RINGING RELUS: HARMONIC DISTORTION ANALYSIS OF NONLINEAR FEEDFORWARD NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

In this paper, we apply harmonic distortion analysis to understand the effect of nonlinearities in the spectral domain. Each nonlinear layer creates higher-frequency harmonics, which we call "blueshift", whose magnitude increases with network depth, thereby increasing the "roughness" of the output landscape. Unlike differential models (such as vanishing gradients, sharpness), this provides a global view of how network architectures behave across larger areas of their parameter domain. For example, the model predicts that residual connections are able to counter the effect by dampening corresponding higher frequency modes. We empirically verify the connection between blueshift and architectural choices and provide evidence for a connection with trainability.

1 INTRODUCTION

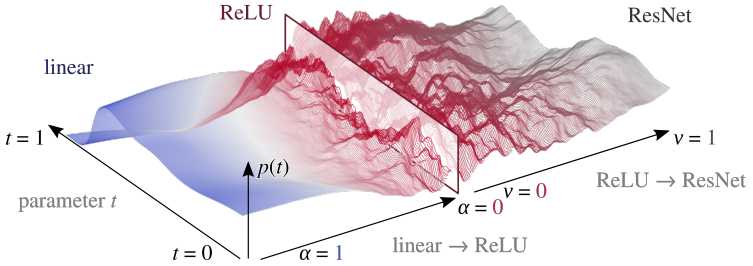

Figure 1: Transition of a loss path between linear activations, ReLU activations, and ResNet with shortcuts. Graph: loss surface of a ResNet56 v2 with LReLUs with negative slope $\alpha \in [0, 1]$ and residual factor $\nu \in [0, 1]$. Left: $\alpha = 0, \nu = 0$, Middle: $\alpha = 1, \nu = 0$, Right: $\alpha = 1, \nu = 1$.

In the past decade, the emergence of practical deep neural networks arguably has had disruptive impact on applications of machine learning. Depth as such appears to be key to expressive models (Raghu et al., 2017). However, depth also comes with challenges concerning training stability. Theoretical problems include vanishing and exploding gradients (Hochreiter, 1991), chaotic feed-forward dynamics (Poole et al., 2016), or decorrelation of gradients (Balduzzi et al., 2017). In practice, a number of "recipes" are widely used, such as specific nonlinearities (Glorot et al., 2011; He et al., 2015), normalization methods such as batch normalization (Ioffe & Szegedy, 2015), shortcut architectures (Srivastava et al., 2015; He et al., 2016a;b), or multi-path architecture with (Huang et al., 2017) and without shortcuts (Szegedy et al., 2016). Broadly speaking, a key research question is to understand how the shape of the network function, i.e., the map from inputs and parameters to outputs, is affected by architectural choices.

Our paper considers specifically the *roughness* of the *weights-to-outputs function* ("*w-o function*") of non-linear feed-forward networks. Motivated by the recent visualizations of (Li et al., 2018), which show how depth increases roughness and residual connections smoothen the output again, our goal is to provide an analytical explanation of this effect, and study its implications on network design and trainability. To this end, we first formalize "roughness" as the decay-rate of the expected power spectrum of a function class. Our main contribution is to then apply *harmonic distortion analysis* to non-linear feedforward networks, which predicts the creation of high-frequency "harmonics" (thereby "blueshifting" the power spectrum) by polynomial non-linearities with large higher-order coefficients.

Based on this model, we discuss how network depth increases blueshift and thus roughness, while shortcut connections, low-degree nonlinearities and parallel computation paths dampen it. In relation to trainability, we show an analytic link between blueshift and vanishing/exploding gradients. Unlike the former model, the spectral view describes a more global behavior of the w-o function over regions in the parameter domain. Experiments confirm the theoretical predictions: We observe the predicted effects of depth, shortcuts and parallel computation on blueshift, and are able to differentiate different types of non-linearities by the decay rate of coefficients of a polynomial approximation. The findings are in-line with known advantages in trainability of the different architectures. We further strengthen the evidence by training a large set of networks with a different amount of non-linearity and depth, which shows a clear correlation between blueshift and training-problems.

In summary, our paper explains how network architecture affects roughness, shows a connection to trainability, and thereby provides a new tool for analyzing the design of deep networks.

2 Related Work

Vanishing or exploding gradients are a central numerical problem (Hochreiter, 1991; Pascanu et al., 2013; Yang et al., 2019): If the the magnitudes of the singular values of layer Jacobians deviates from one, subspaces are attenuated ($|\sigma| < 1$) or amplified ($|\sigma| > 1$), potentially cascading exponentially over multiple layers (Pennington et al., 2017a). Formally, the behavior of stacks of matrices and nonlinear functions can be modeled by random matrix theory or Gaussian mean-field approximations (Poole et al., 2016; Pennington et al., 2017b; 2018). The gist is that at initialization, orthogonal weight matrices are needed, which is challenging for convolutional architectures. A solution for $\tanh$-networks is given by Xiao et al. (2018); for ReLU, there is a negative result (Pennington et al., 2017b). Using mean-field theory, it can be shown that batch normalization (Ioffe & Szegedy, 2015) leads to exploding gradients at initialization (Yang et al., 2019) (which equalize after a few steps, but that might be too late (Frankle et al., 2020)).

A different route is taken by Balduzzi et al. (2017), who observe an increasing decorrelation of gradients in the input space. Similar to our paper, they show that deeper networks lead to spectral whitening (starting from brown noise); however, the analysis is performed with respect to the inputs $\mathbf{x}$, not weights $\mathbf{W}$. The scale-space structure shown by our model might give further hints on the mechanisms behind training difficulties.

Our work explains the observations of Li et al. (2018) by spectral analysis. Fourier analysis of network functions (Candès, 1999) wrt. input (Rahaman et al., 2019; Xu et al., 2019; Xu, 2018; Basri et al., 2019; Yang & Salman, 2019) has been used to show an inductive bias towards low-frequency functions (wrt. input $\mathbf{x}$), as well as a strong anisotropy of this spectrum. Wang et al. (2020) prove under some assumptions that all "bad" local minima of a deep residual network are very shallow.

3 Harmonic Distortion

We consider feedforward networks f consisting of L-layers. We use $\mathbf{x}$ to denote the input to f, $\mathbf{W}$ to denote the weights and ϕ to denote the employed nonlinearity. Li et al. (2018) visualize "roughness" using random 2D-slices in weight space. We follow their basic idea and consider 1D slices

$$p(t) = f(\mathbf{x}, \mathbf{W} + \alpha^{-1} \cdot t \cdot \mathbf{D}) \tag{1}$$

for random directions $\mathbf{D}$ and $t \in [0, 1]$. $\mathbf{D}$ is initialized with entries from $\mathcal{N}_{0,1}$ and normalized to $||\mathbf{D}||_F = 1$. α determines the path length. By varying $\mathbf{D}$, this samples a ball of radius α around a point $\mathbf{W}$ in parameter space. For $\phi = id$, the network f is multi-linear in $\mathbf{W}$ and thus p is polynomial in t; empirically, this yields rather smooth functions (Fig. 1). To understand general non-linearities ϕ better, we represent p by a complex Fourier series (Appendix B.1 shows its convergence):

$$p(t) = \sum_{k=-\infty}^{\infty} z_k \exp\left(2\pi i k t\right), \quad z_k \in \mathbb{C}. \tag{2}$$

Here, the sequence $z : \mathbb{Z} \to \mathbb{C}$ contains the Fourier coefficients for $p : [0, 1] \to \mathbb{R}$. As p is real, z is symmetric in the sense of $z_k = \overline{z_{-k}}$.

3.1 FORMALIZING ROUGHNESS

The roughness of a class of random functions can be characterized by the statistics of their power-spectrum (Musgrave, 1993): Given a random process that generates functions p, we consider mean μ_k and variance σ_k^2 of the Fourier coefficients z_k. Empirically, the mean paths p are zero-functions, which lets us focus on variance: In general, functions where the variance of high-frequency components drop off more quickly will appear smoother. A common model, which often fits natural data well, is fractal Brownian motion (FBM-) noise, where the σ_k drop off according to a power law:

$$\sigma_k \sim \mathcal{O}\left(1/k^h\right) \text{ for some } h > 0. \tag{3}$$

The so-called "fractal coefficient" h describes the roughness of the noise function. In our experiments, we estimate the average power-spectrum $\mathbb{E}_{\mathbf{D}}(|z_k|^2)$ (by sampling $\mathbf{D}$ uniformely on a unit sphere) and fit a power-law to these spectra in order to quantify the roughness in a single number h.

Experiments (Section 4) and analytical arguments (Appendix B.2), show that the FBM/power-law model is a realistic model of the functions computed by the lower layers of a neural network. For higher layers the fit becomes worse, a phenomenon we will explore in the next chapter using harmonic distortion analysis.

3.2 WHY IS THE OUTPUT FUNCTION GETTING ROUGHER?

Intuitively, applying ReLU to function p is reminiscent of clipping an audio signal in amplitude, which is known to produce high-frequent ringing artifacts. Generally, we consider the application of a non-linearity ϕ to a function p. For the analysis, we assume that ϕ is a K-th order polynomial:

$$\phi(x) = \sum_{j=0}^{K} t_j x^j. \tag{4}$$

The effect of polynomial non-linear maps on the spectrum of a function can be understood by harmonic distortion analysis (see e.g. Feynman et al. (1965), Ch. 50.8): We can simply (see Appendix B.1 for details) plug the Fourier expansion of p into the polynomial representation of ϕ:

$$\phi(p(t)) = \sum_{j=0}^{K} t_j \left[\sum_{k=-\infty}^{\infty} z_k \exp\left(2\pi i k t\right) \right]^j \tag{5}$$

The convolution theorem tells us that j-th power of functions corresponds to convolving the spectrum z of the function j-times with itself. This yields a spectrum z' of the output of:

$$z' := \mathcal{F}(\phi(p)) = \sum_{j=0}^{K} t_j \bigotimes_{1=1}^{j} z, \tag{6}$$

Discussion: We can make three important observations: First, each repeated auto-convolutions in Eq. 6 broadens the spectrum by adding higher-frequency terms. The exact magnitude is hard to quantify (see Appendix B.4 for a discussion) and thus left to the experiments. Second and correspondingly, larger coefficients t_j for larger orders j, increase the blueshift. Third, j-fold convolutions correspond to j-th powers of the z_k. Hence, larger magnitudes $|z_k|$ also increase the blueshift (the non-linearity becomes more visible with larger magnitude).

3.3 COMMON NONLINEARITIES

In practice, non-linearities are usually not polynomial (and might not even have a globally convergent Taylor expansion). We therfore resort to Chebyshef approximations (which reduces oscillatory artifacts in polynomial approximations) on a finite interval and validate experimentally that such an approximation still makes qualitatively correct predictions. In this paper, we consider $\mathrm{ReLU}, \tanh$ and $\mathrm{softplus}$ The polynomials approximation of $\mathrm{softplus}$ has the strongest drop-off of higher-order coefficients, followed by $\tanh$ and ReLU. Our experiments show that the magnitude of the blueshift follows this pattern (for all of the examined non-linearities). A visualization of the magnitudes $|t_j|$ for an approximation on $[-5, 5]$ is shown in Appendix D.6 (Figure 18).

3.4 CONNECTION TO EXPLODING GRADIENTS

The creation of higher-frequency harmonics directly affects gradient-based training methods because the gradient operator is a high-pass filter in the spectral domain: Taking derivatives multiplies the Fourier coefficients by $2\pi ik$, amplifying coefficients linearly with frequency k (which is trivial seen by taking the derivatives of Eq. 2). The reciprocal fractal exponent $r := h^{-1}$ has exponential influence on the average gradient magnitude: With $|z_k|^2 \sim \Theta(k^{-h}) = \Theta(k^r)$, we obtain gradients $k \cdot |z_k| \in \Theta(k^{1+r/2})$. Correspondingly, exploding gradients in deep feedforward networks can be explained by repeated blueshifting.

The blueshift effect is independent of the condition number of the weight matrices: blueshift would occur even when only stacking nonlinearities. Although it is possible to find a locally linear region of the loss surface (e.g. CReLU + Looks Linear initialization (Balduzzi et al., 2017)) and locally achieve dynamic isometry (Xiao et al., 2018), blueshift effects become relevant again when this region is eventually left during the training process (ref. appendix Figure 13). In contrast, residual connections affect the entire loss surface and our model is able to capture this effect.

3.5 RESIDUAL AND MULTIPATH NETWORKS

Residual Networks exhibit *two* different dampening effects on the loss surface that can be experimentally isolated:

Exponential downweighting (ED): Residual networks repeatedly compute the sum of the result of a small feedforward network with its input. By construction, this reduces the relative weight of the non-linear output of the residual block by mixing it with the unprocessed signal. Doing this repeatedly creates an ensemble (Huang et al., 2016) in which the weight of contributions shrink exponentially with the number of non-linear processing steps.

Frequency-dependent signal-averaging (FDSA): According to the blueshift model, the outputs of the residual blocks will contain more high-frequency Fourier coefficients (from harmonics). The additional harmonics depend on the weights in the residual block, which are statistically independent of the input (see Appendix B.5 for details). Therefore, we expect to see a reduction in the expected linear correlation between the main and residual branch for higher frequencies. As an effect, when adding the main and residual branch, higher frequencies get dampened even more.

The second effect is weaker than the first (law of large numbers vs. exponential decay), but more broadly applicable: It should generally occur when averaging over multiple computational paths with independent weights, while residual connections require an identity in one path.

4 EXPERIMENTAL RESULTS: MEASURED SPECTRA

The model presented in the previous section gives us only qualitative hints on the magnitude of the blue-shift, the behavior of non-polynomial non-linearities, and the magnitude of FDSA-effects. We therefore now validate these qualitative predictions experimentally. We consider two architectures: a basic Toy-CNN model with a constant number of features in each layer for fair per-layer statistics and ResNet v1 variants for a more "realistic" model. We used Cifar10 as input data, but all figures in this section look qualitatively similar on other datasets; A more detailed network description, parameters and experiments on MNIST can be found in the appendix. In this section, we always measure the effect of blueshift in the region around initialization (i.e., $\mathbf{W}$ is the initialization point); For a visualization during training, see the Appendix (Figure 10).

4.1 SPECTRAL SHIFT

Our first goal is to validate the harmonics creation and in particular measure the blueshift caused by multiple layers, i.e., the effect of deeper networks on the expected power spectrum. We also validate the connection between the blueshift strength and the presence of larger higher-order polynomial coefficients for the non-linearity.

We use a “Toy-CNN” with depth $L = 50$ and track the w-o surfaces of each layer while we walk in a single random direction in weight space. We transform the resulting 1D-paths with discrete FFT

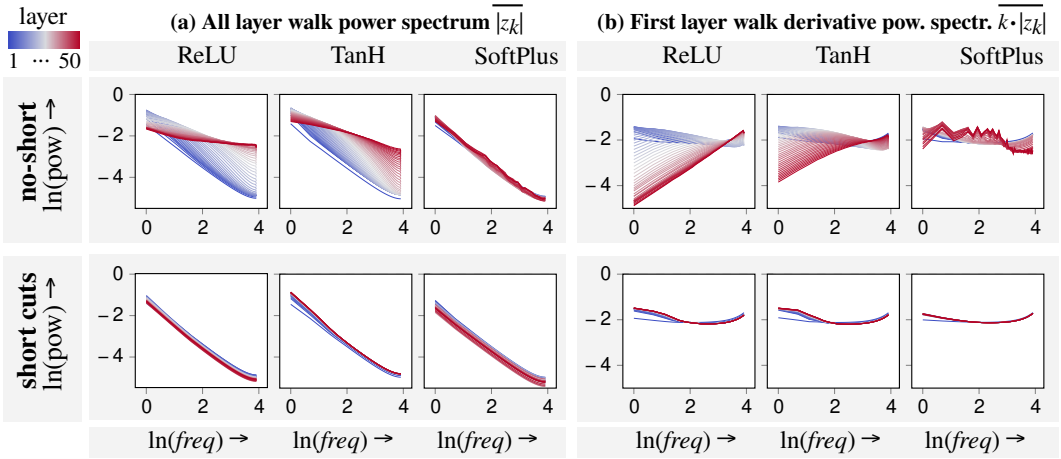

Figure 2: Normalized power spectra z_k per layer in a 50-layer CNN at initialization, averaged over multiple neurons, for various ϕ. (a) On the left, all weights are varied and spectrum of $p(t)$ is shown. (b) On the right, only the first layer is varied and the spectrum of $p'(t)$ is shown. Color indicates layer depth (red is deeper).

and all magnitude spectra are averaged over all neurons in a layer, all input receptive fields and all batch images and then normalized by function norm. Fig. 2a (left side) shows the resulting average magnitude spectrum $\overline{|z_k|}$ for each Fourier coefficients. The right hand side (b) shows the same plot, but with $\mathbf{D}$ restricted to only changing the weights of the first layer of the network, and taking the derivative of the resulting function (by weighting by frequency, $k \cdot \overline{|z_k|}$) before normalization.

Spectral shift over depth: The results in Fig. 2a confirm our predictions quite well: A spectral blueshift is clearly visible and, as expected, ReLU shows a strong blueshift, $\tanh$ a slightly weaker blueshift, and $\mathrm{softplus}$ only a small effect. Harmonics show up as a bump in the spectrum that travels towards higher frequencies with increasing layer. This extends to other non-linearties as well (see Appendix, Fig. 16). The plot also confirms that the power spectrum of the lower layers is well described by a $1/k$-power law, as predicted. A slight upward slant is visible at the high-frequency end of all plots. This is an aliasing phenomenon: As we use only 100 samples without prior bandlimiting, aliasing frequencies are reflected back at the end of the spectrum, bending the curve slightly upward.

Scale shift of a single layer parametrization: By only changing the weights in a single layer and taking the derivative (Fig. 2b), we see that blueshift is responsible for a gradient scale mismatch between earlier and deeper layers in networks without skip-connections: the more nonlinearities are between the modified weight and the output, the higher the gradient magnitude.

Residual connections: Our model predicts that residual connections will reduce the blueshift strongly by **ED** and **SDSA** (ref. Section 3.5). This is also consistent with our observations: higher layers show almost the same $1/k$-spectrum as the initial layer. Looking at the first layer weights, weighted by frequency, the response is almost flat, with a small emphasis on low-frequency weights, and again a bump at the upper end of the spectrum presumably due to aliasing (which is more pronounced for frequency weighting).

4.2 Effect of network depth on network smoothness

The spectral shift in the w-o function causes it to become rougher with network depth. We measure this phenomenon directly by fitting a power law to the average power spectra $\overline{|z_k|^2}$ using non-linear least-squares: $h, b = \arg\min_{h,b \in \mathbb{R}} \left(\frac{b}{t^h} - \overline{|z_k|^2} \right)^2, \quad t \in [0, 1]$.

For averaging, we initialize the network 50 times and sample 20 paths in different random directions. Fig. 3 shows the results for Toy-CNN (every data point corresponds to a different batch, excluding bias by batch selection). The results confirm quantitatively the visual results of Li et al. (2018) that the loss surface becomes rougher with increasing network depth and that shortcut connections have a

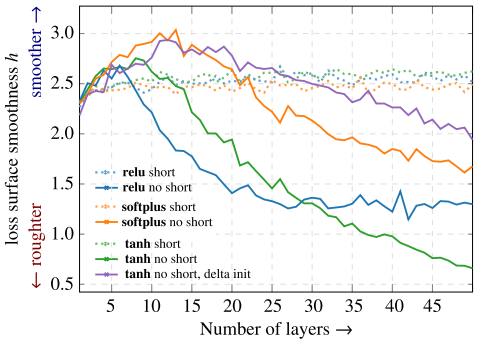

Figure 3: Smoothness of the loss surface at initialization for a Toy-CNN of varying depth with and without shortcuts.

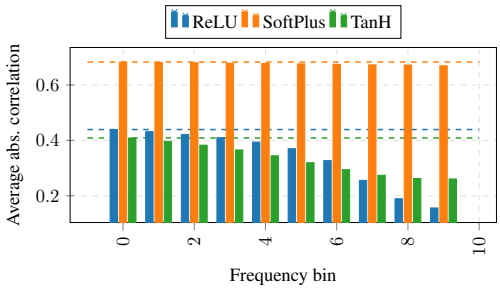

Figure 4: Average absolute path correlation per frequency bin of the main branch and the residual branch of a ResNet56 v2 with different activation functions at initialization. Lower bin number indicates lower frequencies.

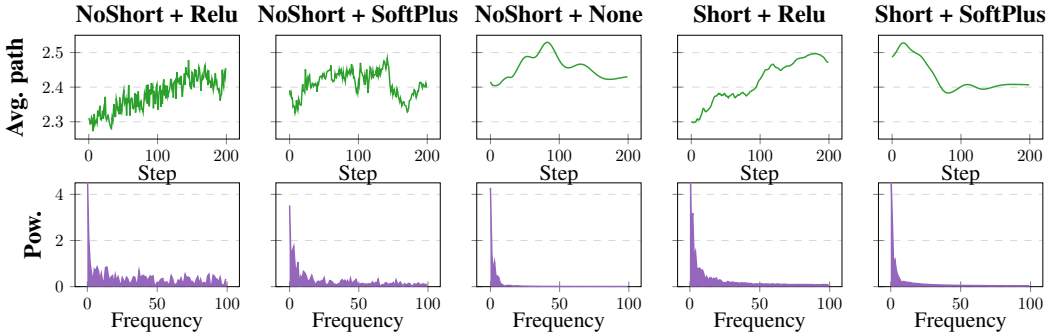

Figure 5: "Average loss path" and its respective magnitude spectrum out of 500 random loss paths for different variants of ResNet56 at initialization.

strong smoothing effect. The power-law fit still has some limitations: As seen in Fig. 2, a power-law fit is imperfect as the higher layers are not straight lines anymore. This shows in curves of Fig. 3 as apparent increase of smoothness in the first layers for the "NoShort" architectures. In Figure 5, we use the power-law fit to the path spectrum to select and visualize the loss function for a path with median roughness value: This way we can validate that the visual results of Li et al. (2018) are consistent over many random directions. We clearly see that nonlinear architectures without skip-connections present most blueshift, the linear architecture has no blueshift and the architectures with skip-connections do also exhibit, albeit dampened, blueshift.

4.3 Residual networks control harmonics / averaging networks

In Section 3.5, we described two different mechanisms that allow ResNets to dampen high-frequency modes. We already saw that residual connections have a dampening effect on the loss surface, but we want to individuate both effects.

SDFA predicts that skip-connections have a dampening effect on high frequencies in the loss surface because the correlation between the main branch and the residual branch is decreasing with higher frequencies. We empirically verify this prediction on a ResNet56 v2 at initialization by walking in 50 random directions and measuring the respective w-o function (we only sample 10% of the outputs for computational reasons) of the two summands in each residual block just before addition. Using Gaussian filters in Fourier domain, we can compute the linear correlation of each pair of paths (block input/output) per frequency bin and average over all batch images, neurons and blocks in the network. In Figure 4, we see that the absolute linear correlation is decaying for all architectures with increasing frequency, with ReLU showing the strongest effect, again followed by tanh and softplus in that order, confirming our prediction. This property seems to allow a residual network to automatically regulate the higher-frequency harmonics content of the output.

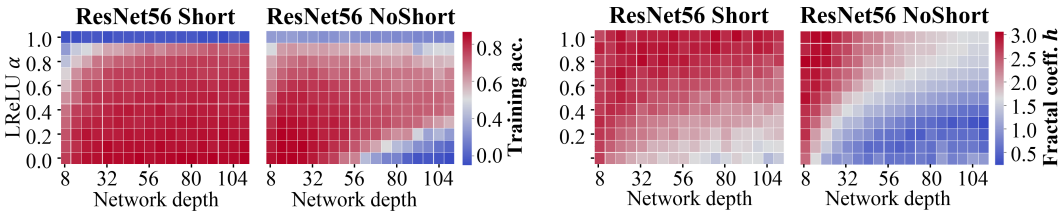

Figure 6: Network trainability correlates with blueshift on a ResNet. The first left side shows training accuracy after 30 epochs of training on Cifar10. The right side shows the blueshift indicated by the fractal coefficient of the network's last layer at initialization.

We want to individuate **SDFA** in a setting where **ED** is not present. We modify our Toy-CNN and ResNet "NoShort" architectures such that we replace every n-feature *convolution / batch normalization / activation* sequence in the network with a $a \cdot n$ feature sequence where the a feature groups get averaged after the activation (in reality we only use summation, since normalizing is taken care of by batch normalization). Since high frequencies added by the nonlinearity are less correlated between features (ref. Fig. 4), we expect that averaging feature groups smooth the w-o function; we see on the right of Figure 8 that this is indeed the case.

Finally, we check that even in a setting with **ED**, **SDFA** is still present. Our model predicts that since ResNet v2 blocks (He et al., 2016b) use summation after ReLU, we will see a stronger averaging effect than in ResNet v1. Indeed, looking at their respective average power spectra in a random direction in Figure 7, we see this confirmed.

4.4 Transitioning Regimes and PreLU / Leaky ReLU

On the left plot of Figure 1 we sampled a loss path from a ResNet56 architecture with shortcuts removed and Leaky ReLU activations. By tuning the slope parameter $\alpha \in [0, 1]$ of Leaky ReLU, we are able to observe the continuous transition of this loss path from a linear to a nonlinear regime, with correspondingly increasing roughness due to the creation of harmonics. By progressively adding in shortcuts again on the right plot (residual branch factor $\nu \in [0, 1]$), we can observe how the loss-surface becomes smooth again.

5 Empirical Results: Training

In this section, we want to demonstrate a correlation between blueshift and network trainability. Again, we use variants of the ResNet v1 network, exact network parameters and training hyperparameters can be found in tabular form in Appendix A. We use Cifar10 as training data, training results for Cifar100 can be found in Appendix C.

Controlling Blueshift with Leaky ReLU: As demonstrated in Figure 1, blueshift can be controlled by tuning Leaky ReLU's negative slope α. We want to show the striking correlation that exists between blueshift and trainability. On Figure 6, we see that for the networks without skip-connections, training becomes more difficult with regard to network depth. We see that this effect can be alleviated by making the network more linear and thus reducing blueshift. For networks with skip-connections, we see that blueshift is greatly reduced and that no deterioration of trainability with regard to network depth is noticeable.

Controlling Blueshift via feature averaging: We now want to investigate the impact of the smoothing of the loss surface via **SDSA** (ref. Section 3.5) on training speeds. We choose a ResNet56 "NoShort" as our architecture, since at this depth simple feedforward networks with batch-normalization start to become difficult to train. We experiment on the averaging network with $a = 3$ since we experimentally determined that for $a > 3$ the performance stops increasing. For fairness, we included a Wide ResNet56 "NoShort" with approximately the same parameter count than the averaging network. We see the averaged results over five runs in Figure 8 (shaded areas represent the standard deviations). We observe that the averaged network outperforms both the original network

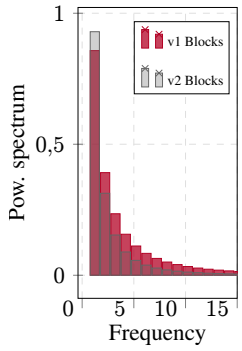

Figure 7: Normalized power spectrum of the last layer of a ResNet194 v1/v2 (pre/post-activation) at initialization.

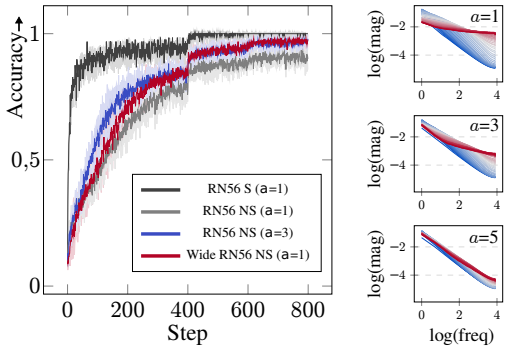

Figure 8: Left: Comparing the training performance on Cifar10 of a feature-averaging ResNet56 NoShort to its vanilla and residual equivalent Right: Spectrum shift of averaging Toy-CNNs with 50 layers at initialization with varying averaging factor a.

and the wide network while still performing worse than the network with skip-connections. This is consistent with our theory since **ED** has a bigger smoothing effect than **SDSA**.

Discussion: Our experimental findings validate the predictions of our model with respect to the roughness of the w-o function: Nonlinearities with larger higher-order polynomial coeffiecnts create larger variances in the high-frequency part of the spectrum. While related results on increasing complexity with depth have been given earlier (Poole et al., 2016; Schoenholz et al., 2017), the harmonics model offers a simple and more global view of the roughness of the loss surface and how it is affected by nonlinearity choice and residual connections. Our observations on training speed are consistent with the hypothesis that spectral blueshift impedes training: Architectural choices (residual connections, ResNet v2 blocks, LReLU) that control harmonics generation are also those that are easier to train to good performance. Further, the results in Figure 6 give strong hints for a correlation between high-frequency content and problems in trainability for deep but very non-linear networks (small α). From an analytical perspective, architectures that are prone to blueshift are prone to vanishing/exploding gradients: We have shown that high magnitudes of high-frequencies in the w-o functions lead to an increase in the expected gradient magnitude. Further, in a simple multi-layer network without shortcuts, we expect more high-frequency content on lower than on higher layers as soon as the non-linearity of the activation function is actually is utilized (and therefore must also creates harmonics). A rescaling of layers can reduce these discrepancies, but the normalization will still necessarily depend on the position in parameter space. As far as accuracy is concerned, there is an obvious trade-off between blueshift and expressivity: networks with low-degree polynomials as nonlinearities contain less high-frequent harmonics but cannot approximate complex functions (Kidger & Lyons, 2020). Future work could consist in further exploring this trade-off. Further, it would be interesting to examine potential architectural alternatives to residual connections with comparable performance and trainability.

6 Conclusion and Future Work

In this paper, we have analyzed the effect of different architectures of non-linear networks on the smoothness of the computed function (various neural outputs and the loss landscape). Specifically, we have linked polynomial non-linearity to harmonics creation and validated experimentally that the mean power spectrum shifts towards higher frequencies ("blueshift"). Further, we have also described two distinct effects that explain the smoothing effect of shortcut connections, which we also confirmed experimentally. Finally, we empirically linked reduced blueshift to increased trainability and training speeds.

In future work, we hope to derive explicit closed-form equations for how architectural designs like network depth, residual connections and nonlinearity choice affect the w-o function. It could be also interesting to further explore the roughness-expressability trade-off that we described earlier.

REFERENCES

David Balduzzi, Marcus Frean, Lennox Leary, J. P. Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 342–350. PMLR, 2017.

Ronen Basri, David W. Jacobs, Yoni Kasten, and Shira Kritchman. The convergence rate of neural networks for learned functions of different frequencies. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 4763–4772, 2019.

Emmanuel J Candès. Harmonic analysis of neural networks. *Applied and Computational Harmonic Analysis*, 6(2):197–218, 1999.

Richard P Feynman, Robert B Leighton, and Matthew Sands. The feynman lectures on physics; vol. i. *American Journal of Physics*, 50(8), 1965.

Jonathan Frankle, David J. Schwab, and Ari S. Morcos. The early phase of neural network training. *CoRR*, abs/2002.10365, 2020.

Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 315–323, 2011.

Guy Gur-Ari, Daniel A. Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *CoRR*, abs/1812.04754, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 1026–1034. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.123.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016a. doi: 10.1109/CVPR.2016.90.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (eds.), *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, volume 9908 of *Lecture Notes in Computer Science*, pp. 630–645. Springer, 2016b. doi: 10.1007/978-3-319-46493-0_38.

S. Hochreiter. Untersuchungen zu dynamischen neuronalen netzen. Diploma thesis, TU Munich, 1991.

Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pp. 646–661. Springer, 2016.

Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 2261–2269. IEEE Computer Society, 2017. doi: 10.1109/CVPR.2017.243.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 448–456. JMLR.org, 2015.

Patrick Kidger and Terry J. Lyons. Universal approximation with deep narrow networks. In Jacob D. Abernethy and Shivani Agarwal (eds.), *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pp. 2306–2327. PMLR, 2020. URL `http://proceedings.mlr.press/v125/kidger20a.html`.

Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 6391–6401, 2018.

F.K. Musgrave. *Methods for Realistic Landscape Imaging*. PhD thesis, Yale University, 1993.

Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, volume 28 of *JMLR Workshop and Conference Proceedings*, pp. 1310–1318. JMLR.org, 2013.

Jeffrey Pennington, Samuel Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through : theory and practice. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 4785–4795. Curran Associates, Inc., 2017a.

Jeffrey Pennington, Samuel S. Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pp. 4785–4795, 2017b.

Jeffrey Pennington, Samuel S. Schoenholz, and Surya Ganguli. The emergence of spectral universality in deep networks. In Amos J. Storkey and Fernando Pérez-Cruz (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1924–1932. PMLR, 2018.

Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3360–3368, 2016.

Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl Dickstein. On the expressive power of deep neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2847–2854. JMLR.org, 2017.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht, Yoshua Bengio, and Aaron C. Courville. On the spectral bias of neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5301–5310. PMLR, 2019.

Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha Sohl-Dickstein. Deep information propagation. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *CoRR*, abs/1505.00387, 2015.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

Lifu Wang, Bo Shen, Ning Zhao, and Zhiyuan Zhang. Is the skip connection provable to reform the neural network loss landscape? In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 2792–2798. ijcai.org, 2020. doi: 10.24963/ijcai.2020/387. URL https://doi.org/10.24963/ijcai.2020/387.

Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5393–5402, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. In Tom Gedeon, Kok Wai Wong, and Minho Lee (eds.), *Neural Information Processing - 26th International Conference, ICONIP 2019, Sydney, NSW, Australia, December 12-15, 2019, Proceedings, Part I*, volume 11953 of *Lecture Notes in Computer Science*, pp. 264–274. Springer, 2019. doi: 10.1007/978-3-030-36708-4_22.

Zhiqin John Xu. Understanding training and generalization in deep learning by fourier analysis. *CoRR*, abs/1808.04295, 2018.

Greg Yang and Hadi Salman. A fine-grained spectral perspective on neural networks. *CoRR*, abs/1907.10599, 2019.

Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. A mean field theory of batch normalization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

A NETWORK DETAILS

A.1 NETWORK ARCHITECTURES

The implementation for our experiments is based on PyTorch 1.5 and will be provided on Github.

Toy-CNN: We create a simple convolutional network consisting of units of 3x3-convolutions with 16 feature channels on all layers ($16 \cdot a$ for averaging), followed by batch normalization and non-linear activation, stacked-up L-times. Optional short-cut connections are added (only if explicitely stated) between each unit (i.e., residual blocks of depth 1). No pooling or striding is used. The input layer uses a stride of 5 to reduce memory demands. The classification layer is composed of a fully-connected layer, softmax and cross-entropy loss.

Residual Network: As a "real-world" example, we use a standard Residual Network v1 (BasicBlock) He et al. (2016a) with varying depth, with ("Short/S") and without ("NoShort/NS") skip-connections. ResNet v2 blocks are used in Figure 1 (v2 blocks show slower transitions) and 4 (correlations are decaying faster in ResNet v2, ref. 3.5). We use the standard number of planes for Cifar-10 training : 16/32/64 (28/56/112 for the wide network).

A.2 FOURIER WALK HYPERPARMETERS

We sample paths consisting of 100 samples each and with path length $\alpha = 1$ (Eq. 1) in all experiments, except for Figures 1, 5 and 4 where we double both values for higher frequency resolution and Figure 13, where we want to walk in a different radius ($\alpha \in \{0.1, 1, 10\}$). All experiments are performed in training mode and with data augmentation at initialization.

A.3 TRAINING HYPERPARMETERS

The hyper-parameters below usually reach the standard test-accuracy of approximately 92-93% for a ResNet56 on Cifar10.

Dataset	Cifar10 (Cifar100 for Figure 9)
Epochs	200
Scheduler	Multistep ($\gamma = 0.1$)
Milestones	100, 150
Learning rate	0.9
Batch size	128
Optimizer	SGD + Momentum
Momentum	0.9
Weight decay	0.0001
Augmentation	Random Flip

B Formal Details Concerning Harmonic Network Analysis

B.1 Series Representation and Convergence

In our paper, we consider feed-forward networks consisting of L layers. Each layer $l = 1...L$ can be one of the following:

- An **affine map**
$$\mathbf{x} \mapsto \mathbf{W}^{(l)}\mathbf{x} + \mathbf{b}^{(l)} \tag{7}$$
that transforms the input using a linear map and an offset vector.
- A **non-linearity**
$$\mathbf{x} \mapsto \phi^{(l)}(\mathbf{x}) \tag{8}$$
that applies a (potentially layer-specific) non-linear function to each input.

All of the discussed architectures (including convolutional networks, striding/average pooling, residual connections, parallel computations, and layer-dependent non-linearitiy) can be expressed as a composition of these two layer types. We currently do not address max-pooling, and understand batch-normalization layers statically as a fixed scaling and shifting (not analyzing their training dynamics). The loss function itself (such as a sequence of softmax and cross-entropy) is also not included in our analysis – we restrict our consideration to layer outputs before any classification. As logistic regression (or similar problems) are numerically well-understood already, this is not a major obstacle. All our experiments show layer outputs before classification (softmax).

The main restriction of our formal model is that we assume that all non-linearities are polynomials of finite degree[1]. Under this condition, all outputs and intermediate results of the whole network can be represented by a large multi-variate polynomial, and any *linear section* $p(t) = f_i^{(l)}(\mathbf{W} + \alpha t \mathbf{D}, \mathbf{x})$ of any output i of any layer l is a univariate polynomial of finite degree.

Fourier analysis: Polynomial functions on a finite domain are Lipschitz-continuous and have finite variation; both conditions are sufficient for the convergence of the Fourier series in an L_2-sense. Further, as the Fourier-basis forms a Schauder-basis of $L_2[0, 1]$, the series expansion is unique.

In our derivation (Equation 5), we plug the series representation into the polynomial and conclude that it can be represented in the spectral domain as a sum of convolutions (Equation 6). To make this more rigorous, we proceed in smaller steps: Let again $p : \mathbb{R}^n \to \mathbb{R}$ denote the linear section along a linear direction in parameter space of any layer output. Our non-linearity expands to

$$\phi(p(t)) = \sum_{j=0}^{K} t_j p(t)^j. \tag{9}$$

As $p(t)$ is a polynomial of finite degree, $p(t)^j$ for any finite j is also a polynomial of finite degree. This implies that there is a unique Fourier series that represents p, characterized by the coefficients

[1]This means that the analytical results are currently not established for practical non-polynomial non-linearities. The results in our paper obtained from polynomial approximations should therefore be considered experimental results at this point. An analysis of infinite series approximations covering a more general class of non-linearities is still left for future work.

sequence $z : \mathbb{Z} \to \mathbb{C}$. According to the convolution theorem for the complex Fourier series, the Fourier series of $p \cdot p$ is given by $z \otimes z$, where

$$[z \otimes u]_k := \sum_{m=-\infty}^{\infty} z_m \cdot u_{k-m}. \tag{10}$$

Because the series must exist (it converges for both p and p^2) and it is unique, the product of the series must be represented by the auto-convolution. Inductively, we obtain p^j corresponding to the j-fold auto-convolution $\bigotimes_{q=1}^{j} z$.

Note: We use the notation $[\mathbf{x}]_i := x_i$ denote the indexing of vectors and sequences that result from a computation.

B.2 Inputs to the Second Layer Non-Linearities are FBM-Noise

For the first layer of a ReLU network, the following calculation also supports the experimental findings of an $\mathcal{O}(1/k)$-power law from an analytical perspective. Let $p^{(1)}$ be a w-o path of the output of a given neuron from the first layer. It follows that:

$$\begin{aligned} p_i^{(1)}(t) &= \mathrm{ReLU}\left[\sum_j \left(w_{ij}^{(1)} + \alpha^{-1} t \cdot d_{ij}^{(1)}\right) \cdot x_j\right] && (11) \\ &= \mathrm{ReLU}\left[m \cdot t + n\right], \text{ for some } m, n \in \mathbb{R}. && (12) \end{aligned}$$

This is a piecewise linear function that is constant zero on one interval and linear with slope m on another interval. With $a < b \in [0, 1]$ bounding the activation interval, the Fourier series becomes:

$$z_k^{(1,i)} = \int_a^b \exp\left(-2\pi i k \left(mt + n\right)\right). \tag{13}$$

By the chain rule of derivations, this implies

$$|z_k^{(1,i)}| \in \mathcal{O}\left(\frac{1}{k}\right). \tag{14}$$

Correspondingly, the preactivations of the second layer will be random linear combinations of FBM-noise with $h = 1$, which again forms FBM-noise with that spectral variance (every Fourier coefficient is a sum of independent coefficients; therefore, variance is additive).

Similar findings where also previously shown in input space by Balduzzi et al. (2017), which, for the first layer, is just the dual to varying the weights (and therefore must yield the same result).

B.3 Blueshift and Depth

We now aim at formally understanding the blueshift effect in a multi-layer network. For simplicity, we consider a simple stack of linear and non-linear layers with a fixed non-linearity:

$$f(\mathbf{x}, \mathbf{W}) = \phi\Big(\mathbf{W}^{(L)} \phi\Big(\mathbf{W}^{(L-1)} \cdots \phi(\mathbf{W}^{(1)}\mathbf{x}) \cdots\Big)\Big). \tag{15}$$

We now consider a single output i of layer l:

$$f_i^{(l)}(\mathbf{x}, \mathbf{W}) = \phi\Bigg(\mathbf{w}_i^{(l)} \cdot \phi\Big(\mathbf{W}^{(l-1)} \phi\Big(\mathbf{W}^{(l-2)} \cdots \phi(\mathbf{W}^{(1)}\mathbf{x}) \cdots\Big)\Big)\Bigg). \tag{16}$$

Next, we consider a linear section $t \mapsto \mathbf{W} + t\mathbf{D}$ with $t \in [0, \alpha]$:

$$p_i^{(l)}(t) = \phi\Bigg(\mathbf{w}_i^{(l)} \cdot \phi\Big((\mathbf{W}^{(l)} + t\mathbf{D}_l)\phi\Big((\mathbf{W}^{(l-1)} + t\mathbf{D}^{(l-1)}) \cdots \phi(\mathbf{W}^{(1)} + t\mathbf{D}^{(1)}\mathbf{x}) \cdots\Big)\Big)\Bigg).$$

By replacing ϕ with finite polynomials (Equation 4), we obtain multi-variate polynomials as outputs of all of these function.

Varying parameters in a single layer: If we now select a layer $l = k$ by setting all $\mathbf{D}^{(i)} = \mathbf{0}$ for all $i \neq k$, and only using the k-th matrix as direction vector, the parameter t will pass through $l - k + 1$ non-linearities:

$$p_i^{(l)}(t) \quad = \quad \phi\Bigg(\mathbf{w}_i^{(l)} \cdot \phi\Bigg(\mathbf{W}^{(l-1)} \cdots \phi\Big((\mathbf{W}^{(i)} + t\mathbf{D}^{(i)})q^{(i-1)} \cdots\Big) \cdots\Bigg)\Bigg). \tag{17}$$

with a constant vector $q^{(i-1)}$.

Correspondingly, each non-linearity will act on the spectrum of $q^{(k)}$ as a blue-shift operator (forming the sum of repeated auto-convolutions of the spectrum, weighted by the polynomial coefficients). Therefore, we see that parameters of earlier layers (small layer index k, closer to the input layer) are blueshifted more frequently than later layers.

The equation also shows that the singular value spectrum of $W^{(l)}$ affects the results as well, in addition to the spectral shifts caused by the non-linearities: the values rescale the input domain (shrink/expand space wrt. parameter t), thereby changing the frequency reciprocal to the scale factor. However, a linear transformation scales the Fourier spectrum as a whole while non-linearities spread the spectrum irreversibly by applying blueshifts (weighted sums of auto-convolutions of the spectrum). The magnitude of this non-linear effect is, nontheless, dependent on the signal magnitude and thus affected by choices of $\mathbf{W}^{(l)}$.

Varying parameters in layer $0..l$: Similar arguments hold if we vary only parameters in the first l layers of the network. However, we obtain a mix of spectra that have been blueshifted (and scaled by weight matrices) a different number of times.

B.4 Spectral Broadening through Auto-Convolutions

The exact amount of broadening of a spectrum $z : \mathbb{Z} \to \mathbb{R}$ by an auto-convolution is non-trivial to quantify; we therefore resort to experiments for determining the effect in practice.

In special cases, we can analyze the effect:

Band limited functions: Assuming that z is bandlimited, i.e., $z_k = z_{-k} = 0$ for all $k > K$, and $|z_k| = |z_{-k}| > 0$ for $k \leq K$ it is trivial to see that an auto-convolution extends the support by k entries. This means, after j convolutions, we obtain a non-zero spectrum within $k = -jK...jK$. The proof is obtained by considering the left-most and right-most terms in the definition of the discrete convolution.

Central limits: If we assume that z is real and positive, and has finite second moments, the central limit theorem guarantees convergence towards a Gaussian function with standard deviation $\mathcal{O}\left(\sqrt{j}\right)$. The first assumption is obviously unrealistic in practice.

In the general case the j-th power of a Fourier series

$$\left[\sum_{k=-K}^{K} z_k \exp\left(2\pi i k t\right)\right]^j \tag{18}$$

expands to the sum

$$\sum_{k_1=-K}^{K} \cdots \sum_{k_j=-K}^{K} z_{k_1} \cdots z_{k_j} \exp\left(2\pi i (k_1 + k_2 + \cdots + k_j)\right). \tag{19}$$

If we only count the number of terms of the same frequency $k = k_1 + \cdots + k_j$, these (occurrence) frequencies of each frequency k will tend towards a normal distribution with variance j. However, as the complex coefficients can cancel out, the shape of the limit distribution is not yet determined by this.

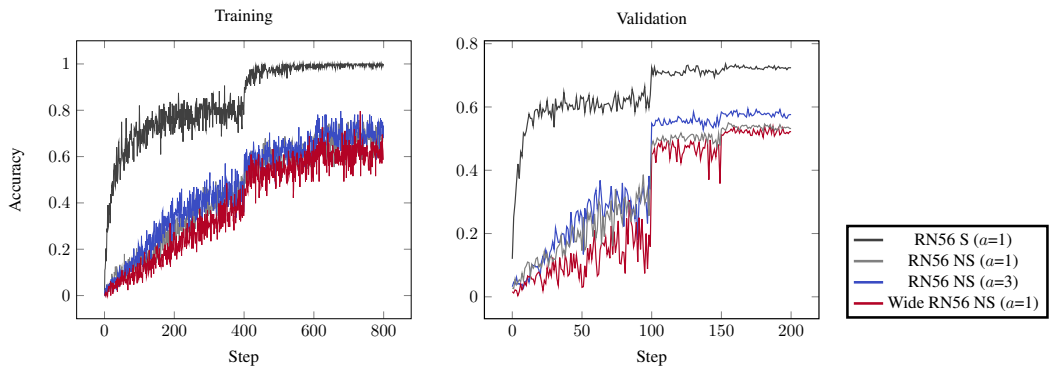

Figure 9: Repeating the experiments of Figure 8 on the Cifar100 dataset.

B.5 Frequency Dependence in Averaging of Multiple Computation Paths

We observe that averaging functions after non-linearities leads to some smoothing of the result by decreasing the correlation between higher-frequency Fourier coefficients. In order to understand the effect, we look at a simple 1D model case: We consider a the composite functions

$$f_{post}(x) = \phi\big(\underbrace{w_1 \cdot g(x)}_{=:h_1'} + \underbrace{w_2 \cdot g(x)}_{=:h_2'}\big) \tag{20}$$

and

$$f_{pre}(x) = \underbrace{\phi(w_1 \cdot g(x))}_{=:h_1} + \underbrace{\phi(w_2 \cdot g(x))}_{=:h_2} \tag{21}$$

where w_1, w_2 are random numbers, drawn from the same normal distribution.

Let z_k be the sequence of Fourier coefficients of g, and u_k that of the result. In the first case, f_{post}, a random linear combintation of the input spectra is performed, which is then blue-shifted:

$$u = \sum_{j=1}^{K} \bigotimes_{q=1}^{j} \big(w_1 z + w_2 z\big) \tag{22}$$

In the second case,

$$u = \sum_{j=1}^{K} \bigotimes_{q=1}^{j} \big(w_1 z\big) + \sum_{j=1}^{K} \bigotimes_{q=1}^{j} \big(w_2 z\big), \tag{23}$$

we first blue-shift the randomly-scale spectra and then form the sum. This creates a frequency dependence because larger powers (larger values of j, corresponding to more repeated convolutions, and thus j-fold products of input Fourier coefficients) behave more non-linearly and thus react more strongly to the random scaling: While a linearly weighted average will be linearly correlated with the original, averages of non-linearly transformed functions will loose linear correlation (and thus appear more random to a simple averaging operation). The harmonic distortion analysis that expresses the powers as auto-convolutions of the spectrum shows that higher-frequency components (created by the blue-shift) are also the ones that behave more non-linearly.

C Training on Cifar100

We repeat the experiment on averaging-networks for the Cifar100 dataset. We leave all hyperparameters untouched. In Figure 9, we still see the advantage of the averaging-network over the regular network, especially in validation accuracy.

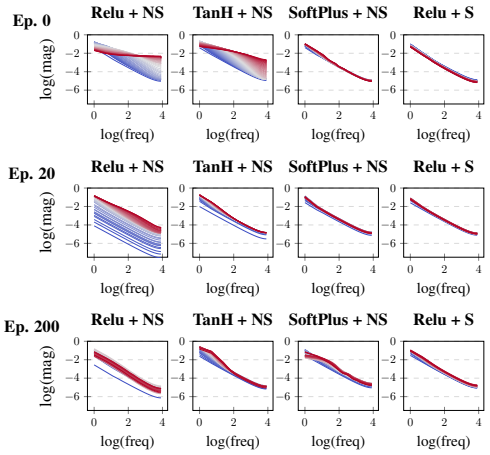

Figure 10: Repeating the experiments of Figure 2 (all layers, unscaled) during training.

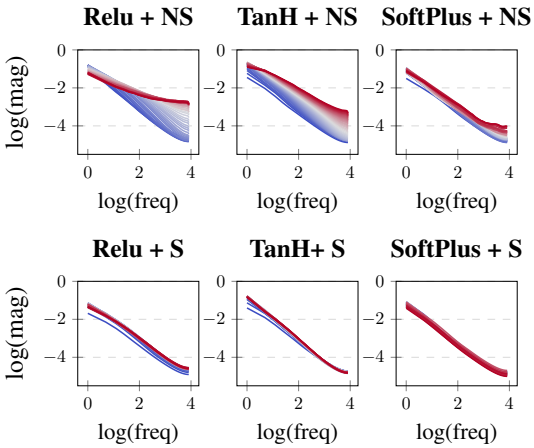

Figure 11: Repeating the experiments of Figure 2 (all layers, unscaled) at initialization in gradient direction.

D OTHER SPECTRAL SHIFT FIGURES

D.1 SPECTRAL SHIFT DURING TRAINING

To show that the blueshift effect is not limited to the initialization, we show the spectral shift of a Toy-CNN 50 at 0/20/200 epochs of training in Figure 10. ResNets still do not show any major blue-shift at any time. The architectures without skip-connections show blueshift at any time, albeit progressively less pronounced, indicating that blueshift is more of a problem in early phases of training (Frankle et al., 2020).

D.2 SPECTRAL SHIFT IN GRADIENT DIRECTION

The subspace of the loss surface where gradient descent operates has been studied (Gur-Ari et al., 2018) and attributed a specific behavior. We want to see how blueshift behaves when slicing the network in the direction of gradient descent. For this, we repeat the experiments of Figure 2, but compute the gradient at initialization, normalize it and use it as direction $\mathbf{D}$ of equation 1. In Figure 11, we see that the effect is still visible although dampened for ReLU and tanh activations. The effect seems a little more visible for softplus activations.

D.3 SPECTRAL SHIFT WITH DELTA-ORTHOGONAL INITIALIZATION

The experiments of Figure 2 with batch normalization are subject to exploding gradients. Xiao et al. (2018) describe an initialization scheme that fixes exploding gradients for TanH activations. Fig 13 shows that for a very small region near the orthogonal initialization point, blueshift vanishes. As soon as we walk further from the initialization point, blueshift is visible again.

D.4 WALK PER LAYER

To see the contribution of each layer to a loss path, we sample a random direction $\mathbf{D}$, but measure the resulting loss path when only modifying weights in a single layer. We realize this for a Toy-CNN with 50 layers at initialization in Figure 12. We see that varying weights in higher layers (red) results in smooth paths whereas varying weights in lower layers (blue) results in rougher paths. In the normalized power-spectra, we can see the increased blueshift of paths resulting from varying earlier layers.

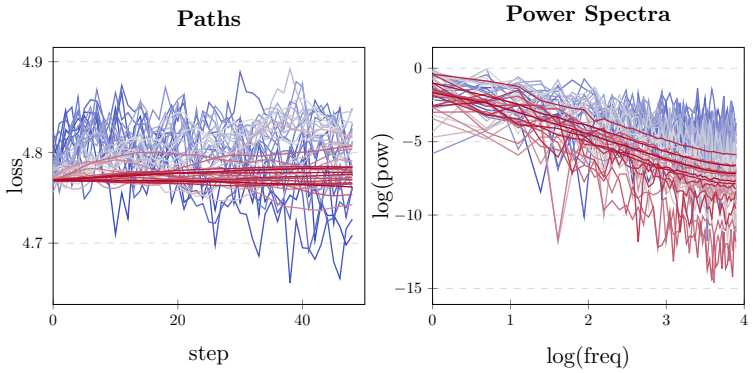

Figure 12: Resulting loss paths and respective path-normalized power spectra when varying parameters in only one layer of a Toy-CNN 50 at initialization. The color of a path indicates in which layer parameters were varied (red is deeper).

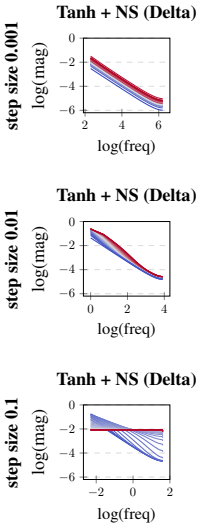

Figure 13: Repeating the experiments of Figure 2 (all layers, unscaled) for TanH activations with delta-orthogonal initialization.

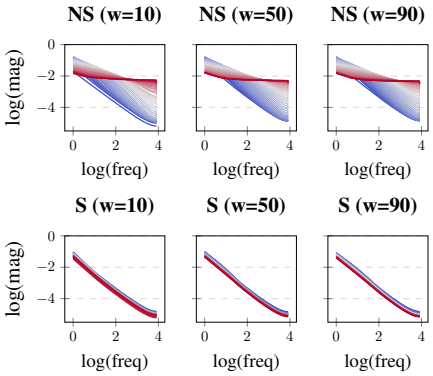

Figure 14: Repeating the experiments of Figure 2 (all layers, unscaled) for variable network width.

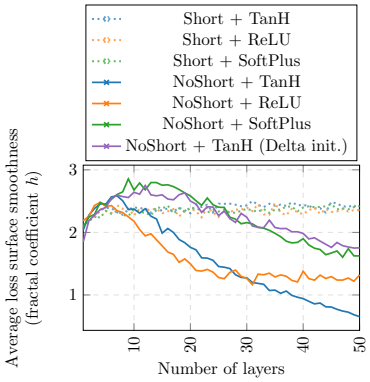

Figure 15: Smoothness of the loss surface at initialization for a Toy-CNN of varying depth with and without shortcuts on MNIST for a batch size of 256.

D.5 Varying width

To see if layer width has any influence on blueshift, we repeat the experiment of Figure 2 for a Toy-CNN 50 with number of filters w. We can see on Figure 14 that layer width doesn't seem to have a notable impact on blueshift strength.

D.6 More activation functions

To demonstrate that blueshift occurs with every nonlinear activation function, we repeat the experiment of Figure 2 for exponential linear units (ELU), gaussian error linear units (GELU), hard tanh (HTANH), leaky relu (LReLU), scaled exponential linear unit (SELU) and sigmoid activations.

In order to demonstrate the that spectral spread depends on the magnitude of higher-order polynomial coefficients, we show the plot of the magnitude of the coefficients of a polynomial approximation

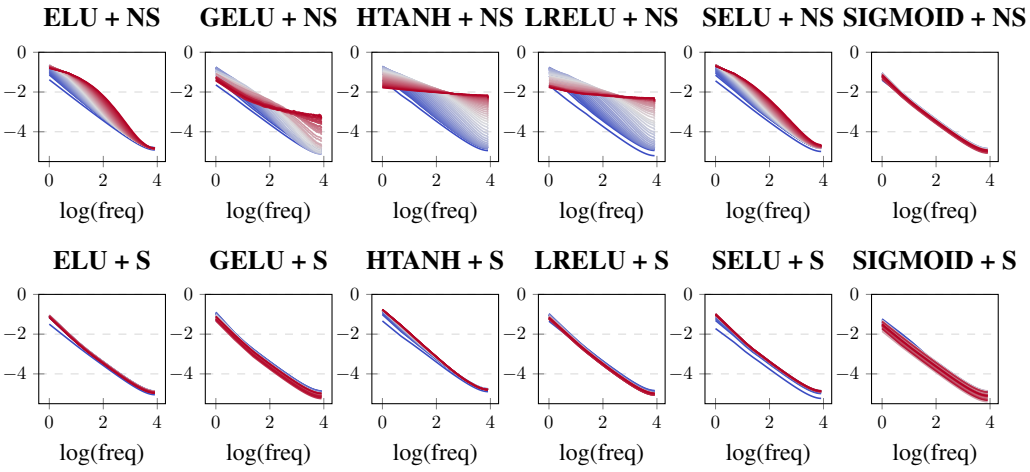

Figure 16: Repeating the experiments of Figure 2 (all layers, unscaled) for more nonlinearities.

Polynomial Chebyshev approximation of various nonlinearities of degree 25 within [-5,5]

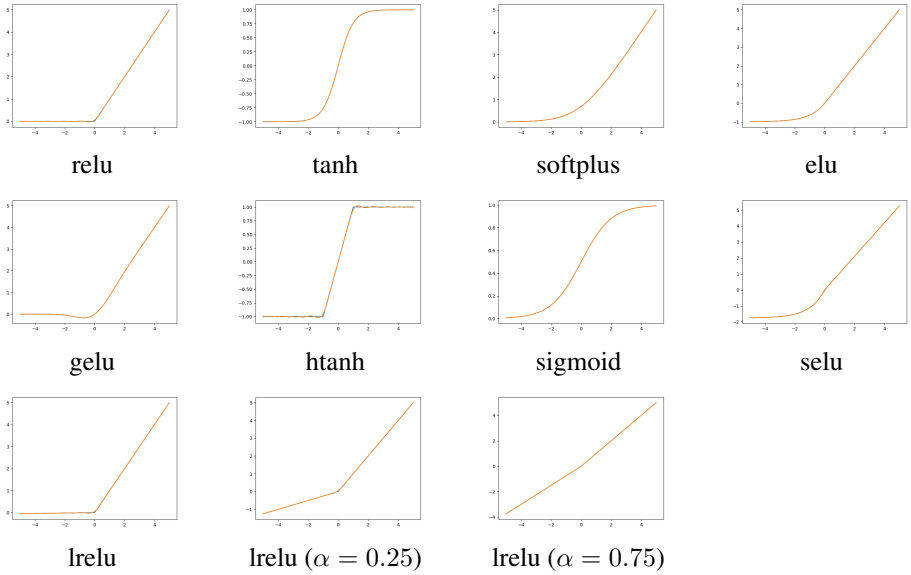

Figure 17: Various Non-Linearities with their Chebyshev-Polynomial Approximation within the interval $[-5, 5]$.

with Chebyshev nodes of degree 25 in Figure 18. The fit is performed for the interval $[-5, 5]$ (slightly different from Figure 2 in the main paper, which shows a fit for [-1,1]). Figure 18(a) also shows again the comparison between ReLU, $\tanh$, and $\mathrm{softplus}$, uniformely with this Chebyshef approximation; this shows the differences more clearly (please also note the logarithmic scale of the y-Axis).

The approximations are quite tight within the interval, see Figure 17 for reference.

E Smoothness measurements on MNIST

To show that our smoothness measurements are mostly independent of batch size and dataset, we revisit Figure 3. This time, we use the MNIST dataset and a batch size of 256. All other hyperparameters are maintained. We see that although the absolute values slightly differ, the relative behavior of the nonlinearities is qualitatively the same.

Comparison of the magnitude of the polynomial coefficients $|t_j|$
(for the degree 25 approximation of Figure 17)

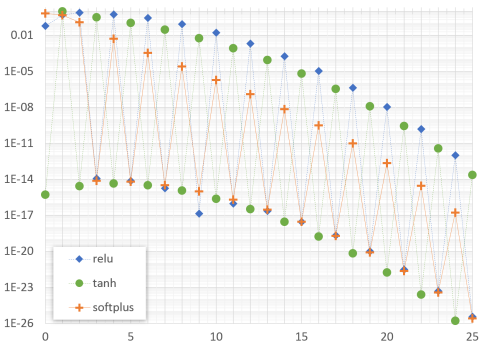

(a) comparison of relu, tanh, softplus
(different interval from main paper!)

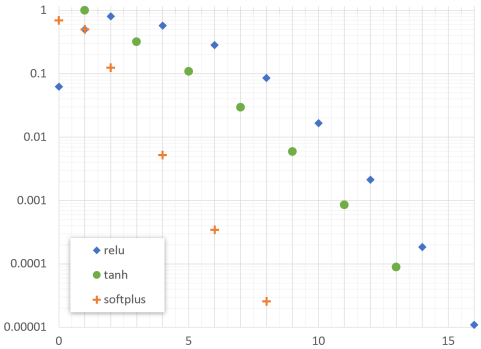

(b) close-up of (a)

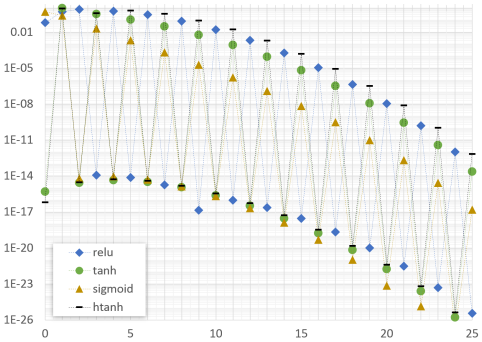

(c) comparison of relu, elu, gelu, selu

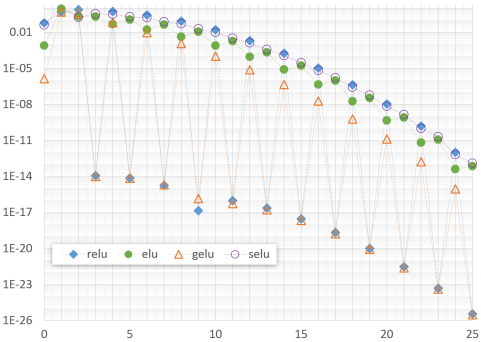

(d) comparison of relu,tanh,sigmoid,hardtanh
(relu for reference)

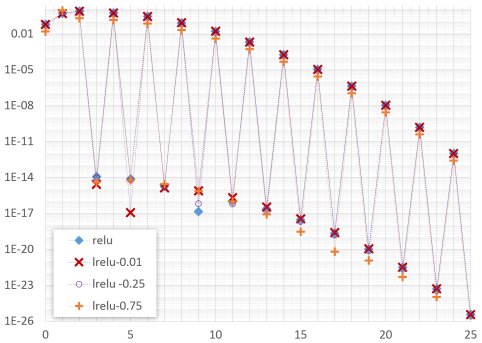

(d) different variants of leaky-relu
(default is $\alpha = 1E-2$)

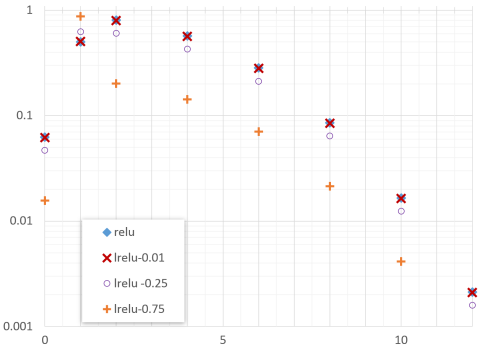

(e) closeup of (d)

Figure 18: Absolute value of the polynomial coeffients of the Chebyshev polynomial approximation. All fits are done within the interval $[-5, 5]$. Note: Different interval from Fig. 2 in the main paper! (to capture the shape of all non-linearities well).