows losses averaged across 16 runs for each activation function. The leftmost column displays the training loss during the first 100 iterations. The middle column shows the training loss during the last 100 iterations out of a total of 2000. The rightmost column compares the validation losses. The observed losses, both training and validation, are nearly identical, with any discrepancies attributable to minor approximation inaccuracies that are absolutely insignificant and negligible for practical purposes.

390

391

392

393

394

395

396

2000 steps of the training process. The variance produced by different random seeds is significantly
 larger than any difference between the two activation functions, suggesting that any discrepancies
 are negligible and on par with machine precision.

402 We extended this experiment to a significantly larger model, Llama v3.1 8B Dubey et al. (2024), 403 which uses a different activation function – SiLU. We fine-tuned a 4-bit quantized version with 404 LoRA Hu et al. (2022) using ORPO on a combination of DPO datasets². The results for Llama 405 training, presented in Fig. 4, show an even closer match between the original SiLU activation and 406 our inverted SiLU activation compared to the BERT experiment. The training and validation losses 407 are virtually indistinguishable throughout the entire training process. This remarkable similarity further supports our conclusion that the approximation inaccuracy of the inverted activation has a 408 negligible impact on training quality and final model performance. 409

Based on this comparison of model training using original and inverted activations, we can conclude
that the approximation inaccuracy of the InvAct does not affect training quality in any way. Combined with the speed measurements presented in Section 3.1, we can state that the inverted activation
nonlinearity layer is a perfect drop-in replacement for GELU and SiLU activation functions.

Furthermore, the seamless integration of our method into existing architectures – without requiring any modifications to pipelines, code, or hyperparameters – highlights its potential for widespread adoption in various deep learning applications. The ability to reduce memory requirements without compromising model accuracy or training dynamics makes our inverted activation approach a promising tool for researchers and practitioners working with memory-intensive models, especially in resource-constrained environments.

420 421

422

3.3 COMPARISON WITH FEWBIT

Another method that reduces the memory required by pointwise nonlinearities is FewBit Novikov 423 et al. (2023), a gradient quantization method that saves a quantized version of the nonlinearity gra-424 dient. FewBit with one bit per element suffers from a considerable accuracy drop, as demonstrated 425 in Fig. 5. The authors report that using 4-bit quantization, which requires four times more memory 426 compared to our standard InvAct implementation, FewBit does not show any performance degrada-427 tion. However, even in that setting, our approach is superior: the approximation of $f'(f^{-1}(x))$ is 428 much more accurate than FewBit quantization up to 8 bits! We show L_2 and L_{∞} approximation 429 errors for FewBit and InvAct in Fig. 5, and our approach demonstrates significantly better approxi-430 mation quality, even when compared to FewBit using 8 times more bits per element.

²https://huggingface.co/datasets/mlabonne/orpo-dpo-mix-40k



Figure 5: Comparison of InvAct (our) nonlinearity layer with the FewBit Novikov et al. (2023) nonlinearity layer. The two left plots show the validation loss for fine-tuning BERT and Llama v3.1 8B, respectively. The BERT model uses the GELU nonlinearity, while the Llama model uses the SiLU nonlinearity. FewBit is used with a 1-bit per element memory budget to match the memory footprint of InvAct. The FewBit approach exhibits significantly worse validation loss, while our InvAct approach is indistinguishable from the standard Torch nonlinearity. The right column shows the L_2 approximation error (top plot) and L_{∞} approximation error (bottom plot) for FewBit GELU with respect to the number of bits used in gradient quantization (blue line), compared to the approximation quality of InvAct (red dashed line). Our approach always uses only 1 bit per element. InvAct demonstrates superior approximation quality, even when compared to 8-bit FewBit GELU (y-axis is in logarithmic scale).

457 458

460

461

444

445

446

447

448

449

450

451

452

4 **RELATED WORK**

The efficient scaling of neural, has necessitated innovations in low-precision training, quantization 459 techniques, and various compression methods. These approaches address computational and memory bottlenecks, enabling the training of ever-larger models on hardware-limited resources.

Quantization methods aim to compress neural network weights and activations by representing 462 them in lower-precision formats, such as 4-bit integers. One can quantize model weights Jacob et al. 463 (2018); Zafrir et al. (2019); Gholami et al. (2021) or activations that are saved for backward Chen 464 et al. (2021); Novikov et al. (2023). Former are used more in during inference, but approaches 465 to utilize it during training are also exist Liu et al. (2024); Shao et al. (2024); Xu et al. (2024). 466 These methods are especially popular for the fine-tuning stage of training, where reducing memory 467 usage without extensive retraining is crucial. While quantization effectively decreases the memory 468 footprint, it can introduce computational inaccuracies. As a result, these errors lead to degradation 469 in model performance, making the balance between memory savings and maintaining model quality 470 a key consideration when applying quantization.

471 Low-precision training techniques have emerged as another powerful approach to reducing both 472 the memory and computational requirements of large-scale models Micikevicius et al. (2018). One 473 widely adopted format is bfloat16 Kalamkar et al. (2019), which retains a wider range of values 474 compared to traditional FP16 while providing faster computation and reduced memory consump-475 tion. More recently work towards even lower precision training in float8 formats have also been 476 explored Wang et al. (2018) for their potential to further reduce memory and computational de-477 mands. These formats are particularly attractive for their extreme compression capability, though their usage requires careful consideration to avoid significant accuracy loss. Moreover, widespread 478 adoption of these lower precision formats relies heavily on GPU manufacturers integrating support 479 for them directly into hardware. Until this support is in place, these formats remain experimental 480 and difficult to implement at scale. 481

482 Custom CUDA and Triton implementations Unsloth (2023); Hsu et al. (2024) have become 483 essential for optimizing critical building blocks in popular architectures, significantly improving memory efficiency and computational speed. For instance, highly fused implementations of MLP 484 blocks Müller et al. (2021) or optimized GeGLU layers are widely used to reduce the overhead of 485 separate operations by merging them into a single, streamlined kernel. Another prominent example

486 is FlashAttention Dao (2024), which significantly improves the efficiency of attention mechanisms 487 by reducing memory bottlenecks while maintaining performance. These custom kernels, often de-488 veloped in CUDA or Triton, enable fine-grained optimizations tailored to specific hardware architec-489 tures, providing substantial performance improvements. However, the development of such highly 490 optimized components demands significant engineering effort and a deep understanding of both the underlying hardware and neural network operations. Each optimization, whether in matrix multipli-491 cations, element-wise operations, or memory management, requires meticulous custom work, and 492 the resulting implementations are typically neural netowrk architecture-specific. Despite the com-493 plexity, continuous innovation in these low-level building blocks is critical, as it provides massive 494 improvements in memory usage and training speed, making it possible to push the boundaries of 495 model scaling and performance. 496

497 498 5

499

5 LIMITATIONS

500 It is important to note that inverted nonlinearities save memory only when the next layer after the 501 nonlinearity also saves its input for the backward pass. Fortunately, this behavior is common in most 502 popular neural network architectures, such as transformers.

Additionally, our method reduces memory usage only for nonlinearity layers. The percentage of memory savings depends on the architecture in use. Furthermore, the presence of complementary techniques may impact these gains. For instance, when training models with checkpointing Griewank & Walther (2000), the memory bottleneck could shift between activations saved during the forward pass and activations recomputed during the backward pass, potentially affecting the benefits of our approach.

Despite these limitations, our highly efficient implementation of inverted activations within the PyTorch framework using Triton kernels makes us confident that our method will see broad adoption.
We believe that it will prove invaluable to deep learning practitioners who seek to reduce memory
usage without compromising speed order and model accuracy.

513 514

515

6 CONCLUSION

In this work, we presented an innovative method to reduce the memory footprint during neural network training by modifying how activation tensors are saved for backward passes in pointwise nonlinearity layers. By saving the output tensor instead of the input tensor, we significantly decreased the memory requirements for transformer-based architectures such as GPT, BERT, Mistral, and Llama.

Our proposed method includes the development of a space-efficient Boolean indicator to handle the non-invertible nature of these functions, enabling a practical implementation within the PyTorch framework. Additionally, we introduced accurate approximations for the derivatives of the inverse functions of GELU and SiLU, ensuring that training performance remains unaffected while achieving substantial memory savings.

The experiments and approximations demonstrated the feasibility and effectiveness of our approach, with results showing that the training process using our modified nonlinearity layers is indistinguishable from using the original layers in terms of accuracy.

We set out to meet three key criteria for a successful replacement of standard nonlinearity layers:
maintaining model speed, preserving model accuracy, and ensuring ease of use. Our proposed
method satisfies all three conditions. It does not slow down the neural network, it does not degrade
model quality, and it is extremely easy to integrate, making memory savings essentially "free."

We have implemented this method for the PyTorch framework, and the code is available at https:
 //github.com/removed/for/anonimity.

536

537 REFERENCES

Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will 540 Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael 541 Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos, 542 Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Chris-543 tian Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, 544 Phil Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou, Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. 546 In 29th ACM International Conference on Architectural Support for Programming Languages and 547 Operating Systems, Volume 2 (ASPLOS '24). ACM, April 2024. doi: 10.1145/3620665.3640366. 548 URL https://pytorch.org/assets/pytorch2-2.pdf. 549

- Jianfei Chen, Lianmin Zheng, Zhewei Yao, Dequan Wang, Ion Stoica, Michael W. Mahoney, and Joseph Gonzalez. Actnn: Reducing training memory footprint via 2-bit activation compressed training. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event,* volume 139 of *Proceedings of Machine Learning Research,* pp. 1803–1813. PMLR, 2021. URL http://proceedings.mlr.press/v139/chen21z.html.
- Miles D. Cranmer. Interpretable machine learning for science with pysr and symbolic regression.jl.
 CoRR, abs/2305.01582, 2023. doi: 10.48550/ARXIV.2305.01582. URL https://doi.org/
 10.48550/arXiv.2305.01582.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.* OpenReview.net, 2024. URL https://openreview.net/forum?id= mZn2Xyh9Ec.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), pp. 4171–4186. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1423. URL https://doi.org/10.18653/v1/n19-1423.

570

588

589

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha 571 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony 572 Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, 573 Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, 574 Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris 575 Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, 576 Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny 577 Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael 578 Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Ander-579 son, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Ko-580 revaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan 581 Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Ma-582 hadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy 583 Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, 584 Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Al-585 wala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. The 586 llama 3 herd of models. CoRR, abs/2407.21783, 2024. doi: 10.48550/ARXIV.2407.21783. URL https://doi.org/10.48550/arXiv.2407.21783.
 - Stefan Elfwing, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018. doi: 10.1016/J. NEUNET.2017.12.012. URL https://doi.org/10.1016/j.neunet.2017.12.012.
- Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *CoRR*, abs/2103.13630, 2021. URL https://arxiv.org/abs/2103.13630.

- Andreas Griewank and Andrea Walther. Algorithm 799: revolve: an implementation of checkpointing for the reverse or adjoint mode of computational differentiation. ACM Trans. Math. Softw., 26(1):19–45, 2000. doi: 10.1145/347837.347846. URL https://doi.org/10.1145/347837.347846.
- Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian
 error linear units. CoRR, abs/1606.08415, 2016. URL http://arxiv.org/abs/1606.
 08415.
- Pin-Lun Hsu, Yun Dai, Vignesh Kothapalli, Qingquan Song, Shao Tang, and Siyu Zhu. Liger-kernel:
 Efficient triton kernels for llm training, 2024. URL https://github.com/linkedin/
 Liger-Kernel.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.*OpenReview.net, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew G. Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pp. 2704– 2713. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR. 2018.00286. URL http://openaccess.thecvf.com/content_cvpr_2018/html/ Jacob_Quantization_and_Training_CVPR_2018_paper.html.
- ⁶¹⁷ Dhiraj D. Kalamkar, Dheevatsa Mudigere, Naveen Mellempudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharma Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, Jiyan Yang, Jongsoo Park, Alexander Heinecke, Evangelos Georganas, Sudarshan Srinivasan, Abhisek Kundu, Misha Smelyanskiy, Bharat Kaul, and Pradeep Dubey. A study of BFLOAT16
 ⁶²¹ for deep learning training. *CoRR*, abs/1905.12322, 2019. URL http://arxiv.org/abs/1905.12322.
- Zechun Liu, Barlas Oguz, Changsheng Zhao, Ernie Chang, Pierre Stock, Yashar Mehdad, Yangyang Shi, Raghuraman Krishnamoorthi, and Vikas Chandra. LLM-QAT: data-free quantization aware training for large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pp. 467–484. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.FINDINGS-ACL.26. URL https://doi.org/10.18653/v1/2024.findings-acl.26.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Diamos, Erich Elsen, David García, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. In 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net, 2018. URL https://openreview.net/forum?id=r1gs9JgRZ.
- Thomas Müller, Fabrice Rousselle, Jan Novák, and Alexander Keller. Real-time neural radiance
 caching for path tracing. *ACM Trans. Graph.*, 40(4):36:1–36:16, 2021. doi: 10.1145/3450626.
 3459812. URL https://doi.org/10.1145/3450626.3459812.
- Georgii Sergeevich Novikov, Daniel Bershatsky, Julia Gusak, Alex Shonenkov, Denis Valerievich Dimitrov, and Ivan V. Oseledets. Few-bit backward: Quantized gradients of activation functions for memory footprint reduction. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 26363–26381. PMLR, 2023. URL https://proceedings. mlr.press/v202/novikov23a.html.
- 646

598

647 NVIDIA Corporation. Triton Inference Server: An Optimized Cloud and Edge Inferencing Solution. URL https://github.com/triton-inference-server/server.

648 649 650 651	Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. Omniquant: Omnidirectionally calibrated quantization for large language models. In <i>The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024</i> . OpenReview.net, 2024. URL https:
652	//openreview.net/forum?id=8Wuvhh0LYW.
653 655	Noam Shazeer. GLU variants improve transformer. CoRR, abs/2002.05202, 2020. URL https://arxiv.org/abs/2002.05202.
656	Ungloth Ungloth: A framework for officient neural network training https://github.com/
657 658	unlslothai/unsloth, 2023. Accessed: 2024-09-26.
659 660 661 662 663 664 665	Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Train- ing deep neural networks with 8-bit floating point numbers. In Samy Bengio, Hanna M. Wal- lach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), Ad- vances in Neural Information Processing Systems 31: Annual Conference on Neural Informa- tion Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, pp. 7686–7695, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/ 335d3d1cd7ef05ec77714a215134914c-Abstract.html.
666 667 668 669	Yuzhuang Xu, Xu Han, Zonghan Yang, Shuo Wang, Qingfu Zhu, Zhiyuan Liu, Weidong Liu, and Wanxiang Che. Onebit: Towards extremely low-bit large language models. <i>CoRR</i> , abs/2402.11295, 2024. doi: 10.48550/ARXIV.2402.11295. URL https://doi.org/10.48550/arXiv.2402.11295.
670 671 672	<pre>Inc. Yelp open dataset. https://www.yelp.com/dataset, 2020. Accessed: 2024-09- 26.</pre>
673 674 675 676 677 678	Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. Q8BERT: quantized 8bit BERT. In Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing - NeurIPS Edition, EMC2@NeurIPS 2019, Vancouver, Canada, December 13, 2019, pp. 36–39. IEEE, 2019. doi: 10.1109/EMC2-NIPS53020.2019.00016. URL https://doi.org/10.1109/ EMC2-NIPS53020.2019.00016.
679 680	A APPENDIX
681 682	A.1 Memory Table
683 684 685 686 687	We perform memory measurements for several popular neural network architectres from Hggging- face model repository: https://huggingface.co/models. Memory saving percentage given in a table are calculated as how much less memory is used for stored activations, without taking model size and optimizer statistics in consideration. The following models were used:
688 689	• BERT: google-bert/bert-base-uncased Measured for sequence length equals 1024
690 691	 Audio Spectrogram Transformer: MIT/ast-finetuned-audioset-10-10-0.4593 Measured for 1024 × 128 spectrogram shape
692	• ViT (Visual Transformer): google/vit-base-patch16-224-in21k
694	• CLIP: openai/clip-vit-large-patch14
695 696	Measured for text length equals 77 and number of image exaples equal to the number of text examples.
698 699	A.2 APPROXIMATION COEFFICIENTS
700 701	Approximations Eqs. (6) to (8) were built by hands. Left part of GELU approximation Eq. (5) were build with an assistance of PySR Cranmer (2023) framework. Parameters of approximation for GELU:

702				
703				
704		c_0	+1.6311011311381	
705		c_1	+0.169972466666667	
706		c_2	-0.06261728	
707	q^{left}	: c_3	+1.2947087	(10)
708	1	c_4	+1.98055565	
709		c_5	+0.22730362	
710		c_6	-0.038978495	
711		c_7	+1.3295193	
712		c_0	-1.383717971214795	
713		c_1	+1.558420184350027	
714	q^{right}	$: c_2$	+0.044045748018110	(11)
715		c_3	+0.032146736769376	
716		c_4	-2.119885089843949	
710				
710	Parameters of approximation for Sil	LU:		
710				
719				
720		c_0	-1.310856402130980	
721		c_1	+0.848589647031652	
722	q^{left}	$: c_2$	-0.162990512595109	(12)
723		c_3	+0.002696163985044	
724		c_4	-5.770613302664509	
725		Co	+0.217177007595768	
726	right	C1	-0.507684370508263	
727	$q^{ m ngm}$	$: \frac{c_1}{c_2}$	+0.079631397669175	(13)
728		C3	+0.357494204859375	
729		-0		
730	A 3 EVDEDIMENT SETUDS			
731	A.J EXFERIMENT SETUPS			
732	Plain Nonlinearity			
733	Batch size = 2^{25}			
734				
735	• Nonlinearity + Linear			
736	Batch size = 2^{13} , features of	limens	$10n = 2^{10}$.	
737	 MLP Block 			
738	Batch size = 2^{15} , input/out	put din	nension = 2^{10} , hidden dimension = $4 * 2^{10}$.	
739	GeGLU Block	•		
740	Botch size $= 2^{15}$ input dim		-210 sutput dimension $-4 + 210$	
741	Batch size = 2° , input diff	lension	1 = 2, output dimension = $4 * 2$.	
7/10	• BERT			
7/2	Batch size = 64 , sequence	length	= 1024	
743	• Llama v3.1 8B			
744	Batch size – 1 sequence le	noth –	- 519	
740	Duten Size – 1, sequence R	.	012	
740				
/4/				
/48				
749				
750				
751				
752				
753				
754				
755				