

---

# Supplementary Material For: Optical Transformers

---

**Anonymous Author(s)**

Affiliation

Address

email

Table 1: Model configurations for optical Transformers.  $M = 10^6$ .

Model	n	d	h	L	Non-emb. Params
Tiny	1024	192	12	12	15M
Small	1024	384	12	12	40.6M
Base	1024	768	12	12	123.7M
Large	1024	1536	12	12	416.3M

Table 2: Pretraining hyperparameters for optical Transformer models. All models were trained with the **AdamW** [15] optimizer.

Model	Steps	Batch	lr	$\beta_1$	$\beta_2$	$\epsilon$	Weight decay	Dropout	Schedule	Warmup	Stop
Tiny	90000	32	2e-4	0.9	0.999	1e-8	0.02	0.1	Cosine	2500	-
Small	90000	32	2e-4	0.9	0.999	1e-8	0.02	0.1	Cosine	2500	-
Base	90000	32	2e-4	0.9	0.999	1e-8	0.02	0.1	Cosine	2500	-
Large	90000	32	2e-4	0.9	0.999	1e-8	0.02	0.1	Cosine	2500	82500

## 1 A Optical Transformer Training Hyperparameters

2 The optical Transformer models were pretrained on the Wikitext-103 [16] dataset and used the same  
 3 tokenizer as GPT2 [20]. All models used **Xavier uniform initialization** [10]. The architectures are  
 4 in Table 1. Embedding layers were initialized with a normal distribution with  $\sigma = 0.02$ . We used the  
 5 AdamW [15] optimizer, with weight decay applied to parameters which were not embedding, gains,  
 6 or biases. Dropout was applied after every linear layer (including those in attention), as well as on the  
 7 attention matrix and after the  $\text{softmax}(\frac{QK^T}{\sqrt{d_h}})V$  product in the attention calculation. The values of the  
 8 parameters used for the training scheme are in Table 2.

9 After pretraining the models were quantized via our 8-bit QAT scheme. For QAT we used the  
 10 RMSProp optimizer [26]. The parameters we used for the training are in Table 3. To clamp weights  
 11 and activations we employ two different approaches: first, we kept running statistics of minimum  
 12 and maximum values with an exponential moving average (EMA, with parameter  $\alpha$ ) for every  
 13 layer and use those to clamp. Second, we recorded the minimum/maximum statistic throughout  
 14 the network for a forward pass to apply a clipping scheme. Specifically, we clamped weights and  
 15 activations to percentiles of the maximum values collected for each layer. The outputs were either  
 16 rounded to the nearest integer during QAT, or stochastically rounded to nearby values. Finally, for the  
 17 Base-sized model we used to run the experiments, we directly used the lookup tables (LUT) instead  
 18 of “simulating” the quantization of inputs and weights (though outputs are still quantized). Table 4  
 19 details our use of these various techniques in the models.

20 For evaluation we used the perplexity (PPL) metric to measure the language modelling performance  
 21 on Wikitext-103. We evaluated the perplexity over the entire validation set, and ran the model with  
 22 context length 1024 (the same as in training) and a 1024-token stride length.

## 23 B ONN Experimental Procedure

### 24 B.1 Experimental Setup

25 Our setup is a SLM-based matrix-vector/vector-vector multiplier. The setup is shown in Figure 1  
 26 with a simplified illustration in Figure 2, and works as follows: Vectors corresponding to the inputs  
 27 and weights are rearranged into squares of pixels and loaded onto the display and SLM respectively.  
 28 They are aligned such that the light from display pixels will reach the corresponding pixels on the  
 29 SLM. First, light from the display enters into the polarizing beam splitter (PBS), and reaches the  
 30 SLM through a half-wave plate (HWP) which rotates its polarization. The phase is then modified  
 31 by the SLM and reflected back through the half-wave plate, rotating the polarization again based  
 32 on the phase difference. Then, the PBS only admits light of a certain polarization along one of its  
 33 arms, aimed at a camera for detection. Summation of the output pixels is performed digitally. This  
 34 SLM-HWP-PBS arrangement effectively creates an amplitude modulating SLM, where the output at  
 35 each pixel is the element-wise product of the input pixel and corresponding weight pixel.

Table 3: Quantization aware training hyperparameters for optical Transformer models. All models were trained with the **RMSProp** [26] optimizer. Quantization parameters are in Table. 4.

Model	Steps	Batch	lr	$\alpha$	$\epsilon$	Weight decay	Dropout	Schedule	Warmup	Stop
Tiny	7327	64	1e-5	0.99	1e-8	1e-5	0.1	Cosine	2500	-
Small	7327	64	1e-5	0.99	1e-8	1e-5	0.1	Cosine	2500	-
Base	7327	64	1e-5	0.99	1e-8	1e-5	0.1	Cosine	2500	5500
Large	7327	32	1e-5	0.99	1e-8	1e-5	0.1	Cosine	2500	5500

Table 4: Hyperparameters for optical Transformer Quantization. We perform QAT with both a percentile-clipping approach and by clamping based on an exponential moving average (EMA) of model statistics with factor  $\gamma$ . For the Base-sized model that is used in our experiments (LUT-Base), we use lookup tables (LUT) for inputs and weights instead of quantization.

Model	Overall Config		EMA $\gamma$	Attention Clipping			Feed-Forward Clipping		
	Precision	Rounding		Input <sub>1</sub>	Input <sub>2</sub>	Output	Input	Weights	Output
Tiny	8-bit	Stochastic	-	99.99%	99.9%	99.9999%	99.99%	99.9%	99.9999%
Small	8-bit	Stochastic	-	99.99%	99.9%	99.9999%	99.99%	99.9%	99.9999%
Base	8-bit	Stochastic	-	99.99%	99.9%	99.9999%	99.99%	99.9%	99.9999%
Large	8-bit	Stochastic	-	99.99%	99.9%	99.9999%	99.99%	99.9%	99.9999%
LUT-Base	LUT	Stochastic	-	99.99%	98%	99.9999%	99.99%	99%	99.9999%
Tiny	8-bit	Deterministic	0.999	-	-	-	-	-	-
Small	8-bit	Deterministic	0.999	-	-	-	-	-	-
Base	8-bit	Deterministic	0.999	-	-	-	-	-	-
Large	8-bit	Deterministic	0.999	-	-	-	-	-	-

36 The OLED display has multiple color channels and a broad spectrum. For easier modulation by the  
37 SLM, we used a band-pass filter and only green light.

38 The components we used are:

- 39 • Organic light-emitting diode (OLED) display (Google Pixel 2016)
- 40 • Reflective liquid-crystal modulator (1920-500-1100-HDMI, Meadowlark Optics)
- 41 • Half-wave plate (PH10ME-532, Thorlabs)
- 42 • Polarizing beam splitter (CCM1-PBS251, Thorlabs)
- 43 • Zoom lens for imaging onto SLM (Resolv4K, Navitar)
- 44 • Zoom lens and objective lens for imaging onto detector (1-81102, Navita and  
45 XLFLUOR4x/340, Olympus)
- 46 • Band-pass filter (FF01-525/15-25, Semroc)
- 47 • Camera for detection (Prime 95B Scientific CMOS Camera, Teledyne Photometrics)

48 This setup works as a good bench for testing the precision of optical Transformers by performing  
49 optical dot products involved in attention and feed-forward layers. Even though the optical dot  
50 products were performed one at a time, it is sufficient for showing that Transformer operations  
51 can run with the accuracy of ONNs, since matrix-vector and matrix-matrix products are merely  
52 collections of many dot products run in parallel.

## 53 B.2 Calibration and Lookup Tables

54 We used several techniques to reduce errors, map inputs to SLM/display values, and to convert  
55 detected outputs back to neural network values.

56 First, we developed a specialized data-pixel encoding scheme to reduce systematic errors. We noticed  
57 that a large source of error was with a limitation of our hardware—in particular the SLM pixels have  
58 cross-talk (pixels may affect their neighbors if they have very different values) and misalignment  
59 in the experimental setup may lead to corrupted outputs. To help with these issues, we created  
60 “macropixels”—each input element (and weight) does not occupy one pixel on the display (SLM) but  
61 rather is mapped to a 3x3 grid of pixels, all with the same value. For the attention layers, we used 5x5  
62 macropixels for the results we report, but later discovered that with 3x3 the performance is essentially  
63 the same. We also rearranged vectors into square blocks of pixels so that significantly nonzero

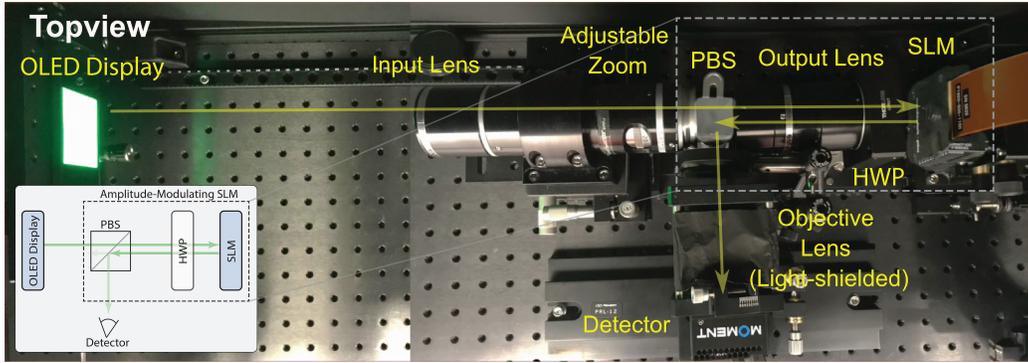


Figure 1: Photo of experimental setup used for running Transformer dot-product operations. Inset: simplified illustration of the experimental system. Spatial light modulator (SLM) + half-wave plate (HWP) + polarizing beam splitter (PBS) arrangement is effectively an amplitude-modulating SLM. The system works as follows: in our experiments, a vector is loaded as pixels on the organic light-emitting diode (OLED) display, and weights on the SLM. The input light enters through the PBS towards the SLM, passing through the HWP twice as the SLM reflects it. The SLM and HWP together rotate the polarization of the light, such that the amount reflected by the PBS towards the detector for each pixel is roughly the product between the pixel value and the corresponding weight on the SLM. The summation of these element-wise products by the detector yields the dot product.

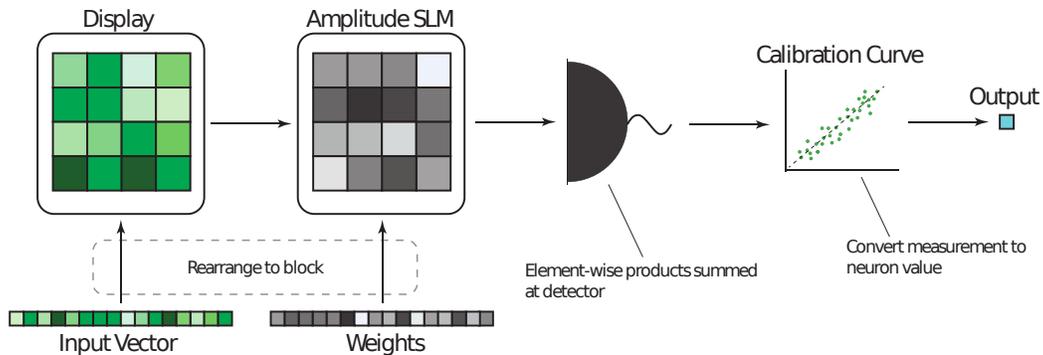


Figure 2: Simplified illustration of experimental setup operation. Weights are loaded and rearranged into a block on spatial light modulator (SLM) to prevent crosstalk between pixels of drastically different values. Data is rearranged on display accordingly. Measurements are looked up against calibration curve to obtain the final output value.

64 weights are nearby each other Figure 2. For the vectors to better fit in the center of the field of view  
 65 (where there is less distortion/misalignment) we computed the dot products using only the 400 largest  
 66 weight elements (the corresponding input elements are loaded). While this may introduce some  
 67 inaccuracy in the final results, we found that the benefits of computing the element-wise products  
 68 more accurately outweigh the drawbacks of pruning the weights; the outputs were still quite accurate  
 69 to the ground-truth dot-product values (see main text, Figure 3). We suspect that this was the case  
 70 because:

- 71 • Transformer weights are not entirely dense; some weights were already zero.
- 72 • Because our setup only supports non-negative data anyway, we use the four-pass approach  
 73 (Section 3.3, main text). This means that for any given dot product, roughly half the weights  
 74 and activations will be zero before considering the previously mentioned sparsity.
- 75 • Meanwhile, a second consequence of this four-pass approach is that roughly half of activa-  
 76 tions will be zero as well, possibly rendering some of the pruned weights irrelevant.

77 • Transformers still perform well when pruned, and luckily larger models can be pruned more  
78 heavily [13]. While our pruning method is quite basic, the number of weights pruned was  
79 light (ie. < 75%) compared to what is possible with more advanced methods.

80 This approach was not necessary for attention calculations, since the dot products were sufficiently  
81 small to fit them entirely (64 elements).

82 Next, we consider the lookup tables (LUTs) of the display and SLM in the setup. In order to optimize  
83 the experimental results, the model used for experiment was trained to be aware of the realistic,  
84 discrete mappable values supported by the system. The display has a LUT with 256 unique levels  
85 (1000 levels total, but many are the same as others) and the SLM has roughly 128 unique levels (256  
86 total). So they are roughly capable of 7 and 8 bit precision. The SLM also cannot fully extinguish  
87 input light—the minimum modulation is 2% of the maximum transmission. Thus, the minimum  
88 absolute values of the weights were mapped to 0.02 instead of 0.

89 After applying these approaches, we finally collected the calibration curve, which maps the output  
90 intensity measurements to neuron values in the neural network and allows us to determine the  
91 experimental setup’s systematic error. To do this we sampled randomly inputs and outputs of  
92 the layers we wished to run, computed their dot products both digitally and in experiment, and  
93 created a data set of experimental measurements and ground-truth-digital dot-product outputs. We  
94 then performed linear regression to find a mapping between experimental output and the correct  
95 values, effectively creating another lookup table. Then when future dot products were computed  
96 experimentally, the output was passed to this linear regression model (or it can literally be stored as a  
97 lookup table) to get the output. We used many photons and averaged outputs across multiple shots  
98 for each input, eliminating shot noise—any remaining error in this calibration scheme we defined as  
99 the system’s systematic error.

100 It is important to note that in general other optical systems might have different causes of error from  
101 ours, but the overall accuracy of our system is representative of a typical ONN nowadays.

### 102 B.3 Model Design Optimization

103 Transformers tend to have large dynamic ranges in their activations and weights [3]. In particular, we  
104 found that systematic error is proportional to some characteristic amplitude of the output. So, because  
105 it scales roughly with the sizes of outputs, having large outlier values can increase the systematic  
106 error and worsen the calibration for all other values in the representable range. Furthermore, after  
107 quantization in a naive, linear scheme, large outliers mean that huge ranges of outputs which are  
108 seldom used are assigned to many of the quantization levels, while the rest of the small, common  
109 outputs are squashed into few buckets—so the model precision is poor. This can be an issue when  
110 quantizing any deep learning model, but was exacerbated here by those systematic errors and the  
111 fact that the lowest levels of the weights are 0.02 and not 0.0. Therefore, we opted for an aggressive  
112 clipping scheme and the clamped activation ReLU6 when training the model to be run (Appendix  
113 A, **LUT-base** model); they reduce the dynamic range of inputs and weights and we found that they  
114 drastically improved the ONN’s ability to run Transformer operations with smaller error. Having  
115 fewer values in the 0.02 bucket of the SLM LUT also improved QAT training stability significantly.  
116 Even though the non-zero light extinction at 0.02 is caused by the specific SLM in our setup, such  
117 issues may happen with other optical implementations made of elements with finite extinction or  
118 resolution, and here we described a method to mitigate such issues by modifying training methods.

### 119 B.4 Transformer Dot Product Samples

120 While the speed and parallelism limitation of our setup made it intractable to run an entire Transformer  
121 model on it, we attempted to sample dot products to run that were representative of the range of  
122 possible activation/weight statistics in the model. That way, our results would be very representative  
123 of what running the full model would be like. In particular, we found two ways in which statistics  
124 throughout the model vary: the statistics change with depth (shallow and deep layers behave differ-  
125 ently) and operation type (matrix-matrix multiplication in attention has different statistics from MLP  
126 layers). So, given our limited ability to run operations on the setup, we sampled roughly 10000 dot  
127 products from the first ( $QK^T$ ) attention operation and second MLP layer of the first and last encoder  
128 layers of the model. The inputs to the whole model were samples from the Wikitext-103 dataset. Our  
129 approach captures the range of statistics throughout a model’s different components, over its depth,

Table 5: Simulated optical Transformer precision ablation. Input precision is degraded by subsampling from lookup table (LUT), while output is quantized. Input precision is approximate, as LUT has 1000 levels, not 1024. Bold: most compressed model found in our ablation with performance very close to the baseline.

Input Precision (LUT)	Output Precision	Val. Loss
~ 10 bits	32 bits	3.0059
~ 9 bits	32 bits	3.0057
~ 8 bits	32 bits	3.0054
~ 7 bits	32 bits	3.0039
~ 6 bits	32 bits	3.0034
~ 5 bits	32 bits	3.0017
~ 4 bits	32 bits	3.0111
~ 3 bits	32 bits	3.1223
~ 5 bits	8 bits	3.0032
<b>~ 5 bits</b>	<b>7 bits</b>	<b>3.0074</b>
~ 5 bits	6 bits	3.0335
~ 5 bits	5 bits	3.3966

130 and when processing a real task’s data. The second MLP layer has dot product size  $4d$ , making it the  
 131 hardest to run experimentally.

132 In sampling the dot products, we tried to sample from both operands equally. For example, one  
 133 could sample 1000 dot products by taking a single input vector and 1000 weight matrix vectors, and  
 134 vice-versa, but choosing random vector pairs captures dot products involving different tokens and  
 135 weights. This is important because Transformer output sizes, particularly the outlier activation values,  
 136 are token-dependent [3]. To maintain this balance, we sample equal rows/columns for both operands.  
 137 For attention layers we sample 100 from each; For linear layers, we sampled 56 rows from the input  
 138 data and 200 columns from the weight matrix  $W^T$ , where the product being computed is  $xW^T$ .

### 139 C Simulated Precision Ablation Study

140 To further study how the optical Transformer can perform inference at lower precisions, we conducted  
 141 a simple ablation on the input and output precisions used at inference, on the 8-bit-QAT base-sized  
 142 model with LUT. We opted to leave the weights at 8-bit precision, since in-place weights are not a  
 143 significant energy cost, and do not take more space/memory in these analog optical systems. In Table  
 144 5 is the performance of the model at lower precisions. With 5-bit input and 7-bit output precision, the  
 145 model performs as well as the baseline. The reported precision values for the LUT are approximate,  
 146 since the LUT has 1000 levels instead of  $2^{10} = 1024$  levels.

147 When using the LUT, it is also not possible to directly change the precision of the input. Instead, we  
 148 employed a subsampling scheme where the precision is degraded by rounding to every  $n$ ’th integer  
 149 level before using the LUT, where  $n$  is a power of 2 and represents a reduction in the effective bit  
 150 precision. The LUT of our display has 1000 levels, some levels have the same value, and we simulate  
 151 the model without added noise. So we say that the original precision is initially *at most* 10 bits  
 152 ( $2^{10} = 1024$ ).

### 153 D ONN Energy Calculation

154 The models we used to estimate the energy use of ONN systems are in Table 6. We used a variety of  
 155 real models that have been introduced by other works, and then designed our family of hypothetical  
 156 future models **FUTURE-\*** in a similar fashion, keeping a reasonable sequence length, increasing the  
 157 embedding dimension drastically, and following the trend of recent large models like PaLM [7] and  
 158 MT-NLG [23] of increasing the ratio  $d/h$ , which results in favorable energy calculations due to the  
 159 lower fraction of memory operations in attention.

160 The calculation of energy costs for ONNs requires consideration of the entire system design and the  
 161 costs of the surrounding electronics—since the optical computation itself is so cheap the electronics

Table 6: Designs of models used for energy estimates. Transformers have embedding dimension  $d$ , process sequence length  $n$ , use  $h$  attention heads, and have  $L$  layers.  $M = 10^6$  parameters.

Model	n	d	h	L	Parameters	Reference
GPT2	1024	768	12	12	117M	[20]
GPT2	1024	1024	16	24	345M	
GPT2	1024	1280	20	36	762M	
GPT2	1024	1600	25	48	1.5B	
Megatron	2048	1536	16	40	1.2B	[22]
Megatron	2048	1920	20	54	2.5B	
Megatron	2048	2304	24	64	4.2B	
Megatron	2048	3072	32	72	8.3B	
GPT3	2048	768	12	32	125M	[4]
GPT3	2048	1024	16	24	350M	
GPT3	2048	1536	16	24	760M	
GPT3	2048	2048	24	24	1.3B	
GPT3	2048	2560	32	32	2.7B	
GPT3	2048	4096	32	32	6.7B	
GPT3	2048	5140	40	40	13B	
GPT3	2048	12288	96	96	175B	
Turing-NLG	1024	4256	28	78	17B	[21]
MT-NLG	2048	20480	128	105	530B	[23]
Chinchilla	2048	640	10	10	73M	[11]
Chinchilla	2048	1024	16	20	305M	
Chinchilla	2048	1280	10	24	552M	
Chinchilla	2048	1792	14	26	1.1B	
Chinchilla	2048	2048	16	28	1.6B	
Chinchilla	2048	3584	28	40	6.8B	
Chinchilla	2048	8192	64	80	70B	
PaLM-like	2048	4096	16	32	8B	[7]
PaLM-like	2048	8192	32	64	62B	
PaLM-like	2048	18432	48	118	540B	
<b>FUTURE</b>	2048	40960	80	120	2.4T	This work
<b>FUTURE</b>	2048	81920	128	200	16T	
<b>FUTURE</b>	2048	163840	160	400	129T	
<b>FUTURE</b>	2048	655360	512	800	4q	

162 account for nearly all of the energy cost. The way the energy is accounted for is as follows: The  
163 energy  $E_{\text{load}}$  can be broken down into three components, related to the energy of the cost of reading  
164 from memory  $E_{\text{read}}$ , digital-to-analog conversion (DAC)  $E_{\text{DAC}}$ , and modulation to generate the light  
165  $E_{\text{mod}}$ :

$$E_{\text{load}} = E_{\text{read}} + E_{\text{DAC}} + E_{\text{mod}}. \quad (1)$$

166 Detection energy consumption  $E_{\text{det}}$  can be broken down in a similar fashion, where

$$E_{\text{det}} = E_{\text{detector}} + E_{\text{amp}} + E_{\text{ADC}} + E_{\text{write}} \quad (2)$$

167 represent the costs of detecting a signal, amplifying the detected signal, performing analog-to-digital  
168 conversion, and writing to memory respectively. There is also a cost of maintaining the weights in  
169 a weights-in-place system, which we call  $E_{\text{maintain}}$ . Because this cost scales per element, it is a  
170 per-MAC cost. But based on values from efficient commercial SLM systems, it is sufficiently small  
171 (and amortized by a large clock rate) that even the largest models we do estimations for are not  
172 bottlenecked. For optical energy, we take 1 eV (single-photon energy at 1240 nm). We started with  
173 using our measured 8-bit-performance photon count of 1500/MAC for the smallest model ( $d = 192$ )  
174 and rescaled the value for larger ones using the constant-per-dot-product trend which we know our  
175 simulated models can match or beat.

176 The assumptions we used were that weights would be loaded from off-chip memory like DRAM (in  
177 the case of a chunked-weights strategy; for a full weights-in-place, one-shot approach this cost does

178 not exist), and that the system uses large amounts of SRAM for activations [9]. We assumed that the  
 179 system only needs 5 bits worth of input precision and 7 bits worth of output precision, per the results  
 180 of our ablation on the base-sized model. We still assumed 8-bit memory accesses for convenience.  
 181 The actual costs for the data access and weight maintenance were assumed to be these values:

- 182 •  $E_{\text{read}} = 1$  pJ/bit for off-chip memory [25], and 0.3 pJ/bit for SRAM. The SRAM estimate  
 183 is based on results for DNN accelerator measurements with 9.55 pJ/32-bit access [19, 8],  
 184 and cutting edge/near-future assumptions for data transport from SRAM/cache [9]. [12]  
 185 estimates 14 pJ per 64-bit access, or roughly 0.22 pJ/bit, for a recent TPU architecture.
- 186 •  $E_{\text{DAC}} = 10$  pJ per 5-bit sample @ 10 GHz—this is achievable with 100 mW at 30.1dB  
 187 SFDR [5].
- 188 •  $E_{\text{mod}} = 1$  fJ/bit @ 110 GHz with thin-film lithium-niobate modulators [28].
- 189 •  $E_{\text{amp}} = 2.4$  pJ per access. A transimpedance amplifier can run at 24 mW at 70 GHz [1]. We  
 190 will just assume 10 GHz.  $24\text{mW} / 10^{10} = 2.4$  pJ per element.
- 191 •  $E_{\text{detector}}$  is negligible compared to  $E_{\text{amp}}$ . For example, [17] calculates the cost of detection  
 192 as the capacitive discharge,  $\frac{1}{2}CV^2$ , with capacitance  $C \sim 1$  fF and voltage  $V = 0.5$  V. This  
 193 results in  $<500$  aJ of energy consumption per detection. The cost is therefore negligible  
 194 compared to amplification ( $E_{\text{amp}}$ ).
- 195 •  $E_{\text{ADC}} = 3.17$  pJ per 7-bit sample. 10 Ghz, need 7-bits of precision, so 128 conversion steps  
 196 per sample – Achievable with 24.8 fJ/c-s [14] ( $24.8 \text{ fJ} \times 128 = 3.17$  pJ per 7-bit sample).
- 197 •  $E_{\text{write}} = E_{\text{read}}$ . Actually, write access was measured to be cheaper than read access in [8],  
 198 but we use  $E_{\text{write}} = E_{\text{read}}$  as a simple, conservative assumption.
- 199 •  $E_{\text{maintain}} = 0.002$  fJ/MAC. Assuming 2W for operation of a 10MP SLM, with inputs  
 200 shone at 10 GHz (each pixel performs one MAC every cycle). There is not much information  
 201 SLM power consumption for maintenance of a fixed pattern on the LCD panel, though more  
 202 typical LCD displays which update can operate in the  $\sim 1$ W regime. For example, [24]  
 203 consumes 30 mW with 180000 pixels, which would scale to 1.67 W with 10MP (at worst,  
 204 multiple SLMs/LCDs could be used in order to scale up).

## 205 E Breakdown-Of-Costs For Estimated ONN Energy Usage

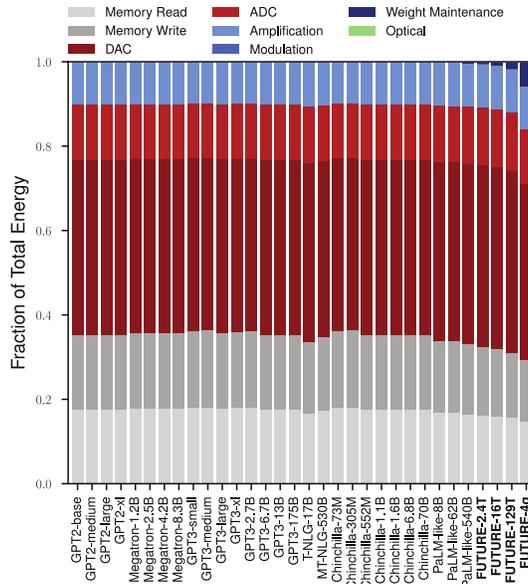


Figure 3: Breakdown of optical Transformer energy costs by energy type at 8-bit operation. Data access costs are dominant due to the high costs of DAC/ADC, but weight maintenance becomes important for large models.

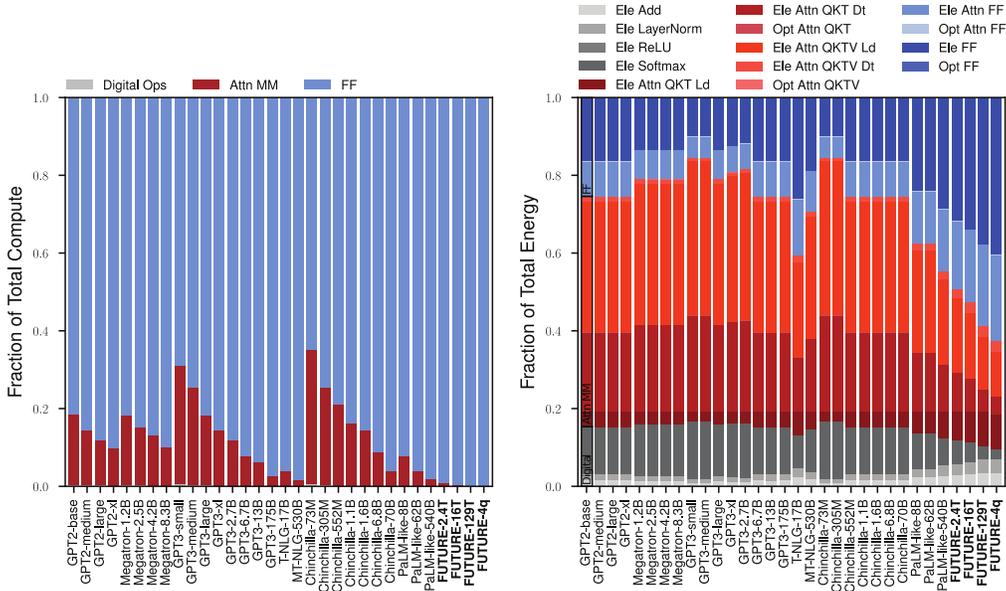


Figure 4: Breakdown of computing costs for optical Transformer models. Left: fraction of total compute used by digital operations, attention, and feed-forward components. Feed-forward layers account for most of the compute. Right: breakdown-of-costs for models by layer. The energy costs of attention operations is expensive. “Ele \*” operations: electrical costs of loading (Ld), detecting (Dt), or both for data for the operation. Operations related to attention computation (ie.  $QK^T$ ) are expensive for little compute. Functions computed digitally have their energy costs estimated as the cost of reading and writing to memory the required data.

206 In Figure 3 we see that data access costs, that is costs per element loaded/stored in memory, are most  
 207 expensive. In particular, the cost of ADC and DAC are the leading contributors to the access costs,  
 208 though since their cost is exponential in the bit precision, one might imagine that a future, optimized  
 209 Transformer running at lower precision than our assumptions would have energy costs dominated  
 210 by the actual SRAM memory costs. Also, for very large models, since the energy from weight  
 211 maintenance scales with the number of MACs, it eventually will dominate if model sizes scale past  
 212 that of FUTURE-4q. But future hardware would reduce  $E_{\text{maintain}}$  through improved electronics or  
 213 higher clock speeds allowing for lower energy per MAC. Finally, the contribution from optical energy  
 214 is vanishingly small, a consequence of the efficient photon usage scaling that we found Transformers  
 215 can leverage. Were it not for this, the cost of actually performing the MACs would be orders of  
 216 magnitude larger than everything else, resulting in energy usage that scales the same way as digital  
 217 systems’.

218 Breaking down the sources of compute and energy costs in Transformer models running optically  
 219 further illustrates how aspects of model/system design affect energy usage. The breakdown of  
 220 compute and energy costs by source is in Figure 4. We find that as models get larger the feed-forward  
 221 layers require most of the computation, but that the energy of data access in attention is still very  
 222 expensive. This is because of the need to save/load many attention matrices from memory, and  
 223 the fact that a weights-in-place scheme cannot be used for the matrix-matrix products because the  
 224 products are between activations. Of course, this also means that there are more activations to load.  
 225 In total, this means that attention layers have high energy costs for small amounts of computation.  
 226 Thankfully, and interestingly, existing model design trends have moved towards focusing much harder  
 227 on feed-forward layers, and so for the largest real (and our hypothetical future) models the fraction of  
 228 energy cost taken by attention is low. Finally, we note that the operations we assume run on digital  
 229 computers - such as nonlinear functions, in gray - do not account for much of the total energy cost  
 230 (though they too are a small fraction of the total compute).

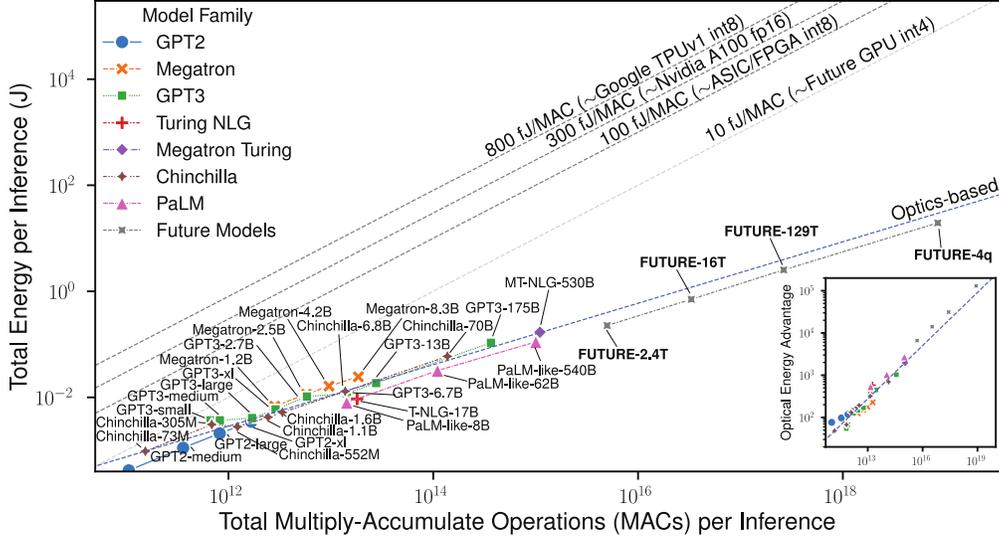


Figure 5: Energy usage estimates of forward pass for Transformers running on optical hardware, under future electronics energy cost assumptions. The energy advantages over our estimate for the current-day NVIDIA A100 GPU are larger than under our original assumptions (main text, Figure 5).  $M = 10^6$ ,  $G = 10^9$ ,  $T = 10^{12}$ ,  $q = 10^{15}$  parameters.

## 231 F Future ONN Energy Consumption

232 As optical accelerators are an emerging technology and as Transformer models continue to scale over  
 233 time, it is worth considering how ONNs might improve over the next several years. For example, an  
 234 interesting question to ask is how well future ONNs will do by the time it is possible to run a large  
 235 model like FUTURE-4q. To investigate this, we estimated the energy costs of various Transformer  
 236 models running optically again, but with the following changes and assumptions:

- 237 •  $E_{\text{maintain}} = 0$ —Future weights-in-place hardware will need effectively no energy to main-  
 238 tain weight information (for example, one might consider the usage of phase change materials  
 239 [27]).
- 240 •  $E_{\text{DAC}}$  and  $E_{\text{ADC}}$  are 1/32 the size—we assume that electronics could achieve a  $2\times$  improve-  
 241 ment in fJ/c-s efficiency, while future advancements in model compression allow for 4-bit  
 242 Transformer models, which are much cheaper since DAC and ADC costs scale exponentially  
 243 with the number of bits [18].
- 244 •  $E_{\text{read}}$  and  $E_{\text{store}}$  are 1/5 the size—there is already a growing recognition of the fact that  
 245 AI accelerators will need high efficiency and large quantities of SRAM and DRAM in the  
 246 future [9, 6].
- 247 •  $E_{\text{amp}}$   $10\times$  cheaper (there are already cheaper trans-impedance amplifiers than our conser-  
 248 vative estimate here, and receiver-less configuration without any amplifier has also been  
 249 demonstrated [2]).

250 Under these assumptions, ONNs become far more efficient, highlighting that improvements to  
 251 electronics will impact ONNs, and not just competing digital hardware. The energy scaling (Figure 5)  
 252 is shifted downward for optics compared to under our previous assumptions, leading to over  $1900\times$   
 253 and  $130,000\times$  advantages over the current A100 GPU for MT-NLG and FUTURE-4q models  
 254 respectively. Of course, by the time this is possible, GPU efficiency will have improved significantly  
 255 as well, and we are comparing a 4-bit accelerator to the 16-bit performance of the A100. It is difficult  
 256 to predict the future efficiency of GPUs at lower precision, but it is clear that ONNs can benefit from  
 257 improvements to electronics and low-precision inference.

## 258 **G Scaling ONNs: System Specifications and Communication Costs in** 259 **Multi-Processor ONN and GPU Setups**

260 Implementation of a real ONN for large models might be difficult because the amount of hardware  
261 needed to maintain all the weights is exceedingly large. In Table 7 are the requirements for hardware  
262 to run the largest future model. To compute the number of weights/elements, we selected the largest  
263 MLP layer in the model, since that requires the most space for weights and activations. While  
264 detector and memory requirements are achievable, the number of required cores—each an optical  
265 component capable of performing 10M multiplications with weights—is enormous. There are some  
266 approaches to remedy this kind of memory issue in both GPUs and ONNs, and we are interested in  
267 their hardware-time-energy tradeoffs for ONNs.

268 One solution is to introduce chunking, where only a portion of the weights are loaded at a time, and  
269 the inputs are passed through. Then, the amount of time it takes to run is increased by a factor of  
270 the number of chunks. This also impacts the optical system’s energy advantage over digital ones in  
271 two ways. First, the weights must be loaded, but the cost can be amortized via reuse with batched  
272 inference. This comes at the expense of latency. This is a new kind of tradeoff, since digital systems  
273 cannot reuse weight data for free. Second, for each weight chunk, all inputs must be reloaded;  
274 changing the chunk number trades energy efficiency for lower hardware requirements. These energy  
275 tradeoffs are illustrated in Figure 6; other factors dominate energy usage until the chunk number is  
276 large and chunking becomes the bottleneck.

277 Realizing large models with GPUs will likely also require a multi-GPU strategy, which will incur  
278 overhead over the peak performance of a single GPU. We find that with a simple model of communi-  
279 cation costs—modelling the activation reloading in both GPU and ONN systems—that ONNs can  
280 retain some of their advantages, dependant on how much system memory (or maximum number  
281 of weight elements) is available per-processor. We created a simple model to estimate the cost of  
282 this approach in GPU systems. In GPU systems, instead of splitting a model over time, the model  
283 may instead be split over multiple GPUs. This introduces an analogous tradeoff to the activation  
284 reloading in ONNs due to communication costs: if each GPU holds some chunk of weights, then  
285 after every layer, the outputs of multiplying the inputs with each chunk must be broadcasted to every  
286 GPU in an all-to-all fashion. This is in essence an all-reduce operation—after every layer, the outputs  
287 from all GPUs must be copied onto all GPUs. In total, this means the total number of activations  
288 is loaded  $k$  times, where  $k$  is the number of GPUs. As a crude but conservative estimate of these  
289 costs, we modeled this by taking the cost of running the entire model on one GPU, and then adding  
290 the energy cost of loading the activations from DRAM, multiplied by the number of chunks (GPUs).  
291 This is likely an underestimate, as broadcasting data across GPUs in a real setup requires sending  
292 data electronically over much longer distances than required for DRAM access, which would be  
293 expensive.

294 To determine the number of chunks, we tested multiple assumptions about device memory. We  
295 assumed a value for the amount of memory that can be used to store weights and take the total number  
296 of weights for each model divided by this memory capacity to determine the number of chunks to be  
297 used.

298 With these models, we found that too much chunking is detrimental to ONN performance, but that  
299 there is still some energy advantage to be had if it is used sparingly (Figure 6). In Figure 7 (top)  
300 are the energy cost estimates assuming a fixed memory of 100M weights (ie. 100MPixel SLM, or  
301 RAM with 100MB capacity if each weight is one byte). We assumed that for GPU, the cost of  
302 communication is at least that of DRAM-level communication due to the physical distances between  
303 GPUs. The curves for GPUs bend upward as the communication costs begin to take over, as do the  
304 largest models running optically. The ONNs still maintain an advantage, but the advantage stops  
305 growing with model size. Looking at the energy advantage illustrates this idea more clearly: up to a  
306 certain model size the advantage is increasing, then as the model size reaches the memory limit it  
307 begins to level off, and then the advantage begins to shrink as the cost of chunking takes over. For  
308 a small range of model sizes near this peak, the advantage is maintained, suggesting that a small  
309 amount of chunking may be useful before it quickly diminishes the energy advantage.

310 The optimal configuration for ONNs, obviously, is to have enough memory (cores which have weights  
311 fixed in place) so that chunking is not necessary. When plotting the advantages for larger memories  
312 (and therefore fewer chunks), the advantage gets better, and larger models become worthwhile to

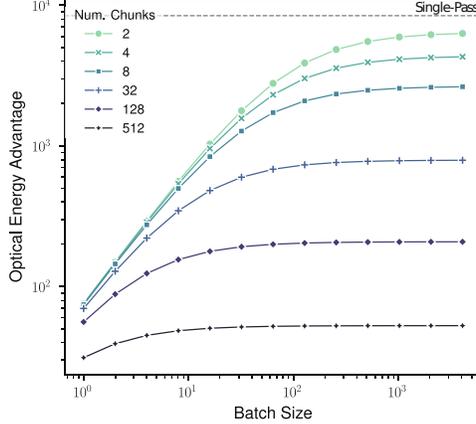


Figure 6: Optical energy advantage vs A100 (FUTURE-4q). When chunking, the cost of loading weights is amortized by increasing batch size, but the overall performance is limited by large numbers of chunks because of input data reloading.

313 run. In hindsight this conclusion makes sense: the benefit of ONNs is their ability to copy data  
 314 (“optical fan-out”) for free for parallel computation, and so reducing this in favor of repeated memory  
 315 accesses removes exactly the mechanism that gives optics-based systems their advantages. This also  
 316 suggests that an “optical memory” from which fixed data can be accessed for free (or significantly less  
 317 than re-access through electronics) may solve this problem, allowing for more scalable ONN design  
 318 without huge amounts of hardware for weights. Currently, optics still has an advantage when using  
 319 multiple cores because in principle the data could be fanned out across cores, while GPUs must pay  
 320 communication costs in multi-processor setups. With a fan-out/fan-in design that can collect/spread  
 321 a vector across cores, the efficiency of an entirely weights-in-place system is fully that of a single,  
 322 large core.

Table 7: Requirements for optical accelerator running feed-forward layer (embedding dimension  $d$ , sequence length  $n$ ) without chunking at 8-bit precision. The requirement of many cores to maintain weights for matrix-vector products (MVM) is high, and we assume the ONN system requires static RAM (SRAM) for saving and loading activations.

Model	Input Vector Elements	Detectors	MVM Cores ( $10^7$ weights each)	SRAM (activations)
FUTURE-4.1q	$2.6 \times 10^6$	$2.6 \times 10^6$	170,000	5.37 GB
FUTURE-129T	$6.55 \times 10^5$	$6.55 \times 10^5$	11,000	1.34 GB
FUTURE-16T	$3.28 \times 10^5$	$3.28 \times 10^5$	2,700	671 MB
FUTURE-2.4T	$1.64 \times 10^5$	$1.64 \times 10^5$	671	336 MB
PaLM-like-540B	$7.37 \times 10^4$	$7.37 \times 10^4$	136	151 MB
MT-NLG-530B	$8.19 \times 10^4$	$8.19 \times 10^4$	168	168 MB
GPT3-175B	$4.91 \times 10^4$	$4.91 \times 10^4$	61	100 MB
<b>General</b>	$4d$	$4d$	$4d^2/10^7$	$4nd$

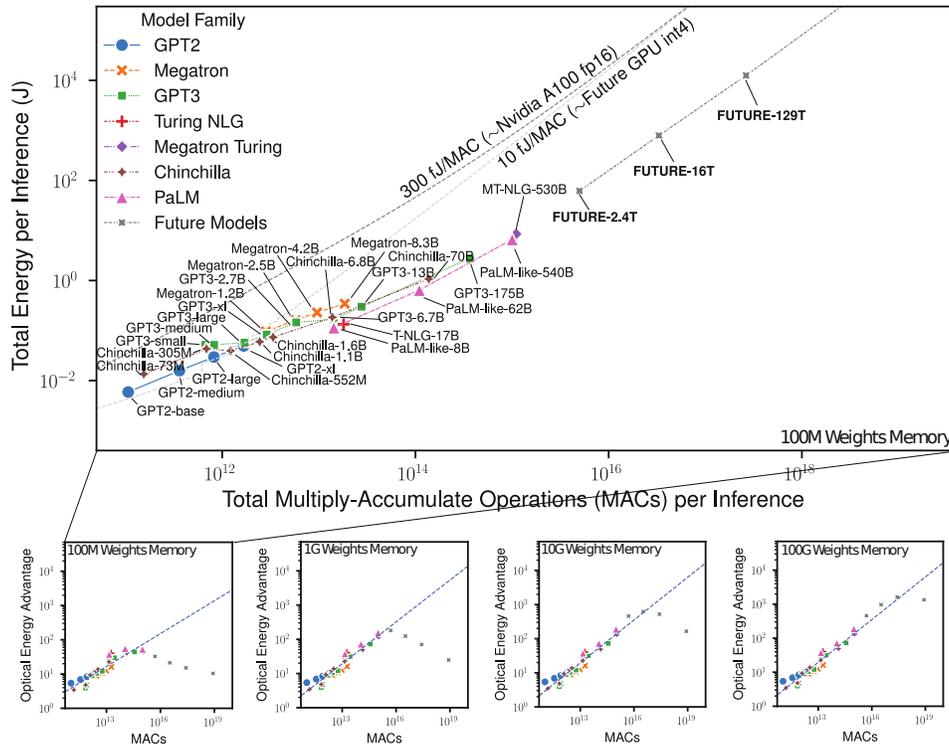


Figure 7: Energy estimates assuming a fixed processor memory size and chunking. Top: estimated energy scaling plot for Transformer models running on optical and digital hardware with 100MB of memory. As models get larger, both optical and digital systems have an upward bend in energy consumption trends, driven by communication/input-reloading-from-chunking costs. Bottom: energy advantage scaling for different memory sizes. As the memory increases, there is a maximum energy advantage for optics over NVIDIA A100 and corresponding model size before chunking costs take over.  $M = 10^6$ ,  $G = 10^9$ ,  $T = 10^{12}$ ,  $q = 10^{15}$  parameters.

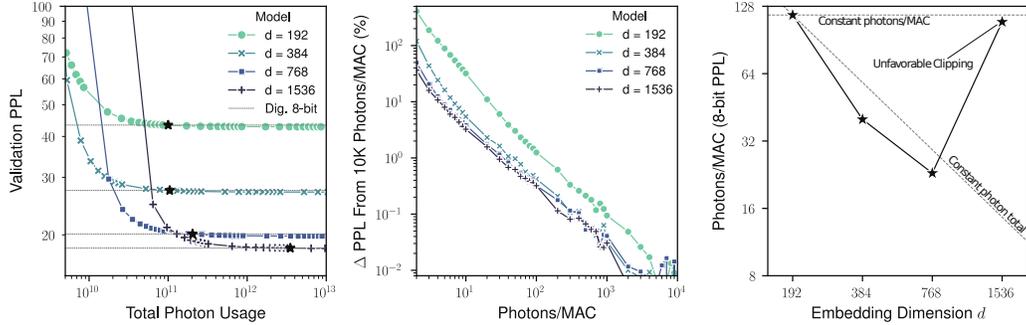


Figure 8: Behavior of optical Transformer models with varying photon usage with percentile clipping scheme. Left: Wikitext-103 validation set perplexity (PPL) versus embedding dimension  $d$  and total photons usage. 8-bit quantized digital model performance levels in dashed lines. Middle: Percent change in perplexity from ideal 10000 photon count performance still exhibits truncated power-law scaling with photons per multiply-accumulate (MAC) operation for all models. Right: Scaling of photon usage for maintaining the 8-bit digital performance versus model size. Dashed lines: constant photons per dot product (optical scaling) and constant photons/MAC analogous to digital scaling. Note that unlike for our results in the main text, smaller models beat the constant-dot-product-total scaling, but the largest model exhibits poor efficiency, as the clipping scheme used here was not well suited for it.

## 323 H Effects of Training and Quantization Scheme on Optical Scaling

324 Our results demonstrating favorable scaling of photon usage in Transformers show that they can  
 325 be optically efficient, but in general the photon usage is affected by the training scheme and other  
 326 factors like quantization. This is because approaches for optimization quantization, regularization,  
 327 etc. affect the statistics of weights and activations in the network, which unlike digital systems, are  
 328 tied to the resource usage. The main example of this is with weights: they are normalized before  
 329 being loaded onto an ONN accelerator, and so large outliers may lead to many weights being near 0  
 330 after normalization—admitting fewer photons through to the detector. This has a direct impact on the  
 331 output SNR, and so depending on weight statistics more or fewer photons may be needed in order to  
 332 run at the same precision.

333 To discover how a different scheme might affect photon usage, we analyzed the optical scaling of our  
 334 quantized optical Transformer models with percentile clipping instead of clamping based on EMA  
 335 statistics. We applied the same clipping to all models (details in Table 4). These clipped models have  
 336 familiar trends in their language modelling performance versus photon numbers, but we notice key  
 337 differences in the photons needed to maintain 8-bit digital performance: first, the absolute number  
 338 of photons needed for the smaller models (120 and 40 versus 340 and 170 of our unclipped scheme  
 339 for  $d = 192, 384$ ) is much lower—this indicates that clipping of large weight values leads to more  
 340 transmission after normalization. Second, the scaling is inconsistent, with smaller models needing  
 341 significantly fewer photons than the expected  $1/d$  scaling, but then requiring many photons again  
 342 for the largest model. The clipping scheme degraded the performance of the large model. Of course,  
 343 this could be improved by designing a better scheme for the largest model such that it requires few  
 344 photons; these results illustrate how differences in the training and quantization recipe could lead to a  
 345 variety of outcomes, and why efficiency is achievable but not an automatic guarantee for any scheme.

## References

- 346
- 347 [1] Maruf N. Ahmed, Joseph Chong, and Dong Sam Ha. A 100 Gb/s transimpedance amplifier in  
348 65 nm CMOS technology for optical communications. In *2014 IEEE International Symposium*  
349 *on Circuits and Systems (ISCAS)*, pages 1885–1888, 2014. doi: 10.1109/ISCAS.2014.6865527.
- 350 [2] Saumil Bandyopadhyay, Alexander Sludds, Stefan Krastanov, Ryan Hamerly, Nicholas Harris,  
351 Darius Bunandar, Matthew Streshinsky, Michael Hochberg, and Dirk Englund. Single chip  
352 photonic deep neural network with accelerated training. *arXiv preprint arXiv:2208.01623*, 2022.  
353 URL <https://doi.org/10.48550/arXiv.2208.01623>.
- 354 [3] Yelysei Bondarenko, Markus Nagel, and Tijmen Blankevoort. Understanding and overcoming  
355 the challenges of efficient Transformer quantization. In *Proceedings of the 2021 Conference*  
356 *on Empirical Methods in Natural Language Processing*, pages 7947–7969, Online and Punta  
357 Cana, Dominican Republic, November 2021. Association for Computational Linguistics. URL  
358 <https://aclanthology.org/2021.emnlp-main.627>.
- 359 [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal,  
360 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel  
361 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M.  
362 Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz  
363 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec  
364 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.  
365 URL <https://arxiv.org/abs/2005.14165>.
- 366 [5] Pietro Caragiulo, Clayton Daigle, and Boris Murmann. DAC performance survey 1996-2020.  
367 URL <https://github.com/pietro-caragiulo/survey-DAC>.
- 368 [6] Cerebras Systems. Cerebras systems: Achieving industry best AI perfor-  
369 mance through a systems approach. Technical report, Apr 2021. URL [https://8968533.fs1.hubspotusercontent-na1.net/hubfs/8968533/whitepapers/  
370 Cerebras-CS-2-Whitepaper.pdf](https://8968533.fs1.hubspotusercontent-na1.net/hubfs/8968533/whitepapers/Cerebras-CS-2-Whitepaper.pdf).
- 371
- 372 [7] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
373 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker  
374 Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes,  
375 Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson,  
376 Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin,  
377 Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier  
378 Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David  
379 Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani  
380 Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat,  
381 Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei  
382 Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei,  
383 Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. PaLM: Scaling  
384 language modeling with Pathways, 2022. URL <https://arxiv.org/abs/2204.02311>.
- 385 [8] Benoît W. Denkinger, Flavio Ponzina, Soumya S. Basu, Andrea Bonetti, Szabolcs Balási,  
386 Martino Ruggiero, Miguel Peón-Quirós, Davide Rossi, Andreas Burg, and David Atenza.  
387 Impact of memory voltage scaling on accuracy and resilience of deep learning based edge  
388 devices. *IEEE Design & Test*, 37(2):84–92, 2020. doi: 10.1109/MDAT.2019.2947282.
- 389 [9] Yaosheng Fu, Evgeny Bolotin, Niladrish Chatterjee, David Nellans, and Stephen W. Keckler.  
390 GPU domain specialization via composable on-package architecture. *ACM Trans. Archit.*  
391 *Code Optim.*, 19(1), dec 2021. ISSN 1544-3566. doi: 10.1145/3484505. URL <https://doi.org/10.1145/3484505>.
- 392
- 393 [10] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward  
394 neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth*  
395 *International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of*  
396 *Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May  
397 2010. PMLR. URL <https://proceedings.mlr.press/v9/glorot10a.html>.

- 398 [11] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza  
399 Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom  
400 Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia  
401 Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent  
402 Sifre. Training compute-optimal large language models, 2022. URL [https://arxiv.org/  
403 abs/2203.15556](https://arxiv.org/abs/2203.15556).
- 404 [12] Norman P. Jouppi, Doe Hyun Yoon, Matthew Ashcraft, Mark Gottscho, Thomas B. Jablin,  
405 George Kurian, James Laudon, Sheng Li, Peter Ma, Xiaoyu Ma, Thomas Norrie, Nishant  
406 Patil, Sushma Prasad, Cliff Young, Zongwei Zhou, and David Patterson. Ten lessons from  
407 three generations shaped Google’s TPUv4i : Industrial product. In *2021 ACM/IEEE 48th  
408 Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14, 2021. doi:  
409 10.1109/ISCA52012.2021.00010.
- 410 [13] Zhuohan Li, Eric Wallace, Sheng Shen, Kevin Lin, Kurt Keutzer, Dan Klein, and Joseph E.  
411 Gonzalez. Train large, then compress: Rethinking model size for efficient training and inference  
412 of transformers. In *Proceedings of the 37th International Conference on Machine Learning,  
413 ICML’20*. JMLR.org, 2020.
- 414 [14] Juzheng Liu, Mohsen Hassanpourghadi, and Mike Shuo-Wei Chen. A 10GS/s 8b 25fJ/c-s  
415 2850um<sup>2</sup> two-step time-domain ADC using delay-tracking pipelined-SAR TDC with 500fs  
416 time step in 14nm CMOS technology. In *2022 IEEE International Solid- State Circuits  
417 Conference (ISSCC)*. IEEE, February 2022. doi: 10.1109/isscc42614.2022.9731625. URL  
418 <https://doi.org/10.1109/isscc42614.2022.9731625>.
- 419 [15] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. 2019. URL  
420 <https://openreview.net/forum?id=Bkg6RiCqY7>.
- 421 [16] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture  
422 models. In *International Conference on Learning Representations (ICLR)*, 2017. URL [https:  
423 //openreview.net/forum?id=Byj72udxe](https://openreview.net/forum?id=Byj72udxe).
- 424 [17] David A B Miller. Attojoule optoelectronics for low-energy information processing and  
425 communications. *Journal of Lightwave Technology*, 35(3):346–396, 2017. URL [https:  
426 //doi.org/10.1109/JLT.2017.2647779](https://doi.org/10.1109/JLT.2017.2647779).
- 427 [18] Boris Murmann. ADC performance survey 1997-2020. [http://web.stanford.edu/  
428 ~murmam/adcsurvey.html](http://web.stanford.edu/~murmam/adcsurvey.html), 2020. Online Accessed: 2021-02-18.
- 429 [19] Flavio Ponzina, Miguel Peon-Quiros, Andreas Burg, and David Atienza. E<sup>2</sup>CNNs: Ensem-  
430 bles of convolutional neural networks to improve robustness against memory errors in edge-  
431 computing devices. *IEEE Transactions on Computers*, 70(8):1199–1212, August 2021. doi:  
432 10.1109/tc.2021.3061086. URL <https://doi.org/10.1109/tc.2021.3061086>.
- 433 [20] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Lan-  
434 guage models are unsupervised multitask learners. 2019. URL [https://openai.com/blog/  
435 better-language-models/](https://openai.com/blog/better-language-models/).
- 436 [21] Corby Rosset. Turing-NLG: A 17-billion-parameter language model by Mi-  
437 crosoft, Feb 2020. URL [https://www.microsoft.com/en-us/research/blog/  
438 turing-nlg-a-17-billion-parameter-language-model-by-microsoft/](https://www.microsoft.com/en-us/research/blog/turing-nlg-a-17-billion-parameter-language-model-by-microsoft/).
- 439 [22] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan  
440 Catanzaro. Megatron-LM: Training multi-billion parameter language models using model  
441 parallelism, 2019. URL <https://arxiv.org/abs/1909.08053>.
- 442 [23] Shaden Smith, Mostofa Patwary, Brandon Norrick, Patrick LeGresley, Samyam Rajbhandari,  
443 Jared Casper, Zhun Liu, Shrimai Prabhunoye, George Zerveas, Vijay Korthikanti, Elton Zhang,  
444 Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi,  
445 Yuxiong He, Michael Houston, Saurabh Tiwary, and Bryan Catanzaro. Using DeepSpeed and  
446 Megatron to train Megatron-Turing NLG 530B, a large-scale generative language model, 2022.  
447 URL <https://arxiv.org/abs/2201.11990>.

- 448 [24] Sony Corporation. *LCX027AKB Datasheet (PDF) - Sony Corporation*. Sony Corpo-  
449 ration. URL [https://pdf1.alldatasheet.com/datasheet-pdf/view/47550/SONY/](https://pdf1.alldatasheet.com/datasheet-pdf/view/47550/SONY/LCX027AKB.html)  
450 [LCX027AKB.html](https://pdf1.alldatasheet.com/datasheet-pdf/view/47550/SONY/LCX027AKB.html).
- 451 [25] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep  
452 neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.  
453 URL <https://doi.org/10.1109/JPROC.2017.2761740>.
- 454 [26] Tijmen Tieleman, Geoffrey Hinton, et al. Lecture 6.5-rmsprop: Divide the gradient by a running  
455 average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):  
456 26–31, 2012.
- 457 [27] Matthias Wuttig, Harish Bhaskaran, and Thomas Taubner. Phase-change materials for non-  
458 volatile photonic applications. *Nature photonics*, 11(8):465–476, 2017. URL [https://doi.](https://doi.org/10.1038/nphoton.2017.126)  
459 [org/10.1038/nphoton.2017.126](https://doi.org/10.1038/nphoton.2017.126).
- 460 [28] Mengyue Xu, Yuntao Zhu, Fabio Pittalà, Jin Tang, Mingbo He, Wing Chau Ng, Jingyi Wang,  
461 Ziliang Ruan, Xuefeng Tang, Maxim Kuschnerov, et al. Dual-polarization thin-film lithium  
462 niobate in-phase quadrature modulators for terabit-per-second transmission. *Optica*, 9(1):61–62,  
463 2022.