Effective Theory of Transformers at Initialization

Emily Dinan,[*] Sho Yaida,[†] and Susan Zhang[‡]

Meta AI

Meta Platforms, Inc.[§]

We perform an effective-theory analysis of forward–backward signal propagation in wide and deep Transformers, i.e., residual neural networks with multi-head self-attention blocks and multilayer perceptron blocks. This analysis suggests particular width scalings of initialization and training hyperparameters for these models. We then take up such suggestions, training Vision and Language Transformers in practical setups.

———
[*]Electronic address: edinan@meta.com
[†]Electronic address: shoyaida@meta.com
[‡]Electronic address: susanz@meta.com
[§]The author ordering was determined by the hypothetical coin toss that 100%-respects the alphabetical ordering.

Contents

I.   Theoretical Foundations

Autobots, transform and roll out!

Optimus Prime

We live in an exciting time in the history of artificial intelligence science and technology. On the one hand, our state-of-the-art models are getting larger and larger in the blink of an eye [1–5], with the number of their model parameters in the order of trillions at the time of this writing. On the other hand, we have an effective theoretical description that becomes asymptotically more accurate and more relevant as models get wider [6–9] and deeper [10–12]. Thus, we now have no excuse but to bring the theory and practice closer together.

Transformers [13, 14] provide an ideal case in point, not only because they are driving numerous technical innovations of this era but also because – to achieve these innovations – they are getting so wide and deep that the effective theory should in principle be getting more apt and relevant. With that optimistic conviction, in Part I of this note, we'll roll out the general-purpose blueprint laid out in Ref. [15] and fine-tune it to develop the effective theory of Transformers; in Part II, we'll then see the impacts of the resulting theoretical suggestions in practice.[1]

The rest of Part I is organized as follows. We'll kick off our forward-path analysis in §I 0 by briefly reviewing the building blocks of Transformers while setting up our notations. Then in §I 1 we'll calculate the statistics of preactivations at initialization, all the way from inputs to outputs: this analysis helps us determine how to scale initialization hyperparameters with width (1.25)–(1.33). Shifting our gear to the backward path, after a brief detour on neural tangent kernels in §I 2, we'll calculate the statistical means of squared gradients in §I 3: this analysis helps us determine how to scale group-wise learning-rate factors with width, both for the stochastic gradient descent optimizer (1.99)–(1.107) and for the AdamW optimizer (1.108)–(1.116).[2]

---

[1] Realistically speaking, Transformers won't forever be the pinnacle of the neural-network evolution but will be replaced by other model families. The broader goal of this note is to provide a *meta-blueprint* to roll out the said effective-theory blueprint – which was illustrated in Ref. [15] for the simple case of multilayer perceptrons – for more complicated neural-network architectures that come our way. To that end, we'll be as verbose as possible in our presentation so that it should be relatively straightforward to mirror the way we roll out the blueprint.

[2] For those in the know: our effective-theory analysis stays at a meta level. That is, there are two things we *won't* do: *(a)* the criticality analysis – because the putative exploding or vanishing gradient problems are mostly taken care of by normalization layers and residual connections – and *(b)* the finite-width analysis – because such a treatment would quadruple the length of our note and such labor is not necessary in reading off the hyperparameter scalings that we'll test in Part II. That said, hopefully this note lays the foundations for those who dare to push on these extensions, which can be useful, e.g., in investigating how to judiciously scale up architecture hyperparameters.

0.   A Crash Course on Transformers

Given a dataset $\mathcal{D}$, we denote its tokenized inputs as $x_{\alpha;t;i}$ with a sample index $\alpha = 1, \ldots, |\mathcal{D}|$, a token index $t = 1, \ldots, T$, and a vector index $i = 1, \ldots, n_{\text{in}}$ (see §I 0 A for concrete examples). Then a standard Transformer [13, 14] with an embedding dimension n recursively transforms them as

$$z_{\alpha;t;i}^{(1)} = \text{R}_{t;i}^{(1)} \left(x_\alpha; \theta^{(1)} \right) \qquad \text{for} \quad i = 1, \ldots, n, \tag{1.1}$$

$$s_{\alpha;t;i}^{(\ell)} = \text{LN}_i \left(z_{\alpha;t}^{(\ell)} \right) \qquad \text{for} \quad i = 1, \ldots, n; \quad \ell = 1, \ldots, L-1, \tag{1.2}$$

$$z_{\alpha;t;i}^{(\ell+1)} = \text{R}_{t;i}^{(\ell+1)} \left(s_\alpha^{(\ell)}; \theta^{(\ell+1)} \right) + z_{\alpha;t;i}^{(\ell)} \quad \text{for} \quad i = 1, \ldots, n; \quad \ell = 1, \ldots, L-2, \tag{1.3}$$

$$z_{\alpha;t;i}^{(L)} = \text{R}_{t;i}^{(L)} \left(s_\alpha^{(L-1)}; \theta^{(L)} \right) \qquad \text{for} \quad i = 1, \ldots, n_{\text{out}}, \tag{1.4}$$

where $\theta^{(\ell)}$ are the model parameters for the ℓ-th block. Specifically, first, the stem-block operation $\text{R}^{(1)}$ (1.1) converts the inputs $x_{\alpha;t;i}$ into the first-block preactivations $z_{\alpha;t;i}^{(1)}$, which are T-sequenced n-dimensional vectors; then, these preactivations recursively go through a series of the layer-normalization operations (1.2) [16] (§I 0 B) and block operations (1.3) – with each residual path $\text{R}^{(\ell)}$ being either a multi-head self-attention block (§I 0 C) or a multilayer perceptron block (§I 0 D) – that are sandwiched by skip connections; and, last, the head-block operation $\text{R}^{(L)}$ (1.4) converts them into T-sequenced n_{out}-dimensional vectors $z_{\alpha;t;i}^{(L)}$ (see §I 0 E for concrete examples).[3]

Now that we have an overall schematic, let us define each operation in detail, block by block.

A.   Stem Block

Vision: patchify embedding and positional embedding

For Vision Transformers [14], input images are typically tokenized into non-overlapping patches. To be very concrete, in typical ImageNet [18] training setups, given a 3-colored 224-by-224-pixelated

---

[3] More generally, each skip connection can come with its own parameters as

$$z_{\alpha;t;i}^{(\ell+1)} = \text{R}_{t;i}^{(\ell+1)} \left(s_\alpha^{(\ell)}; \theta^{(\ell+1)} \right) + \xi_i^{(\ell+1)} z_{\alpha;t;i}^{(\ell)}, \tag{1.5}$$

where element-wise affine parameters $\xi_i^{(\ell)}$ dictate how much of the signal we copy and paste in each channel. But, in this note, we'll keep them fixed as $\xi_i^{(\ell)} = 1$ for simplicity – as is also standard in practice. Similarly, each layer-normalization operation can come with its own trainable element-wise affine parameters (§I 0 B), but we'll again keep them fixed in our theoretical and practical treatments – which is less standard but not unheard of: `https://pytorch.org/docs/stable/generated/torch.nn.LayerNorm.html`.

Incidentally, normalization layers were originally placed differently [13], but the above positioning is more standard today and, for residual neural networks, such placement ensures that block-to-block signal propagation is *everywhere critical* [17], that is, it makes the success of training less sensitive to the choices of initialization hyperparameters.

image, we can for instance partition it into ($T = 14^2 = 196$) tokens of non-overlapping patches, each of which can be seen as an ($n_{\text{in}} \equiv n_{\text{patch}} = 16 \cdot 16 \cdot 3 = 768$)-dimensional vector. To implement this, we can use patchify weights – 16-by-16 convolutional weights with stride 16, intaking 3 in-channels and outputting n out-channels – which act as

$$\tilde{z}_{\alpha;t;i} = \sum_{j=1}^{n_{\text{patch}}} W_{ij}^{\text{patch}} x_{\alpha;t;j} \,, \tag{1.6}$$

with the embedding index $i = 1, \ldots, n$. Note that, due to convolutional weight tying, these weights act in the same way for all tokens $t = 1, \ldots, T$. Then they are often amended by positional-embedding parameters as

$$z_{\alpha;t;i}^{(1)} = b_{t;i}^{\text{PE}} + \tilde{z}_{\alpha;t;i} = b_{t;i}^{\text{PE}} + \sum_{j=1}^{n_{\text{patch}}} W_{ij}^{\text{patch}} x_{\alpha;t;j} \,, \tag{1.7}$$

where the T-by-n-dimensional tensor $b_{t;i}^{\text{PE}}$ acts like bias parameters and is folklored to be useful for distinguishing different patches.[4]

Language: word embedding and positional embedding

For Language Transformers [13], inputs are typically given in the form of – or at least can be interpreted as – one-hot vectors $x_{\alpha;t;j} = \delta_{jj_\star}$ where the vector index $j = 1, \ldots, n_{\text{vocab}}$ runs over all possible tokens in the vocabulary and, for a given pair of sample–token indices (α, t), the input vector takes the unit value at a specific index $j_\star = j_\star(\alpha; t)$ and otherwise returns zero. Each word is then embedded into an n-dimensional space as

$$\tilde{z}_{\alpha;t;i} = \sum_{j=1}^{n_{\text{vocab}}} W_{ij}^{\text{WE}} x_{\alpha;t;j} = W_{ij_\star(\alpha;t)}^{\text{WE}} \,, \tag{1.9}$$

where the word-embedding parameters W_{ij}^{WE} act like an n-by-n_{vocab} weight matrix. As in the vision case (1.7), they are then amended by the positional-embedding parameters as $z_{\alpha;t;i}^{(1)} = \tilde{z}_{\alpha;t;i} + b_{t;i}^{\text{PE}}$.

---

[4] There exist several variants of the stem block in Vision Transformers. As one variant, we can attach an n-dimensional class token c_i at, say, $t = 0$ such that we have $(T + 1)$-by-n-dimensional first-block preactivations

$$z_{\alpha;t;i}^{(1)} = \begin{cases} c_i + b_{0;i}^{\text{PE}} & \text{for } t = 0 \,, \\ b_{t;i}^{\text{PE}} + \sum_{j=1}^{n_{\text{patch}}} W_{ij}^{\text{patch}} x_{\alpha;t;j} & \text{for } t = 1, \ldots, T \,, \end{cases} \tag{1.8}$$

for each input α [14] (see also Ref. [19]). Another variant is to have a stack of several convolutional layers [14, 20]. While we'll focus on the stem block as described in the main text, it is straightforward to account for these variants.

B. Layer Normalization

In general, a layer-normalization operation [16] is defined as

$$s_{\alpha;t;i}^{(\ell)} = \gamma_i^{(\ell)} \left[\frac{z_{\alpha;t;i}^{(\ell)} - \left(\frac{1}{n} \sum_{j=1}^n z_{\alpha;t;j}^{(\ell)} \right)}{\sqrt{\frac{1}{n} \sum_{j=1}^n \left(z_{\alpha;t;j}^{(\ell)} \right)^2 - \left(\frac{1}{n} \sum_{j=1}^n z_{\alpha;t;j}^{(\ell)} \right)^2 + \epsilon}} \right] + \beta_i^{(\ell)} \,, \qquad (1.10)$$

with the regularization parameter ϵ and element-wise affine parameters $\gamma_i^{(\ell)}$ and $\beta_i^{(\ell)}$. For simplicity, in this note, we'll keep these element-wise affine parameters fixed at $\gamma_i^{(\ell)} = 1$ and $\beta_i^{(\ell)} = 0$, that is,

$$s_{\alpha;t;i}^{(\ell)} = \text{LN}_i \left(z_{\alpha;t}^{(\ell)} \right) = \frac{z_{\alpha;t;i}^{(\ell)} - \left(\frac{1}{n} \sum_{j=1}^n z_{\alpha;t;j}^{(\ell)} \right)}{\sqrt{\frac{1}{n} \sum_{j=1}^n \left(z_{\alpha;t;j}^{(\ell)} \right)^2 - \left(\frac{1}{n} \sum_{j=1}^n z_{\alpha;t;j}^{(\ell)} \right)^2 + \epsilon}} \,. \qquad (1.11)$$

In particular, when $\epsilon = 0$ each signal is normalized exactly to unity as $(1/n) \sum_{i=1}^n \left(s_{\alpha;t;i}^{(\ell)} \right)^2 = 1$ and nearly to unity for sufficiently small ϵ.

C. Multi-Head Self-Attention Block

A residual path of a multi-head self-attention (MHSA) block takes in a T-sequenced n-dimensional signal $s_{t;i}$ and outputs another (T, n)-dimensional tensor $r_{t;i}$, acting nontrivially in the sequence direction t. Specifically – suppressing the block index ℓ – the standard MHSA-residual-block operation with H heads,

$$r_{\alpha;t;i} = \text{R}_{t;i}^{\text{MHSA}}(s_\alpha; Q, K, V, U) \,, \qquad (1.12)$$

is defined as follows.

1. Define query, key, and value vectors as

$$q_{\alpha;t;c}^h \equiv \sum_{i=1}^n Q_{ci}^h s_{\alpha;t;i} \,, \qquad (1.13)$$

$$k_{\alpha;t;c}^h \equiv \sum_{i=1}^n K_{ci}^h s_{\alpha;t;i} \,, \qquad (1.14)$$

$$v_{\alpha;t;c}^h \equiv \sum_{i=1}^n V_{ci}^h s_{\alpha;t;i} \,, \qquad (1.15)$$

for $t = 1, \ldots, T$, $c = 1, \ldots, C$, and $h = 1, \ldots, H$, where the number of channels per head, $C \equiv n/H$, must be an integer.

2. Define the query–key dot product as

$$\widetilde{\Omega}^h_{\alpha;tt'} \equiv \frac{1}{\sqrt{C}} \sum_{c=1}^{C} q^h_{\alpha;t;c} k^h_{\alpha;t';c} = \frac{1}{\sqrt{C}} \sum_{c=1}^{C} \sum_{i_1,i_2=1}^{n} Q^h_{ci_1} K^h_{ci_2} s_{\alpha;t;i_1} s_{\alpha;t';i_2} \,. \tag{1.16}$$

(We'll justify the factor of $1/\sqrt{C}$ first in §I 1 C and then again in §I 3 C.)

3. Define the self-attention matrix $\Omega^h_{\alpha;tt'}$ as softmax based on the query–key dot product. Specifically, for *bidirectional* MHSA blocks used in *encoders*, we typically set

$$\Omega^h_{\alpha;tt'} \equiv \frac{\exp\left(\widetilde{\Omega}^h_{\alpha;tt'}\right)}{\sum_{t''=1}^{T} \exp\left(\widetilde{\Omega}^h_{\alpha;tt''}\right)} \,, \tag{1.17}$$

while, for *masked* MHSA blocks used in *decoders*, we typically set

$$\Omega^h_{\alpha;tt'} \equiv \begin{cases} \frac{\exp\left(\widetilde{\Omega}^h_{\alpha;tt'}\right)}{\sum_{t''=1}^{t} \exp\left(\widetilde{\Omega}^h_{\alpha;tt''}\right)} & \text{for } t' \leq t \,, \\[2ex] 0 & \text{for } t' > t \,, \end{cases} \tag{1.18}$$

that causally masks the future tokens from queriers.[5]

4. Weigh the value vector with the self-attention matrix in the token direction for each head, and then tie H heads and C channels with the unifying – a.k.a. `out_proj` – matrix U^h_{ic} as

$$r_{\alpha;t;i} \equiv \sum_{h=1}^{H} \sum_{c=1}^{C} U^h_{ic} \left(\sum_{t'=1}^{T} \Omega^h_{\alpha;tt'} v^h_{\alpha;t';c} \right) = \sum_{h=1}^{H} \sum_{t'=1}^{T} \Omega^h_{\alpha;tt'} \sum_{c=1}^{C} \sum_{j=1}^{n} U^h_{ic} V^h_{cj} s_{\alpha;t';j} \,. \tag{1.19}$$

D. Multilayer Perceptron Block

A residual path of a multilayer perceptron (MLP) block in Transformers often takes the form

$$w_{\alpha;t;i} = \sum_{j=1}^{n} W_{ij} s_{\alpha;t;j} \qquad \text{for} \quad i = 1, \ldots, Mn \,, \tag{1.20}$$

$$r_{\alpha;t;i} = \sum_{j=1}^{Mn} X_{ij} \sigma\left(w_{\alpha;t;j}\right) \quad \text{for} \quad i = 1, \ldots, n \,, \tag{1.21}$$

where σ is an activation function and M is an MLP multiplier, typically set to $M = 4$ for sufficiently-large Transformers.[6]

---

[5] In general, the self-attention matrix can be any order-one matrix-valued function of the query–key dot product, $\Omega^h_{\alpha;tt'} = \mathcal{F}_{tt'}\left(\widetilde{\Omega}^h_{\alpha}\right)$, that maps a T-by-T matrix to a T-by-T matrix. Indeed, in this note, we won't use any particular property of the function $\mathcal{F}$ beyond that it is of order one.

[6] For some variant of the MLP block, see, e.g., Ref. [21].

E. Head Block

Vision: linear classification layer

For Vision Transformers, the head block often simply consists of a standard linear layer as

$$z_{\alpha;t;i}^{(L)} = b_i^{\text{head}} + \sum_{j=1}^{n} W_{ij}^{\text{head}} s_{\alpha;t;j}^{(L-1)} \quad \text{for} \quad i = 1, \ldots, n_{\text{out}}, \tag{1.22}$$

where, e.g., for image classification tasks, $n_{\text{out}} \equiv n_{\text{class}}$ is the number of the classes.[7]

Language: word embedding, transposed (and rescaled)

For Language Transformers, we typically tie weights between the stem and head blocks [22], which means that we multiply the signals by a *transpose* of word-embedding parameters as

$$z_{\alpha;t;i}^{(L)} = \mathcal{N}_{\text{rescale}} \sum_{j=1}^{n} (W^{\text{WE}})^{\top}{}_{ij} s_{\alpha;t;j}^{(L-1)} = \mathcal{N}_{\text{rescale}} \sum_{j=1}^{n} W_{ji}^{\text{WE}} s_{\alpha;t;j}^{(L-1)} \quad \text{for} \quad i = 1, \ldots, n_{\text{out}}, \tag{1.24}$$

with $n_{\text{out}} \equiv n_{\text{vocab}}$. Here, we've also introduced the rescaling factor $\mathcal{N}_{\text{rescale}}$, with the reason of its existence – and its width scaling – to be elucidated first in §I 1 E and vindicated again in §I 3 E.

---

[7] In one implementation – which we'll follow in our Vision-Transformer experiments in §II 1 – the actual outputs are given by mean-pooling in the token direction as

$$\widetilde{z}_{\alpha;i}^{(L)} = \frac{1}{T} \sum_{t=1}^{T} z_{\alpha;t;i}^{(L)} = b_i^{\text{head}} + \sum_{j=1}^{n} W_{ij}^{\text{head}} \left(\frac{1}{T} \sum_{t=1}^{T} s_{\alpha;t;j}^{(L-1)} \right) = b_i^{\text{head}} + \sum_{j=1}^{n} W_{ij}^{\text{head}} \widetilde{s}_{\alpha;j}^{(L-1)}, \tag{1.23}$$

where $\widetilde{s}_{\alpha;j}^{(L-1)} \equiv (1/T) \sum_{t=1}^{T} s_{\alpha;t;j}^{(L-1)}$. In another implementation – which we mentioned in footnote 4 but won't follow – the class token is attached at $t = 0$ in the stem block and the output would be given by $\widetilde{z}_{\alpha;i}^{(L)} = z_{\alpha;0;i}^{(L)}$.

1. Statistics of Preactivations

Before any training happens, the model parameters need to be initialized in some way. Typically, they are initialized by independently and identically drawing them from *mean-zero* uniform, normal, or truncated-normal distributions. For notational housekeeping, let us stipulate up front that their covariances should be scaled as

$$\mathbb{E}\left[W_{i_1 j_1}^{\text{patch}} W_{i_2 j_2}^{\text{patch}}\right] = \left(\frac{C_{\text{patch}}}{n_{\text{patch}}}\right) \delta_{i_1 i_2} \delta_{j_1 j_2} \qquad \Big\| \qquad \mathbb{E}\left[W_{i_1 j_1}^{\text{WE}} W_{i_2 j_2}^{\text{WE}}\right] = \left(C_{\text{WE}}\right) \delta_{i_1 i_2} \delta_{j_1 j_2}, \quad (1.25)$$

$$\mathbb{E}\left[b_{t_1;i_1}^{\text{PE}} b_{t_2;i_2}^{\text{PE}}\right] = \left(C_{\text{PE}}\right) \delta_{t_1 t_2} \delta_{i_1 i_2}, \tag{1.26}$$

$$\mathbb{E}\left[Q_{c_1 i_1}^{h_1} Q_{c_2 i_2}^{h_2}\right] = \left(\frac{C_Q}{n}\right) \delta_{c_1 c_2} \delta_{i_1 i_2} \delta^{h_1 h_2}, \tag{1.27}$$

$$\mathbb{E}\left[K_{c_1 i_1}^{h_1} K_{c_2 i_2}^{h_2}\right] = \left(\frac{C_K}{n}\right) \delta_{c_1 c_2} \delta_{i_1 i_2} \delta^{h_1 h_2}, \tag{1.28}$$

$$\mathbb{E}\left[V_{c_1 i_1}^{h_1} V_{c_2 i_2}^{h_2}\right] = \left(\frac{C_V}{n}\right) \delta_{c_1 c_2} \delta_{i_1 i_2} \delta^{h_1 h_2}, \tag{1.29}$$

$$\mathbb{E}\left[U_{i_1 c_1}^{h_1} U_{i_2 c_2}^{h_2}\right] = \left(\frac{C_U}{n}\right) \delta_{i_1 i_2} \delta_{c_1 c_2} \delta^{h_1 h_2}, \tag{1.30}$$

$$\mathbb{E}\left[W_{i_1 j_1} W_{i_2 j_2}\right] = \left(\frac{C_W}{n}\right) \delta_{i_1 i_2} \delta_{j_1 j_2}, \tag{1.31}$$

$$\mathbb{E}\left[X_{i_1 j_1} X_{i_2 j_2}\right] = \left(\frac{C_X}{Mn}\right) \delta_{i_1 i_2} \delta_{j_1 j_2}, \tag{1.32}$$

$$\mathbb{E}\left[W_{i_1 j_1}^{\text{head}} W_{i_2 j_2}^{\text{head}}\right] = \left(\frac{C_{\text{head}}}{n}\right) \delta_{i_1 i_2} \delta_{j_1 j_2}, \quad b_i^{\text{head}} = 0 \qquad \Big\| \qquad \mathcal{N}_{\text{rescale}} = \sqrt{\frac{1}{n}}, \tag{1.33}$$

where $\mathbb{E}\left[\cdot\right]$ denotes an expectation value with respect to the initialization distribution, δ_{ij} is the Kronecker delta (i.e., $\delta_{ij} = 1$ when $i = j$ and $\delta_{ij} = 0$ when $i \neq j$), and all the initialization hyperparameters C_G are order-one numbers for each group G of model parameters. Here, by "order-one," we mean that these hyperparameters are fixed when we scale up the width n, depth L, and any other architecture hyperparameters of the networks; we'll also often say "observables stay of order one" to mean that the said observable numbers don't blow up to infinity or vanish down to zero as networks are scaled up indefinitely.[8]

---

[8] More generally, if we'd like to employ generic meta-principled scaling strategies [23] that interpolate the neural-tangent scaling strategy [9] (at $s = 0$) and maximal-update scaling strategy [24] (at $s = 1$), then the head scalings (1.33) should be modified to

$$\mathbb{E}\left[W_{i_1 j_1}^{\text{head}} W_{i_2 j_2}^{\text{head}}\right] = \left(\frac{C_{\text{head}}}{n^{1+s}}\right) \delta_{i_1 i_2} \delta_{j_1 j_2}, \quad b_i^{\text{head}} = 0 \qquad \Big\| \qquad \mathcal{N}_{\text{rescale}} = \sqrt{\frac{1}{n^{1+s}}}, \tag{1.34}$$

In the main text, we'll focus on the neural-tangent scaling strategy with $s = 0$, and we'll mention the corresponding meta-change for learning-rate factors only in future footnote 26.

For the rest of this section, we'll see how these width scalings of initialization hyperparameters (1.25)–(1.33) are chosen to ensure that the preactivations $z^{(\ell)}_{\alpha;t;i}$ stay of order one as we widen Transformers. More specifically, we'll recursively show that their covariances take the form

$$\mathbb{E}\left[z^{(\ell)}_{\alpha_1;t_1;i_1} z^{(\ell)}_{\alpha_2;t_2;i_2}\right] = \delta_{i_1 i_2} G^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}, \tag{1.35}$$

with order-one kernels $G^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}$ and we'll also show in §I 1 B that the expected squared norms of signals right after the layer-normalization operation are given by

$$F^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)} \equiv \mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n s^{(\ell)}_{\alpha_1;t_1;i} s^{(\ell)}_{\alpha_2;t_2;i}\right] = \frac{G^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}}{\sqrt{G^{(\ell)}_{(\alpha_1;t_1)(\alpha_1;t_1)} + \epsilon}\sqrt{G^{(\ell)}_{(\alpha_2;t_2)(\alpha_2;t_2)} + \epsilon}} + O\left(\frac{1}{n}\right), \tag{1.36}$$

at leading order in $1/n$. All in all, we'll show that, with the stipulated initialization hyperparameter scalings (1.25)–(1.33), Transformers respect the principle of criticality [15], that is, signals stay of order one on forward paths.[9]

With those goals in mind, let's analyze the statistics of preactivations forwardly, block by block.

A.   Stem Block

Vision: patchify embedding and positional embedding

In the stem block of Vision Transformers, which outputs the first-block preactivations $z^{(1)}_{\alpha;t;i} = b^{\mathrm{PE}}_{t;i} + \sum_{j=1}^{n_{\mathrm{patch}}} W^{\mathrm{patch}}_{ij} x_{\alpha;t;j}$, patchify weights and positional-embedding parameters are initialized with mean-zero distributions whose covariances are scaled as

$$\mathbb{E}\left[W^{\mathrm{patch}}_{i_1 j_1} W^{\mathrm{patch}}_{i_2 j_2}\right] = \left(\frac{C_{\mathrm{patch}}}{n_{\mathrm{patch}}}\right)\delta_{i_1 i_2}\delta_{j_1 j_2}, \quad \mathbb{E}\left[b^{\mathrm{PE}}_{t_1;i_1} b^{\mathrm{PE}}_{t_2;i_2}\right] = (C_{\mathrm{PE}})\,\delta_{t_1 t_2}\delta_{i_1 i_2}, \tag{1.37}$$

with order-one initialization hyperparameters C_{patch} and C_{PE}, respectively. This way, we have the vanishing mean,

$$\mathbb{E}\left[z^{(1)}_{\alpha;t;i}\right] = 0, \tag{1.38}$$

---

[9] That is, at a meta level. If there were no layer normalization, then we would have further followed the non-meta principle of criticality [10–12, 15] to fine-tune initialization hyperparameters C_G's and (initial) element-wise affine parameters $(\gamma^{(\ell)}_i, \beta^{(\ell)}_i, \xi^{(\ell)}_i)$ so as to avoid exponentially exploding or vanishing signal problems; since there typically *are* normalization layers, we'll see in §I 1 C and §I 1 D that the kernel $G^{(\ell)}$ grows only linearly with depth ℓ – until it finally gets further normalized right before entering the head block at $\ell = L$ (§I 1 E).

and order-one covariance,

$$\mathbb{E}\left[z_{\alpha_1;t_1;i_1}^{(1)}z_{\alpha_2;t_2;i_2}^{(1)}\right] = \mathbb{E}\left[\left(b_{t_1;i_1}^{\text{PE}} + \sum_{j_1=1}^{n_{\text{patch}}} W_{i_1j_1}^{\text{patch}}x_{\alpha_1;t_1;j_1}\right)\left(b_{t_2;i_2}^{\text{PE}} + \sum_{j_2=1}^{n_{\text{patch}}} W_{i_2j_2}^{\text{patch}}x_{\alpha_2;t_2;j_2}\right)\right] \quad (1.39)$$

$$= C_{\text{PE}}\delta_{i_1i_2}\delta_{t_1t_2} + \delta_{i_1i_2}C_{\text{patch}}\left(\frac{1}{n_{\text{patch}}}\sum_{j=1}^{n_{\text{patch}}} x_{\alpha_1;t_1;j}x_{\alpha_2;t_2;j}\right)$$

$$= \delta_{i_1i_2}\left(C_{\text{PE}}\delta_{t_1t_2} + C_{\text{patch}}G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)}\right)$$

$$= \delta_{i_1i_2}G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(1)}.$$

Here, in the penultimate line, we introduced the input kernel,

$$G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)} \equiv \frac{1}{n_{\text{patch}}}\sum_{j=1}^{n_{\text{patch}}} x_{\alpha_1;t_1;j}x_{\alpha_2;t_2;j}, \quad (1.40)$$

which is of order one since all pixel values are usually preprocessed to be of order one, and in the last line, we introduced the stem-block kernel,

$$G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(1)} \equiv C_{\text{patch}}G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)} + C_{\text{PE}}\delta_{t_1t_2}, \quad (1.41)$$

which is consequently of order one. In particular, each component of the first-block preactivations is expected to be of order one: $\mathbb{E}\left[\left(z_{\alpha;t;i}^{(1)}\right)^2\right] = G_{(\alpha;t)(\alpha;t)}^{(1)} = C_{\text{patch}}\left(\frac{1}{n_{\text{patch}}}\sum_{j=1}^{n_{\text{patch}}} x_{\alpha;t;j}^2\right) + C_{\text{PE}} = O(1)$.

Language: word embedding and positional embedding

In the stem block of Language Transformers, which outputs the first-block preactivations $z_{\alpha;t;i}^{(1)} = b_{t;i}^{\text{PE}} + \sum_{j=1}^{n_{\text{vocab}}} W_{ij}^{\text{WE}}x_{\alpha;t;j}$, the positional-embedding parameters are initialized as in the vision case with zero mean and the covariance $\mathbb{E}\left[b_{t_1;i_1}^{\text{PE}}b_{t_2;i_2}^{\text{PE}}\right] = C_{\text{PE}}\delta_{t_1t_2}\delta_{i_1i_2}$, while the word-embedding parameters are initialized with zero mean and the covariance

$$\mathbb{E}\left[W_{i_1j_1}^{\text{WE}}W_{i_2j_2}^{\text{WE}}\right] = (C_{\text{WE}})\,\delta_{i_1i_2}\delta_{j_1j_2}, \quad (1.42)$$

where the initialization hyperparameter C_{WE} is of order one. Here, note that we did *not* divide the covariance by the vocabulary size n_{vocab}; this way, we have

$$\mathbb{E}\left[z_{\alpha_1;t_1;i_1}^{(1)}z_{\alpha_2;t_2;i_2}^{(1)}\right] = \mathbb{E}\left[\left(b_{t_1;i_1}^{\text{PE}} + \sum_{j_1=1}^{n_{\text{vocab}}} W_{i_1j_1}^{\text{WE}}x_{\alpha_1;t_1;j_1}\right)\left(b_{t_2;i_2}^{\text{PE}} + \sum_{j_2=1}^{n_{\text{vocab}}} W_{i_2j_2}^{\text{WE}}x_{\alpha_2;t_2;j_2}\right)\right]$$

$$= \delta_{i_1i_2}\left[C_{\text{PE}}\delta_{t_1t_2} + C_{\text{WE}}\left(\sum_{j=1}^{n_{\text{vocab}}} x_{\alpha_1;t_1;j}x_{\alpha_2;t_2;j}\right)\right]$$

$$= \delta_{i_1i_2}\left[C_{\text{PE}}\delta_{t_1t_2} + C_{\text{WE}}G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)}\right]$$

$$= \delta_{i_1i_2}G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(1)}, \quad (1.43)$$

where the input kernel

$$G^{(0)}_{(\alpha_1;t_1)(\alpha_2;t_2)} \equiv \sum_{j=1}^{n_{\text{vocab}}} x_{\alpha_1;t_1;j} x_{\alpha_2;t_2;j} \,, \tag{1.44}$$

is of order one due to the one-hot structure of word inputs, and the stem-block kernel

$$G^{(1)}_{(\alpha_1;t_1)(\alpha_2;t_2)} \equiv C_{\text{WE}} G^{(0)}_{(\alpha_1;t_1)(\alpha_2;t_2)} + C_{\text{PE}} \delta_{t_1 t_2} \,, \tag{1.45}$$

is consequently of order one as well. In particular the input kernel $G^{(0)}_{(\alpha_1;t_1)(\alpha_2;t_2)}$ is one if and only if $(\alpha_1;t_1)$ and $(\alpha_2;t_2)$ correspond to the same token in the vocabulary, and zero otherwise.

B. Layer Normalization

Here, we'll analyze how each layer-normalization operation (1.11) transforms the kernel. Before doing so, we note that layer normalization operates only over the embedding direction i and in particular acts independently on each pair of sample–token indices $(\alpha;t)$, so, to declutter our analysis, we'll drop these latter indices along with the block index ℓ for now; we'll add these indices back after all the dust settles. With those in mind, we'll be analyzing the baby version of the object (1.36),

$$F \equiv \mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n s_i s_i\right] = \frac{\left(\frac{1}{n}\sum_{i=1}^n z_i^2\right) - \left(\frac{1}{n}\sum_{i=1}^n z_i\right)^2}{\left[\sqrt{\left(\frac{1}{n}\sum_{i=1}^n z_i^2\right) - \left(\frac{1}{n}\sum_{i=1}^n z_i\right)^2 + \epsilon}\,\right]^2} \,, \tag{1.46}$$

for a vectorial random variable z_i, with zero mean and order-one covariance

$$\mathbb{E}\left[z_{i_1} z_{i_2}\right] = \delta_{i_1 i_2} G \,. \tag{1.47}$$

Intuitively speaking, for sufficiently large n, we should have $\frac{1}{n}\sum_{i=1}^n z_i^2 \approx G$ and – as z_i's are mean-zero order-one random numbers – $\frac{1}{n}\sum_{i=1}^n z_i \sim \frac{\sqrt{n}}{n} \approx 0$. To make precise this intuition, let us introduce two types of finite-n corrections: the instantiation-to-instantiation magnitude fluctuations,

$$\widehat{\Delta G} \equiv \left(\frac{1}{n}\sum_{i=1}^n z_i^2\right) - G \,, \tag{1.48}$$

and the square of the mean of the preactivations,

$$\widehat{\nabla G} \equiv \left(\frac{1}{n}\sum_{i=1}^n z_i\right)^2 \,. \tag{1.49}$$

With these definitions, we can concisely rewrite the baby object (1.46) as

$$F = \frac{G + \widehat{\Delta G} - \widehat{\nabla G}}{\left[\sqrt{G + \epsilon + \widehat{\Delta G} - \widehat{\nabla G}}\right]^2} \cdot \tag{1.50}$$

We'll now show that

$$\mathbb{E}\left[\widehat{\Delta G}^p \widehat{\nabla G}^q\right] = O\left(\frac{1}{n}\right) \tag{1.51}$$

for any integers $p, q \geq 0$ with $p + q > 0$ so that, at leading order, we can neglect these finite-n corrections.[10]

Let's slowly roll out the expression in eight steps (a wordy description follows these equation):

$$\mathbb{E}\left[\widehat{\Delta G}^p \widehat{\nabla G}^q\right] \tag{1.52}$$

$$= \mathbb{E}\left\{\left[\sum_{r=0}^{p} \binom{p}{r} \frac{(-G)^{p-r}}{n^r} \sum_{i_1,\ldots,i_r=1}^{n} z_{i_1}^2 \cdots z_{i_r}^2\right]\left[\frac{1}{n^{2q}} \sum_{j_1,\ldots,j_{2q}=1}^{n} z_{j_1} \cdots z_{j_{2q}}\right]\right\}$$

$$= \sum_{r=0}^{p} \binom{p}{r} (-G)^{p-r} \frac{1}{n^{r+2q}} \sum_{i_1,\ldots,i_r=1}^{n} \sum_{j_1,\ldots,j_{2q}=1}^{n} \mathbb{E}\left[z_{i_1}^2 \cdots z_{i_r}^2 z_{j_1} \cdots z_{j_{2q}}\right]$$

$$= \sum_{r=0}^{p} \binom{p}{r} (-G)^{p-r} \frac{n(n-1)\cdots(n-r-2q+1)}{n^{r+2q}} \mathbb{E}\left[z_1^2 \cdots z_r^2 z_{r+1} \cdots z_{r+2q}\right] + O\left(\frac{1}{n}\right)$$

$$= \sum_{r=0}^{p} \binom{p}{r} (-G)^{p-r} \mathbb{E}\left[z_1^2\right] \cdots \mathbb{E}\left[z_r^2\right] \mathbb{E}\left[z_{r+1}\right] \cdots \mathbb{E}\left[z_{r+2q}\right] + O\left(\frac{1}{n}\right)$$

$$= \sum_{r=0}^{p} \binom{p}{r} (-G)^{p-r} G^r 0^q + O\left(\frac{1}{n}\right)$$

$$= G^p \sum_{r=0}^{p} \binom{p}{r} (-1)^{p-r} 0^q + O\left(\frac{1}{n}\right)$$

$$= G^p (1-1)^p 0^q + O\left(\frac{1}{n}\right)$$

$$= O\left(\frac{1}{n}\right) .$$

In the first step, we explicitly wrote out the expressions for $\widehat{\Delta G}$ – using the binomial formula – and $\widehat{\nabla G}$; in the second step, we pulled out the non-random coefficients and the summation marks outside the expectation; in the third step, we separated the sum into the terms with the embedding indices $i_1, \ldots, i_r, j_1, \ldots, j_{2q}$ all distinct, which give rise to the putatively leading-order contribution written there, and the rest of the terms that have at least one coincident embedding index, which combinatorially give rise to at most $O(1/n)$ contributions; in the fourth step, we picked up the

---

[10] The result we'll derive was used in Ref. [17] on intuitive ground; what follows is a slightly more careful analysis.

leading contribution from the first term and dumped the rest into $O\left(1/n\right);$[11] in the fifth step, we simply evaluated the expectations; in the sixth step, we pulled the common factor of G^p out of the sum; in the seventh step, we used the binomial formula in reverse; and in the eighth step, we used the fact that $0^p 0^q = 0$ for any $p, q \geq 0$ with $p + q > 0$.

With all the dust settled, judiciously putting back sample–token–block indices, we get

$$F^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)} = \mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n s^{(\ell)}_{\alpha_1;t_1;i} s^{(\ell)}_{\alpha_2;t_2;i}\right] = \frac{G^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}}{\sqrt{G^{(\ell)}_{(\alpha_1;t_1)(\alpha_1;t_1)} + \epsilon}\sqrt{G^{(\ell)}_{(\alpha_2;t_2)(\alpha_2;t_2)} + \epsilon}} + O\left(\frac{1}{n}\right). \tag{1.53}$$

As a quick sanity check, when $(\alpha_1;t_1) = (\alpha_2;t_2)$ and $\epsilon = 0$, we have $F^{(\ell)}_{(\alpha;t)(\alpha;t)} = 1$, as should be the case because $\frac{1}{n}\sum_{i=1}^n s_{\alpha;t;i} s_{\alpha;t;i} = 1$ by construction. Most importantly, even if we got an exponentially large kernel $G^{(\ell)} \sim e^{\text{large}}$, they cancel between the numerator and denominator to yield order-one numbers as $F^{(\ell)} \sim \frac{e^{\text{large}}}{\sqrt{e^{\text{large}}+\epsilon}\sqrt{e^{\text{large}}+\epsilon}} = O(1)$. This way, normalization layers prevent exponentially exploding signal problems from ever happening, at least from block to block.

C. Multi-Head Self-Attention Block

As explained in §I 0 C, the residual path of the MHSA block outputs

$$r_{\alpha;t;i} = \sum_{h=1}^H \sum_{t'=1}^T \sum_{c=1}^C \sum_{j=1}^n \Omega^h_{\alpha;tt'} U^h_{ic} V^h_{cj} s_{\alpha;t';j}, \tag{1.54}$$

where the self-attention matrix $\Omega^h_{\alpha;tt'} = \Omega^h_{\alpha;tt'}\left[\widetilde{\Omega}\left(Q,K;s\right)\right]$ is an order-one function of the query–key dot product matrix $\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}$ (1.16) given by

$$\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'} \equiv \frac{1}{\sqrt{C}}\sum_{c=1}^C q^h_{\alpha;\tilde{t};c} k^h_{\alpha;\tilde{t}';c} = \frac{1}{\sqrt{C}}\sum_{c=1}^C \sum_{i_1,i_2=1}^n Q^h_{ci_1} K^h_{ci_2} s_{\alpha;\tilde{t};i_1} s_{\alpha;\tilde{t}';i_2}. \tag{1.55}$$

---

[11] To be more verbose, this fourth step involved two mini-steps: one is a simple algebra, $n(n-1)\cdots(n-r-2q+1)/n^{r+2q} = 1 + O\left(1/n\right)$; the other is to neglect correlations among distinct neurons as subleading, that is, $\mathbb{E}\left[z_1^2\cdots z_r^2 z_{r+1}\cdots z_{r+2q}\right] = \mathbb{E}\left[z_1^2\right]\cdots\mathbb{E}\left[z_r^2\right]\mathbb{E}\left[z_{r+1}\right]\cdots\mathbb{E}\left[z_{r+2q}\right] + O\left(1/n\right)$. See Ref. [15] for more on the latter.

We initialize the associated weights Q_{ci}^h, K_{ci}^h, V_{ci}^h, and U_{ic}^h by drawing them from mean-zero distributions with the covariances

$$\mathbb{E}\left[Q_{c_1i_1}^{h_1}Q_{c_2i_2}^{h_2}\right] = \left(\frac{C_Q}{n}\right)\delta_{c_1c_2}\delta_{i_1i_2}\delta^{h_1h_2}\,, \tag{1.56}$$

$$\mathbb{E}\left[K_{c_1i_1}^{h_1}K_{c_2i_2}^{h_2}\right] = \left(\frac{C_K}{n}\right)\delta_{c_1c_2}\delta_{i_1i_2}\delta^{h_1h_2}\,, \tag{1.57}$$

$$\mathbb{E}\left[V_{c_1i_1}^{h_1}V_{c_2i_2}^{h_2}\right] = \left(\frac{C_V}{n}\right)\delta_{c_1c_2}\delta_{i_1i_2}\delta^{h_1h_2}\,, \tag{1.58}$$

$$\mathbb{E}\left[U_{i_1c_1}^{h_1}U_{i_2c_2}^{h_2}\right] = \left(\frac{C_U}{n}\right)\delta_{i_1i_2}\delta_{c_1c_2}\delta^{h_1h_2}\,, \tag{1.59}$$

with order-one initialization hyperparameters C_Q, C_K, C_V, and C_U. (The associated bias parameters – when they exist – are typically initialized to zero.) To see why these are the right scalings, let's calculate the statistics of the query–key dot product (1.55) and residual-path output (1.54).

First, the mean of the query–key dot product vanishes at initialization as

$$\mathbb{E}\left[\widetilde{\Omega}_{\alpha;tt'}^h\right] = \mathbb{E}\left[\frac{1}{\sqrt{C}}\sum_{c=1}^{C}\sum_{i,j=1}^{n}Q_{ci}^h K_{cj}^h s_{\alpha;t;i}s_{\alpha;t';j}\right] \tag{1.60}$$

$$= \frac{1}{\sqrt{C}}\sum_{c=1}^{C}\sum_{i,j=1}^{n}\mathbb{E}\left[Q_{ci}^h\right]\mathbb{E}\left[K_{cj}^h\right]\mathbb{E}\left[s_{\alpha;t;i}s_{\alpha;t';j}\right] = 0\,,$$

since these weights are independently drawn from mean-zero distributions. As for the covariance,

$$\mathbb{E}\left[\widetilde{\Omega}_{\alpha_1;t_1t_1'}^{h_1}\widetilde{\Omega}_{\alpha_2;t_2t_2'}^{h_2}\right] = \frac{1}{C}\sum_{c_1,c_2=1}^{C}\sum_{i_1,i_2,j_1,j_2=1}^{n}\mathbb{E}\left[Q_{c_1i_1}^{h_1}K_{c_1j_1}^{h_1}s_{\alpha_1;t_1;i_1}s_{\alpha_1;t_1';j_1}Q_{c_2i_2}^{h_2}K_{c_2j_2}^{h_2}s_{\alpha_2;t_2;i_2}s_{\alpha_2;t_2';j_2}\right]$$

$$= \frac{1}{C}\sum_{c_1,c_2=1}^{C}\sum_{i_1,i_2,j_1,j_2=1}^{n}\frac{C_Q}{n}\frac{C_K}{n}\delta_{c_1c_2}\delta_{i_1i_2}\delta_{j_1j_2}\delta^{h_1h_2}\mathbb{E}\left[s_{\alpha_1;t_1;i_1}s_{\alpha_1;t_1';j_1}s_{\alpha_2;t_2;i_2}s_{\alpha_2;t_2';j_2}\right]$$

$$= \delta^{h_1h_2}C_QC_K\mathbb{E}\left[\left(\frac{1}{n}\sum_{i=1}^{n}s_{\alpha_1;t_1;i}s_{\alpha_2;t_2;i}\right)\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;t_1';j}s_{\alpha_2;t_2';j}\right)\right]$$

$$= \delta^{h_1h_2}C_QC_K F_{(\alpha_1;t_1)(\alpha_2;t_2)}F_{(\alpha_1;t_1')(\alpha_2;t_2')} + O\left(\frac{1}{n}\right)\,, \tag{1.61}$$

where in the last step we used our eightfold result (1.52) to truncate away the $1/n$ corrections. Together, we expect $\widetilde{\Omega}_{tt'}^h$ to have zero mean and order-one covariance. Here, the rescaling of the query–key dot product with $1/\sqrt{C} = \sqrt{H/n}$ was critical as otherwise the covariance would have scaled nontrivially with the width n and/or the number of heads H.[12] Consequently, we expect

---

[12] The intuition behind this lengthy math – that the sum over C mean-zero order-one random numbers scales like $\sqrt{C}$ – was explained in the footnote 4 of the original Transformer paper [13]. In contrast, the Tensor Program V [25] puts forth a different scaling strategy – $1/C$ instead of $1/\sqrt{C}$ – based on heuristics that query and key vectors might get correlated at some point in training.

the self-attention matrix to be of order one.[13]

Second, turning our attention to the residual-path output, it has vanishing mean at initialization,

$$\mathbb{E}\left[r_{\alpha;t;i}\right] = \sum_{h=1}^{H}\sum_{t'=1}^{T}\sum_{c=1}^{C}\sum_{j=1}^{n}\mathbb{E}\left[U_{ic}^{h}\right]\mathbb{E}\left[V_{cj}^{h}\right]\mathbb{E}\left[\Omega_{\alpha;tt'}^{h}s_{\alpha;j;t'}\right] = 0\,, \tag{1.62}$$

and order-one covariance,

$$\mathbb{E}\left[r_{\alpha_1;t_1;i_1}r_{\alpha_2;t_2;i_2}\right] \tag{1.63}$$

$$= \sum_{h_1,h_2=1}^{H}\sum_{t_1',t_2'=1}^{T}\sum_{c_1,c_2=1}^{C}\sum_{j_1,j_2=1}^{n}\mathbb{E}\left[\Omega_{\alpha_1;t_1t_1'}^{h_1}U_{i_1c_1}^{h_1}V_{c_1j_1}^{h_1}s_{\alpha_1;t_1';j_1}\Omega_{\alpha_2;t_2t_2'}^{h_2}U_{i_2c_2}^{h_2}V_{c_2j_2}^{h_2}s_{\alpha_2;t_2';j_2}\right]$$

$$= \sum_{h_1,h_2=1}^{H}\sum_{t_1',t_2'=1}^{T}\sum_{c_1,c_2=1}^{C}\sum_{j_1,j_2=1}^{n}\frac{C_U}{n}\frac{C_V}{n}\delta_{c_1c_2}\delta_{i_1i_2}\delta_{j_1j_2}\delta^{h_1h_2}\mathbb{E}\left[\Omega_{\alpha_1;t_1t_1'}^{h_1}s_{\alpha_1;t_1';j_1}\Omega_{\alpha_2;t_2t_2'}^{h_2}s_{\alpha_2;t_2';j_2}\right]$$

$$= \delta_{i_1i_2}C_UC_V\frac{C}{n^2}\sum_{h=1}^{H}\sum_{t_1',t_2'=1}^{T}\sum_{j=1}^{n}\mathbb{E}\left[\Omega_{\alpha_1;t_1t_1'}^{h}\Omega_{\alpha_2;t_2t_2'}^{h}s_{\alpha_1;t_1';j}s_{\alpha_2;t_2';j}\right]$$

$$= \delta_{i_1i_2}C_UC_V\sum_{t_1',t_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\Omega_{\alpha_1;t_1t_1'}^{h}\Omega_{\alpha_2;t_2t_2'}^{h}\right)\left(\frac{1}{n}\sum_{i=1}^{n}s_{\alpha_1;t_1';i}s_{\alpha_2;t_2';i}\right)\right]\,.$$

Here, the term in each pair of the parentheses is expected to be of order one.[14]

Overall, adding the skip path back in and inking the block indices, we have the following kernel transformation through the MHSA block:

$$G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(\ell+1)} = G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(\ell)} \tag{1.65}$$

$$+ C_UC_V\sum_{t_1',t_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\Omega_{\alpha_1;t_1t_1'}^{h}\Omega_{\alpha_2;t_2t_2'}^{h}\right)\left(\frac{1}{n}\sum_{i=1}^{n}s_{\alpha_1;t_1';i}^{(\ell)}s_{\alpha_1;t_2';i}^{(\ell)}\right)\right]\,.$$

In particular, we expect the second additive piece to be order-one functional of the order-one layer-normalized kernel $F^{(\ell)}$ (1.53) and thus expect the overall kernel $G^{(\ell)}$ to grow linearly.

---

[13] See Appendix A for more fun with the statistics of the self-attention.

[14] To be more specific, we can use the results of Appendix A to factor the expectation as

$$\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\Omega_{\alpha_1;t_1t_1'}^{h}\Omega_{\alpha_2;t_2t_2'}^{h}\right)\left(\frac{1}{n}\sum_{i=1}^{n}s_{\alpha_1;t_1';i}s_{\alpha_2;t_2';i}\right)\right] \tag{1.64}$$

$$= \frac{1}{H}\sum_{h=1}^{H}\left(\mathbb{E}\left[\Omega_{\alpha_1;t_1t_1'}^{h}\Omega_{\alpha_2;t_2t_2'}^{h}\right]\right)F_{(\alpha_1;t_1')(\alpha_2;t_2')} + O\left(\frac{1}{C}\right)\,,$$

and further express the remaining expectation as a $\left(|\mathcal{D}|T^2\right)$-dimensional Gaussian integral.

D. Multilayer Perceptron Block

As explained in §I0 D, the residual path of the MLP block outputs

$$r_{\alpha;t;i} = \sum_{j=1}^{Mn} X_{ij}\sigma\left(w_{\alpha;t;j}\right) = \sum_{j=1}^{Mn} X_{ij}\sigma\left(\sum_{k=1}^{n} W_{jk}s_{\alpha;t;k}\right).\tag{1.66}$$

We initialize the associated weights W_{ij} and X_{ij} by drawing them from mean-zero distributions with covariances

$$\mathbb{E}\left[W_{i_1 j_1}W_{i_2 j_2}\right] = \left(\frac{C_W}{n}\right)\delta_{i_1 i_2}\delta_{j_1 j_2},\tag{1.67}$$

$$\mathbb{E}\left[X_{i_1 j_1}X_{i_2 j_2}\right] = \left(\frac{C_X}{Mn}\right)\delta_{i_1 i_2}\delta_{j_1 j_2}.\tag{1.68}$$

Let's see why these are the right scalings.

First, in the middle layer, we get vanishing mean and order-one covariance as

$$\mathbb{E}\left[w_{\alpha;t;i}\right] = \sum_{j=1}^{n}\mathbb{E}\left[W_{ij}\right]\mathbb{E}\left[s_{\alpha;t;j}\right] = 0,\tag{1.69}$$

$$\mathbb{E}\left[w_{\alpha_1;t_1;i_1}w_{\alpha_2;t_2;i_2}\right] = \sum_{j_1,j_2=1}^{n}\mathbb{E}\left[W_{i_1 j_1}W_{i_2 j_2}s_{\alpha_1;t_1;j_1}s_{\alpha_2;t_2;j_2}\right]$$

$$= C_W\delta_{i_1 i_2}\mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;t_1;j}s_{\alpha_2;t_2;j}\right]\tag{1.70}$$

$$= \delta_{i_1 i_2}C_W F_{(\alpha_1;t_1)(\alpha_2;t_2)}.$$

Indeed, at leading order in the $1/n$ expansion, we can further show that $w_{\alpha;t;i}$ is governed by mean-zero Gaussian distributions with covariance $\delta_{i_1 i_2}C_W F_{(\alpha_1;t_1)(\alpha_2;t_2)}$ [15].[15] In particular, to evaluate various expectation values, we can use the factorization formula

$$\mathbb{E}\left[f_1(w_{\alpha_1;t_1;i_1})\cdots f_p(w_{\alpha_p;t_p;i_p})\right] = \left\langle f_1\left(\widetilde{w}_{\alpha_1;t_1}\right)\right\rangle_{C_W F}\cdots\left\langle f_p\left(\widetilde{w}_{\alpha_p;t_p}\right)\right\rangle_{C_W F} + O\left(\frac{1}{n}\right),\tag{1.73}$$

for distinct embedding indices $i_1,\ldots,i_p$ and any functions $f_1,\ldots,f_p$. Here, $\langle\cdot\rangle_K$ with the kernel $K_{(\alpha_1;t_1)(\alpha_2;t_2)}$ in general denotes a $(|\mathcal{D}|T)$-dimensional Gaussian integral

$$\left\langle f\left(\widetilde{w}\right)\right\rangle_K \equiv \frac{1}{\sqrt{\det\left(2\pi K\right)}}\int\left[\prod_{\alpha\in\mathcal{D}}\prod_{t=1}^{T}d\widetilde{w}_{\alpha;t}\right]f\left(\widetilde{w}\right)\tag{1.74}$$

$$\times\exp\left[-\frac{1}{2}\sum_{\alpha_1,\alpha_2\in\mathcal{D}}\sum_{t_1,t_2=1}^{T}\widetilde{w}_{\alpha_1;t_1}\left(K^{-1}\right)^{(\alpha_1;t_1)(\alpha_2;t_2)}\widetilde{w}_{\alpha_2;t_2}\right].$$

---

[15] What the MLP block does is extensively covered in Ref. [15] – both at infinite width and at finite width – for those interested in more details than necessary here. That said, to give a flavor of (non-)Gaussianity, let's examine the

Second, moving onto the residual-path output, we again have vanishing mean and order-one covariance as

$$\mathbb{E}\left[r_{\alpha;t;i}\right] = \sum_{j=1}^{Mn} \mathbb{E}\left[X_{ij}\right] \mathbb{E}\left[\sigma\left(w_{\alpha;t;j}\right)\right] = 0\,, \tag{1.75}$$

$$\mathbb{E}\left[r_{\alpha_1;t_1;i_1} r_{\alpha_2;t_2;i_2}\right] = C_X \mathbb{E}\left[\frac{1}{Mn} \sum_{j=1}^{Mn} \sigma\left(w_{\alpha_1;t_1;j}\right) \sigma\left(w_{\alpha_2;t_2;j}\right)\right] \tag{1.76}$$

$$= C_X \left\langle \sigma\left(\widetilde{w}_{\alpha_1;t_1}\right) \sigma\left(\widetilde{w}_{\alpha_2;t_2}\right)\right\rangle_{C_W F} + O\left(\frac{1}{n}\right),$$

where in the last step we used the factorization formula (1.73) just mentioned.

Combining two operations, adding the skip path back in, and inking the block indices, we have the kernel transformation through the MLP block

$$G^{(\ell+1)}_{(\alpha_1;t_1)(\alpha_2;t_2)} = G^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)} + C_X \left\langle \sigma\left(\widetilde{w}_{\alpha_1;t_1}\right) \sigma\left(\widetilde{w}_{\alpha_2;t_2}\right)\right\rangle_{C_W F^{(\ell)}} + O\left(\frac{1}{n}\right), \tag{1.77}$$

where the layer-normalized kernel $F^{(\ell)}$ (1.53) is given in terms of the previous-block kernel $G^{(\ell)}$. In particular, as in the case of the MHSA block, we expect the second additive piece to be of order one due to the layer normalization and thus expect the kernel to linearly grow.[16]

---

fourth moment of the middle-layer preactivations,

$$\mathbb{E}\left[w_{\alpha_1;t_1;i_1} w_{\alpha_2;t_2;i_2} w_{\alpha_3;t_3;i_3} w_{\alpha_4;t_4;i_4}\right] \tag{1.71}$$

$$= \sum_{j_1,j_2,j_3,j_4=1}^{n} \mathbb{E}\left[W_{i_1 j_1} W_{i_2 j_2} W_{i_3 j_3} W_{i_4 j_4} s_{\alpha_1;t_1;j_1} s_{\alpha_2;t_2;j_2} s_{\alpha_3;t_3;j_3} s_{\alpha_4;t_4;j_4}\right]$$

$$= C_W^2 \left\{ \delta_{i_1 i_2} \delta_{i_3 i_4} \mathbb{E}\left[\left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t_1;j} s_{\alpha_2;t_2;j}\right)\left(\frac{1}{n}\sum_{k=1}^{n} s_{\alpha_3;t_3;k} s_{\alpha_4;t_4;k}\right)\right] \right.$$

$$+ \delta_{i_1 i_3} \delta_{i_2 i_4} \mathbb{E}\left[\left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t_1;j} s_{\alpha_3;t_3;j}\right)\left(\frac{1}{n}\sum_{k=1}^{n} s_{\alpha_2;t_2;k} s_{\alpha_4;t_4;k}\right)\right]$$

$$\left. + \delta_{i_1 i_4} \delta_{i_2 i_3} \mathbb{E}\left[\left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t_1;j} s_{\alpha_4;t_4;j}\right)\left(\frac{1}{n}\sum_{k=1}^{n} s_{\alpha_2;t_2;k} s_{\alpha_3;t_3;k}\right)\right] \right\} + O\left(\frac{1}{n}\right)$$

$$= \mathbb{E}\left[w_{\alpha_1;t_1;i_1} w_{\alpha_2;t_2;i_2}\right] \mathbb{E}\left[w_{\alpha_3;t_3;i_3} w_{\alpha_4;t_4;i_4}\right] + \mathbb{E}\left[w_{\alpha_1;t_1;i_1} w_{\alpha_3;t_3;i_3}\right] \mathbb{E}\left[w_{\alpha_2;t_2;i_2} w_{\alpha_4;t_4;i_4}\right]$$

$$+ \mathbb{E}\left[w_{\alpha_1;t_1;i_1} w_{\alpha_4;t_4;i_4}\right] \mathbb{E}\left[w_{\alpha_2;t_2;i_2} w_{\alpha_3;t_3;i_3}\right] + O\left(\frac{1}{n}\right).$$

In the second equality, we used the i.i.d.-ness of the initialization distribution; specifically, the first term comes from when $(i_1, j_1) = (i_2, j_2)$ and $(i_3, j_3) = (i_4, j_4)$ and similarly for the other two terms. [For the normal distribution, this step is exact and is a special case of Wick's theorem; for non-normal distributions, there is a missing non-Gaussian term when $i_1 = i_2 = i_3 = i_4$ and $j_1 = j_2 = j_3 = j_4$, which leads to the contribution proportional to $(1/n)\delta_{i_1 i_2 i_3 i_4} = (1/n)\delta_{i_1 i_2}\delta_{i_1 i_3}\delta_{i_1 i_4}$.] In the last equality, we used our eightfold result (1.52) to truncate away the $1/n$ correction. Overall, this result on the fourth moment can be concisely summarized as

$$\mathbb{E}\left[w_{\alpha_1;t_1;i_1} w_{\alpha_2;t_2;i_2} w_{\alpha_3;t_3;i_3} w_{\alpha_4;t_4;i_4}\right]\Big|_{\text{connected}} = O\left(\frac{1}{n}\right). \tag{1.72}$$

[16] In contrast, if there were no normalization layer, then we would have instead gotten – say, e.g., for the ReLU

E.   Head Block

Vision: linear classification layer

In the head block of Vision Transformers, which outputs the last-block preactivations $z^{(L)}_{\alpha;t;i} = b^{\text{head}}_i + \sum_{j=1}^n W^{\text{head}}_{ij} s^{(L-1)}_{\alpha;t;j}$, we typically zero-initialize the head biases as $b^{\text{head}}_i = 0$ while we draw the head weights W^{head}_{ij} from a mean-zero distribution with the covariance

$$\mathbb{E}\left[W^{\text{head}}_{i_1 j_1} W^{\text{head}}_{i_2 j_2}\right] = \left(\frac{C_{\text{head}}}{n}\right)\delta_{i_1 i_2}\delta_{j_1 j_2}\,. \tag{1.78}$$

As usual, the mean of the output vanishes, while their covariance is given by

$$\mathbb{E}\left[z^{(L)}_{\alpha_1;t_1;i_1} z^{(L)}_{\alpha_2;t_2;i_2}\right] = \delta_{i_1 i_2} C_{\text{head}}\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^n s^{(L-1)}_{\alpha_1;t_1;i} s^{(L-1)}_{\alpha_2;t_2;i}\right] = \delta_{i_1 i_2} C_{\text{head}} F^{(L-1)}_{(\alpha_1;t_1)(\alpha_2;t_2)}\,. \tag{1.79}$$

That is, the output kernel

$$G^{(L)}_{(\alpha_1;t_1)(\alpha_2;t_2)} = C_{\text{head}} F^{(L-1)}_{(\alpha_1;t_1)(\alpha_2;t_2)}\,, \tag{1.80}$$

is manifestly of order one.[17] Note that even the linear growth of the kernel $G^{(\ell)}$ with the depth ℓ cancels out in this last step because of the very last normalization layer.

Language: word embedding, transposed (and rescaled)

In the head block of Language Transformers, which outputs the last-block preactivations $z^{(L)}_{\alpha;t;i} = \mathcal{N}_{\text{rescale}} \sum_{j=1}^n W^{\text{WE}}_{ji} s^{(L-1)}_{\alpha;t;j}$, we've already specified the covariance of the word-embedding parameters to be of order one as $\mathbb{E}\left[W^{\text{WE}}_{i_1 j_1} W^{\text{WE}}_{i_2 j_2}\right] = (C_{\text{WE}})\delta_{i_1 i_2}\delta_{j_1 j_2}$, so let's focus on the rescaling factor $\mathcal{N}_{\text{rescale}}$. Intuitively, the need for this rescaling factor $\mathcal{N}_{\text{rescale}}$ should be clear: since we are summing n mean-zero random order-one numbers, it calls for the rescaling by

$$\mathcal{N}_{\text{rescale}} = \sqrt{\frac{1}{n}}\,, \tag{1.82}$$

---

activation – something like $C_X \langle \sigma(\widetilde{w})\sigma(\widetilde{w})\rangle_{C_W G^{(\ell)}} \sim \sharp G^{(\ell)}$, which would have resulted in $G^{(\ell+1)} \sim (1+\sharp)G^{(\ell)}$ and hence an exponential explosion $\sim (1+\sharp)^\ell$.

[17] If we mean pool in the token direction as $\widetilde{z}^{(L)}_{\alpha;i} = \frac{1}{T}\sum_{t=1}^T z^{(L)}_{\alpha;t;i}$ (1.23), then the actual output kernel is given by

$$\widetilde{G}^{(L)}_{\alpha_1 \alpha_2} = \frac{1}{T^2}\sum_{t_1,t_2=1}^T G^{(L)}_{(\alpha_1;t_1)(\alpha_2;t_2)}\,. \tag{1.81}$$

If we attach a class token at $t=0$ and take $\widetilde{z}^{(L)}_{\alpha;i} = z^{(L)}_{\alpha;0;i}$ as the output, then the output kernel is given by $\widetilde{G}^{(L)}_{\alpha_1 \alpha_2} = G^{(L)}_{(\alpha_1;0)(\alpha_2;0)}$.

so as to maintain the outputs to be of order one.[18] Non-intuitively, we have

$$\mathbb{E}\left[z_{\alpha_1;t_1;i_1}^{(L)} z_{\alpha_2;t_2;i_2}^{(L)}\right] = \mathcal{N}_{\text{rescale}}^2 \delta_{i_1 i_2} C_{\text{WE}} \sum_{j=1}^{n} \mathbb{E}\left[s_{\alpha_1;t_1;j}^{(L-1)} s_{\alpha_2;t_2;j}^{(L-1)}\right] \tag{1.83}$$

$$= \delta_{i_1 i_2} C_{\text{WE}} \mathbb{E}\left[\frac{1}{n} \sum_{j=1}^{n} s_{\alpha_1;t_1;j}^{(L-1)} s_{\alpha_2;t_2;j}^{(L-1)}\right] = \delta_{i_1 i_2} C_{\text{WE}} F_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(L-1)},$$

where in the second equality we used the rescaling prescription (1.82). We thus again have the order-one output kernel

$$G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(L)} = C_{\text{WE}} F_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(L-1)}. \tag{1.84}$$

Actually this entire discussion neglected the subtle interlayer correlation created by weight tying between the stem and head blocks, which is discussed in the footnote here.[19]

---

[18] That said, this rescaling factor was absent in the original paper [22]. This output rescaling factor of $1/\sqrt{n}$ has been implemented in, for example, Ref. [4] (though we could not track the original reference) and the maximal-update version – $\mathcal{N}_{\text{rescale}} = 1/n$ – was first implemented in Ref. [25].

[19] The discussion of general case below in the next paragraph will be condensed and not particularly illuminating, so we first recommend working out a simple toy model consisting of three linear layers with a skip connection, $z_{\alpha;t;i}^{(3)} = \mathcal{N}_{\text{rescale}} \sum_{j,k}^{n} \sum_{m=1}^{n_{\text{vocab}}} W_{ji}^{\text{WE}} \left(W_{jk}^{(2)} + \delta_{jk}\right) W_{km}^{\text{WE}} x_{\alpha;t;m}$ where normally-distributed word-embedding weights W_{ij}^{WE} have the covariance C_{WE} and middle-block weights $W_{ij}^{(2)}$ have the covariance $C_{W^{(2)}}/n$. Then we get

$$\mathbb{E}\left[z_{\alpha_1;t_1;i_1}^{(3)} z_{\alpha_2;t_2;i_2}^{(3)}\right] \tag{1.85}$$

$$= \delta_{i_1 i_2} C_{\text{WE}}^2 \left(C_{W^{(2)}} + 1\right) G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)} + C_{\text{WE}}^2 \left(n + \frac{C_{W^{(2)}}}{n}\right) x_{\alpha_1;t_1;i_1} x_{\alpha_2;t_2;i_2} + \frac{1}{n} \left(C_{W^{(2)}} + 1\right) \left(x_{\alpha_1;t_1;i_2} x_{\alpha_2;t_2;i_1}\right).$$

Here the first term is the one discussed in the main text and the rest is largely ignorable, except the contribution $n C_{\text{WE}}^2 x_{\alpha_1;t_1;i_1} x_{\alpha_2;t_2;i_2}$ which comes from the skip path $\mathcal{N}_{\text{rescale}} \sum_{j=1}^{n} \sum_{m=1}^{n_{\text{vocab}}} W_{ji}^{\text{WE}} W_{jm}^{\text{WE}} x_{\alpha;t;m}$. At first sight, this contact term seems the most dominant contribution, but note that the main contribution with $\delta_{i_1 i_2}$ has n_{vocab} nonzero components of order one, while this contact term has only one nonzero component of order n so, as long as $n \lesssim n_{\text{vocab}}$, it is not obviously problematic. Nonetheless, if we'd like to get rid of this contact term, then the easiest solution would be to just drop the skip path in the second block.

To deal with more general cases, we can first study correlators $\mathbb{E}\left[W_{j_1 i_1}^{\text{WE}} W_{j_2 i_2}^{\text{WE}} z_{\alpha_1;t_1;k_1}^{(1)} \cdots z_{\alpha_p;t_p;k_p}^{(1)}\right]$ and show that, for any function $\mathcal{F}$ that depends on W_{km}^{WE} only through the first-block preactivations $z_{\alpha;t;k}^{(1)}$,

$$\mathbb{E}\left[W_{j_1 i_1}^{\text{WE}} W_{j_2 i_2}^{\text{WE}} \mathcal{F}\right] = \frac{1}{n} C_{\text{WE}} \delta_{i_1 i_2} \delta_{j_1 j_2} \mathbb{E}\left[\mathcal{F}\right] + \frac{1}{n} C_{\text{WE}}^2 \sum_{\tilde{\alpha}_1, \tilde{\alpha}_2 \in \mathcal{D}} \sum_{\tilde{t}_1, \tilde{t}_2 = 1}^{T} x_{\tilde{\alpha}_1;\tilde{t}_1;i_1} x_{\tilde{\alpha}_2;\tilde{t}_2;i_2} \mathbb{E}\left[\frac{\partial^2 \mathcal{F}}{\partial z_{\tilde{\alpha}_1;\tilde{t}_1;j_1}^{(1)} \partial z_{\tilde{\alpha}_2;\tilde{t}_2;j_2}^{(1)}}\right].$$

So, putting in $\mathcal{F} = s_{\alpha_1;t_1;j_1}^{(L-1)} s_{\alpha_2;t_2;j_2}^{(L-1)}$, the first term gives the naive contribution (1.83) discussed in the main text

while the second term gives rise to four more contributions

$$\frac{1}{n}C_{\mathrm{WE}}^2\sum_{\tilde{t}_1,\tilde{t}_2=1}^{T}x_{\alpha_1;\tilde{t}_1;i_1}x_{\alpha_2;\tilde{t}_2;i_2}\mathbb{E}\left[\left(\sum_{j_1=1}^{n}\frac{\partial s^{(L-1)}_{\alpha_1;t_1;j_1}}{\partial z^{(1)}_{\alpha_1;\tilde{t}_1;j_1}}\right)\left(\sum_{j_2=1}^{n}\frac{\partial s^{(L-1)}_{\alpha_2;t_2;j_2}}{\partial z^{(1)}_{\alpha_2;\tilde{t}_2;j_2}}\right)\right] \tag{1.86}$$

$$+\frac{1}{n}C_{\mathrm{WE}}^2\sum_{j_1,j_2=1}^{n}\sum_{\tilde{t}_1,\tilde{t}_2=1}^{T}x_{\alpha_1;\tilde{t}_1;i_2}x_{\alpha_2;\tilde{t}_2;i_1}\mathbb{E}\left[\frac{\partial s^{(L-1)}_{\alpha_1;t_1;j_2}}{\partial z^{(1)}_{\alpha_1;\tilde{t}_1;j_1}}\frac{\partial s^{(L-1)}_{\alpha_2;t_2;j_2}}{\partial z^{(1)}_{\alpha_2;\tilde{t}_2;j_1}}\right]$$

$$+\frac{1}{n}C_{\mathrm{WE}}^2\sum_{j_1,j_2=1}^{n}\sum_{\tilde{t}_1,\tilde{t}_2=1}^{T}x_{\alpha_1;\tilde{t}_1;i_1}x_{\alpha_1;\tilde{t}_2;i_2}\mathbb{E}\left[\frac{\partial^2 s^{(L-1)}_{\alpha_1;t_1;j_1}}{\partial z^{(1)}_{\alpha_1;\tilde{t}_1;j_1}\partial z^{(1)}_{\alpha_1;\tilde{t}_2;j_2}}s^{(L-1)}_{\alpha_2;t_2;j_2}\right]$$

$$+\frac{1}{n}C_{\mathrm{WE}}^2\sum_{j_1,j_2=1}^{n}\sum_{\tilde{t}_1,\tilde{t}_2=1}^{T}x_{\alpha_2;\tilde{t}_1;i_1}x_{\alpha_2;\tilde{t}_2;i_2}\mathbb{E}\left[s^{(L-1)}_{\alpha_1;t_1;j_1}\frac{\partial^2 s^{(L-1)}_{\alpha_2;t_2;j_2}}{\partial z^{(1)}_{\alpha_2;\tilde{t}_1;j_1}\partial z^{(1)}_{\alpha_2;\tilde{t}_2;j_2}}\right].$$

Out of these four terms, the first one contains the contact term. Specifically, to see that, we use the chain rule $\partial s^{(L-1)}_{\alpha;t;j}/\partial z^{(1)}_{\alpha;\tilde{t};j}=\sum_{k=1}^{n}\left(\partial s^{(L-1)}_{\alpha;t;j}/\partial z^{(L-1)}_{\alpha;t;k}\right)\left(\partial z^{(L-1)}_{\alpha;t;k}/\partial z^{(1)}_{\alpha;\tilde{t};j}\right)$, use the future mnemonic (1.140) to pick up the leading term $\partial s^{(L-1)}_{\alpha;t;j}/\partial z^{(L-1)}_{\alpha;t;k}\approx\delta_{jk}/\sqrt{G^{(L-1)}_{(\alpha;t)(\alpha;t)}+\epsilon}$, and then finally note that the dominant contribution to $\partial z^{(L-1)}_{\alpha;t;j}/\partial z^{(1)}_{\alpha;\tilde{t};j}$ comes from the skip path (essentially because the "j" in the j-th component of the residual-path output $\mathtt{R}^{(\ell)}_{t;j}$ has no meaningful association with the "j" in the j-th component of the first-block preactivation $z^{(1)}_{\alpha;t;j}$ and hence its contribution is diluted and is down by powers of $1/n$).

2. A Crash Course on Neural Tangent Kernels

In this short crash course, we'll give a condensed introduction to the concept of neural tangent kernels [9], both for the vanilla stochastic gradient descent (SGD) optimizer (§I 2 A) and for the AdamW optimizer (§I 2 B); for a vaporized introduction, see, e.g., Ref. [15]. This introduction paves the way for our effective-theory analysis of the backward path in §I 3, where we'll figure out how to scale a relative learning-rate factor for each group of model parameters in Transformers.

A. Vanilla SGD

The SGD update equation is given by

$$\theta_\mu(t) = \theta_\mu(t-1) - \eta_t \frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial \theta_\mu}\bigg|_{\theta=\theta(t-1)}, \tag{1.87}$$

where the model-parameter index μ runs over all the P model parameters θ_μ in the architecture, η_t is a learning rate at iteration t, $\mathcal{L}_{\mathcal{A}_t}$ denotes a loss function evaluated on a minibatch $\mathcal{A}_t$ at iteration t, and $\theta_\mu(0)$ are drawn from the initialization distribution that was extensively discussed in §I 1.[20] In this standard form, we assign the single learning rate η_t for all the model parameters, but in theory we'll soon find that the learning rate for each group G of model parameters must be scaled differently as we embiggen Transformers. [To see how model parameters in a Transformer can get partitioned into groups, look back at equations (1.25)–(1.33) in §I 1 or look ahead to equations (1.99)–(1.116) in §I 3.] To that end, we generalize the SGD update equation (1.87) to

$$\theta_\mu(t) = \theta_\mu(t-1) - \eta_t \lambda_{G(\mu)} \frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial \theta_\mu}\bigg|_{\theta=\theta(t-1)}, \tag{1.88}$$

where $G(\mu)$ denotes a group G to which the μ-th model parameter belongs, and relative learning-rate factors λ_G's can be used to balance the degrees of updates among various groups in the architecture. In particular, the standard SGD naively sets $\lambda_G = 1$ for all G's, i.e., it uses a uniform learning rate for all the model parameters, while we'll figure out how to sophisticatedly scale them with widths for all the model-parameter groups in Transformers.

To do so, we look at how the network function $f_{\alpha;i} \equiv z_{\alpha;i}^{(L)}$ gets updated (where we temporarily suppress the token index t to mitigate the notational conflict). That is, Taylor-expanding the

---

[20] We apologize to the letter t for our using it to represent both the iteration index t and the token index t.

change in the function to the first order in the model-parameter update (1.88), we get

$$
\begin{aligned}
f_{\alpha;i}(t) =& f_{\alpha;i}(t-1) - \eta_t \sum_{\mu=1}^{P} \lambda_{G(\mu)} \left(\frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial \theta_\mu} \frac{\partial f_{\alpha;i}}{\partial \theta_\mu} \right) \Bigg|_{\theta=\theta(t-1)} + \dots \\
=& f_{\alpha;i}(t-1) - \eta_t \sum_{\mu=1}^{P} \lambda_{G(\mu)} \left[\sum_{\widetilde\alpha \in \mathcal{A}_t} \sum_{j=1}^{n_{\text{out}}} \frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial f_{\widetilde\alpha;j}} \left(\frac{\partial f_{\widetilde\alpha;j}}{\partial \theta_\mu} \frac{\partial f_{\alpha;i}}{\partial \theta_\mu} \right) \right] \Bigg|_{\theta=\theta(t-1)} + \dots \\
=& f_{\alpha;i}(t-1) - \eta_t \sum_{\widetilde\alpha \in \mathcal{A}_t} \sum_{j=1}^{n_{\text{out}}} \left(\frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial f_{\widetilde\alpha;j}} H_{(\widetilde\alpha;j)(\alpha;i)} \right) \Bigg|_{\theta=\theta(t-1)} + \dots .
\end{aligned}
\tag{1.89}
$$

Here, in the last line, we've defined the neural tangent kernel as

$$
H_{(\alpha_1;i_1)(\alpha_2;i_2)} \equiv \sum_{\mu=1}^{P} \lambda_{G(\mu)} \frac{\partial f_{\alpha_1;i_1}}{\partial \theta_\mu} \frac{\partial f_{\alpha_2;i_2}}{\partial \theta_\mu} .
\tag{1.90}
$$

In particular, as the derivative of the loss with respect to the network outputs, $\partial \mathcal{L}_{\mathcal{A}_t}/\partial f_{\widetilde\alpha;j}$, is expected to be of order one for generic losses, our objective in §I 3 will be to scale relative learning-rate factors λ_G's such that the neural tangent kernel $H_{(\alpha_1;i_1)(\alpha_2;i_2)}$ – and hence the change in the network function – stays of order one as Transformers get wider.

In passing, we mention that, in principle, the higher-order terms "…" in the Taylor expansion (1.89) can matter, and the neural tangent kernel itself can dynamically change too. However, these effects are all $1/n$-suppressed and hence can be neglected at the (leading) order we are working in in this note.[21]

B.   AdamW

The update equations for AdamW [28, 29] – here again generalized with the relative learning-rate factors $\widetilde\lambda_G$'s – can be written as

$$
g^\mu(t) \equiv \frac{\partial \mathcal{L}_{\mathcal{A}_t}}{\partial \theta_\mu} \Bigg|_{\theta=\theta(t-1)} ,
\tag{1.91}
$$

$$
v^\mu(t) = \beta_1 v^\mu(t-1) + (1-\beta_1) g^\mu(t) ,
\tag{1.92}
$$

$$
u^{\mu\mu}(t) = \beta_2 u^{\mu\mu}(t-1) + (1-\beta_2) [g^\mu(t)]^2 ,
\tag{1.93}
$$

$$
\theta_\mu(t) = (1 - \eta_t \cdot \mathtt{wd}) \theta_\mu(t-1) - \widetilde\lambda_{G(\mu)} \frac{\eta_t}{\sqrt{\frac{u^{\mu\mu}(t)}{(1-\beta_2^t)}} + \epsilon} \frac{v^\mu(t)}{(1-\beta_1^t)} ,
\tag{1.94}
$$

---

[21] For a more thorough treatment of these $1/n$ effects, see Ref. [15] (see also Refs. [26, 27] for some original work). More generally, the leading perturbative corrections scale with width as $1/n^{1-s}$ [23], which reduces to $1/n$ for the neural-tangent scaling strategy at $s = 0$ while being never perturbative for the maximal-update scaling strategy at $s = 1$.

where $v^\mu(0) = u^{\mu\mu}(0) = 0$ while $\theta_\mu(0)$ are again drawn from the initialization distribution. Here, the optimizer hyperparameters β_1 and β_2 respectively set the decay rates for the running averages of the first and second moments of per-parameter gradients – roughly taking averages over the past $1/(1 - \beta_\sharp)$ iterations – while the hyperparameters η_t and $\mathtt{wd}$ are the global learning rate and weight decay, respectively. In particular, in words, the last update equation (1.94) can be viewed as adapting the *per-parameter* learning rate as $\eta\widetilde{\lambda}_{G(\mu)}/\sqrt{u^{\mu\mu}}$, i.e., according to the inverse square root of the *per-parameter* running average of the squared gradients.[22]

Let us gedanken-Taylor-expand the change in the function to the first order in the model-parameter update (1.94), as we did before for SGD. We then note that the essential difference between SGD and AdamW is the factor of $\sqrt{u^{\mu\mu}} \sim |g^\mu|$ in the denominator of the parameter update, that is, we get the (AdamW-modified) neural tangent kernel of the form

$$\widetilde{H}_{(\alpha_1;i_1)(\alpha_2;i_2)} \equiv \sum_{\mu=1}^{P} \frac{\widetilde{\lambda}_{G(\mu)}}{|g^\mu|} \frac{\partial f_{\alpha_1;i_1}}{\partial \theta_\mu} \frac{\partial f_{\alpha_2;i_2}}{\partial \theta_\mu} . \tag{1.95}$$

We'll accordingly need to scale $\widetilde{\lambda}_G$'s to keep this modified neural tangent kernel of order one so that, again, the change in the network function stays of order one.

$$\blacklozenge\blacklozenge\blacklozenge\blacklozenge$$

In our analysis in §I 3, to be economical, we'll take the following strategy to deal with λ_G for SGD and $\widetilde{\lambda}_G$ for AdamW together: for each group G of model parameters, we'll *(i)* first figure out the appropriate scaling of the relative learning-rate factor λ_G for SGD so that the neural tangent kernel (1.90) stays of order one, *(ii)* then estimate the width scaling of the expected per-parameter gradient magnitude $|g^\mu|$ for model parameters in the group G, and *(iii)* accordingly adjust the relative learning-rate factor $\widetilde{\lambda}_{G(\mu)} \sim \lambda_{G(\mu)}|g^\mu|$ for AdamW so that the modified neural tangent kernel (1.95) stays of order one.

---

[22] Let us explain other factors in the last AdamW update equation (1.94) in more detail: the dividing factors of $(1 - \beta_1^t)$ and $(1 - \beta_2^t)$ are the bias-correction terms for the moment estimators v^μ and $u^{\mu\mu}$, respectively [28], and the regularization factor of ϵ in the denominator sets the cutoff scale for gradients, that is, for the model parameter component μ with $|g^\mu| \gg \epsilon$ the update is basically $\eta \cdot \mathrm{sign}(g^\mu)$ whereas for $|g^\mu| \ll \epsilon$ the magnitude of the update starts to diminish in proportion to the magnitude of the gradient, albeit with an inflated learning rate η/ϵ.

While we are on the topic of the AdamW optimizer hyperparameters, intuitively, with the neural-tangent scalings we employ in this note, gradient information doesn't change at leading order, so we can view gradients as fixed at initialization and forget about the effects of β_1 and β_2 as far as the leading-order theoretical analysis goes. As for the weight decay, properly carrying through the exercise on the regularization proposed in the footnote of §10 in Ref. [15], we find that the product of the global learning rate and weight decay, $\eta_t \cdot \mathtt{wd}$, stays of order one: roughly, in their notation, the learning-rate tensor has lower indices as $\lambda_{\mu\nu}$ while the weight-decay tensor comes in as $\sum_{\mu,\nu=1}^{P} a^{\mu\nu}\theta_\mu\theta_\nu$ with upper indices, and $\lambda_{\mu\mu}$ scales like $1/\mathtt{fan\_in}$ while $a^{\mu\mu}$ scales like $\mathtt{fan\_in}$, so their product $\lambda_{\mu\mu}a^{\mu\mu}$ is of order one.

3. Statistics of Neural Tangent Kernels

In the previous few pages, we've identified the neural tangent kernel,

$$\widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)} \equiv \sum_{\mu=1}^{P} \lambda_{G(\mu)} \frac{\partial z_{\alpha_1;t_1;i_1}^{(L)}}{\partial \theta_\mu} \frac{\partial z_{\alpha_2;t_2;i_2}^{(L)}}{\partial \theta_\mu} \,, \tag{1.96}$$

as the central object of interest in identifying the proper scalings of per-group learning rates λ_G's. That is, the neural tangent kernel dictates the update in the network function, and we would like to keep it of order one so that it doesn't explode or vanish in scaling up the models. (Here, we've decorated the neural tangent kernel with a *hat* [15] to connote that we'll be studying the statistics – or, in this note, just the mean – of this object *at initialization* and also restored the token indices t_1, t_2 as there won't be any conflict with the iteration index t at this point.[23])

To carry out our analysis, it is useful to introduce the ℓ-th-block neural tangent kernel

$$\widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)}^{(\ell)} \equiv \sum_{\mu=1}^{P} \lambda_{G(\mu)} \frac{\partial z_{\alpha_1;t_1;i_1}^{(\ell)}}{\partial \theta_\mu} \frac{\partial z_{\alpha_2;t_2;i_2}^{(\ell)}}{\partial \theta_\mu} \,, \tag{1.97}$$

which equals the actual neural tangent kernel (1.96) when $\ell = L$. In particular, these objects enable us to carry out a *forward* recursive analysis that very much mirrors our analysis of the statistics of preactivations in §I 1. Specifically, by recalling the main forward equation (1.3) and using the chain rule to roll out derivatives, we have

$$\widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)}^{(\ell+1)} \tag{1.98}$$

$$= \sum_{\mu \in (\ell+1)\text{-th block}} \lambda_{G(\mu)} \frac{\partial \mathcal{R}_{\alpha_1;t_1;i_1}^{(\ell+1)}}{\partial \theta_\mu^{(\ell+1)}} \frac{\partial \mathcal{R}_{\alpha_2;t_2;i_2}^{(\ell+1)}}{\partial \theta_\mu^{(\ell+1)}}$$

$$+ \sum_{t_1',t_2'=1}^{T} \sum_{j_1,j_2=1}^{n} \sum_{k_1,k_2=1}^{n} \left(\frac{\partial \mathcal{R}_{\alpha_1;t_1;i_1}^{(\ell+1)}}{\partial s_{\alpha_1;t_1';j_1}^{(\ell)}} \frac{\partial \mathcal{R}_{\alpha_2;t_2;i_2}^{(\ell+1)}}{\partial s_{\alpha_2;t_2';j_2}^{(\ell)}} \right) \left(\frac{\partial s_{\alpha_1;t_1';j_1}^{(\ell)}}{\partial z_{\alpha_1;t_1';k_1}^{(\ell)}} \frac{\partial s_{\alpha_2;t_2';j_2}^{(\ell)}}{\partial z_{\alpha_2;t_2';k_2}^{(\ell)}} \right) \widehat{H}_{(\alpha_1;t_1';k_1)(\alpha_2;t_2';k_2)}^{(\ell)}$$

$$+ \widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)}^{(\ell)}$$

$$+ \sum_{t_1'=1}^{T} \sum_{j_1=1}^{n} \sum_{k_1=1}^{n} \left(\frac{\partial \mathcal{R}_{\alpha_1;t_1;i_1}^{(\ell+1)}}{\partial s_{\alpha_1;t_1';j_1}^{(\ell)}} \right) \left(\frac{\partial s_{\alpha_1;t_1';j_1}^{(\ell)}}{\partial z_{\alpha_1;t_1';k_1}^{(\ell)}} \right) \widehat{H}_{(\alpha_1;t_1';k_1)(\alpha_2;t_2;i_2)}^{(\ell)}$$

$$+ \sum_{t_2'=1}^{T} \sum_{j_2=1}^{n} \sum_{k_2=1}^{n} \left(\frac{\partial \mathcal{R}_{\alpha_2;t_2;i_2}^{(\ell+1)}}{\partial s_{\alpha_2;t_2';j_2}^{(\ell)}} \right) \left(\frac{\partial s_{\alpha_2;t_2';j_2}^{(\ell)}}{\partial z_{\alpha_2;t_2';k_2}^{(\ell)}} \right) \widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2';k_2)}^{(\ell)} \,.$$

In words, the first term is the additive contribution from the $(\ell+1)$-th-block model parameters, the second term is the cumulative contribution from the residual path, the third term is the xeroxed

---

term from the skip path, and the last two terms are the cross terms between the skip and residual paths. (For $\ell + 1 = L$, the last three terms are absent as there is no skip path.) In more words, within the cumulative term, the multiplicative factor in the first set of the parentheses describes how the gradient changes as signals pass through the $(\ell+1)$-th block while the multiplicative factor in the second set of the parentheses describes how the normalization layer affects the gradient.[24]

Going forward – block by block – our strategy will be to *(i)* choose the relative learning-rate factor for SGD, λ_G, so that the first additive term is of order one, *(ii)* estimate the width scaling of the expected per-parameter gradient magnitude $|g^\mu|$ for model parameters in the group G, *(iii)* adjust the relative learning-rate factor for AdamW as $\widetilde{\lambda}_{G(\mu)} \sim \lambda_{G(\mu)}|g^\mu|$, and further *(iv)* recursively show that our choices of the initialization hyperparameters (1.25)–(1.33) lead to order-one multiplicative factors in the second cumulative term while cross terms vanish in expectation.

For SGD, such neural-tangent analysis of gradients will lead to – much like our analysis of preactivations led to the scalings (1.25)–(1.33) of the initialization hyperparameters – the following set of the scalings for the relative learning-rate factors:

$$\lambda_{W^{\text{patch}}} = \left(\frac{1}{n_{\text{patch}}}\right)\Lambda_{W^{\text{patch}}} \quad \Big|\Big| \quad \lambda_{W^{\text{WE}}} = (1)\cdot\Lambda_{\text{WE}}, \tag{1.99}$$

$$\lambda_{b^{\text{PE}}} = (1)\cdot\Lambda_{\text{PE}}, \tag{1.100}$$

$$\lambda_Q = \left(\frac{1}{n}\right)\Lambda_Q, \tag{1.101}$$

$$\lambda_K = \left(\frac{1}{n}\right)\Lambda_K, \tag{1.102}$$

$$\lambda_V = \left(\frac{1}{n}\right)\Lambda_V, \tag{1.103}$$

$$\lambda_U = \left(\frac{1}{n}\right)\Lambda_U, \tag{1.104}$$

$$\lambda_W = \left(\frac{1}{n}\right)\Lambda_W, \tag{1.105}$$

$$\lambda_X = \left(\frac{1}{Mn}\right)\Lambda_X, \tag{1.106}$$

$$\lambda_{W^{\text{head}}} = \left(\frac{1}{n}\right)\Lambda_{W^{\text{head}}}, \quad \lambda_{b^{\text{head}}} = (1)\cdot\Lambda_{b^{\text{head}}}, \tag{1.107}$$

where – like the order-one initialization hyperparameters C_G's – Λ_G's are order-one training hyperparameters which we could in principle tune but won't.[25] For AdamW, accounting for the factor

---

[24] These multiplicative factors are identified as the partial Jacobians in Ref. [17] – here squared and decomposed.
[25] Had we employed the $\lambda_{\mu\nu}$ notation of Refs. [15, 23] instead of the λ_G notation used herein, we would have denoted, e.g., the first equation (1.99) as $\lambda_{W^{\text{patch}}_{i_1 j_1} W^{\text{patch}}_{i_2 j_2}} = \lambda_{W^{\text{patch}}}\delta_{i_1 i_2}\delta_{j_1 j_2} = (\Lambda_{W^{\text{patch}}}/n_{\text{patch}})\delta_{i_1 i_2}\delta_{j_1 j_2}$. Slightly more

of $\sim |g^\mu|$ modifies the above relative learning-rate factors to

$$\widetilde{\lambda}_{W^{\mathrm{patch}}} = \left(\frac{1}{n_{\mathrm{patch}}\sqrt{n}}\right)\widetilde{\Lambda}_{W^{\mathrm{patch}}} \quad \Big\| \quad \widetilde{\lambda}_{W^{\mathrm{WE}}} = \left(\frac{1}{\sqrt{n}}\right)\widetilde{\Lambda}_{W^{\mathrm{WE}}}, \tag{1.108}$$

$$\widetilde{\lambda}_{b^{\mathrm{PE}}} = \left(\frac{1}{\sqrt{n}}\right)\widetilde{\Lambda}_{b^{\mathrm{PE}}}, \tag{1.109}$$

$$\widetilde{\lambda}_Q = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_Q, \tag{1.110}$$

$$\widetilde{\lambda}_K = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_K, \tag{1.111}$$

$$\widetilde{\lambda}_V = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_V, \tag{1.112}$$

$$\widetilde{\lambda}_U = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_U, \tag{1.113}$$

$$\widetilde{\lambda}_W = \left(\frac{1}{n\sqrt{Mn}}\right)\widetilde{\Lambda}_W, \tag{1.114}$$

$$\widetilde{\lambda}_X = \left(\frac{1}{Mn\sqrt{n}}\right)\widetilde{\Lambda}_X, \tag{1.115}$$

$$\widetilde{\lambda}_{W^{\mathrm{head}}} = \left(\frac{1}{n\sqrt{n_{\mathrm{out}}}}\right)\widetilde{\Lambda}_{W^{\mathrm{head}}}, \quad \widetilde{\lambda}_{b^{\mathrm{head}}} = \left(\frac{1}{\sqrt{n_{\mathrm{out}}}}\right)\widetilde{\Lambda}_{b^{\mathrm{head}}}, \tag{1.116}$$

so as to ensure that the network update stays of order one for AdamW.[26] With these choices, for the rest of this section, we'll recursively show that the means of the neural tangent kernels at initialization take the form

$$\mathbb{E}\left[\widehat{H}^{(\ell)}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)}\right] = \delta_{i_1 i_2}\Theta^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}, \tag{1.117}$$

with order-one kernels $\Theta^{(\ell)}_{(\alpha_1;t_1)(\alpha_2;t_2)}$.

With those goals in mind, let's proceed forward with our effective-theory analysis of gradients, block by block.

A. Stem Block

To stem the recursive analysis, we analyze the first-block neural tangent kernel,

$$\widehat{H}^{(1)}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)} = \sum_{\mu \in 1\text{st block}} \lambda_{G(\mu)} \frac{\partial z^{(1)}_{\alpha_1;t_1;i_1}}{\partial \theta^{(1)}_\mu} \frac{\partial z^{(1)}_{\alpha_2;t_2;i_2}}{\partial \theta^{(1)}_\mu}, \tag{1.118}$$

---

confusing in translation is that order-one training hyperparameters such as $\lambda_W^{(\ell)}$ in Ref. [15] would be denoted as $\Lambda_{W^{(\ell)}}$ herein.

[26] Let's harken back to footnote 8 and discuss generic meta-principled scaling strategies [23]. For learning rates, there is gauge redundancy [23, 24], but one way to implement it is, for SGD, to multiply the right-hand sides of (1.99)–(1.106) by n^s *except* the head one (1.107) and, for AdamW, to multiply the right-hand sides of (1.108)–(1.115) by $n^{\frac{s}{2}}$ *except* the head one (1.116). Here the factor of $n^{\frac{s}{2}}$ difference between SGD and AdamW arises from non-head gradient magnitude $|g^\mu|$'s picking up a factor of $\sqrt{1/n^s}$ from the head initialization scaling (1.34).

first for Vision Transformers and then for Language Transformers.

Vision: patchify embedding and positional embedding

Taking derivatives of the stem-block preactivations

$$z_{\alpha;t;i}^{(1)} = b_{t;i}^{\text{PE}} + \sum_{j=1}^{n_{\text{patch}}} W_{ij}^{\text{patch}} x_{\alpha;t;j} \tag{1.119}$$

with respect to the stem-block model parameters, we have

$$\frac{\partial z_{\alpha;t;i}^{(1)}}{\partial W_{jk}^{\text{patch}}} = \delta_{ij} x_{\alpha;t;k} \,, \tag{1.120}$$

$$\frac{\partial z_{\alpha;t;i}^{(1)}}{\partial b_{t';j}^{\text{PE}}} = \delta_{tt'} \delta_{ij} \,. \tag{1.121}$$

The squared norms of these gradients are then given by

$$\sum_{i=1}^{n} \sum_{j=1}^{n_{\text{patch}}} \frac{\partial z_{\alpha_1;t_1;i_1}^{(1)}}{\partial W_{ij}^{\text{patch}}} \frac{\partial z_{\alpha_2;t_2;i_2}^{(1)}}{\partial W_{ij}^{\text{patch}}} = \delta_{i_1 i_2} n_{\text{patch}} \left(\frac{1}{n_{\text{patch}}} \sum_{j=1}^{n_{\text{patch}}} x_{\alpha_1;t_1;j} x_{\alpha_2;t_2;j} \right) \tag{1.122}$$

$$= \delta_{i_1 i_2} \left[(n_{\text{patch}}) \cdot G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)} \right] \,,$$

$$\sum_{t=1}^{T} \sum_{i=1}^{n} \frac{\partial z_{\alpha_1;t_1;i_1}^{(1)}}{\partial b_{t;i}^{\text{PE}}} \frac{\partial z_{\alpha_2;t_2;i_2}^{(1)}}{\partial b_{t;i}^{\text{PE}}} = \delta_{i_1 i_2} \left[(1) \cdot \delta_{t_1 t_2} \right] \,, \tag{1.123}$$

where in the middle we recalled the definition of the order-one input kernel (1.40).[27] Note that, as promised (1.117), these additive contributions are diagonal in the embedding direction.

For SGD, the additive contributions (1.122) and (1.123) from these gradients to the stem-block neural tangent kernel (1.118) are multiplied by $\lambda_{W^{\text{patch}}}$ and $\lambda_{b^{\text{PE}}}$, respectively. Thus, in order to make these contributions stay of order one, we need to scale relative learning-rate factors as

$$\lambda_{W^{\text{patch}}} = \left(\frac{1}{n_{\text{patch}}} \right) \Lambda_{W^{\text{patch}}} \,, \quad \lambda_{b^{\text{PE}}} = (1) \cdot \Lambda_{b^{\text{PE}}} \quad \text{for} \quad \text{SGD} \,, \tag{1.124}$$

with order-one hyperparameters $\Lambda_{W^{\text{patch}}}$ and $\Lambda_{b^{\text{PE}}}$. For AdamW, we need to account for the magnitude of the per-parameter gradient $|g^{\mu}|$. To estimate it for the patchify weights, we note that there are $n \cdot n_{\text{patch}}$ components adding up to the squared norm of the order n_{patch}, so each component of the gradient on average scales as $|g^{\mu}| \sim \sqrt{n_{\text{patch}}/(n \cdot n_{\text{patch}})} = 1/\sqrt{n}$. Similarly for

---

[27] In passing, we note that these relations hold even without taking expectation values, i.e., the first-block neural tangent kernel is deterministic.

the positional-embedding parameters, we estimate $|g^\mu| \sim \sqrt{1/n} = 1/\sqrt{n}$.[28] Thus, to offset these factors in the AdamW optimizer, we set

$$\widetilde{\lambda}_{W^{\mathrm{patch}}} = \left(\frac{1}{n_{\mathrm{patch}}\sqrt{n}}\right) \widetilde{\Lambda}_{W^{\mathrm{patch}}} , \quad \widetilde{\lambda}_{b^{\mathrm{PE}}} = \left(\frac{1}{\sqrt{n}}\right) \widetilde{\Lambda}_{b^{\mathrm{PE}}} \quad \text{for} \quad \text{AdamW} , \tag{1.125}$$

with order-one hyperparameters $\widetilde{\Lambda}_{W^{\mathrm{patch}}}$ and $\widetilde{\Lambda}_{b^{\mathrm{PE}}}$.

Language: word embedding and positional embedding

Very similarly to the vision case, taking derivatives of the stem-block preactivations

$$z_{\alpha;t;i}^{(1)} = b_{t;i}^{\mathrm{PE}} + \sum_{j=1}^{n_{\mathrm{vocab}}} W_{ij}^{\mathrm{WE}} x_{\alpha;t;j} , \tag{1.126}$$

with respect to the stem-block model parameters, we have

$$\left.\frac{\partial z_{\alpha;t;i}^{(1)}}{\partial W_{jk}^{\mathrm{WE}}}\right|_{\mathrm{stem}} = \delta_{ij} x_{\alpha;t;k} , \tag{1.127}$$

$$\frac{\partial z_{\alpha;t;i}^{(1)}}{\partial b_{t';j}^{\mathrm{PE}}} = \delta_{tt'} \delta_{ij} . \tag{1.128}$$

Here, for the derivative with respect to the word-embedding parameters, we emphasized that these gradients come from the stem block because there will be another contribution from the head block, to be discussed in §I 3 E.

Anyhow, the squared norms of these gradients are then given by

$$\sum_{i=1}^{n} \sum_{j=1}^{n_{\mathrm{vocab}}} \left.\frac{\partial z_{\alpha_1;t_1;i_1}^{(1)}}{\partial W_{ij}^{\mathrm{WE}}}\right|_{\mathrm{stem}} \left.\frac{\partial z_{\alpha_2;t_2;i_2}^{(1)}}{\partial W_{ij}^{\mathrm{WE}}}\right|_{\mathrm{stem}} = \delta_{i_1 i_2} \left(\sum_{j=1}^{n_{\mathrm{vocab}}} x_{\alpha_1;t_1;j} x_{\alpha_2;t_2;j}\right) \tag{1.129}$$

$$= \delta_{i_1 i_2} \left[(1) \cdot G_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(0)}\right] ,$$

$$\sum_{t=1}^{T} \sum_{i=1}^{n} \frac{\partial z_{\alpha_1;t_1;i_1}^{(1)}}{\partial b_{t;i}^{\mathrm{PE}}} \frac{\partial z_{\alpha_2;t_2;i_2}^{(1)}}{\partial b_{t;i}^{\mathrm{PE}}} = \delta_{i_1 i_2} \left[(1) \cdot \delta_{t_1 t_2}\right] , \tag{1.130}$$

where in the middle we recalled the definition of the order-one input kernel (1.44). Note that, as promised (1.117), these additive contributions are diagonal in the embedding direction.

For SGD, to ensure order-one additive contributions to the stem-block neural tangent kernel (1.118), we scale the relative learning-rate factors as

$$\lambda_{W^{\mathrm{WE}}} = (1) \cdot \Lambda_{W^{\mathrm{WE}}} , \quad \lambda_{b^{\mathrm{PE}}} = (1) \cdot \Lambda_{b^{\mathrm{PE}}} \quad \text{for} \quad \text{SGD} , \tag{1.131}$$

---

[28] In this note, we implicitly focus on the scalings of hyperparameters with widths and do *not* investigate the scalings with the batch size $|\mathcal{A}_t|$, sequence length T, etc.

i.e., we keep both of them of order one. For AdamW, we adjust for the magnitude of the per-parameter gradient $|g^\mu| \sim \sqrt{1/n} = 1/\sqrt{n}$ both for word-embedding and positional-embedding parameters, yielding

$$\widetilde{\lambda}_{W^{\text{WE}}} = \left(\frac{1}{\sqrt{n}}\right) \widetilde{\Lambda}_{W^{\text{WE}}}, \quad \widetilde{\lambda}_{b^{\text{PE}}} = \left(\frac{1}{\sqrt{n}}\right) \widetilde{\Lambda}_{b^{\text{PE}}} \quad \text{for} \quad \text{AdamW}. \tag{1.132}$$

Here, in estimating the gradient magnitude for each component of the word-embedding parameters, we didn't divide by a factor of n_{vocab}, for a reason: unlike in the vision case where gradients are spread across all the n_{patch} input components, here in the language case the gradients are mostly zero due to the one-hot structure of inputs, i.e., $x_{\alpha;t;j} = \delta_{jj_\star(\alpha;t)}$.[29]

B. Layer Normalization

Mirroring §I 1 B, we'll temporarily drop the sample–token–block indices until the dust settles.

To analyze how normalization layers affect the gradients, looking back at the forward equation for the neural tangent kernels (1.98), it would appear that we need to analyze the multiplicative factor

$$\frac{\partial s_{j_1}}{\partial z_{k_1}} \frac{\partial s_{j_2}}{\partial z_{k_2}}, \tag{1.133}$$

for general combinations of the embedding indices (j_1, j_2, k_1, k_2). Fortunately, because of the recursively-expected diagonal form (1.117) of the neural tangent kernel in the embedding direction, we can focus on the ones with $k_1 = k_2$. Further we'll recursively show in the following §§I 3 C–I 3 E [in particular equations (1.161), (1.163), (1.173), and (1.182)] that the (putatively) leading-order contributions always come in the forms of the following two particular linear combinations:

$$\frac{1}{n} \sum_{j=1}^{n} \sum_{k=1}^{n} \frac{\partial s_j}{\partial z_k} \frac{\partial s_j}{\partial z_k}, \tag{1.134}$$

and

$$\frac{1}{n^2} \sum_{j_1, j_2 = 1}^{n} \sum_{k=1}^{n} s_{j_1} s_{j_2} \frac{\partial s_{j_1}}{\partial z_k} \frac{\partial s_{j_2}}{\partial z_k}. \tag{1.135}$$

---

[29] This reasoning, however, becomes more dubious as the batch size $|\mathcal{A}_t|$ grows larger: for sufficiently large batch sizes, all the gradients will in principle become nonzero, albeit inhomogeneously due to non-uniform token distributions in typical vocabularies [30] (which incidentally forbids us from naively scaling $\widetilde{\lambda}_{W^{\text{WE}}}$ with a simple power of n_{vocab}). There is another subtlety that arises in the head block: see footnote 33.

Overall, we leave the analysis of the interplay between the batch size $|\mathcal{A}_t|$, sequence length T, vocabulary size n_{vocab}, and token distribution to future work.

Let's start with the first one (1.134). Mindlessly rolling it out, we get

$$\frac{1}{n}\sum_{j=1}^{n}\sum_{k=1}^{n}\frac{\partial s_j}{\partial z_k}\frac{\partial s_j}{\partial z_k} \tag{1.136}$$

$$=\frac{1}{n}\sum_{j=1}^{n}\sum_{k=1}^{n}\left\{\frac{\partial}{\partial z_k}\left[\frac{z_j-\frac{1}{n}\sum_{i=1}^{n}z_i}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}z_i^2-\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2+\epsilon}}\right]\right\}^2$$

$$=\frac{1}{n}\sum_{j=1}^{n}\sum_{k=1}^{n}\left[\frac{1}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}z_i^2-\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2+\epsilon}}\left(\delta_{jk}-\frac{1}{n}-\frac{1}{n}s_js_k\right)\right]^2$$

$$=\frac{1}{\left[\sqrt{\frac{1}{n}\sum_{i=1}^{n}z_i^2-\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2+\epsilon}\right]^2}\left[1-\frac{1}{n}-\frac{2}{n}\left(\frac{1}{n}\sum_{j=1}^{n}s_j^2\right)+\frac{2}{n}\left(\frac{1}{n}\sum_{j=1}^{n}s_j\right)^2+\frac{1}{n}\left(\frac{1}{n}\sum_{j=1}^{n}s_j^2\right)^2\right].$$

Recalling from §I 0 B that $\frac{1}{n}\sum_{i=1}^{n}z_i^2=G+\Delta G$ and $\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2=\nabla G$ and that ΔG and ∇G are ignorable at leading order, we get

$$\frac{1}{n}\sum_{j=1}^{n}\sum_{k=1}^{n}\frac{\partial s_j}{\partial z_k}\frac{\partial s_j}{\partial z_k}=\frac{1}{\left(\sqrt{G+\epsilon}\right)^2}+O\left(\frac{1}{n}\right). \tag{1.137}$$

Next, mindlessly rolling out the second one (1.135), we get

$$\frac{1}{n^2}\sum_{j_1,j_2=1}^{n}\sum_{k=1}^{n}s_{j_1}s_{j_2}\frac{\partial s_{j_1}}{\partial z_k}\frac{\partial s_{j_2}}{\partial z_k} \tag{1.138}$$

$$=\sum_{k=1}^{n}\left\{\frac{1}{n}\sum_{j=1}^{n}s_j\frac{\partial}{\partial z_k}\left[\frac{z_j-\frac{1}{n}\sum_{i=1}^{n}z_i}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}z_i^2-\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2+\epsilon}}\right]\right\}^2$$

$$=\sum_{k=1}^{n}\left\{\frac{1}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}z_i^2-\left(\frac{1}{n}\sum_{i=1}^{n}z_i\right)^2+\epsilon}}\left[\frac{s_k}{n}-\frac{1}{n}\left(\frac{1}{n}\sum_{j=1}^{n}s_j\right)-\frac{1}{n}\left(\frac{1}{n}\sum_{j=1}^{n}s_j^2\right)s_k\right]\right\}^2$$

$$=O\left(\frac{1}{n}\right).$$

That is, we can neglect this type of contributions at the leading order.

As the dust has settled, carefully putting all the indices back in, we get

$$\frac{1}{n}\sum_{j=1}^{n}\sum_{k=1}^{n}\frac{\partial s^{(\ell)}_{\alpha_1;t_1;j}}{\partial z^{(\ell)}_{\alpha_1;t_1;k}}\frac{\partial s^{(\ell)}_{\alpha_2;t_2;j}}{\partial z^{(\ell)}_{\alpha_2;t_2;k}}=\frac{1}{\sqrt{G^{(\ell)}_{(\alpha_1;t_1)(\alpha_1;t_1)}+\epsilon}\sqrt{G^{(\ell)}_{(\alpha_2;t_2)(\alpha_2;t_2)}+\epsilon}}+O\left(\frac{1}{n}\right). \tag{1.139}$$

Note that, since we expect the kernel $G^{(\ell)}$ to grow linearly with the depth ℓ, we expect this multiplicative factor to be linearly suppressed [17].

Overall, in retrospect and in prospect, we can simply use the quick intuitive mnemonics

$$s_{\alpha;t;j}\approx\frac{z_{\alpha;t;j}}{\sqrt{G_{(\alpha;t)(\alpha;t)}+\epsilon}}\quad\text{and}\quad\frac{\partial s_{\alpha;t;j}}{\partial z_{\alpha;t;k}}\approx\frac{1}{\sqrt{G_{(\alpha;t)(\alpha;t)}+\epsilon}}\delta_{jk}, \tag{1.140}$$

from which all of our results here in I 3 B and there in I 1 B follow immediately.

C.   Multi-Head Self-Attention Block

Taking derivatives of the residual MHSA-block outputs

$$r_{\alpha;t;i} = \sum_{h=1}^{H}\sum_{t'=1}^{T}\sum_{c=1}^{C}\sum_{j=1}^{n} \Omega^h_{\alpha;tt'}\left[\widetilde{\Omega}\left(Q,K;s\right)\right]U^h_{ic}V^h_{cj}s_{\alpha;t';j} \tag{1.141}$$

with respect to the MHSA-block model parameters and the incoming signals, we have

$$\frac{\partial r_{\alpha;t;i}}{\partial U^h_{jc}} = \delta_{ij}\sum_{t'=1}^{T}\sum_{k=1}^{n}V^h_{ck}\Omega^h_{\alpha;tt'}s_{\alpha;t';k} \,, \tag{1.142}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial V^h_{cj}} = U^h_{ic}\sum_{t'=1}^{T}\Omega^h_{\alpha;tt'}s_{\alpha;t';j} \,, \tag{1.143}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial K^h_{cj}} = \sum_{t'=1}^{T}\sum_{c'=1}^{C}\sum_{k=1}^{n}U^h_{ic'}V^h_{c'k}\left(\sum_{\tilde{t},\tilde{t}'=1}^{T}\frac{\partial\Omega^h_{\alpha;tt'}}{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial K^h_{cj}}\right)s_{\alpha;t';k} \,, \tag{1.144}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial Q^h_{cj}} = \sum_{t'=1}^{T}\sum_{c'=1}^{C}\sum_{k=1}^{n}U^h_{ic'}V^h_{c'k}\left(\sum_{\tilde{t},\tilde{t}'=1}^{T}\frac{\partial\Omega^h_{\alpha;tt'}}{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial Q^h_{cj}}\right)s_{\alpha;t';k} \,, \tag{1.145}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial s_{\alpha;t'';k}} = \sum_{t'=1}^{T}\sum_{h=1}^{H}\sum_{c=1}^{C}U^h_{ic}\sum_{j=1}^{n}V^h_{cj}\left[\delta_{jk}\Omega^h_{\alpha;tt''} + \left(\sum_{\tilde{t},\tilde{t}'=1}^{T}\frac{\partial\Omega^h_{\alpha;tt'}}{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial s_{\alpha;t'';k}}\right)s_{\alpha;t';j}\right] \,, \tag{1.146}$$

where for the last three we used the chain rule to roll out the derivatives of the self-attention matrix $\Omega^h_{\alpha;tt'}$. A little bit more explicitly, the derivatives of the query–key dot product (1.16),

$$\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'} = \frac{1}{\sqrt{C}}\sum_{c=1}^{C}q^h_{\alpha;\tilde{t};c}k^h_{\alpha;\tilde{t}';c} = \frac{1}{\sqrt{C}}\sum_{c=1}^{C}\sum_{i_1,i_2=1}^{n}Q^h_{ci_1}K^h_{ci_2}s_{\alpha;\tilde{t};i_1}s_{\alpha;\tilde{t}';i_2} \,, \tag{1.147}$$

are given by

$$\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial K^h_{cj}} = \frac{1}{\sqrt{C}}\sum_{m=1}^{n}Q^h_{cm}s_{\alpha;\tilde{t};m}s_{\alpha;\tilde{t}';j} = \frac{1}{\sqrt{C}}q^h_{\alpha;\tilde{t};c}s_{\alpha;\tilde{t}';j} \,, \tag{1.148}$$

$$\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial Q^h_{cj}} = \frac{1}{\sqrt{C}}\sum_{m=1}^{n}K^h_{cm}s_{\alpha;\tilde{t};j}s_{\alpha;\tilde{t}';m} = \frac{1}{\sqrt{C}}s_{\alpha;\tilde{t};j}k^h_{\alpha;\tilde{t}';c} \,, \tag{1.149}$$

$$\frac{\partial\widetilde{\Omega}^h_{\alpha;\tilde{t}\tilde{t}'}}{\partial s_{\alpha;t'';k}} = \frac{1}{\sqrt{C}}\sum_{c=1}^{C}\sum_{m=1}^{n}\left(\delta_{t''\tilde{t}}Q^h_{ck}K^h_{cm}s_{\alpha;\tilde{t}';m} + \delta_{t''\tilde{t}'}Q^h_{cm}K^h_{ck}s_{\alpha;\tilde{t};m}\right) \tag{1.150}$$

$$= \frac{1}{\sqrt{C}}\sum_{c=1}^{C}\left(\delta_{t''\tilde{t}}Q^h_{ck}k^h_{\alpha;\tilde{t}';c} + \delta_{t''\tilde{t}'}q^h_{\alpha;\tilde{t};c}K^h_{ck}\right) \,.$$

Expectation values for the squared gradient norms are then given by – recalling $n = HC$ –

$$\mathbb{E}\left[\sum_{h=1}^{H}\sum_{c=1}^{C}\sum_{j=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial U_{jc}^h}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial U_{jc}^h}\right] \tag{1.151}$$

$$=\delta_{i_1 i_2}\left\{(n)\cdot C_V\sum_{t_1',t_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\Omega_{\alpha_1;t_1 t_1'}^h\Omega_{\alpha_2;t_2 t_2'}^h\right)\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;t_1';j}s_{\alpha_2;t_2';j}\right)\right]\right\},$$

$$\mathbb{E}\left[\sum_{h=1}^{H}\sum_{c=1}^{C}\sum_{j=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial V_{cj}^h}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial V_{cj}^h}\right] \tag{1.152}$$

$$=\delta_{i_1 i_2}\left\{(n)\cdot C_U\sum_{t_1',t_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\Omega_{\alpha_1;t_1 t_1'}^h\Omega_{\alpha_2;t_2 t_2'}^h\right)\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;t_1';j}s_{\alpha_2;t_2';j}\right)\right]\right\},$$

$$\mathbb{E}\left[\sum_{h=1}^{H}\sum_{c=1}^{C}\sum_{j=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial K_{cj}^h}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial K_{cj}^h}\right] \tag{1.153}$$

$$=\delta_{i_1 i_2}\left\{(n)\cdot C_U C_V\sum_{t_1',t_2',\tilde{t}_1,\tilde{t}_2,\tilde{t}_1',\tilde{t}_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\frac{\partial\Omega_{\alpha_1;t_1 t_1'}^h}{\partial\widetilde{\Omega}_{\alpha_1;\tilde{t}_1\tilde{t}_1'}^h}\frac{\partial\Omega_{\alpha_2;t_2 t_2'}^h}{\partial\widetilde{\Omega}_{\alpha_2;\tilde{t}_2\tilde{t}_2'}^h}\right)\left(\frac{1}{C}\sum_{c=1}^{C}q_{\alpha_1;\tilde{t}_1;c}^h q_{\alpha_2;\tilde{t}_2;c}^h\right)\right.\right.$$

$$\left.\left.\times\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;\tilde{t}_1';j}s_{\alpha_2;\tilde{t}_2';j}\right)\left(\frac{1}{n}\sum_{k=1}^{n}s_{\alpha_1;t_1';k}s_{\alpha_2;t_2';k}\right)\right]\right\},$$

$$\mathbb{E}\left[\sum_{h=1}^{H}\sum_{c=1}^{C}\sum_{j=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial Q_{cj}^h}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial Q_{cj}^h}\right] \tag{1.154}$$

$$=\delta_{i_1 i_2}\left\{(n)\cdot C_U C_V\sum_{t_1',t_2',\tilde{t}_1,\tilde{t}_2,\tilde{t}_1',\tilde{t}_2'=1}^{T}\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\frac{\partial\Omega_{\alpha_1;t_1 t_1'}^h}{\partial\widetilde{\Omega}_{\alpha_1;\tilde{t}_1\tilde{t}_1'}^h}\frac{\partial\Omega_{\alpha_2;t_2 t_2'}^h}{\partial\widetilde{\Omega}_{\alpha_2;\tilde{t}_2\tilde{t}_2'}^h}\right)\left(\frac{1}{C}\sum_{c=1}^{C}k_{\alpha_1;\tilde{t}_1';c}^h k_{\alpha_2;\tilde{t}_2';c}^h\right)\right.\right.$$

$$\left.\left.\times\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;\tilde{t}_1;j}s_{\alpha_2;\tilde{t}_2;j}\right)\left(\frac{1}{n}\sum_{k=1}^{n}s_{\alpha_1;t_1';k}s_{\alpha_2;t_2';k}\right)\right]\right\}.$$

Here, the term in each pair of the parentheses in expectation is expected to be of order one and thus the squared gradient norms all scale as n.[30] Note that, as promised (1.117), these additive

---

[30] In particular, for instance, we can use the results of Appendix A to factor the expectation as

$$\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\frac{\partial\Omega_{\alpha_1;t_1 t_1'}^h}{\partial\widetilde{\Omega}_{\alpha_1;\tilde{t}_1\tilde{t}_1'}^h}\frac{\partial\Omega_{\alpha_2;t_2 t_2'}^h}{\partial\widetilde{\Omega}_{\alpha_2;\tilde{t}_2\tilde{t}_2'}^h}\right)\left(\frac{1}{C}\sum_{c=1}^{C}k_{\alpha_1;\tilde{t}_1 c}^h k_{\alpha_2;\tilde{t}_2 c}^h\right)\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;\tilde{t}_1';j}s_{\alpha_2;\tilde{t}_2';j}\right)\left(\frac{1}{n}\sum_{k=1}^{n}s_{\alpha_1;t_1';k}s_{\alpha_2;t_2';k}\right)\right]$$

$$=\mathbb{E}\left[\left(\frac{1}{H}\sum_{h=1}^{H}\frac{\partial\Omega_{\alpha_1;t_1 t_1'}^h}{\partial\widetilde{\Omega}_{\alpha_1;\tilde{t}_1\tilde{t}_1'}^h}\frac{\partial\Omega_{\alpha_2;t_2 t_2'}^h}{\partial\widetilde{\Omega}_{\alpha_2;\tilde{t}_2\tilde{t}_2'}^h}\right)\right]\mathbb{E}\left[\frac{1}{C}\sum_{c=1}^{C}k_{\alpha_1;\tilde{t}_1 c}^h k_{\alpha_2;\tilde{t}_2 c}^h\right]F_{(\alpha_1;\tilde{t}_1')(\alpha_2;\tilde{t}_2')}F_{(\alpha_1;t_1')(\alpha_2;t_2')}+O\left(\frac{1}{C}\right) \tag{1.155}$$

where

$$\mathbb{E}\left[\frac{1}{C}\sum_{c=1}^{C}k_{\alpha_1;\tilde{t}_1 c}^h k_{\alpha_2;\tilde{t}_2 c}^h\right]=\mathbb{E}\left[\frac{1}{C}\sum_{c=1}^{C}\sum_{i_1,i_2=1}^{n}K_{ci_1}^h K_{ci_2}^h s_{\alpha_1;\tilde{t}_1;i_1}s_{\alpha_2;\tilde{t}_2;i_2}\right]=C_K F_{(\alpha_1;\tilde{t}_1)(\alpha_2;\tilde{t}_2)}, \tag{1.156}$$

and further express the remaining expectation as a $\left(|\mathcal{D}|T^2\right)$-dimensional Gaussian integral.

contributions are diagonal in the embedding direction.

For SGD, to ensure order-one additive contributions to the neural tangent kernel, we thus should scale relative learning-rate factors λ_G as

$$\lambda_Q = \left(\frac{1}{n}\right)\Lambda_Q\,,\quad \lambda_K = \left(\frac{1}{n}\right)\Lambda_K\,,\quad \lambda_V = \left(\frac{1}{n}\right)\Lambda_V\,,\quad \lambda_U = \left(\frac{1}{n}\right)\Lambda_U \quad \text{for}\quad \text{SGD}\,. \tag{1.157}$$

For AdamW, we need to adjust for the magnitude of the per-parameter gradient. Since all these weights have n^2 components adding up to the squared norm of order n, each component of the gradient scales as $|g^\mu| \sim \sqrt{n/(n^2)} = 1/\sqrt{n}$. Thus, to offset these factors in the AdamW optimizer, we set

$$\widetilde{\lambda}_Q = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_Q\,,\quad \widetilde{\lambda}_K = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_K\,,\quad \widetilde{\lambda}_V = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_V\,,\quad \widetilde{\lambda}_U = \left(\frac{1}{n\sqrt{n}}\right)\widetilde{\Lambda}_U \quad \text{for}\quad \text{AdamW}\,. \tag{1.158}$$

Finally, moving onto the recursive contributions, the cross terms vanish as

$$\mathbb{E}\left[\frac{\partial r_{\alpha;t;i}}{\partial s_{\alpha;t'';k}}\mathcal{F}^{(\text{previous})}\right] = 0\,, \tag{1.159}$$

due to the dangling mean-zero weights U and V in the signal derivative (1.146), where here $\mathcal{F}^{(\text{previous})}$ is a function of observables from the preceding blocks. For the cumulative contribution, we have

$$\mathbb{E}\left[\sum_{k_1,k_2=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial s_{\alpha_1;t_1'';k_1}}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial s_{\alpha_2;t_2'';k_2}}\mathcal{F}^{(\text{previous})}_{k_1k_2}\right] \tag{1.160}$$

$$=\delta_{i_1i_2}C_U C_V \frac{1}{H}\sum_{h=1}^{H}\sum_{t_1',t_2'=1}^{T}\frac{1}{n}\sum_{k_1,k_2=1}^{n}$$

$$\times\,\mathbb{E}\Bigg\{\Bigg[\delta_{k_1k_2}\Omega^h_{\alpha_1;t_1t_1''}\Omega^h_{\alpha_2;t_2t_2''}$$

$$+\,s_{\alpha_2;t_2';k_1}\Omega^h_{\alpha_1;t_1t_1''}\sum_{\tilde{t}_2,\tilde{t}_2'=1}^{T}\frac{\partial\Omega^h_{\alpha_2;t_2t_2'}}{\partial\widetilde{\Omega}^h_{\alpha_2;\tilde{t}_2\tilde{t}_2'}}\frac{1}{\sqrt{C}}\sum_{c_2=1}^{C}\left(\delta_{t_2''\tilde{t}_2}Q^h_{c_2k_2}k^h_{\alpha_2;\tilde{t}_2';c_2}+\delta_{t_2''\tilde{t}_2'}K^h_{c_2k_2}q^h_{\alpha_2;\tilde{t}_2;c_2}\right)$$

$$+\,s_{\alpha_1;t_1';k_2}\Omega^h_{\alpha_2;t_2t_2''}\sum_{\tilde{t}_1,\tilde{t}_1'=1}^{T}\frac{\partial\Omega^h_{\alpha_1;t_1t_1'}}{\partial\widetilde{\Omega}^h_{\alpha_1;\tilde{t}_1\tilde{t}_1'}}\frac{1}{\sqrt{C}}\sum_{c_1=1}^{C}\left(\delta_{t_1''\tilde{t}_1}Q^h_{c_1k_1}k^h_{\alpha_1;\tilde{t}_1';c_1}+\delta_{t_1''\tilde{t}_1'}K^h_{c_1k_1}q^h_{\alpha_1;\tilde{t}_1;c_1}\right)$$

$$+\,n\left(\frac{1}{n}\sum_{j=1}^{n}s_{\alpha_1;t_1';j}s_{\alpha_2;t_2';j}\right)\sum_{\tilde{t}_1,\tilde{t}_1',\tilde{t}_2,\tilde{t}_2'=1}^{T}\frac{\partial\Omega^h_{\alpha_1;t_1t_1'}}{\partial\widetilde{\Omega}^h_{\alpha_1;\tilde{t}_1\tilde{t}_1'}}\frac{\partial\Omega^h_{\alpha_2;t_2t_2'}}{\partial\widetilde{\Omega}^h_{\alpha_2;\tilde{t}_2\tilde{t}_2'}}\frac{1}{C}\sum_{c_1,c_2=1}^{C}\sum_{m_1,m_2=1}^{n}$$

$$\times\left(\delta_{t_1''\tilde{t}_1}Q^h_{c_1k_1}K^h_{c_1m_1}s_{\alpha_1;\tilde{t}_1';m_1}+\delta_{t_1''\tilde{t}_1'}Q^h_{c_1m_1}K^h_{c_1k_1}s_{\alpha_1;\tilde{t}_1;m_1}\right)$$

$$\times\left(\delta_{t_2''\tilde{t}_2}Q^h_{c_2k_2}K^h_{c_2m_2}s_{\alpha_2;\tilde{t}_2';m_2}+\delta_{t_2''\tilde{t}_2'}Q^h_{c_2m_2}K^h_{c_2k_2}s_{\alpha_2;\tilde{t}_2;m_2}\right)\Bigg]\mathcal{F}^{(\text{previous})}_{k_1k_2}\Bigg\}\,.$$

Now there are four terms in expectation. The first term gives rise to the contribution

$$\delta_{i_1 i_2} C_U C_V \frac{1}{H} \sum_{h=1}^{H} \sum_{t_1', t_2'=1}^{T} \mathbb{E}\left[\Omega^h_{\alpha_1; t_1 t_1''} \Omega^h_{\alpha_2; t_2 t_2''} \frac{1}{n} \left(\sum_{k=1}^{n} \mathcal{F}^{\text{(previous)}}_{kk} \right) \right], \qquad (1.161)$$

which has the first promised form (1.134) of the layer-normalization multiplicative factor and also is diagonal in the embedding direction (1.117). The next two cross terms actually vanish at this

order.[31] For the last term, integrating over Q and K weights, we get

$$\delta_{i_1 i_2} C_U C_V C_Q C_K \frac{1}{H} \sum_{h=1}^{H} \sum_{t'_1,t'_2,\tilde{t}_1,\tilde{t}'_1,\tilde{t}_2,\tilde{t}'_2=1}^{T} \tag{1.163}$$

$$\times \mathbb{E}\left[\frac{\partial \Omega^h_{\alpha_1;t_1 t'_1}}{\partial \widetilde{\Omega}^h_{\alpha_1;\tilde{t}_1 \tilde{t}'_1}} \frac{\partial \Omega^h_{\alpha_2;t_2 t'_2}}{\partial \widetilde{\Omega}^h_{\alpha_2;\tilde{t}_2 \tilde{t}'_2}} \left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t'_1;j} s_{\alpha_2;t'_2;j}\right) \frac{1}{n^2} \sum_{k_1,k_2,m_1,m_2=1}^{n}\right.$$

$$\times \left(\delta_{t''_1 \tilde{t}_1} \delta_{t''_2 \tilde{t}_2} \delta_{k_1 k_2} \delta_{m_1 m_2} s_{\alpha_1;\tilde{t}'_1;m_1} s_{\alpha_2;\tilde{t}'_2;m_2} + \delta_{t''_1 \tilde{t}_1} \delta_{t''_2 \tilde{t}_2} \delta_{k_1 m_2} \delta_{k_2 m_1} s_{\alpha_1;\tilde{t}'_1;m_1} s_{\alpha_2;\tilde{t}_2;m_2}\right.$$

$$\left.\left. + \delta_{t''_1 \tilde{t}'_1} \delta_{t''_2 \tilde{t}_2} \delta_{k_1 m_2} \delta_{k_2 m_1} s_{\alpha_1;\tilde{t}_1;m_1} s_{\alpha_2;\tilde{t}'_2;m_2} + \delta_{t''_1 \tilde{t}'_1} \delta_{t''_2 \tilde{t}_2} \delta_{k_1 k_2} \delta_{m_1 m_2} s_{\alpha_1;\tilde{t}_1;m_1} s_{\alpha_2;\tilde{t}_2;m_2}\right) \mathcal{F}^{(\mathrm{previous})}_{k_1 k_2}\right]$$

$$=\delta_{i_1 i_2} C_U C_V C_Q C_K \frac{1}{H} \sum_{h=1}^{H} \sum_{t'_1,t'_2,\tilde{t}_1,\tilde{t}'_1,\tilde{t}_2,\tilde{t}'_2=1}^{T}$$

$$\times \mathbb{E}\left\{\frac{\partial \Omega^h_{\alpha_1;t_1 t'_1}}{\partial \widetilde{\Omega}^h_{\alpha_1;\tilde{t}_1 \tilde{t}'_1}} \frac{\partial \Omega^h_{\alpha_2;t_2 t'_2}}{\partial \widetilde{\Omega}^h_{\alpha_2;\tilde{t}_2 \tilde{t}'_2}} \left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t'_1;j} s_{\alpha_2;t'_2;j}\right)\right.$$

$$\left.\times \left[\delta_{t''_1 \tilde{t}_1} \delta_{t''_2 \tilde{t}_2} \left(\frac{1}{n}\sum_{m=1}^{n} s_{\alpha_1;\tilde{t}'_1;m} s_{\alpha_2;\tilde{t}'_2;m}\right) + \delta_{t''_1 \tilde{t}'_1} \delta_{t''_2 \tilde{t}'_2} \left(\frac{1}{n}\sum_{m=1}^{n} s_{\alpha_1;\tilde{t}_1;m} s_{\alpha_2;\tilde{t}_2;m}\right)\right] \left(\frac{1}{n}\sum_{k=1}^{n} \mathcal{F}^{(\mathrm{previous})}_{kk}\right)\right\}$$

$$+\delta_{i_1 i_2} C_U C_V C_Q C_K \frac{1}{H} \sum_{h=1}^{H} \sum_{t'_1,t'_2,\tilde{t}_1,\tilde{t}'_1,\tilde{t}_2,\tilde{t}'_2=1}^{T}$$

$$\times \mathbb{E}\left\{\frac{\partial \Omega^h_{\alpha_1;t_1 t'_1}}{\partial \widetilde{\Omega}^h_{\alpha_1;\tilde{t}_1 \tilde{t}'_1}} \frac{\partial \Omega^h_{\alpha_2;t_2 t'_2}}{\partial \widetilde{\Omega}^h_{\alpha_2;\tilde{t}_2 \tilde{t}'_2}} \left(\frac{1}{n}\sum_{j=1}^{n} s_{\alpha_1;t'_1;j} s_{\alpha_2;t'_2;j}\right)\right.$$

$$\times \left[\delta_{t''_1 \tilde{t}_1} \delta_{t''_2 \tilde{t}'_2} \left(\frac{1}{n^2}\sum_{m_1,m_2=1}^{n} s_{\alpha_1;\tilde{t}'_1;m_1} s_{\alpha_2;\tilde{t}_2;m_2} \mathcal{F}^{(\mathrm{previous})}_{m_2 m_1}\right)\right.$$

$$\left.\left.\left. + \delta_{t''_1 \tilde{t}'_1} \delta_{t''_2 \tilde{t}_2} \left(\frac{1}{n^2}\sum_{m_1,m_2=1}^{n} s_{\alpha_1;\tilde{t}_1;m_1} s_{\alpha_2;\tilde{t}'_2;m_2} \mathcal{F}^{(\mathrm{previous})}_{m_2 m_1}\right)\right]\right]\right\}.$$

---

[31] The argument essentially boils down to doing algebra of the form

$$\frac{1}{n}\sum_{k_1,k_2=1}^{n} \mathbb{E}\left[\left(\widetilde{\Omega}^h_{\alpha;tt'}\right) s_{\alpha_2;t'_2;k_1} \frac{1}{\sqrt{C}} Q^h_{c_2 k_2} k^h_{\alpha_2;\tilde{t}'_2;c_2} \mathcal{F}^{(\mathrm{previous})}_{k_1 k_2}\right] \tag{1.162}$$

$$=\frac{1}{n}\sum_{k_1,k_2=1}^{n} \frac{1}{C}\sum_{c_2=1}^{C} \mathbb{E}\left[\left(\sum_{c=1}^{C}\sum_{i_1,i_2=1}^{n} Q^h_{c i_1} K^h_{c i_2} s_{\alpha;t;i_1} s_{\alpha;t';i_2}\right) s_{\alpha_2;t'_2;k_1} Q^h_{c_2 k_2} \sum_{m_2=1}^{n} K^h_{c_2 m_2} s_{\alpha_2;\tilde{t}'_2;m_2} \mathcal{F}^{(\mathrm{previous})}_{k_1 k_2}\right]$$

$$=C_Q C_K \frac{1}{n^3}\sum_{i_1,i_2,k_1,k_2,m_2=1}^{n} \mathbb{E}\left[\delta_{i_1 k_2} \delta_{i_2 m_2} s_{\alpha;t;i_1} s_{\alpha;t';i_2} s_{\alpha_2;t'_2;k_1} s_{\alpha_2;\tilde{t}'_2;m_2} \mathcal{F}^{(\mathrm{previous})}_{k_1 k_2}\right]$$

$$=C_Q C_K \mathbb{E}\left[\left(\frac{1}{n}\sum_{m_2=1}^{n} s_{\alpha;t';m_2} s_{\alpha_2;\tilde{t}'_2;m_2}\right)\left(\frac{1}{n^2}\sum_{k_1,k_2=1}^{n} s_{\alpha;t;k_2} s_{\alpha_2;t'_2;k_1} \mathcal{F}^{(\mathrm{previous})}_{k_1 k_2}\right)\right] = O\left(\frac{1}{n}\right),$$

where in the very last equality we noticed that it results in the second promised form (1.135) of the layer-normalization multiplicative factor, which can be dropped as $O(1/n)$. The rest of the argument is then similar to the one in Appendix A, doing the same algebra with – instead of just one query–key dot product $\widetilde{\Omega}^h_{\alpha;tt'}$ – any odd number of the query–key dot products.

The first contribution has the first promised form (1.134) of the layer-normalization multiplicative factor and also is diagonal in the embedding direction (1.117). Meanwhile, the second contribution has the second promised form (1.135), which we can drop as $O(1/n)$.

D.   Multilayer Perceptron Block

Taking derivatives of the residual MLP-block outputs

$$r_{\alpha;t;i} = \sum_{j=1}^{Mn} X_{ij}\sigma\left(w_{\alpha;t;j}\right) = \sum_{j=1}^{Mn} X_{ij}\sigma\left(\sum_{k=1}^{n} W_{jk}s_{\alpha;t;k}\right), \tag{1.164}$$

with respect to the MLP-block model parameters and the incoming signals, we have

$$\frac{\partial r_{\alpha;t;i}}{\partial X_{jk}} = \delta_{ij}\sigma\left(w_{\alpha;t;k}\right), \tag{1.165}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial W_{jk}} = X_{ij}\sigma'\left(w_{\alpha;t;j}\right)s_{\alpha;t;k}, \tag{1.166}$$

$$\frac{\partial r_{\alpha;t;i}}{\partial s_{\alpha;t;k}} = \sum_{j=1}^{Mn} X_{ij}\sigma'\left(w_{\alpha;t;j}\right)W_{jk}. \tag{1.167}$$

Expectation values for the squared gradient norms are then given by

$$\mathbb{E}\left[\sum_{j=1}^{n}\sum_{k=1}^{Mn}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial X_{jk}}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial X_{jk}}\right] = Mn\delta_{i_1i_2}\mathbb{E}\left[\frac{1}{Mn}\sum_{k=1}^{Mn}\sigma\left(w_{\alpha_1;t_1;k}\right)\sigma\left(w_{\alpha_2;t_2;k}\right)\right] \tag{1.168}$$

$$= \delta_{i_1i_2}\left\{(Mn)\cdot\left[\langle\sigma\left(\widetilde{w}_{\alpha_1;t_1}\right)\sigma\left(\widetilde{w}_{\alpha_2;t_2}\right)\rangle_{C_WF}+O\left(\frac{1}{n}\right)\right]\right\},$$

$$\mathbb{E}\left[\sum_{j=1}^{Mn}\sum_{k=1}^{n}\frac{\partial r_{\alpha_1;t_1;i_1}}{\partial W_{jk}}\frac{\partial r_{\alpha_2;t_2;i_2}}{\partial W_{jk}}\right] = nC_X\delta_{i_1i_2}\mathbb{E}\left\{\left[\frac{1}{Mn}\sum_{j=1}^{Mn}\sigma'\left(w_{\alpha_1;t_1;j}\right)\sigma'\left(w_{\alpha_2;t_2;j}\right)\right]\right. \tag{1.169}$$

$$\left.\times\left(\frac{1}{n}\sum_{k=1}^{n}s_{\alpha_1;t_1;k}s_{\alpha_2;t_2;k}\right)\right\}$$

$$= \delta_{i_1i_2}\left\{(n)\cdot\left[C_X\langle\sigma'\left(\widetilde{w}_{\alpha_1;t_1}\right)\sigma'\left(\widetilde{w}_{\alpha_2;t_2}\right)\rangle_{C_WF}F_{(\alpha_1;t_1)(\alpha_2;t_2)}+O\left(\frac{1}{n}\right)\right]\right\},$$

where we in particular used the factorization formula (1.73) to simplify the expressions. Note that, as promised (1.117), these additive contributions are diagonal in the embedding direction.

For SGD, to ensure order-one additive contributions to the neural tangent kernel, we thus should scale relative learning-rate factors λ_G as

$$\lambda_W = \left(\frac{1}{n}\right)\Lambda_W, \quad \lambda_X = \left(\frac{1}{Mn}\right)\Lambda_X \quad \text{for} \quad \text{SGD}. \tag{1.170}$$

For AdamW, we need to adjust for the magnitude of the per-parameter gradient. For the W weights, each component of the gradient scales as $|g^{\mu}| \sim \sqrt{n/(Mn\cdot n)} = 1/\sqrt{Mn}$ and, for the X

weights, $|g^\mu| \sim \sqrt{Mn/(n \cdot Mn)} = 1/\sqrt{n}$. Thus, to offset these factors in the AdamW optimizer, we set

$$\widetilde{\lambda}_W = \left(\frac{1}{n\sqrt{Mn}}\right)\widetilde{\Lambda}_W, \quad \widetilde{\lambda}_X = \left(\frac{1}{Mn\sqrt{n}}\right)\widetilde{\Lambda}_X \quad \text{for} \quad \text{AdamW}. \tag{1.171}$$

Finally, moving onto the recursive contributions, the cross terms vanish as

$$\mathbb{E}\left[\frac{\partial r_{\alpha;t;i}}{\partial s_{\alpha;t;j}}\mathcal{F}^{(\text{previous})}\right] = 0, \tag{1.172}$$

due to the dangling mean-zero weight X in the signal derivative (1.167), where here $\mathcal{F}^{(\text{previous})}$ is a function of observables from the preceding blocks. For the cumulative contribution, we have

$$\mathbb{E}\left[\sum_{k_1,k_2=1}^{n} \frac{\partial r_{\alpha_1;t_1;i_1}}{\partial s_{\alpha_1;t_1;k_1}} \frac{\partial r_{\alpha_2;t_2;i_2}}{\partial s_{\alpha_2;t_2;k_2}} \mathcal{F}^{(\text{previous})}_{k_1 k_2}\right] \tag{1.173}$$

$$= C_X \delta_{i_1 i_2} \mathbb{E}\left\{\left[\frac{1}{Mn}\sum_{j=1}^{Mn} \sigma'(w_{\alpha_1;t_1;j})\,\sigma'(w_{\alpha_2;t_2;j})\right]\left[\sum_{k_1,k_2=1}^{n} W_{jk_1} W_{jk_2} \mathcal{F}^{(\text{previous})}_{k_1 k_2}\right]\right\}$$

$$= C_X C_W \delta_{i_1 i_2} \mathbb{E}\left\{\left[\frac{1}{Mn}\sum_{j=1}^{Mn} \sigma'(w_{\alpha_1;t_1;j})\,\sigma'(w_{\alpha_2;t_2;j})\right]\left[\frac{1}{n}\sum_{k=1}^{n} \mathcal{F}^{(\text{previous})}_{kk}\right]\right\} + O\left(\frac{1}{n}\right)$$

$$= \delta_{i_1 i_2}\left[C_X C_W \left\langle \sigma'(\widetilde{w}_{\alpha_1;t_1})\,\sigma'(\widetilde{w}_{\alpha_2;t_2})\right\rangle_{C_W F}\right]\mathbb{E}\left[\frac{1}{n}\sum_{k=1}^{n} \mathcal{F}^{(\text{previous})}_{kk}\right] + O\left(\frac{1}{n}\right).$$

Here, in the first equality, we integrated out the X weights; in the second inequality, we integrated out the W weights (where we ignored the interlayer correlations – see Ref. [15] – as subleading $1/n$ corrections); in the last equality we used the factorization formula (1.73) again to simplify the expression. This results in the first promised form (1.134) of the layer-normalization multiplicative factor and also is diagonal in the embedding direction (1.117).

E.   Head Block

Vision: linear classification layer

Taking derivatives of the network outputs

$$z^{(L)}_{\alpha;t;i} = b^{\text{head}}_i + \sum_{j=1}^{n} W^{\text{head}}_{ij} s^{(L-1)}_{\alpha;t;j} \tag{1.174}$$

with respect to the head-block model parameters and the incoming signals, we have

$$\frac{\partial z^{(L)}_{\alpha;t;i}}{\partial b^{\text{head}}_j} = \delta_{ij} \,, \tag{1.175}$$

$$\frac{\partial z^{(L)}_{\alpha;t;i}}{\partial W^{\text{head}}_{jk}} = \delta_{ij} s^{(L-1)}_{\alpha;t;k} \,, \tag{1.176}$$

$$\frac{\partial z^{(L)}_{\alpha;t;i}}{\partial s^{(L-1)}_{\alpha;t;j}} = W^{\text{head}}_{ij} \,. \tag{1.177}$$

Expectation values for the squared gradient norms are then given by

$$\mathbb{E}\left[\sum_{j=1}^{n_{\text{out}}} \frac{\partial z^{(L)}_{\alpha_1;t_1;i_1}}{\partial b^{\text{head}}_j} \frac{\partial z^{(L)}_{\alpha_2;t_2;i_2}}{\partial b^{\text{head}}_j}\right] = \delta_{i_1 i_2} \cdot (1) \,, \tag{1.178}$$

$$\mathbb{E}\left[\sum_{j=1}^{n_{\text{out}}}\sum_{k=1}^{n} \frac{\partial z^{(L)}_{\alpha_1;t_1;i_1}}{\partial W^{\text{head}}_{jk}} \frac{\partial z^{(L)}_{\alpha_2;t_2;i_2}}{\partial W^{\text{head}}_{jk}}\right] = \delta_{i_1 i_2}\mathbb{E}\left[\sum_{k=1}^{n} s^{(L-1)}_{\alpha_1;t_1;k} s^{(L-1)}_{\alpha_2;t_2;k}\right] = \delta_{i_1 i_2} n\,\mathbb{E}\left[\frac{1}{n}\sum_{k=1}^{n} s^{(L-1)}_{\alpha_1;t_1;k} s^{(L-1)}_{\alpha_2;t_2;k}\right]$$

$$= \delta_{i_1 i_2}\left[(n) \cdot F^{(L-1)}_{(\alpha_1;t_1)(\alpha_2;t_2)}\right] \,. \tag{1.179}$$

Note that, as promised (1.117), these additive contributions are diagonal in the embedding direction.

For SGD, to ensure order-one additive contributions to the neural tangent kernel, we thus should scale relative learning-rate factors as

$$\lambda_{b^{\text{head}}} = (1) \cdot \Lambda_{b^{\text{head}}} \,, \quad \lambda_{W^{\text{head}}} = \left(\frac{1}{n}\right)\Lambda_{W^{\text{head}}} \quad \text{for} \quad \text{SGD} \,. \tag{1.180}$$

For AdamW, we need to adjust for the magnitude of the per-parameter gradient. For the biases, each component of the gradient scales as $|g^\mu| \sim 1/\sqrt{n_{\text{out}}}$ and, similarly for the weights, $|g^\mu| \sim \sqrt{n/(n \cdot n_{\text{out}})} = 1/\sqrt{n_{\text{out}}}$. Thus, to offset these factors in the AdamW optimizer, we set

$$\widetilde{\lambda}_{b^{\text{head}}} = \left(\frac{1}{\sqrt{n_{\text{out}}}}\right) \cdot \widetilde{\Lambda}_{b^{\text{head}}} \,, \quad \widetilde{\lambda}_{W^{\text{head}}} = \left(\frac{1}{n\sqrt{n_{\text{out}}}}\right)\widetilde{\Lambda}_{W^{\text{head}}} \quad \text{for} \quad \text{AdamW} \,. \tag{1.181}$$

Finally, moving onto the recursive contributions, as parenthetically mentioned right after our block-to-block recursive formula for neural tangent kernels (1.98), we don't have the skip–residual cross terms in the last block. For the cumulative contribution, we have

$$\mathbb{E}\left[\sum_{j_1,j_2=1}^{n} \frac{\partial z^{(L)}_{\alpha_1;t_1;i_1}}{\partial s^{(L-1)}_{\alpha_1;t_1;j_1}} \frac{\partial z^{(L)}_{\alpha_2;t_2;i_2}}{\partial s^{(L-1)}_{\alpha_2;t_2;j_2}} \mathcal{F}^{(L-1)}_{j_1 j_2}\right] = \sum_{j_1,j_2=1}^{n}\mathbb{E}\left[W^{\text{head}}_{i_1 j_1} W^{\text{head}}_{i_2 j_2} \mathcal{F}^{(L-1)}_{j_1 j_2}\right] \tag{1.182}$$

$$= \delta_{i_1 i_2} C_{\text{head}}\mathbb{E}\left[\frac{1}{n}\sum_{j=1}^{n} \mathcal{F}^{(L-1)}_{jj}\right] \,.$$

This results in the first promised form (1.134) of the layer-normalization multiplicative factor and also is diagonal in the embedding direction (1.117).

<p style="text-align:center">Language: word embedding, transposed (and rescaled)</p>

Taking derivatives of the network outputs

$$z_{\alpha;t;i}^{(L)} = \mathcal{N}_{\text{rescale}} \sum_{j=1}^{n} W_{ji}^{\text{WE}} s_{\alpha;t;j}^{(L-1)}, \tag{1.183}$$

with respect to the word-embedding parameters and the incoming signals, we have

$$\left. \frac{\partial z_{\alpha;t;i}^{(L)}}{\partial W_{jk}^{\text{WE}}} \right|_{\text{head}} = \mathcal{N}_{\text{rescale}} \delta_{ik} s_{\alpha;t;j}^{(L-1)}, \tag{1.184}$$

$$\frac{\partial z_{\alpha;t;i}^{(L)}}{\partial s_{\alpha;t;j}^{(L-1)}} = \mathcal{N}_{\text{rescale}} W_{ji}^{\text{WE}}. \tag{1.185}$$

Note that there was another contribution from the stem block (1.127) and here we are disambiguating the one from the head.

Then – recalling $\mathcal{N}_{\text{rescale}} = \sqrt{1/n}$ (1.82) – we have the expectation value for the squared gradient norm given by

$$\mathbb{E}\left[\sum_{i=1}^{n_{\text{vocab}}} \sum_{j=1}^{n} \left. \frac{\partial z_{\alpha_1;t_1;i_1}^{(L)}}{\partial W_{ji}^{\text{WE}}} \right|_{\text{head}} \left. \frac{\partial z_{\alpha_2;t_2;i_2}^{(L)}}{\partial W_{ji}^{\text{WE}}} \right|_{\text{head}} \right] = \delta_{i_1 i_2} \mathbb{E}\left[\left(\frac{1}{n} \sum_{j=1}^{n} s_{\alpha_1;t_1;j}^{(L-1)} s_{\alpha_2;t_2;j}^{(L-1)} \right) \right] \tag{1.186}$$

$$= \delta_{i_1 i_2} \left[(1) \cdot F_{(\alpha_1;t_1)(\alpha_2;t_2)}^{(L-1)} \right],$$

Note that, as promised (1.117), these additive contributions are diagonal in the embedding direction.[32]

For SGD, to ensure the order-one additive contributions to the neural tangent kernel, we thus should scale the relative learning-rate factor as

$$\lambda_{W^{\text{WE}}} = (1) \cdot \Lambda_{W^{\text{WE}}} \quad \text{for} \quad \text{SGD}. \tag{1.188}$$

---

[32] To be complete, in addition to this additive term (1.186) and the cumulative term, we need to consider stem–head cross gradient terms. As in long footnote 19, the most efficient way to get at the crux of the matter is to work out the toy model $z_{\alpha;t;i}^{(3)} = \mathcal{N}_{\text{rescale}} \sum_{j,k}^{n} \sum_{m=1}^{n_{\text{vocab}}} W_{ji}^{\text{WE}} \left(W_{jk}^{(2)} + \delta_{jk} \right) W_{km}^{\text{WE}} x_{\alpha;t;m}$, for which we find

$$\mathbb{E}\left[\sum_{i=1}^{n_{\text{vocab}}} \sum_{j=1}^{n} \left. \frac{\partial z_{\alpha_1;t_1;i_1}^{(3)}}{\partial W_{ji}^{\text{WE}}} \right|_{\text{head}} \left. \frac{\partial z_{\alpha_2;t_2;i_2}^{(3)}}{\partial W_{ji}^{\text{WE}}} \right|_{\text{stem}} \right] = C_{\text{WE}} \left(1 + \frac{C_{W^{(2)}}}{n} \right) x_{\alpha_1;t_1;i_2} x_{\alpha_2;t_2;i_1}, \tag{1.187}$$

with the leading term coming from the skip path. Similarly to our observations in footnote 19, this contact contribution is zero except for one specific component and the same contribution arises in more general cases.

For AdamW, we need to adjust for the magnitude of the per-parameter gradient $|g^\mu| \sim 1/\sqrt{n}$. Thus, to offset these factors in the AdamW optimizer, we set

$$\widetilde{\lambda}_{W^{\mathrm{WE}}} = \left(\frac{1}{\sqrt{n}}\right) \widetilde{\Lambda}_{W^{\mathrm{WE}}} \quad \text{for} \quad \text{AdamW}. \tag{1.189}$$

Importantly, both of these scalings are consistent with the scalings (1.131) and (1.132) we decided to set for them by analyzing the gradient contribution from the stem block.[33]

Finally, moving onto the recursive contributions, as parenthetically mentioned right after our block-to-block recursive formula for neural tangent kernels (1.98), we don't have the skip–residual cross terms in the last layer. For the cumulative contribution, we have

$$\mathbb{E}\left[\sum_{j_1,j_2=1}^{n} \frac{\partial z_{\alpha_1;t_1;i_1}^{(L)}}{\partial s_{\alpha_1;t_1;j_1}^{(L-1)}} \frac{\partial z_{\alpha_2;t_2;i_2}^{(L)}}{\partial s_{\alpha_2;t_2;j_2}^{(L-1)}} \mathcal{F}_{j_1 j_2}^{(L-1)}\right] = \mathcal{N}_{\mathrm{rescale}}^2 \sum_{j_1,j_2=1}^{n} \mathbb{E}\left[W_{j_1 i_1}^{\mathrm{WE}} W_{j_2 i_2}^{\mathrm{WE}} \mathcal{F}_{j_1 j_2}^{(L-1)}\right] \tag{1.190}$$

$$= \delta_{i_1 i_2} C_{\mathrm{WE}} \mathbb{E}\left[\frac{1}{n} \sum_{j=1}^{n} \mathcal{F}_{jj}^{(L-1)}\right].$$

This results in the first promised form (1.134) of the layer-normalization multiplicative factor and also is diagonal in the embedding direction (1.117).[34]

---

[33] Further consistently, as discussed in footnote 29 for the stem block, the scaling of the relative AdamW learning-rate factor $\widetilde{\lambda}_{W^{\mathrm{WE}}}$ with the vocabulary size n_{vocab} is subtle – if not subtler – for the head block. Specifically, when using the cross-entropy loss, the gradient respect to the word-embedding parameter W_{ij}^{WE} gets a factor of $p_{\alpha;t;j} - q_{\alpha;t;j}$ where $p_{\alpha;t;j}$ is the softmax (in the embedding direction j) distribution of the output $z_{\alpha;t;j}^{(L)}$ and $q_{\alpha;t;j}$ is the target (typically one-hot) distribution for the task at hand. There is then a subtlety even for a small batch size $|\mathcal{A}_t|$ as then most of these gradients are small but nonzero (roughly of order $\sim 1/n_{\mathrm{vocab}}$ at initialization) and their noisy signals get amplified in the AdamW optimizer – unless one decides to explicitly focus on the top-few components in the j direction or implicitly regularize with appropriately chosen ϵ.

[34] Continuing – and ending – our usual operation of hiding the subtle interlayer correlation due to the stem–head weight tying, we here simply note that the first-block neural tangent kernel $\widehat{H}_{(\alpha_1;t_1;i_1)(\alpha_2;t_2;i_2)}^{(1)}$ doesn't depend on the word-embedding parameters (1.129), and using the leading-order layer-normalization mnemonics (1.140), there is nothing to worry about here.

II.   Practical Applications

<div align="right">

Such heroic nonsense.

Megatron

</div>

In Part I of this note, we've developed an effective theory of Transformers at leading order. Such an analysis in particular suggested proper width scalings of initialization and training hyperparameters that collectively ensure benevolent limiting behaviors when widening Transformers. These suggestions, however, would be theoretical nonsense if they don't better the practice in one way or the other. In this Part II, we'll thus put our theoretical suggestions to practical tests, training both Vision and Language Transformers.

In §II 1, we'll train Vision Transformers for an image classification task. For them, initialization hyperparameters are often – though not always – scaled well with width, so we'll focus on comparing the standard uniform learning rate – that is, $\widetilde{\lambda}_G = 1$ in our language – with our non-uniform neural-tangent scalings (1.108)–(1.116) of learning rates. In our training setup, we'll find that the neural-tangent scaling strategy can not only improve the model performance but also reduce the frequency of mid-training spikes when compared with the uniform learning rate.[35]

In §II 2, we'll pretrain Language Transformers for a span denoising task. For them, we'll make three changes to the standard uniform (non-)scaling of hyperparameters: *(i)* the initialization hyperparameter for the word-embedding parameters will be cranked up from the standard $C_{\mathrm{WE}} = (0.02)^2$ to $C_{\mathrm{WE}} = 1$; *(ii)* concomitantly the network output is rescaled with $\mathcal{N}_{\mathrm{rescale}} = 1/\sqrt{n}$ instead of the standard non-rescaling $\mathcal{N}_{\mathrm{rescale}} = 1$; and *(iii)* again the relative learning-rate factors $\widetilde{\lambda}_G$ are set according to the neural-tangent scaling strategy. Here, our results are slightly more mixed than those in §II 1, but we'll nonetheless include them as encouraging anecdata for larger scales.

Programming note 1: mirroring our theoretical treatment in Part I of this note, for all models, we'll drop bias parameters from MHSA and MLP blocks – as practiced, e.g., by some large language models such as PaLM [4] – and keep element-wise affine parameters in normalization layers fixed.
Programming note 2: for all experiments, we'll use PyTorch's automatic mixed precision [31], to better reflect the current practices and speed up training.

---

[35] We'll also test the maximal-update scaling strategy [24] and the hybrid of neural-tangent and maximal-update scaling strategies [23] for Vision Transformers. We'll find that they both reduce the frequency of mid-training spikes and – among all the four scaling strategies tested herein – the hybrid scaling strategy yields the best performance while the maximal-update scaling strategy yields the worst.

1. Image Classification with Encoder-Only Transformers

Here we'll train encoder-only Vision Transformers with an image classification objective.[36] We'll describe our dataset and task in §II 1 A, architecture hyperparameters in §II 1 B, initialization hyperparameters in §II 1 C, and optimizer and training hyperparameters in §II 1 D. In §II 1 E, we'll then compare the standard uniform scaling of learning rates with the neural-tangent scalings.

A. Dataset and Task

As for the dataset, we use the ILSVRC challenge version [33] of the ImageNet-1k dataset [18], containing RGB images of 1000 different categories. The dataset is split into the training set containing 1,281,167 images and the validation set containing 50,000 images. Both for the training set and for the validation set, as is customary, each image is preprocessed by subtracting the $\texttt{mean} = [0.485, 0.456, 0.406]$ and dividing by the standard deviation $\texttt{std} = [0.229, 0.224, 0.225]$.

As for the training objective, we augment images first with the by-now standard random resizing and cropping to 224-by-224 and random horizontal flip, then with RandAugment [34] ($n = 2$, $m = 9$), one of Mixup [35] ($\alpha = 0.8$) or CutMix [36] ($\alpha = 1.0$) chosen with equal probability 50%, and apply a single Random Erasing [37] rectangle with probability 25% and with values set randomly per-pixel within the block.[37] We then optimize the cross-entropy loss between the model outputs and target labels, with label smoothing [39] of 0.1.

At evaluation, images in the validation set are scaled to 256-by-256 and center cropped to 224-by-224, without any data argumentation. We then evaluate the top-one prediction accuracy.

B. Architecture Hyperparameters

Overall, our architectural design follows the original Vision Transformers [14] with one difference: as noted in the introduction of this Part II, we turn off all the bias parameters in MHSA and MLP blocks [4] and don't train element-wise affine parameters in normalization layers.

With that difference in mind, our Vision-Transformer architecture is exactly that described in

---

[36] Following Ref. [32], all models were trained by using the codebase available at `https://github.com/facebookresearch/convnext`; to it, we added our own model code for Vision Transformers and also introduced the relative learning-rate factors $\widetilde{\lambda}_G$'s to the AdamW optimizer code.

[37] As a full specification for RandAugment, we use `rand-m9-mstd0.5-inc1` in `timm` [38], which applies magnitude noise of the standard deviation $\texttt{std} = 0.5$.

the *Vision* track of §I 0, with the patch dimension $n_{\text{patch}} = 16 \cdot 16 \cdot 3 = 768$, the number of tokenized patches $T = 196$, the embedding dimension $n = 768$, the normalization layer regularization $\epsilon = 10^{-6}$, the number of MHSA heads $H = 12$, the MLP multiplier $M = 4$, `GELU` activation functions in the MLP blocks, and output dimension $n_{\text{out}} = 1000$; in the bulk, 12 encoders – each encoder consisting of one bidirectional MHSA block followed by one MLP block – are stacked.

To make sure we are on the same footing, please check that those architectural choices result in $P = 768 \cdot 768 + 196 \cdot 768 + 768^2 \cdot (4 + 2 \cdot 4) \cdot 12 + 768 \cdot 1000 + 1000 \approx 86 \cdot 10^6$ model parameters.

C.   Initialization Hyperparameters

As for the initialization, the practices vary in the literature. In our experiments, for both standard and neural-tangent runs, we decided to more or less follow the PyTorch default at some points in time, that is, in the bulk, for the weights in the MHSA blocks we use mean-zero uniform distributions with $C_Q = C_K = C_V = 2/(3+1) = 1/2$ and $C_U = 1/3$ and for the weights in the MLP blocks we use mean-zero uniform distributions with $C_W = 2/(4+1) = 2/5$, and $C_X = 2 \cdot 4/(4+1) = 8/5$.[38] These are acceptable as order-one numbers. As for the stem and head blocks, we simply use a mean-zero normal distribution with $C_{\text{patch}} = C_{\text{head}} = 1$ and $b_i^{\text{head}} = 0$ while $C_{\text{PE}} = (0.02)^2$ for the positional-embedding parameters as practiced sometimes.[39]

D.   Optimizer and Training Hyperparameters

Overall, our training recipe almost exactly follows the one used in Ref. [32] – except that we here ablate *(i)* `DropPath`s [40] and *(ii)* exponential moving averaging [41], in order to make the comparison slightly less confounded on their respective hyperparameters.

As for the optimizer, we use AdamW [28, 29] (1.91)–(1.94) with $(\beta_1, \beta_2, \epsilon) = (0.9, 0.999, 10^{-8})$; also note that weight decay is *not* applied to the bias parameters in the head block [32].

As for the learning schedule, we use a batch size of 4096 – which results in 312 iterations per epoch – with linear warmup [42] for the first 20 epochs, followed by a cosine learning-rate decay [43]

---

[38] The most of them derives directly from applying the standard Xavier initialization – where we in particular note that Q–K–V weights are often coded as one linear layer instead of three linear layers – while for U the factor of $1/3$ in the covariance derives from uniformly distributing weights in the interval $[-1/\sqrt{n}, 1/\sqrt{n}]$.

[39] After our theoretical exposition, this last choice may look unnaturally small but – since the positional-embedding parameters act like bias parameters (i.e., not multiplicative but additive) and hence we could in principle set this hyperparameter to any order-one number including $C_{\text{PE}} = 0$ – we adopt it as acceptable. In general, we tried to minimize the amount of changes with respect to the literature.

over the next 280 epochs, that is,

$$
\eta_t = \begin{cases} \mathtt{lr} \cdot \left(\frac{t}{312 \cdot 20} \right) & \text{for} \quad t \le 312 \cdot 20 \,, \\[2mm] \mathtt{lr}_{\min} + (\mathtt{lr} - \mathtt{lr}_{\min}) \cdot \left\{ \frac{1}{2} + \frac{1}{2} \cos \left[\frac{(t - 312 \cdot 20)\pi}{312 \cdot 280} \right] \right\} & \text{for} \quad t > 312 \cdot 20 \,, \end{cases} \tag{2.1}
$$

with $\mathtt{lr}_{\min} = 10^{-6}$.[40] We train the models for the full 300 epochs; we'll discuss the settings of the overall learning rate $\mathtt{lr}$ and the weight decay $\mathtt{wd}$ shortly in §II 1 E.

As for the per-group learning-rate factors $\widetilde{\lambda}_G$'s, for the standard runs, we use the standard uniform scaling $\widetilde{\lambda}_G = 1$, while for the neural-tangent runs, we follow our theoretical suggestions (1.108)–(1.116), except that we ignore the factor of $M = 4$: see Table II 1 E.[41] Overall, with $n = 768 \sim 1000$, our neural-tangent scaling essentially boils down to cranking up the learning rate for positional-embedding parameters (and bias parameters in the head block) by a factor of $n \sim 1000$ with respect to other model parameters.

E.   Comparison of Scaling Strategies

We here compare the performances of the models trained with the standard uniform learning rate against the ones trained with the neural-tangent scalings of learning rates: see Table II 1 E below for a concise summary.

AdamW–standard runs			AdamW–neural-tangent runs		
	initial std	relative lr factors		initial std	relative lr factors
patchify W_{ij}^{patch}	$768^{-\frac{1}{2}}$	1	patchify W_{ij}^{patch}	$768^{-\frac{1}{2}}$	$768^{-1} \cdot n^{-\frac{1}{2}}$
positional embedding $b_{t;i}^{\text{PE}}$	0.02	1	positional embedding $b_{t;i}^{\text{PE}}$	0.02	$n^{-\frac{1}{2}}$
Q–K–V–U–W–X weights	$\sharp n^{-\frac{1}{2}}$	1	Q–K–V–U–W–X weights	$\sharp n^{-\frac{1}{2}}$	$n^{-1} \cdot n^{-\frac{1}{2}}$
head weights W_{ij}^{head}	$n^{-\frac{1}{2}}$	1	head weights W_{ij}^{head}	$n^{-\frac{1}{2}}$	$n^{-1} \cdot 1000^{-\frac{1}{2}}$
head biases b_i^{head}	0	1	head biases b_i^{head}	0	$1000^{-\frac{1}{2}}$
AdamW–$\sqrt{\text{neural-tangent} \cdot \text{maximal-update}}$ runs			AdamW–maximal-update runs		
	initial std	relative lr factors		initial std	relative lr factors
patchify W_{ij}^{patch}	$768^{-\frac{1}{2}}$	$768^{-1} \cdot n^{-\frac{1}{4}}$	patchify W_{ij}^{patch}	$768^{-\frac{1}{2}}$	768^{-1}
positional embedding $b_{t;i}^{\text{PE}}$	0.02	$n^{-\frac{1}{4}}$	positional embedding $b_{t;i}^{\text{PE}}$	0.02	1.0
Q–K–V–U–W–X weights	$\sharp n^{-\frac{1}{2}}$	$n^{-1} \cdot n^{-\frac{1}{4}}$	Q–K–V–U–W–X weights	$\sharp n^{-\frac{1}{2}}$	n^{-1}
head weights W_{ij}^{head}	$n^{-\frac{3}{4}}$	$n^{-1} \cdot 1000^{-\frac{1}{2}}$	head weights W_{ij}^{head}	n^{-1}	$n^{-1} \cdot 1000^{-\frac{1}{2}}$
head biases b_i^{head}	0	$1000^{-\frac{1}{2}}$	head biases b_i^{head}	0	$1000^{-\frac{1}{2}}$

---

[40] In retrospect, the effect of $\mathtt{lr}_{\min} = 10^{-6}$ kicks in roughly only in the last ~ 5 epochs for the standard scaling runs (at $\mathtt{lr}^\star = 0.001$) and in the last ~ 10 iterations for the neural-tangent scaling runs (at $\mathtt{lr}^\star = 20$).

[41] To closely follow the theoretical suggestions, we should've set $\widetilde{\lambda}_W = 1/(n\sqrt{Mn})$ and $\widetilde{\lambda}_X = 1/(Mn\sqrt{n})$ but we instead caved to $\widetilde{\lambda}_W = \widetilde{\lambda}_X = 1/(n\sqrt{n})$, given that it is easier to implement for larger language models – where sharding and flattening sometimes complicate the coding. However, if anyone scales M large in the future, then $\widetilde{\lambda}_W$ and $\widetilde{\lambda}_X$ should be scaled properly with M.

In order to provide rigorous comparisons – within reason – we tune the overall learning rate `lr` and the weight decay `wd` for each scaling strategy: see Fig. 1. Specifically, we search for the optimal $(\mathtt{lr}, \mathtt{wd})$ in $\log_2$ grid space; once we find the candidate optimum $(\mathtt{lr}^\star, \mathtt{wd}^\star)$, we make sure that the top-one validation accuracies are lower for $(\mathtt{lr}, \mathtt{wd}) = (2 \cdot \mathtt{lr}^\star, \mathtt{wd}^\star)$, $(\mathtt{lr}^\star, 2 \cdot \mathtt{wd}^\star)$, $(2 \cdot \mathtt{lr}^\star, 0.5 \cdot \mathtt{wd}^\star)$, $(0.5 \cdot \mathtt{lr}^\star, 2 \cdot \mathtt{wd}^\star)$, $(\mathtt{lr}^\star, 0.5 \cdot \mathtt{wd}^\star)$, $(0.5 \cdot \mathtt{lr}^\star, \mathtt{wd}^\star)$.[42] Once the optimal pair $(\mathtt{lr}^\star, \mathtt{wd}^\star)$ is found for each scaling strategy, we ran the experiments for three different seeds: see Fig. 2.
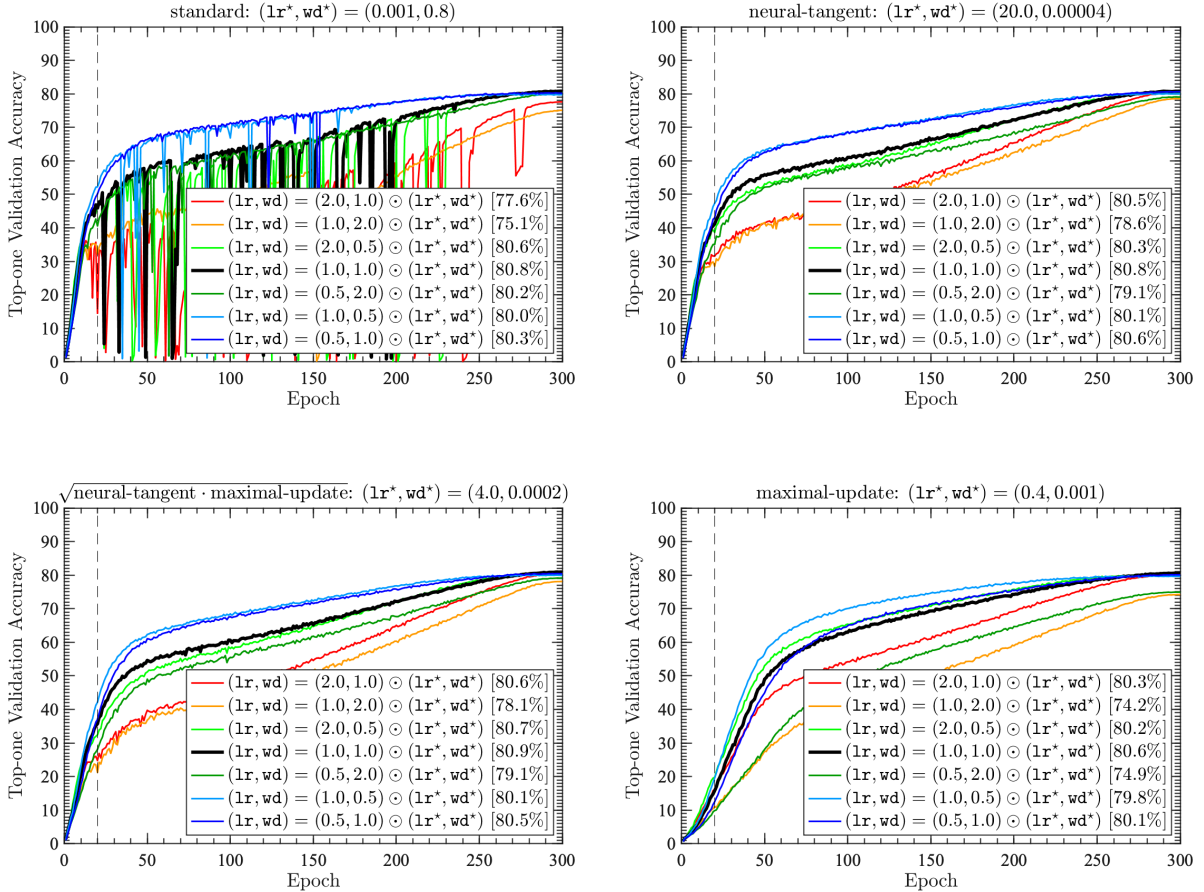
FIG. 1: Hyperparameter searches for standard uniform (top left), neural-tangent (top right), hybrid neural-tangent–maximal-update (bottom left), and maximal-update (bottom right) scaling strategies for Vision Transformers trained by $\text{AdamW}_{(\beta_1, \beta_2, \epsilon) = (0.9, 0.999, 1\text{e-}8)}$. For each, the top-one validation accuracy on the ImageNet-1k dataset is plotted as a function of training epochs. In the legend, we record the training hyperparameter pair $(\mathtt{lr}, \mathtt{wd})$ [with the max top-one validation accuracy along each trajectory].

---

[42] To save compute, we here omit $(\mathtt{lr}, \mathtt{wd}) = (2 \cdot \mathtt{lr}^\star, 2 \cdot \mathtt{wd}^\star)$, $(0.5 \cdot \mathtt{lr}^\star, 0.5 \cdot \mathtt{wd}^\star)$ because, if doubling/halving `lr` and `wd` *each individually* degrades the performance, then we expect that doubling/halving *both together* degrades the performance as well.

Overall, in our experimental setup, we see that the neural-tangent scaling strategy not only statistically improves the performance but also reduces the frequency of mid-training spikes when compared to the standard scaling strategy.[43]

We additionally note that, for the neural-tangent runs, the effective overall learning rate, $\mathtt{lr} \cdot \widetilde{\lambda}_G$, for Q–K–V–U–W–X weights at optimality is given by $20 \cdot 768^{-3/2} \approx 0.94 \cdot 10^{-4} \sim 0.001$, which is essentially equal to the optimal learning rate for the standard runs; we further note that $\mathtt{lr}^\star \cdot \mathtt{wd}^\star = 8 \cdot 10^{-4}$ for both scaling strategies. We thus conclude that the stability and performance improvement of the neural-tangent scaling strategy essentially arises from *increasing* the learning rate for the positional-embedding parameters by a factor of $n \sim 1000$.[44]
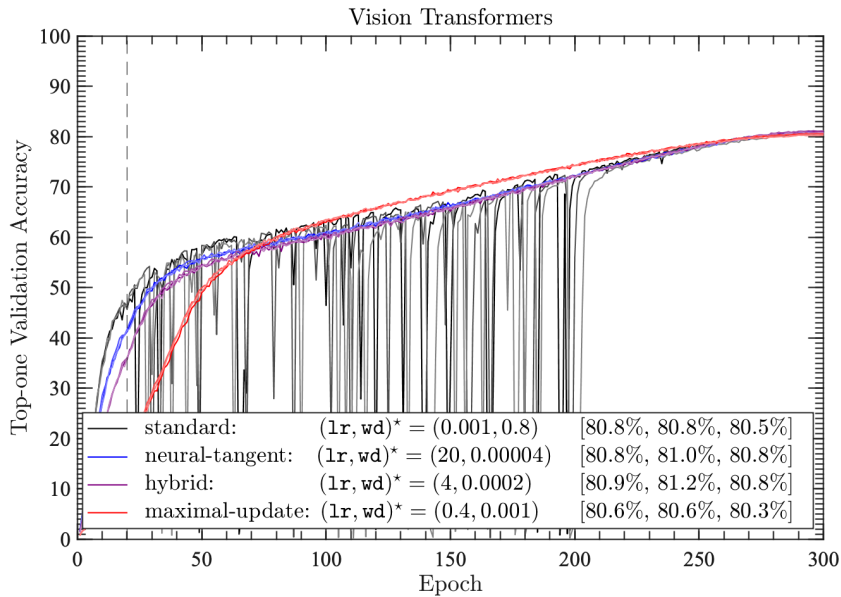
FIG. 2: Comparison of the standard (black), neural-tangent (blue), hybrid neural-tangent–maximal-update (purple), and maximal-update (red) scaling strategies for Vision Transformers trained by $\mathrm{AdamW}_{(\beta_1,\beta_2,\epsilon)=(0.9,0.999,1\text{e-}8)}$. For each, the validation accuracy is plotted as a function of training epochs, for three different seeds (different whiteness). In the legend, we record the optimal training hyperparameter pair $(\mathtt{lr}^\star, \mathtt{wd}^\star)$ for each scaling strategy [with the max top-one validation accuracy along each trajectory].

---

[43] We also observe the same reduction of spikes for the neural-tangent–maximal-update hybrid and maximal-update scaling strategies: see Table II 1 E and Figs. 1 and 2. Incidentally, in our experimental setup, among four scaling strategies, the hybrid scaling strategy yields the best performance while the maximal-update scaling strategy yields the worst performance.

[44] Technically speaking, we are also increasing the learning rate for the biases in the head block. However, we also performed experiments with the biases in the head block turned off and observed the qualitatively similar results, so we believe that the claimed benefits arise from our proper treatment of the positional-embedding parameters.

2. Span Denoising with Encoder–Decoder Transformers

Here we'll pretrain encoder–decoder Language Transformers with a span denoising objective.[45] We'll describe our dataset and task in §II 2 A, architecture hyperparameters in §II 2 B, initialization hyperparameters in §II 2 C, and optimizer and training hyperparameters in §II 2 D. In §II 2 E, we'll then compare the standard uniform scalings of hyperparameters with the neural-tangent scalings.

A. Dataset and Task

As for the dataset, following Ref. [45], we use the dataset that combines the corpora used by RoBERTa [47] – which by themselves are amalgamation of BookCorpus [48] plus English Wikipedia [49], CommonCrawl News [50], OpenWebText [51, 52], and Stories [53] – with the English subset of the CC-100 corpus [54] and a dump of pushshift.io Reddit compiled by a third party [55, 56]. We employ the same byte-pair encoding tokenizer as GPT-2 [52]. Both for the training set and for the validation set, each input sample α is a set of sentences taken from these documents with its total sequence length up to 512, always starting and ending at the beginning of some sentence and at the end of some – same or different – sentence, where we also add a `[start]` token and an `[end]` token, respectively; we further pad `[pad]` tokens to make the full sequence length always be 514.

As for the training objective, following Ref. [44], we first corrupt documents by randomly permuting sentences and then repeatedly infilling a span of tokens with a `[MASK]` token – where each span length is randomly sampled from a Poisson distribution with its mean and variance 3 – until 30% of input sequence tokens are corrupted. We then optimize the cross-entropy loss between the model output and the original uncorrupted sample – discarding the `[pad]`s.

At evaluation, we corrupt documents and measure the cross-entropy in the same way, albeit now with the validation set.

B. Architecture Hyperparameters

Overall, our architectural design follows those of BART-large [44] and R2C2 [45], again with one difference: as noted in the introduction of this Part II, we turn off all the bias parameters in MHSA

---

and MLP blocks [4] and don't train element-wise affine parameters in normalization layers.[46]

With that difference in mind, our BART-large (R2C2) architecture is *almost* – see the next paragraph – that described in the *Language* track of §I 0, with the vocabulary size $n_{\text{vocab}} = 50265$, the sequence length $T = 514$, the width $n = 1024$ (2048), the normalization layer regularization $\epsilon = 10^{-5}$, the number of attention heads $H = 16$ (32), the MLP multiplier $M = 4$, and GELU activation functions in the MLP blocks; in the bulk, 12 (22) encoders – each encoder consisting of one bidirectional MHSA block followed by one MLP block – are stacked first and then 12 (22) decoders – each decoder consisting of one bidirectional MHSA block, followed by one soon-to-be-described encoder–decoder multi-headed *mixed*-attention block, further followed by one MLP block – are stacked.

Compared to what's outlined in the *Language* track of §I 0, there are a few additional gadgets for the BART-large and R2C2 architecture:

- Right after the stack of the encoder blocks and right before the stack of the decoder blocks, one additional set of positional-embedding parameters is placed, which is often considered as a stem block in the decoder stack.

- In addition to the normalization layers applied at the beginning of each residual path, there are three additional ones placed *(1)* right after the first positional-embedding parameters in the stem block of the encoder stack, *(2)* right after the stack of all the encoder blocks, and *(3)* right after the second positional-embedding parameters in junction just mentioned above, i.e., in the stem block of the decoder stack.

- In the encoder–decoder multi-headed mixed-attention blocks in the decoders, instead of the usual query–key–value vectors (1.13)–(1.15) for bidirectional/masked MHSA blocks, we use

$$q^h_{\alpha;t;c} \equiv \sum_{i=1}^{n} Q^h_{ci} s_{\alpha;t;i} \,, \tag{2.2}$$

$$k^h_{\alpha;t;c} \equiv \sum_{i=1}^{n} K^h_{ci} \widetilde{s}_{\alpha;t;i} \,, \tag{2.3}$$

$$v^h_{\alpha;t;c} \equiv \sum_{i=1}^{n} V^h_{ci} \widetilde{s}_{\alpha;t;i} \,, \tag{2.4}$$

---

[46] Like the original Language Transformer [13], the original BART-large [44] – but not R2C2 [45] – placed normalization layers at the end of each block; as mentioned in §I 0, as is standard nowadays, we here place normalization layers at the beginning of each residual path for both BART-large and R2C2.

where – as before – $s_{\alpha;t;i}$ is the layer-normalized signal from the preceding (masked MHSA) block while – *not* as before – $\widetilde{s}_{\alpha;t;i}$ is the layer-normalized signal from right before it enters the decoder stack, i.e., right before the second positional-embedding parameters are added.[47] Then, similarly to the bidirectional MHSA, we use the softmax (1.17) over all the tokens except that we mask out the [pad]ed tokens from considerations.

To make sure we are on the same footing, please check that those architectural choices result in $P = 50265 \cdot 1024 + 514 \cdot 1024 + 1024^2 \cdot (4 + 2 \cdot 4) \cdot 12 + 514 \cdot 1024 + 1024^2 \cdot (4 + 4 + 2 \cdot 4) \cdot 12 \approx 405 \cdot 10^6$ model parameters for BART-large and $P = 50265 \cdot 2048 + 514 \cdot 2048 + 2048^2 \cdot (4 + 2 \cdot 4) \cdot 22 + 514 \cdot 2048 + 2048^2 \cdot (4 + 4 + 2 \cdot 4) \cdot 22 \approx 2.7 \cdot 10^9$ model parameters for R2C2.

C.   Initialization Hyperparameters

As for the initialization, for standard runs, following Refs. [44, 45], we initialize all the model paprameters uniformly across all the groups by using the mean-zero normal distribution with the standard deviation std $= 0.02$.[48] For BART-large with $n = 1024$ and $M = 4$, this means in our language that $C_Q = C_K = C_V = C_U = C_W = 0.4096$, $C_X = 1.6384$, and $C_{\text{PE}} = C_{\text{WE}} = 0.02$; for R2C2 with $n = 2048$ and $M = 4$, this means that $C_Q = C_K = C_V = C_U = C_W = 0.8192$, $C_X = 3.2768$, and $C_{\text{PE}} = C_{\text{WE}} = (0.02)^2$. These are mostly acceptable as order-one numbers – however aesthetically displeasing they are – *except* that the covariance $C_{\text{WE}} = (0.02)^2 = 0.0004$ for word-embedding parameters is far from being of order one and hence, for the neural-tangent runs, we adjust it to $C_{\text{WE}} = 1$ and, concomitantly, change $\mathcal{N}_{\text{rescale}} = 1$ for the standard runs to $\mathcal{N}_{\text{rescale}} = \sqrt{1/n}$ for the neural-tangent runs: attend to green color in Table II 2 E.[49]

D.   Optimizer and Training Hyperparameters

Overall, our training recipe mostly follows the ones used in Refs. [44, 45].[50]

---

[47] In theory, the scaling analysis stays intact and the inner products $(1/n)\sum_{i=1}^{n} s_{\alpha_1;t_1;i} s_{\alpha_2;t_2;i}$ just get replaced by $(1/n)\sum_{i=1}^{n} s_{\alpha_1;t_1;i} \widetilde{s}_{\alpha_2;t_2;i}$ in some places and $(1/n)\sum_{i=1}^{n} \widetilde{s}_{\alpha_1;t_1;i} \widetilde{s}_{\alpha_2;t_2;i}$ in other places.

[48] To be precise, we zero-initialize the word-embedding parameters W_{ij}^{WE} when j corresponds to the [pad] token.

[49] In contrast, we take $C_{\text{PE}} = (0.02)^2$ as acceptable for the same reason as described in footnote 39: the positional-embedding parameters act like bias parameters (i.e., not multiplicative but additive) and hence we could in principle set this hyperparameter to any order-one number including $C_{\text{PE}} = 0$.

[50] The most significant difference is that, from compute considerations for our study, our warmup and *target* training length of 5,000 iterations and 200,000 iterations are respectively shorter than 10,000 (15,000) iterations and 500,000 (500,000) iterations used in Ref. [44] (Ref. [45]).

As for the optimizer, we use AdamW [28, 29] (1.91)–(1.94) with $(\beta_1, \beta_2, \epsilon) = (0.9, 0.998, 10^{-6})$; we also set `--dropout` and `--attention-dropout` to 0.1 and `--clip-norm` to 0.1.

As for the learning schedule, we use an effective batch size of 512 with linear warmup [42] for the first 5,000 iterations, followed by a linear learning-rate decay over the next 195,000 iterations, that is,

$$\eta_t = \begin{cases} \texttt{lr} \cdot \left(\frac{t}{5000} \right) & \text{for} \quad t \leq 5000\,, \\ \texttt{lr} \cdot \left\{ \frac{200000-t}{195000} \right\} & \text{for} \quad t > 5000\,. \end{cases} \tag{2.5}$$

We train the models for the full 200,000 iterations for BART-large models while we stop the runs at 15,000 iterations for R2C2 models due to the compute consideration; we'll discuss the settings of the overall learning rate `lr` and the weight decay `wd` shortly in §II 2 E.

As for the per-group learning-rate factors $\widetilde{\lambda}_G$'s, for the standard runs, we use the standard uniform scaling $\widetilde{\lambda}_G = 1$, while for the neural-tangent runs, we follow our theoretical suggestions (1.108)–(1.116), except, again, that we ignore the factor of $M = 4$: attend to blue color in Table II 2 E. Overall, our neural-tangent scalings of learning rates essentially boil down to cranking up the learning rates for both word-embedding and positional-embedding parameters by a factor of the width – $n = 1024$ for BART-large and $n = 2048$ for R2C2 – with respect to other model parameters.

E.   Comparison of Scaling Strategies

We here compare the performances of the models trained with the standard uniform scalings of hyperparameters against the ones trained with the neural-tangent scalings of hyperparameters: see Table II 2 E below for a concise summary.

AdamW–standard runs				AdamW–neural-tangent runs			
	initial `std`	rescale	relative `lr` factors		initial `std`	rescale	relative `lr` factors
word embedding W_{ij}^{WE}	0.02		1	word embedding W_{ij}^{WE}	1		$n^{-\frac{1}{2}}$
positional embedding $b_{t;i}^{\text{PE}}$	0.02		1	positional embedding $b_{t;i}^{\text{PE}}$	0.02		$n^{-\frac{1}{2}}$
Q–K–V–U–W–X weights	0.02		1	Q–K–V–U–W–X weights	$0.02_{\text{(yuck)}}$		$n^{-\frac{3}{2}}$
output $\mathcal{N}$		1.0		output $\mathcal{N}$		$n^{-\frac{1}{2}}$	

In order to provide semi-rigorous comparisons – within reason – for BART-large, we tune the overall learning rate `lr` but now that the models are five-fold larger than Vision Transformers treated in §II 1, we won't tune the weight decay `wd`. Specifically, for the standard runs, we fix the weight decay at `wd` $= 0.01$ and search for the optimal overall learning rate `lr` in $\log_2$ grid

space; for the neural-tangent runs, given the remarks at the end of §II 1 E, we shift the search space to $(\texttt{lr},\texttt{wd}) \to (n^{3/2}\texttt{lr}, n^{-3/2}\texttt{wd})$, now with fixed $\texttt{wd} = 0.01 \cdot 1024^{-3/2} \approx$ 3e-7: see Fig. 3. In fact, under this map, we see that the optimal $\texttt{lr}$ match between the standard uniform scaling strategy – for which we find $\texttt{lr}^\star = $ 1e-3 – and neural-tangent scaling strategy – for which we find $\texttt{lr}^\star = 1024^{3/2} \cdot$ 1e-3 $= 32.768$. We then compare the optimal runs: see Fig. 4.

For R2C2, since the models are seven-fold larger than BART-large – and thirty-fold larger than Vision Transformers – we won't even try to tune any of the hyperparameters. Instead, for the standard run, we follow the $(\texttt{lr},\texttt{wd}) = (\text{7e-4}, 0.01)$ used in Ref. [45] – albeit here with shorter warmup and target training lengths – and, for the neural-tangent run, we shift it to $(n^{3/2}\texttt{lr}, n^{-3/2}\texttt{wd}) = (2048^{3/2} \cdot \text{7e-4}, 2048^{-3/2} \cdot 0.01) \approx (64.88, 1.08\text{e-7})$, motivated by our tuning results for both Vision Transformers and BART-large models: see Fig. 5

For BART-large, the standard uniform scaling strategy seems to yield a marginally better model than the neural-tangent scaling strategy, though the latter strategy appears more robust against the change in the hyperparameter – at least against lowering the global learning rate $\texttt{lr}$ – and it would have been nice to run more experiments with distinct seeds. On a more encouraging note, for R2C2, the neural-tangent scaling strategy seems to improve the convergence, at least at the early stage in training. It would be interesting to further extend the experiments to the scale where Language Transformers start to suffer from mid-training spikes and see if the neural-tangent scaling strategy ameliorates them and if it results in further performance boost.
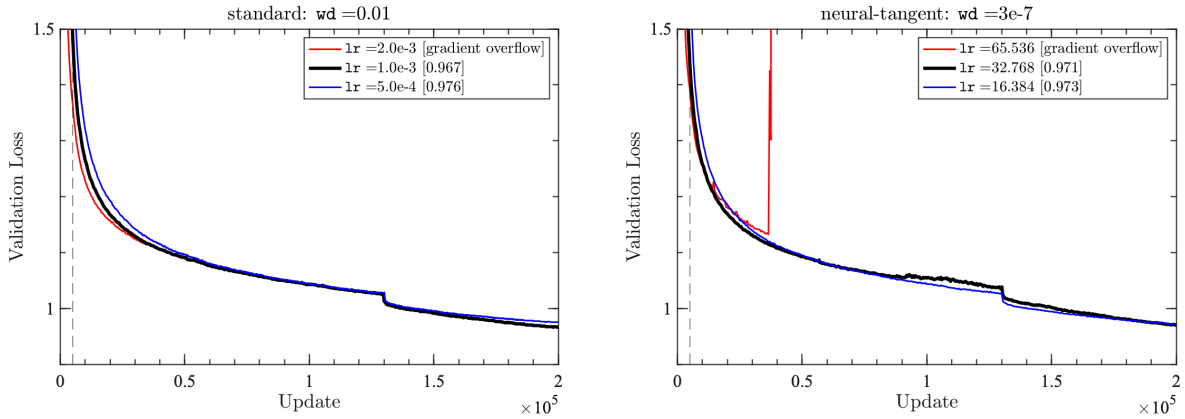
FIG. 3: Hyperparameter searches for standard uniform (left) and neural-tangent (right) scaling strategies for BART-large trained by $\text{AdamW}_{(\beta_1,\beta_2,\epsilon)=(0.9,0.998,1\text{e-}6)}$. For each, the validation loss is plotted as a function of training updates. In the legend, we record the global learning rate `lr` [with the minimum validation loss along each trajectory]. Note that *(i)* for both scaling strategies, the higher `lr` runs (red) experienced gradient overflow, *(ii)* lowering `lr` (going from black to blue) degrades the model performance more for the standard uniform scaling strategy than for the neural-tangent scaling strategy, and *(iii)* for the neural-tangent scaling strategy, the `lr` $= 32.768$ run – despite its "hiccup" at around 90,000 iterations – caught up with the `lr` $= 16.384$ run in the end.
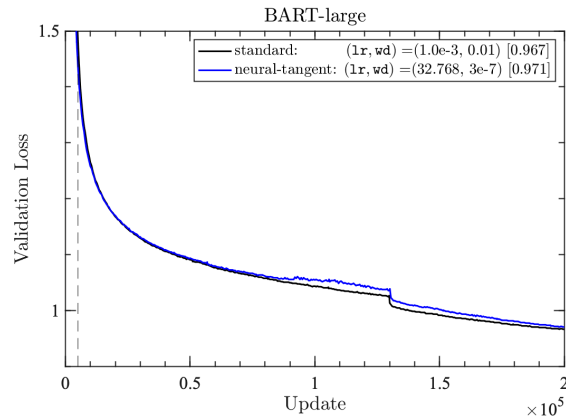
FIG. 4: Comparison of the standard uniform (black) and neural-tangent (blue) scaling strategies for BART-large trained by $\text{AdamW}_{(\beta_1,\beta_2,\epsilon)=(0.9,0.998,1\text{e-}6)}$. For each, the validation loss is plotted as a function of training updates. In the legend, we record the selected training hyperparameter pair (lr, wd) for each scaling strategy [with the minimum validation loss along each trajectory].
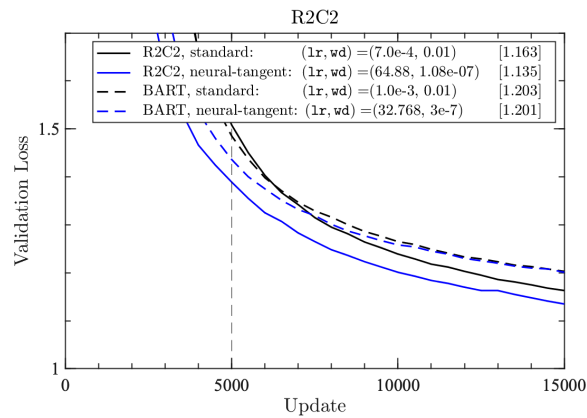

FIG. 5: Comparison of the standard uniform (black) and neural-tangent (blue) scaling strategies for R2C2 trained by $\text{AdamW}_{(\beta_1,\beta_2,\epsilon)=(0.9,0.998,1\text{e-}6)}$; for further comparison, we also include the similar runs for BART-large (dashed). For each, the validation loss is plotted as a function of training updates. In the legend, we record the zero-shotted training hyperparameter pair (lr, wd) for each scaling strategy [with the minimum validation loss along each trajectory].

Acknowledgements

On a technical note, to bring to the fore the future direction buried in footnote 29, we collectively acknowledge that – given our time and energy constraints – we didn't properly take into account the interplay between the batch size $|\mathcal{A}_t|$, sequence length T, vocabulary size n_{vocab}, and token distribution, but nonetheless hope that our work will be a helpful anecdatapoint for future work.

A. Attention on Statistics of Self-Attention

In this Appendix, we'll determine the statistics of the query–key dot product (1.16),

$$\widetilde{\Omega}^h_{\alpha;tt'} \equiv \frac{1}{\sqrt{C}} \sum_{c=1}^{C} q^h_{\alpha;t;c} k^h_{\alpha;t';c} = \frac{1}{\sqrt{C}} \sum_{c=1}^{C} \sum_{i_1,i_2=1}^{n} Q^h_{ci_1} K^h_{ci_2} s_{\alpha;t;i_1} s_{\alpha;t';i_2} , \tag{A1}$$

which in turn dictates the statistics of the self-attention matrices. In the main text, we've seen that the mean of this product (1.60) vanishes as

$$
\begin{aligned}
\mathbb{E}\left[\widetilde{\Omega}^h_{\alpha;tt'}\right] =& \mathbb{E}\left[\frac{1}{\sqrt{C}} \sum_{c=1}^{C} \sum_{i,j=1}^{n} Q^h_{ci} K^h_{cj} s_{\alpha;t;i} s_{\alpha;t';j}\right] \\
=& \frac{1}{\sqrt{C}} \sum_{c=1}^{C} \sum_{i,j=1}^{n} \mathbb{E}\left[Q^h_{ci}\right] \mathbb{E}\left[K^h_{cj}\right] \mathbb{E}\left[s_{\alpha;t;i} s_{\alpha;t';j}\right] = 0 ,
\end{aligned}
\tag{A2}
$$

while its covariance (1.61) is given by

$$
\begin{aligned}
& \mathbb{E}\left[\widetilde{\Omega}^{h_1}_{\alpha_1;t_1t_1'} \widetilde{\Omega}^{h_2}_{\alpha_2;t_2t_2'}\right] \tag{A3}\\
=& \frac{1}{C} \sum_{c_1,c_2=1}^{C} \sum_{i_1,i_2,j_1,j_2=1}^{n} \mathbb{E}\left[Q^{h_1}_{c_1i_1} K^{h_1}_{c_1j_1} s_{\alpha_1;t_1;i_1} s_{\alpha_1;t_1';j_1} Q^{h_2}_{c_2i_2} K^{h_2}_{c_2j_2} s_{\alpha_2;t_2;i_2} s_{\alpha_2;t_2';j_2}\right] \\
=& \frac{1}{C} \sum_{c_1,c_2=1}^{C} \sum_{i_1,i_2,j_1,j_2=1}^{n} \frac{C_Q}{n} \frac{C_K}{n} \delta_{c_1c_2} \delta_{i_1i_2} \delta_{j_1j_2} \delta^{h_1h_2} \mathbb{E}\left[s_{\alpha_1;t_1;i_1} s_{\alpha_1;t_1';j_1} s_{\alpha_2;t_2;i_2} s_{\alpha_2;t_2';j_2}\right] \\
=& \delta^{h_1h_2} C_Q C_K \mathbb{E}\left[\left(\frac{1}{n} \sum_{i=1}^{n} s_{\alpha_1;t_1;i} s_{\alpha_2;t_2;i}\right)\left(\frac{1}{n} \sum_{j=1}^{n} s_{\alpha_1;t_1';j} s_{\alpha_2;t_2';j}\right)\right] \\
=& \delta^{h_1h_2} C_Q C_K F_{(\alpha_1;t_1)(\alpha_2;t_2)} F_{(\alpha_1;t_1')(\alpha_2;t_2')} + O\left(\frac{1}{n}\right) \\
\equiv& \delta^{h_1h_2} A_{(\alpha_1;t_1t_1')(\alpha_2;t_2t_2')} , \tag{A4}
\end{aligned}
$$

where in the last line we've introduced the kernel $A_{(\alpha_1;t_1t_1')(\alpha_2;t_2t_2')}$ for the query–key dot products.

Assuming that the query and key weights are distributed symmetrically around zero – as is usually the case – we can clearly see that all the odd-point correlators vanish as

$$\mathbb{E}\left[\widetilde{\Omega}^{h_1}_{\alpha_1;t_1t_1'} \cdots \widetilde{\Omega}^{h_{2m-1}}_{\alpha_{2m-1};t_{2m-1}t_{2m-1}'}\right] = 0 . \tag{A5}$$

Thus we'll focus on the even-point correlators. Going ahead and rolling them out just as we did

for the covariance (a step-by-step instruction will follow), we get

$$\mathbb{E}\left[\widetilde{\Omega}^{h_1}_{\alpha_1;t_1t'_1}\cdots\widetilde{\Omega}^{h_{2m}}_{\alpha_{2m};t_{2m}t'_{2m}}\right] \tag{A6}$$

$$=\frac{1}{C^m}\sum_{c_1,\ldots,c_{2m}=1}^{C}\sum_{i_1,\ldots,i_{2m}=1}^{n}\sum_{j_1,\ldots,j_{2m}=1}^{n}\mathbb{E}\left[Q^{h_1}_{c_1i_1}\cdots Q^{h_{2m}}_{c_{2m}i_{2m}}\right]\mathbb{E}\left[K^{h_1}_{c_1j_1}\cdots K^{h_{2m}}_{c_{2m}j_{2m}}\right]$$

$$\times\mathbb{E}\left[\left(s_{\alpha_1;t_1;i_1}\cdots s_{\alpha_{2m};t_{2m};i_{2m}}\right)\left(s_{\alpha_1;t'_1;j_1}\cdots s_{\alpha_{2m};t'_{2m};j_{2m}}\right)\right]$$

$$=\frac{1}{C^m}\sum_{c_1,\ldots,c_{2m}=1}^{C}\sum_{i_1,\ldots,i_{2m}=1}^{n}\sum_{j_1,\ldots,j_{2m}=1}^{n}\left(\frac{C_Q}{n}\frac{C_K}{n}\right)^m$$

$$\times\left[\sum_{\text{all pairing}}\left(\delta_{c_{p_1}c_{p_2}}\cdots\delta_{c_{p_{2m-1}}c_{p_{2m}}}\right)\left(\delta_{i_{p_1}i_{p_2}}\cdots\delta_{i_{p_{2m-1}}i_{p_{2m}}}\right)\left(\delta^{h_{p_1}h_{p_2}}\ldots\delta^{h_{p_{2m-1}}h_{p_{2m}}}\right)\right]$$

$$\times\left[\sum_{\text{all pairing}}\left(\delta_{c_{q_1}c_{q_2}}\cdots\delta_{c_{q_{2m-1}}c_{q_{2m}}}\right)\left(\delta_{j_{q_1}j_{q_2}}\cdots\delta_{j_{q_{2m-1}}j_{q_{2m}}}\right)\left(\delta^{h_{q_1}h_{q_2}}\ldots\delta^{h_{q_{2m-1}}h_{q_{2m}}}\right)\right]$$

$$\times\mathbb{E}\left[\left(s_{\alpha_1;t_1;i_1}\cdots s_{\alpha_{2m};t_{2m};i_{2m}}\right)\left(s_{\alpha_1;t'_1;j_1}\cdots s_{\alpha_{2m};t'_{2m};j_{2m}}\right)\right]+O\left(\frac{1}{n}\right)$$

$$=\frac{1}{C^m}C^m\sum_{i_1,\ldots,i_{2m}=1}^{n}\sum_{j_1,\ldots,j_{2m}=1}^{n}\left(\frac{C_Q}{n}\frac{C_K}{n}\right)^m$$

$$\times\left[\sum_{\text{all pairing}}\left(\delta_{i_{p_1}i_{p_2}}\cdots\delta_{i_{p_{2m-1}}i_{p_{2m}}}\right)\left(\delta^{h_{p_1}h_{p_2}}\ldots\delta^{h_{p_{2m-1}}h_{p_{2m}}}\right)\left(\delta_{j_{p_1}j_{p_2}}\cdots\delta_{j_{p_{2m-1}}j_{p_{2m}}}\right)\right]$$

$$\times\mathbb{E}\left[\left(s_{\alpha_1;t_1;i_1}\cdots s_{\alpha_{2m};t_{2m};i_{2m}}\right)\left(s_{\alpha_1;t'_1;j_1}\cdots s_{\alpha_{2m};t'_{2m};j_{2m}}\right)\right]+O\left(\frac{1}{C}\right)$$

$$=(C_QC_K)^m\sum_{\text{all pairing}}\left(\delta^{h_{p_1}h_{p_2}}\ldots\delta^{h_{p_{2m-1}}h_{p_{2m}}}\right)$$

$$\times\mathbb{E}\left[\left(\frac{1}{n}\sum_{i_1=1}^{n}s_{\alpha_{p_1};t_{p_1};i_1}s_{\alpha_{p_2};t_{p_2};i_1}\right)\cdots\left(\frac{1}{n}\sum_{i_m=1}^{n}s_{\alpha_{p_{2m-1}};t_{p_{2m-1}};i_m}s_{\alpha_{p_{2m}};t_{p_{2m}};i_m}\right)\right.$$

$$\left.\times\left(\frac{1}{n}\sum_{j_1=1}^{n}s_{\alpha_{p_1};t'_{p_1};j_1}s_{\alpha_{p_2};t'_{p_2};j_1}\right)\cdots\left(\frac{1}{n}\sum_{j_m=1}^{n}s_{\alpha_{p_{2m-1}};t'_{p_{2m-1}};j_m}s_{\alpha_{p_{2m}};t'_{p_{2m}};j_m}\right)\right]+O\left(\frac{1}{C}\right)$$

$$=\sum_{\text{all pairing}}\left(\delta^{h_{p_1}h_{p_2}}\ldots\delta^{h_{p_{2m-1}}h_{p_{2m}}}\right)$$

$$\times A_{(\alpha_{p_1};t_{p_1}t'_{p_1})(\alpha_{p_2};t_{p_2}t'_{p_2})}\cdots A_{(\alpha_{p_{2m-1}};t_{p_{2m-1}}t'_{p_{2m-1}})(\alpha_{p_{2m}};t_{p_{2m}}t'_{p_{2m}})}+O\left(\frac{1}{C}\right).$$

Here, in the first equality, we explicitly wrote out the query–key dot products and separated the expectations for the query and key weights using their statistical independence; in the second equality, we used the Wick's theorem and expressed the expectation of query weights as the sum over all the $(2m-1)!!$ pairings of the auxiliary indices $1,\ldots,2m$, and did the same for the key weights (this step is exact if the query and key weights are drawn from normal distributions while

it involves the $1/n$ corrections for non-normal distributions); in the third equality we summed over the per-head-channel indices, noticing that the contributions would be $(1/C = H/n)$-suppressed when the pairings of those indices differ between that for the query weights and that for the key weights (if this explanation is cryptic, then we recommend performing this step explicitly for the case of $2m = 4$); in the fourth step, we simply performed the summation over embedding indices; in the fifth step, we used our eightfold result (1.52) to truncate away the $1/n$ corrections and used the definition of the newly-introduced kernel $A_{(\alpha_1;t_1 t_1')(\alpha_2;t_2 t_2')}$.

All in all, at the leading order in $1/C = H/n$, the query–key dot product $\widetilde{\Omega}^h_{\alpha;tt'}$ obeys Gaussian statistics with zero mean and order-one covariance $\delta^{h_1 h_2} A_{(\alpha_1;t_1 t_1')(\alpha_2;t_2 t_2')}$. As such, any order-one function of the query–key dot product $\widetilde{\Omega}^h_{\alpha;tt'}$, in particular the self-attention matrix $\Omega^h_{\alpha;tt'}$, has the expectation value of order one (unless the said function involves *only* odd powers of the query–key dot products). We also emphasize that all the distinct heads are statistically independent – as manifested by the Kronecker delta $\delta^{h_1 h_2}$ – and, e.g., the expectation value $\mathbb{E}\left[\Omega^h_{\alpha_1;t_1 t_1'} \Omega^h_{\alpha_2;t_2 t_2'}\right]$ that appeared in footnote 14 takes the same order-one value for all the heads h and can be expressed as a $\left(|\mathcal{D}|T^2\right)$-dimensional Gaussian integral with the kernel $A_{(\alpha_1;t_1 t_1')(\alpha_2;t_2 t_2')}$.

Speaking of footnote 14, to show the factorization therein, we just have to note that the fivefold derivation (A6) of the Gaussian statistics above goes through verbatim to yield

$$
\begin{aligned}
&\mathbb{E}\left[\widetilde{\Omega}^{h_1}_{\alpha_1;t_1 t_1'} \cdots \widetilde{\Omega}^{h_{2m}}_{\alpha_{2m};t_{2m} t_{2m}'} \mathcal{F}(s)\right] \\
&= (C_Q C_K)^m \sum_{\text{all pairing}} \left(\delta^{h_{p_1} h_{p_2}} \ldots \delta^{h_{p_{2m-1}} h_{p_{2m}}}\right) \\
&\quad \times \mathbb{E}\Bigg[\left(\frac{1}{n}\sum_{i_1=1}^n s_{\alpha_{p_1};t_{p_1};i_1} s_{\alpha_{p_2};t_{p_2};i_1}\right) \cdots \left(\frac{1}{n}\sum_{i_m=1}^n s_{\alpha_{p_{2m-1}};t_{p_{2m-1}};i_m} s_{\alpha_{p_{2m}};t_{p_{2m}};i_m}\right) \\
&\quad \times \left(\frac{1}{n}\sum_{j_1=1}^n s_{\alpha_{p_1};t_{p_1}';j_1} s_{\alpha_{p_2};t_{p_2}';j_1}\right) \cdots \left(\frac{1}{n}\sum_{j_m=1}^n s_{\alpha_{p_{2m-1}};t_{p_{2m-1}}';j_m} s_{\alpha_{p_{2m}};t_{p_{2m}}';j_m}\right) \mathcal{F}(s)\Bigg] + O\left(\frac{1}{C}\right)
\end{aligned}
\tag{A7}
$$

with an insertion of any function $\mathcal{F}(s)$ of the signals $s_{\alpha;t;i}$, and we can then use the eightfold result (1.52) once again to seal the deal (1.64).

If you've been paying attention, then there is one more deal (1.155) to seal, that is, footnote 30.

That sealing essentially boils down to realizing that

$$
\frac{1}{C^m} \sum_{c_1,\dots,c_{2m}=1}^{C} \mathbb{E}\left[Q_{c_1 i_1}^{h_1} \cdots Q_{c_{2m} i_{2m}}^{h_{2m}} \right] \mathbb{E}\left[K_{c_1 j_1}^{h_1} \cdots K_{c_{2m} j_{2m}}^{h_{2m}} \left(K_{c_{2m+1} j_{2m+1}}^{h_{2m+1}} K_{c_{2m+2} j_{2m+2}}^{h_{2m+2}} \right) \right] \tag{A8}
$$

$$
= \frac{1}{C^m} \sum_{c_1,\dots,c_{2m}=1}^{C} \left(\frac{C_Q}{n} \frac{C_K}{n} \right)^m
$$

$$
\times \left[\sum_{\text{all pairing}} \left(\delta_{c_{p_1} c_{p_2}} \cdots \delta_{c_{p_{2m-1}} c_{p_{2m}}} \right) \left(\delta_{i_{p_1} i_{p_2}} \cdots \delta_{i_{p_{2m-1}} i_{p_{2m}}} \right) \left(\delta^{h_{p_1} h_{p_2}} \dots \delta^{h_{p_{2m-1}} h_{p_{2m}}} \right) \right]
$$

$$
\times \left[\sum_{\text{all pairing}} \left(\delta_{c_{q_1} c_{q_2}} \cdots \delta_{c_{q_{2m+1}} c_{q_{2m+2}}} \right) \left(\delta_{j_{q_1} j_{q_2}} \cdots \delta_{j_{q_{2m+1}} j_{q_{2m+2}}} \right) \left(\delta^{h_{q_1} h_{q_2}} \dots \delta^{h_{q_{2m+1}} h_{q_{2m+2}}} \right) \right]
$$

$$
= \left(\frac{C_Q}{n} \frac{C_K}{n} \right)^m \delta_{c_{2m+1} c_{2m+2}} \delta_{j_{2m+1} j_{2m+2}} \delta^{h_{2m+1} h_{2m+2}}
$$

$$
\times \left[\sum_{\text{all pairing}} \left(\delta_{i_{p_1} i_{p_2}} \cdots \delta_{i_{p_{2m-1}} i_{p_{2m}}} \right) \left(\delta^{h_{p_1} h_{p_2}} \dots \delta^{h_{p_{2m-1}} h_{p_{2m}}} \right) \left(\delta_{j_{p_1} j_{p_2}} \cdots \delta_{j_{p_{2m-1}} j_{p_{2m}}} \right) + O\left(\frac{1}{C} \right) \right].
$$

For a neater derivation that generalizes better, search for interlayer correlations in Ref. [15].

---

[1] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, *et al.*, "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901. 2020.

[2] J. W. Rae, S. Borgeaud, T. Cai, K. Millican, J. Hoffmann, F. Song, J. Aslanides, S. Henderson, R. Ring, S. Young, *et al.*, "Scaling Language Models: Methods, Analysis & Insights from Training Gopher," `arXiv:2112.11446 [cs.CL]`.

[3] S. Smith, M. Patwary, B. Norick, P. LeGresley, S. Rajbhandari, J. Casper, Z. Liu, S. Prabhumoye, G. Zerveas, V. Korthikanti, *et al.*, "Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model," `arXiv:2201.11990 [cs.CL]`.

[4] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, *et al.*, "PaLM: Scaling Language Modeling with Pathways," `arXiv:2204.02311 [cs.CL]`.

[5] S. Zhang, S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, *et al.*, "OPT: Open Pre-trained Transformer Language Models," `arXiv:2205.01068 [cs.CL]`.

[6] R. M. Neal, "Priors for Infinite Networks," in *Bayesian Learning for Neural Networks*, pp. 29–53. Springer, 1996.

[7] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein, "Deep Neural Networks as Gaussian Processes," in *International Conference on Learning Representations*. 2018.

[8] A. G. d. G. Matthews, M. Rowland, J. Hron, R. E. Turner, and Z. Ghahramani, "Gaussian Process Behaviour in Wide Deep Neural Networks," in *International Conference on Learning Representations*. 2018.

[9] A. Jacot, F. Gabriel, and C. Hongler, "Neural Tangent Kernel: Convergence and Generalization in Neural Networks," in *Advances in Neural Information Processing Systems*, vol. 31, pp. 8571–8580. 2018.

[10] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli, "Exponential expressivity in deep neural networks through transient chaos," in *Advances in Neural Information Processing Systems*, vol. 29, pp. 3360–3368. 2016.

[11] M. Raghu, B. Poole, J. Kleinberg, S. Ganguli, and J. Sohl-Dickstein, "On the Expressive Power of Deep Neural Networks," in *International Conference on Machine Learning*, pp. 2847–2854. 2017.

[12] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, "Deep Information Propagation," in *International Conference on Learning Representations*. 2017.

[13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems*, vol. 30, pp. 5998–6008. 2017.

[14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," in *International Conference on Learning Representations*. 2021.

[15] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory*. Cambridge University Press, 2022. `arXiv:2106.10165 [cs.LG]`. `https://deeplearningtheory.com`.

[16] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer Normalization," `arXiv:1607.06450 [stat.ML]`.

[17] D. Doshi, T. He, and A. Gromov, "Critical Initialization of Wide and Deep Neural Networks through Partial Jacobians: General Theory and Applications," `arXiv:2111.12143 [cs.LG]`.

[18] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. 2009.

[19] X. Zhai, A. Kolesnikov, N. Houlsby, and L. Beyer, "Scaling Vision Transformers," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12104–12113. 2022.

[20] T. Xiao, M. Singh, E. Mintun, T. Darrell, P. Dollár, and R. Girshick, "Early Convolutions Help Transformers See Better," in *Advances in Neural Information Processing Systems*, vol. 34, pp. 30392–30400. 2021.

[21] N. Shazeer, "GLU Variants Improve Transformer," `arXiv:2002.05202 [cs.CL]`.

[22] O. Press and L. Wolf, "Using the Output Embedding to Improve Language Models," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 157–163. 2017.

[23] S. Yaida, "Meta-Principled Family of Hyperparameter Scaling Strategies," `arXiv:2210.04909`

`[cs.LG]`.

[24] G. Yang and E. J. Hu, "Tensor Programs IV: Feature Learning in Infinite-Width Neural Networks," in *International Conference on Machine Learning*, pp. 11727–11737. 2021.

[25] G. Yang, E. J. Hu, I. Babuschkin, S. Sidor, X. Liu, D. Farhi, N. Ryder, J. Pachocki, W. Chen, and J. Gao, "Tensor Programs V: Tuning Large Neural Networks via Zero-Shot Hyperparameter Transfer," `arXiv:2203.03466 [cs.LG]`.

[26] E. Dyer and G. Gur-Ari, "Asymptotics of Wide Networks from Feynman Diagrams," in *International Conference on Learning Representations*. 2020.

[27] B. Hanin and M. Nica, "Finite Depth and Width Corrections to the Neural Tangent Kernel," in *International Conference on Learning Representations*. 2019.

[28] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," `arXiv:1412.6980 [cs.LG]`.

[29] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," `arXiv:1711.05101 [cs.LG]`.

[30] G. K. Zipf, *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Ravenio Books, 2016.

[31] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, *et al.*, "Mixed Precision Training," `arXiv:1710.03740 [cs.AI]`.

[32] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A ConvNet for the 2020s," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11976–11986. 2022.

[33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, *et al.*, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision* **115** no. 3, (2015) 211–252.

[34] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "Randaugment: Practical Automated Data Augmentation With a Reduced Search Space," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 702–703. 2020.

[35] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond Empirical Risk Minimization," in *International Conference on Learning Representations*. 2018.

[36] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features," in *IEEE/CVF International Conference on Computer Vision*, pp. 6023–6032. 2019.

[37] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random Erasing Data Augmentation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 13001–13008. 2020.

[38] R. Wightman *et al.*, "PyTorch Image Models," 2019.

[39] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception Architecture for Computer Vision," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826. 2016.

[40] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger, "Deep Networks with Stochastic Depth,"

in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pp. 646–661. 2016.

[41] B. T. Polyak and A. B. Juditsky, "Acceleration of Stochastic Approximation by Averaging," *SIAM Journal on Control and Optimization* **30** no. 4, (1992) 838–855.

[42] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, "Accurate, Large Minibatch SGD: Training ImageNet in 1 hour," `arXiv:1706.02677 [cs.CV]`.

[43] I. Loshchilov and F. Hutter, "SGDR: Stochastic Gradient Descent with Warm Restarts," in *International Conference on Learning Representations*. 2017.

[44] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880. 2020.

[45] K. Shuster, M. Komeili, L. Adolphs, S. Roller, A. Szlam, and J. Weston, "Language Models that Seek for Knowledge: Modular Search & Generation for Dialogue and Prompt Completion," in *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pp. 373–393. 2022.

[46] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A Fast, Extensible Toolkit for Sequence Modeling," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pp. 48–53. 2019.

[47] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," `arXiv:1907.11692 [cs.CL]`.

[48] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler, "Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books," in *IEEE International Conference on Computer Vision*, pp. 19–27. 2015.

[49] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," `arXiv:1810.04805 [cs.CL]`.

[50] S. Nagel, "CC-News," 2016. `http://commoncrawl.org/2016/10/news-dataset-available`.

[51] A. Gokaslan and V. Cohen, "Openwebtext corpus," 2019. `http://Skylion007.github.io/OpenWebTextCorpus`.

[52] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, *et al.*, "Language Models are Unsupervised Multitask Learners," *OpenAI blog* **1** no. 8, (2019) 9.

[53] T. H. Trinh and Q. V. Le, "A Simple Method for Commonsense Reasoning," `arXiv:1806.02847 [cs.AI]`.

[54] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott,

L. Zettlemoyer, and V. Stoyanov, "Unsupervised Cross-lingual Representation Learning at Scale," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 8440–8451. 2020.

[55] J. Baumgartner, S. Zannettou, B. Keegan, M. Squire, and J. Blackburn, "The Pushshift Reddit Dataset," in *Proceedings of the AAAI International Conference on Web and Social Media*, vol. 14, pp. 830–839. 2020.

[56] S. Roller, E. Dinan, N. Goyal, D. Ju, M. Williamson, Y. Liu, J. Xu, M. Ott, K. Shuster, E. M. Smith, *et al.*, "Recipes for Building an Open-Domain Chatbot," in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 300–325. 2021.