

BEYOND GAUSSIAN INITIALIZATIONS: SIGNAL PRESERVING WEIGHT INITIALIZATION FOR ODD-SIGMOID ACTIVATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

Activation functions critically influence trainability and expressivity, and recent work has therefore explored a broad range of nonlinearities. However, widely used Gaussian i.i.d. initializations are designed to preserve activation variance under wide or infinite width assumptions. In deep and relatively narrow networks with sigmoidal nonlinearities, these schemes often drive preactivations into saturation, and collapse gradients. To address this, we introduce an odd-sigmoid activation and propose an activation aware initialization tailored to any function in this class. Our method remains robust over a wide band of variance scales, preserving both forward signal variance and backpropagated gradient norms even in very deep and narrow networks. Empirically, across standard image benchmarks we find that the proposed initialization is substantially less sensitive to depth, width, and activation scale than Gaussian initializations. In physics informed neural networks (PINNs), scaled odd-sigmoid activations combined with our initialization achieve lower losses than Gaussian based setups, suggesting that diagonal-plus-noise weights provide a practical alternative when Gaussian initialization breaks down.

1 INTRODUCTION

Deep learning has advanced substantially across diverse domains (LeCun et al., 2015; He et al., 2016; Hwang et al., 2022). A central ingredient in these successes is the choice of activation function, which controls both the expressive power of a network and the way signals and gradients propagate through depth (Poole et al., 2016). Recent work has proposed a wide variety of nonlinearities beyond classical sigmoid and ReLU (e.g., GELU, Swish, Mish, SELU) to improve trainability, stability, and accuracy (Hendrycks & Gimpel, 2016; Ramachandran et al., 2017; Misra, 2019; Klambauer et al., 2017; Murray et al., 2022; Bingham & Miikkulainen, 2023; Zhang et al., 2024). However, activation functions and weight initialization are tightly coupled (He et al., 2015; Lee et al., 2024). The scale and shape of the nonlinearity determine how information propagates with depth. As a result, an initialization that is not tuned to the chosen activation can drive activations into saturation and cause gradients to vanish or explode, even when the activation itself is reasonable. These issues are especially pronounced for sigmoidal activations, which are widely used in sequence models and physics informed neural network (PINN) (Raissi et al., 2019).

Standard Gaussian i.i.d. initializations (Xavier (Glorot & Bengio, 2010), He (He et al., 2015), EOC (Hayou et al., 2019) are derived under wide or infinite width assumptions and choose a single variance parameter to preserve signal propagation. In practice, we find that this design breaks down in deep and relatively narrow networks and under activation rescaling: small deviations in the variance can trigger saturation, or extreme learning rate sensitivity. Motivated by these




Figure 1: odd-sigmoid activations

054 observations, we propose an activation aware initialization for fully feedforward neural net-
 055 works (FFNNs) with odd-sigmoid activations in the class \mathcal{F} (Figure 1). For any $f \in \mathcal{F}$ with
 056 $\omega := 1/f'(0)$, we study the scalar dynamics of $x \mapsto f(ax)$ and choose the layerwise gains a from a
 057 distribution with mean ω so that trajectories neither collapse to zero nor saturate. This gain design is
 058 then implemented via a simple initial weight matrix whose effective gain statistics match the desired
 059 behavior. Unlike Gaussian i.i.d. schemes, whose trainability is highly sensitive to the exact vari-
 060 ance, the proposed initialization keeps both forward activations and backpropagated gradients in a
 061 well conditioned range across depth (Raghu et al., 2017). Mean field analysis further shows that the
 062 corresponding gradient amplification factor χ_ℓ remains close to 1 over a significantly wider band of
 063 noise scales (Schoenholz et al., 2016).

064 We evaluate the proposed initialization in extensive experiments with odd-sigmoid activations, cover-
 065 ing both the regime $\omega \approx 1$ and the more challenging regime $\omega \gg 1$ or $\omega \ll 1$. On standard image
 066 classification benchmarks, Gaussian initializations often fail or degrade sharply in deep, narrow, or
 067 rescaled-activation networks, whereas our scheme trains stably and is more data-efficient than these
 068 Gaussian baselines. We also study scaled and composite activations (e.g., $\alpha f(x)$, $f(\alpha x)$, and posi-
 069 tive linear combinations) and observe that our initialization remains stable even for extreme scales
 070 (up to $\alpha \approx 10^9$). In PINN settings, appropriately scaled odd-sigmoid activations combined with our
 071 initialization consistently achieve lower losses than EOC-based baselines and do so without batch
 072 normalization (BN) (Ioffe, 2015), indicating that diagonal-plus-noise weights provide a practical
 073 alternative when Gaussian schemes become brittle. Our main contributions are as follows:

- 074 • We define the odd-sigmoid function class and characterize its properties (Section 4.1).
- 075 • We view a deep feedforward network as a parallel ensemble of scalar dynamical systems,
 076 and use the resulting gain statistics to design a proposed initialization (Section 4.2).
- 077 • We show that the proposed initialization is more variance robust than standard Gaussian
 078 i.i.d. initializations and better preserves forward and backward signals (Section 4.3).
- 079 • We empirically validate the proposed initialization across a wide range of activations on
 080 standard image classification benchmarks and PINNs (Section 5.1 and 5.2).

083 2 RELATED WORK

084
 085 Classic initialization schemes were originally designed to stabilize layerwise variance, and later
 086 work showed that good performance also hinges on well-tuned statistics and momentum, with poor
 087 initialization causing saturation, variance collapse, and learning-rate fragility (Sutskever et al., 2013;
 088 Narkhede et al., 2022). For sigmoidal networks, follow-up studies adjust the mean or support of the
 089 weights or use more spread-out mappings to keep units responsive and accelerate convergence (Yilmaz
 090 & Poli, 2022; Qiao et al., 2016; Sodhi et al., 2014). Beyond first-order variance matching,
 091 curvature- and correlation-based criteria (Hessian-norm control, depth \times width bounds, and propa-
 092 gation analyses in constrained architectures such as input-convex networks and transformers) and
 093 architectural tweaks such as ReZero and dynamical-isometry-preserving setups all point toward ini-
 094 tializations that are explicitly calibrated to the activation and the network’s depth and shape (Skorski
 095 et al., 2021; Iyer et al., 2023; Hoedt & Klambauer, 2023; Noci et al., 2022; Bachlechner et al.,
 096 2021; Lee et al., 2024; Price et al., 2024; Blumenfeld et al., 2020). Our work makes this activa-
 097 tion-initializer coupling explicit by formalizing an odd-sigmoid class and, for any f in this class,
 098 deriving a closed-form noise scale that keeps preactivations in a high-gain regime up to a target
 099 depth, yielding dispersed forward signals and stable gradients without normalization (Noci et al.,
 200 2022; Skorski et al., 2021).

100 On the activation side, early bounded nonlinearities (Sigmoid, Tanh) suffered from vanishing gradi-
 101 ents, whereas ReLU variants improved efficiency but can produce dying neurons (Adeli et al., 2025).
 102 Surveys emphasize that no single activation is universally optimal and highlight desirable traits such
 103 as continuity, boundedness, monotonicity, symmetry, and smoothness (Alcaide et al., 2025; Adeli
 104 et al., 2025), with bounded, odd-symmetric forms helping to center signals and stabilize optimiza-
 105 tion—directly motivating our initializer. In parallel, adaptive and learned activation families (e.g.,
 106 parametric activations for DL and PINNs, error-function-corrected ReLU variants, per-neuron self-
 107 activating networks, and tailored activations for zeroing neural networks) demonstrate that carefully
 structured nonlinearities can improve convergence and robustness (Jagtap et al., 2020; Wang et al.,

2023; Ullah et al., 2025; Tutuncuoglu, 2025; Liu, 2025), further supporting the need for initialization strategies aligned with activation geometry.

3 PRELIMINARIES

In this section, we introduce the notation used throughout the paper, review a simplified elementwise analysis of signal propagation (Lee et al., 2025), and recall the classical edge of chaos mean field theory (Schoenholz et al., 2016) and its associated Gaussian i.i.d. EOC initialization ((Hayou et al., 2019)).

Notation. We consider a feedforward neural network with L layers. The dataset comprises K pairs $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^K$, where $\mathbf{x}_i \in \mathbb{R}^{N_x}$ is the input and $\mathbf{y}_i \in \mathbb{R}^{N_y}$ is the corresponding target. For $\ell = 1, \dots, L$, the layerwise update is

$$\mathbf{h}^\ell = \mathbf{W}^\ell \mathbf{x}^{\ell-1} + \mathbf{b}^\ell, \quad \mathbf{x}^\ell = f(\mathbf{h}^\ell) \in \mathbb{R}^{N_\ell},$$

where $\mathbf{W}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ is the weight matrix, $\mathbf{b}^\ell \in \mathbb{R}^{N_\ell}$ is the bias vector, and $f(\cdot)$ denotes an activation function. We write $\mathbf{W}^\ell = [w_{ij}^\ell]$.

Signal Propagation Analysis. Following the simplified framework of Lee et al. (2025), we analyze signal propagation in feedforward networks with an elementwise activation f . For convenience, all layers, including input and output, are assumed to have dimension n (i.e., $N_\ell = n$ for all ℓ). A key observation of this approach is that the usual matrix–vector multiplication $\mathbf{W}^\ell \mathbf{x}^{\ell-1}$ can be reformulated elementwise, such that each coordinate is expressed in terms of an effective self-scaling factor. Formally, for $\ell = 0, \dots, L-1$ and $i = 1, \dots, n$,

$$x_i^{\ell+1} = f(a_i^{\ell+1} x_i^\ell), \quad a_i^{\ell+1} = w_{ii}^{\ell+1} + \sum_{j \neq i} \frac{w_{ij}^{\ell+1} x_j^\ell}{x_i^\ell}, \quad (1)$$

assuming that $\mathbf{b}^\ell = \mathbf{0}$. This elementwise decomposition highlights the role of $a_i^{\ell+1}$ as an effective gain that combines both diagonal and off-diagonal terms and provides a tractable basis for studying signal propagation in deep networks.

Edge of chaos. We briefly recall the “edge of chaos” viewpoint based on mean field theory for fully connected networks with standard i.i.d. Gaussian initialization, where the weights are sampled as $w_{ij}^\ell \sim \mathcal{N}(0, \sigma_w^2/N_\ell)$ (Poole et al., 2016; Schoenholz et al., 2016). In this setting, the preactivations h_i^ℓ are modeled as i.i.d. Gaussians with layerwise variance q^ℓ , and both forward signals and backward gradients admit simple recursions in the infinite width limit. In particular, the mean field analysis of backpropagation shows that the squared gradient norm obeys

$$\chi_{\ell+1} \approx \sigma_w^2 \mathbb{E}[f'(\mathbf{h}^{\ell+1})^2],$$

where $\mathbf{h}^{\ell+1}$ denotes a typical preactivation at layer $\ell+1$. The scalar factor $\chi_{\ell+1}$ is the average gradient amplification of layer $\ell+1$: in the ordered phase ($\chi_{\ell+1} < 1$) gradients vanish exponentially with depth, while in the chaotic phase ($\chi_{\ell+1} > 1$) they explode. The critical curve defined by $\chi_{\ell+1} \approx 1$ is called the edge of chaos and separates these two regimes. At this edge, information and gradients can propagate through many layers (Schoenholz et al., 2016). Building on this perspective, (Hayou et al., 2019) refined the EOC analysis and showed that choosing (σ_w^2, σ_b^2) on the EOC curve leads to better information propagation and faster training in deep networks. Throughout this paper, we refer to Gaussian i.i.d. initializations with (σ_w^2, σ_b^2) on the EOC curve as EOC initialization.

4 METHODOLOGY

We first define the odd–sigmoid function class (Section 4.1). Building on this, we propose a weight initialization for feedforward networks with odd–sigmoid activations (Section 4.2). In Section 4.3, we highlight how the proposed initialization differs from standard Gaussian i.i.d. initializations, and show these differences through theoretical analysis and numerical experiments. The proofs of the theoretical results are provided in Appendix A, and additional numerical results are presented in Appendix B.

4.1 ODD-SIGMOID FUNCTION CLASS

In practical neural networks with sigmoidal activations, both forward signals and backpropagated gradients tend to saturate or vanish as the depth increases. Our goal in this work is to move beyond a single sigmoid and consider a broad class of odd-sigmoid activation functions, and to design a weight initialization that prevents such vanishing for any activation in this class. To this end, Section 4.1 introduces the odd-sigmoid function class and establishes its basic properties, which will be used in our subsequent analysis. Formal definitions of the activation functions used in this paper are provided in Appendix B.1.

Definition 4.1. A function $f : \mathbb{R} \rightarrow \mathbb{R}$ is an *odd-sigmoid function* if it satisfies the following:

- (i) **Regularity:** $f \in C^1(\mathbb{R})$.
- (ii) **Odd symmetry:** $f(-x) = -f(x)$ for all $x \in \mathbb{R}$.
- (iii) **Boundedness:** $\sup_{x \in \mathbb{R}} |f(x)| < \infty$.
- (iv) **Strict monotonicity:** $f'(x) > 0$ for all $x \in \mathbb{R}$.
- (v) **Slope decay:** f' is strictly decreasing on $[0, \infty)$.

Denote the class of all odd-sigmoid functions as \mathcal{F} . In the following, we establish several basic properties of functions in \mathcal{F} .

Pitchfork Bifurcation. Recall that $x^* \in \mathbb{R}$ is a fixed point of $f : \mathbb{R} \rightarrow \mathbb{R}$ if $f(x^*) = x^*$. For $f \in \mathcal{F}$ define $\omega_f := 1/f'(0) > 0$ and, unless stated otherwise, write ω simply. Consider $\phi_a(x) := f(ax)$ for $a > 0$. We say that ϕ_a has a pitchfork bifurcation at $a = \omega$ if it has exactly one fixed point $\{0\}$ for $0 < a < \omega$ and exactly three fixed points $\{0, \pm\xi_a\}$ for $a > \omega$, with the nonzero points occurring as a symmetric pair. We show that this holds for every $f \in \mathcal{F}$.

Proposition 4.2. Suppose $f \in \mathcal{F}$ with $\omega := 1/f'(0)$, and for a fixed $a > 0$ define $\phi_a(x) := f(ax)$. Then

- (i) If $0 < a \leq \omega$, then $\phi_a(x)$ has a unique fixed point $x^* = 0$.
- (ii) If $a > \omega$, then $\phi_a(x)$ has three distinct fixed points: $x^* = -\xi_a, 0, \xi_a$ such that $\xi_a > 0$.

The following theorem establishes convergence properties of $x_{n+1} = f(ax_n)$ for all $x_0 > 0$.

Theorem 4.3. Suppose $f \in \mathcal{F}$ with $\omega := 1/f'(0)$, and for a fixed $a > 0$ define

$$x_0 > 0, \quad x_{n+1} = \phi_a(x_n), \quad n = 0, 1, 2, \dots$$

Then the sequence $\{x_n\}$ converges for every $x_0 > 0$. Furthermore,

- (1) if $0 < a \leq \omega$, then $x_n \rightarrow 0$ as $n \rightarrow \infty$.
- (2) if $a > \omega$, then $x_n \rightarrow \xi_a$ as $n \rightarrow \infty$.

According to Theorem 4.3, when $a > \omega$, for any initial value $x_0 > 0$ (resp. $x_0 < 0$), the sequence defined by $x_{n+1} = f(ax_n)$ converges to ξ_a (resp. $-\xi_a$) as $n \rightarrow \infty$. From a signal propagation viewpoint in FFNN (Equation 1), this implies that activations do not vanish as depth increases. See Figures 8 and 9 for the convergence of the iterates $x_{n+1} = \phi_a(x_n)$ as a function of a and the initial value x_0 . Proposition A.2 and Corollary A.3 analyze the iteration with a coefficient sequence $\{a_m\}$ that may vary across steps. They show that excessively large or small a_m drives the dynamics into saturation, highlighting the importance of appropriately scaling $\{a_m\}$ when designing the initialization (Figure 2).

Corollary 4.4. Let $f_1, f_2 \in \mathcal{F}$ and let $c_1, c_2 \geq 0$ with $(c_1, c_2) \neq (0, 0)$. If $g = c_1 f_1 + c_2 f_2$, then $g \in \mathcal{F}$. Furthermore, it holds that

$$\frac{1}{\omega_g} = \frac{c_1}{\omega_{f_1}} + \frac{c_2}{\omega_{f_2}}.$$

Since \mathcal{F} is closed under addition, more generally, any finite sum $\sum_{j=1}^M c_j f_j \in \mathcal{F}$ for all $c_j \geq 0$ with $(c_1, \dots, c_M) \neq \mathbf{0}$ and $f_j \in \mathcal{F}$.

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269




Figure 2: Iterated dynamics of the tanh activation under randomly varying gains. For each panel, we fix a sequence (a_1, \dots, a_{1000}) sampled i.i.d. from $\mathcal{U}[1 - \delta, 1 + \delta]$ and iterate $x_{n+1} = \tanh(a_{n+1} x_n)$ from 60 initial points $x_0 \in [-1, 1]$. Panels (a)–(d) correspond to $\delta = 0.0001, 0.5, 1, 3$.

4.2 PROPOSED WEIGHT INITIALIZATION METHOD

We build on the theoretical results in Section 4.1 to design an initialization scheme that keeps forward signals well dispersed and avoids saturation even in very deep networks. We then compare the proposed initialization with standard Gaussian i.i.d. schemes (e.g., Xavier, He, EOC) from both forward and backward pass perspectives, studying signal and gradient propagation through theoretical analysis and experiments.

Proposed Weight Initialization Consider $f \in \mathcal{F}$ with $\omega := 1/f'(0) > 0$ and a target depth L . We initialize each layer as $\mathbf{W}^\ell = \mathbf{D}^\ell + \mathbf{Z}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$, where $(\mathbf{D}^\ell)_{ij} = \omega$ if $i \equiv j \pmod{N_{\ell-1}}$ and 0 otherwise. The noise matrix \mathbf{Z}^ℓ is sampled with noise scale σ_z as

$$(\mathbf{Z}^\ell)_{ij} \sim \mathcal{N}\left(0, \sigma_z^2 / N_{\ell-1}\right).$$

The value of σ_z is determined via a calibrated scalar surrogate model. Let $p_{\text{real}} = 0.4$ (see Section B.7) be the desired negative rate at depth L in the full network, and let $p_{\text{sur}}(L)$ denote the surrogate target used in the scalar model. As shown in Appendix B.4, for shallow depths the optimal surrogate target remains close to p_{real} , whereas for larger depths it decays approximately exponentially. In particular, a least-squares fit for $L \geq 10$ yields

$$p_{\text{sur}}(L) \approx 2.05 e^{-0.133L}.$$

In practice we therefore set

$$p(L) := \begin{cases} p_{\text{real}}, & L \leq L_{th}, \\ p_{\text{sur}}(L), & L > L_{th}, \end{cases}$$

with $L_{th} = 10$, and use this $p(L)$ in the closed-form calibration

$$\sigma^*(p(L), L, \omega) = -\frac{\omega}{\Phi^{-1}\left(\frac{1 - (1 - 2p(L))^{1/L}}{2}\right)},$$

where Φ denotes the standard normal cumulative distribution function. We then set $\sigma_z := \sigma^*(p(L), L, \omega)$.

For the learning rate, we use a band proportional to ω ; for Adam (Kingma & Ba, 2014),

$$\eta \in [10^{-5}\omega, 10^{-3}\omega].$$

We discuss the learning-rate range η and the surrogate calibration in Section 5.2 and Appendix B.4.

Derivation of the initialization. Using the elementwise formulation from Section 4.1, the usual matrix–vector product $\mathbf{W}^\ell \mathbf{x}^{\ell-1}$ can be rewritten as

$$x_i^{\ell+1} = f(a_i^{\ell+1} x_i^\ell), \quad a_i^{\ell+1} = w_{ii}^{\ell+1} + \sum_{j \neq i} \frac{w_{ij}^{\ell+1} x_j^\ell}{x_i^\ell}, \quad (2)$$

so that $a_i^{\ell+1}$ plays the role of an effective gain for neuron i in layer $\ell + 1$. Under our proposed initialization, Lemma 4.5 shows that $a_i^{\ell+1}$ is approximately Gaussian with mean ω and a data-dependent variance.

Lemma 4.5. *Using the elementwise formulation in equation 1 and employing the proposed weight initialization, fix an arbitrary layer ℓ and index i such that $x_i^\ell \neq 0$. Then, conditionally on x^ℓ ,*

$$a_i^{\ell+1} \sim \mathcal{N}\left(\omega, \frac{\sigma_z^2}{N_\ell} \left(1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2\right)\right). \quad (3)$$

Moreover, if $|x_j^\ell| \leq M$ for all j and $|x_i^\ell| \geq \varepsilon > 0$, then

$$\frac{\sigma_z^2}{N_\ell} \leq \text{Var}(a_i^{\ell+1} | x^\ell) \leq \sigma_z^2 \frac{M^2}{\varepsilon^2}.$$

The distribution of the effective gain a_i^ℓ is crucial for understanding signal propagation. We first analyze its mean, denoted by μ_a . This mean is determined by the diagonal entries of the proposed initialization matrix \mathbf{D} . In particular, we are interested in the supercritical regime $\mu_a > \omega$, which is analyzed in Theorem 4.6 below.

Theorem 4.6. *Let $f \in \mathcal{F}$ be an odd–sigmoid activation with $\omega := 1/f'(0)$, and fix any $\varepsilon > 0$. Consider the feedforward network and proposed initialization, except that the diagonal element is set to $a_0 := \omega + \varepsilon$, and let $a_i^{\ell+1}$ be the effective gain defined in equation 1. Fix a tolerance $\gamma \in (0, 1)$ and a finite depth $L \in \mathbb{N}$. Then there exist a threshold depth $L_0 \leq L$ and a noise threshold $\sigma_0 > 0$ such that, for all $0 < \sigma_z \leq \sigma_0$,*

$$\mathbb{P}\left((1-\gamma)\sigma_z^2 \leq \text{Var}(a_i^{\ell+1} | x^\ell) \leq (1+\gamma)\sigma_z^2 \text{ for all } L_0 \leq \ell < L, 1 \leq i \leq N_\ell\right) \geq 1-\gamma. \quad (4)$$

The data dependency of the variance term can be characterized as follows. For each layer ℓ and neuron i with $x_i^\ell(\sigma_z) \neq 0$, define $R_\ell(\sigma_z; i) := \frac{\|\mathbf{x}^\ell(\sigma_z)\|_2^2}{N_\ell (x_i^\ell(\sigma_z))^2}$. Here, $x_i^\ell(\sigma_z)$ denotes the activation at layer ℓ under noise scale σ_z . Theorem 4.6 implies that when the diagonal mean $\mu_a > \omega$, there exist a depth threshold L_0 and a noise threshold $\sigma_0 > 0$ such that $R_\ell(\sigma_z; i) \approx 1$ for all $\ell \geq L_0$ and $\sigma_z \leq \sigma_0$. This means that all neurons in deep layers converge to nearly identical values, indicating saturation of activations and loss of data dependence (Figure 8). Similarly, for $\mu_a < \omega$, Theorem 4.6 yields the same qualitative behavior. Deeper networks require smaller initialization variance to remain trainable (Poole et al., 2016; Schoenholz et al., 2016). Motivated by this, we set the diagonal mean to the critical value $a_0 = \omega$ when initializing deep networks.

Proposition A.2 and Corollary A.3 suggest that, for most gains a_i^ℓ across layers, if the variance of the gains is either too small or too large, the activations tend to saturate. We define the negative rate at depth L as $\pi_L(\sigma) := \mathbb{P}(x_L < 0 | x_0 > 0)$, i.e., the probability that a neuron that starts with a positive activation ends up with a negative sign at layer L . Under the surrogate model introduced below, it suffices to track a single neuron with $x_i^1 > 0$, since for any $f \in \mathcal{F}$ the activation is sign-preserving and odd. Consequently, the probability that a positive entry becomes negative at layer L is equal to the probability that a negative entry becomes positive, and does not depend on the particular choice of x_i^1 .

To handle this dependence on the current activations, we approximate the gain distribution by a scalar surrogate model, $a_i^{\ell+1} \sim \mathcal{N}(\omega, \sigma_z^2)$, and estimate the negative rate from this Gaussian approximation. Our initialization has a nonzero mean gain ω , which makes sign changes along a given coordinate well defined across layers. We interpret frequent sign flips during forward propagation

324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377




Figure 3: Last layer activation histograms for tanh networks with depth $L = 1000$ and varying width under the proposed initialization (**top row**) and the EOC initialization (**bottom row**). Each column corresponds to a different hidden width.

as a form of information loss and therefore use the surrogate calibration to control the negative rate at the final layer. In practice, we set the noise level so that the empirical negative rate at depth L is close to $p_{\text{real}} = 0.4$, preserving most sign information while still retaining sufficient randomness for feature learning. The next theorem shows how, given a target negative rate at depth L , we can compute the corresponding noise scale σ^* in this surrogate model.

Theorem 4.7. Fix a target $p \in [0, \frac{1}{2}]$, a depth $\ell \in \mathbb{N}$, and $\omega > 0$. There exists a unique $\sigma^* = \sigma^*(p, \ell, \omega) > 0$ such that $\pi_\ell(\sigma^*) = p$, and it is given by

$$\sigma^*(p, \ell, \omega) = - \frac{\omega}{\Phi^{-1}\left(\frac{1 - (1 - 2p)^{1/\ell}}{2}\right)}. \tag{5}$$

Figure 10 shows how the negative rate varies with L and p . As shown in Figure 11, for the proposed initialization χ_ℓ at the $\sigma^*(p, L, \omega)$ stays within a few percent of 1 for both $p = 0.01$ and $p = 0.49$ over depths up to 2×10^5 , suggesting that the calibration preserves trainable gradient scales. The relationship between the FFNN-level negative rate and our scalar surrogate, as well as the corresponding validation experiments, is detailed in Appendix B.4.

4.3 COMPARATIVE ANALYSIS OF GAUSSIAN AND PROPOSED INITIALIZATIONS

Gaussian i.i.d. schemes such as Xavier, He, and EOC all draw weights independently from $\mathcal{N}(0, \sigma_w^2/N_\ell)$, differing only in the choice of σ_w . In this section, we contrast our method with EOC, examining their forward and backward signal propagation both theoretically and empirically.

Forward Pass. As shown in Figures 19, 20, and 21, our proposed initialization keeps the activations well dispersed even at depth 10,000, whereas the EOC initialization drives them toward saturation near zero. Moreover, Figures 3, 22, and 23 show that, for EOC, signal propagation degrades as the width decreases, reflecting that this initialization is derived under an infinite width (or sufficiently wide) assumption. In contrast, our initialization is obtained by directly controlling the effective gain a_i^ℓ and does not rely on any large width approximation, which leads to stable forward propagation even in deep and relatively narrow networks. The theoretical motivation for forward signal propagation with odd-sigmoid activations is discussed in Sections 4.1 and 4.2.

Backward Pass. Although our initialization is derived without mean field assumptions, for the backward pass we adopt the standard mean field framework as an analytical tool to study gradient propagation. Let $\mathbf{g}^\ell = \partial\mathcal{L}/\partial\mathbf{x}^\ell$ denote the gradient at depth ℓ for a scalar loss \mathcal{L} . Under the assumptions, the backpropagated gradients satisfy the form $\frac{1}{N_\ell} \mathbb{E}\|\mathbf{g}^\ell\|_2^2 \approx \chi_{\ell+1} \frac{1}{N_{\ell+1}} \mathbb{E}\|\mathbf{g}^{\ell+1}\|_2^2$, where $\chi_{\ell+1}$ is the average gradient amplification factor of layer $\ell + 1$. Values $\chi_{\ell+1} \ll 1$ correspond to vanishing gradients, whereas $\chi_{\ell+1} \gg 1$ lead to exploding gradients; thus, keeping $\chi_{\ell+1}$ close to 1 across layers is essential for stable training in very deep networks. For our proposed initialization we show in Appendix A.3 that $\chi_{\ell+1}$ can be written as $\chi_{\ell+1} \approx (\omega^2 + \sigma_z^2) \mathbb{E}[f'(\mathbf{h}^{\ell+1})^2]$. For

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431




Figure 5: Heatmaps of validation accuracy on MNIST as a function of depth L and noise scale σ for three activations $\tanh(x)$, $0.1 \tanh(x)$, and $10 \tanh(x)$. Each model is a FFNN with width 128 trained with Adam. Each cell shows the mean validation accuracy over 3 runs. **(Top row)** Proposed. **(Bottom row)** EOC. Darker colors indicate higher validation accuracy.

comparison, under a standard i.i.d. Gaussian initialization, the corresponding factor takes the form $\chi_{\ell+1}^* \approx \sigma_w^2 \mathbb{E}[f'(\mathbf{h}^{\ell+1})^2]$.

For our proposed initialization, as σ_z grows, $\mathbb{E}[f'(\mathbf{h}^{\ell+1})^2]$ decreases, so the product stays close to 1 over a relatively wide range of σ_z . In contrast, for standard Gaussian i.i.d. initialization $\chi_{\ell+1}^*(\sigma_w) \approx \sigma_w^2 \mathbb{E}[f'(\mathbf{h}^{\ell+1})^2]$ has no analogous ω^2 term and is much more sensitive to perturbations of σ_w around its critical value. Therefore, $\chi_{\ell+1}$ remains close to 1 over a wide range of noise scales σ_z , leading to much more robust training (Figure 24). Consistent with this analysis, Figure 5 indicates that our initialization supports accurate training over a much broader region in (L, σ) and across activation scales, whereas Gaussian initializations remain effective only in a narrow band of σ and at relatively shallow depths. Figure 4 shows the gradient norm for the quadratic loss $L = \frac{1}{2} \|\mathbf{y}\|_2^2$ as a function of depth. Under our proposed initialization, the gradient norm is effectively preserved over very deep networks and remains stable across different widths. In contrast, EOC and Gaussian initializations exhibit rapid gradient decay at moderate widths, consistent with their reliance on infinite width assumptions, whereas our scheme maintains well scaled forward activations and backward gradients even in finite width regimes.




Figure 4: Layerwise gradient norms at initialization for a tanh FFNN of depth $L = 1000$. We set $\delta^\ell := \partial L / \partial \mathbf{h}^\ell$.

5 EXPERIMENTS

In this section, we evaluate the proposed initialization method. Section 5.1 studies widely used activations with $\omega \approx 1$, comparing our method against Gaussian initializations (Xavier, He, EOC) in terms of data efficiency and the dependence on network width and depth. Section 5.2 then turns to activations with $\omega \not\approx 1$, where we analyze performance, the ability to train effectively without batch normalization, learning-rate sensitivity, and the trainability of networks across a wide range

Table 1: Mean of the best validation accuracies within 50 epochs over 10 independent runs for a 50 layer FFNN (512 units) on MNIST and Fashion MNIST, trained on 100 or 500 sample subsets.

Dataset	Method	tanh(x)		erf(x)		arctan(x)		gd(x)		softsign ₂ (x)		softsign ₁ (x) + softsign ₂ (x)	
		100	500	100	500	100	500	100	500	100	500	100	500
MNIST	Xavier	64.00	79.85	66.15	84.63	64.58	81.30	60.87	84.18	60.52	81.38	32.63	53.97
	He	41.65	72.55	21.05	48.37	51.40	77.68	42.50	72.75	48.35	76.18	11.35	10.02
	EOC	59.83	82.92	64.58	82.97	64.52	83.73	62.30	84.05	66.57	82.48	59.60	71.57
	Proposed	68.23	86.75	66.53	87.13	67.63	86.82	66.38	86.23	69.51	86.43	65.65	84.02
FMNIST	Xavier	67.65	74.53	68.92	76.13	66.75	73.78	67.10	72.65	66.85	74.47	49.38	69.88
	He	61.13	74.60	44.60	65.25	62.92	76.75	62.88	76.55	64.97	75.22	11.50	11.70
	EOC	67.22	76.87	66.85	73.83	66.93	76.45	67.00	77.63	67.08	75.67	65.48	73.07
	Proposed	70.67	77.43	70.63	78.17	71.55	77.92	68.62	77.80	68.17	78.33	67.93	76.43

of activation scales, including experiments on physicsinformed neural networks (PINNs). The experimental setting is described in Appendix C.

5.1 EXPERIMENTS WITH $\omega \approx 1$ ACTIVATIONS

Dataset Efficiency. Table 1 presents the best validation accuracy within 50 epochs for a 50 layer FFNN (512 units), trained on small subsets of size 100 or 500. We compare four initializations across odd-sigmoid activations $f \in \mathcal{F}$. The proposed initialization attains the top accuracy in all settings, across both datasets and both sample regimes, with the improvements most pronounced at 100 samples. These results indicate that our method is data efficient, achieving higher validation accuracy with limited data.

Network Size Independence. We assess how independent the proposed initialization is of depth and width in networks. Figures 25, 26, 27, and 28 show validation accuracy versus depth for FFNNs (width 64) on MNIST, Fashion MNIST, CIFAR-10, and CIFAR-100, using odd sigmoid activations and four initializations. Each panel fixes one activation and shows the best validation accuracy over 10 epochs for depths $L \in \{20, 50, 100, 150, 200\}$. Across all datasets, activations, and depths, the proposed initialization consistently achieves the highest validation accuracy among all Gaussian initializations. Figures 6 and 32 further examine the effect of width in deep networks and show that standard Gaussian initializations struggle when the network is either too narrow or too wide, whereas the proposed scheme maintains strong performance over a broad range of widths.

5.2 EXPERIMENTS WITH ACTIVATIONS FAR FROM $\omega \approx 1$

In this section we investigate activations $f, g \in \mathcal{F}$ whose $\omega = 1/f'(0)$ is not close to 1, including rescaled variants $\alpha f(x)$, input scaled variants $f(\alpha x)$, and positive linear combinations $\alpha f(x) + \beta g(x)$ with $\alpha, \beta > 0$.

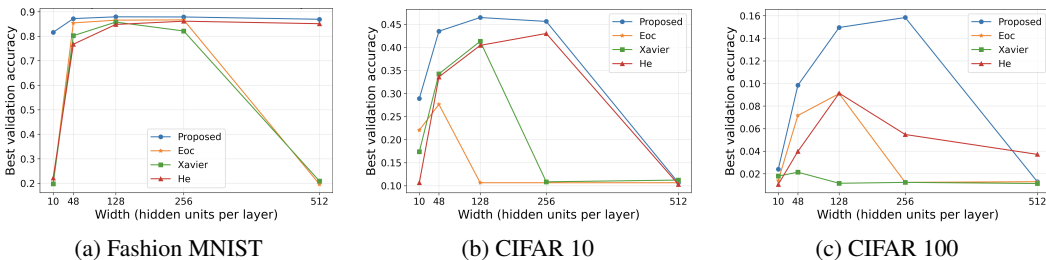


Figure 6: Best validation accuracy versus width $\{10, 48, 128, 150, 200, 512\}$ for a 100 layer tanh FFNN. Each curve shows the Proposed, EOC, Xavier, and He initialization schemes, and each point corresponds to the best validation accuracy over 20 training epochs.




Figure 7: Best PINN loss versus activation scale for the Black–Scholes PINN (depth 50, width 64), comparing the proposed and EOC initializations. **(Left)** $a \tanh(x)$ with $a \in \{0.1, 0.3, 0.5, 0.8, 1.0, 1.3, 1.5, 1.8, 2.0\}$. **(Right)** $b \text{softsign}_2(x)$ with the same set of scales.

Batch Normalization Free Training. For

$$f(x) = \tanh(ax) + \text{erf}(bx) + \frac{x}{1 + |cx|} + \text{gd}(dx), \quad (6)$$

by Corollary 4.4 this f is an odd–sigmoid activation. Empirical results for various (a, b, c, d) appear in Figure 29. Across all settings, the proposed initialization achieves the highest validation accuracy even without batch normalization, outperforming Gaussian i.i.d. initializations both with and without batch normalization. We further investigate the effect of applying batch normalization, including training 800 layer networks, in Figures 30 and 31.

Learnable Learning Rate. To study optimization stability, we plot learning rate versus validation accuracy curves on MNIST and Fashion MNIST for 20 layer, width 512 FFNNs with activations $f(x) = \frac{2}{\pi} \arctan(ax)$ and $f(x) = \tanh(ax)$, using scales $\alpha \in \{10^2, 10^1, 1, 10^{-1}, 10^{-2}, 10^{-3}\}$ (Figures 33, 34, 35, and 36). The proposed method consistently yields a learnable, typically wider LR window across a broad range of α , demonstrating that training remains feasible across diverse ω scales. If the learning rate is chosen too large or too small relative to the ω -scaled band, the first parameter update either destroys the calibrated noise spread or becomes negligible, so that training effectively fails from the second pass onward. Motivated by this, we use the practical learning rate range

$$\eta \in [10^{-5} \omega, 10^{-3} \omega].$$

Scale Preserving Odd–Sigmoid Activations. For $f \in \mathcal{F}$, we investigate whether the proposed and EOC initializations preserve the activation range of the scaled activation αf (Figure 37). Unlike EOC, the proposed initialization maintains an activation range proportional to α even as the network depth increases. Networks with $\alpha \tanh$, $\alpha \arctan$, and $\alpha \text{softsign}_2$ activations are trained on MNIST and Fashion MNIST for α ranging from 10^{-2} to 10^9 , and the proposed initialization enables stable training across all α scales for every activation (Table 2–7). Since an α scaled activation directly controls the scale of the output y_i , we also evaluate our initialization on regression problems using physics informed neural networks (PINNs). As shown in Figures 7 and 38, these results indicate that scaled odd–sigmoid activations are well suited for regression tasks, where controlling the output range is important. We defer detailed PINN setups and PDE formulations to Appendix C.4

6 CONCLUSION

We introduced an activation aware initialization scheme for FFNNs with odd–sigmoid activations that does not rely on infinite or very wide width assumptions. Unlike standard Gaussian i.i.d. initializations, the proposed method is more robust to variance choice and maintains stable signal and gradient propagation even in very deep and relatively narrow networks. In physics informed neural networks, we further observed that appropriately scaled odd–sigmoid activations, combined with our initialization, can achieve lower PINN losses than standard choices. These results suggest that, with a suitable initialization, scaled odd–sigmoid activations can be used as practical design knobs to improve performance beyond what is possible with conventional Gaussian initializations.

540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

ETHICS STATEMENT

This research adheres to the ICLR Code of Ethics. A large language model was used only to refine the writing and assist in preparing visualizations; all core ideas, methods, and results were developed independently by the authors. We adhere to the highest standards of research integrity and transparency as set out in the ICLR guidelines.

REFERENCES

- Behtom Adeli, John P. McLinden, Pankaj Pandey, Ming Shao, and Yalda Shahriari. Toward improving fnirs classification: A study on activation functions in deep neural architectures. *arXiv preprint*, 2025.
- Daniel Alcaide, Marta Gonzalez, Juan Ramirez, and Luis Ortega. Telu activation function for fast and stable deep learning. *Information Sciences*, 690:120345, 2025.
- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Henry Mao, Gary Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI 2021)*, pp. 1352–1361. PMLR, 2021.
- Garrett Bingham and Risto Miikkulainen. Efficient activation function optimization through surrogate modeling. *Advances in Neural Information Processing Systems*, 36:6634–6661, 2023.
- Yaniv Blumenfeld, Dar Gilboa, and Daniel Soudry. Beyond signal propagation: Is feature diversity necessary in deep neural network initialization? In *Proceedings of the 37th International Conference on Machine Learning*, pp. PMLR 119, 2020. PMLR, 2020.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- Soufiane Hayou, Arnaud Doucet, and Judith Rousseau. On the impact of the activation function on deep neural networks training. In *International conference on machine learning*, pp. 2672–2680. PMLR, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016.
- Pieter-Jan Hoedt and G"unter Klambauer. Principled weight initialisation for input-convex neural networks. In *37th Conference on Neural Information Processing Systems (NeurIPS 2023)*, 2023.
- Jiyoung Hwang, Byeongjoon Noh, Zhixiong Jin, and Hwasoo Yeo. Asymmetric long-term graph multi-attention network for traffic speed prediction. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1498–1503. IEEE, 2022.
- Sergey Ioffe. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Gaurav Iyer, Boris Hanin, and David Rolnick. Maximal initial learning rates in deep relu networks. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 1–30. PMLR, 2023.
- Ameya D. Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Adaptive activation functions accelerate convergence in deep and physics-informed neural networks. *Journal of Computational Physics*, 404:109136, 2020.

- 594 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
595 *arXiv:1412.6980*, 2014.
- 596
597 Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing
598 neural networks. *Advances in neural information processing systems*, 30, 2017.
- 599 Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444,
600 2015.
- 601
602 Hyunwoo Lee, Yunho Kim, Seung Yeop Yang, and Hayoung Choi. Improved weight initialization
603 for deep and narrow feedforward neural network. *Neural Networks*, 176:106362, 2024.
- 604 Hyunwoo Lee, Hayoung Choi, and Hyunju Kim. Robust weight initialization for tanh neural net-
605 works with fixed point analysis. In *The Thirteenth International Conference on Learning Repre-*
606 *sentations*, 2025.
- 607
608 Xun Liu. A novel zeroing neural network with activation function-enhanced convergence for effi-
609 cient data matrix decomposition. *Neurocomputing*, 616:128391, 2025.
- 610 Diganta Misra. Mish: A self regularized non-monotonic activation function. *arXiv preprint*
611 *arXiv:1908.08681*, 2019.
- 612
613 Michael Murray, Vinayak Abrol, and Jared Tanner. Activation function design for deep networks:
614 linearity and effective initialisation. *Applied and Computational Harmonic Analysis*, 59:117–154,
615 2022.
- 616
617 Meenal V. Narkhede, Prashant P. Bartakke, and Mukul S. Sutaone. A review on weight initialization
618 strategies for neural networks. *Artificial Intelligence Review*, 55:291–322, 2022.
- 619
620 Lorenzo Noci, Luca Biggio, Sotiris Anagnostidis, Antonio Orvieto, Sidak Pal Singh, and Aurelien
621 Lucchi. Signal propagation in transformers: Theoretical perspectives and the role of rank collapse.
622 In *36th Conference on Neural Information Processing Systems (NeurIPS 2022)*, 2022.
- 623
624 Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponen-
625 tial expressivity in deep neural networks through transient chaos. *Advances in neural information*
626 *processing systems*, 29, 2016.
- 627
628 Ilan Price, Nicholas Daultry Ball, Samuel C.H. Lam, Adam C. Jones, and Jared Tanner. Deep neural
629 network initialization with sparsity inducing activations. *Published as a conference paper at ICLR*
630 *2024*, 2024.
- 631
632 Junfei Qiao, Sanyi Li, and Wenjing Li. Mutual information based weight initialization method for
633 sigmoidal feedforward neural networks. *Neurocomputing*, 207:676–683, 2016. doi: 10.1016/j.
634 neucom.2016.05.054.
- 635
636 Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the ex-
637 pressive power of deep neural networks. In *International Conference on Machine Learning*, pp.
638 2847–2854, 2017.
- 639
640 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A
641 deep learning framework for solving forward and inverse problems involving nonlinear partial
642 differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- 643
644 Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv*
645 *preprint arXiv:1710.05941*, 2017.
- 646
647 Samuel S Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha Sohl-Dickstein. Deep information
648 propagation. *arXiv preprint arXiv:1611.01232*, 2016.
- 649
650 Maciej Skorski, Alessandro Temperoni, and Martin Theobald. Revisiting weight initialization of
651 deep neural networks. *Proceedings of Machine Learning Research*, 157, 2021.
- 652
653 Sartaj Singh Sodhi, Pravin Chandra, and Sharad Tanwar. A new weight initialization method for sig-
654 moidal feedforward artificial neural networks. In *2014 International Joint Conference on Neural*
655 *Networks (IJCNN)*, pp. 291–298. IEEE, 2014.

648 Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initial-
649 ization and momentum in deep learning. In *Proceedings of the 30th International Conference on*
650 *Machine Learning*, JMLR: W&CP volume 28, pp. –, Atlanta, Georgia, USA, 2013. JMLR.org.
651
652 Bekir Tolga Tutuncuoglu. Neuron-level activation learning in neural networks: A self-configuring
653 ai approach. *SSRN*, 2025.
654
655 Abaid Ullah, Muhammad Imran, Muhammad Abdul Basit, Madeeha Tahir, and Jihad Younis. Aher-
656 656 relu: A novel adaptive activation function enhancing deep neural network performance. *Com-
657 plexity*, 2025:8233876, 2025.
658
659 Honghui Wang, Lu Lu, Shiji Song, and Gao Huang. Learning specialized activation functions for
660 physics-informed neural networks. *arXiv preprint*, 2023.
661
662 Ahmet Yilmaz and Riccardo Poli. Successfully and efficiently training deep multi-layer perceptrons
663 with logistic activation function simply requires initializing the weights with an appropriate neg-
664 ative mean. *Neural Networks*, 153:87 – 103, 2022. URL [https://doi.org/10.1016/j.
665 neunet.2022.05.030](https://doi.org/10.1016/j.neunet.2022.05.030).
666
667 Shijun Zhang, Jianfeng Lu, and Hongkai Zhao. Deep network approximation: Beyond relu to
668 diverse activation functions. *Journal of Machine Learning Research*, 25(35):1–39, 2024.
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 SUPPLEMENTARY MATERIAL
703

704 The supplementary material is structured as follows.
705

- 706 • Appendix A provides the theoretical results and their proofs.
- 707 • Appendix B contains additional simulation results without neural networks.
- 708 • Appendix C reports extra experimental details and results with neural networks.

710
711 A THEORETICAL RESULTS
712

713 In this section, we provide proofs for the statements presented in Section 4.
714

715 A.1 THEORETICAL RESULTS FOR SECTION 4.1
716

717 **Lemma A.1.** *If $f \in \mathcal{F}$, then the following holds.*

- 718 (i) $f(0) = 0$.
- 719 (ii) $0 = \arg \max |f'(x)|$.
- 720 (iii) $\lim_{x \rightarrow \infty} f'(x) = 0$.

721
722
723
724 *Proof.* (i) It is trivial that odd symmetry implies $f(0) = 0$. (ii) Since f' is an even function, by
725 Definition 4.1(v) $f'(x)$ has the maximum value at $x = 0$. (iii) Since f' is strictly decreasing on
726 $[0, \infty)$ and $f'(x) > 0$ for all $x \in \mathbb{R}$, the Monotone Convergence Theorem for real functions implies
727 that $\ell := \lim_{x \rightarrow \infty} f'(x)$ exists for some $\ell \in [0, \infty)$. Suppose that $\ell > 0$. Then there exists $R > 0$
728 such that

$$729 f'(x) \geq \frac{\ell}{2} \quad \text{for all } x \geq R.$$

730 Integrating from R to x , we obtain
731

$$732 f(x) - f(R) = \int_R^x f'(t) dt \geq \frac{\ell}{2}(x - R) \quad \text{for all } x > R.$$

733 Hence $f(x) \rightarrow \infty$ as $x \rightarrow \infty$, which contradicts the boundedness of f . Therefore $\ell = 0$. \square
734
735
736

737 **Lemma A.1 summarizes the basic regularity and saturation properties of odd-sigmoid activations**
738 **in \mathcal{F} . In particular, it shows that every $f \in \mathcal{F}$ has a unique global slope maximum at the origin**
739 **and becomes arbitrarily flat in the tails. These simple but structural features will be crucial for**
740 **understanding how the gain parameter a reshapes the fixed-point structure of the scalar map $x \mapsto$**
741 **$f(ax)$.**
742

743 **Proposition 4.2** Suppose $f \in \mathcal{F}$ with $\omega := 1/f'(0)$, and for a fixed $a > 0$ define $\phi_a(x) := f(ax)$.
744 Then
745

- 746 (i) If $0 < a \leq \omega$, then $f(ax)$ has a unique fixed point $x^* = 0$.
- 747 (ii) If $a > \omega$, then $f(ax)$ has three distinct fixed points: $x^* = -\xi_a, 0, \xi_a$ such that $\xi_a > 0$.

748
749 *Proof.* For $a > 0$, consider $g(x, a) := f(ax) - x$. We have $g(0, a) = 0$ and $g'(x, a) = a f'(ax) - 1$.
750

751 Case (i): $0 < a \leq \omega$. For $x > 0$ we have $ax > 0$, and $f'(ax) < f'(0)$; hence
752

$$753 g'(x, a) = a f'(ax) - 1 < a f'(0) - 1 \leq 0,$$

754 for all $x > 0$. Thus $g(\cdot, a)$ is strictly decreasing on $(0, \infty)$, $g(0, a) = 0$, and $\lim_{x \rightarrow \infty} g(x, a) =$
755 $L - x = -\infty$, so $g(x, a) < 0$ for all $x > 0$ and there is no positive root. Since $g(\cdot, a)$ is odd, there
is no negative root either. Hence $x = 0$ is the unique solution.

756 Case (ii): $a > \omega$. Then $g'(0, a) = af'(0) - 1 > 0$. $g'(\cdot, a)$ is strictly decreasing on $[0, \infty)$, and by
 757 the Lemma A.1

$$758 \lim_{x \rightarrow \infty} g'(x, a) = a \lim_{x \rightarrow \infty} f'(ax) - 1 = -1 < 0.$$

759
 760 By the intermediate value theorem, there exists a unique $\hat{x} > 0$ with $g'(\hat{x}, a) = 0$. Hence $g(\cdot, a)$ is
 761 strictly increasing on $[0, \hat{x}]$ and strictly decreasing on $[\hat{x}, \infty)$. Since $g(0, a) = 0$ and $g'(0, a) > 0$,
 762 we have $g(x, a) > 0$ for all $0 < x \leq \hat{x}$. Using (3), $\lim_{x \rightarrow \infty} g(x, a) = -\infty$; because g is strictly
 763 decreasing on $[\hat{x}, \infty)$ and continuous, there exists a unique $\xi_a > \hat{x}$ with $g(\xi_a, a) = 0$. Thus, there is
 764 exactly one positive, nonzero root. Oddness of g yields the symmetric negative root $-\xi_a$, so the full
 765 set of real solutions is $\{-\xi_a, 0, \xi_a\}$. \square

766 **Proposition 4.2** provides a precise characterization of how the fixed points of the scalar map $x \mapsto$
 767 $f(ax)$ bifurcate as the gain a crosses the critical value $\omega = 1/f'(0)$. For $a \leq \omega$ the origin is the
 768 only fixed point, while for $a > \omega$ a symmetric nonzero pair $\pm\xi_a$ emerges, exhibiting a classical
 769 pitchfork structure. We next lift this static picture to the dynamical setting, showing that the iterates
 770 of $x_{n+1} = f(ax_n)$ converge to the corresponding fixed point for every positive initial condition.

771
 772 **Theorem 4.3** Suppose $f \in \mathcal{F}$ with $\omega := 1/f'(0)$, and for a fixed $a > 0$ define

$$773 x_0 > 0, \quad x_{n+1} = \phi_a(x_n), \quad n = 0, 1, 2, \dots$$

774 Then the sequence $\{x_n\}$ converges for every $x_0 > 0$. Furthermore,

- 775
 776
 777 (1) if $0 < a \leq \omega$, then $x_n \rightarrow 0$ as $n \rightarrow \infty$.
 778 (2) if $a > \omega$, then $x_n \rightarrow \xi_a$ as $n \rightarrow \infty$.

779
 780 *Proof.* (1) Since f is odd and f' is strictly decreasing on $[0, \infty)$, we have $0 < f'(x) < f'(0)$ for
 781 all $x \neq 0$; hence, for any $a \in (0, \omega)$ and any $x_n > 0$, it follows that $x_{n+1} = f(ax_n) < x_n$ for all
 782 $n \in \mathbb{N}$. By the monotone convergence theorem, it converges to the fixed point $x^* = 0$.

783 (2) Let $x_0 < \xi_a$. Since $\phi'(x)$ is decreasing for $x \geq 0$, with ξ_a is the unique fixed point for $x > 0$,
 784 it holds that $x_n < x_{n+1} < \xi_a$ for all $n \in \mathbb{N}$. Thus, by the monotone convergence theorem, the
 785 sequence converges to the fixed point ξ_a . The proof is similar when $x_0 > \xi_a$. By the monotone
 786 convergence theorem, the sequence also converges to the fixed point ξ_a . \square

787
 788
 789 **Corollary 4.4** Let $f_1, f_2 \in \mathcal{F}$ and let $c_1, c_2 \geq 0$ with $(c_1, c_2) \neq (0, 0)$. If $g = c_1f_1 + c_2f_2$, then
 790 $g \in \mathcal{F}$. Furthermore, it holds that

$$791 \frac{1}{\omega_g} = \frac{c_1}{\omega_{f_1}} + \frac{c_2}{\omega_{f_2}}.$$

792
 793
 794 *Proof.* Let $g = c_1f_1 + c_2f_2$. Since $c_1, c_2 \geq 0$, the linear combination preserves oddness and bounded
 795 saturation by linearity. Strict monotonicity on \mathbb{R} follows from $g'(x) = c_1f_1'(x) + c_2f_2'(x) > 0$,
 796 and slope decay on $[0, \infty)$ is preserved because a positive linear combination of strictly decreasing
 797 functions is strictly decreasing. Thus $g \in \mathcal{F}$. Evaluating at the origin gives

$$798 g'(0) = c_1f_1'(0) + c_2f_2'(0) = \frac{c_1}{\omega_{f_1}} + \frac{c_2}{\omega_{f_2}},$$

799 so by definition $1/\omega_g = g'(0)$, which yields the claimed identity. \square

800
 801
 802 In practice this means that we can build richer odd-sigmoid activations by mixing simpler ones
 803 without losing the pitchfork structure described above. We now move from the constant-gain setting
 804 to the more general case where the layerwise gains (a_n) are allowed to vary with depth.

Proposition A.2. Let $f \in \mathcal{F}$ and $\{a_n\}_{n=1}^{\infty}$ be a positive real sequence, i.e., $a_n > 0$ for all $n \in \mathbb{N}$, such that only finitely many elements are greater than $\omega = 1/f'(0)$. For a positive sequence $\{a_n\}_{n \geq 1}$, set

$$\Phi^m := \phi_{a_m} \circ \phi_{a_{m-1}} \circ \cdots \circ \phi_{a_1}.$$

Then for any $x \in \mathbb{R}$

$$\lim_{m \rightarrow \infty} \Phi^m(x) = 0.$$

Proof. Let $N := \max\{n : a_n > \omega\}$ (take $N = 0$ if the set is empty) and define

$$b_n = \begin{cases} a_n, & n \leq N, \\ 0, & n > N, \end{cases} \quad c_n = \begin{cases} a_n, & n \leq N, \\ \omega, & n > N. \end{cases}$$

Let $\hat{\Phi}^m := \phi_{b_m} \circ \cdots \circ \phi_{b_1}$ and $\tilde{\Phi}^m := \phi_{c_m} \circ \cdots \circ \phi_{c_1}$. For $x \geq 0$ and $n > N$, by definition of \mathcal{F} , we obtain

$$\hat{\Phi}^m(x) \leq \Phi^m(x) \leq \tilde{\Phi}^m(x).$$

Oddness yields the same in absolute value for all $x \in \mathbb{R}$:

$$|\hat{\Phi}^m(x)| \leq |\Phi^m(x)| \leq |\tilde{\Phi}^m(x)|.$$

By Proposition 4.3, $\hat{\Phi}^m(x) \rightarrow 0$ and $\tilde{\Phi}^m(x) \rightarrow 0$. The squeeze theorem gives $\Phi^m(x) \rightarrow 0$. \square

Corollary A.3. Let $\epsilon > 0$ be given and set $\omega := 1/f'(0)$. Suppose that $\{a_n\}_{n=1}^{\infty}$ be a positive real sequence such that only finitely many elements are lower than $\omega + \epsilon$. Then for any $x \in \mathbb{R} \setminus \{0\}$

$$\liminf_{m \rightarrow \infty} |\Phi^m(x)| \geq \xi_{\omega+\epsilon}.$$

Proof. Let $N := \max\{n : a_n < \omega + \epsilon\}$ (take $N = 0$ if the set is empty) and define

$$b_n = \begin{cases} a_n, & n \leq N, \\ \omega + \epsilon, & n > N. \end{cases} \quad \hat{\Phi}_m := \phi_{b_m} \circ \cdots \circ \phi_{b_1}, \quad \Phi^m := \phi_{a_m} \circ \cdots \circ \phi_{a_1},$$

where $\phi_a(x) := f(ax)$. By definition of \mathcal{F} ,

$$|\hat{\Phi}^m(x)| \leq |\Phi^m(x)| \quad (\forall x \in \mathbb{R}, \forall m).$$

Taking \liminf in the inequality yields

$$\liminf_{m \rightarrow \infty} |\Phi^m(x)| \geq \xi_{\omega+\epsilon}.$$

\square

Proposition A.2 and Corollary A.3 together describe two opposite extremes of how layerwise gains affect the one-dimensional dynamics. Roughly speaking, Proposition A.2 says that if, after some finite depth, all gains a_n stay below the critical value ω , then the composed map Φ^m always drives the signal back to zero, no matter what happened in the earlier layers. In contrast, Corollary A.3 shows that if the gains are eventually bounded away from ω by a fixed margin $\epsilon > 0$, then the compositions cannot collapse to zero: for any nonzero input, the iterates stay at least as large (in absolute value) as the positive fixed point $\xi_{\omega+\epsilon}$.

A.2 THEORETICAL RESULTS FOR SECTION 4.2

Lemma 4.5 Using the elementwise formulation in equation 1 and employing the proposed weight initialization, fix an arbitrary layer ℓ and index i such that $x_i^\ell \neq 0$. Then, conditionally on x^ℓ ,

$$a_i^{\ell+1} \sim \mathcal{N}\left(\omega, \frac{\sigma_z^2}{N_\ell} \left(1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2\right)\right).$$

Moreover, if $|x_j^\ell| \leq M$ for all j and $|x_i^\ell| \geq \varepsilon > 0$, then

$$\frac{\sigma_z^2}{N_\ell} \leq \text{Var}(a_i^{\ell+1} | x^\ell) \leq \sigma_z^2 \frac{M^2}{\varepsilon^2}.$$

Proof. From equation 1 and the proposed initialization, we write

$$W^{\ell+1} = D^{\ell+1} + Z^{\ell+1},$$

where the diagonal of $D^{\ell+1}$ equals ω and $Z^{\ell+1}$ has independent entries $(Z^{\ell+1})_{ij} \sim \mathcal{N}(0, \sigma_z^2/N_\ell)$. The pre-activation at coordinate i reads

$$s_i^{\ell+1} = \sum_{j=1}^{N_\ell} ((D^{\ell+1})_{ij} + (Z^{\ell+1})_{ij}) x_j^\ell = \omega x_i^\ell + (Z^{\ell+1})_{ii} x_i^\ell + \sum_{j \neq i} (Z^{\ell+1})_{ij} x_j^\ell.$$

On the event $x_i^\ell \neq 0$, we define the effective gain

$$a_i^{\ell+1} := \frac{s_i^{\ell+1}}{x_i^\ell} = \omega + (Z^{\ell+1})_{ii} + \sum_{j \neq i} (Z^{\ell+1})_{ij} \frac{x_j^\ell}{x_i^\ell}.$$

Conditionally on x^ℓ , the coefficients $\{x_j^\ell/x_i^\ell\}$ are deterministic, whereas the random variables $\{(Z^{\ell+1})_{ij}\}_{j=1}^{N_\ell}$ are independent, centered Gaussians with variance σ_z^2/N_ℓ . Therefore $a_i^{\ell+1} | x^\ell$ is a linear combination of independent Gaussians, hence Gaussian:

$$a_i^{\ell+1} | x^\ell \sim \mathcal{N}\left(\omega, \frac{\sigma_z^2}{N_\ell} \left(1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2\right)\right).$$

This yields the conditional variance

$$\text{Var}(a_i^{\ell+1} | x^\ell) = \frac{\sigma_z^2}{N_\ell} \left(1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2\right).$$

Since the summation term is nonnegative, we immediately obtain the conditional lower bound $\text{Var}(a_i^{\ell+1} | x^\ell) \geq \sigma_z^2/N_\ell$ and hence the unconditional lower bound $\text{Var}(a_i^{\ell+1}) \geq \sigma_z^2/N_\ell$.

For the upper bound, note that

$$1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2 = \frac{(x_i^\ell)^2 + \sum_{j \neq i} (x_j^\ell)^2}{(x_i^\ell)^2} = \frac{\|x^\ell\|_2^2}{(x_i^\ell)^2},$$

whence

$$\text{Var}(a_i^{\ell+1} | x^\ell) = \frac{\sigma_z^2}{N_\ell} \frac{\|x^\ell\|_2^2}{(x_i^\ell)^2}.$$

If $|x_j^\ell| \leq M$ for all j , then $\|x^\ell\|_2^2 = \sum_{j=1}^{N_\ell} (x_j^\ell)^2 \leq N_\ell M^2$, and if $|x_i^\ell| \geq \varepsilon > 0$, we obtain

$$\frac{\|x^\ell\|_2^2}{(x_i^\ell)^2} \leq \frac{N_\ell M^2}{\varepsilon^2},$$

which implies

$$\text{Var}(a_i^{\ell+1} | x^\ell) \leq \frac{\sigma_z^2}{N_\ell} \frac{N_\ell M^2}{\varepsilon^2} = \sigma_z^2 \frac{M^2}{\varepsilon^2}.$$

The unconditional upper bound follows from $\text{Var}(a_i^{\ell+1}) = \mathbb{E}[\text{Var}(a_i^{\ell+1} | x^\ell)] + \text{Var}(\mathbb{E}[a_i^{\ell+1} | x^\ell])$, and the fact that $\mathbb{E}[a_i^{\ell+1} | x^\ell] = \omega$ is constant and the first term is bounded by the conditional upper bound. \square

Lemma 4.5 shows that, under our structured initialization, the effective gain $a_i^{\ell+1}$ is approximately Gaussian with mean ω and a variance that scales like σ_z^2 times a data dependent factor. In particular, the lower bound $\text{Var}(a_i^{\ell+1} | x^\ell) \geq \sigma_z^2/N_\ell$ guarantees a nontrivial amount of gain noise at every layer, while the upper bound prevents the variance from blowing up when the activations remain bounded away from zero. These properties will be crucial when we study how the gain variance behaves as depth and noise scale vary.

Theorem 4.6 Let $f \in \mathcal{F}$ be an odd-sigmoid activation with $\omega := 1/f'(0)$, and fix any $\varepsilon > 0$. Consider the feedforward network and proposed initialization, except that the diagonal element is set to $a_0 := \omega + \varepsilon$, and let $a_i^{\ell+1}$ be the effective gain defined in equation 1. Fix a tolerance $\gamma \in (0, 1)$ and a finite depth $L \in \mathbb{N}$. Then there exist a threshold depth $L_0 \leq L$ and a noise threshold $\sigma_0 > 0$ such that, for all $0 < \sigma_z \leq \sigma_0$,

$$\mathbb{P}\left((1-\gamma)\sigma_z^2 \leq \text{Var}(a_i^{\ell+1} | x^\ell) \leq (1+\gamma)\sigma_z^2 \text{ for all } L_0 \leq \ell < L, 1 \leq i \leq N_\ell\right) \geq 1-\gamma. \quad (7)$$

Proof. We first consider $\sigma_z = 0$. In this case we set $\mathbf{W}^\ell = \mathbf{D}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ with $(\mathbf{D}^\ell)_{ij} = a_0$ if $i \equiv j \pmod{N_{\ell-1}}$ and 0 otherwise, so that the layerwise update reduces to

$$x^{\ell+1} = f(a_0 x^\ell).$$

Hence each coordinate x_i^ℓ evolves independently according to the scalar recurrence $x_{n+1} = f(a_0 x_n)$. Since $a_0 f'(0) > 1$, Theorem 4.3 implies that this map has three fixed points $\{0, \pm \xi_{a_0}\}$ and, for any $x_0 \neq 0$,

$$x_n \rightarrow \pm \xi_{a_0} \text{ as } n \rightarrow \infty,$$

with the sign determined by $\text{sign}(x_0)$. For every $\delta > 0$ and every initial value $x_0 \neq 0$ there exists an integer $N_i(\delta)$ such that

$$|x_n(x_0)| \in [\xi_{a_0} - \delta, \xi_{a_0} + \delta] \text{ for all } n \geq N_i(\delta).$$

Since each layer contains only finitely many neurons, we can take

$$L_0(\delta) := \max_{1 \leq i \leq N_\ell} N_i(\delta),$$

so that

$$|x_i^\ell(0)| \in [\xi_{a_0} - \delta, \xi_{a_0} + \delta] \text{ for all } 1 \leq i \leq N_\ell, \ell \geq L_0(\delta), \quad (8)$$

where $x^\ell(0)$ denotes the activations at depth ℓ in the noiseless case $\sigma_z = 0$.

From these bounds we obtain, for all $\ell \geq L_0(\delta)$,

$$N_\ell(\xi_{a_0} - \delta)^2 \leq \|x^\ell(0)\|_2^2 \leq N_\ell(\xi_{a_0} + \delta)^2, \quad (x_i^\ell(0))^2 \in [(\xi_{a_0} - \delta)^2, (\xi_{a_0} + \delta)^2].$$

Therefore the deterministic ratio defined in Lemma 4.5

$$R_\ell(0; i) := \frac{\|x^\ell(0)\|_2^2}{N_\ell(x_i^\ell(0))^2}$$

satisfies

$$R_\ell(0; i) \in \left[\frac{(\xi_{a_0} - \delta)^2}{(\xi_{a_0} + \delta)^2}, \frac{(\xi_{a_0} + \delta)^2}{(\xi_{a_0} - \delta)^2} \right] \text{ for all } \ell \geq L_0(\delta), i.$$

As $\delta \rightarrow 0$, this interval shrinks to 1. Hence, given any $\gamma \in (0, 1)$ we can choose $\delta > 0$ such that

$$|R_\ell(0; i) - 1| \leq \frac{\gamma}{2} \text{ for all } \ell \geq L_0(\delta), i.$$

Now we consider $\sigma_z > 0$. Recall that in our initialization we write

$$\mathbf{Z}^{(\ell)} = \frac{\sigma_z}{\sqrt{N_{\ell-1}}} \mathbf{G}^{(\ell)}, \quad (\mathbf{G}^{(\ell)})_{ij} \sim \mathcal{N}(0, 1) \text{ i.i.d.},$$

so that the randomness is entirely carried by the Gaussian matrices $\mathbf{G}^{(1)}, \dots, \mathbf{G}^{(L)}$. For any deterministic family of matrices $\widehat{\mathbf{G}}^{(1)}, \dots, \widehat{\mathbf{G}}^{(L)}$ we define, for $\sigma_z \geq 0$, the corresponding activations $\widehat{\mathbf{x}}^\ell(\sigma_z)$ recursively by

$$\begin{aligned} \widehat{\mathbf{x}}^0(\sigma_z) &:= \mathbf{x}^0, \\ \widehat{\mathbf{x}}^\ell(\sigma_z) &:= f\left((a_0 \mathbf{D}^\ell + \frac{\sigma_z}{\sqrt{N_{\ell-1}}} \widehat{\mathbf{G}}^{(\ell)}) \widehat{\mathbf{x}}^{\ell-1}(\sigma_z)\right), \quad \ell = 1, 2, \dots, L, \end{aligned}$$

where f is applied coordinatewise. For each fixed choice of $(\widehat{\mathbf{G}}^{(1)}, \dots, \widehat{\mathbf{G}}^{(L)})$, each layer ℓ and each coordinate i , this defines a map

$$\sigma_z \mapsto \widehat{x}_i^\ell(\sigma_z).$$

Since f is continuous and, for fixed $\widehat{\mathbf{G}}^{(\ell)}$, the map $\sigma_z \mapsto (a_0 \mathbf{D}^{(\ell)} + \frac{\sigma_z}{\sqrt{N_{\ell-1}}} \widehat{\mathbf{G}}^{(\ell)}) \widehat{\mathbf{x}}^{\ell-1}(\sigma_z)$ is affine in σ_z , it follows by induction on ℓ that

$$\sigma_z \mapsto \widehat{x}_i^\ell(\sigma_z) \quad \text{is continuous on } [0, \sigma_1] \text{ for any finite } \sigma_1 > 0. \quad (9)$$

For each ℓ and i with $x_i^\ell(\sigma_z) \neq 0$, define

$$R_\ell(\sigma_z; i) := \frac{\|\mathbf{x}^\ell(\sigma_z)\|_2^2}{N_\ell (x_i^\ell(\sigma_z))^2}, \quad R_\ell(0; i) := \frac{\|\mathbf{x}^\ell(0)\|_2^2}{N_\ell (x_i^\ell(0))^2},$$

so that $R_\ell(0; i)$ is exactly the deterministic ratio considered above. On the event that $x_i^\ell(\sigma_z) \neq 0$ for all $\ell \leq L$ and all $0 \leq \sigma_z \leq \sigma_1$, equation 9 implies that $\sigma_z \mapsto R_\ell(\sigma_z; i)$ is continuous. Hence, for each fixed pair (ℓ, i) and any $\gamma \in (0, 1)$ there exists $\sigma_0(\ell, i) > 0$ such that

$$|R_\ell(\sigma_z; i) - R_\ell(0; i)| \leq \frac{\gamma}{2} \quad \text{for all } 0 \leq \sigma_z \leq \sigma_0(\ell, i). \quad (10)$$

Combining equation 10 with the deterministic bound $|R_\ell(0; i) - 1| \leq \gamma/2$ (valid for all $\ell \geq L_0(\delta)$ from equation 8), we obtain

$$|R_\ell(\sigma_z; i) - 1| \leq |R_\ell(\sigma_z; i) - R_\ell(0; i)| + |R_\ell(0; i) - 1| \leq \gamma$$

for all $0 \leq \sigma_z \leq \sigma_0(\ell, i)$ and all $\ell \geq L_0(\delta)$.

Since we only consider a finite set of layers $\ell < L$ and indices $1 \leq i \leq N_\ell$, we may define, for each fixed $(\widehat{\mathbf{G}}^{(1)}, \dots, \widehat{\mathbf{G}}^{(L)})$,

$$\sigma_0(\widehat{\mathbf{G}}^{(1)}, \dots, \widehat{\mathbf{G}}^{(L)}) := \min_{L_0(\delta) \leq \ell < L} \min_{1 \leq i \leq N_\ell} \sigma_0(\ell, i) > 0,$$

so that the bound

$$|R_\ell(\sigma_z; i) - 1| \leq \gamma \quad (11)$$

holds simultaneously for all $L_0(\delta) \leq \ell < L$ and all $1 \leq i \leq N_\ell$ whenever $0 \leq \sigma_z \leq \sigma_0(\widehat{\mathbf{G}}^{(1)}, \dots, \widehat{\mathbf{G}}^{(L)})$.

Now regard $(\mathbf{G}^{(1)}, \dots, \mathbf{G}^{(L)})$ as random Gaussian matrices and set

$$S := \sigma_0(\mathbf{G}^{(1)}, \dots, \mathbf{G}^{(L)}).$$

By the construction above we have $S > 0$ almost surely. Hence, for a given $\gamma \in (0, 1)$ we can choose a deterministic constant $\sigma_0 > 0$ such that $\mathbb{P}(S \geq \sigma_0) \geq 1 - \gamma$ (for example, take σ_0 to be the $(1 - \gamma)$ -quantile of S). On the event $\{S \geq \sigma_0\}$ the estimate equation 11 therefore holds for all $0 \leq \sigma_z \leq \sigma_0$, all $L_0(\delta) \leq \ell < L$ and all $1 \leq i \leq N_\ell$.

Finally, Lemma 4.5 (with a_0 in place of ω) gives the conditional variance formula

$$\text{Var}(a_i^{\ell+1} | x^\ell) = \frac{\sigma_z^2}{N_\ell} \sum_{j=1}^{N_\ell} \left(\frac{x_j^\ell(\sigma_z)}{x_i^\ell(\sigma_z)} \right)^2 = \sigma_z^2 R_\ell(\sigma_z; i).$$

Combining this with equation 11 and the choice of σ_0 yields, for all $0 < \sigma_z \leq \sigma_0$ and all $L_0(\delta) \leq \ell < L$, $1 \leq i \leq N_\ell$,

$$(1 - \gamma) \sigma_z^2 \leq \text{Var}(a_i^{\ell+1} | x^\ell) \leq (1 + \gamma) \sigma_z^2,$$

on an event of probability at least $1 - \gamma$. \square

The previous results describe how the magnitude of the effective gains behaves across layers (Figures 8 and 10). To control the sign dynamics, we now turn to a scalar surrogate model in which the gains are i.i.d. Gaussian with mean ω and variance σ^2 . In this simplified setting, the only quantity that matters is the probability that the scalar iterate becomes negative at a given depth. The next lemma provides a closed-form recursion for this negative rate.

Lemma A.4. *Let $f \in \mathcal{F}$ and $x_0 > 0$, for every $j \geq 1$,*

$$\pi_j = \frac{1}{2} \left(1 - (1 - 2p_-)^j \right), \quad p_- = \mathbb{P}(A_1 < 0) = \Phi\left(-\frac{k}{\sigma}\right). \quad (12)$$

Proof. Since f is odd and strictly increasing, $\text{sign}(f(u)) = \text{sign}(u)$ for all u , hence

$$\{X_j < 0\} = \{A_j X_{j-1} < 0\} = (\{A_j < 0\} \cap \{X_{j-1} > 0\}) \cup (\{A_j > 0\} \cap \{X_{j-1} < 0\}),$$

up to null sets (because $\mathbb{P}(A_j = 0) = 0$ for Gaussian A_j). Independence of A_j and X_{j-1} yields

$$\pi_j = \mathbb{P}(A_j < 0) \mathbb{P}(X_{j-1} > 0) + \mathbb{P}(A_j > 0) \mathbb{P}(X_{j-1} < 0) = p_- (1 - \pi_{j-1}) + (1 - p_-) \pi_{j-1}.$$

Thus $\pi_j = (1 - 2p_-) \pi_{j-1} + p_-$. With $\pi_0 = \mathbb{P}(X_0 < 0) = 0$, the first-order linear recursion solves to

$$\pi_j - \frac{1}{2} = (1 - 2p_-) \left(\pi_{j-1} - \frac{1}{2} \right) \implies \pi_j - \frac{1}{2} = (1 - 2p_-)^j \left(\pi_0 - \frac{1}{2} \right) = -\frac{1}{2} (1 - 2p_-)^j.$$

The value of p_- follows from $A_1 \sim \mathcal{N}(\omega, \sigma^2)$. \square

Theorem 4.7 Fix a target $p \in [0, \frac{1}{2})$, a depth $\ell \in \mathbb{N}$, and $\omega > 0$. There exists a unique $\sigma^* = \sigma^*(p, \ell, \omega) > 0$ such that $\pi_\ell(\sigma^*) = p$, and it is given by

$$\sigma^*(p, \ell, \omega) = - \frac{\omega}{\Phi^{-1}\left(\frac{1 - (1 - 2p)^{1/\ell}}{2}\right)}. \quad (13)$$

Proof. $\pi_\ell(\cdot)$ is continuous and strictly increasing from 0 (at $\sigma \downarrow 0$) to $1/2$ (as $\sigma \uparrow \infty$). Hence for any $p \in [0, 1/2)$ there exists a unique $\sigma^* > 0$ with $\pi_\ell(\sigma^*) = p$.

To obtain the explicit form, set $q := p_-(\sigma^*) = \Phi(-\omega/\sigma^*) \in (0, 1/2)$. From Lemma A.4 we have

$$p = \pi_\ell(\sigma^*) = \frac{1}{2} \left(1 - (1 - 2q)^\ell \right) \implies q = \frac{1 - (1 - 2p)^{1/\ell}}{2}.$$

Applying the inverse CDF Φ^{-1} to $q = \Phi(-\omega/\sigma^*)$ yields

$$-\frac{k}{\sigma^*} = \Phi^{-1}(q) \implies \sigma^* = -\frac{k}{\Phi^{-1}(q)}.$$

Since $q \in (0, 1/2)$, the quantile $\Phi^{-1}(q) < 0$, so the right-hand side is positive. \square

1080 A.3 THEORETICAL RESULTS FOR SECTION 4.3

1081 Before turning to the empirical comparisons in Section 4.3, we connect our gain calibration to gra-
 1082 dient propagation. In the main text we argued that the trainability of very deep networks is governed
 1083 by how the norm of the backpropagated gradient evolves with depth. In this appendix we make
 1084 this precise by computing, under standard mean-field assumptions, the layerwise gradient amplifi-
 1085 cation factor χ_ℓ for our structured initialization and contrasting it with the Gaussian i.i.d. case. This
 1086 will justify the claims in Section 4.3 about the robustness of the proposed scheme to the choice of
 1087 variance.

1088 Let $\mathbf{g}^\ell := \partial\mathcal{L}/\partial\mathbf{x}^\ell \in \mathbb{R}^{N_\ell}$ denote the gradient at layer ℓ for a loss \mathcal{L} . By the chain rule and the
 1089 layerwise relation $\mathbf{h}^{\ell+1} = \mathbf{W}^{\ell+1}\mathbf{x}^\ell + \mathbf{b}^{\ell+1}$, we can write

$$1090 \mathbf{g}^\ell = (\mathbf{W}^{\ell+1})^\top (f'(\mathbf{h}^{\ell+1}) \odot \mathbf{g}^{\ell+1}) = (\mathbf{J}^{\ell+1})^\top \mathbf{g}^{\ell+1},$$

1091 where \odot denotes the Hadamard product and

$$1092 \mathbf{J}^{\ell+1} := \text{diag}(f'(\mathbf{h}^{\ell+1})) \mathbf{W}^{\ell+1}$$

1093 is the Jacobian matrix of layer $\ell + 1$.

1094 In the wide layer regime ($N_\ell \rightarrow \infty$) and under standard mean field assumptions given $\mathbf{h}^{\ell+1}$, the
 1095 squared gradient norms satisfy the approximate recursion

$$1096 \frac{1}{N_\ell} \mathbb{E} \|\mathbf{g}^\ell\|_2^2 \approx \chi_{\ell+1} \frac{1}{N_{\ell+1}} \mathbb{E} \|\mathbf{g}^{\ell+1}\|_2^2, \quad (14)$$

1097 where the layerwise amplification factor $\chi_{\ell+1}$ is defined by

$$1098 \chi_{\ell+1} := \frac{1}{N_{\ell+1}} \mathbb{E} \|\mathbf{J}^{\ell+1} \mathbf{u}\|_2^2,$$

1099 with $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{N_{\ell+1}})$ independent of $\mathbf{J}^{\ell+1}$. Iterating equation 14 over $\ell = 0, \dots, L-1$ gives

$$1100 \frac{1}{N_0} \mathbb{E} \|\mathbf{g}^0\|_2^2 \approx \left(\prod_{\ell=0}^{L-1} \chi_{\ell+1} \right) \frac{1}{N_L} \mathbb{E} \|\mathbf{g}^L\|_2^2.$$

1101 We now compute $\chi_{\ell+1}$ for the proposed structured initialization $\mathbf{W}^\ell = \omega \mathbf{D}^\ell + \mathbf{Z}^\ell$. By definition of
 1102 $\chi_{\ell+1}$,

$$1103 \chi_{\ell+1} = \frac{1}{N_{\ell+1}} \mathbb{E} \|\mathbf{J}^{\ell+1} \mathbf{u}\|_2^2 = \frac{1}{N_{\ell+1}} \mathbb{E} \sum_{i=1}^{N_{\ell+1}} \left(\sum_{j=1}^{N_\ell} w_{ij}^{\ell+1} f'(h_i^{\ell+1}) u_j \right)^2.$$

1104 Conditioning on $\mathbf{W}^{\ell+1}$ and $\mathbf{h}^{\ell+1}$, the inner sum is a centred Gaussian in \mathbf{u} with variance

$$1105 \sum_{j=1}^{N_\ell} (w_{ij}^{\ell+1})^2 f'(h_i^{\ell+1})^2.$$

1106 Taking the expectation over \mathbf{u} and then over the weights and preactivations yields

$$1107 \chi_{\ell+1} = \frac{1}{N_{\ell+1}} \mathbb{E} \sum_{i=1}^{N_{\ell+1}} f'(h_i^{\ell+1})^2 \sum_{j=1}^{N_\ell} (w_{ij}^{\ell+1})^2$$

$$1108 \approx \mathbb{E} \left[f'(h_1^{\ell+1})^2 \sum_{j=1}^{N_\ell} (w_{1j}^{\ell+1})^2 \right],$$

1109 where we used exchangeability of the rows of $\mathbf{W}^{\ell+1}$ and of the coordinates of $\mathbf{h}^{\ell+1}$.

1110 For the proposed initialization, $(\mathbf{Z}_{ij}^{\ell+1}) \sim \mathcal{N}(0, \sigma_z^2/N_\ell)$, so

$$1111 \sum_{j=1}^{N_\ell} (w_{1j}^{\ell+1})^2 = \omega^2 + \sum_{j=1}^{N_\ell} (\mathbf{Z}_{1j}^{\ell+1})^2.$$

1134 Taking expectation over $\mathbf{Z}^{\ell+1}$ and using $\mathbb{E}[(Z_{1j}^{\ell+1})^2] = \sigma_z^2/N_\ell$ gives
 1135

$$1136 \mathbb{E}\left[\sum_{j=1}^{N_\ell} (w_{1j}^{\ell+1})^2\right] = \omega^2 + \sum_{j=1}^{N_\ell} \mathbb{E}[(Z_{1j}^{\ell+1})^2] = \omega^2 + \sigma_z^2.$$

1139 Moreover, by construction of the initialization, $\mathbf{W}^{\ell+1}$ and $\mathbf{h}^{\ell+1}$ are independent, and the coordinates
 1140 $h_i^{\ell+1}$ are exchangeable. Hence
 1141

$$1142 \mathbb{E}\left[f'(h_1^{\ell+1})^2 \sum_{j=1}^{N_\ell} (w_{1j}^{\ell+1})^2\right] \approx \mathbb{E}[f'(h^{\ell+1})^2] \mathbb{E}\left[\sum_{j=1}^{N_\ell} (w_{1j}^{\ell+1})^2\right] = (\omega^2 + \sigma_z^2) \mathbb{E}[f'(h^{\ell+1})^2].$$

1144 Thus,

$$1145 \chi_{\ell+1} \approx (\omega^2 + \sigma_z^2) \mathbb{E}[f'(h^{\ell+1})^2]. \quad (15)$$

1148 Equation equation 15 shows that, for the proposed initialization, the mean-field amplification factor
 1149 $\chi_{\ell+1}$ depends on the combined scale $\omega^2 + \sigma_z^2$ rather than on a bare variance parameter alone. In par-
 1150 ticular, increasing the noise level σ_z decreases $\mathbb{E}[f'(h^{\ell+1})^2]$ while increasing the prefactor $\omega^2 + \sigma_z^2$,
 1151 so that $\chi_{\ell+1}$ remains close to one over a much wider range of σ_z than in the Gaussian i.i.d. case,
 1152 where $\chi_{\ell+1}^* \approx \sigma_w^2 \mathbb{E}[f'(h^{\ell+1})^2]$ has no analogous ω^2 term. This analytic behavior underlies the em-
 1153 pirical observations in Section 4.3 that our initialization preserves gradient norms more effectively
 1154 across depth and is substantially more robust to variance misspecification.
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

B ADDITIONAL EXPERIMENTAL RESULTS WITHOUT TRAINED NETWORKS

B.1 DEFINITIONS OF ACTIVATION FUNCTIONS

This section introduces the activation functions considered in this paper. The following functions belong to the odd-sigmoid function class.

Gudermannian function The Gudermannian function $\text{gd} : \mathbb{R} \rightarrow \mathbb{R}$ is defined by

$$\text{gd}(x) = \int_0^x \frac{dt}{\cosh t} = 2 \arctan\left(\tanh\left(\frac{x}{2}\right)\right).$$

Error Function. The error function $\text{erf} : \mathbb{R} \rightarrow \mathbb{R}$ is defined as

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt.$$

Softsign-Type Functions. For $k \geq 1$, we define the generalized softsign function

$$\text{softsign}_k(x) = \frac{x}{(1 + |x|^k)^{1/k}}, \quad x \in \mathbb{R}.$$

This family interpolates between several commonly used smooth odd activations:

$$\text{softsign}_1(x) = \frac{x}{1 + |x|}, \quad \text{softsign}_2(x) = \frac{x}{\sqrt{1 + x^2}}, \quad \text{softsign}_3(x) = \frac{x}{(1 + |x|^3)^{1/3}}.$$

We also use the combined variant

$$\text{softsign}_{1+3}(x) := \text{softsign}_1(x) + \text{softsign}_3(x).$$

Hyperbolic Tangent. The hyperbolic tangent function $\tanh : \mathbb{R} \rightarrow \mathbb{R}$ is defined by

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}.$$

Arctangent. The (scaled) arctangent function $\arctan : \mathbb{R} \rightarrow \mathbb{R}$ is given by

$$\arctan(x) = \int_0^x \frac{dt}{1 + t^2}.$$

In practice, we often use the normalized form $\frac{2}{\pi} \arctan(x)$ so that its range matches that of $\tanh(x)$.

B.2 PROPERTIES OF THE ODD-SIGMOID FUNCTIONS

This section empirically investigates the properties of the iteration $x_{n+1} = f(ax_n)$. When the gain a is fixed across all iterations, Figure 8 shows how the dynamics depend on the initial data: for any nonzero input, the iterates converge to the same nonzero fixed point ξ_a . Figure 9 further demonstrates that this limit depends only on the gain, not on the input, by comparing distinct gains $a \neq a' > 0$ and observing convergence to different fixed points ξ_a and $\xi_{a'}$. These results indicate that, by appropriately choosing a , one can control how long the signal remains informative across layers and thus preserve information up to a desired depth.

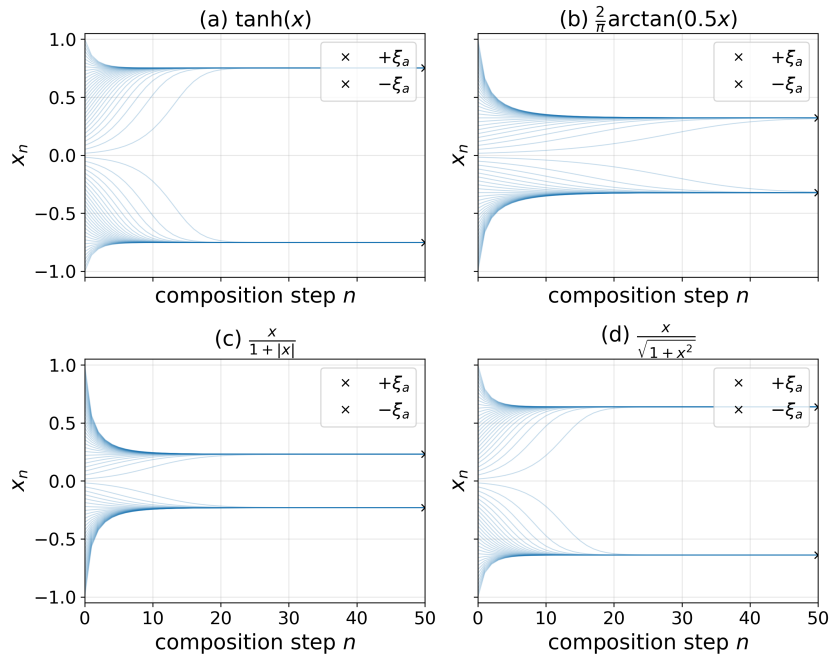


Figure 8: Iterated scalar dynamics for four odd-sigmoid activations under a fixed supercritical gain. For each activation f , we compute $\omega = 1/f'(0)$ and set $a = \omega + 0.3$. We then iterate the one-dimensional map $x_{n+1} = f(ax_n)$ for $n = 0, \dots, 50$ starting from 60 initial values $x_0 \in [-1, 1]$. The limiting fixed points $\pm\xi_a$ satisfying $f(a\xi_a) = \xi_a$ are approximated by iterating from ± 1 and are marked at $n = 50$ with “x”.




Figure 9: Convergence of the iteration $x_{n+1} = f(ax_n)$ for odd-sigmoid f with $a = \omega + \epsilon$ ($\epsilon \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$); curves show x_n up to $n = 50$ from $x_0 = 0.1$, and “x” marks the positive fixed point ξ_a solving $f(a\xi_a) = \xi_a$.

B.3 NEGATIVE RATE FUNCTION

We calibrate the noise scale σ_z in the surrogate model by targeting a desired negative rate at a specified depth. For the scalar recursion with gains $A_j \sim \mathcal{N}(\omega, \sigma^2)$ and depth L , let $\pi_L(\sigma)$ denote the probability that the iterate is negative at layer L . Given a target $p \in [0, \frac{1}{2}]$, Theorem 4.7 yields a unique $\sigma^*(p, L, \omega)$ such that $\pi_L(\sigma^*) = p$. Figure 10 illustrates this calibration: panel (a) shows the closed-form scale $\sigma^*(p, L, \omega)$ as a function of p for $L = 100$ and $\omega = 1$, while panel (b) plots $\pi_L(\sigma)$ over (L, σ) for $\omega = 1$. As expected, $\pi_L(\sigma) \rightarrow 0$ as $\sigma \rightarrow 0$ (no sign flips) and $\pi_L(\sigma) \rightarrow \frac{1}{2}$ as $\sigma \rightarrow \infty$ (full symmetry), so the negative rate traces a narrow band between these two extremes.

Figure 11 evaluates the resulting calibration from a mean-field perspective. For $f(x) = \tanh(x)$ and $\omega = 1$, we choose $\sigma_z = \sigma^*(p, L, \omega)$ for target negative rates $p = 0.01$ and $p = 0.49$, and compute the corresponding gradient amplification factor $\chi_\ell \approx (\omega^2 + \sigma_z^2) \mathbb{E}[f'(h^\ell)^2]$ across depths. The curves show that χ_ℓ remains very close to 1 over a wide range of depths, up to $L = 2 \times 10^5$, for both target negative rates, indicating that the negative-rate calibration keeps gradients in a trainable regime even in extremely deep networks.




Figure 10: Closed-form characterization of the scalar surrogate. (a) Closed form scale $\sigma^*(p, L, \omega)$ as a function of the target negative rate p for $L = 100$ and $\omega = 1$, with reference lines at $p = 0.10$ and $p = 0.49$. (b) Closed form scalar surrogate negative rate $\pi_L(\sigma)$ for $\omega = 1$, network depths $L = 1, 2, \dots, 1000$, and $\sigma \in [0.01, 1.0]$.




Figure 11: χ_ℓ for the proposed initialization with $f(x) = \tanh(x)$, evaluated at the closed form noise scales $\sigma^*(p, L, \omega)$ corresponding to target negative rates $p = 0.01$ and $p = 0.49$. For each depth ℓ , we estimate $\chi_\ell(\sigma_z) \approx (\omega^2 + \sigma_z^2) \mathbb{E}[f'(h^\ell)^2]$.

B.4 NEGATIVE RATE SURROGATE ANALYSIS

We begin from the elementwise representation of a fully connected network with odd-sigmoid activation $f \in \mathcal{F}$. For each layer ℓ and neuron index i , the forward update can be written as

$$x_i^{\ell+1} = f(a_i^{\ell+1} x_i^\ell), \quad a_i^{\ell+1} = w_{ii}^{\ell+1} + \sum_{j \neq i} \frac{w_{ij}^{\ell+1} x_j^\ell}{x_i^\ell}, \quad (16)$$

so that $a_i^{\ell+1}$ plays the role of an effective scalar gain for neuron i in layer $\ell + 1$. Under our proposed initialization $W^{\ell+1} = \omega D^{\ell+1} + Z^{\ell+1}$ with $(Z_{ij}^{\ell+1}) \sim \mathcal{N}(0, \sigma_z^2/N_\ell)$, Lemma 4.5 shows that, conditional on x^ℓ ,

$$a_i^{\ell+1} \sim \mathcal{N}\left(\omega, \frac{\sigma_z^2}{N_\ell} \left(1 + \sum_{j \neq i} \left(\frac{x_j^\ell}{x_i^\ell}\right)^2\right)\right). \quad (17)$$

To handle this dependence on the current activations, we approximate the gain distribution by a scalar surrogate model $a_i^{\ell+1} \sim \mathcal{N}(\omega, \sigma_z^2)$ and estimate sign statistics from this Gaussian approximation.

We define the negative rate at depth L as

$$\pi_L(\sigma) := \mathbb{P}(x_L < 0 \mid x_0 > 0),$$

that is, the probability that a neuron whose initial activation is positive ends up with a negative sign at layer L . Because every $f \in \mathcal{F}$ is odd and strictly increasing, the sign of x_i^ℓ is entirely controlled by the product of the effective gains along the path, and the probability that a positive entry becomes negative at depth L equals the probability that a negative entry becomes positive. It therefore suffices to track a single scalar chain starting from $x_0 > 0$.

Our initialization has a nonzero mean gain ω , which makes sign changes along a given coordinate well defined across layers. We interpret frequent sign flips during the forward pass as a form of information loss, and we therefore use the surrogate calibration to control the negative rate at the final layer. In the scalar surrogate, for a given depth L and gain variance σ^2 we can compute $\pi_L(\sigma)$ in closed form. We then define, for a target $p \in (0, \frac{1}{2})$, the calibrated scale $\sigma^*(p, L, \omega)$ as the unique solution to $\pi_L(\sigma^*) = p$ (see Theorem 4.7). In practice, we choose a desired “real” negative rate $p_{\text{real}} = 0.4$ and select σ_z so that the empirical FFNN-driven negative rate at depth L is close to p_{real} , which preserves most sign information while still leaving enough randomness for learning.




Figure 12: Best surrogate target $p_{\text{target}}(L)$ producing an FFNN-driven negative rate $\tilde{\pi}_L \approx 0.4$ as a function of depth L for a tanh network with width 64. The curve shows that $p_{\text{target}}(L)$ stays near 0.3–0.4 for shallow depths and then decays approximately exponentially with L , providing an empirical calibration rule for choosing the surrogate negative rate in deep networks.

To understand how the scalar surrogate relates to the actual network, we compare $\pi_L(\sigma^*)$ with the empirical negative rate obtained from an initialized FFNN. For each p_{target} and depth L , we first compute $\sigma^*(p_{\text{target}}, L, \omega)$, initialize an L -layer tanh network with weights $W^\ell = \omega I + (\sigma^*/\sqrt{N})G^\ell$, extract the effective gains a_i^ℓ , and form FFNN-driven scalar chains $y_0 = 1$,

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457




Figure 13: Negative rate at depth $L = 1, 10, 50$ as a function of the target negative rate p_{target} . The curve $\pi_1(\sigma^*)$ denotes the scalar surrogate prediction at the calibrated scale $\sigma^*(p_{\text{target}})$, $\tilde{\pi}_1$ is the empirical negative rate obtained from the FFNN-driven scalar chains, and the dashed line shows the identity p_{target} .

$y_{\ell+1} = \tanh(a^\ell y_\ell)$. Measuring $\tilde{\pi}_L = \mathbb{P}(y_L < 0)$ and comparing it to the analytic $\pi_L(\sigma^*)$ reveals that the discrepancy grows with depth (see Figure 13), as expected from the increasing influence of higher-order correlations. Nevertheless, the relationship between the surrogate target p_{target} and the FFNN-driven negative rate $\tilde{\pi}_L$ remains systematic.

In particular, for each depth L we can invert this relationship numerically: we search over p_{target} and find the value $p_{\text{sur}}(L)$ such that the FFNN-driven negative rate $\tilde{\pi}_L$ is as close as possible to $p_{\text{real}} = 0.4$ (see Section B.7).

The resulting sequence $p_{\text{sur}}(L)$ is reported in Figure 12. For small depths L , the surrogate target remains near 0.3–0.4, but for larger L it decays approximately exponentially. Fitting a simple model $\log p_{\text{sur}}(L) \approx c_0 + c_1 L$ on the tail (e.g., $L \geq 10$) yields

$$p_{\text{sur}}(L) \approx C e^{-\alpha L},$$

for some $C > 0$ and $\alpha > 0$. This empirical law expresses how aggressively the surrogate target negative rate must shrink with depth in order for the actual network to maintain a fixed, moderate negative rate at its final layer.

These calibration insights are consistent with our training experiments. Figures 14 and 15 plot validation accuracy as a function of the surrogate target p for deep ($L = 50$) and shallow ($L = 3$) networks. For the deep case, performance is maximized when p_{sur} is on the order of 0.01, in line with the exponentially decayed target predicted by the surrogate analysis. In contrast, for shallow networks the best performance is attained near $p_{\text{sur}} \approx 0.4$, matching the regime where the fitted decay has not yet taken effect. Finally, Figures 16 and 17 show that, for depths $L = 50$ and $L = 100$,




Figure 14: Validation accuracy as a function of the target negative rate p for a 50-layer fully connected network (width 64) with activation $f(x) = \tanh(x)$ under the proposed initialization. Each point shows the mean \pm standard deviation over 5 runs, where each run is trained for 600 iterations, with $\sigma^*(p, L, \omega)$ computed from the scalar surrogate calibration.

the learning curves are optimized near the σ^* values obtained from the calibrated surrogate, with accuracy degrading as we move away from these scales. Together, these results validate that the negative-rate surrogate provides a useful, quantitatively accurate guideline for choosing the noise level σ_z across depths, and that the resulting diagonal-plus-noise initialization indeed preserves signal sign statistics in deep networks.




Figure 15: Validation accuracy as a function of the target negative rate p for a 3-layer fully connected network (width 512) with activation $f(x) = \tanh(x)$ under the proposed initialization. Each point shows the mean \pm standard deviation over 5 runs, where each run is trained for 600 iterations, with $\sigma^*(p, L, \omega)$ computed from the scalar surrogate calibration.




Figure 16: Validation accuracy as a function of the noise scale σ_z for 50 layer fully connected networks (width 64) with activation $f(x) = \tanh(x)$ under the proposed initialization, on MNIST, Fashion MNIST, and CIFAR-10. For each $\sigma_z \in \{0, 0.03, 0.06\}$, we report the mean \pm standard deviation over 5 training runs.




Figure 17: Validation accuracy as a function of the noise scale σ_z for 100 layer fully connected networks (width 64) with activation $f(x) = \tanh(x)$ under the proposed initialization, on MNIST, Fashion MNIST, and CIFAR-10. For each $\sigma_z \in \{0, 0.03, 0.06\}$, we report the mean \pm standard deviation over 5 training runs.

B.5 FORWARD SIGNAL PROPAGATION

In Section 4.3 we study forward signal propagation under the proposed initialization. Figure 10 (a) reports how well the last-layer activation distribution is preserved as the depth L increases (with fixed width $W = 64$), while panel (b) varies the width N_ℓ at fixed depth $L = 1000$ to test whether the distribution remains well spread even in relatively narrow networks. In both settings, the proposed initialization maintains a dispersed last-layer distribution up to depth $L = 10,000$ and down to width $N_\ell = 20$, whereas the EOC (Gaussian i.i.d.) initialization quickly collapses toward a more concentrated distribution.

To summarize how well the last-layer distribution is dispersed, we employ a normalized histogram-entropy score

$$\text{Spread}(x) = \frac{-\sum_{i=1}^B p_i \log p_i}{\log B}, \quad p_i = \int_{\text{bin } i} \hat{f}_x(t) dt, \quad \sum_{i=1}^B p_i = 1, \quad (18)$$

where \hat{f}_x is the empirical density over $[-1, 1]$ using B bins. Values close to 1 indicate a highly dispersed (near-uniform) last-layer distribution, whereas values near 0 correspond to a highly concentrated distribution.


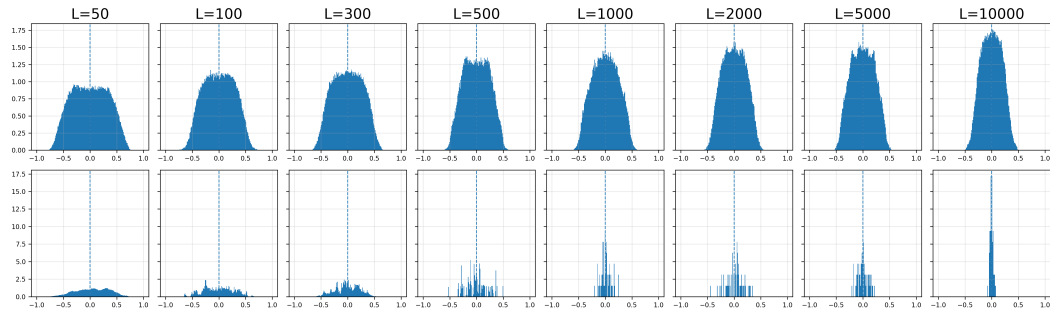


Figure 18: Entropy of the last-layer activation distribution for tanh networks under the proposed initialization and the EOC initialization. (a) Entropy as a function of depth L with fixed width $W = 64$. (b) Entropy as a function of width N_ℓ with fixed depth $L = 1000$, using widths $W \in \{20, 32, 48, 64, 100, 128, 256, 512\}$ as shown in the panels. The entropy is defined in Equation 18.


Figures 19, 20, and 21 visualize the last-layer activation histograms obtained after the initial forward pass in 1000-layer FFNNs of width 64 with activations $0.1 \tanh(x)$, $\tanh(x)$, and $10 \tanh(x)$, respectively. Under the proposed initialization, the activation distribution remains well dispersed in all three cases, whereas Gaussian i.i.d. initializations quickly collapse toward a narrow band around zero. Figures 22 and 23 show a similar comparison as the width is reduced: even for very narrow networks, the proposed scheme preserves a spread-out last layer distribution, while Gaussian initializations drive the last layer activations to saturate near zero.

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576



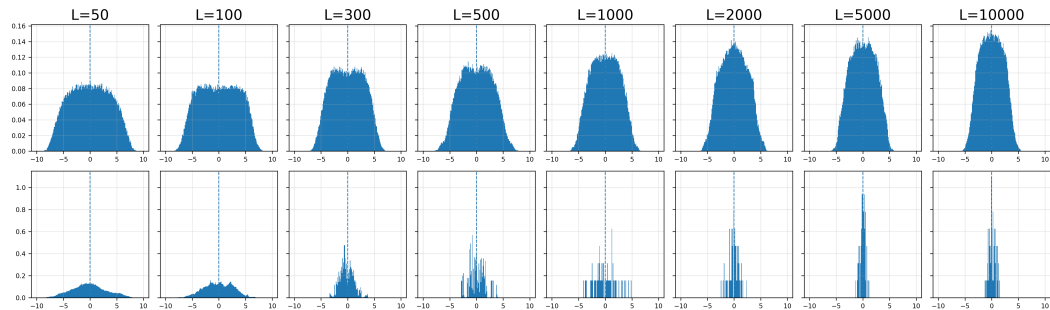
1577
1578 **Figure 19: Last layer activation histograms for tanh networks with width $W = 64$ and varying**
1579 **depth under the proposed initialization (top row) and the EOC initialization (bottom row). Each**
1580 **column corresponds to a different depth L .**

1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592



1593
1594 **Figure 20: Last layer activation histograms for 0.1 tanh networks with width $W = 64$ and varying**
1595 **depth under the proposed initialization (top row) and the EOC initialization (bottom row). Each**
1596 **column corresponds to a different depth L .**


1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608



1609
1610 **Figure 21: Last layer activation histograms for 10 tanh networks with width $W = 64$ and varying**
1611 **depth under the proposed initialization (top row) and the EOC initialization (bottom row). Each**
1612 **column corresponds to a different depth L .**


1613
1614
1615
1616
1617
1618
1619

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630



1631 **Figure 22:** Last layer activation histograms for 0.1 tanh networks with depth $L = 1000$ and varying
1632 width under the proposed initialization (top row) and the EOC initialization (bottom row).
1633

1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645




1646 **Figure 23:** Last layer activation histograms for 10 tanh networks with depth $L = 1000$ and varying
1647 width under the proposed initialization (top row) and the EOC initialization (bottom row).
1648

1649
1650
1651
1652
1653
1654
1655
1656
1657

1650 B.6 BACKWARD SIGNAL PROPAGATION

1652 **Figure 24** visualizes $|\chi_L(\sigma) - 1|$ for the proposed initialization and a Gaussian initialization with
1653 tanh activations, as a function of depth L and scale σ . For the proposed scheme, the region where
1654 $\chi_L(\sigma) \approx 1$ occupies a broad band in the (L, σ) plane, whereas for the Gaussian initialization it is
1655 confined to a narrow strip around a single variance. This indicates that our initialization keeps χ_L
1656 close to 1 over a much wider range of noise scales and depths, and is therefore more robust and
1657 better suited for stable training in deep, narrow networks.

1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673



1671 **Figure 24:** Heatmaps of the deviation $|\chi_L(\sigma) - 1|$ for the proposed initialization (**left**) and a Gaussian
1672 i.i.d. initialization (**right**) with tanh activations, as a function of depth L and scale σ . Brighter bands
1673 indicate near critical regimes where forward and backward signals are approximately preserved.

1674 B.7 CHOICE OF THE TARGET NEGATIVE RATE $p_{\text{real}} = 0.4$.

1675
1676 Because $f \in \mathcal{F}$ is odd and strictly increasing, the sign of each coordinate x_i^ℓ is entirely determined
1677 by the product of the effective gains along that coordinate. We therefore use the coordinate-wise
1678 sign flip probability

$$1679 \tilde{\pi}_L := \mathbb{P}(x_L < 0 \mid x_0 > 0)$$

1680 as a simple indicator for how much “sign information” about the input is preserved at depth L .
1681 Two extreme regimes are undesirable. If $\tilde{\pi}_L \approx 0$, almost all coordinates preserve their initial sign,
1682 so the network behaves nearly like an identity map at initialization; this preserves information but
1683 yields very limited expressiveness and weak exploration of the odd-sigmoid nonlinearity. At the
1684 other extreme, if $\tilde{\pi}_L \approx 0.5$, the final sign is essentially a fair coin flip regardless of the initial sign,
1685 meaning that the directional information carried by the input has been almost completely randomized
1686 and we interpret this as a form of information loss.

1687 In practice, we therefore target an intermediate regime in which most coordinates keep their initial
1688 sign, but a non-negligible fraction flip so that the representation can change meaningfully. Con-
1689 cretely, we fix a desired “real” negative rate $p_{\text{real}} = 0.4$ and choose σ_z so that the empirical FFNN-
1690 driven negative rate at depth L satisfies $\tilde{\pi}_L \approx p_{\text{real}}$. This calibration preserves the sign of the
1691 majority of coordinates ($\approx 60\%$) while still allowing a substantial minority ($\approx 40\%$) to flip, which
1692 provides enough randomness for learning without completely destroying the initial sign structure.

1693 We do not claim that $p_{\text{real}} = 0.4$ is an information-theoretically optimal value. Rather, it is an
1694 empirically grounded target: across a wide range of depths, widths, datasets, and activation scales,
1695 we consistently observe that the σ_z obtained from the $p_{\text{real}} = 0.4$ calibration yields the best or
1696 near-best validation performance, with accuracy degrading as we move to significantly smaller or
1697 larger negative rates (see Figures 14, 15, 16, and 17). This suggests that the proposed negative rate
1698 criterion provides a practically useful operating point for odd-sigmoid networks.

1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

C ADDITIONAL EXPERIMENTAL RESULTS WITH NEURAL NETWORKS

Experimental Setting. We evaluate the proposed initialization using the Adam optimizer with a batch size of 128 and reserve 15% of the training data for validation. All experiments are implemented in PyTorch without skip connections and without learning rate decay. We set the learning rate to 10^{-4} , ω and we use the same learning rate for all Gaussian i.i.d. baseline initializations (Xavier, He, and EOC) for a fair comparison.

C.1 NETWORK SIZE INDEPENDENCE

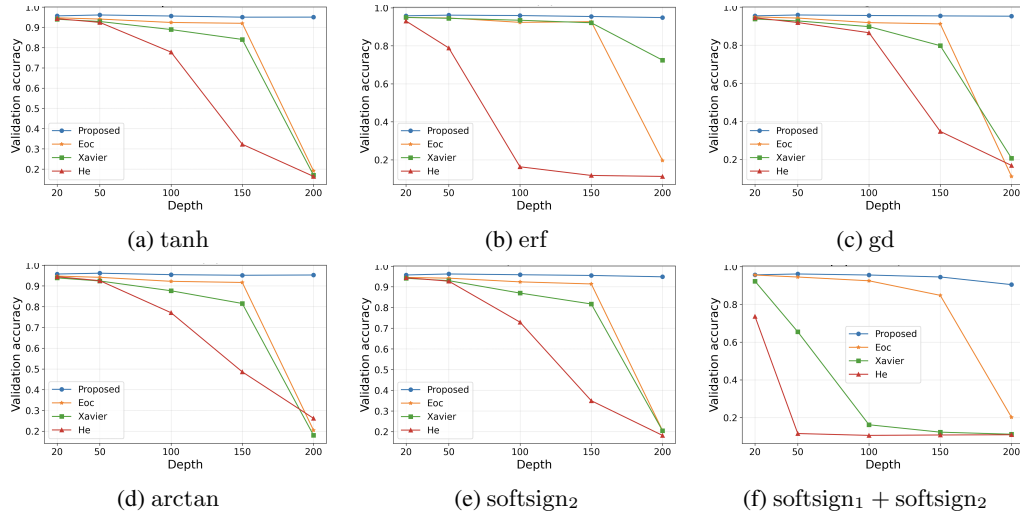


Figure 25: MNIST validation accuracy versus depth for FFNNs (width 64) with odd-sigmoid activations and four initializations (Proposed, EOC, Xavier, He). Each panel fixes one activation and shows the best validation accuracy over 10 epochs for depths $L \in \{20, 50, 100, 150, 200\}$.

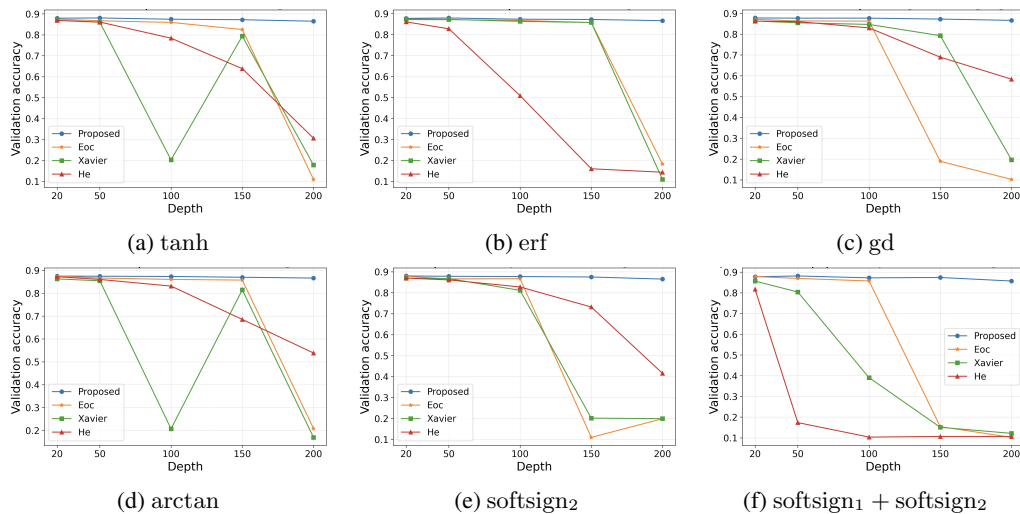


Figure 26: Fashion MNIST validation accuracy versus depth for FFNNs (width 64) with odd-sigmoid activations and four initializations (Proposed, EOC, Xavier, He). Each panel fixes one activation and shows the best validation accuracy over 10 epochs for depths $L \in \{20, 50, 100, 150, 200\}$.

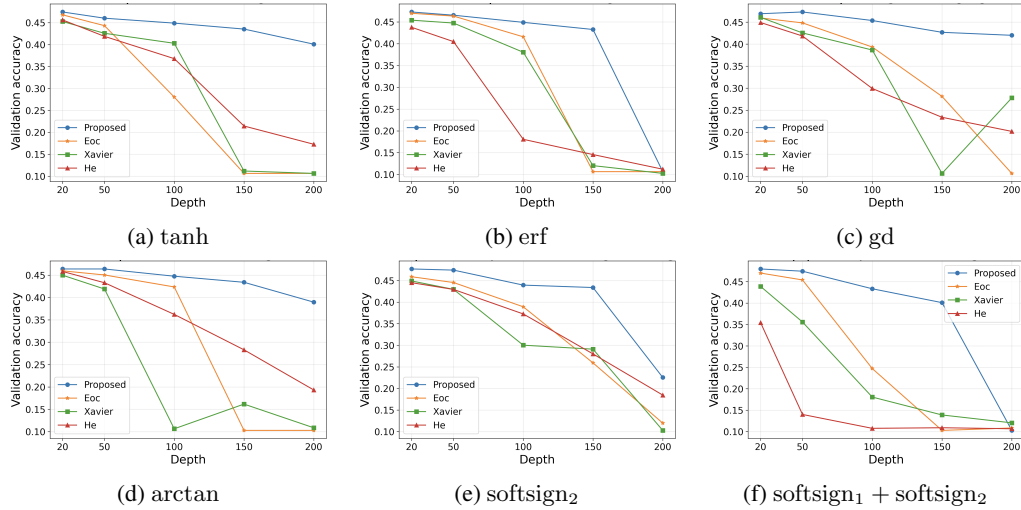


Figure 27: CIFAR 10 validation accuracy versus depth for FFNNs (width 64) with odd-sigmoid activations and four initializations (Proposed, EOC, Xavier, He). Each panel fixes one activation and shows the best validation accuracy over 10 epochs for depths $L \in \{20, 50, 100, 150, 200\}$.

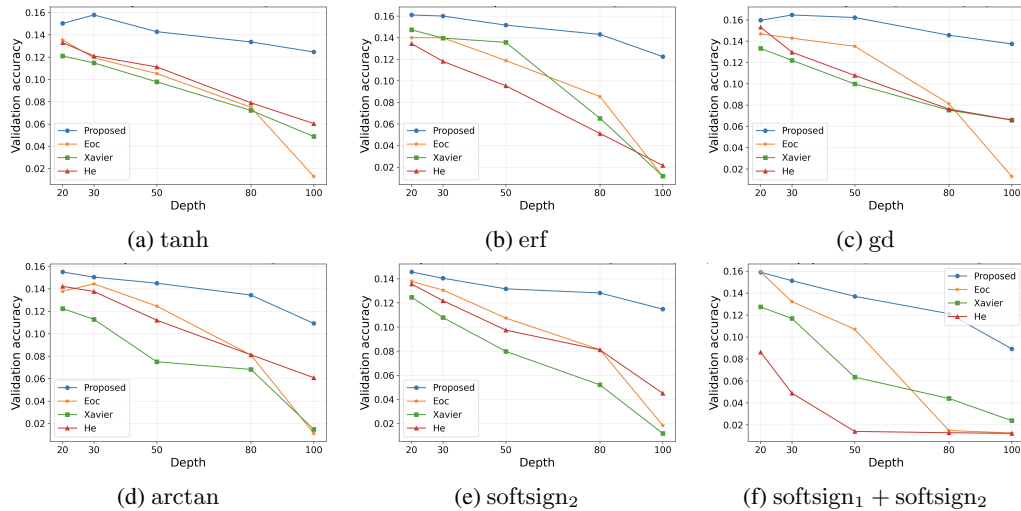


Figure 28: CIFAR 100 validation accuracy versus depth for FFNNs (width 64) with odd-sigmoid activations and four initializations (Proposed, EOC, Xavier, He). Each panel fixes one activation and shows the best validation accuracy over 10 epochs for depths $L \in \{20, 50, 100, 150, 200\}$.

Figures 25, 26, 27, and 28 report the validation accuracy as a function of depth for networks with six different activation functions. On the more complex CIFAR-10 and CIFAR-100 datasets, even the proposed initialization exhibits some change in performance once the depth becomes sufficiently large, whereas on MNIST and Fashion-MNIST the validation accuracy is essentially preserved across all depths considered. These results indicate that, across diverse datasets and for odd-sigmoid activations with effective gain parameter ω close to one, the proposed initialization behaves substantially more depth invariant than standard Gaussian initializations.

C.2 BATCH NORMALIZATION FREE TRAINING

Batch normalization (BN) has made training deep networks substantially easier, but it also incurs a significant computational overhead. This raises the question of whether our initialization can match or surpass BN-based methods while avoiding this overhead. Figure 30 evaluates tanh networks and shows that, across all datasets, the proposed initialization without BN outperforms or matches competing schemes that rely on BN. Figures 29 and 31 extend this comparison to activations with $\omega = 1/f'(0)$ far from 1, including an 800 layer network with width 32. In this challenging deep and narrow regime, our method still trains reliably without BN, whereas Gaussian initializations (with or without BN) either fail to converge or achieve worse accuracy. These results suggest that the proposed initialization can substitute for BN in many settings, enabling stable training of very deep and narrow networks without the cost of normalization layers.

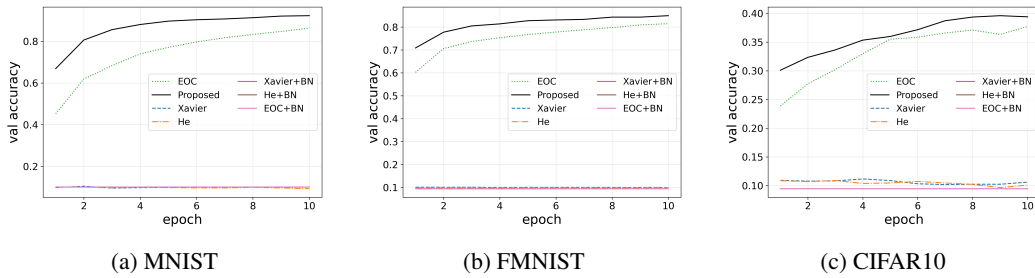


Figure 29: Validation accuracy for the $f(x) = \tanh(ax) + \text{erf}(bx) + x/(1 + |cx|) + \text{gd}(dx)$ in FFNN with 100 hidden layers of width 64. We compare seven strategies: Proposed, Xavier, He, EOC, and their BN variants. Each panel uses a different dataset and a different choice of coefficients (a, b, c, d) : (a) (10,1000,10,1), (b) (100,1000,10,1000), (c) (1000,100,0.1,0.01).

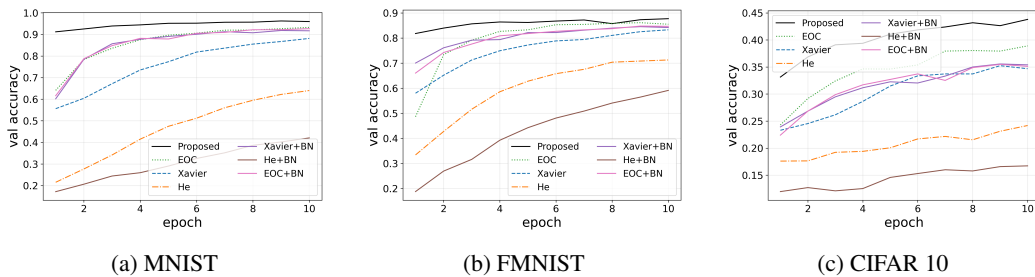


Figure 30: Validation accuracy over 10 epochs for fully connected 100 layer networks with \tanh activations and width 64. The curves compare the Proposed, EOC, Xavier, and He initializations, with and without batch normalization.

1890
 1891
 1892
 1893
 1894
 1895
 1896
 1897
 1898
 1899
 1900
 1901
 1902
 1903
 1904
 1905
 1906
 1907
 1908
 1909
 1910
 1911
 1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943

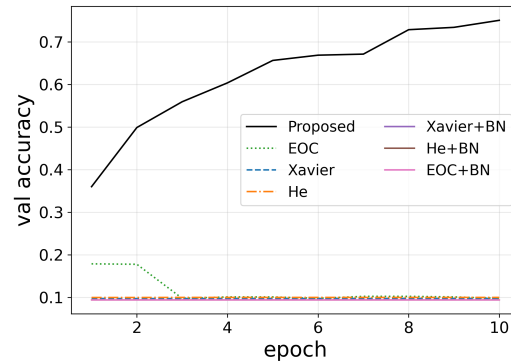


Figure 31: Validation accuracy on Fashion MNIST for a fully connected network with 800 hidden layers of width 32 and $f(x) = \tanh(ax) + \text{erf}(bx) + x/(1 + |cx|) + \text{gd}(dx)$, using coefficients $(a, b, c, d) = (100, 1000, 10, 1000)$.




Figure 32: Best validation accuracy on MNIST versus width $\{10, 48, 128, 150, 200, 512\}$ for 100 layer tanh FFNN. Each curve compares the Proposed, EOC, Xavier, and He initialization schemes, and each point reports the best validation accuracy over 20 training epochs.

C.3 LEARNABLE LEARNING RATE

Our proposed initialization depends explicitly on the $\omega = 1/f'(0)$, so rescaling the activation also rescales the initialized weights. In particular, when we replace f by a scaled activation $\alpha f(x)$ or $f(\alpha x)$, the corresponding value of ω changes and the scale of the initial weight matrix is adjusted accordingly. As a consequence, the learning rate that yields comparable gradient updates should also depend on ω . In our experiments we therefore choose η from an ω -scaled band (e.g., $\eta \in [10^{-5}\omega, 10^{-3}\omega]$) when using the proposed initialization. Across all settings, the proposed initializer remains trainable over a wider interval of η than Gaussian initializations, indicating that it is more robust to the choice of learning rate even when the activation scale varies significantly.

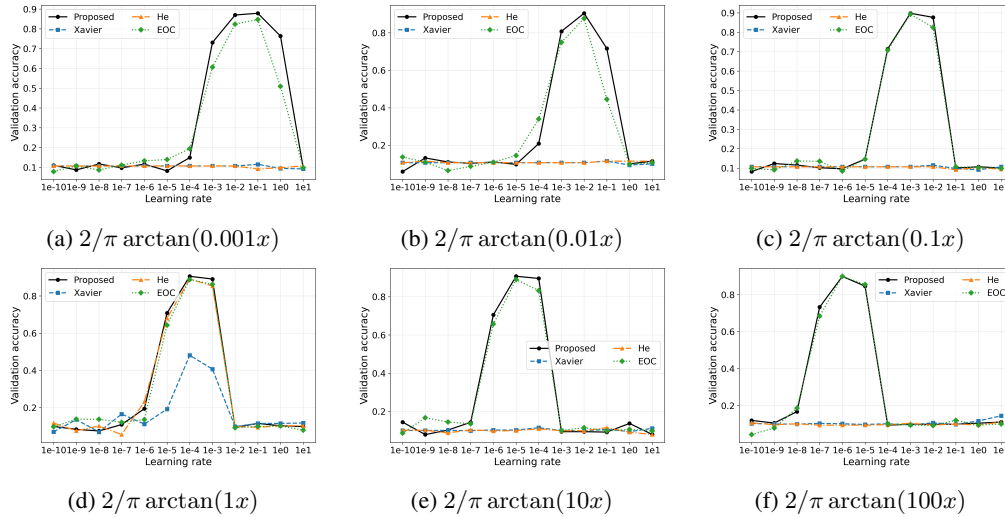


Figure 33: Learning rate accuracy curves on MNIST for a 20 layer, width 512 feedforward network with activation $f(x) = \frac{2}{\pi} \arctan(\alpha x)$. Each panel corresponds to a different activation scale $\alpha \in \{10^2, 10^1, 1, 10^{-1}, 10^{-2}, 10^{-3}\}$. For each learning rate, we train for 200 iterations on a 10k training subset and report the validation accuracy.

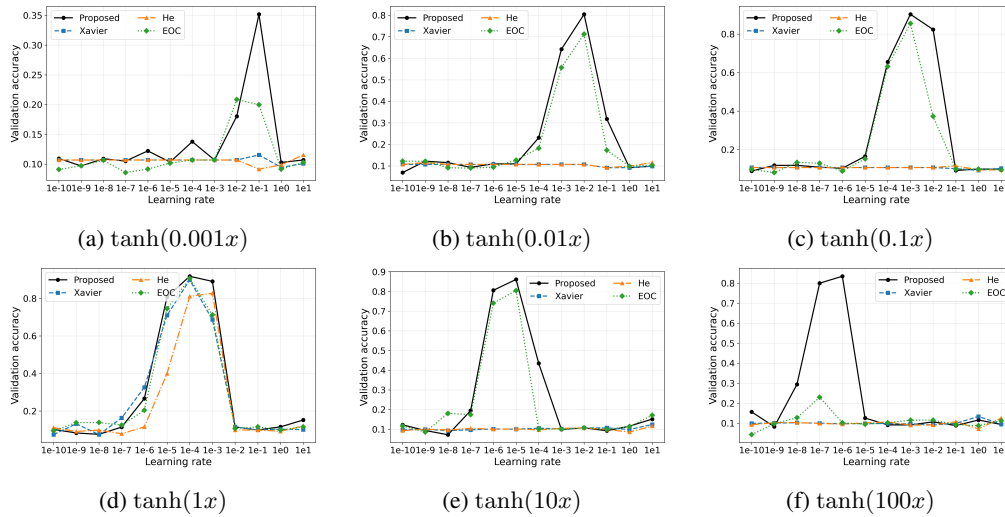


Figure 34: Learning rate accuracy curves on MNIST for a 20 layer, width 512 feedforward network with activation $f(x) = \tanh(\alpha x)$. Each panel corresponds to a different activation scale $\alpha \in \{10^2, 10^1, 1, 10^{-1}, 10^{-2}, 10^{-3}\}$. For each learning rate, we train for 200 iterations on a 1k training subset and report the validation accuracy.

1998

1999

2000

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017







(a) $\tanh(0.01x)$ (b) $\tanh(0.1x)$ (c) $\tanh(x)$ (d) $\tanh(10x)$ (e) $\tanh(100x)$ (f) $\tanh(1000x)$

Figure 35: Learning rate accuracy curves on Fashion MNIST for a 20 layer, width 512 feedforward network with activation $f(x) = \tanh(\alpha x)$. Each panel corresponds to a different activation scale $\alpha \in \{10^3, 10^2, 10^1, 1, 10^{-1}, 10^{-2}\}$. For each learning rate, we train for 200 iterations on a 10k training subset and report the validation accuracy. Curves compare four initializations: Proposed, Xavier, He, and EOC.

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045







(a) $2/\pi \arctan(0.01x)$ (b) $2/\pi \arctan(0.1x)$ (c) $2/\pi \arctan(x)$ (d) $2/\pi \arctan(10x)$ (e) $2/\pi \arctan(100x)$ (f) $2/\pi \arctan(1000x)$

Figure 36: Learning rate accuracy curves on Fashion MNIST for a 20 layer, width 512 feedforward network with activation $f(x) = 2/\pi \arctan(\alpha x)$. Each panel corresponds to a different activation scale $\alpha \in \{10^3, 10^2, 10^1, 1, 10^{-1}, 10^{-2}\}$. For each learning rate, we train for 200 iterations on a 10k training subset and report the validation accuracy. Curves compare four initializations: Proposed, Xavier, He, and EOC.

2049

2050

2051

C.4 SCALE PRESERVING ODD-SIGMOID ACTIVATIONS.

Using a scaled activation function amounts to adjusting the effective range of both the input and output axes of the nonlinearity. For example, $\tanh(x)$ has output range $[-1, 1]$, whereas $\alpha \tanh(x)$ has range $[-\alpha, \alpha]$. With our initialization we therefore ask whether signals remain well propagated over the full interval $[-\alpha, \alpha]$ when the activation is scaled. As shown in Figure 37, for several values of α the proposed initialization keeps the last layer activations of $\alpha \tanh(x)$ spread over $[-\alpha, \alpha]$ even in very deep networks (up to $L = 10^5$), while Gaussian i.i.d. initialization rapidly drives the activations to saturate near zero. In all α -scale experiments (Tables 2–7), each initializer was given its own LR grid search. Even under this per initializer tuning, Gaussian and EOC schemes fail to train for large α , while our method remains stable

Since the proposed initialization lets us target the output scale via the choice of α , it is natural to expect benefits on regression tasks where the range of the target y matters. To test this, we evaluate our method in physics informed neural networks (PINNs), using scaled odd-sigmoid activations for Burgers’ equation and the Black–Scholes equation. The PINN setup and the precise PDE formulations are described in the following subsections.




Figure 37: Activation range as a function of depth $L \in \{50, 500, 5,000, 50,000, 10^5\}$ for fully connected width 64 networks with activation $\alpha \tanh(x)$ under the Proposed and EOC initializations. Panel (a) uses $\alpha = 10^{-4}$ and panel (b) uses $\alpha = 10^3$.

2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

Table 2: Validation accuracy on MNIST (left) and Fashion MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \tanh(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

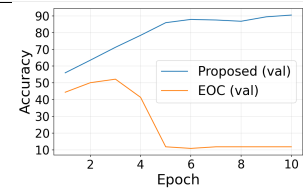
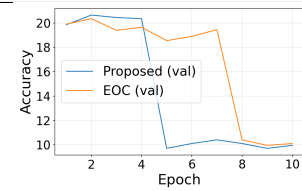
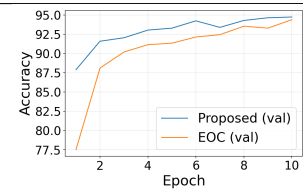
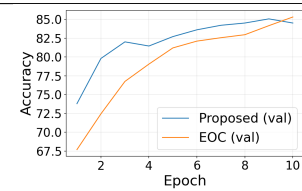
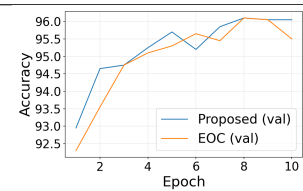
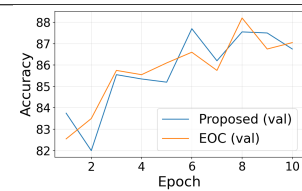
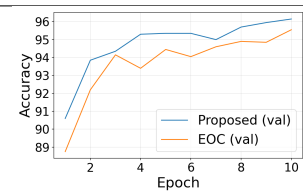
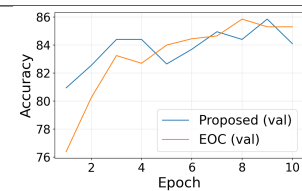
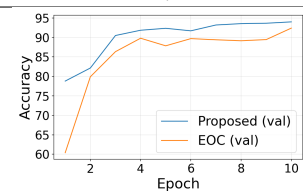
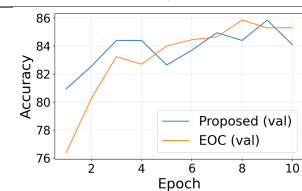
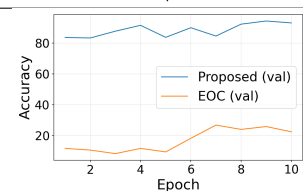
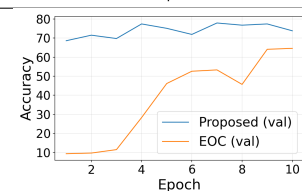
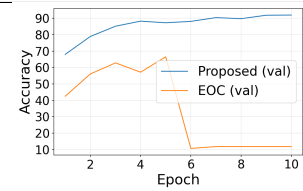
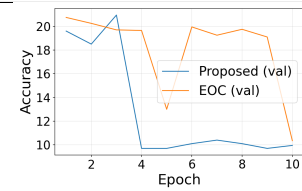
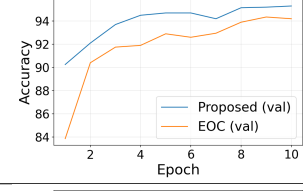
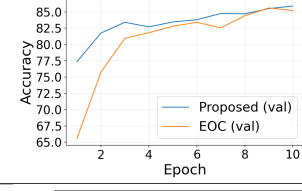
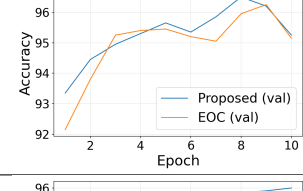
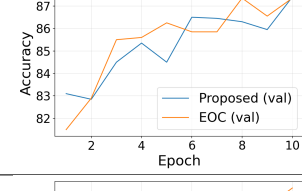
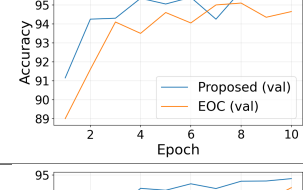
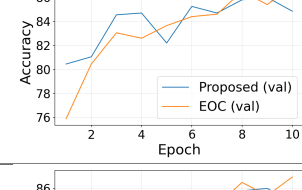
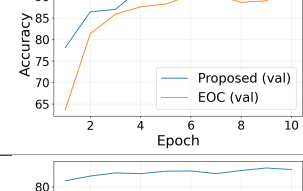
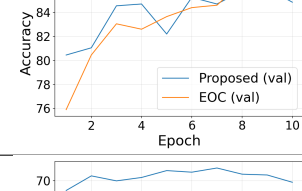
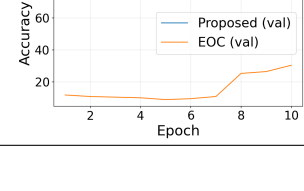
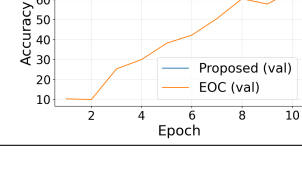
a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^{-2}		
10^{-1}		
1		
10^1		
10^2		
10^3		

Table 3: Validation accuracy on MNIST (left) and Fashion-MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \tanh(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^4		
10^5		
10^6		
10^7		
10^8		
10^9		

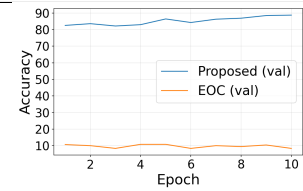
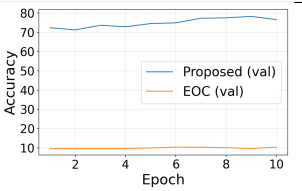
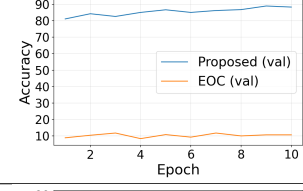
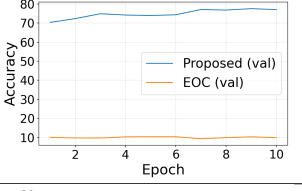
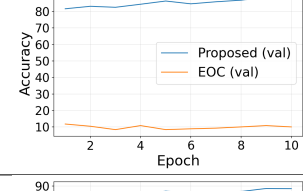
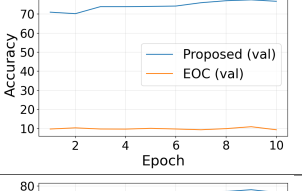
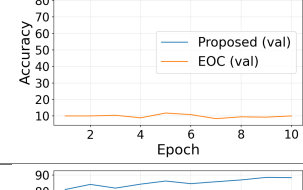
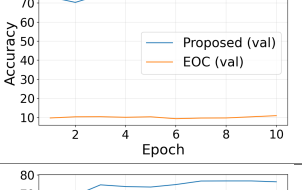
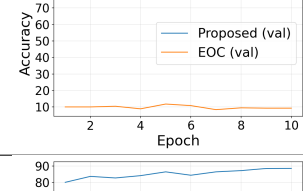
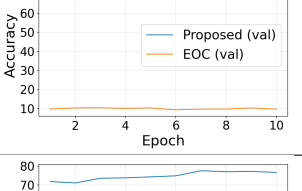
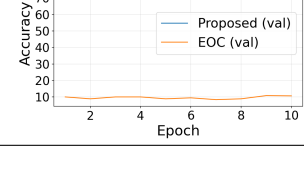
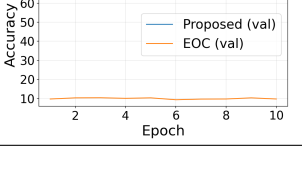
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267

Table 4: Validation accuracy on MNIST (left) and Fashion MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \arctan(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^{-2}		
10^{-1}		
1		
10^1		
10^2		
10^3		

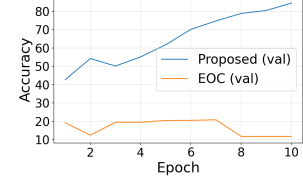
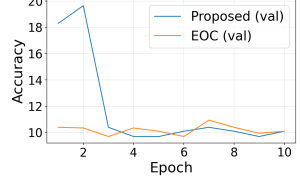
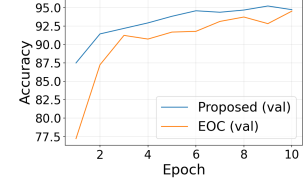
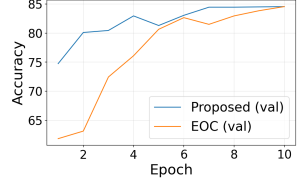
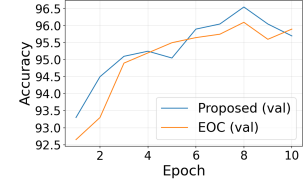
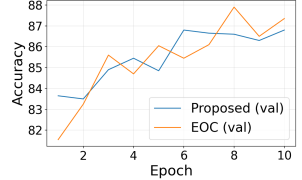
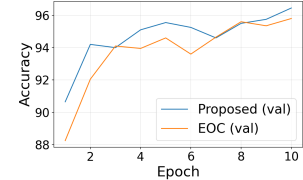
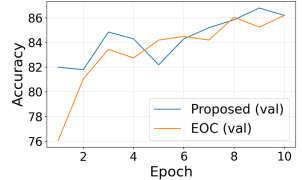
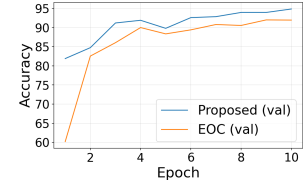
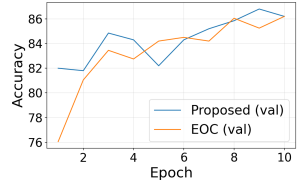
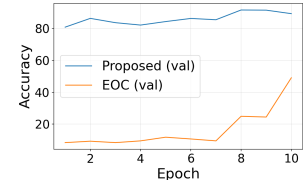
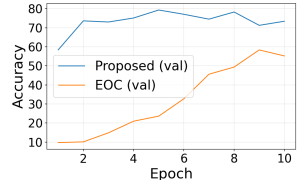
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321

Table 5: Validation accuracy on MNIST (left) and Fashion MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \arctan(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^4		
10^5		
10^6		
10^7		
10^8		
10^9		

2322
 2323
 2324
 2325
 2326
 2327
 2328
 2329
 2330
 2331
 2332
 2333
 2334
 2335
 2336
 2337
 2338
 2339
 2340
 2341
 2342
 2343
 2344
 2345
 2346
 2347
 2348
 2349
 2350
 2351
 2352
 2353
 2354
 2355
 2356
 2357
 2358
 2359
 2360
 2361
 2362
 2363
 2364
 2365
 2366
 2367
 2368
 2369
 2370
 2371
 2372
 2373
 2374
 2375

Table 6: Validation accuracy on MNIST (left) and Fashion MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \text{softsign}_2(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^{-2}		
10^{-1}		
1		
10^1		
10^2		
10^3		

2376
 2377
 2378
 2379
 2380
 2381
 2382
 2383
 2384
 2385
 2386
 2387
 2388
 2389
 2390
 2391
 2392
 2393
 2394
 2395
 2396
 2397
 2398
 2399
 2400
 2401
 2402
 2403
 2404
 2405
 2406
 2407
 2408
 2409
 2410
 2411
 2412
 2413
 2414
 2415
 2416
 2417
 2418
 2419
 2420
 2421
 2422
 2423
 2424
 2425
 2426
 2427
 2428
 2429

Table 7: Validation accuracy on MNIST (left) and Fashion MNIST (right) for a 50 layer, width 128 fully connected neural network with activation $a \text{softsign}_2(x)$. Each row corresponds to a different activation scale a , and for every a the learning rate is set to $\eta = 10^{-4}/a$ for both initializations.

a	MNIST (accuracy vs. epoch)	Fashion-MNIST (accuracy vs. epoch)
10^4		
10^5		
10^6		
10^7		
10^8		
10^9		

Common PINN configuration. For both PDE benchmarks we use fully connected PINNs with input (ξ_1, ξ_2) corresponding to either (S, τ) or (x, t) , 50 hidden layers, and width 64 in every hidden layer. The hidden activation is either $a \tanh(x)$ or $b \text{softsign}_2(x) = bx/\sqrt{1+x^2}$, and we compare the proposed initialization against the EOC initialization for scales $a, b \in \{0.1, 0.3, 0.5, 0.8, 1.0, 1.3, 1.5, 1.8, 2.0\}$. For the proposed scheme we set the structured weights using the negative-rate calibration with target $p = 0.49$ at depth $L = 50$, and for EOC we compute (σ_b, σ_w) using the Gauss–Hermite fixed point equations for the corresponding activation. In all experiments we train with Adam for iterations (600 for Black–Scholes, 600 for Burgers) followed by L-BFGS for iterations (1000 and 1000, respectively). The learning rate is scaled by the local gain of the activation via $\eta = 10^{-4}/f'(0)$.

Black–Scholes equation. We consider the Black–Scholes PDE in time-to-maturity form

$$-V_\tau + \frac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV = 0,$$

with volatility $\sigma = 0.2$, interest rate $r = 0.05$, asset price $S \in [S_{\min}, S_{\max}] = [0, 1]$, and time-to-maturity $\tau \in [0, T_{\max}] = [0, 1]$. The initial condition at $\tau = 0$ (maturity) is

$$V(S, 0) = \max(S - K, 0), \quad K = 0.5,$$

and the spatial boundary conditions are

$$V(0, \tau) = 0, \quad V(S_{\max}, \tau) = S_{\max} - Ke^{-r\tau}.$$

We sample 30,000 collocation points (S, τ) uniformly from $[0, 1] \times [0, 1]$, $n_{\text{IC}} = 4,000$ points along $\tau = 0$, and $n_{\text{SB}} = 4,000$ boundary points along $S = 0$ and $S = S_{\max}$.

Burgers’ equation. For Burgers’ equation we solve

$$u_t + uu_x - \nu u_{xx} = 0,$$

with viscosity $\nu = 0.01/\pi$, spatial domain $x \in [X_{\min}, X_{\max}] = [-1, 1]$, and time interval $t \in [0, T_{\max}] = [0, 1]$. The initial condition is

$$u(x, 0) = -\sin(\pi x),$$

and we impose homogeneous Dirichlet boundary conditions

$$u(X_{\min}, t) = 0, \quad u(X_{\max}, t) = 0.$$

We draw 20,000 collocation points (x, t) uniformly from $[-1, 1] \times [0, 1]$, use $n_{\text{IC}} = 1,000$ points along $t = 0$, and $n_{\text{SB}} = 1,000$ boundary points along $x = X_{\min}$ and $x = X_{\max}$.




Figure 38: Best PINN loss versus activation scale for the Burgers PINN (depth 50, width 64), comparing the proposed and EOC initializations. Left: $a \tanh(x)$ with $a \in \{0.1, 0.3, 0.5, 0.8, 1.0, 1.3, 1.5, 1.8, 2.0\}$. Right: $b \text{softsign}_2(x)$ with the same set of scales.