
The Shaped Transformer: Attention Models in the Infinite Depth-and-Width Limit

Lorenzo Noci^{*†}

Chuning Li^{*‡}

Mufan (Bill) Li^{*‡}

Bobby He[§]

Thomas Hofmann[†]

Chris Maddison[‡]

Daniel M. Roy[‡]

Abstract

In deep learning theory, the covariance matrix of the representations serves as a proxy to examine the network’s trainability. Motivated by the success of Transformers, we study the covariance matrix of a modified Softmax-based attention model with skip connections in the proportional limit of infinite-depth-and-width. We show that at initialization the limiting distribution can be described by a stochastic differential equation (SDE) indexed by the depth-to-width ratio. To achieve a well-defined stochastic limit, the Transformer’s attention mechanism is modified by centering the Softmax output at identity, and scaling the Softmax logits by a width-dependent temperature parameter. We examine the stability of the network through the corresponding SDE, showing how the scale of both the drift and diffusion can be elegantly controlled with the aid of residual connections. The existence of a stable SDE implies that the covariance structure is well-behaved, even for very large depth and width, thus preventing the notorious issues of rank degeneracy in deep attention models. Finally, we show, through simulations, that the SDE provides a surprisingly good description of the corresponding finite-size model. We coin the name *shaped Transformer* for these architectural modifications.

1 Introduction

Pre-trained large language models have experienced a remarkable increase in popularity due to their eerily human-like ability to puzzle through complex reasoning tasks, solve coding challenges, and produce pages of logically sound text [1]. Arguably, the Transformer is the foundation of these successes [2]. Recent research has found evidence for scaling laws, linking the performance of these architectures to their parameter counts and the quantity of training data, fueling the desire to train deeper and wider models on ever larger datasets in order to unlock new levels of performance [3–7].

Bundled with the increased expressivity of deep architectures, however, is increased numerical instability, both in the forward pass and gradients, which hinders training. One of the clearest examples of instability is the so-called rank collapse phenomenon [8, 9] – the observation that, in Softmax-based attention models, the network’s representation of different tokens tend to perfectly align at large depth. The resulting poorly conditioned covariance and correlation between tokens leads to exploding and/or vanishing gradients at initialization, disrupting gradient updates of the affected parameters. This situation violates a well-known guiding principle from the literature of deep signal propagation: a stable covariance is a necessary condition for stable training [10–15]. In fact,

*Equal contribution. Correspondence to:

lorenzo.noci@inf.ethz.ch, chuning.li@mail.utoronto.ca, mufan.li@mail.utoronto.ca

†ETH Zurich

‡University of Toronto and Vector Institute

§University of Oxford

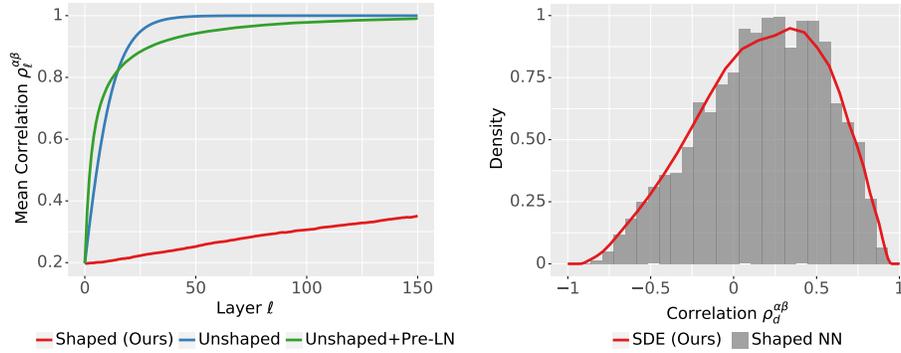


Figure 1: Our shaped Transformer prevents token representations from becoming perfectly aligned, i.e. rank collapse. Left: mean correlation $\rho_\ell^{\alpha\beta}$ of Transformers (Eq. 11) with and without shaped attention (Eq. 9) and Pre-LN [48]. Right: kernel density estimate and histogram of correlations from covariance SDE in Theorem 4.2 and shaped attention NN. Here we note correlation converging to 1 implies a poorly conditioned covariance matrix. Simulated with $n = 200$, $d = 150$, $\gamma = 1/\sqrt{8}$, $\tau_0 = 1$, $\rho_0^{\alpha\beta} = 0.2$, SDE step size 0.01, and 2^{12} samples.

the instability of Transformers is evident when considering the critical role of hyperparameter tuning and the judicious use of normalization layers. In this work, we study Transformers in a novel infinite limit, rectify sources of instability with a novel modification, and derive the SDEs characterizing the covariance and output distribution.

Scaling limits have been used successfully to provide guidance on architecture [16–18] and tuning hyperparameters settings [19]. Our work represents a contribution in this direction. The ability to use such limits to diagnose instabilities depends on their tractability and faithfulness to real-world (finite) networks. In this regard, not all limits are created equal. In particular, the faithfulness of scaling limits depends critically on how other parameters are scaled with width. One of the simplest (and thus most popular) limits to work with – the “NTK” limit [20–24] – treats the depth of the network as fixed. As a result, at initialization, this limit does not accumulate sufficient random fluctuations over the depth of the network, leading to deterministic covariance matrices that do not agree with those of standard (finite) networks. Such networks have another defect: they are incapable of learning features in the limit [25]. Various other limits have been studied, towards identifying tractable yet faithful models of initialization and/or training. These include mean field limits [26–29] and the perturbative regime [30–36].

This work operates in a relatively new regime – the *proportional* infinite depth-and-width limit – where depth d and width n diverge as the ratio d/n tends to a positive constant. This limit, first analyzed by Hanin and Nica [37], has been the recent subject of study in the context of neural network [38–41, 18]. A related line of work also studied the Lyapunov exponent for products of random matrices [42–45]. This regime retains the network’s stochasticity and, at initialization, has been shown to closely resemble the behaviour of finite architectures, yet still yield a relatively simple limiting description, expressible in terms of stochastic differential equations [40, 18]. In this work, we fully characterize the initial output distributions of a network with skip connections and Softmax-based attention mechanisms, in the proportional infinite-depth-and-width limit.

Inspired by the idea of shaping activation functions [16–18, 46], our theoretical approach finds an adequately modified attention mechanism via its SDE limit. Our modification involves making the attention matrix closer to the identity, and appropriately choosing the temperature parameter τ , which re-scales the logits of the Softmax. Similar to shaping activation functions, the temperature scaling we devise linearizes and reduces the saturation of the Softmax, a known source of training instability in Transformers [47]. In order to model the feedforward layer of a Transformer’s block, we extend existing results [18] to derive an SDE for the proportional limit of shaped-ReLU feedforward multi-layer perceptrons (MLPs) with skip connections. Combined, we fully characterize the output distribution of a Transformer with shaped non-linearities (Corollary 4.3).

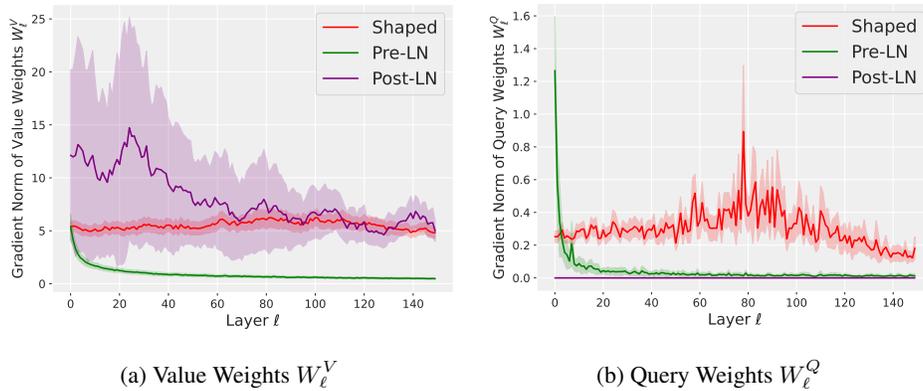


Figure 2: Comparing gradients norms at initialization for different parameters as a function of depth, with and without shaped attention. The architecture is the same as in Figure 1 but with autoregressive causal masking, and the task is next-token prediction on code data. Left: Value weights W_ℓ^V for shaped attention, standard Pre-LN, and the original Post-LN block [2]. Right: the same gradient norm plot but for Query weights W_ℓ^Q . We find that shaping the attention mechanism successfully prevents gradients from vanishing, while unshaped Transformers suffer from rapidly vanishing gradients. Interestingly, only the Post-LN query gradients vanish, but value gradients are stable across depths, which is consistent with the findings of Noci et al. [9]. On the other hand, shaped attention has stable gradients for both parameters inside and outside the Softmax nonlinearity.

Notably, our modification successfully prevents a poorly conditioned covariance matrix, whereas the vanilla Softmax-based attention model without LayerNorm [49] fails in this regard, and the corresponding Pre-LN architecture provides only marginal improvements (see Figure 1). Given that our modification is inspired by previous work on shaping activation functions, we coin the terms *shaped attention* for the proposed attention mechanism and *shaped Transformer* for the overall architecture that includes the MLP block and residual connections. Through simulations (e.g., Figure 1), we show that the limiting neural covariance SDE approximates the distribution of finite-size Transformers with shaped attention mechanism surprisingly well. We also provide preliminary training experiments for our proposed shaped attention architecture on standard language modeling tasks, demonstrating the feasibility of the new architecture in practice (see Section 5 and Appendix D).

In summary, our contributions are as follows:

1. We study the effect of skip connections in the proportional limit, showing that under a precise relation between the scaling parameters of the shortcut and residual branches, the feature covariance converges to the solution of a weighted version of the neural covariance SDE for MLPs (Theorem 3.2). The dependence on the depth-to-width ratio implies the existence of a stable non-commutative limit for residual networks, complementing the commutative limit studied in Hayou and Yang [50].
2. We propose *shaped attention*, where we modify the Softmax-based attention mechanism to be a perturbation of the identity. We demonstrate that shaped attention successfully prevents the degeneracy of correlation in contrast to existing Transformer architectures (Figure 1). The enhanced stability in the forward pass is reflected in the gradients, which are also stable with depth, as we empirically show in Figure 2.
3. For the proposed shaped attention architecture, we derive the neural covariance SDE characterizing the initial distribution in the proportional limit (Theorem 4.2). Consequently, we provide the first characterization of Transformer-type architectures, i.e. the shaped Transformer, in the large depth-and-width regime (Corollary 4.3).
4. We provide simulations to validate the theory and to interpret the effects of network hyperparameters on the covariance matrix of the shaped Transformer. Specifically, we study finite time stability of the SDE and provide explicit guidance on hyperparameters to prevent numerical instability.

The paper is organized as follows: In Section 2, we provide the basic setup and some background on existing results. In Section 3, we generalize the SDE results of M. Li et al. [18] to include skip connections. This serves as a model to understand the effect of skip connections in isolation from the attention model. In Section 4, we present our main result, first pinpointing the origins of instability in the Softmax, then showing how the modifications underlying *shaped attention* allow us to derive a non-trivial SDE limit. Finally, in Section 5, we discuss the implications of our results and some future directions. Proofs of all theorems and additional experiments are deferred to the Appendix.

2 Background

Setup. Let $X_\ell \in \mathbb{R}^{m \times n}$ be the data matrix representing a sequence of m tokens embedded in n dimensions at layer $\ell \in [d]$, where d is the depth of the network. We elide the explicit dependence on ℓ when it is clear from the context, and use superscript Greek letters to indicate specific tokens' representations, for instance $x_\ell^\alpha \in \mathbb{R}^n$ is the α -th row of X_ℓ . We consider the following attention model with residual connections:

$$X_{\ell+1} = \lambda X_\ell + \gamma A_\ell X_\ell \frac{1}{\sqrt{n}} W_\ell^V \quad (1)$$

where $\gamma, \lambda \in [0, 1]$ are parameters that control the strength of the shortcut and residual branch, respectively, $W_\ell^V \in \mathbb{R}^{n \times n}$ is the weight matrix of the values, and $A_\ell \in \mathbb{R}^{m \times m}$ is the attention matrix. We consider Softmax-based scaled dot-product attention, where A_ℓ has the form:

$$A_\ell = \text{Softmax} \left(\frac{1}{\tau} X_\ell \frac{1}{\sqrt{n}} W_\ell^Q \frac{1}{\sqrt{n}} W_\ell^{K, \top} X_\ell^\top \right), \quad (2)$$

where the Softmax is applied row-wise, $W_\ell^Q, W_\ell^K \in \mathbb{R}^{n \times n_k}$ are additional random weights, and τ is a temperature parameter, which controls the entropy of the distribution. Here we let all the weight matrices $W_\ell^Q, W_\ell^K, W_\ell^V$ have $\mathcal{N}(0, 1)$ -iid entries. In the case where $\lambda, \gamma = 1$, with the application of LayerNorm on the residual branch [48], and with $\tau = \sqrt{n_k}$, we recover the attention block of the vanilla "Pre-LN" Transformer architecture [2]. Here we note that we pull the conventional $n^{-1/2}$ factor outside of the weight matrices, which preserves the forward pass, and yields equivalent training dynamics up to a reparameterization of the learning rate [25]. In this work, we consider unnormalized architectures, and control the variance propagation with the condition $\lambda^2 + \gamma^2 = 1$ [40]. We are interested in studying the so-called *neural covariance* for the attention model (Eq. 1) in the proportional limit.

Neural Covariance. In deep learning theory, researchers have long sought to understand how networks internally represent different inputs and how different architectural choices affect these representations. The approach followed by work on signal propagation has been to study how the relative alignment of different inputs evolves across the network, as measured by the neural covariance $V_\ell^{\alpha\beta} := \frac{1}{n} \langle x_\ell^\alpha, x_\ell^\beta \rangle$ (or $\rho^{\alpha\beta} := (V_\ell^{\alpha\alpha} V_\ell^{\beta\beta})^{-1/2} V_\ell^{\alpha\beta}$ if interested only in the correlation). At initialization, characterizations of this covariance structure have been exploited to infer important properties of neural networks [10, 11]. As an example, in the sequential infinite-width-then-depth limit, the correlation $\rho_d^{\alpha\beta}$ of MLPs is known to converge to a fixed point independent of the input [11, 16]. In this regime, the model is not able to discriminate different data points, which severely hinders training, as the gradient step for the deep layers is taken in the same direction regardless of the input. In the context of Softmax-based attention models, Dong et al. [8] proved that the feature matrix X_ℓ loses rank doubly exponentially fast with depth, and Noci et al. [9] showed how this leads to vanishing gradients of the queries and keys parameters, thus further highlighting how the stability of forward and backward passes are deeply entangled (see also Figure 2).

Stabilizing the Effect of Non-Linear Layers. Central to the issue of degeneracy of the neural covariance are commonly used non-linear activation functions that severely deviate from the identity. The recent line of work of Deep Kernel Shaping (DKS) [16–18] addresses the issue by considering the cumulative amount of non-linearity throughout layers, and *shaping* the activation function by making it closer to the identity map. Inspired by this line of work, B. He et al. [46] devise an initialization for Transformers that avoid the rank collapse problem without the aid of skip connections or LayerNorm.

In an alternative approach, the line of work behind Stable ResNets [50–53] considers scaling the residual branches by $\gamma = 1/\sqrt{\text{depth}}$, and postulates this scaling is sufficient to stabilize the neural

covariance with minimal assumptions on the activation function. Noci et al. [9] adopts this scaling to give precise formulas on the expected covariance of a Transformer at initialization. In this work, we consider γ constant in width and depth, and derive a complementary limiting result.

The Proportional Infinite-Depth-and-Width Limit. In the context of feed-forward MLPs, the output distribution with respect to a single input was studied in [37, 41], where it was shown that for the ReLU nonlinearity, the norm of the activations $V^{\alpha\alpha}$ converges to a log-normal random variable. To resolve the degeneracy of covariance and provide a characterization of output distributions for *multiple inputs*, M. Li et al. [18] shapes the ReLU by setting its slope $1/\sqrt{\text{width}}$ -away from linearity. In the proportional limit, the effect of the non-linearity accumulates over the d layers, and the covariance matrix $V_\ell = [V_\ell^{\alpha\beta}]_{\alpha\beta}$ converges weakly to the solution of the SDE

$$dV_t = b_{\text{ReLU}}(V_t) dt + \Sigma_{\text{lin}}^{1/2}(V_t) dB_t, \quad (3)$$

where the formulae for coefficients $b_{\text{ReLU}}, \Sigma_{\text{lin}}$ can be found in Theorem 3.2.

We note that the output neuron distributions are directly recovered as a conditional Gaussian with covariance V_T for $T = \frac{d}{n}$, in a similar spirit as the neural network Gaussian process (NNGP) results [20–22]. For example, the i -th output $X_{\text{out},i}$ conditioned on V_d are asymptotically iid. $\mathcal{N}(0, V_T)$ as $d, n \rightarrow \infty$. The reader is referred to Appendix A for more technical background on the covariance SDE and the convergence result.

While the existing results are limited to initialization, we remind the reader that this is a necessary step before we can study training dynamics. In particular, the NNGP techniques developed for infinite-width networks at initialization were directly used to study the training dynamics in the same limit [24, 54]. We will provide further discussions on this topic in Section 5.

3 Warm-Up: a Neural Covariance SDE for ResNets

To understand the effect of skip connections, it is helpful to look at a simplified model composed of a shaped ReLU-activated layer and skip connections:

$$X_{\ell+1} = \lambda X_\ell + \gamma \sigma_s \left(X_\ell \frac{1}{\sqrt{n}} W_\ell^{\text{pre}} \right) \sqrt{\frac{c}{n}} W_\ell^{\text{post}}, \quad (4)$$

where $\sigma_s(x) := s_+ \max(x, 0) + s_- \min(x, 0)$ is the shaped ReLU with slopes $s_\pm := 1 + c_\pm n^{-1/2}$ for some constants $c_+, c_- \in \mathbb{R}$. We assume i.i.d weights $(W_\ell^{\text{pre}})_{ij}, (W_\ell^{\text{post}})_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, and $c^{-1} = \mathbb{E} \sigma_s(g)^2$ for $g \sim \mathcal{N}(0, 1)$ is a constant that ensures that the activations are normalized [55]. Notice that this is the form of the feedforward layer in a Transformer [2].

We will next define the notion of convergence for our covariance matrices and state our first main result. We refer the reader to Appendix A for more precise details on the Skorohod topology.

Definition 3.1 (Convergence of Covariance). *Let $X_\ell \in \mathbb{R}^{m \times n}$ be the ℓ -th layer matrix of representations, and define the feature covariance as $V_\ell = \frac{1}{n} X_\ell X_\ell^\top$. Let $V_t^{(n)} = V_{\lfloor tn \rfloor} \in \mathbb{R}^{m(m+1)/2}$ be the the continuous time interpolation of the upper triangular entries as a vector. We say the covariance $V^{(n)}$ converges to V , if in the limit as $n, d \rightarrow \infty, \frac{d}{n} \rightarrow T$, the process $\{V_t^{(n)}\}_{t \in [0, T]}$ converges to $\{V_t\}_{t \in [0, T]}$ weakly in the Skorohod topology.*

Theorem 3.2. *Let X_ℓ be the hidden layers of a ResNet defined in Eq. 4 with $\lambda^2 + \gamma^2 = 1$, where both λ and γ do not depend on d, n . Then the feature covariance V_ℓ converges to the solution of the following SDE (in the sense of Definition 3.1)*

$$dV_t = b_{\text{res}}(V_t) dt + \Sigma_{\text{res}}(V_t)^{1/2} dB_t, \quad V_0 = \frac{1}{n} X_0 X_0^\top, \quad (5)$$

where $b_{\text{res}}(V) = \gamma^2 b_{\text{ReLU}}(V) = \gamma^2 [\nu(\rho^{\alpha\beta}) \sqrt{V^{\alpha\alpha} V^{\beta\beta}}]_{\alpha \leq \beta}$ with $\rho^{\alpha\beta} = V^{\alpha\beta} (V^{\alpha\alpha} V^{\beta\beta})^{-1/2}$ and

$$\nu(\rho) = \frac{(c_+ - c_-)^2}{2\pi} \left(\sqrt{1 - \rho^2} - \rho \arccos \rho \right), \quad (6)$$

furthermore, $\Sigma_{\text{res}}(V) = 2\gamma^2 \Sigma_{\text{lin}}(V) = 2\gamma^2 [V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}]_{\alpha \leq \beta, \delta \leq \omega}$.

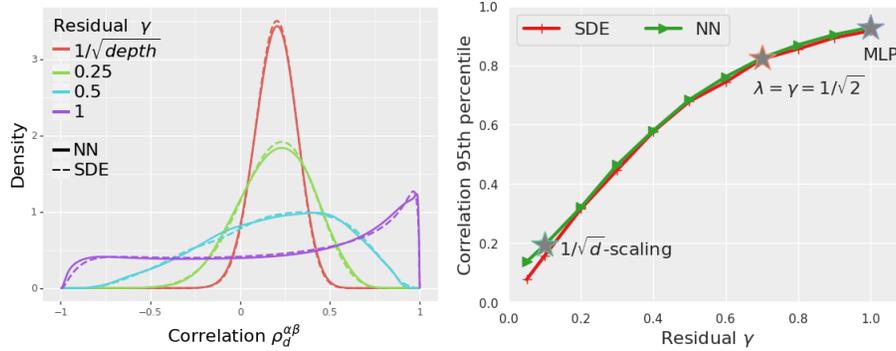


Figure 3: Left: Kernel density estimates of correlation $\rho_d^{\alpha\beta}$ for various values of the residual strength parameter γ . In particular, $\gamma = 1$ recovers a shaped-ReLU MLP without skip connections, and $\gamma = 1/\sqrt{d}$ is the setting studied in Noci et al. [9] and Hayou and Yang [50]. Solid lines represent finite networks, while our SDE simulations are presented in dashed lines. Right: 95th percentile of the absolute value of the correlation distribution as a function of γ . Note reducing γ reduces the concentration around $\rho^{\alpha\beta} = 1$, and our SDE reliably approximates finite size networks. Simulated with $n = 300$, $d = 100$, $\rho_0^{\alpha\beta} = 0.2$, $c_+ = 0$, $c_- = -1$, and 2^{13} samples.

Notice how the limiting SDE closely resembles the MLP case (Eq. 3), which is recovered exactly when $\gamma = 1$. The only difference is the extra 2 factor, which comes from the fact that in our definition each layer has effectively two times the number of weight matrices than the standard formulation for MLPs. As the drift depends solely on the nonlinearity, and the diffusion depends solely on the random weights, only the diffusion variance is doubled. The residual branch parameter $\gamma < 1$ dampens both the drift and the variance of the Brownian motion by γ^2 , thus it can be interpreted as a time change. In other words, the effect of γ at initialization is equivalent to reducing depth-to-width ratio, in line with existing intuitions that ResNets have a lower “effective-depth” [56]. To visualize the stabilizing effect of γ on the distribution, in Figure 3 (right) we plot the 95th percentile correlation as a function of γ . The increasing trend indicates a larger probability of perfect alignment between two tokens. In Figure 3 (left) we plot the densities of both the residual SDE and the corresponding residual network for various values of γ . Notice how the samples from the SDE well-approximates the histogram of a finite network.

A Stable Non-Commutative Limit. Our results complement those of Hayou and Yang [50], where the authors have shown that for a similar ResNet under the parameter scaling $\lambda = 1, \gamma = 1/\sqrt{d}$, the depth and width limits *commute*. More precisely, the covariance $V^{\alpha\beta}$ converges to the same limit regardless of the order with respect to which the limit is taken or the depth-to-width ratio. Furthermore, the limit is *deterministic*, and can be described by an ordinary differential equation (ODE). Intuitively, the convergence to a deterministic quantity occurs because $\gamma = 1/\sqrt{d}$ suppresses the random fluctuations enough to vanish in the limit. On the other hand, our results show that for λ, γ constant in n, d , the random fluctuations are on the right order of $O(n^{-1/2})$ as in the MLP case (Eq. 3), hence they do not vanish in the simultaneous limit. The most notable difference is that our limiting regime is *non-commutative* as it depends on the depth-to-width ratio of the network. We remark that both regimes prevents degeneracy of covariance for residual architectures, forming two theories that complement each other.

4 Neural Covariance SDE for Softmax-Based Attention

4.1 Unshaped Attention and Its Taylor Expansion

A central piece to the neural covariance SDE theory for MLPs [18] is identifying the exact scaling of shaped activation functions. In particular, the effect of the activations on the covariance Markov chain V_ℓ must be on the same order as the random weights in an MLP, thus forming an approximate

Euler-discretization

$$V_{\ell+1} = V_{\ell} + \frac{b(V_{\ell})}{n} + \frac{\Sigma(V_{\ell})^{1/2}\xi_{\ell}}{\sqrt{n}} + O(n^{-3/2}), \quad (7)$$

where b, Σ are deterministic coefficients, and ξ_{ℓ} are random vectors with zero mean and identity covariance. From here onwards, we use $O(n^{-p})$ to denote a random variable Z such that $n^p Z$ has all moments bounded by universal constants (i.e. independent of n). Since the update can be interpreted as discretization with step size n^{-1} , naturally the Markov chain converges to an SDE. We again note that a stable SDE implies a stable covariance structure for finite size networks.

To achieve the same goal for modified attention mechanisms, we consider a similar approach as M. Li et al. [18] for smooth activation functions, and Taylor expand the Softmax function in terms of a large temperature parameter τ . To this end, let Y_{ℓ} to be the matrix of dot-products between queries, and keys, i.e. $Y_{\ell} := X_{\ell} \frac{1}{\sqrt{n}} W_{\ell}^Q \frac{1}{\sqrt{n}} W_{\ell}^{K, \top} X_{\ell}^{\top}$.

More specifically, given a row $y^{\alpha} \in \mathbb{R}^{1 \times m}$ of the logits $Y_{\ell} \in \mathbb{R}^{m \times m}$, we can Taylor expand the row-wise Softmax in terms of τ^{-1} :

$$\text{Softmax}(\tau^{-1}y^{\alpha}) = \frac{1}{m} \mathbf{1}^{\top} + \frac{1}{\tau m} (y^{\alpha} - \bar{y}^{\alpha}) + \frac{1}{2\tau^2 m} \left[(y^{\alpha} - \bar{y}^{\alpha})^2 - (\bar{y}^{\alpha^2} - (\bar{y}^{\alpha})^2) \right] + O(\tau^{-3}), \quad (8)$$

where $\bar{y}^{\alpha} := \frac{1}{m} \sum_{\beta} y^{\alpha\beta} \mathbf{1}^{\top}$ and $(y^{\alpha})^2$ is the vector with squared entries of y^{α} , and $\mathbf{1} \in \mathbb{R}^{m \times 1}$ is the (column) vector of ones. We note in practice τ is often set to $\sqrt{nm_k}$, which is often quite large and allows for asymptotic analysis [9].

We observe that the zero-th order term $m^{-1} \mathbf{1}^{\top}$ is independent of τ . Considering the form of the attention block as $A_{\ell} X_{\ell} \frac{1}{\sqrt{n}} W_{\ell}^V$, this yields an update that is no longer a small perturbation of V_{ℓ} , regardless of how τ is chosen. Therefore, to form a Markov chain like Eq. 7, we actually require A_{ℓ} to be approximately the identity matrix.

4.2 Shaped Attention

To shape the Softmax-attention mechanism as a perturbation of the identity matrix, we propose the following modifications which we call the *shaped attention*⁵

$$A_{\ell} = I + \text{Softmax}(\tau^{-1}Y_{\ell}) - \frac{1}{m} \mathbf{1} \mathbf{1}^{\top}, \quad \tau = \tau_0 \sqrt{nm_k}. \quad (9)$$

The shaped attention presents three modifications to the Softmax attention in Eq. 2. Firstly, the zero-order term $m^{-1} \mathbf{1} \mathbf{1}^{\top}$ of the Taylor expansion (Eq. 8) is removed as it causes a non-infinitesimal drift in the Markov Chain that ultimately leads to instability in the covariance (see Section 4.1). Secondly, we also observe that when τ is very large, the centered Softmax is a perturbation around zero. To recover an approximate Euler-update like in Eq. 7, we simply add back the identity matrix. By biasing the attention matrix towards the identity, we encourage each token to self-attend. This type of modification was also previously considered by B. He et al. [46]. Finally, the Softmax's temperature is chosen to scale as $\tau = \tau_0 \sqrt{nm_k}$, for some constant $\tau_0 > 0$, which guarantees a non-degenerate limit as $(d, n) \rightarrow \infty$ (Theorem 4.2). Note that the extra \sqrt{n} term is a departure from the standard parameterization.

In Figure 4, we show how removing any of the proposed changes individually alters the neural covariance structure, which becomes degenerate for large depths, while the proposed modifications remain stable. We stress that here for simplicity we focus on attention without masking. Shaped attention can be extended to include masking (e.g. casual masking) by centering each i -th row of the Softmax matrix by a different factor $1/m_i$, where m_i is the number of un-masked tokens in the i -th row.

⁵In principle, it could be possible to have a close-to-identity Softmax matrix when the logits are large. However, this regime also corresponds to a very saturated Softmax, thus making training unstable [57]. As a result, we will avoid this direction in this work.

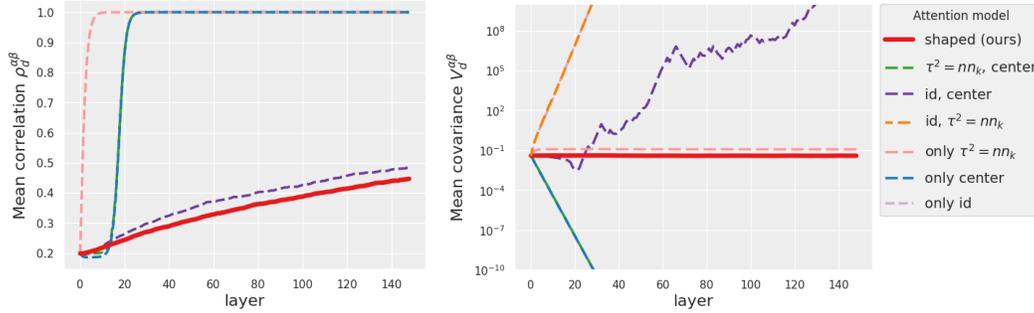


Figure 4: Mean correlation (left) and covariance (right) (in absolute value) under various interventions on the proposed shaped attention. In particular, we remove either one or two of the three modifications from the shaped attention in Eq. 9. For instance " $\tau^2 = nn_k$, center" indicates that we use the proposed temperature, and we center by m^{-1} , but we do not add the identity matrix, while in "only id" we add the identity matrix but use $\tau = \sqrt{nn_k}$ and do not center. We note in this "only id" case, the covariance remains unstable due to incorrect scaling. Due to exploding covariance, we choose to not include the cases "id, $\tau^2 = nn_k$ " and "only id" in the correlation plot (but only in the covariance plot). Simulated with $n = 300$, $d = 150$, $\rho_0^{\alpha\beta} = 0.2$, $\gamma = 1/\sqrt{2}$ and 2^{13} samples.

4.3 Main Result – Neural Covariance SDEs for Shaped Attention Models and Shaped Transformers

Before we state our main results, we will first define a weakened notion of convergence, which is required whenever the drift and covariance coefficients are not Lipschitz. This was also required for the case of shaped MLPs with smooth activations [18].

Definition 4.1 (Local Convergence). *We say the covariance $V^{(n)}$ converges locally to V if the stopped process $\{V_{t \wedge T_r}^{(n)}\}_{t \geq 0}$ converges to $\{V_{t \wedge T_r}\}_{t \geq 0}$ in the sense of Definition 3.1 for all stopping times of the form $T_r = \inf\{t > 0 : \|V_t\| \geq r\}$ with $r > 0$.*

Let the covariance with respect to the average token be defined as $V^{\alpha\bar{x}} := m^{-1} \sum_{\nu=1}^m V^{\alpha\nu}$, and the average trace be $\bar{V} := m^{-1} \sum_{\nu=1}^m V^{\nu\nu}$. We will need to compute a couple of important moments from the Taylor expansion terms of the Softmax (Lemma C.2)

$$\begin{aligned} S_1^{\alpha\delta, \beta\omega} &:= n_k^{-1} \mathbb{E}(Y^{\alpha\delta} - \bar{y}^\alpha)(Y^{\beta\omega} - \bar{y}^\beta) = V^{\alpha\beta} (V^{\delta\omega} - V^{\delta\bar{x}} - V^{\omega\bar{x}} + V^{\bar{x}\bar{x}}), \\ S_2^{\alpha\delta} &:= n_k^{-1} \mathbb{E} \left[(Y^{\alpha\delta} - \bar{y}^\alpha)^2 - ((Y^\alpha)^2 - \bar{y}^{\alpha^2}) \right] = V^{\alpha\alpha} (V^{\delta\delta} - 2V^{\delta\bar{x}} + 2V^{\bar{x}\bar{x}} - \bar{V}). \end{aligned} \quad (10)$$

We are now ready to state our main result.

Theorem 4.2. *Let X_ℓ be the hidden layers of a residual attention network defined in Eq. 1 with shaped attention in Eq. 9, parameters $\lambda^2 + \gamma^2 = 1$ and $\tau = \tau_0 \sqrt{nn_k}$, where λ, γ, τ_0 all do not depend on d, n . Then the feature covariance V_ℓ converges locally to the solution of the following SDE (in the sense of Definition 4.1)*

$$dV_t = b(V_t)dt + \Sigma(V_t)^{1/2}dB_t, \quad V_0 = \frac{1}{n}X_0X_0^\top,$$

where the drift has the following form

$$b(V) = \frac{\gamma^2}{\tau_0^2} \left[\frac{1}{m^2} \sum_{\nu, \kappa=1}^m V^{\nu\kappa} S_1^{\alpha\nu, \beta\kappa} + \frac{1}{2m} \sum_{\nu=1}^m (V^{\beta\nu} S_2^{\alpha\nu} + V^{\alpha\nu} S_2^{\beta\nu}) \right]_{\alpha \leq \beta},$$

the diffusion coefficient is defined by $\Sigma(V) = \gamma^2(2 - \gamma^2)\Sigma_{\text{lin}}(V) + \gamma^4\tau_0^{-2}[\mathcal{A}^{\alpha\beta, \delta\omega}]_{\alpha \leq \beta, \delta \leq \omega}$, and

$$\mathcal{A}^{\alpha\beta, \delta\omega} := \frac{1}{m^2} \sum_{\nu, \kappa=1}^m \left(V^{\alpha\kappa} V^{\delta\nu} S_1^{\beta\kappa, \omega\nu} + V^{\alpha\kappa} V^{\omega\nu} S_1^{\beta\kappa, \delta\nu} + V^{\beta\nu} V^{\delta\kappa} S_1^{\alpha\nu, \omega\kappa} + V^{\beta\nu} V^{\omega\kappa} S_1^{\alpha\nu, \delta\kappa} \right).$$

The drift depends on the shaped attention mechanism through $S_1^{\alpha\delta, \beta\omega}$ and $S_2^{\alpha\delta}$, the moments of the first and second order terms of the Softmax's Taylor expansion. On the other hand, the diffusion term depends on the attention solely through S_1 , present in the additional term $\mathcal{A}^{\alpha\beta, \delta\omega}$. The presence of $\mathcal{A}^{\alpha\beta, \delta\omega}$ is an intriguing difference compared to shaped ReLU networks, where the diffusion is not affected by the activation function. Both components of the SDE depend on averages over the tokens, reflecting the mixing property of the self-attention mechanism, in which every pair of tokens is compared through dot products to form the attention weights. Finally, notice how the residual branch parameter γ^2 has a dampening effect on the scale of both the drift and the diffusion in a similar way as in fully-connected residual network.

We are now ready to introduce the full shaped Transformer architecture, where we combine the attention and residual layers:

$$Z_\ell = \lambda X_\ell + \gamma A_\ell X_\ell \frac{1}{\sqrt{n}} W_\ell^V, \quad X_{\ell+1} = \lambda Z_\ell + \gamma \sigma_s \left(Z_\ell \frac{1}{\sqrt{n}} W_\ell^{\text{pre}} \right) \sqrt{\frac{c}{n}} W_\ell^{\text{post}}, \quad (11)$$

where A_ℓ is the shaped attention defined by Eq. 9. We note that covariance SDE handle stacking of different layer types very conveniently by simply adding the drift and covariance of the diffusion coefficients, which we summarize in the Corollary below.

Corollary 4.3 (Shaped Transformer Covariance SDE). *Let X_ℓ be the hidden layers of a shaped transformer defined in Eq. 11 with parameters $\lambda^2 + \gamma^2 = 1$ and $\tau = \tau_0 \sqrt{nn_k}$, where λ, γ, τ_0 all do not depend on d, n . Then the feature covariance V_ℓ converges locally to the solution of the following SDE (in the sense of Definition 4.1)*

$$dV_t = [b(V_t) + b_{res}(V_t)] dt + [\Sigma(V_t) + \Sigma_{res}(V_t)]^{1/2} dB_t, \quad (12)$$

where the coefficients are defined in Theorem 3.2 and Theorem 4.2.

4.4 On Finite Time Stability of the SDE and Shaped Attention Networks

Although we did not observe numerical instability in majority of our simulations of the shaped attention networks and the corresponding SDE, we did observe that the drift component $b(V_t)$ in Theorem 4.2 is cubic in the entries of V_t . Whenever the drift is not Lipschitz as in this case, we do not have general guarantees for the existence of a solution for all time (see the Feller test for explosions [58, Theorem 5.5.29]). In fact, MLPs with smooth activations also yield non-Lipschitz drift coefficients as seen in M. Li et al. [18].

However, locally Lipschitz coefficients are sufficient to guarantee the existence of local solutions, in the sense of up to a stopping time [59, Proposition 6.9]. Not only does this fact help us establish a precise notion of convergence (Definition 4.1), we can also study the practical implications of this for finite sized attention networks. More specifically, we can inspect the effect of architectural changes to a stopping time.

To demonstrate the potential numerical instabilities, we had to choose an *adversarial* set of parameters: in particular, an unrealistically large norm (approx. $10\sqrt{n}$) for the initial tokens X_0 , which enlarges the eigenvalues of V_0 to the order of 100. Given these initial conditions and a large residual connection weight γ , we were able to consistently generate numerically unstable behaviour in shaped attention networks (see Figure 5 (left)).

That being said, it is very straight forward to stabilize the network by tweaking parameters such as γ, τ_0 and the depth-to-width ratio of the network. We demonstrate the effect of tuning γ on both sample trajectories of the maximum eigenvalue of V_ℓ and the stopping time in Figure 5. As we may intuitively expect, tuning γ smaller will delay the time scale of numerical instabilities, hence allowing for larger depth networks to remain stable.

5 Discussion

Architecture Design and Hyperparameter Tuning. Previous work have demonstrated the practical impact scaling limits can have on designing activation functions [16, 17] and tuning hyperparameters [19]. We follow this line of motivations and proposed a novel attention mechanism, which successfully stabilizes the covariance structure in arbitrarily deep Transformers (e.g. Figure 1). The natural next

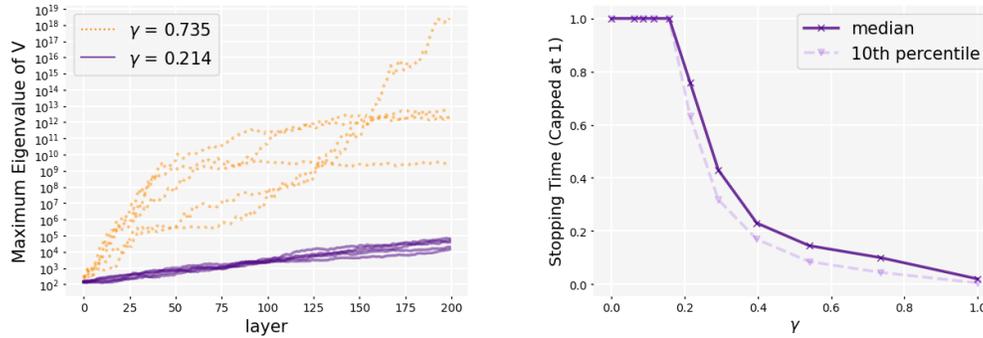


Figure 5: Left: Trajectories of the maximum eigenvalue of the covariance matrix in a shaped attention network, with *adversarially* large initial condition. Right: Stopping time of the shaped attention neural network, capped at 1. Stopping time is defined as $t^* = d^*/n$ with d^* the maximum depth beyond which one of the eigenvalues of the covariance matrix exceeds 10^4 or drops below 10^{-4} . Simulated with $n = d = 200$, $\tau_0 = 1$, and 100 samples used for median and 10th percentile.

step is to investigate the scaling of gradients in the infinite-depth-and-width limit. As Yang et al. [19] illustrated, the existence of an infinite-width limit for the gradient implies the optimal hyperparameters for the training algorithm will also converge. This type of results allows for tuning of hyperparameters on networks with a much smaller width, yet extends easily to arbitrarily large networks that approximates the same limit, saving massive amounts of computing cost in the process. Given the existence of an infinite-depth-and-width limit for the forward pass, we believe it's possible to extract optimal hyperparameters from networks with not only a much smaller width, but *smaller depth* as well.

Preliminary Experiments. Although this work is primarily theoretical, it is important to consider whether or not the proposed architecture is useful in practice. Given limited computing resources, we chose to only briefly test the feasibility of training the shaped Transformer. Nevertheless, our preliminary experiments show promising results when it comes to training stability. In particular, the shaped Transformer (without LayerNorm) does indeed train at approximately the same speed as well tuned Transformer architectures. Full details of the experiment and results can be found in Appendix D. A more comprehensive set of experiments with different tasks, datasets, and larger networks will be required to confidently determine the practical feasibility of the shaped Transformer, which we defer to future work.

Training Dynamics and Generalization. As mentioned in the introduction, the limitations of infinite-width NTK theories motivates our study of the proportional infinite-depth-and-width limit. In particular, to address many of the open problems in deep learning theory, we need a faithful and tractable description of training dynamics. Given the results at initialization, the proportional limit holds the potential for such a theory of training as well. Another promising indicator is that deep networks learn features in the proportional regime [38], which has been identified as a key advantage of neural networks over kernel methods [25, 60–66]. A precise theory of training will help us understand other types of instabilities during training and improve existing optimization methods. Furthermore, determining the network which training converges to is a necessary step towards a theory of generalization, as demonstrated by the infinite-width approach [67]. In light of our results, we believe that our theory sets the stage for future work on training and generalization in deep learning.

Acknowledgement

CL and ML would like to thank Keiran Paster for insightful discussions. LN would like to thank Sotiris Anagnostidis for support in pre-processing the dataset used for the training experiments of this manuscript. ML is supported by the Ontario Graduate Scholarship and Vector Institute. DMR is supported in part by Canada CIFAR AI Chair funding through the Vector Institute, an NSERC Discovery Grant, Ontario Early Researcher Award, a stipend provided by the Charles Simonyi Endowment, and a New Frontiers in Research Exploration Grant.

References

- [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* 33 (2020), pp. 1877–1901.
- [2] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. *Attention is all you need*. 2017. arXiv: [1706.03762](https://arxiv.org/abs/1706.03762).
- [3] J. Hestness, S. Narang, N. Ardalani, G. Diamos, H. Jun, H. Kianinejad, M. Patwary, M. Ali, Y. Yang, and Y. Zhou. *Deep learning scaling is predictable, empirically*. 2017. arXiv: [1712.00409](https://arxiv.org/abs/1712.00409).
- [4] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. *Scaling laws for neural language models*. 2020. arXiv: [2001.08361](https://arxiv.org/abs/2001.08361).
- [5] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark, et al. *Training compute-optimal large language models*. 2022. arXiv: [2203.15556](https://arxiv.org/abs/2203.15556).
- [6] Y. Tay, M. Dehghani, S. Abnar, H. W. Chung, W. Fedus, J. Rao, S. Narang, V. Q. Tran, D. Yogatama, and D. Metzler. *Scaling Laws vs Model Architectures: How does Inductive Bias Influence Scaling?* 2022. arXiv: [2207.10551](https://arxiv.org/abs/2207.10551).
- [7] A. Aghajanyan, L. Yu, A. Conneau, W.-N. Hsu, K. Hambarzumyan, S. Zhang, S. Roller, N. Goyal, O. Levy, and L. Zettlemoyer. *Scaling Laws for Generative Mixed-Modal Language Models*. 2023. arXiv: [2301.03728](https://arxiv.org/abs/2301.03728).
- [8] Y. Dong, J.-B. Cordonnier, and A. Loukas. “Attention is not all you need: Pure attention loses rank doubly exponentially with depth”. *International Conference on Machine Learning*. PMLR. 2021, pp. 2793–2803.
- [9] L. Noci, S. Anagnostidis, L. Biggio, A. Orvieto, S. P. Singh, and A. Lucchi. *Signal Propagation in Transformers: Theoretical Perspectives and the Role of Rank Collapse*. 2022. arXiv: [2206.03126](https://arxiv.org/abs/2206.03126).
- [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* 29 (2016).
- [11] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein. “Deep Information Propagation”. *ICLR*. 2017.
- [12] G. Yang and S. S. Schoenholz. “Mean field residual networks: on the edge of chaos”. *Advances in Neural Information Processing Systems*. 2017, pp. 2865–2873.
- [13] S. Hayou, A. Doucet, and J. Rousseau. “On the impact of the activation function on deep neural networks training”. *International Conference on Machine Learning*. PMLR. 2019, pp. 2672–2680.
- [14] M. Murray, V. Abrol, and J. Tanner. “Activation function design for deep networks: linearity and effective initialisation”. In: *Applied and Computational Harmonic Analysis* 59 (2022), pp. 117–154.
- [15] L. Xiao, J. Pennington, and S. Schoenholz. “Disentangling trainability and generalization in deep neural networks”. *International Conference on Machine Learning*. PMLR. 2020, pp. 10462–10472.
- [16] J. Martens, A. Ballard, G. Desjardins, G. Swirszcz, V. Dalibard, J. Sohl-Dickstein, and S. S. Schoenholz. *Rapid training of deep neural networks without skip connections or normalization layers using deep kernel shaping*. 2021. arXiv: [2110.01765](https://arxiv.org/abs/2110.01765).
- [17] G. Zhang, A. Botev, and J. Martens. *Deep Learning without Shortcuts: Shaping the Kernel with Tailored Rectifiers*. 2022. arXiv: [2203.08120](https://arxiv.org/abs/2203.08120).
- [18] M. Li, M. Nica, and D. Roy. “The neural covariance SDE: Shaped infinite depth-and-width networks at initialization”. In: *Advances in Neural Information Processing Systems* 35 (2022), pp. 10795–10808.
- [19] 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*. 2022. arXiv: [2203.03466](https://arxiv.org/abs/2203.03466).
- [20] R. M. Neal. “Bayesian learning for neural networks”. PhD thesis. University of Toronto, 1995.

- [21] J. Lee, Y. Bahri, R. Novak, S. S. Schoenholz, J. Pennington, and J. Sohl-Dickstein. “Deep Neural Networks as Gaussian Processes”. *ICLR*. 2018.
- [22] A. G. d. G. Matthews, M. Rowland, J. Hron, R. E. Turner, and Z. Ghahramani. *Gaussian process behaviour in wide deep neural networks*. 2018. arXiv: [1804.11271](https://arxiv.org/abs/1804.11271).
- [23] J. Hron, Y. Bahri, J. Sohl-Dickstein, and R. Novak. “Infinite attention: NNGP and NTK for deep attention networks”. *International Conference on Machine Learning*. PMLR. 2020, pp. 4376–4386.
- [24] A. Jacot, F. Gabriel, and C. Hongler. “Neural tangent kernel: Convergence and generalization in neural networks”. *Advances in Information Processing Systems (NeurIPS)*. 2018. arXiv: [1806.07572](https://arxiv.org/abs/1806.07572).
- [25] G. Yang and E. J. Hu. *Feature learning in infinite-width neural networks*. 2020. arXiv: [2011.14522](https://arxiv.org/abs/2011.14522).
- [26] J. Sirignano and K. Spiliopoulos. “Mean field analysis of neural networks: A law of large numbers”. In: *SIAM Journal on Applied Mathematics* 80.2 (2020), pp. 725–752.
- [27] S. Mei, A. Montanari, and P.-M. Nguyen. “A mean field view of the landscape of two-layer neural networks”. In: *Proceedings of the National Academy of Sciences* 115.33 (2018), E7665–E7671.
- [28] L. Chizat and F. Bach. “On the global convergence of gradient descent for over-parameterized models using optimal transport”. In: *Advances in Neural Information Processing Systems* 31 (2018).
- [29] G. M. Rotskoff and E. Vanden-Eijnden. *Trainability and Accuracy of Neural Networks: An Interacting Particle System Approach*. 2018. arXiv: [1805.00915](https://arxiv.org/abs/1805.00915).
- [30] S. Yaida. “Non-Gaussian processes and neural networks at finite widths”. *Mathematical and Scientific Machine Learning*. PMLR. 2020, pp. 165–192.
- [31] E. Dyer and G. Gur-Ari. “Asymptotics of wide networks from feynman diagrams”. In: *arXiv preprint arXiv:1909.11304* (2019).
- [32] D. A. Roberts, S. Yaida, and B. Hanin. *The principles of deep learning theory*. Cambridge University Press, 2022.
- [33] J. Zavatone-Veth and C. Pehlevan. “Exact marginal prior distributions of finite Bayesian neural networks”. In: *Advances in Neural Information Processing Systems* 34 (2021).
- [34] J. Zavatone-Veth, A. Canatar, B. Ruben, and C. Pehlevan. “Asymptotics of representation learning in finite Bayesian neural networks”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 24765–24777.
- [35] B. Hanin. *Correlation Functions in Random Fully Connected Neural Networks at Finite Width*. 2022. arXiv: [2204.01058](https://arxiv.org/abs/2204.01058).
- [36] E. Dinan, S. Yaida, and S. Zhang. *Effective Theory of Transformers at Initialization*. 2023. arXiv: [2304.02034](https://arxiv.org/abs/2304.02034).
- [37] B. Hanin and M. Nica. “Products of many large random matrices and gradients in deep neural networks”. In: *Communications in Mathematical Physics* (2019), pp. 1–36.
- [38] B. Hanin and M. Nica. *Finite depth and width corrections to the neural tangent kernel*. 2019. arXiv: [1909.05989](https://arxiv.org/abs/1909.05989).
- [39] Z. Hu and H. Huang. “On the Random Conjugate Kernel and Neural Tangent Kernel”. *International Conference on Machine Learning*. PMLR. 2021, pp. 4359–4368.
- [40] M. B. Li, M. Nica, and D. Roy. “The future is log-Gaussian: ResNets and their infinite-depth-and-width limit at initialization”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 7852–7864.
- [41] L. Noci, G. Bachmann, K. Roth, S. Nowozin, and T. Hofmann. “Precise characterization of the prior predictive distribution of deep ReLU networks”. In: *Advances in Neural Information Processing Systems* 34 (2021).
- [42] B. Hanin and G. Paouris. “Non-asymptotic results for singular values of Gaussian matrix products”. In: *Geometric and Functional Analysis* 31.2 (2021), pp. 268–324.
- [43] G. Akemann, Z. Burda, and M. Kieburg. “From integrable to chaotic systems: Universal local statistics of Lyapunov exponents”. In: *Europhysics Letters* 126.4 (2019), p. 40001.
- [44] J. Ipsen. “Lyapunov exponents for products of rectangular real, complex and quaternionic Ginibre matrices”. In: *Journal of Physics A: Mathematical and Theoretical* 48.15 (2015), p. 155204.

- [45] T. Jiang and Y. Qi. “Spectral radii of large non-Hermitian random matrices”. In: *Journal of Theoretical Probability* 30.1 (2017), pp. 326–364.
- [46] B. He, J. Martens, G. Zhang, A. Botev, A. Brock, S. L. Smith, and Y. W. Teh. *Deep Transformers without Shortcuts: Modifying Self-attention for Faithful Signal Propagation*. 2023. arXiv: [2302.10322](https://arxiv.org/abs/2302.10322).
- [47] M. Dehghani, J. Djolonga, B. Mustafa, P. Padlewski, J. Heek, J. Gilmer, A. Steiner, M. Caron, R. Geirhos, I. Alabdulmohsin, et al. *Scaling vision transformers to 22 billion parameters*. 2023. arXiv: [2302.05442](https://arxiv.org/abs/2302.05442).
- [48] R. Xiong, Y. Yang, D. He, K. Zheng, S. Zheng, C. Xing, H. Zhang, Y. Lan, L. Wang, and T. Liu. “On layer normalization in the transformer architecture”. *International Conference on Machine Learning*. PMLR. 2020, pp. 10524–10533.
- [49] J. L. Ba, J. R. Kiros, and G. E. Hinton. “Layer normalization”. In: (2016). arXiv: [1607.06450](https://arxiv.org/abs/1607.06450).
- [50] S. Hayou and G. Yang. *Width and Depth Limits Commute in Residual Networks*. 2023. arXiv: [2302.00453](https://arxiv.org/abs/2302.00453).
- [51] S. Hayou, E. Clerico, B. He, G. Deligiannidis, A. Doucet, and J. Rousseau. “Stable resnet”. *International Conference on Artificial Intelligence and Statistics*. PMLR. 2021, pp. 1324–1332.
- [52] W. Tarnowski, P. Warchoł, S. Jastrzebski, J. Tabor, and M. Nowak. “Dynamical isometry is achieved in residual networks in a universal way for any activation function”. *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 2019, pp. 2221–2230.
- [53] S. Hayou. *On the infinite-depth limit of finite-width neural networks*. 2022. arXiv: [2210.00688](https://arxiv.org/abs/2210.00688).
- [54] G. Yang. “Tensor programs ii: Neural tangent kernel for any architecture”. In: *arXiv preprint arXiv:2006.14548* (2020).
- [55] K. He, X. Zhang, S. Ren, and J. Sun. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification”. *Proceedings of the IEEE international conference on computer vision*. 2015, pp. 1026–1034.
- [56] A. Veit, M. J. Wilber, and S. Belongie. “Residual networks behave like ensembles of relatively shallow networks”. In: *Advances in neural information processing systems* 29 (2016).
- [57] S. Zhai, T. Likhomanenko, E. Littwin, D. Busbridge, J. Ramapuram, Y. Zhang, J. Gu, and J. Susskind. *Stabilizing Transformer Training by Preventing Attention Entropy Collapse*. 2023. arXiv: [2303.06296](https://arxiv.org/abs/2303.06296).
- [58] I. Karatzas and S. Shreve. *Brownian motion and stochastic calculus*. Vol. 113. Springer Science & Business Media, 2012.
- [59] J. Miller. *Stochastic Calculus (Lecture Notes)*. http://www.statslab.cam.ac.uk/~jpm205/teaching/lent2016/lecture_notes.pdf. 2015.
- [60] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari. “When do neural networks outperform kernel methods?” In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 14820–14830.
- [61] E. Abbe, E. B. Adsera, and T. Misiakiewicz. “The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks”. *Conference on Learning Theory*. PMLR. 2022, pp. 4782–4887.
- [62] J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang. “High-dimensional asymptotics of feature learning: How one gradient step improves the representation”. *Int. Conf. Learning Representations (ICLR)*. 2022. URL: <https://openreview.net/forum?id=akddwRG6EGi>.
- [63] A. Damian, J. Lee, and M. Soltanolkotabi. “Neural networks can learn representations with gradient descent”. *Conference on Learning Theory*. PMLR. 2022, pp. 5413–5452.
- [64] A. Mousavi-Hosseini, S. Park, M. Girotti, I. Mitliagkas, and M. A. Erdogdu. *Neural Networks Efficiently Learn Low-Dimensional Representations with SGD*. 2022. arXiv: [2209.14863](https://arxiv.org/abs/2209.14863).
- [65] E. Abbe, E. Boix-Adsera, and T. Misiakiewicz. *SGD learning on neural networks: leap complexity and saddle-to-saddle dynamics*. 2023. arXiv: [2302.11055](https://arxiv.org/abs/2302.11055).
- [66] R. Berthier, A. Montanari, and K. Zhou. *Learning time-scales in two-layers neural networks*. 2023. arXiv: [2303.00055](https://arxiv.org/abs/2303.00055).

- [67] P. L. Bartlett, A. Montanari, and A. Rakhlin. “Deep learning: a statistical viewpoint”. In: *Acta numerica* 30 (2021), pp. 87–201.
- [68] O. Kallenberg. *Foundations of Modern Probability*. Probability theory and stochastic modelling. Springer, 2021. ISBN: 9783030618728.
- [69] HuggingFace. *Wikipedia data set*. <https://huggingface.co/datasets/wikipedia>.
- [70] HuggingFace. *Bookcorpus data set*. <https://huggingface.co/datasets/bookcorpus>.
- [71] D. P. Kingma and J. Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [72] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. “GLUE: A multi-task benchmark and analysis platform for natural language understanding”. In: *arXiv preprint arXiv:1804.07461* (2018).
- [73] HuggingFace. *Hugging face Bert implementation*. https://github.com/huggingface/transformers/blob/main/src/transformers/models/bert/modeling_bert.py.

A Preliminaries: Covariance SDE Framework

In this section, we will review existing results on Markov chain convergence to an SDE, as well as existing results for the covariance SDE. See also the Appendix of M. Li et al. [18] for more details.

Firstly, we will define the Skorohod J_1 -topology, or just the Skorohod topology for short [68, Appendix 5]. The topology is used to describe convergence of continuous time processes with discontinuities, in particular, Markov chains with a continuous time interpolation fits in this category. Let S be a complete separable metric space, and $D_{\mathbb{R}_+, S}$ be the space of càdlàg functions (right continuous with left limits) from $\mathbb{R}_+ \rightarrow S$. We use $x_n \xrightarrow{ul} x$ to denote locally uniform convergence (uniform on compact subsets of \mathbb{R}_+), and consider the class of bijections λ on \mathbb{R}_+ so that λ is strictly increasing with $\lambda_0 = 0$ (can be interpreted as a time change).

Definition A.1. We define *Skorohod convergence* $x_n \xrightarrow{s} x$ on $D_{\mathbb{R}_+, S}$ if there exists a sequence of bijections λ_n satisfying above conditions and

$$\lambda_n \xrightarrow{ul} Id, \quad x_n \circ \lambda_n \xrightarrow{ul} x. \quad (13)$$

In particular, we call the topology that induces this convergence the *Skorohod topology* [68, Theorem A5.3].

On a heuristic level, if we have sequences of Markov chains Y^n that satisfy the following type of Euler updates

$$Y_{\ell+1}^n = Y_\ell^n + \frac{b(Y_\ell^n)}{n} + \frac{\sigma(Y_\ell^n) \xi_\ell}{\sqrt{n}}, \quad (14)$$

where ξ_ℓ are iid random variables with zero mean and identity covariance, then we can interpolate this process in continuous time with $X_t^n = Y_{\lfloor tn \rfloor}^n$ and show that as $n \rightarrow \infty$, we have that X^n converges to the solution of the following SDE (weakly with respect to the Skorohod topology)

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \quad X_0 = \lim_{n \rightarrow \infty} Y_0^n. \quad (15)$$

Our next theorem essentially weakens this result in several ways. Firstly, we don't need to take the step size n^{-1} , but instead replace it with n^{-2p} for all $p > 0$. Next, we can allow the update to contain higher order terms that vanish, in particular, any terms of order $O(n^{-2p-\delta})$ for all $\delta > 0$. Here, we remind the readers that $O(n^{-p})$ denotes a random variable Z such that $n^p Z$ has all moment bounded by a universal constant (independent of n). Thirdly, we can allow $b(y)$ to be random, or more precisely replace it with $\widehat{b}(y, \omega_n)$ such that $\mathbb{E} \widehat{b}(y, \omega_n) = b(y)$. Fourthly, we can also allow b, σ to be a sequence b_n, σ_n such that they converge to b, σ on uniformly on compact sets. Finally, we can also weaken the topology of convergence to locally, that is all processes are stopped by a stopping time as in Definition 4.1.

Now we will state the main technical result in this section.

Proposition A.2 (Convergence of Markov Chains to SDE, Proposition A.6, [18]). *Let Y^n be a discrete time Markov chain on \mathbb{R}^N defined by the following update for $p, \delta > 0$*

$$Y_{\ell+1}^n = Y_\ell^n + \frac{\widehat{b}_n(Y_\ell^n, \omega_\ell^n)}{n^{2p}} + \frac{\sigma_n(Y_\ell^n)}{n^p} \xi_\ell^n + O(n^{-2p-\delta}), \quad (16)$$

where $\xi_\ell^n \in \mathbb{R}^N$ are iid random variables with zero mean, identity covariance, and moments uniformly bounded in n . Furthermore, ω_ℓ^n are also iid random variables such that $\mathbb{E}[\widehat{b}_n(Y_\ell^n, \omega_\ell^n) | Y_\ell^n = y] = b_n(y)$ and $\widehat{b}_n(y, \omega_\ell^n)$ has uniformly bounded moments in n . Finally, σ_n is a deterministic function, and the remainder terms in $O(n^{-2p-\delta})$ have uniformly bounded moments in n .

Suppose b_n, σ_n are uniformly Lipschitz functions in n and converges to b, σ uniformly on compact sets, then in the limit as $n \rightarrow \infty$, the process $X_t^n = Y_{\lfloor tn^{2p} \rfloor}^n$ converges in distribution to the solution of the following SDE in the Skorohod topology of $D_{\mathbb{R}_+, \mathbb{R}^N}$

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \quad X_0 = \lim_{n \rightarrow \infty} Y_0^n. \quad (17)$$

Suppose otherwise b_n, σ_n are only locally Lipschitz (but still uniform in n), then X^n converges locally to X in the same topology (see Definition 4.1). More precisely, for any fixed $r > 0$, we consider the stopping times

$$\tau^n := \inf \{t \geq 0 : |X_t^n| \geq r\}, \quad \tau := \inf \{t \geq 0 : |X_t| \geq r\}, \quad (18)$$

then the stopped process $X_{t \wedge \tau^n}^n$ converges in distribution to the stopped solution $X_{t \wedge \tau}$ of the above SDE in the same topology.

We will briefly recall the main result of M. Li et al. [18] next. As mentioned earlier in the text, the setting is for MLPs defined as follows

$$X_{\ell+1} = \sigma_s(X_\ell) \sqrt{\frac{c}{n}} W_\ell, \quad (19)$$

where $\sigma_s(x) = s_+ \max(x, 0) + s_- \min(x, 0)$ with $s_\pm = 1 + \frac{c_\pm}{\sqrt{n}}$, $c^{-1} = \mathbb{E} \sigma_s(g)^2$ for $g \sim \mathcal{N}(0, 1)$, n is the width of the network, and $W_{\ell, ij} \sim \mathcal{N}(0, 1)$ are iid random weights.

Then it was shown that in the limit as $n, d \rightarrow \infty$ with $\frac{d}{n} = T > 0$, the covariance matrix $V_\ell = \frac{1}{n} X_\ell X_\ell^\top$ (or equivalent the post activation $\frac{c}{n} \sigma_s(X_\ell) \sigma_s(X_\ell)^\top$ since the activation becomes infinitesimal) converges the solution of the following SDE in the Skorohod topology [18, Theorem 3.2]

$$dV_t = b_{\text{ReLU}}(V_t) dt + \Sigma_{\text{lin}}(V_t)^{1/2} dB_t, \quad V_0 = \frac{1}{n} X_0 X_0^\top, \quad (20)$$

where the coefficients were given in Theorem 3.2.

We remark that if the output is defined as $X_{\text{out}} = \frac{1}{\sqrt{n}} W_{\text{out}} X_d \in \mathbb{R}^{n_{\text{out}}}$, then we can recover the output distribution as

$$X_{\text{out}} \sim N(0, V_T \otimes I_{n_{\text{out}}}), \quad (21)$$

where we treated X_{out} as a vector in $\mathbb{R}^{m_{\text{out}}}$.

B SDE for Residual Network

Recall that we adopt the following model:

$$X_{\ell+1} = \lambda X_\ell + \gamma \frac{1}{\sqrt{n}} \sigma_s \left(\sqrt{\frac{c}{n}} X_\ell W_\ell^{\text{pre}} \right) W_\ell^{\text{post}}, \quad (22)$$

where $\sigma_s(x) := s_+ \max(x, 0) + s_- \min(x, 0)$ is the shaped ReLU with slopes $s_\pm := 1 + \frac{c_\pm}{\sqrt{n}}$ for some constants $c_+, c_- \in \mathbb{R}$. We let the weights be $(W_\ell^{\text{pre}})_{ij} (W_\ell^{\text{post}})_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$ and $c^{-1} = \mathbb{E} \sigma_s(g)^2$ for $g \sim \mathcal{N}(0, 1)$ is the He initialization constant [55].

From here onwards, we will define the filtration $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$, where $\sigma(\cdot)$ denotes the sigma-algebra generated by the random variable. Furthermore, we will define the conditional expectation $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$.

We are interested in studying the *neural covariance*, i.e. $V_\ell^{\alpha\beta} := \frac{c}{n} \langle x_\ell^\alpha, x_\ell^\beta \rangle$ where $\ell \in [d]$ indexes the layers and d is the depth. In particular, we are interested in understanding the simultaneous limit $(n, d) \rightarrow \infty$, where the ratio $t = d/n$ remains constant.

From now on, we will remove the dependence on ℓ . Defining the residual branch as $f(X, W) := \frac{1}{\sqrt{n}} \sigma_s \left(\sqrt{\frac{c}{n}} X W^{\text{pre}} \right) W^{\text{post}}$, we can write $V^{\alpha\beta}$ as:

$$V^{\alpha\beta} = \lambda^2 X + \lambda \gamma \langle x^\alpha, f(X, W)^\beta \rangle + \lambda \gamma \langle x^\beta, f(X, W)^\alpha \rangle + \gamma^2 \langle f(X, W)^\alpha, f(X, W)^\beta \rangle.$$

For the cross product terms, we get:

$$\begin{aligned} \langle x^\alpha, f(X, W)^\beta \rangle &= \frac{\sqrt{c}}{n} \langle x^\alpha, (\sigma_s(X W^{\text{pre}}) W^{\text{post}})^\beta \rangle = \frac{\sqrt{c}}{n} \sum_{i,j} x_i^\alpha W_{ji}^{\text{post}} \sigma_s \left(\sum_{j'} x_{j'}^\beta W_{j'j}^{\text{pre}} \right), \\ \langle x^\beta, f(X, W)^\alpha \rangle &= \frac{\sqrt{c}}{n} \sum_{i,j} x_i^\beta W_{ji}^{\text{post}} \sigma_s \left(\sum_{j'} x_{j'}^\alpha W_{j'j}^{\text{pre}} \right). \end{aligned}$$

For the term in γ^2 :

$$\begin{aligned} \frac{c}{n^2} \langle (\sigma_s (XW^{\text{pre}}) W^{\text{post}})^\alpha, (\sigma_s (XW^{\text{pre}}) W^{\text{post}})^\beta \rangle &= \frac{c}{n^2} \sum_{i=1}^n (\sigma_s (XW^{\text{pre}}) W^{\text{post}})_{\alpha i} (\sigma_s (XW^{\text{pre}}) W^{\text{post}})_{\beta i} \\ &= \frac{c}{n^2} \sum_{i,j,j'=1}^n W_{ji}^{\text{post}} W_{j'i}^{\text{post}} \sigma_s \left(\sum_k x_k^\alpha W_{kj}^{\text{pre}} \right) \sigma_s \left(\sum_{k'} x_{k'}^\beta W_{k'j'}^{\text{pre}} \right). \end{aligned}$$

We will define the following terms

$$\mathcal{T}_1^{\alpha\beta} := \frac{c\sqrt{c}}{n\sqrt{n}} \sum_{i,j=1}^n W_{ji}^{\text{post}} \left(x_i^\alpha \sigma_s \left(\sum_{j'} x_{j'}^\beta W_{j'j}^{\text{pre}} \right) + x_i^\beta \sigma_s \left(\sum_{j'} x_{j'}^\alpha W_{j'j}^{\text{pre}} \right) \right),$$

and

$$\mathcal{T}_2^{\alpha\beta} := \frac{c^2}{n^2\sqrt{n}} \sum_{i,j,j'=1}^n W_{ji}^{\text{post}} W_{j'i}^{\text{post}} \sigma_s \left(\sum_k x_k^\alpha W_{kj}^{\text{pre}} \right) \sigma_s \left(\sum_{k'} x_{k'}^\beta W_{k'j'}^{\text{pre}} \right).$$

Hence we get the following update for $V_\ell^{\alpha\beta} := \frac{c}{n} \langle x_\ell^\alpha, x_\ell^\beta \rangle$:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \mathcal{T}_2^{\alpha\beta}.$$

It is easy to see that the cross product terms have (conditional) mean zero. For the term with γ^2 :

$$\mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] = \sqrt{n} \sqrt{V^{\alpha\alpha} V^{\beta\beta}} cK_1(\rho^{\alpha\beta}), \quad (23)$$

where $K_1(\rho^{\alpha\beta}) := \mathbb{E}[\sigma_s(g^\alpha) \sigma_s(g^\beta)]$ where $g^\alpha, g^\beta \sim \mathcal{N}(0, 1)$ and $\mathbb{E}[g^\alpha g^\beta] = \frac{\langle x^\alpha, x^\beta \rangle}{\|x^\alpha\| \|x^\beta\|}$. Here we recall $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ is the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$.

M. Li et al. [18] showed that if the non linearity scaling exponent is $p = 1/2$, then:

$$cK_1(\rho) = \rho + \frac{\nu(\rho)}{n} + \mathcal{O}(n^{-3/2}),$$

where $\nu(\rho) = \frac{(c_+ + c_-)^2}{2\pi} \left(\sqrt{1 - \rho^2} + \rho \arccos(\rho) \right)$. Using this result, and summing and subtracting the mean of \mathcal{T}_2 :

$$\begin{aligned} V_{\ell+1}^{\alpha\beta} &= \lambda^2 V_\ell^{\alpha\beta} + \gamma^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta}} cK_1(\rho^{\alpha\beta}) + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} (\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}]) \\ &= \lambda^2 V_\ell^{\alpha\beta} + \gamma^2 V_\ell^{\alpha\beta} + \gamma^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta}} \frac{\nu(\rho^{\alpha\beta})}{n} + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} (\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}]) + \mathcal{O}(n^{-3/2}). \end{aligned}$$

Using $\lambda^2 + \gamma^2 = 1$:

$$V_{\ell+1}^{\alpha\beta} = V_\ell^{\alpha\beta} + \gamma^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta}} \frac{\nu(\rho^{\alpha\beta})}{n} + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} (\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}]) + \mathcal{O}(n^{-3/2}).$$

Furthermore, we need second order moments of \mathcal{T}_1 and \mathcal{T}_2 . We derive them in the following Lemmas:

Lemma B.1.

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega}] &= V^{\alpha\delta} \sqrt{V^{\beta\beta} V^{\omega\omega}} cK_1(\rho^{\beta\omega}) + V^{\alpha\omega} \sqrt{V^{\beta\beta} V^{\delta\delta}} cK_1(\rho^{\beta\delta}) \\ &\quad + V^{\beta\delta} \sqrt{V^{\alpha\alpha} V^{\omega\omega}} cK_1(\rho^{\alpha\omega}) + V^{\beta\omega} \sqrt{V^{\alpha\alpha} V^{\delta\delta}} cK_1(\rho^{\alpha\delta}) \\ &= 2V^{\alpha\delta} V^{\beta\omega} + 2V^{\alpha\omega} V^{\beta\delta} + \mathcal{O}(n^{-1}). \end{aligned}$$

Proof. Recall the definition of $\mathcal{T}_1^{\alpha\beta}$:

$$\mathcal{T}_1^{\alpha\beta} := \frac{c\sqrt{c}}{n\sqrt{n}} \sum_{i,j=1}^n W_{ji}^{\text{post}} \left(x_i^\alpha \sigma_s \left(\sum_{j'} x_{j'}^\beta W_{j'j}^{\text{pre}} \right) + x_i^\beta \sigma_s \left(\sum_{j'} x_{j'}^\alpha W_{j'j}^{\text{pre}} \right) \right),$$

We have that:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega}] &= \frac{c^3}{n^3} \sum_{ij'i'j'} \mathbb{E}[W_{ji} W_{j'i'}] \mathbb{E}_\ell \left[\left(x_i^\alpha \|x^\beta\| \sigma_s(g_j^\beta) + x_i^\beta \|x^\alpha\| \sigma_s(g_j^\alpha) \right) \left(x_i^\delta \|x^\omega\| \sigma_s(g_j^\omega) + x_i^\omega \|x^\delta\| \sigma_s(g_j^\delta) \right) \right] \\ &= \frac{c^3}{n^3} \sum_{ij} \mathbb{E}_\ell \left[x_i^\alpha x_i^\delta \|x^\beta\| \|x^\omega\| \sigma_s(g_j^\beta) \sigma_s(g_j^\omega) + x_i^\alpha x_i^\omega \|x^\beta\| \|x^\delta\| \sigma_s(g_j^\beta) \sigma_s(g_j^\delta) \right. \\ &\quad \left. + x_i^\beta x_i^\delta \|x^\alpha\| \|x^\omega\| \sigma_s(g_j^\alpha) \sigma_s(g_j^\omega) + x_i^\beta x_i^\omega \|x^\alpha\| \|x^\delta\| \sigma_s(g_j^\alpha) \sigma_s(g_j^\delta) \right] \\ &= V^{\alpha\delta} \sqrt{V^{\beta\beta} V^{\omega\omega}} c K_1(\rho^{\beta\omega}) + V^{\alpha\omega} \sqrt{V^{\beta\beta} V^{\delta\delta}} c K_1(\rho^{\beta\delta}) \\ &\quad + V^{\beta\delta} \sqrt{V^{\alpha\alpha} V^{\omega\omega}} c K_1(\rho^{\alpha\omega}) + V^{\beta\omega} \sqrt{V^{\alpha\alpha} V^{\delta\delta}} c K_1(\rho^{\alpha\delta}) \end{aligned}$$

□

Lemma B.2.

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] &= c\sqrt{n} \sqrt{V^{\alpha\alpha} V^{\beta\beta}} K_1(\rho^{\alpha\beta}) = \sqrt{n} V^{\alpha\beta} + \frac{1}{\sqrt{n}} \sqrt{V^{\alpha\alpha} V^{\beta\beta}} \nu(\rho) + \mathcal{O}(n^{-1}), \\ \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] \mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}] &= 2V^{\alpha\delta} V^{\beta\omega} + 2V^{\alpha\omega} V^{\beta\delta} + \mathcal{O}(n^{-1}). \end{aligned}$$

Proof. Recall the definition of $\mathcal{T}_2^{\alpha\beta}$:

$$\begin{aligned} \mathcal{T}_2^{\alpha\beta} &:= \frac{c^2}{n^2 \sqrt{n}} \sum_{i,j,j'=1}^n W_{ji}^{\text{post}} W_{j'i}^{\text{post}} \sigma_s \left(\sum_k x_k^\alpha W_{kj}^{\text{pre}} \right) \sigma_s \left(\sum_{k'} x_{k'}^\beta W_{k'j'}^{\text{pre}} \right) \\ &= \frac{c^2}{n^2 \sqrt{n}} \|x^\alpha\| \|x^\beta\| \sum_{i,j,j'=1}^n W_{ji}^{\text{post}} W_{j'i}^{\text{post}} \sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \end{aligned}$$

For the mean, using the independence between W^{pre} and W^{post} , we have that:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] &= \frac{c^2}{n^2 \sqrt{n}} \|x^\alpha\| \|x^\beta\| \sum_{i,j,j'=1}^n \mathbb{E}[W_{ji}^{\text{post}} W_{j'i}^{\text{post}}] \mathbb{E}[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta)] \\ &= \frac{c^2}{n\sqrt{n}} \|x^\alpha\| \|x^\beta\| \sum_{j=1}^n \mathbb{E}[\sigma_s(g_j^\alpha) \sigma_s(g_j^\beta)] \\ &= \frac{c^2}{\sqrt{n}} \|x^\alpha\| \|x^\beta\| K_1(\rho^{\alpha\beta}) \\ &= c\sqrt{n} \sqrt{V^{\alpha\alpha} V^{\beta\beta}} K_1(\rho^{\alpha\beta}), \end{aligned}$$

which is the desired result. The final expression is the result of the aforementioned expansion for K_1 [18]. For the covariance, we have that:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] &= \frac{c^4}{n^5} \|x^\alpha\| \|x^\beta\| \|x^\delta\| \|x^\omega\| \sum_{i,j,j',i',k,k'} \mathbb{E}[W_{ji}^{\text{post}} W_{j'i}^{\text{post}} W_{ki'}^{\text{post}} W_{k'i'}^{\text{post}}] \mathbb{E}[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \sigma_s(g_k^\delta) \sigma_s(g_{k'}^\omega)] \\ &= \frac{c^2}{n^3} \sqrt{V^{\alpha\alpha} V^{\beta\beta} V^{\delta\delta} V^{\omega\omega}} \sum_{i,j,j',i',k,k'} (\delta_{jj'} \delta_{kk'} + \delta_{jk} \delta_{ii'} \delta_{j'k'} + \delta_{jk'} \delta_{ii'} \delta_{j'k}) \mathbb{E}[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \sigma_s(g_k^\delta) \sigma_s(g_{k'}^\omega)]. \end{aligned}$$

Let's look at each term of the sum separately. For the first term we have that:

$$\begin{aligned} & \sum_{i,j,j',i',k,k'} \delta_{jj'} \delta_{kk'} \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \sigma_s(g_k^\delta) \sigma_s(g_{k'}^\omega) \right] \\ &= n^2 \sum_{j,k} \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_j^\beta) \sigma_s(g_k^\delta) \sigma_s(g_k^\omega) \right] \\ &= n^2 \sum_j \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_j^\beta) \sigma_s(g_j^\delta) \sigma_s(g_j^\omega) \right] + n^2 \sum_{j \neq k} \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_j^\beta) \right] \mathbb{E} \left[\sigma_s(g_k^\delta) \sigma_s(g_k^\omega) \right] \\ &= n^3 K_2(\rho^{\alpha\beta\delta\omega}) + n^3(n-1) K_1(\rho^{\alpha\beta}) K_1(\rho^{\delta\omega}), \end{aligned}$$

where we have defined the fourth moment of the shaped activation: $K_2(\rho^{\alpha\beta\delta\omega}) := \mathbb{E} \left[\sigma_s(g^\alpha) \sigma_s(g^\beta) \sigma_s(g^\delta) \sigma_s(g^\omega) \right]$ for which it holds that [18, Lemma C.2]:

$$K_2(\rho^{\alpha\beta\delta\omega}) = \mathbb{E}[g^\alpha g^\beta g^\delta g^\omega] + \mathcal{O}(n^{-1/2}) = \rho^{\alpha\beta} \rho^{\delta\omega} + \rho^{\alpha\delta} \rho^{\beta\omega} + \rho^{\alpha\omega} \rho^{\beta\delta} + \mathcal{O}(n^{-1/2}).$$

The other two summands can be solved similarly:

$$\sum_{i,j,j',i',k,k'} \delta_{jk} \delta_{ii'} \delta_{j'k'} \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \sigma_s(g_k^\delta) \sigma_s(g_{k'}^\omega) \right] = n^2 K_2(\rho^{\alpha\beta\delta\omega}) + n^2(n-1) K_1(\rho^{\alpha\delta}) K_1(\rho^{\beta\omega}),$$

and

$$\sum_{i,j,j',i',k,k'} \delta_{jk'} \delta_{ii'} \delta_{j'k} \mathbb{E} \left[\sigma_s(g_j^\alpha) \sigma_s(g_{j'}^\beta) \sigma_s(g_k^\delta) \sigma_s(g_{k'}^\omega) \right] = n^2 K_2(\rho^{\alpha\beta\delta\omega}) + n^2(n-1) K_1(\rho^{\alpha\omega}) K_1(\rho^{\beta\delta}).$$

Hence, summing the three terms:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] &= c^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta} V^{\delta\delta} V^{\omega\omega}} \left(K_2(\rho^{\alpha\beta\delta\omega}) + (n-1) K_1(\rho^{\alpha\beta}) K_1(\rho^{\delta\omega}) \right) \\ &\quad + \frac{1}{n} \left(K_2(\rho^{\alpha\beta\delta\omega}) + (n-1) K_1(\rho^{\alpha\delta}) K_1(\rho^{\beta\omega}) \right) + \frac{1}{n} \left(K_2(\rho^{\alpha\beta\delta\omega}) + (n-1) K_1(\rho^{\alpha\omega}) K_1(\rho^{\beta\delta}) \right) \\ &= c^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta} V^{\delta\delta} V^{\omega\omega}} \left(K_2(\rho^{\alpha\beta\delta\omega}) + n K_1(\rho^{\alpha\beta}) K_1(\rho^{\delta\omega}) - K_1(\rho^{\alpha\beta}) K_1(\rho^{\delta\omega}) \right) \\ &\quad + K_1(\rho^{\alpha\delta}) K_1(\rho^{\beta\omega}) + K_1(\rho^{\alpha\omega}) K_1(\rho^{\beta\delta}) + \mathcal{O}(n^{-1}). \end{aligned}$$

Now, subtracting $\mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] \mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}]$, using the aforementioned expansions for K_1 and K_2 :

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] \mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}] &= c^2 \sqrt{V^{\alpha\alpha} V^{\beta\beta} V^{\delta\delta} V^{\omega\omega}} \left(\rho^{\alpha\beta} \rho^{\delta\omega} + \rho^{\alpha\delta} \rho^{\beta\omega} + \rho^{\alpha\omega} \rho^{\beta\delta} \right) \\ &\quad - V^{\alpha\beta} V^{\delta\omega} + V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta} + \mathcal{O}(n^{-1}) \\ &= 2V^{\alpha\delta} V^{\beta\omega} + 2V^{\alpha\omega} V^{\beta\delta} + \mathcal{O}(n^{-1}), \end{aligned}$$

where in the final step we have used the fact that $c = 1 + \mathcal{O}(n^{-1/2})$. This completes the proof. \square

Finally we also have that \mathcal{T}_1 and \mathcal{T}_2 are uncorrelated, i.e.

Lemma B.3.

$$\mathbb{E}_\ell[\mathcal{T}_1^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] = 0.$$

Proof. It is easy to see that $\mathbb{E}_\ell \mathcal{T}_1^{\alpha\beta} \mathcal{T}_2^{\delta\omega}$ involve odd standard Gaussian moments (in particular third order moments), which vanish due to the parity of the centered Gaussian measure. \square

Theorem 3.2. Let X_ℓ be the hidden layers of a ResNet defined in Eq. 4 with $\lambda^2 + \gamma^2 = 1$, where both λ and γ do not depend on d, n . Then the feature covariance V_ℓ converges to the solution of the following SDE (in the sense of Definition 3.1)

$$dV_t = b_{res}(V_t) dt + \Sigma_{res}(V_t)^{1/2} dB_t, \quad V_0 = \frac{1}{n} X_0 X_0^\top, \quad (5)$$

where $b_{res}(V) = \gamma^2 b_{ReLU}(V) = \gamma^2 [\nu(\rho^{\alpha\beta}) \sqrt{V^{\alpha\alpha} V^{\beta\beta}}]_{\alpha \leq \beta}$ with $\rho^{\alpha\beta} = V^{\alpha\alpha} (V^{\alpha\alpha} V^{\beta\beta})^{-1/2}$ and

$$\nu(\rho) = \frac{(c_+ - c_-)^2}{2\pi} \left(\sqrt{1 - \rho^2} - \rho \arccos \rho \right), \quad (6)$$

furthermore, $\Sigma_{res}(V) = 2\gamma^2 \Sigma_{lin}(V) = 2\gamma^2 [V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}]_{\alpha \leq \beta, \delta \leq \omega}$.

Proof. We know that:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \mathcal{T}_2^{\alpha\beta}.$$

Using $\lambda^2 + \gamma^2 = 1$, and summing and subtracting the mean of \mathcal{T}_2 , we have that:

$$V_{\ell+1}^{\alpha\beta} = V_\ell^{\alpha\beta} - \gamma^2 V_\ell^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta}] + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} (\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta}]).$$

Using Lemma B.2 for the mean of \mathcal{T}_2 , we have that:

$$V_{\ell+1}^{\alpha\beta} = V_\ell^{\alpha\beta} + \frac{1}{n} \sqrt{V^{\alpha\alpha} V^{\beta\beta}} \nu(\rho) + \frac{\gamma\lambda}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} (\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta}]) + \mathcal{O}(n^{-3/2}),$$

which gives us an expression in the right Markov Chain form with drift term $\sqrt{V^{\alpha\alpha} V^{\beta\beta}} \nu(\rho)$. For the diffusion term, we need to compute the covariance between two different neural covariances $V_{\ell+1}^{\alpha\beta}, V_{\ell+1}^{\delta\omega}$. Using Lemmas B.1 to B.3, we have

$$\begin{aligned} \text{Cov}_\ell (V_{\ell+1}^{\alpha\beta}, V_{\ell+1}^{\delta\omega}) &= \mathbb{E}_\ell \left[\left(\lambda\gamma \mathcal{T}_1^{\alpha\beta} + \gamma^2 \mathcal{T}_2^{\alpha\beta} - \gamma^2 \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta}] \right) \left(\lambda\gamma \mathcal{T}_1^{\delta\omega} + \gamma^2 \mathcal{T}_2^{\delta\omega} - \gamma^2 \mathbb{E}_\ell [\mathcal{T}_2^{\delta\omega}] \right) \right] \\ &= \lambda^2 \gamma^2 \mathbb{E}_\ell [\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega}] + \gamma^4 \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] - \gamma^4 \mathbb{E}_\ell [\mathcal{T}_2^{\alpha\beta}] \mathbb{E}_\ell [\mathcal{T}_2^{\delta\omega}] \\ &= 2\lambda^2 \gamma^2 (V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + 2\gamma^4 (V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + \mathcal{O}(n^{-1}) \\ &= 2\gamma^2 (V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + \mathcal{O}(n^{-1}), \end{aligned}$$

where we use Cov_ℓ to denote the covariance conditioned on the sigma-algebra $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$. Finally, applying Proposition A.2, we get the desired result. \square

C SDE for Softmax-based Attention

C.1 Dot-Product Attention

Dot-product attention applies the Softmax row-wise to the following matrix:

$$Y_\ell = \frac{1}{n} X_\ell \underbrace{W_\ell^K W_\ell^{Q,\top}}_{W_\ell^B} X_\ell^\top,$$

where $W_\ell^K, W_\ell^Q \in \mathbb{R}^{n \times n_k}$ are Gaussian matrices with unit variance entries, and n_k is the queries and keys' dimension. Here we study the first two moments of Y_ℓ . In particular, note that $\mathbb{E}_\ell[Y_\ell] = 0$, where $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ denotes the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$.

For the second moment we have the following Lemma.

Lemma C.1. *Let*

$$Y_\ell = \frac{1}{n} X_\ell W_\ell^K W_\ell^{Q,\top} X_\ell^\top,$$

be the dot-product attention parametrized by $W_\ell^K, W_\ell^Q \in \mathbb{R}^{n \times n_k}$. Then:

$$\mathbb{E}_\ell[Y^{\alpha\beta} Y^{\delta\omega}] = n_k V^{\alpha\delta} V^{\beta\omega}. \quad (24)$$

Proof.

$$\mathbb{E}_\ell[Y^{\alpha\beta} Y^{\delta\omega}] = \frac{1}{n^2} \sum_{kk',jj'} X_k^\alpha X_{k'}^\beta X_j^\delta X_{j'}^\omega \mathbb{E}[W_{kk'}^B W_{jj'}^B].$$

For the expectation, we have that:

$$\begin{aligned}\mathbb{E}[W_{kk'}^B, W_{jj'}^B] &= \sum_{ii'}^{n_k} \mathbb{E} \left[W_{ki}^K W_{ji'}^K W_{k'i}^Q W_{j'i'}^Q \right] \\ &= \sum_{ii'} \mathbb{E} \left[W_{ki}^K W_{ji'}^K \right] \mathbb{E} \left[W_{k'i}^Q W_{j'i'}^Q \right] \\ &= \sum_{ii'} \delta_{kj} \delta_{k'j'} \delta_{ii'} \\ &= n_k \delta_{kj} \delta_{k'j'},\end{aligned}$$

where we recall $W^B = W^K W^{Q, \top}$.

Hence:

$$\mathbb{E}_\ell[Y^{\alpha\beta} Y^{\delta\omega}] = \frac{n_k}{n^2} \sum_{kk'} X_k^\alpha X_{k'}^\beta X_k^\delta X_{k'}^\omega = n_k V^{\alpha\delta} V^{\beta\omega}.$$

□

C.2 Shaping the Softmax

Recall that we have the following model:

$$X_{\ell+1} = \lambda X_\ell + \gamma A_\ell X_\ell \frac{1}{\sqrt{n}} W_\ell^V$$

For the above model, V_ℓ has the following form:

$$V_{\ell+1} = \lambda^2 V_\ell + \frac{\lambda\gamma}{n\sqrt{n}} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top + A_\ell X_\ell W_\ell X_\ell^\top) + \frac{\gamma^2}{n^2} A_\ell X_\ell W_\ell W_\ell^\top X_\ell^\top A_\ell^\top.$$

In order to have infinitesimal updates in V_ℓ , we would intuitively like the attention matrix to be of the form $A^{\alpha\beta} = \delta_{\alpha\beta} + \mathcal{O}(n^{-1})$. In the case of the usual Softmax based attention, we have:

$$A_\ell := \text{Softmax}(\tau^{-1} Y_\ell),$$

where τ^{-1} is a temperature parameter that regulates the entropy of the resulting categorical distribution; τ is often chosen to be large (scale as a power of n or d), as low temperature results in unstable training.

To transform the Softmax into the desired form of $A^{\alpha\beta} = \delta_{\alpha\beta} + \mathcal{O}(n^{-1})$, we first center the attention matrix around the identity matrix:

$$A_\ell = I + \text{Softmax}(\tau^{-1} Y_\ell),$$

and then examine the Taylor-expansion of the Softmax w.r.t τ^{-1} :

$$A_\ell = I + \frac{1}{m} 11^\top + \mathcal{O}(\tau^{-1}). \quad (25)$$

The first few terms of the expansion provides a good approximation of A_ℓ when τ is large. Observe that for fixed m , the zero order term $\frac{1}{m} 11^\top$ is not of $\mathcal{O}(n^{-1})$. In particular, suppose that we use the attention model in Eq. 25. We can show that the expectation of the term in γ^2 (responsible for the drift) has the following form:

$$\gamma^2 \sum_{\delta, \omega=1}^m \mathbb{E}_\ell[A^{\alpha\delta} A^{\beta\omega}] V^{\delta\omega} = \gamma^2 V^{\alpha\beta} + \frac{\gamma^2}{m} \sum_\omega V^{\alpha\omega} + \frac{\gamma^2}{m} \sum_\delta V^{\delta\beta} + \frac{\gamma^2}{m^2} \sum_{\delta\omega} V^{\delta\omega} + \mathcal{O}(\tau^{-1}),$$

where the terms scaled by $\frac{1}{m}$ leads to large incremental updates in expected value of V_ℓ w.r.t the previous layer, and precludes the possibility of convergence to an SDE. Hence, we choose to re-center the Softmax by removing the term $\frac{1}{m} 11^\top$, as follows:

$$A_\ell = I + \text{Softmax}(\tau^{-1} Y_\ell) - \frac{1}{m} 11^\top,$$

which admits the following Taylor-expansion:

$$[\text{Softmax}(\tau^{-1}Y_\ell) - m^{-1}\mathbf{1}\mathbf{1}^\top] = \frac{1}{\tau m} [Y^{\alpha\beta} - \bar{Y}^\alpha]_{\alpha,\beta} + \frac{1}{2\tau^2 m} \left[(Y^{\alpha\beta} - \bar{Y}^\alpha)^2 - ((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2}) \right]_{\alpha,\beta} + O(\tau^{-3}),$$

where $\bar{Y}^\alpha := \frac{1}{m} \sum_\nu Y^{\alpha\nu}$, and $(\bar{Y}^\alpha)^2 := \frac{1}{m} \sum_\nu (Y^{\alpha\nu})^2$.

Hence, up to third order:

$$A_\ell = I + \frac{1}{\tau m} [Y^{\alpha\beta} - \bar{Y}^\alpha]_{\alpha,\beta} + \frac{1}{2\tau^2 m} \left[(Y^{\alpha\beta} - \bar{Y}^\alpha)^2 - ((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2}) \right]_{\alpha,\beta} + O(\tau^{-3}).$$

For sufficiently large τ , the formulation above allows infinitesimal updates in V_ℓ , and as we will show rigorously in the rest of this section, permits convergence to an SDE.

C.3 Lemmas on Moments of Shaped Attention

Define:

$$\begin{aligned} F_1^{\alpha\beta} &= Y^{\alpha\beta} - \bar{Y}^\alpha, \\ F_2^{\alpha\beta} &= (Y^{\alpha\beta} - \bar{Y}^\alpha)^2 - ((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2}). \end{aligned}$$

Hence, A_ℓ can be written as:

$$A_\ell^{\alpha\beta} = \delta_{\alpha\beta} + \frac{1}{\tau m} F_1^{\alpha\beta} + \frac{1}{2\tau^2 m} F_2^{\alpha\beta} + O(\tau^{-3}).$$

We now compute the moments of A_ℓ . We define the following quantities:

$$\begin{aligned} S_1^{\alpha\delta,\beta\omega} &:= \frac{1}{n_k} \mathbb{E}_\ell (Y^{\alpha\delta} - \bar{Y}^\alpha)(Y^{\beta\omega} - \bar{Y}^\beta) = \frac{1}{n_k} \mathbb{E}_\ell F_1^{\alpha\delta} F_1^{\beta\omega}, \\ S_2^{\alpha\delta} &:= \frac{1}{n_k} \mathbb{E}_\ell \left[(Y^{\alpha\delta} - \bar{Y}^\alpha)^2 - ((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2}) \right] = \frac{1}{n_k} \mathbb{E}_\ell F_2^{\alpha\delta}, \end{aligned}$$

where we recall $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ is the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$.

Lemma C.2 (Moments of Taylor Expansion).

$$\begin{aligned} S_1^{\alpha\delta,\beta\omega} &= V^{\alpha\beta} (V^{\delta\omega} - V^{\delta\bar{x}} - V^{\omega\bar{x}} + V^{\bar{x}\bar{x}}), \\ S_2^{\alpha\delta} &= V^{\alpha\alpha} (V^{\delta\delta} - 2V^{\delta\bar{x}} + 2V^{\bar{x}\bar{x}} - \bar{V}), \end{aligned}$$

where $\bar{V} = \frac{1}{m} \sum_\nu V^{\nu\nu}$ and $\bar{x} = \frac{1}{m} \sum_\nu x^\nu$ is the average token.

Proof. Using Lemma C.1 and linearity of expectation:

$$\begin{aligned} S_1^{\alpha\delta,\beta\omega} &= \frac{1}{n_k} \left(\mathbb{E}_\ell [Y^{\alpha\delta} Y^{\beta\omega}] - \mathbb{E}_\ell [Y^{\alpha\delta} \bar{Y}^\beta] - \mathbb{E}_\ell [Y^{\beta\omega} \bar{Y}^\alpha] + \mathbb{E}_\ell [\bar{Y}^\alpha \bar{Y}^\beta] \right) \\ &= (V^{\alpha\beta} V^{\delta\omega} - V^{\alpha\beta} V^{\delta\bar{x}} - V^{\alpha\beta} V^{\omega\bar{x}} + V^{\alpha\beta} V^{\bar{x}\bar{x}}) \\ &= V^{\alpha\beta} (V^{\delta\omega} - V^{\delta\bar{x}} - V^{\omega\bar{x}} + V^{\bar{x}\bar{x}}), \end{aligned}$$

and:

$$\begin{aligned} S_2^{\alpha\delta} &= \frac{1}{n_k} \left(\mathbb{E}_\ell \left[(Y^{\alpha\delta} - \bar{Y}^\alpha)^2 \right] - \mathbb{E}_\ell \left[((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2}) \right] \right) \\ &= \frac{1}{n_k} \left(\mathbb{E}_\ell [(Y^{\alpha\delta})^2] - 2\mathbb{E}_\ell [Y^{\alpha\delta} \bar{Y}^\alpha] + 2\mathbb{E}_\ell [\bar{Y}^{\alpha^2}] - \mathbb{E}_\ell [(\bar{Y}^\alpha)^2] \right) \\ &= (V^{\alpha\alpha} V^{\delta\delta} - 2V^{\alpha\alpha} V^{\delta\bar{x}} + 2V^{\alpha\alpha} V^{\bar{x}\bar{x}} - V^{\alpha\alpha} \bar{V}) \\ &= V^{\alpha\alpha} (V^{\delta\delta} - 2V^{\delta\bar{x}} + 2V^{\bar{x}\bar{x}} - \bar{V}), \end{aligned}$$

where $\bar{V} = \frac{1}{m} \sum_\beta V^{\beta\beta}$. □

Lemma C.3.

$$\mathbb{E}_\ell[A^{\alpha\delta}A^{\beta\omega}] = \delta_{\alpha\delta}\delta_{\beta\omega} + \frac{n_k}{\tau^2m^2}S_1^{\alpha\delta,\beta\omega} + \frac{n_k}{2\tau^2m}(\delta_{\beta\omega}S_2^{\alpha\delta} + \delta_{\alpha\delta}S_2^{\beta\omega}) + O(n_k\tau^{-3}).$$

Proof.

$$\begin{aligned} \mathbb{E}_\ell[A^{\alpha\delta}A^{\beta\omega}] &= \delta_{\alpha\delta}\delta_{\beta\omega} + \frac{\delta_{\beta\omega}}{\tau m}\mathbb{E}_\ell[Y^{\alpha\delta} - \bar{Y}^\alpha] + \frac{\delta_{\alpha\delta}}{\tau m}\mathbb{E}_\ell[Y^{\beta\omega} - \bar{Y}^\beta] + \frac{1}{\tau^2m^2}\mathbb{E}_\ell(Y^{\alpha\delta} - \bar{Y}^\alpha)(Y^{\beta\omega} - \bar{Y}^\beta) \\ &\quad + \frac{\delta_{\beta\omega}}{2\tau^2m}\mathbb{E}_\ell\left[(Y^{\alpha\delta} - \bar{Y}^\alpha)^2 - ((\bar{Y}^\alpha)^2 - \bar{Y}^{\alpha^2})\right] \\ &\quad + \frac{\delta_{\alpha\delta}}{2\tau^2m}\mathbb{E}_\ell\left[(Y^{\beta\omega} - \bar{Y}^\beta)^2 - ((\bar{Y}^\beta)^2 - \bar{Y}^{\beta^2})\right] + O(n_k\tau^{-3}) \\ &= \delta_{\alpha\delta}\delta_{\beta\omega} + \frac{n_k}{\tau^2m^2}S_1^{\alpha\delta,\beta\omega} + \frac{n_k}{2\tau^2m}(\delta_{\beta\omega}S_2^{\alpha\delta} + \delta_{\alpha\delta}S_2^{\beta\omega}) + O(n_k\tau^{-3}). \end{aligned}$$

□

Lemma C.4.

$$\begin{aligned} \mathbb{E}_\ell\left[A^{\alpha\alpha'}A^{\beta\beta'}A^{\delta\delta'}A^{\omega\omega'}\right] &= \delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}\delta_{\omega\omega'} \\ &\quad + \frac{n_k}{\tau^2m^2}\left(\delta_{\alpha\alpha'}\delta_{\beta\beta'}S_1^{\delta\delta',\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\delta\delta'}S_1^{\beta\beta',\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\omega\omega'}S_1^{\beta\beta',\delta\delta'}\right. \\ &\quad + \delta_{\beta\beta'}\delta_{\delta\delta'}S_1^{\alpha\alpha',\omega\omega'} + \delta_{\beta\beta'}\delta_{\omega\omega'}S_1^{\alpha\alpha',\delta\delta'} + \delta_{\omega\omega'}\delta_{\delta\delta'}S_1^{\alpha\alpha',\beta\beta'}\left.)\right) \\ &\quad + \frac{n_k}{2\tau^2m}\left(\delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}S_2^{\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\omega\omega'}S_2^{\delta\delta'}\right. \\ &\quad + \left.\delta_{\alpha\alpha'}\delta_{\omega\omega'}\delta_{\delta\delta'}S_2^{\beta\beta'} + \delta_{\omega\omega'}\delta_{\beta\beta'}\delta_{\delta\delta'}S_2^{\alpha\alpha'}\right) + O(n_k\tau^{-3}). \end{aligned}$$

Proof.

$$\begin{aligned} \mathbb{E}_\ell\left[A^{\alpha\alpha'}A^{\beta\beta'}A^{\delta\delta'}A^{\omega\omega'}\right] &= \mathbb{E}_\ell\left[\left(\delta_{\alpha\alpha'} + \frac{1}{\tau m}F_1^{\alpha\alpha'} + \frac{1}{2\tau^2m}F_2^{\alpha\alpha'}\right)\left(\delta_{\beta\beta'} + \frac{1}{\tau m}F_1^{\beta\beta'} + \frac{1}{2\tau^2m}F_2^{\beta\beta'}\right)\right. \\ &\quad \left.\left(\delta_{\delta\delta'} + \frac{1}{\tau m}F_1^{\delta\delta'} + \frac{1}{2\tau^2m}F_2^{\delta\delta'}\right)\left(\delta_{\omega\omega'} + \frac{1}{\tau m}F_1^{\omega\omega'} + \frac{1}{2\tau^2m}F_2^{\omega\omega'}\right)\right] + O(\tau^{-3}) \\ &= \mathbb{E}_\ell\left[\delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}\delta_{\omega\omega'} + \frac{1}{\tau^2m^2}\delta_{\alpha\alpha'}\delta_{\beta\beta'}F_1^{\delta\delta'}F_1^{\omega\omega'} + \frac{1}{\tau^2m^2}\delta_{\alpha\alpha'}\delta_{\delta\delta'}F_1^{\beta\beta'}F_1^{\omega\omega'}\right. \\ &\quad + \frac{1}{\tau^2m^2}\delta_{\alpha\alpha'}\delta_{\omega\omega'}F_1^{\delta\delta'}F_1^{\beta\beta'} + \frac{1}{\tau^2m^2}\delta_{\beta\beta'}\delta_{\delta\delta'}F_1^{\alpha\alpha'}F_1^{\omega\omega'} \\ &\quad + \frac{1}{\tau^2m^2}\delta_{\beta\beta'}\delta_{\omega\omega'}F_1^{\alpha\alpha'}F_1^{\delta\delta'} + \frac{1}{\tau^2m^2}\delta_{\delta\delta'}\delta_{\omega\omega'}F_1^{\alpha\alpha'}F_1^{\beta\beta'} \\ &\quad + \frac{1}{2\tau^2m}\delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}F_2^{\omega\omega'} + \frac{1}{2\tau^2m}\delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\omega\omega'}F_2^{\delta\delta'} \\ &\quad + \left.\frac{1}{2\tau^2m}\delta_{\alpha\alpha'}\delta_{\omega\omega'}\delta_{\delta\delta'}F_2^{\beta\beta'} + \frac{1}{2\tau^2m}\delta_{\omega\omega'}\delta_{\beta\beta'}\delta_{\delta\delta'}F_2^{\alpha\alpha'}\right] + O(\tau^{-3}). \end{aligned}$$

Using the linearity of expectation:

$$\begin{aligned} \mathbb{E}_\ell\left[A^{\alpha\alpha'}A^{\beta\beta'}A^{\delta\delta'}A^{\omega\omega'}\right] &= \delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}\delta_{\omega\omega'} \\ &\quad + \frac{n_k}{\tau^2m^2}\left(\delta_{\alpha\alpha'}\delta_{\beta\beta'}S_1^{\delta\delta',\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\delta\delta'}S_1^{\beta\beta',\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\omega\omega'}S_1^{\beta\beta',\delta\delta'}\right. \\ &\quad + \delta_{\beta\beta'}\delta_{\delta\delta'}S_1^{\alpha\alpha',\omega\omega'} + \delta_{\beta\beta'}\delta_{\omega\omega'}S_1^{\alpha\alpha',\delta\delta'} + \delta_{\omega\omega'}\delta_{\delta\delta'}S_1^{\alpha\alpha',\beta\beta'}\left.)\right) \\ &\quad + \frac{n_k}{2\tau^2m}\left(\delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\delta\delta'}S_2^{\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\beta\beta'}\delta_{\omega\omega'}S_2^{\delta\delta'}\right. \\ &\quad + \left.\delta_{\alpha\alpha'}\delta_{\omega\omega'}\delta_{\delta\delta'}S_2^{\beta\beta'} + \delta_{\omega\omega'}\delta_{\beta\beta'}\delta_{\delta\delta'}S_2^{\alpha\alpha'}\right) + O(n_k\tau^{-3}). \end{aligned}$$

Note that the terms above can be computed using Lemma C.2. □

Lemma C.5.

$$\begin{aligned} & \mathbb{E}_\ell \left[A^{\alpha\alpha'} A^{\beta\beta'} A^{\delta\delta'} A^{\omega\omega'} \right] - \mathbb{E}_\ell \left[A^{\alpha\alpha'} A^{\beta\beta'} \right] \mathbb{E}_\ell \left[A^{\delta\delta'} A^{\omega\omega'} \right] \\ &= \frac{n_k}{\tau^2 m^2} \left(\delta_{\alpha\alpha'} \delta_{\delta\delta'} S_1^{\beta\beta', \omega\omega'} + \delta_{\alpha\alpha'} \delta_{\omega\omega'} S_1^{\beta\beta', \delta\delta'} + \delta_{\beta\beta'} \delta_{\delta\delta'} S_1^{\alpha\alpha', \omega\omega'} + \delta_{\beta\beta'} \delta_{\omega\omega'} S_1^{\alpha\alpha', \delta\delta'} \right) + \mathcal{O}(n_k \tau^{-3}). \end{aligned}$$

Proof. The results is an immediate consequence of Lemma C.4 and Lemma C.3, where only the terms that do not cancel out are kept. □

C.4 Neural Covariance SDE for Stable Attention

Recall that we have the following model:

$$X_{\ell+1} = \lambda X_\ell + \gamma A_\ell X_\ell \frac{1}{\sqrt{n}} W_\ell^Y$$

For the above model, V_ℓ has the following form:

$$V_{\ell+1} = \lambda^2 V_\ell + \frac{\lambda\gamma}{n\sqrt{n}} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top + A_\ell X_\ell W_\ell X_\ell^\top) + \frac{\gamma^2}{n^2} A_\ell X_\ell W_\ell W_\ell^\top X_\ell^\top A_\ell^\top.$$

We define:

$$\begin{aligned} \mathcal{T}_1^{\alpha\beta} &:= \frac{1}{n} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top + A_\ell X_\ell W_\ell X_\ell^\top)^{\alpha\beta}, \\ \mathcal{T}_2^{\alpha\beta} &:= \frac{1}{n\sqrt{n}} (A_\ell X_\ell W_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta}. \end{aligned}$$

Hence, the expression for V_ℓ simplifies to:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \frac{\lambda\gamma}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \mathcal{T}_2^{\alpha\beta}.$$

We need to compute the moments for these two quantities:

Lemma C.6 (Moments of \mathcal{T}_1).

$$\mathbb{E}_\ell[\mathcal{T}_1^{\alpha\beta}] = 0,$$

and

$$\mathbb{E}_\ell \left[\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega} \right] = 2(V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + \mathcal{O}(n_k \tau^{-3}).$$

Proof.

$$\begin{aligned} \mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega} &= \frac{1}{n^2} \left[(X_\ell W_\ell^\top X_\ell^\top A_\ell^\top + A_\ell X_\ell W_\ell X_\ell^\top)^{\alpha\beta} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top + A_\ell X_\ell W_\ell X_\ell^\top)^{\delta\omega} \right] \\ &= \frac{1}{n^2} \left[(X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\delta\omega} + (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta} (A_\ell X_\ell W_\ell X_\ell^\top)^{\delta\omega} \right. \\ &\quad \left. + (A_\ell X_\ell W_\ell X_\ell^\top)^{\alpha\beta} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\delta\omega} + (A_\ell X_\ell W_\ell X_\ell^\top)^{\alpha\beta} (A_\ell X_\ell W_\ell X_\ell^\top)^{\delta\omega} \right]. \end{aligned}$$

Let's look at the first summand:

$$\frac{1}{n^2} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta} (X_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\delta\omega} = \frac{1}{n^2} \sum_{\nu\kappa} \sum_{kk'jj'} X_k^\alpha W_{k'k} X_{k'}^\nu A^{\beta\nu} X_j^\delta W_{j'j} X_{j'}^\kappa A^{\omega\kappa}.$$

Hence, in expectation with respect to W :

$$\begin{aligned} \frac{1}{n^2} \sum_{\nu\kappa} \sum_{kk'jj'} X_k^\alpha \mathbb{E}[W_{k'k} W_{j'j}] X_{k'}^\nu A^{\beta\nu} X_j^\delta X_{j'}^\kappa A^{\omega\kappa} &= \frac{1}{n^2} \sum_{\nu\kappa} \sum_{kk'jj'} X_k^\alpha \delta_{kj} \delta_{k'j'} X_{k'}^\nu A^{\beta\nu} X_j^\delta X_{j'}^\kappa A^{\omega\kappa} \\ &= \frac{1}{n^2} \sum_{\nu\kappa} \sum_{kk'} X_k^\alpha X_{k'}^\nu A^{\beta\nu} X_k^\delta X_{k'}^\kappa A^{\omega\kappa} \\ &= \sum_{\nu\kappa} V^{\alpha\delta} V^{\nu\kappa} A^{\beta\nu} A^{\omega\kappa}. \end{aligned}$$

An identical argument can be made for the remaining three summands. Hence, taking expectation with respect to the Softmax weights:

$$\mathbb{E}_\ell [\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega}] = \sum_{\nu\kappa} (V^{\alpha\delta} V^{\nu\kappa} \mathbb{E}_\ell[A^{\beta\nu} A^{\omega\kappa}] + V^{\alpha\omega} V^{\nu\kappa} \mathbb{E}_\ell[A^{\beta\nu} A^{\delta\kappa}] + V^{\beta\delta} V^{\nu\kappa} \mathbb{E}_\ell[A^{\alpha\nu} A^{\omega\kappa}] + V^{\beta\omega} V^{\nu\kappa} \mathbb{E}_\ell[A^{\alpha\nu} A^{\delta\kappa}]).$$

Now, using Lemma C.3:

$$\begin{aligned} \mathbb{E}_\ell [\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega}] &= \sum_{\nu\kappa} V^{\nu\kappa} (V^{\alpha\delta} \mathbb{E}_\ell[A^{\beta\nu} A^{\omega\kappa}] + V^{\alpha\omega} \mathbb{E}_\ell[A^{\beta\nu} A^{\delta\kappa}] + V^{\beta\delta} \mathbb{E}_\ell[A^{\alpha\nu} A^{\omega\kappa}] + V^{\beta\omega} \mathbb{E}_\ell[A^{\alpha\nu} A^{\delta\kappa}]) \\ &= \sum_{\nu\kappa} V^{\nu\kappa} (V^{\alpha\delta} \delta_{\beta\nu} \delta_{\omega\kappa} + V^{\alpha\omega} \delta_{\beta\nu} \delta_{\delta\kappa} + V^{\beta\delta} \delta_{\alpha\nu} \delta_{\omega\kappa} + V^{\beta\omega} \delta_{\alpha\nu} \delta_{\delta\kappa}) + \mathcal{O}(n_k \tau^{-2}) \\ &= V^{\beta\omega} V^{\alpha\delta} + V^{\alpha\omega} V^{\beta\delta} + V^{\beta\delta} V^{\alpha\omega} + V^{\beta\omega} V^{\alpha\delta} + \mathcal{O}(n_k \tau^{-2}) \\ &= 2(V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + \mathcal{O}(n_k \tau^{-2}). \end{aligned}$$

□

Lemma C.7 (Moments of \mathcal{T}_2).

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] &= \sqrt{n} \sum_{\nu\kappa} V^{\nu\kappa} \mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa}], \\ \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] &= \sum_{\nu\kappa\nu'\kappa'} \mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa} A^{\delta\nu'} A^{\omega\kappa'}] (n V^{\nu\kappa} V^{\nu'\kappa'} + V^{\nu\nu'} V^{\kappa\kappa'} + V^{\nu\kappa'} V^{\nu'\kappa}). \end{aligned}$$

Proof.

$$\mathcal{T}_2^{\alpha\beta} := \frac{1}{n\sqrt{n}} (A_\ell X_\ell W_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta} = \frac{1}{n\sqrt{n}} \sum_{\nu\kappa} \sum_{kk'j} A^{\alpha\nu} X_k^\nu W_{kk'} W_{jk'} X_j^\kappa A^{\beta\kappa}.$$

Taking expectation with respect to W :

$$\begin{aligned} \frac{1}{n\sqrt{n}} (A_\ell X_\ell W_\ell W_\ell^\top X_\ell^\top A_\ell^\top)^{\alpha\beta} &= \frac{1}{n\sqrt{n}} \sum_{\nu\kappa} \sum_{kk'j} A^{\alpha\nu} X_k^\nu \mathbb{E}[W_{kk'} W_{jk'}] X_j^\kappa A^{\beta\kappa} \\ &= \frac{1}{n\sqrt{n}} \sum_{\nu\kappa} \sum_{kk'j} A^{\alpha\nu} X_k^\nu \delta_{kj} X_j^\kappa A^{\beta\kappa} \\ &= \frac{1}{n\sqrt{n}} \sum_{\nu\kappa} \sum_{kk'} A^{\alpha\nu} A^{\beta\kappa} X_k^\nu X_k^\kappa \\ &= \frac{1}{\sqrt{n}} \sum_{\nu\kappa} \sum_k A^{\alpha\nu} A^{\beta\kappa} X_k^\nu X_k^\kappa \\ &= \sqrt{n} \sum_{\nu\kappa} V^{\nu\kappa} A^{\alpha\nu} A^{\beta\kappa}. \end{aligned}$$

Taking expectation w.r.t the Softmax weights, we get the desired result.

For second moment, we can take the conditional expectation:

$$\mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] = \frac{1}{n^3} \sum_{\nu\kappa\nu'\kappa'} \sum_{kk'jj'ii'} A^{\alpha\nu} A^{\beta\kappa} A^{\delta\nu'} A^{\omega\kappa'} X_k^\nu X_j^\kappa X_{i'}^{\nu'} X_{j'}^{\kappa'} \mathbb{E}[W_{kk'} W_{jk'} W_{ii'} W_{j'i'}],$$

where we recall $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ is the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [l]})$.

Using Isserlis Theorem, we have that:

$$\mathbb{E}[W_{kk'}W_{jk'}W_{ii'}W_{j'i'}] = \delta_{kj}\delta_{ij'} + \delta_{ki}\delta_{k'i'}\delta_{jj'} + \delta_{kj'}\delta_{k'i'}\delta_{ji}.$$

Hence:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}\mathcal{T}_2^{\delta\omega}] &= \frac{1}{n^3} \sum_{\nu\kappa\nu'\kappa'} A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'} \sum_{kk'jj'ii'} X_k^\nu X_j^\kappa X_i^{\nu'} X_{j'}^{\kappa'} \left(\delta_{kj}\delta_{ij'} + \delta_{ki}\delta_{k'i'}\delta_{jj'} + \right. \\ &\quad \left. + \delta_{kj'}\delta_{k'i'}\delta_{ji} \right) \\ &= \frac{1}{n^3} \sum_{\nu\kappa\nu'\kappa'} A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'} \left(\sum_{kk'ii'} X_k^\nu X_k^\kappa X_i^{\nu'} X_i^{\kappa'} + \sum_{kk'j} X_k^\nu X_j^\kappa X_k^{\nu'} X_j^{\kappa'} \right. \\ &\quad \left. + \sum_{kk'j} X_k^\nu X_j^\kappa X_j^{\nu'} X_k^{\kappa'} \right) \\ &= \frac{1}{n^3} \sum_{\nu\kappa\nu'\kappa'} A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'} \left(n^4 V^{\nu\kappa} V^{\nu'\kappa'} + n^3 V^{\nu\nu'} V^{\kappa\kappa'} + n^3 V^{\nu\kappa'} V^{\nu'\kappa} \right) \\ &= \sum_{\nu\kappa\nu'\kappa'} A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'} \left(n V^{\nu\kappa} V^{\nu'\kappa'} + V^{\nu\nu'} V^{\kappa\kappa'} + V^{\nu\kappa'} V^{\nu'\kappa} \right). \end{aligned}$$

By taking expectation w.r.t the Softmax parameters, we get the desired result. \square

Lemma C.8 (Covariance of \mathcal{T}_2).

$$\mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}\mathcal{T}_2^{\delta\omega}] - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}]\mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}] = \frac{nm_k}{\tau^2} A^{\alpha\beta\delta\omega} + V^{\alpha\delta}V^{\beta\omega} + V^{\alpha\omega}V^{\beta\delta} + \mathcal{O}\left(\frac{nm_k}{\tau^3} + \frac{n_k}{\tau^2}\right),$$

where:

$$A^{\alpha\beta\delta\omega} := \frac{1}{m^2} \sum_{\nu\kappa} \left(V^{\alpha\nu}V^{\delta\nu}S_1^{\beta\kappa,\omega\nu} + V^{\alpha\nu}V^{\omega\nu}S_1^{\beta\kappa,\delta\nu} + V^{\beta\nu}V^{\delta\kappa}S_1^{\alpha\nu,\omega\kappa} + V^{\beta\nu}V^{\omega\kappa}S_1^{\alpha\nu,\delta\kappa} \right).$$

Proof. Using Lemma C.7, we have that:

$$\begin{aligned} \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}\mathcal{T}_2^{\delta\omega}] - \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}]\mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}] &= \sum_{\nu\kappa\nu'\kappa'} \mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'}] \left(n V^{\nu\kappa} V^{\nu'\kappa'} + V^{\nu\nu'} V^{\kappa\kappa'} + V^{\nu\kappa'} V^{\nu'\kappa} \right) \\ &\quad - n \sum_{\nu\kappa\nu'\kappa'} V^{\nu\kappa} V^{\nu'\kappa'} \mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}] \mathbb{E}_\ell[A^{\delta\nu'}A^{\omega\kappa'}] \\ &= n \sum_{\nu\kappa\nu'\kappa'} V^{\nu\kappa} V^{\nu'\kappa'} \left(\mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'}] - \mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}] \mathbb{E}_\ell[A^{\delta\nu'}A^{\omega\kappa'}] \right) \\ &\quad + \sum_{\nu\kappa\nu'\kappa'} \mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'}] \left(V^{\nu\nu'} V^{\kappa\kappa'} + V^{\nu\kappa'} V^{\nu'\kappa} \right). \end{aligned}$$

Now we can use Lemma C.3 and Lemma C.4 to compute the moments of A . For the second summand, we simply have:

$$\mathbb{E}_\ell \left[A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'} \right] = \delta_{\alpha\nu}\delta_{\beta\kappa}\delta_{\delta\nu'}\delta_{\omega\kappa'} + \mathcal{O}(n_k\tau^{-2}),$$

hence:

$$\sum_{\nu\kappa\nu'\kappa'} \mathbb{E}_\ell[A^{\alpha\nu}A^{\beta\kappa}A^{\delta\nu'}A^{\omega\kappa'}] \left(V^{\nu\nu'} V^{\kappa\kappa'} + V^{\nu\kappa'} V^{\nu'\kappa} \right) = V^{\alpha\delta}V^{\beta\omega} + V^{\alpha\omega}V^{\beta\delta} + \mathcal{O}(n_k\tau^{-2}).$$

For the first summand, recall from Lemma C.5 that:

$$\begin{aligned} \mathbb{E}_\ell \left[A^{\alpha\alpha'}A^{\beta\beta'}A^{\delta\delta'}A^{\omega\omega'} \right] - \mathbb{E}_\ell \left[A^{\alpha\alpha'}A^{\beta\beta'} \right] \mathbb{E}_\ell \left[A^{\delta\delta'}A^{\omega\omega'} \right] \\ = \frac{n_k}{\tau^2 m^2} \left(\delta_{\alpha\alpha'}\delta_{\delta\delta'}S_1^{\beta\beta',\omega\omega'} + \delta_{\alpha\alpha'}\delta_{\omega\omega'}S_1^{\beta\beta',\delta\delta'} + \delta_{\beta\beta'}\delta_{\delta\delta'}S_1^{\alpha\alpha',\omega\omega'} + \delta_{\beta\beta'}\delta_{\omega\omega'}S_1^{\alpha\alpha',\delta\delta'} \right) \\ + \mathcal{O}(n_k\tau^{-3}). \end{aligned}$$

Hence:

$$\begin{aligned}
 & n \sum_{\nu\kappa\nu'\kappa'} V^{\nu\kappa} V^{\nu'\kappa'} \left(\mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa} A^{\delta\nu'} A^{\omega\kappa'}] - \mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa}] \mathbb{E}_\ell[A^{\delta\nu'} A^{\omega\kappa'}] \right) \\
 &= \frac{nn_k}{\tau^2 m^2} \sum_{\nu\kappa\nu'\kappa'} V^{\nu\kappa} V^{\nu'\kappa'} \left(\delta_{\alpha\nu} \delta_{\delta\nu'} S_1^{\beta\kappa, \omega\kappa'} + \delta_{\alpha\nu} \delta_{\omega\kappa'} S_1^{\beta\kappa, \delta\nu'} + \delta_{\beta\kappa} \delta_{\delta\nu'} S_1^{\alpha\nu, \omega\kappa'} + \delta_{\beta\kappa} \delta_{\omega\kappa'} S_1^{\alpha\nu, \delta\nu'} \right) \\
 &+ \mathcal{O}(nn_k \tau^{-3}) \\
 &= \frac{nn_k}{\tau^2} \frac{1}{m^2} \underbrace{\sum_{\nu\kappa} \left(V^{\alpha\kappa} V^{\delta\nu} S_1^{\beta\kappa, \omega\nu} + V^{\alpha\kappa} V^{\omega\nu} S_1^{\beta\kappa, \delta\nu} + V^{\beta\nu} V^{\delta\kappa} S_1^{\alpha\nu, \omega\kappa} + V^{\beta\nu} V^{\omega\kappa} S_1^{\alpha\nu, \delta\kappa} \right)}_{A^{\alpha\beta\delta\omega}} \\
 &+ \mathcal{O}(nn_k \tau^{-3}).
 \end{aligned}$$

□

We are now ready to re-state and proof of Theorem 4.2.

Theorem 4.2. Let X_ℓ be the hidden layers of a residual attention network defined in Eq. 1 with shaped attention in Eq. 9, parameters $\lambda^2 + \gamma^2 = 1$ and $\tau = \tau_0 \sqrt{nn_k}$, where λ, γ, τ_0 all do not depend on d, n . Then the feature covariance V_ℓ converges locally to the solution of the following SDE (in the sense of Definition 4.1)

$$dV_t = b(V_t)dt + \Sigma(V_t)^{1/2}dB_t, \quad V_0 = \frac{1}{n}X_0X_0^\top,$$

where the drift has the following form

$$b(V) = \frac{\gamma^2}{\tau_0^2} \left[\frac{1}{m^2} \sum_{\nu, \kappa=1}^m V^{\nu\kappa} S_1^{\alpha\nu, \beta\kappa} + \frac{1}{2m} \sum_{\nu=1}^m (V^{\beta\nu} S_2^{\alpha\nu} + V^{\alpha\nu} S_2^{\beta\nu}) \right]_{\alpha \leq \beta},$$

the diffusion coefficient is defined by $\Sigma(V) = \gamma^2(2 - \gamma^2)\Sigma_{lim}(V) + \gamma^4\tau_0^{-2}[A^{\alpha\beta, \delta\omega}]_{\alpha \leq \beta, \delta \leq \omega}$, and

$$A^{\alpha\beta, \delta\omega} := \frac{1}{m^2} \sum_{\nu, \kappa=1}^m \left(V^{\alpha\kappa} V^{\delta\nu} S_1^{\beta\kappa, \omega\nu} + V^{\alpha\kappa} V^{\omega\nu} S_1^{\beta\kappa, \delta\nu} + V^{\beta\nu} V^{\delta\kappa} S_1^{\alpha\nu, \omega\kappa} + V^{\beta\nu} V^{\omega\kappa} S_1^{\alpha\nu, \delta\kappa} \right).$$

Proof. Recall that:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \frac{\lambda\gamma}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \mathcal{T}_2^{\alpha\beta}.$$

Drift. Summing and subtracting $\mathbb{E}_\ell[\mathcal{T}_2]$ (and using Lemma C.7), we have that:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \gamma^2 \sum_{\nu\kappa} V^{\nu\kappa} \mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa}] + \frac{\lambda\gamma}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \left(\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell[\mathcal{T}_2] \right),$$

where we recall $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ is the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$.

Re-stating Lemma C.3, we have that:

$$\mathbb{E}_\ell[A^{\alpha\nu} A^{\beta\kappa}] = \delta_{\alpha\nu} \delta_{\beta\kappa} + \frac{n_k}{\tau^2 m^2} S_1^{\alpha\nu, \beta\kappa} + \frac{n_k}{2\tau^2 m} (\delta_{\beta\kappa} S_2^{\alpha\nu} + \delta_{\alpha\nu} S_2^{\beta\kappa}) + \mathcal{O}(n_k \tau^{-3}).$$

Plugging in the expression, and using $\lambda^2 + \gamma^2 = 1$, we have that the drift is:

$$\begin{aligned}
 & \lambda^2 V_\ell^{\alpha\beta} + \gamma^2 V_\ell^{\alpha\beta} + \gamma^2 \frac{n_k}{\tau^2 m} \sum_{\nu\kappa} V^{\nu\kappa} \left(\frac{1}{m} S_1^{\alpha\nu, \beta\kappa} + \frac{1}{2} (\delta_{\beta\kappa} S_2^{\alpha\nu} + \delta_{\alpha\nu} S_2^{\beta\kappa}) \right) + \mathcal{O}(n_k \tau^{-3}) \\
 &= V_\ell^{\alpha\beta} + \gamma^2 \frac{n_k}{\tau^2} \left(\frac{1}{m^2} \sum_{\nu\kappa} V^{\nu\kappa} S_1^{\alpha\nu, \beta\kappa} + \frac{1}{2m} \sum_{\nu} (V^{\beta\nu} S_2^{\alpha\nu} + V^{\alpha\nu} S_2^{\beta\nu}) \right) + \mathcal{O}(n_k \tau^{-3}).
 \end{aligned}$$

From here, it is evident that in order to have the drift scaling as $\mathcal{O}(1/n)$ we need to choose:

$$\tau^2 = \tau_0^2 n n_k, \quad (26)$$

where $\tau_0 > 0$ is a constant.

Covariance. Recall that:

$$V_{\ell+1}^{\alpha\beta} = \lambda^2 V_\ell^{\alpha\beta} + \gamma^2 \sum_{\nu\kappa} V^{\nu\kappa} \mathbb{E}_\ell [A^{\alpha\nu} A^{\beta\kappa}] + \frac{\lambda\gamma}{\sqrt{n}} \mathcal{T}_1^{\alpha\beta} + \frac{\gamma^2}{\sqrt{n}} \left(\mathcal{T}_2^{\alpha\beta} - \mathbb{E}_\ell[\mathcal{T}_2] \right).$$

Furthermore, we have set $\lambda^2 + \gamma^2 = 1$ and $\tau^2 = \tau_0^2 n n_k$ to have the right scaling for the drift.

What's left is to compute the conditional covariance for $V_{\ell+1}$ given \mathcal{F}_ℓ . Noting that $\mathcal{T}_1^{\alpha\beta}, \mathcal{T}_2^{\delta\omega}$ are uncorrelated, i.e. $\mathbb{E}_\ell [\mathcal{T}_1^{\alpha\beta} \mathcal{T}_2^{\delta\omega}] = 0$ (similarly to the case of Resnet with shaped ReLU Lemma B.3), we have that:

$$\begin{aligned} \text{Cov}_\ell \left(V_{\ell+1}^{\alpha\beta}, V_{\ell+1}^{\delta\omega} \right) &= \mathbb{E}_\ell \left[\left(\lambda\gamma \mathcal{T}_1^{\alpha\beta} + \gamma^2 \mathcal{T}_2^{\alpha\beta} - \gamma^2 \mathbb{E}_\ell[\mathcal{T}_2^{\alpha\beta}] \right) \left(\lambda\gamma \mathcal{T}_1^{\delta\omega} + \gamma^2 \mathcal{T}_2^{\delta\omega} - \gamma^2 \mathbb{E}_\ell[\mathcal{T}_2^{\delta\omega}] \right) \right] \\ &= \lambda^2 \gamma^2 \mathbb{E}_\ell \left[\mathcal{T}_1^{\alpha\beta} \mathcal{T}_1^{\delta\omega} \right] + \gamma^4 \mathbb{E}_\ell \left[\mathcal{T}_2^{\alpha\beta} \mathcal{T}_2^{\delta\omega} \right] - \gamma^4 \mathbb{E}_\ell \left[\mathcal{T}_2^{\alpha\beta} \right] \mathbb{E}_\ell \left[\mathcal{T}_2^{\delta\omega} \right], \end{aligned}$$

where we use Cov_ℓ to denote the conditional covariance given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$.

Using Lemma C.6 and Lemma C.8, we have that:

$$\begin{aligned} \text{Cov}_\ell \left(V_{\ell+1}^{\alpha\beta}, V_{\ell+1}^{\delta\omega} \right) &= \lambda^2 \gamma^2 (2(V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta})) + \gamma^4 \left(\frac{1}{\tau_0^2} \mathcal{A}^{\alpha\beta\delta\omega} + V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta} \right) \\ &= \gamma^2 (2 - \gamma^2) (V^{\alpha\delta} V^{\beta\omega} + V^{\alpha\omega} V^{\beta\delta}) + \frac{\gamma^4}{\tau_0^2} \mathcal{A}^{\alpha\beta\delta\omega}. \end{aligned}$$

Now we can apply Proposition A.2 for locally Lipschitz drift and covariance coefficients, which gives us the desired result in local convergence in the Skorohod topology. \square

We will also restate and prove Corollary 4.3.

Corollary 4.3 (Shaped Transformer Covariance SDE). *Let X_ℓ be the hidden layers of a shaped transformer defined in Eq. 11 with parameters $\lambda^2 + \gamma^2 = 1$ and $\tau = \tau_0 \sqrt{n n_k}$, where λ, γ, τ_0 all do not depend on d, n . Then the feature covariance V_ℓ converges locally to the solution of the following SDE (in the sense of Definition 4.1)*

$$dV_t = [b(V_t) + b_{res}(V_t)] dt + [\Sigma(V_t) + \Sigma_{res}(V_t)]^{1/2} dB_t, \quad (12)$$

where the coefficients are defined in Theorem 3.2 and Theorem 4.2.

Proof. To combine the results of Theorem 3.2 and Theorem 4.2, it is sufficient to combine the following (simplified) iterated Markov updates into one Markov chain

$$U_\ell = V_\ell + \frac{b(V_\ell)}{n} + \frac{\Sigma(V_\ell)^{1/2} \xi_\ell}{\sqrt{n}}, \quad V_{\ell+1} = U_\ell + \frac{b_{\text{ReLU}}(U_\ell)}{n} + \frac{\Sigma_{\text{lin}}(U_\ell)^{1/2} \xi'_\ell}{\sqrt{n}}, \quad (27)$$

where ξ_ℓ, ξ'_ℓ are independent zero mean and identity covariance random vectors.

Since in the limit, we have that either updates are infinitesimal, i.e.

$$|U_\ell - V_\ell| \xrightarrow{n \rightarrow \infty} 0 \quad \text{almost surely}, \quad (28)$$

then we can write $b_{\text{ReLU}}(U_\ell) = \widehat{b}_{\text{ReLU}, n}(V_\ell, \omega_\ell)$, where eventually we have that

$$\lim_{n \rightarrow \infty} \mathbb{E}_\ell \left[\widehat{b}_n(V_\ell, \omega_\ell) \right] = b(V_\ell), \quad (29)$$

so it will not contribute towards affecting the limit. Here we recall $\mathbb{E}_\ell[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_\ell]$ is the conditional expectation given the sigma-algebra generated by $\mathcal{F}_\ell = \sigma(\{X_k\}_{k \in [\ell]})$. We can treat $\Sigma_{\text{lin}}(U_\ell)$ similarly to get the Markov chain update

$$V_{\ell+1} = V_\ell + \frac{b(V_\ell) + \widehat{b}_{\text{ReLU},n}(V_\ell, \omega_\ell)}{n} + \frac{\Sigma(V_\ell)^{1/2} \xi_\ell}{\sqrt{n}} + \frac{\widehat{\Sigma}_{\text{lin},n}(V_\ell)^{1/2} \xi'_\ell}{\sqrt{n}}, \quad (30)$$

which converges to the following SDE with two Brownian motions using Proposition A.2

$$dV_t = [b(V_t) + b_{\text{res}}(V_t)] dt + \Sigma(V_t)^{1/2} dB_t + \Sigma_{\text{res}}(V_t)^{1/2} dB'_t. \quad (31)$$

Observe that since the two Brownian motions B_t, B'_t are independent, it's equivalent to write

$$\Sigma(V_t)^{1/2} dB_t + \Sigma_{\text{res}}(V_t)^{1/2} dB'_t \stackrel{d}{=} [\Sigma(V_t) + \Sigma_{\text{res}}(V_t)]^{1/2} dB_t. \quad (32)$$

We recover the desired results from considering a more general form of the iterated Markov updates as in Proposition A.2, which do not hinder the above derivation. □

D Preliminary Experiments

We perform preliminary experiments to understand the effect of *shaped attention* on training deep Transformer architectures. In particular, we consider a pre-training masked language modeling task, where we mask 15% of the tokens. We use a subset of the English Wikipedia 20220301.en and English *bookcorpus* datasets [69, 70]. As a baseline, we adopt a Pre-LN Transformer encoder architecture with 18 or 24 blocks. For the residual feedforward layer, we shape the ReLU activation according to Eq. 4, by changing its negative slope to $s_- = 1 - 1/\sqrt{n}$ instead of 0. We then incorporate our shaped attention by replacing the attention mechanism as dictated in Eq. 9. We also add scalar multipliers $\gamma_1, \gamma_2 \in \mathbb{R}$ both to the identity and centering terms of the shaped attention:

$$A_\ell = \gamma_1 I + \text{Softmax}(\tau^{-1} Y_\ell) - \gamma_2 \frac{1}{m} \mathbf{1}\mathbf{1}^\top, \quad \tau = \tau_0 \sqrt{nn_k}, \quad (33)$$

and propose two ways to set γ_1, γ_2 during training. In both alternatives we initialize $\gamma_1, \gamma_2 = 1$, thus leveraging the stability properties of shaped attention at initialization. During training, we either:

1. **Recover.** Linearly decrease γ_1, γ_2 to zero with in the first 4000 steps, thus recovering the standard attention layer. Apply the same schedule for the shaped-ReLU slope s_- , recovering the usual ReLU activation. This approach recovers the vanilla Transformer architecture (without LayerNorm).
2. **Learn.** Learn all the shaping parameters γ_1, γ_2 and s_- .

The intuitive rationale behind these choices is that at initialization we want good signal propagation and a non-degenerate covariance (according to our theory, this requires the shaped attention and ReLU). On the other hand, we also allow the model to more dynamically make use of the nonlinearity during training to modify the correlation structure with **Recover** or **Learn**. We report that without either adjustment, *shaped attention* is still trainable but at much slower rates.

We also incorporate the $\frac{1}{\sqrt{n}}$ factor into the initialization of the queries and keys weights by decreasing the variance of their entries by a factor of $\frac{1}{n}$. This allows us to not re-tune the learning rate for the queries and keys. We stress that at initialization, the two formulations ($\tau = \sqrt{nn_k}$ and $\tau = \sqrt{n_k}$ with decreased weight's variance) are equivalent. In both alternatives we train all the skip connection parameters λ, γ , and initialize γ in the grid (0.05, 0.1, 0.2) and set $\tau_0 = 1$. We report training instabilities (loss divergence) for larger values of γ . All models—including the baselines—use Adam [71] with learning rate warmup of 4000 steps, the learning rate is tuned in the grid (0.0001, 0.0005, 0.001, 0.005). We report the train/test loss after 100K optimization steps, averaged over 4 random seeds. All the other experimental details can be found in Appendix D.1.

The Shaped Transformer is Trainable. In Table 1, we compare the train/test loss of the two variant of shaped attention with the baseline Pre-LN model. Notice that our model (in both variants) achieves comparable performance to standard Transformers across all the reported values of γ .

GLUE evaluation Furthermore, we evaluate the trained models on three datasets from the GLUE benchmark [72] and summarize the results in Table 2. These demonstrate that our shaped Transformers holds promise by outperforming our pre-LN baseline.

| | $d = 18$ | | | $d = 24$ | |
|----------|----------|------------------------|------------------------|------------------------|-----------------|
| | γ | Train Loss | Test Loss | Train Loss | Test Loss |
| Learn | 0.05 | 2.09 \pm 0.02 | 2.07 \pm 0.02 | 2.03 \pm 0.02 | 2.03 \pm 0.01 |
| | 0.1 | 2.10 \pm 0.02 | 2.07 \pm 0.01 | 2.02 \pm 0.01 | 2.02 \pm 0.03 |
| | 0.2 | 2.07 \pm 0.05 | 2.06 \pm 0.02 | 2.09 \pm 0.04 | 2.09 \pm 0.03 |
| Recover | 0.05 | 2.05 \pm 0.03 | 2.04 \pm 0.02 | 2.02 \pm 0.04 | 2.00 \pm 0.01 |
| | 0.1 | 2.05 \pm 0.01 | 2.04 \pm 0.01 | 2.06 \pm 0.03 | 2.00 \pm 0.02 |
| | 0.2 | 2.01 \pm 0.03 | 2.03 \pm 0.01 | 2.05 \pm 0.04 | 2.00 \pm 0.01 |
| Baseline | | 2.08 \pm 0.03 | 2.07 \pm 0.01 | 2.08 \pm 0.03 | 2.01 \pm 0.01 |

Table 1: Training and test loss of the proposed *shaped attention*, both in the "Recover" and "Learn" alternatives. Baseline refers to the vanilla Pre-LN Transformer. We report the best run under the learning rates (0.0001, 0.0005, 0.001, 0.005), averaged over 4 random seeds. We include the confidence interval of \pm one standard deviation.

| | Model | COLA | MRPC | RTE |
|----------|--------------------------|-------------------|-------------------|-------------------|
| $d = 18$ | Baseline | 0.139 \pm 0.008 | 0.797 \pm 0.010 | 0.504 \pm 0.024 |
| | Learn, $\gamma = 0.2$ | 0.182 \pm 0.041 | 0.813 \pm 0.005 | 0.528 \pm 0.019 |
| | Recover, $\gamma = 0.2$ | 0.221 \pm 0.042 | 0.809 \pm 0.009 | 0.520 \pm 0.029 |
| $d = 24$ | Baseline | 0.022 \pm 0.055 | 0.785 \pm 0.009 | 0.48 \pm 0.046 |
| | Learn, $\gamma = 0.05$ | 0.150 \pm 0.026 | 0.812 \pm 0.004 | 0.513 \pm 0.019 |
| | Recover, $\gamma = 0.05$ | 0.211 \pm 0.039 | 0.803 \pm 0.012 | 0.529 \pm 0.019 |
| | BERT (Geiping et al.) | 0.103 | 74.8 | 0.509 |

Table 2: Evaluation of the pretrained models on a subset of the GLUE benchmark. We take the checkpoints of the baselines and selected shaped Transformers ($\gamma = 0.2$ for the shallower model $d = 18$, $\gamma = 0.05$ for the deeper model $d = 24$). Then, we fine-tune on each of three datasets from GLUE (COLA, MRPC, RTE) for 5 epochs and Adam optimizer $lr = 5e - 4$ with batch size 16, and report test evaluation metric for the corresponding task.

Entropy Collapse for Large Learning rates. To understand the sources of training instability, we keep track of the entropy of the probability distribution induced by the Softmax, as it has been observed that the Transformer’s training is unstable in the low-entropy regime [57]. The entropy is calculated for each row of the Softmax matrix, and it is averaged across rows and heads. The results are in Fig. 6. Notice how for the large learning rate regime observed in Fig. ??, the entropy collapses for the baseline model, but not for the proposed *shaped Transformer*. Entropy collapse indicates that the Softmax distribution degenerates to a point mass, which is itself caused by large logits. Remarkably, this phenomenon does not affect the recover setting, despite recovering the Transformer architecture (without layer normalization) after the warm-up period.

D.1 Experimental details

Dataset. We use a subset of the English Wikipedia 20220301.en and English *bookcorpus* datasets [69, 70]. The sentences are tokenized using by pre-training a tokenizer on the training set. We use a vocabulary size of 32000, and a maximum sequence length of 128 tokens.

Model parameters. We use an embedding size of $n = 768$ and 8 multi-attention heads. The batch size is fixed to 32 sequences. All the initial weights are sampled from $\mathcal{N}(0, n^{-1})$, with the exception of the queries and keys’ weights W^K, W^Q in the *shaped attention* case, that are sampled from $\mathcal{N}(0, n^{-3/2})$ (as explained in Appendix D). The feedforward layer maps the $n = 768$ -dimensional embedding to the larger dimension 3072, as in the Hugging face implementation of Bert [73].

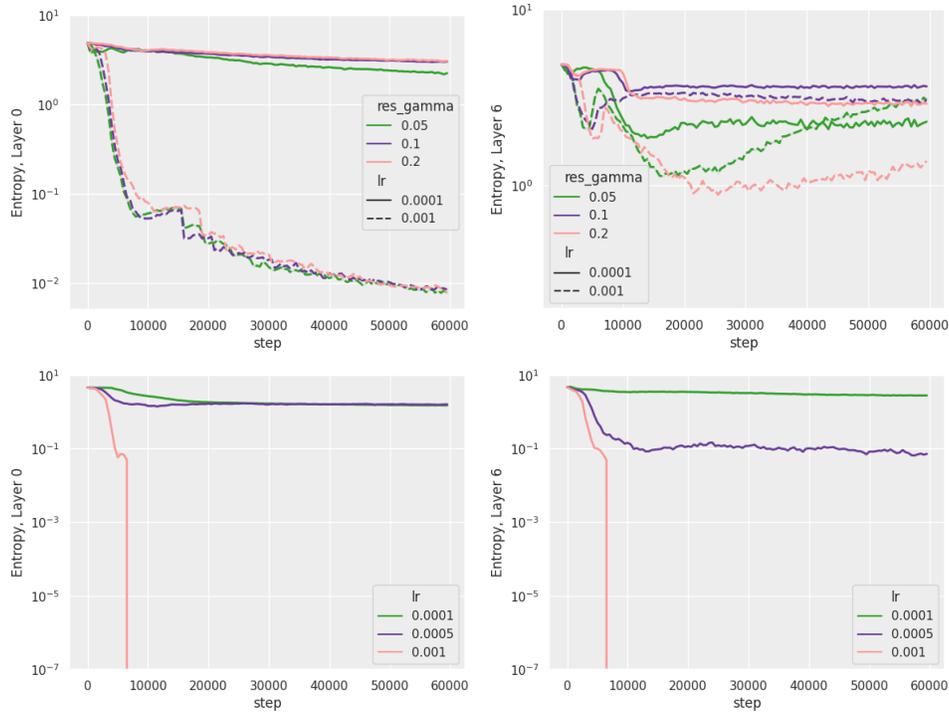


Figure 6: Dynamics of the mean entropy across heads for selected layers (first and seventh) and learning rates. (Above): *shaped attention* in *recover* setting. (Below): Baseline transformers. Notice that at $lr = 0.001$ the entropy collapses for the baseline model.

Optimization. We train using Adam [71] with betas parameters (0.9, 0.999) and learning rate chosen in the grid (0.0001, 0.0005, 0.001, 0.005). We do not use weight decay.

Computational Resources. The experiments are executed on Nvidia DGX-1 GPU nodes equipped with 4 20-core Xeon E5-2698v4 processors, 512 GB of memory and 8 Nvidia V100 GPUs.

E Additional Figures

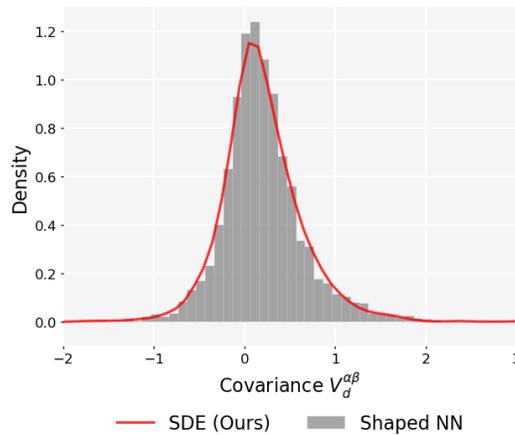


Figure 7: Kernel density estimate and histogram of covariances from the covariance SDE in Theorem 4.2 and shaped attention NN (Eq. 9). Simulated with $n = 200$, $d = 150$, $\gamma = 1/\sqrt{8}$, $\tau_0 = 1$, $V_0^{\alpha\beta} = 0.2$, SDE step size 0.01, and 2^{12} samples.

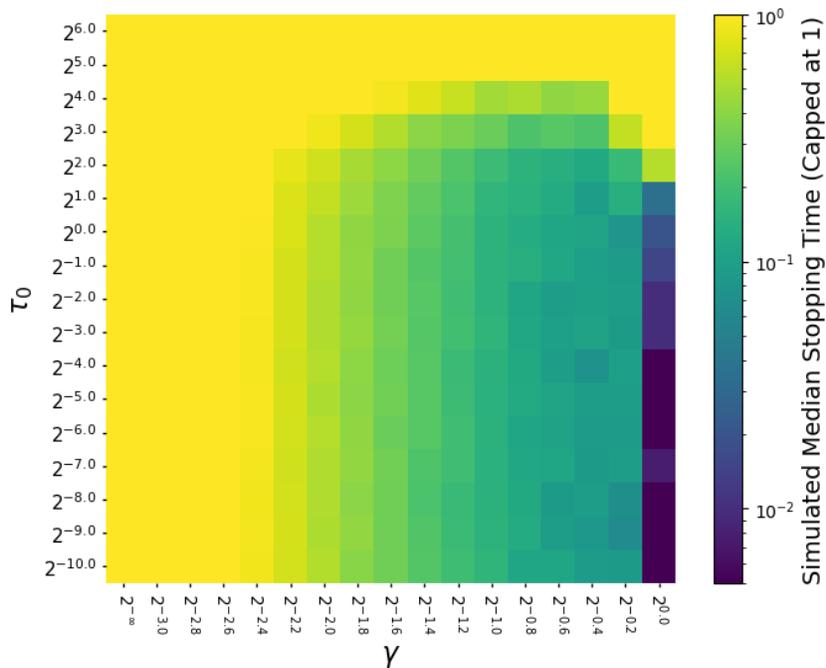


Figure 8: Median of the stopping time capped at 1, of the shaped attention neural network with respect to its parameters γ and τ_0 . Stopping time is defined as $t^* = d^*/n$ with d^* the maximum depth beyond which one of the eigenvalues of the covariance matrix exceeds 10^4 or drops below 10^{-4} . Simulated with $n = d = 200$, and 100 samples used to estimate the median. To demonstrate the potential numerical instabilities, we had to choose an *adversarial* set of parameters: in particular, an unrealistically large norm (approx. $10\sqrt{n}$) for the initial tokens X_0 , which enlarges the eigenvalues of V_0 to the order of 100.