



HAL
open science

Training Integrable Parameterizations of Deep Neural Networks in the Infinite-Width Limit

Karl Hajjar, Lénaïc Chizat, Christophe Giraud

► **To cite this version:**

Karl Hajjar, Lénaïc Chizat, Christophe Giraud. Training Integrable Parameterizations of Deep Neural Networks in the Infinite-Width Limit. 2021. hal-03388815v3

HAL Id: hal-03388815

<https://hal.science/hal-03388815v3>

Preprint submitted on 6 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Training Integrable Parameterizations of Deep Neural Networks in the Infinite-Width Limit

Karl Hajjar^{*,1}, Lénaïc Chizat², and Christophe Giraud¹

¹Laboratoire de Mathématiques d’Orsay, Université Paris-Saclay, 91405 Orsay, France

²Institut de Mathématiques, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

December 6, 2021

Abstract

To theoretically understand the behavior of trained deep neural networks, it is necessary to study the dynamics induced by gradient methods from a random initialization. However, the nonlinear and compositional structure of these models make these dynamics difficult to analyze. To overcome these challenges, large-width asymptotics have recently emerged as a fruitful viewpoint and led to practical insights on real-world deep networks. For two-layer neural networks, it has been understood via these asymptotics that the nature of the trained model radically changes depending on the scale of the initial random weights, ranging from a kernel regime (for large initial variance) to a feature learning regime (for small initial variance). For deeper networks more regimes are possible, and in this paper we study in detail a specific choice of “small” initialization corresponding to “*mean-field*” limits of neural networks, which we call *integrable parameterizations* (IPs).

First, we show that under standard i.i.d. zero-mean initialization, integrable parameterizations of neural networks with more than four layers start at a stationary point in the infinite-width limit and no learning occurs. We then propose various methods to avoid this trivial behavior and analyze in detail the resulting dynamics. In particular, one of these methods consists in using large initial learning rates, and we show that it is equivalent to a modification of the recently proposed maximal update parameterization μP . We confirm our results with numerical experiments on image classification tasks, which additionally show a strong difference in behavior between various choices of activation functions that is not yet captured by theory.

1 Introduction

While artificial neural networks routinely achieve state-of-the art performance in various real-world machine learning tasks, it is still a theoretical challenge to understand why and under which conditions they perform so well. The training algorithm—typically a variant of stochastic gradient descent (SGD) with random initialization—plays a central role in this performance but is difficult to analyze for general neural network architectures, because of their highly non-linear and compositional structure. Large-width asymptotics, which have previously been considered for other purposes (Neal, 1995; Bengio et al., 2006), have recently been proposed to overcome some of these difficulties and have brought numerous insights on the training behavior of neural networks (Nitanda and Suzuki, 2017; Mei et al., 2018;

*Corresponding author: hajjarkarl@gmail.com

Jacot et al., 2018; Rotskoff and Vanden-Eijnden, 2018; Chizat and Bach, 2018; Sirignano and Spiliopoulos, 2020).

One of these insights is that the magnitude of the random weights at initialization has a dramatic impact on the learning behavior of neural networks (Chizat et al., 2019). For two-layer networks and with suitable learning rates, initializing the output layer weights with a standard deviation of $1/m$, where m is the width of the network, leads to *feature learning* when m is large, while the same network initialized with a standard deviation of $1/\sqrt{m}$ leads to the Neural Tangent Kernel (NTK) regime, a.k.a. *lazy* regime, where the network simply learns a linear predictor on top of fixed features. This observation suggests that *parameterizations*—that is, the choice of the scaling factors, with the width m , of the initial magnitude and of the learning rates of each layer of a neural network—are of fundamental importance in the theory of neural networks. While standard deep learning packages offer various choices of scale at initialization (Glorot and Bengio, 2010; He et al., 2015), those have been designed with the sole criterion in mind to have a non-vanishing first forward and backward passes for arbitrary depths. Theory now offers the tools to explore a larger space of parameterizations and study their dynamics beyond the first forward and backward passes in the infinite-width limit.

With more than two layers, the categorization of parameterizations is more subtle and there are disparate lines of work. On the one hand, some parameterizations still lead to the kernel regime, which is subject to an intense research activity (e.g., Jacot et al., 2018, 2019; Allen-Zhu et al., 2019; Du et al., 2019; Arora et al., 2019; Geiger et al., 2020a,c; Yang, 2020a). Since this regime reduces to learning a linear predictor on top of fixed features in the large width limit, this parameterization is of limited relevance to understand representation learning in networks used in practice (although it should be noted that non-asymptotic analyses reveal interesting effects, e.g., Hanin and Nica, 2019). On the other hand, there is a growing literature around parameterizations where weights are initialized with a standard deviation of $1/m$ (except for the first layer). These are often called “*mean-field*” models but we prefer to call them *integrable parameterizations* (IPs) in this work¹, in reference to the fact that sums of m terms with standard deviation of order of $1/m$ are absolutely convergent. There already exists mathematical tools to describe the evolution of the parameters of IPs in the infinite-width limit but they are not fully satisfactory to understand the properties of the learned function in the standard setting used in practice (see review in Section 1.2).

Going beyond the dichotomy between the scales $1/m$ and $1/\sqrt{m}$, Yang and Hu (2021) have exhibited, using a technique called the *Tensor Program* (Yang, 2019, 2020a,b), a general categorization of parameterizations, in particular between those which allow feature learning and those which do not. As a result from their analysis, they singled out a *maximal update* parameterization μP where, as for the NTK parameterization, the intermediate layers’ weights are initialized with a standard deviation of $1/\sqrt{m}$, but the last layer weights are initialized with a standard deviation of $1/m$: they show that with appropriate learning rates, this leads to maximal feature learning (in a certain sense). This parameterization had been previously considered in (Geiger et al., 2020b) where the authors study empirically the effect of the *scale* (Chizat et al., 2019) on learning.

In (Yang and Hu, 2021), IPs have been excluded from the analysis on the basis that they are *trivial*: if one follows the usual training procedure—which we refer to as *Naive-IP*—the network starts on a stationary point in the infinite-width limit and the learned function remains at its initial value.

¹For deep neural networks, it is somewhat arbitrary to associate the term *mean-field* with a specific choice of scaling so we believe that this term lacks precision when it comes to discussing various parameterizations.

1.1 Contributions

Our goal is to draw connections between the various lines of research discussed above, and to improve our understanding of integrable parameterizations: when and why are they trivial? How can we avoid triviality and actually learn features? What are the salient properties of the resulting networks in the infinite-width limit? To answer these questions rigorously, we leverage the Tensor Program technique developed in (Yang, 2019, 2020a,b; Yang and Hu, 2021). Specifically, our contributions are the following:

- We first show in [Theorem 3.1](#) that with learning rates constant in time, the functions learned using SGD for integrable parameterizations of neural networks with four layers or more either remain at their value at initialization or explode in the infinite-width limit when the weights are initialized using the standard zero-mean i.i.d. schemes used in practice.
- We show in [Theorem 4.1](#) that using large learning rates, which grow as a power of m , for the first gradient step—and that step only—allows SGD to escape the initial stationary point for integrable parameterizations and to initiate a non-trivial learning phase. In fact, we prove in [Theorem 4.2](#) that the resulting dynamic is equivalent to a modification of the dynamic of μP where, after the first gradient step, one subtracts the initial weights from the learned weights of the intermediate layers.
- We study two alternative ways to escape the initial stationary point for integrable parameterizations and analyze the corresponding dynamics. Removing the scale factor in $1/m$ on the bias terms allows to escape the initial stationary point when using moderately large initial learning rates. A drawback of the resulting dynamics is that its updates only depend weakly on the input data (see [Theorem 5.2](#)). On the other hand, using a non-centered law also allows to escape the initial stationary point for i.i.d. initializations without having to use large learning rates, but the dynamics become degenerate as the updates of the entries of the weight matrix in a given layer are all equal to the same fixed quantity in the infinite-width limit (see [Theorem 5.1](#)). We investigate numerically the performance of those two models and show that the aforementioned behaviors are detrimental to learning.

The code to reproduce the results of the numerical experiments can be found at:

<https://github.com/karl-hajjar/wide-networks>.

1.2 Related Work

While the study of infinitely wide neural networks has a long history (Barron, 1993; Neal, 1995, 1996; Kurková and Sanguineti, 2001; Mhaskar, 2004; Bengio et al., 2006; Bach, 2017), it is only recently that their training dynamics have been investigated. Two-layer neural networks with IP enjoy some global convergence properties (Chizat and Bach, 2018) and favorable guarantees in terms of generalization (Bach, 2017; Chizat and Bach, 2020). Going beyond two layers, Nguyen and Pham (2020) and Pham and Nguyen (2020) study the infinite-width limit of IPs and also prove global convergence results for networks with three layers or more. However, those results hold for standard zero-mean i.i.d. initialization schemes only for networks with two or three layers (which is consistent with the results of Section 3.1): for deeper networks they require non-standard (correlated) initializations.

Several other works describe the infinite-width limit of multi-layer IPs: Araújo et al. (2019) characterize the infinite-width dynamics via a model of McKean-Vlasov type, for which they prove existence and uniqueness of solutions, and Sirignano and Spiliopoulos (2021) prove a global convergence result for three-layer networks. They take the number of

units in each layer to infinity sequentially and describe the dynamics of the limit as a system of differential equations over the weights/parameters. On the other hand, Fang et al. (2020) take the infinite-width limit for all layers at once (as in Araújo et al., 2019; Nguyen and Pham, 2020; Pham and Nguyen, 2020) and describe the resulting dynamics as an ODE over functions of the features (pre-activations) of the network. It is interesting to note that Araújo et al. (2019); Sirignano and Spiliopoulos (2021); Pham and Nguyen (2020) all discuss the difficulties associated with describing the dynamics of the infinite-width of IPs with more than three layers. As noted in (Araújo et al., 2019), and appropriately addressed by Nguyen and Pham (2020); Fang et al. (2020); Sirignano and Spiliopoulos (2021), there is a separation of time scales as soon as there are two hidden layers or more, where the gradients of the intermediate layers appear to scale as m^{-2} whereas the gradients of the input and output layers appear to scale as m^{-1} , requiring separate learning rate values which can make the analysis of the infinite-width limit more difficult.

In a separate line of work, Yang and Hu (2021) provide with the Tensor Program a theoretical tool to describe the infinite-width limit of different parameterizations of neural networks and categorize them between feature learning and kernel-like behavior. However, IPs with three layers or more are left out of this categorization. Using the same tools, we show that IPs with more than four layers are indeed trivial at any time step if the initial learning rates are not appropriately scaled with m under standard zero-mean i.i.d. initializations. This closes the gap with (Nguyen and Pham, 2020) which proves global convergence results for IPs with two or three layers initialized using those standard schemes. We also demonstrate in Section 4 how scaling the initial learning rates appropriately allows to properly train an IP—inducing a feature learning regime as defined in (Yang and Hu, 2021)—and connect the resulting model with a version of the maximal update parameterization μP (Yang and Hu, 2021) where the initial weights of the intermediate layers are replaced by zero in the first update.

The setting where non-centered i.i.d. initialization laws are used is covered in (Nguyen and Pham, 2020), where it is shown that a certain collapse phenomenon occurs, namely that the updates of the entries of the weight matrix in a given layer are all equal to the same deterministic quantity in the large-width limit. We obtain a similar result in Section 5.1 using different theoretical tools.

Tensor Program vs. other formalisms. In contrast to prior literature on IPs, we do not use the description of the infinite-width limit as a composition of integral transforms. With the standard (centered i.i.d.) initializations considered in this paper, that description does not offer much insight about the limit beyond the fact that it starts on a stationary point. In order to escape this initial stationary point, we propose in this paper to amplify the random fluctuations around the limit using large initial learning rates. The strength of the Tensor Program formalism (Yang, 2019, 2020a,b; Yang and Hu, 2021) is precisely that it is able to describe rigorously the magnitudes of these fluctuations and allows us to analyze the functions learned with various choices of learning rates. This formalism relies on techniques initiated in the statistical physics literature (Bayati and Montanari, 2011; Bolthausen, 2014) that use the Gaussian conditioning technique to describe the behavior of algorithms (such as message passing) involving random matrices and nonlinearities.

1.3 Organisation of the Paper and Notations

We define and analyze integrable parameterizations in Section 3 and show that they are trivial for common choices of learning rates. In Section 4, we describe how a specific scaling of the learning rates allows to escape the initial stationary point, and further investigate the

connection between IPs with large initial learning rates and μP . In Section 5, we present two alternative modifications of IPs to escape the initial stationary point and discuss the impact of each on the learning dynamics.

We defer all the rigorous proofs of our theoretical results to the Appendix, so as to make the core message of our work stand out more clearly, and keep the flow of the results structured and easy to follow. Among other things, this prevents us from diving too deep into the Tensor Program formalism and calculations (which can be somewhat tedious and abstruse) in the main part of our work. Most proofs require heavy inductions on the time step t , and proving the induction step itself often involves inductions on l in the forward pass (from $l = 1$ to $l = L$) and in the backward pass (from $l = L$ to $l = 1$). Breaking down all these steps makes for a lengthy Appendix, but the ideas of the proof are relatively straightforward, only their proper formal writing is tedious.

Throughout the paper, for two integers p, q , we denote by $[p]$ the set $\{1, \dots, p\}$ and by $[p, q]$ the set $\{p, \dots, q\}$. We write $u \odot v$ for the Hadamard (*i.e.*, element-wise) product of two vectors u and v . We use Landau notations for comparing two real sequences (u_m) and (v_m) : we write $u_m = O(v_m)$ when there exists a constant $C > 0$ such that $|u_m| \leq C|v_m|$ for large enough m , and $u_m = \Theta(v_m)$ when we both have $u_m = O(v_m)$ and $v_m = O(u_m)$. We similarly use the O (respectively Θ) notation for two sequences of real-valued random variables (u_m) and (v_m) when, almost surely, $u_m = O(v_m)$ (respectively $u_m = \Theta(v_m)$).

2 General Setting

In this section, we introduce the general setting we consider for this work, as well as the corresponding notations. We also define precisely the notion of parameterization of a neural network and discuss examples of parameterizations commonly found in the literature.

2.1 Network and Data

Training data. We consider a training dataset $\{(\xi^{(i)}, y^{(i)})\}_{i \in [n]}$ containing n (input, output) pairs with $\xi^{(i)} \in \mathbb{R}^d$ and $y^{(i)} \in \mathbb{R}$. We will use $\xi^{(i)}$ or $y^{(i)}$ when we refer to the i -th sample in the training dataset, but use ξ_t and y_t to denote the sample(s) fed to train the network at time step t , that is for the $(t + 1)$ -th step of optimization.

Width and depth. Throughout this work, we consider a feed-forward fully connected neural network, with L hidden layers and a common width m . The total number of layers, *i.e.*, weight matrices and bias vectors will thus be $L + 1$, and most of our results are concerned with four or more layers, that is $L \geq 3$, and in the limit $m \rightarrow \infty$. The integer $l \in [L + 1]$ will always be used to index the layers of a network, and we call the **intermediate layers** of a network the layers indexed by $l \in [2, L]$ (*i.e.*, excluding input and output layers).

Activation function. We assume that all the neurons in the network share the same activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$. The activation is always taken entry-wise and for any vector $h \in \mathbb{R}^m$, we denote by $\sigma(h)$ the vector $(\sigma(h_p))_{p \in [m]} \in \mathbb{R}^m$.

Weights and forward pass. We denote by $W^l(t)$ and $B^l(t)$ respectively the weight matrix and bias vector of layer l at time step t (*i.e.*, after t steps of SGD), and thus have $W^1(t) \in \mathbb{R}^{m \times d}$, $W^l(t) \in \mathbb{R}^{m \times m}$ for $l \in [2, L]$ and $W^{L+1}(t) \in \mathbb{R}^m$. At any time step t we denote by $h_t^l(\xi)$ and $x_t^l(\xi)$ the pre-activations and activations respectively coming out of the l -th layer

when feeding input ξ to the network (with the convention that $x_t^0(\xi) = \xi$). That is

$$h_t^l(\xi) := W^l(t)x_t^{l-1}(\xi) + B^l(t), \quad \text{and} \quad x_t^l(\xi) := \sigma(h_t^l(\xi)), \quad \text{for } l \in [1, L]. \quad (2.1)$$

Output. We denote the output of the network by

$$f_t(\xi) = f(\theta(t); \xi) := (W^{L+1}(t))^\top x_t^L(\xi) + B^{L+1}(t), \quad (2.2)$$

where $\theta(t)$ denotes the set of all network parameters at time t . We often drop the dependency of the forward pass on the input ξ for brevity and simply use h_t^l, x_t^l instead of $h_t^l(\xi), x_t^l(\xi)$ as it should always be clear from the context which input is being fed to the network. Note that the weights and biases as well as all the (pre-)activations depend on the width m of the network (through their dimensions) but we omit this dependency for clarity.

Loss. We denote by ℓ the loss function used to train the network, which is a function from \mathbb{R}^2 to \mathbb{R} . The fit of a prediction \hat{y} is thus measured by $\ell(y, \hat{y})$ where y is the desired output. In all this work, we make the following assumption on the loss function ℓ , which is met by most common loss functions:

Assumption 1 (Smooth loss *w.r.t* second argument). The loss ℓ is differentiable with respect to its second argument and $\partial_2 \ell(y, \cdot)$ is a continuous function for any $y \in \mathbb{R}$.

Assumption 1 is essentially here to guarantee that if the sequence $(\hat{y}^{(m)})_{m \in \mathbb{N}^*}$ converges almost surely to some $\hat{y}^{(\infty)}$, then $\partial_2 \ell(y, \hat{y}^{(m)})$ also converges almost surely to $\partial_2 \ell(y, \hat{y}^{(\infty)})$.

2.2 Parameterizations of Neural Networks

The fact that the magnitude of the initialization of the weights and of the scale pre-factor for the weights are key quantities that determine the learning regime achieved by neural networks—and more generally by differentiable models—was pointed out in (Chizat et al., 2019). In this paper, we are interested in the behavior of neural networks when their width m goes to infinity, and we refer to as a *parameterization* of a neural network the choice of how (a) the pre-factor of the weights, (b) the variance at initialization and (c) the learning rates, evolve as a function of m . This concept was called an abc-parameterization by Yang and Hu (2021), because these dependencies are given by m^{-a} , m^{-b} and m^{-c} .

As explained by these authors, one of those three choices is actually redundant, and one can do with only the choice of two among those three scales. We take the point view considering a parameterization as a choice of scale for the pre-factor of the weights (a) and a choice of scale for the learning rates (c) while the random weights are always initialized (b) with standard i.i.d. Gaussians $\mathcal{N}(0, 1)$. We make this (arbitrary) choice as typically in the literature, different models of the infinite-width limit correspond to different choices of scales for the weights’ pre-factors, *e.g.*, NTK corresponds to a pre-factor in $1/\sqrt{m}$ while “mean-field” models correspond to a choice of pre-factor in $1/m$ for the weights. We thus define below ac-parameterizations which are a slight variation of the abc-parameterizations introduced in (Yang and Hu, 2021).

Definition 2.1. (ac-parameterization). An ac-parameterization of an L -hidden layer fully-connected neural network is a choice of scalar exponents (a_1, \dots, a_{L+1}) , and (c_1, \dots, c_{L+1}) such that for any layer $l \in [L + 1]$,

- (i) the **learnable weights** (*i.e.*, those over which we optimize) are initialized with independent standard Gaussian random variables $w_{j,q}^l(0) \sim \mathcal{N}(0, 1)$, i.i.d. over (l, j, q) , *i.e.*, $w^l(0) = U^l$ with $(U^l)_{l \in [L+1]}$ independent random matrices with i.i.d. standard Gaussian entries,

- (ii) the **learnable biases** are initialized independently of the weights, with $b_j^l(0) \sim \mathcal{N}(0, 1)$, i.i.d. over (l, j) , i.e., $b^l(0) = v^l$ with $(v^l)_{l \in [L+1]}$ independent standard Gaussian random vectors, independent of U^l ,
- (iii) the **effective weights** $W^l(t)$ used to compute the pre-activations at time t are $W^l(t) = m^{-a_l} w^l(t)$, and the **effective biases** are $B^l(t) = m^{-a_l} b^l(t)$, so that the pre-activations are

$$h_t^l = W^l(t) x_t^{l-1} + B^l(t) = m^{-a_l} \left(w^l(t) \sigma(h_t^{l-1}) + b^l(t) \right), \quad l \in [1, L],$$

and the output is

$$f(\theta(t); \xi) = m^{-a_{L+1}} \left(w^{L+1}(t)^T \sigma(h_t^L(\xi)) + b^{L+1}(t) \right),$$

- (iv) the $(t+1)$ -th update of learnable weights and biases is given by the update rules

$$\begin{aligned} \Delta w^l(t+1) &:= w^l(t+1) - w^l(t) = -\eta m^{-c_l} \nabla_{w^l} \ell(y_t, f(\theta(t); \xi_t)), \\ \Delta b^l(t+1) &:= b^l(t+1) - b^l(t) = -\eta m^{-c_l} \nabla_{b^l} \ell(y_t, f(\theta(t); \xi_t)), \end{aligned}$$

where $\theta(t) = \{(w^1(t), b^1(t)), \dots, (w^{L+1}(t), b^{L+1}(t))\}$ is the full set of all network parameters, (ξ_t, y_t) represent the input(s) and target(s) to the network at step t and $\eta \in \mathbb{R}_+^*$ is the scalar part of the learning rate which does not depend on m and which we call the **base learning rate**. We denote by $\eta_l := \eta m^{-c_l}$ the full learning rate for layer l .

Remark 2.1.

1. Compared to the definition of (Yang and Hu, 2021), we allow for different values of c_l at different layers and remove the redundant initialization scale (the b in abc-parameterizations). Any abc-parameterization with constant c for all layers (as presented in Yang and Hu, 2021) can be recovered (same effective weights and biases at any time step) with an ac-parameterization with individual learning rates at each layer via the re-parameterization $a_l \leftarrow a_l + b_l$, $b_l \leftarrow 0$, $c_l := c - 2b_l$.
2. As we study the infinite-width limit $m \rightarrow \infty$, we need to consider an infinite number of random weights at initialization. To this end, we consider for any $l \in [2, L]$, two infinite lists of i.i.d. standard Gaussian variables, independent of each other: $(U_{jq}^l)_{j,q \in \mathbb{N}^*}$ and $(v_j^l)_{p \in \mathbb{N}^*}$, and often simply call, by an abuse of notations, $U^l = (U_{jq}^l)_{1 \leq j,q \leq m}$ for the corresponding matrix at width m and $v^l = (v_j^l)_{1 \leq j \leq m}$ the corresponding bias vector at width m . We proceed similarly at initialization for the input weights U^1 and the output vector U^{L+1} .
3. The $(t+1)$ -th update of the effective weights is given by $\Delta W^l(t+1) := W^l(t+1) - W^l(t) = -\eta m^{-(2a_l+c_l)} \nabla_{w^l} \ell(y_t, f(\theta(t); \xi_t))$, and the update of the effective biases by $\Delta B^l(t+1) := B^l(t+1) - B^l(t) = -\eta m^{-(2a_l+c_l)} \nabla_{b^l} \ell(y_t, f(\theta(t); \xi_t))$

Examples of ac-parameterizations:

NTK parameterization. For the NTK parametrization (Jacot et al., 2018) the scaling is $a_1 = 0$ for the input layer, and $a_l = 1/2$ for all the other layers $l \in [2, L+1]$. The scaling of the learning rates is $c_l = 0$ for all layers. Neural networks in the NTK parametrization have been shown to behave as kernel methods in the infinite-width limit (Jacot et al., 2018; Yang, 2020a) and there is no feature learning in that limit.

μP . To avoid the lazy training phenomenon arising in the NTK parameterization, [Yang and Hu \(2021\)](#) propose to adjust the scale of the output layer by setting $a_{L+1} = 1$, while keeping $a_1 = 0$ and $a_l = 1/2$ for the intermediate layers $l \in [2, L]$. The learning rates are appropriately adjusted: $c_l = -1$ for any layer l . With this parameterization, [Yang and Hu \(2021\)](#) show that feature learning (see Definition B.1 in Appendix B.3 for a precise statement) occurs at every layer.

Integrable Parameterizations (IPs). The limits investigated in [Araújo et al. \(2019\)](#); [Sirignano and Spiliopoulos \(2021\)](#); [Pham and Nguyen \(2020\)](#); [Weinan and Wojtowytsch \(2020\)](#) are associated to a scale multiplier in $1/m$ for all layers except the first one. This corresponds to the choice $a_1 = 0$ and $a_l = 1$ for $l \in [2, L + 1]$. We choose the adjective “*integrable*” in reference to the absolute convergence of sums of the form $(1/m) \sum_q x_q$ for i.i.d. random variables with finite expectation. Integrable parameterizations really refer to a class of ac-parameterizations, because various choices for the learning rate exponents c_l are admissible.

Naive-IP. In the mean-field literature, integrable parameterizations often come with the standard learning rates corresponding to $c_1 = c_{L+1} = -1$ for the input/output layers and $c_l = -2$ for the intermediate layers $l \in [2, L]$, see *e.g.*, ([Araújo et al., 2019](#), Remark 3.4), ([Fang et al., 2020](#), Algorithm 1), ([Weinan and Wojtowytsch, 2020](#), Lemma 5.1), and ([Sirignano and Spiliopoulos, 2021](#), Equation 4.3). Mean-field models with these learning rates are the natural counterparts of the infinite-width limits where sums are replaced by integrals, and we call the integrable parameterization with this specific choice of learning rates the *Naive Integrable Parameterization*.

When $L = 1$, μP and the Naive-IP coincide. For deeper networks, in the setting of abc-parameterizations described in ([Yang and Hu, 2021](#)), μP and Naive-IP correspond to the same parameterization (same values for a and c) except that the weights of the intermediate layers are initialized with a standard deviation of $1/m$ for Naive-IP instead of $1/\sqrt{m}$ for μP , that is they are downscaled by $1/\sqrt{m}$ compared to μP . In Section 4.2, we show that there is also a close relationship between μP and IP with large initial learning rates.

We give below an intuitive explanation for the choice $c_1 = c_{L+1} = -1$ and $c_l = -2$ for $l \in [2, L]$ for the scaling of the learning rates in Naive-IP. For $l \in [2, L]$, we have $h_t^l = m^{-1}(w^l(t)x_t^{l-1} + b^l(t))$, so that $\nabla_{w^l} f_t(\xi_t) = m^{-1}(\nabla_{h^l} f_t(\xi_t))(x_t^{l-1})^\top$. In addition $\nabla_{w^{L+1}} f_t(\xi_t) = x_t^L/m$ and $\nabla_{w^1} f_t(\xi_t) = (\nabla_{h^1} f_t(\xi_t))(\xi_t)^\top$. So for one step of SGD:

$$\begin{aligned} \Delta W^1(t+1)\xi_{t+1} &= -\eta \partial_2 \ell(y_t, f_t(\xi_t))(\xi_t^\top \xi_{t+1}) m^{-(1+c_1)} (m \nabla_{h^1} f_t(\xi_t)). \\ \Delta W^l(t+1)x_{t+1}^{l-1} &= -\eta \partial_2 \ell(y_t, f_t(\xi_t)) m^{-(2+c_l)} \frac{(x_t^{l-1})^\top x_{t+1}^{l-1}}{m} (m \nabla_{h^l} f_t(\xi_t)), \quad \text{for } l \in [2, L] \\ (\Delta W^{L+1}(t+1))^\top x_{t+1}^L &= -\eta \partial_2 \ell(y_t, f_t(\xi_t)) m^{-(1+c_l)} \frac{(x_t^L)^\top x_{t+1}^L}{m}. \end{aligned} \tag{2.3}$$

In addition, from the equations of backpropagation, we get

$$\nabla_{h_t^L} f_t(\xi_t) = \frac{1}{m} w^{L+1}(t) \odot \sigma'(h_t^L) \quad \text{and} \quad \nabla_{h^l} f_t(\xi_t) = \frac{(w^l(t))^\top \nabla_{h_t^{l+1}} f_t(\xi_t)}{m} \odot \sigma'(h_t^l),$$

for $l \in [1, L - 1]$, so that, by a simple induction, $\nabla_{h^l} f_t(\xi_t) = O(1/m)$ for $l \in [1, L]$. In addition, the averaged inner products $(x_t^{l-1})^\top x_{t+1}^{l-1}/m$ in Equation (2.3) converge as $m \rightarrow \infty$. This point is somewhat technical and is handled within the framework of the Tensor Program. The choice of c_l in Naive-IP thus ensures that the updates are $O(1)$ when m goes to infinity.

We conclude this section by giving the definition of a training routine which consists in the combination of the base learning rate, the sequence of training samples and a loss function:

Definition 2.2 (Training routine). A training routine is the list consisting of the base learning rate $\eta > 0$, $(a_l, c_l)_{l \in [L+1]}$ in the ac-parameterization, the loss ℓ and the sequence of training samples $(\xi_0, y_0), \dots, (\xi_{T-1}, y_{T-1})$ used to train a network for T steps.

3 Deep Networks with Naive Integrable Parameterization are Trivial

In this section, we point out that, in the wide limit, neural networks in the Naive-IP remain at their initial value. We then prove that no choice for the learning rates exponents $(c_l)_{l \in [L+1]}$ which is constant in time can induce non-degenerate learning.

3.1 No learning in Deep Networks with Naive Integrable Parameterization

To start with, we show that the functions learned by networks with more than four layers in the naive integrable parameterization, as described in prior work (Araújo et al., 2019; Rotskoff and Vanden-Eijnden, 2019; Fang et al., 2020; Nguyen and Pham, 2020; Weinan and Wojtowytsch, 2020; Sirignano and Spiliopoulos, 2021), remain at their value at initialization in the infinite-width limit: they are identically equal to zero at any time step. Our proof of this result is based on the Tensor Program framework (Yang, 2020b; Yang and Hu, 2021), which requires some regularity assumptions on the activation function.

Definition 3.1. (Pseudo-Lipschitz functions). A function $\psi : \mathbb{R}^k \rightarrow \mathbb{R}$ is pseudo-Lipschitz of degree $p > 0$ if there exists a constant $K > 0$, such that, for any $x, y \in \mathbb{R}^k$,

$$|\psi(x) - \psi(y)| \leq K \|x - y\| \left(1 + \sum_{r=1}^k |x_r|^p + \sum_{r=1}^k |y_r|^p \right).$$

A function is pseudo-Lipschitz, if it is pseudo-Lipschitz of degree p for some $p > 0$.

In particular, functions with polynomially bounded weak derivatives are pseudo-Lipschitz. In the next proposition, we require the activation function σ and its derivative to be pseudo-Lipschitz.

Assumption 2 (Smooth activation). The activation function σ is differentiable and both σ and its derivative σ' are pseudo-Lipschitz and not identically zero.

Proposition 3.1 (Naive-IP is trivial). Let $L \geq 3$ and consider the naive integrable parameterization of a network with L -hidden layers, and an activation function satisfying Assumption 2 and $\sigma(0) = 0$. Then, for any training routine which has a loss satisfying Assumption 1, the function learned by SGD remains at its value at initialization in the infinite-width limit:

$$\forall t \geq 0, \quad \forall \xi \in \mathbb{R}^d, \quad \lim_{m \rightarrow \infty} f_t(\xi) = \lim_{m \rightarrow \infty} f_0(\xi) = 0 \quad \text{almost surely.}$$

Remark 3.1.

1. In the above statement, “almost surely” is relative to the randomness of the initialization.

2. The smoothness Assumption 2 on σ is met by common activation functions such as GeLU (Hendrycks and Gimpel, 2016), ELU (Clevert et al., 2016), tanh and the sigmoid activations, but it excludes ReLU and all the other variants of Leaky ReLU. This assumption is required to apply (Yang and Hu, 2021, Theorem 7.4) (which we recall in Appendix B.2) which is the main theoretical result of the Tensor Program series (Yang, 2019, 2020a,b; Yang and Hu, 2021), but the result is likely to hold with weaker assumptions, as observed numerically in Section 6, and we leave this for future work.
3. The assumption $\sigma(0) = 0$ is met by the activation functions mentioned above (except the sigmoid) and is necessary to prove that the network does not move at any layer. Without this assumption, learning is degenerate but not trivial at all layers. It is trivial at step $t = 1$ at all layers except the last two: the coordinates of h_1^L and $f_1(\xi)$ converge, with m , to quantities which are not 0 but which are independent of the input ξ to the network, similarly to the effect described in Section 5.2.

The proof of Proposition 3.1, presented in Appendix D, proceeds by induction over t to show that the forward and backward passes vanish at any time step. For any time t , we proceed again by induction over l (from $l = 1$ to $l = L + 1$ for the forward pass and from $l = L + 1$ to $l = 1$ for the backward pass) to prove this vanishing occurs given the magnitudes of the previous forward and backward passes. The informal idea of the proof is the following: essentially, the multiplications of the activation vectors by $m^{-1/2}U^l$ yield vectors whose coordinates are distributed as a Gaussian with finite variance as $m \rightarrow \infty$ for $l \geq 2$ (see Appendix B.1.1 for more details). At initialization, since $w^l(0) = m^{-1}U^l$ for $l \geq 2$ for IPs, the coordinates of h_0^l converge towards 0 as fast as $m^{-1/2}$ and that of x_0^l towards $\sigma(0)$ for σ continuous at 0. For the same reasons, $f_0(\xi_0)$ converges to 0. In the first backward pass, multiplications by $(W^l(0))^T$ also yield vectors whose coordinates are in $O(m^{-1/2})$. In contrast to the forward pass, these scales propagate from $l = L$ to $l = 1$ and thus compound with depth, and since the last layer’s gradient x_0^L/m is in $O(m^{-1})$, all the gradients’ coordinates vanish as $m \rightarrow \infty$ and there is no learning. This reasoning can be repeated at later time steps as there are no correlations between the initial weight matrices and the vectors they multiply because of the degeneracy of the (pre)-activations (their coordinates become equal to the constant $\sigma(0)$ as $m \rightarrow \infty$). Those informal calculations are made rigorous by the Tensor Program.

Proposition 3.1 shows that the parameters of neural networks in the integrable parameterization are stuck in a stationary point of the objective function in the infinite-width limit, and no learning occurs. It might appear obvious that using larger learning rates to correct the scale with m of the weight updates can avoid this pitfall, but as discussed in the following Section 3.2—where we study which choices of learning rates can lead to stable learning with homogeneous activation functions—the issue is more subtle.

3.2 No stable learning with learning rates constant over time

As m grows, to compensate the vanishing gradients in the first SGD step, one can use larger learning rates than in the Naive-IP. Yet, as explained below, exponents $(c_l)_{l \in [L+1]}$ for the learning rates which allow to escape the stationary point at initialization will induce an explosion of the pre-activations, if the same values of the exponents are used in the subsequent gradient steps. Indeed, the next informal statement of Theorem 3.2 shows that, with IPs, one cannot have non-trivial and stable learning with learning rate scales c_l constant in time.

Theorem 3.1 (Informal). *Consider an L -hidden layer fully-connected neural network with $L \geq 3$ in the integrable parameterization. Assume that the contributions of the first and*

second updates $\Delta W^l(1)x_1^{l-1}$ and $\Delta W^l(2)x_2^{l-1}$ are non-vanishing and non-exploding with m at every layer l . Then, the learning rates scales c_l cannot have the same value at $t = 0$ and $t = 1$.

In a nutshell, one needs large learning rates to escape the initial stationary point, but keeping those initial values at later time steps would make the pre-activations blow-up as $m \rightarrow \infty$. The formal version of the previous [Theorem 3.1](#) is given in [Theorem 3.2](#) below. For this formal statement, we introduce some definitions and assumptions.

Assumption 3 (Smooth non-negative homogeneous activation). The activation function σ is non-negative, not identically zero and it is positively p -homogeneous with $p \geq 2$, i.e., $\sigma(\lambda z) = \lambda^p \sigma(z)$ for any $\lambda > 0$ and $z \in \mathbb{R}$. Additionally, σ has faster growth on the positive part of the real line: $\exists z > 0$ s.t. $\sigma(z) > \sigma(-z)$.

Remark 3.2.

1. While the homogeneity assumption is core to the calculation of scales with integrable parameterization, the fact that $p \geq 2$, and that σ is non-negative and has faster growth on the positive part of the real line are simply here to avoid cumbersome technical difficulties in the proofs. It is clear that ReLU^p satisfies [Assumption 3](#) for any $p \geq 2$.
2. With the assumption that $p \geq 2$, σ also satisfies [Assumption 2](#), so that the rules of the Tensor Program can be applied.

Definition 3.2 (Scales of first updates with homogeneity). Let $p > 0$. We define the following exponents:

$$\gamma_1(p) = \gamma_{L+1}(p) = -\frac{1}{2} \left(1 + \sum_{k=0}^{L-1} p^k \right),$$

$$\text{and } \gamma_l(p) = -1 - \frac{1}{2} \sum_{k=0}^{L-1} p^k, \quad \text{for } l \in [2, L].$$

Theorem 3.2 (Formal version). Consider an L -hidden layer fully-connected neural network with $L \geq 3$ in the integrable parameterization, and with no bias terms, except for the first layer. Assume that the activation function σ satisfies [Assumption 3](#), the loss ℓ satisfies [Assumption 1](#) and that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) \neq 0$, and $\lim_{m \rightarrow \infty} \partial_2 \ell(y_1, f_1(\xi_1)) \neq 0$ almost surely. Assume further that $\xi_0, \xi_1, \xi_2 \in \mathbb{R}^d$ are all distinct vectors such that $\xi_0^\top \xi_1 \neq 0$ and $\xi_1^\top \xi_2 \neq 0$. Finally assume that:

$$\begin{cases} \frac{1}{m} \|\Delta W^l(1)x_1^{l-1}\|^2 = \Theta(1), & l \in [1, L] \\ (\Delta W^{L+1}(1))^\top x_1^L = \Theta(1) \end{cases} \quad (3.1)$$

and

$$\begin{cases} \frac{1}{m} \|\Delta W^l(2)x_2^{l-1}\|^2 = \Theta(1), & l \in [1, L] \\ (\Delta W^{L+1}(2))^\top x_2^L = \Theta(1) \end{cases} \quad (3.2)$$

Then, one necessarily has that:

- (i) at $t = 0$, $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ (see [Definition 3.2](#)),
- (ii) at $t = 1$, $c_1 = c_{L+1} = -1$, and $c_l = -2$ for $l \in [2, L]$.

Let us comment briefly on the hypotheses of Theorem 3.2. The proof of Theorem 3.2 relies on an analysis of the SGD steps involving both (Yang and Hu, 2021, Theorem 7.4) and the homogeneity property of the activation function. The requirement that $p \geq 2$ allows to satisfy the smoothness assumption of (Yang and Hu, 2021, Theorem 7.4) and the removal of the bias terms allows to fully exploit homogeneity. In Section 6, we numerically check that the result still holds with $\sigma = \text{ReLU}$, which is $p = 1$ homogeneous. The corresponding scales for the learning rates in the ReLU case are $\gamma_1(1) = -(L+1)/2$, $\gamma_l(1) = -(L+2)/2$ and $\gamma_{L+1}(1) = -(L+1)/2$.

We give below an informal explanation for the values of the learning rates appearing in Theorem 3.2 in the case of a positively 1-homogeneous activation function. As previously mentioned in Section 3.1, each multiplication by $W^l(0) = m^{-1}U^l$ or its transpose yields a factor in $m^{-1/2}$ for $l \geq 2$. Because of the homogeneity property, this scale propagates from layer to layer starting from layer 2, and the coordinates of h_0^l and x_0^l are thus in $\Theta(m^{-(l-1)/2})$ for $l \in [1, L]$. For the backward pass, the first gradient $\nabla_{x^L} f_0(\xi_0) = U^{L+1}/m$ has coordinates in $\Theta(m^{-1})$, and, as already discussed in Section 3.1, from $l = L$ to $l = 2$, each multiplication by $(W^l(0))^T$ yields an additional factor in $m^{-1/2}$ and those compound with depth so that the coordinates of $\nabla_{h^l} f_0(\xi_0)$ are in $\Theta(m^{-1}m^{-(L-l)/2})$. Therefore, calling $\tilde{x}_0^l := m^{(l-1)/2}x_0^l$, and $d\tilde{h}_0^l := m^{1+(L-l)/2}\nabla_{h^l} f_0(\xi_0)$, we have after the first weight update

$$\begin{aligned} \Delta W^1(1)\xi_1 &= -\eta\partial_2\ell(y_0, f_t(\xi_0))(\xi_0^T\xi_1)m^{-c_1}m^{-(L+1)/2}d\tilde{h}_0^1, \\ \Delta W^l(1)x_1^{l-1} &= -\eta\partial_2\ell(y_0, f_0(\xi_0))m^{-c_l}m^{-2}m^{-(L-l)/2-(l-2)/2}\frac{(\tilde{x}_0^{l-1})^T x_1^{l-1}}{m}d\tilde{h}_0^l, \quad l \in [2, L], \\ (\Delta W^{L+1}(1))^T x_1^L &= -\eta\partial_2\ell(y_0, f_0(\xi_0))m^{-c_l}m^{-1}m^{-(L-1)/2}\frac{(\tilde{x}_0^L)^T x_1^L}{m}. \end{aligned}$$

Since $d\tilde{h}_0^l$ and \tilde{x}_0^l have coordinates in $\Theta(1)$ by design, and since averaged inner products of the type $(\tilde{x}_0^{l-1})^T x_1^{l-1}/m$ converge to finite expectations (by the rules of the Tensor Program, see Yang and Hu, 2021, Theorem 7.4), we see that the choice $c_1 = -(L+1)/2$, $c_l = -(L+2)/2$ for $l \in [2, L]$, and $c_{L+1} = -(L+1)/2$ is the only way to ensure that the updates induce contributions which have coordinates in $\Theta(1)$ at $t = 1$. Given this choice for the learning rate scales c_1, \dots, c_{L+1} at $t = 0$, we readily get that the coordinates of h_1^l and x_1^l are in $\Theta(1)$ because the contributions $W^l(0)x_1^{l-1}$ have coordinates in $O(m^{-1/2})$ for intermediate layers, and in $O(1)$ for the input and output layers. From the Equations (2.3) with $t = 1$, we see that for the second gradient step, $m\nabla_{h^l} f_1(\xi_1)$ has coordinates in $\Theta(1)$ because the multiplications by $(W^l(1))^T$ do not yield a factor in $m^{-1/2}$ due to the scale correction introduced in the first update. At $t = 1$, this leads to the choice $c_1 = c_{L+1} = -1$, and $c_l = -2$ for $l \in [2, L]$, in order to have update contributions with coordinates in $\Theta(1)$ at $t = 2$. These informal calculations are made rigorous in the proof of Theorem 3.2 using the Tensor Program (Yang, 2020b).

4 Large Initial Learning Rates Induce Learning

In this section, we show that with positively homogeneous activation functions, using large initial learning rates (polynomial in m) allows the network to escape from the initial stationary point and to initiate a non-trivial training phase in the infinite-width limit. Because we use the homogeneity property extensively for our results, in all this section, as in Section 3.2, we consider a version of integrable parameterizations where the bias terms are removed except for the first layer.

As observed in Section 3.2, beyond the fact that IPs require large learning rates (for the first gradient step) to be trained, one crucial characteristic of IPs is that no choice of learning rate scales (c_l) which are constant in time can induce a favorable learning behavior: one has

to first use large learning rates to escape the stationary point at initialization ($t = 0$) and then revert to the Naive-IP learning rates for $t \geq 1$ to induce stable learning.

Definition 4.1 (IP with large initial learning rates). Let σ be a positively p -homogeneous activation function with $p > 0$. We define *the integrable parameterization with large initial learning rates* (IP-LLR) as the integrable parameterization of an L -hidden layer fully connected-network with activation σ such that:

- (i) At $t = 0$: $c_l = \gamma_l(p)$, for $l \in [1, L + 1]$;
- (ii) At $t \geq 1$: $c_1 = c_{L+1} = -1$ and $c_l = -2$, for $l \in [2, L]$,

where the values of the $\gamma_l(p)$ are given in Definition 3.2.

Remark 4.1.

1. The definition means that $\Delta w^l(1) = -\eta m^{-\gamma_l(p)} \nabla_{w^l} \ell(y_0, f_0(\xi_0))$ for the first weight update after the forward-backward pass at time $t = 0$, and for $t \geq 1$, the $(t + 1)$ -th weight update is $\Delta w^l(t + 1) = -\eta m^{-2} \nabla_{w^l} \ell(y_t, f_t(\xi_t))$ for $l \in [2, L]$, and $\Delta w^1(t + 1) = -\eta m^{-1} \nabla_{w^1} \ell(y_t, f_t(\xi_t))$, $\Delta w^{L+1}(t + 1) = -\eta m^{-1} \nabla_{w^{L+1}} \ell(y_t, f_t(\xi_t))$ after the forward-backward pass at time t .
2. We give the definition with an arbitrary degree of homogeneity p (the values of the $\gamma_l(p)$ are given in Definition 3.2) as for some theorems where we use the Tensor Program for the proof, we need sufficient smoothness of the activation function, which is achieved only when $p \geq 2$, but we always use $\sigma = \text{ReLU}$ (which corresponds to $p = 1$) in our informal derivations and numerical experiments. Note that since the values of c_1, \dots, c_{L+1} at $t = 0$ depend on p , the definition of an IP-LLR parameterization also implicitly depends on the degree of homogeneity p .
3. Since $a_1 = 0$ for IPs, we leverage the homogeneity property only for layers $l \in [2, L]$ (see Appendix F.2 for more details), so that we might as well assume $L \geq 2$ whenever we study IP-LLR.

4.1 Non-trivial and Stable Learning for Integrable Parameterizations

Theorem 4.1 (Non-trivial and non-exploding learning of IP-LLR). *Consider the IP-LLR parameterization of an L -hidden layer neural network with no bias terms, except for the first layer, and with an activation function σ satisfying Assumption 3 and a loss function ℓ satisfying Assumption 1. Let $\xi \in \mathbb{R}^d$ be an input to the network, and assume $\partial_2 \ell(y_0, 0) \neq 0$. Then, one has:*

- (i) $f_0(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} 0$.
- (ii) $f_1(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{f}_1(\xi)$, $0 < |\overset{\circ}{f}_1(\xi)| < \infty$ a.s.
- (iii) $f_2(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{f}_2(\xi)$, $|\overset{\circ}{f}_2(\xi)| < \infty$ a.s.

Remark 4.2.

1. We show in our numerical experiments (see Section 6) that with $\sigma = \text{ReLU}$ (*i.e.*, $p = 1$), the choice of learning rates for IP-LLR is indeed able to induce learning for networks deeper than four layers without creating instabilities.

2. A similar result could be obtained with more general assumptions on the activation function σ , namely that σ is twice differentiable almost everywhere and that $\sigma(0) = 0$ and $\sigma'(0) \neq 0$ (which is the case for many activation functions such as GeLU, ELU, tanh), but at the cost of a more technical proof. The idea in this case is that because of the scaling in $1/m$ which makes the forward pass vanish at initialization, one can recover the homogeneity property by linearizing σ around 0: $\sigma(h) \simeq \sigma'(0)h$. This linearization also provides the right value $|\sigma'(0)|^{-1}$ for the standard deviation of the initial Gaussians in order to avoid vanishing or explosion at initialization with the depth L . See more details in Remark F.3.
3. For positively p -homogeneous activations with $p \geq 2$, we have $\sigma'(0) = 0$ and the behavior of the network is inherently different from that of a network where the first forward pass can effectively be linearized (the setting described in the previous point). This difference appears clearly in the numerical experiments presented in Section 6 where we also discuss the reasons for such a qualitatively different behavior.
4. In IP-LLR, the initial gradient direction will be determined by the first sample (ξ_0, y_0) fed to the network. To avoid giving too much importance to a single sample, one can in practice average the gradients over a batch of many training samples instead, which is what we do in our numerical experiments in Section 6.

The idea of the proof essentially lies in the informal calculations of Section 3.2 which are made rigorous using the framework of the Tensor Program. Point (ii) stems from the fact that at $t = 1$, the output is the difference between two expectations in the limit $m \rightarrow \infty$, which can both be shown to be different from 0 and of opposite signs.

4.2 IP-LLR is a Modified μ P

In this section, we analyze the behavior of IP-LLR more in detail and show that this model is actually equivalent to a modification of μ P where the initial weights are removed from the first weight update for all of the intermediate layers. We first show an equivalence at finite-width in Section 4.2.1 with mild assumptions, and then extend those results to the infinite-width limit in Section 4.2.2 with slightly more restrictive assumptions on the activation function σ . Since we study the IP-LLR parameterization, we consider positively p -homogeneous activation functions, and only the degree of homogeneity allowed will vary between Sections 4.2.1 and 4.2.2. In short, the main idea behind this equivalence is that since IP-LLR and μ P are both designed to have maximal update contributions at $t = 0$, they will induce the same update at initialization, and the only difference at later time steps is that the initial weights of IP-LLR contribute vanishingly to the pre-activations whereas those of μ P contribute in $\Theta(1)$.

4.2.1 Finite-Width Equivalence

As explained in Section 2.2 in the examples of ac-parameterizations, from the point of view of abc-parameterizations (see Yang and Hu, 2021), both μ P and Naive-IP follow the same training procedure for the effective weights W^l , the only difference being the standard deviation at initialization which is downscaled by $1/\sqrt{m}$ for Naive-IP compared to μ P. In this regard, since IP-LLR is a modification of Naive-IP where large learning rates are used at initialization, it comes as no surprise that the learning dynamics of IP-LLR and μ P are closely related. We detail this relationship in this section.

Recall that for μ P one has $W_{\mu\text{P}}^1(0) = U^1$, $W_{\mu\text{P}}^l(0) = m^{-1/2}U^l$ for $l \in [2, L]$, and $W_{\mu\text{P}}^{L+1}(0) = m^{-1}U^{L+1}$ whereas for any integrable parameterization, one has $W_{\text{IP}}^1(0) = U^1$,

$W_{\text{IP}}^l(0) = m^{-1}U^l$ for $l \in [2, L + 1]$. Consider the following *hybrid parameterization* (HP) which consists in training with the maximal update parameterization μP all along, but simply replacing, for all intermediate layers $l \in [2, L]$, the first update $W^l(1) = W^l(0) + \Delta W^l(1)$ by $W^l(1) = m^{-1}U^l + \Delta W^l(1)$. In other words, this simply consists in using the weight pre-factors of μP for the intermediate layers in the initial forward and backward passes, and then using the pre-factors from IP for the initial weights of the intermediate layers in any subsequent update.

Proposition 4.1 (Finite width equivalence between IP-LLR and HP). Consider the IP-LLR and HP parameterizations with a p -homogeneous activation function σ with $p \geq 1$ and without any bias term except at the first layer. Let us sub/super-script the variables of each model with IP and HP respectively. Assume the full sequence of training samples $(\xi_0, y_0), \dots, (\xi_s, y_s), \dots$ and the loss ℓ are the same for both parameterizations. Assume further that $\partial_2 \ell(y_0, f_0^{\text{HP}}(\xi_0)) \neq 0$, and denote by η the base learning rate of the IP-LLR parameterization. Finally consider the following schedule for the base learning rate of HP:

$$\begin{aligned} \eta_{\text{HP}}(0) &= \frac{\partial_2 \ell(y_0, f_0^{\text{IP}}(\xi_0))}{\partial_2 \ell(y_0, f_0^{\text{HP}}(\xi_0))} \eta, \\ \eta_{\text{HP}}(s) &= \eta, \quad s \geq 1. \end{aligned}$$

Then one has:

$$\forall t \geq 1, \quad \forall \xi \in \mathbb{R}^d, \quad f_t^{\text{HP}}(\xi) = f_t^{\text{IP}}(\xi).$$

The proof, presented in Appendix J.1, simply shows inductively that the effective weight matrices for both models are equal for all $t \geq 1$. Since the Tensor Program is not needed here as we consider only finite-width networks, we can work with any positively homogeneous activation function (not necessarily smooth, so that $p = 1$ is not precluded).

4.2.2 Infinite-Width Equivalence

Similarly to HP, we now consider another hybrid parameterization where the initial weights $W^l(0)$ are simply replaced by 0 in the first update of the intermediate layers. We thus consider the following *hybrid parameterization with zero re-initialization* (HPZ): we train with μP all along, but simply replace, for all intermediate layers $l \in [2, L]$, the first update $W^l(1) = W^l(0) + \Delta W^l(1)$ by $W^l(1) = \Delta W^l(1)$. In other words, this simply consists in using the weight pre-factors of μP for the intermediate layers in the initial forward and backward passes, and then forgetting the contribution of the initial weights of the intermediate layers in any subsequent update. As already discussed in Section 3.1, the contribution of the initial weights of the intermediate layers $m^{-1}U^l$ vanishes as $m \rightarrow \infty$ for IP, so that HPZ is simply the infinite-width equivalent of HP.

Theorem 4.2 (HPZ and IP-LLR are equivalent). Consider the IP-LLR and HPZ parameterizations with a p -homogeneous activation function σ with $p \geq 2$, and with no bias terms except at the first layer. Let us sub/super-script the variables of each models with IP and HPZ respectively. Assume that the training routine is the same for both parameterizations, and assume further that the loss ℓ satisfies Assumption 1. Then, one has:

$$\forall t \geq 0, \quad \forall \xi \in \mathbb{R}^d, \quad \lim_{m \rightarrow \infty} f_t^{\text{HPZ}}(\xi) = \lim_{m \rightarrow \infty} f_t^{\text{IP}}(\xi) \quad \text{almost surely.}$$

The proof, presented in Appendix J.2, proceeds by induction to show that the quantities appearing in the forward and backward passes at every layer are the same for both models

at every time step in the infinite-width limit. We use the Tensor Program framework for this proof so we need smoothness of σ ($p \geq 2$) for this result.

In essence, [Theorem 4.2](#) shows that the IP-LLR parameterization is equivalent to μP where we simply forget the initialization after the first forward and backward passes. Said differently, IP-LLR is the same as μP , except that IP-LLR re-initializes the weights of the intermediate layers $l \in [2, L]$ at $t = 1$ with $W^l(1) = \Delta W^l(1)$, *i.e.*, with the first update computed after the first forward-backward pass. It is not entirely clear whether forgetting the initial weights in one step is beneficial or detrimental to learning. On the one hand, it would seem like forgetting the random initialization could make the network learn faster and be more robust to perturbations (but this is only speculative at this point, and we leave this open for future work), on the other hand the large rank of the initial weight matrices with i.i.d. Gaussian entries might increase the stability of the training dynamics. In other words, while the randomness from initialization propagates to every layer at every time step for μP , it is forgotten in one step of SGD for IP-LLR in the infinite-width limit. We explore the comparative performance of μP and IP-LLR in [Section 6](#) but there appears to be no clear-cut indication towards one model or the other.

Another interesting difference between IP-LLR and μP is that for any intermediate layer $l \in [2, L]$, while $(W_{jq}^l(t) - W_{jq}^l(0))/W_{jq}^l(0) = \Theta(m^{-1/2})$ for μP , so that the effective weights only move infinitesimally (in the infinite-width limit) relatively to their initial values, we have $(W_{jq}^l(t) - W_{jq}^l(0))/W_{jq}^l(0) = \Theta(1)$ for IP-LLR so that the effective weights actually move in the infinite-width limit (see more details in [Remark F.6](#)).

5 Alternative Methods for Escaping the Initial Stationary Point

As discussed in [Section 4](#), using large initial learning rates in combination with a positively homogeneous activation function allows escaping the initial stationary point and induces stable learning. In this section, we introduce two alternatives to escape this initial stationary point and discuss the properties of the resulting models. In contrast to the setting of [Section 4](#), in all this section, we consider IPs with bias terms at every layer.

A first alternative to escape the initial stationary point, which we discuss in [Section 5.1](#), is to simply initialize the weight matrices with i.i.d. Gaussian distributions which are not-centered around 0, as suggested by [Nguyen and Pham \(2020\)](#). This method is able to escape the stationary point without large initial learning rates and without any homogeneity assumption on the activation function. It turns out that the computations in that setting are well described within the Tensor Program framework and we show that, as highlighted in ([Nguyen and Pham, 2020](#), [Corollary 37](#)), a collapse phenomenon occurs, where all the individual entries in the weight matrix of an intermediate layer evolve by the same deterministic quantity in the infinite-width limit.

Another alternative is to remove the pre-factor m^{-1} in front of the bias terms of layers $l \geq 2$. Indeed, as observed in [Section 3.1](#), the vanishing of the forward pass and the weight updates in integrable parameterizations is mostly due to the multiplications by the weight matrices $m^{-1}U^l$ which results in pre-activations whose coordinates are $\Theta(m^{-1/2})$ for $l \in [2, L]$. Since the bias terms are decoupled from the input to the layer, re-scaling them appropriately avoids vanishing of the forward pass for IPs. Escaping the initial stationary point can then be achieved without any homogeneity assumption on the activation function σ . However, one issue which arises then is that the bias terms have the dominant contribution to the pre-activations, and since the input signal propagates through the network via the weight multiplications, the output of the trained network is only “weakly” dependent on its

input and the training data. Let us now study in more details these two alternatives.

5.1 Using Non-Centered i.i.d. Initialization

In this section, we consider the following modified version of IPs which we call *IP-non-centered* : the forward pass is computed exactly as in IPs but the weight matrices of layers $l \geq 2$ are initialized with $w_{jq}^l(0) = U_{jq}^l + u_l \sim \mathcal{N}(u_l, 1)$ i.i.d. over (j, q) with $u_l \neq 0$. This simply consists in setting $w^l(0) = U^l + u_l J$ for $l \in [2, L]$ and $w^{L+1}(0) = U^{L+1} + u_{L+1} \mathbf{1}$ where J is the square matrix full of ones (whose variable size is the same as U^2 and thus equal to m) and $\mathbf{1}$ is the vector (of variable size equal to m) full of ones. As we will see shortly, the effect of this type of initialization is similar to removing the pre-factor in m^{-1} on the bias terms in that the vanishing of the matrix multiplications $m^{-1}U^l x_t^{l-1}$ is offset by the appearance of an additional term in the expression of h_t^l whose coordinates are all equal and *depend* on the input data.

5.1.1 First forward Pass

As for any IP, h_0^1 is a Gaussian vector with i.i.d. coordinates following $\mathcal{N}(0, \|\xi\|^2 + 1)$ at any width, and for the second layer we have

$$h_0^2 = m^{-1}(U^2 x_0^1 + v^2) + u_2 m^{-1} J x_0^1.$$

The coordinates of $m^{-1}Jx_0^1$ are all equal to $(1/m) \sum_{q=1}^m \sigma(h_{0,q}^1)$, which converges almost surely, by the law of large numbers, towards $\mathbb{E}[\sigma(Z_1)]$ where $Z_1 \sim \mathcal{N}(0, \|\xi\|^2 + 1)$. When $\sigma = \text{ReLU}$ this expectation is tractable as shown in Appendix M and equal to $\sqrt{(\|\xi\|^2 + 1)/(2\pi)}$. On the other hand, the coordinates of $m^{-1}(U^2 x_0^1 + v^2)$ simply converge to 0. The term $u_2 m^{-1} J x_0^1$ thus offsets the vanishing of the term $m^{-1}(U^2 x_0^1 + v^2)$. In the infinite-width limit, we thus have that $x_{0,j}^2 \simeq \sigma(u_2 \mathbb{E}[\sigma(Z_1)])$ for any $j \in [1, m]$. We thus already see that the coordinates of h_0^2 all converge almost surely to the same deterministic constant $C_2 = u_2 \mathbb{E}[\sigma(Z_1)]$ and the coordinates of x_0^2 towards $\sigma(C_2)$.

Degeneracy in intermediate layers. An easy induction gives that for any $l \in [2, L]$, for any coordinate j , and for large m

$$\begin{aligned} h_{0,j}^l &\simeq u_l \sigma(u_{l-1} \sigma(\dots \sigma(u_2 \mathbb{E}[\sigma(Z_1)]))) =: C_l, \\ x_{0,j}^l &\simeq \sigma(u_l \sigma(u_{l-1} \sigma(\dots \sigma(u_2 \mathbb{E}[\sigma(Z_1)])))) = \sigma(C_l), \end{aligned} \tag{5.1}$$

so that the coordinates of the (pre-)activations of any intermediate layer are all equal to the same deterministic constant for large m . Finally, the output of the first forward pass is $f_0(\xi) = m^{-1}((U^{L+1})^\top x_0^L + v^{L+1}) + u_{L+1} (1/m) \sum_{q=1}^m x_{0,q}^L$ and converges almost surely towards the constant $u_{L+1} \sigma(u_L \sigma(\dots \sigma(u_2 \mathbb{E}[\sigma(Z_1)])))$ (this is made rigorous within the framework of the Tensor Program).

If $\sigma = \text{ReLU}$ we see that to avoid vanishing of the first forward pass, one must set $u_l > 0$ for $l \in [2, L]$, and we then get that the coordinates of h_0^l are roughly all equal to $u_l u_{l-1} \dots u_2 \sqrt{(\|\xi\|^2 + 1)/(2\pi)}$. This suggests that to avoid vanishing or explosion with the depth L , one should set $u_l = 1$ for $l \in [2, L]$.

5.1.2 First Backward Pass

We show here that the same degeneracy as in the first forward pass is also at play in the first backward pass. We have $\nabla_{x^L} f_0(\xi_0) = W^{L+1}(0) = m^{-1}(U^{L+1} + u_{L+1} \mathbf{1})$, so that the coordinates of $m \nabla_{x^L} f_0(\xi_0)$ are not deterministic in the infinite-width limit and simply follow

$\mathcal{N}(u_{L+1}, 1)$ i.i.d. We have $m\nabla_{h^L} f_0(\xi_0) = m\nabla_{x^L} f_0(\xi_0) \odot \sigma'(h_0^L)$ and as shown in Section 5.1.1 the coordinates of h_0^L are roughly all equal to the same constant for large m , so that the coordinates of $m\nabla_{h^L} f_0(\xi_0)$ are in $\Theta(1)$.

Degeneracy for layers $l \in [1, L-1]$. Using the equations of backpropagation, we have

$$m\nabla_{x^{L-1}} f_0(\xi_0) = m^{-1}(U^L)^\top (m\nabla_{h^L} f_0(\xi_0)) + u_L m^{-1} J(m\nabla_{h^L} f_0(\xi_0)).$$

The multiplication by $m^{-1}(U^L)^\top$ yield a vector whose coordinates converge to 0, and the coordinates of $m^{-1}J(m\nabla_{h^L} f_0(\xi_0))$ are all equal to $(1/m) \sum_{q=1}^m d\tilde{h}_{0,q}^L$ where $d\tilde{h}_0^L := m\nabla_{h^L} f_0(\xi_0) = m\nabla_{x^L} f_0(\xi_0) \odot \sigma'(h_0^L)$. We thus have that $(1/m) \sum_{q=1}^m d\tilde{h}_{0,q}^L$ converges almost surely to the constant $u_{L+1}\sigma'(C_L)$, where C_L is defined in Equation (5.1). Because $m\nabla_{h^{L-1}} f_0(\xi_0) = m\nabla_{x^{L-1}} f_0(\xi_0) \odot \sigma'(h_0^{L-1})$, we get that the coordinates of $m\nabla_{h^{L-1}} f_0(\xi_0)$ are roughly all equal to the constant $u_{L+1}u_L\sigma'(C_L)\sigma'(C_{L-1})$ for large m . An easy induction then yields that for any $l \in [1, L-1]$, and for any coordinate j

$$\begin{aligned} d\tilde{x}_{0,j}^l &\simeq u_{L+1} \dots u_{l+1} \sigma'(C_L) \dots \sigma'(C_{l+1}), \\ d\tilde{h}_{0,j}^l &\simeq u_{L+1} \dots u_{l+1} \sigma'(C_L) \dots \sigma'(C_l), \end{aligned}$$

as $m \rightarrow \infty$, where $d\tilde{x}_0^l := m\nabla_{x^l} f_0(\xi_0)$, $d\tilde{h}_0^l := m\nabla_{h^l} f_0(\xi_0)$, and C_l is defined in Equation (5.1). Note that all the C_l depend on ξ_0 through $Z_1 \sim \mathcal{N}(0, \|\xi_0\|^2 + 1)$, and the coordinates of $m\nabla_{h^l} f_0(\xi_0)$ are in $\Theta(1)$ for all $l \in [1, L]$. Again, the products of u_l which appear in the backward pass strongly suggest setting $u_l = 1$ for any $l \in [2, L]$ to avoid issues with increasing depth L .

5.1.3 First parameter updates

Now that we have described the first forward and backward passes, we can give the formulas for the first weight updates of IP-non-centered. We have:

$$\begin{aligned} \Delta W^1(1) &= -\eta m^{-(1+c_1)} \partial_2 \ell(y_0, f_0(\xi_0)) (m\nabla_{h^1} f_0(\xi_0)) \xi_0^\top \\ \Delta B^1(1) &= -\eta m^{-(1+c_1)} \partial_2 \ell(y_0, f_0(\xi_0)) (m\nabla_{h^1} f_0(\xi_0)) \\ \Delta W^l(1) &= -\eta m^{-(2+c_l)} \partial_2 \ell(y_0, f_0(\xi_0)) \frac{(m\nabla_{h^l} f_0(\xi_0)) (x_0^{l-1})^\top}{m}, & l \in [2, L] \\ \Delta B^l(1) &= -\eta m^{-(3+c_l)} \partial_2 \ell(y_0, f_0(\xi_0)) (m\nabla_{h^l} f_0(\xi_0)) & l \in [2, L] \\ \Delta W^{L+1}(1) &= -\eta m^{-(1+c_{L+1})} \partial_2 \ell(y_0, f_0(\xi_0)) \frac{x_0^L}{m}, \\ \Delta B^{L+1}(1) &= -\eta m^{-(2+c_{L+1})} \partial_2 \ell(y_0, f_0(\xi_0)). \end{aligned}$$

Choice of learning rates and update contributions. To ensure non-vanishing and non-exploding updates for both the weights and the bias terms, one must choose different learning rate exponents c_l for the weights and for the bias terms for layers $l \geq 2$. To make things simpler, we simply choose $c_l = -2$ for $l \in [2, L]$ and $c_1 = c_{L+1} = -1$ (which are the learning rates of Naive-IP) for both the weights and the bias terms, which implies that the updates of the bias terms contribute vanishingly to the second forward pass as $m \rightarrow \infty$ for layers $l \in [2, L+1]$, but this is offset by the non-centered initialization.

Degeneracy of the weight updates. With the choice of learning rate exponents of the Naive-IP, all the entries of $\Delta w^l(1)$ are equal to the same deterministic constant for large m for $l \in [3, L - 1]$. In other words, for those layers $l \in [3, L - 1]$, there is a *collapse to a single parameter per layer* (since the contribution of the centered initialization vanishes for large m) which evolves by a deterministic quantity. We recover a result proved by [Nguyen and Pham \(2020\)](#) (see [Nguyen and Pham, 2020](#), Corollary 37). In Section 3.1 and Proposition 3.1, we have additionally shown that this translation is 0 when the i.i.d. initialization is centered around 0. In fact, a slightly more precise statement can be made: although the coordinates of $\Delta w^2(1)$ do not become equal to deterministic constants for large m , the coordinates of $\Delta W^2(1)x_1^1$ all become equal to the same deterministic constant in the large-width limit because the term $(x_0^1)^\top x_1^1/m$ converges to a finite expectation.

5.1.4 Collapse to Deterministic Dynamics

Repeating the same calculations as in Sections 5.1.1 and 5.1.2 shows that the choice of learning rate exponents as in Naive-IP (see Section 2.2), *i.e.*, $c_1 = c_{L+1} = -1$, and $c_l = -2$ for $l \in [2, L]$ leads to non-vanishing and non-exploding updates for the weights at any time step for IP-non-centered, and deterministic dynamics as summarized in the following informal theorem:

Theorem 5.1 (Informal). *Consider IP-non-centered with the Naive-IP learning rates at every time step, and let $t \geq 0$ and $\xi \in \mathbb{R}^d$ be an input to the network. Then, one has that:*

- (i) *for any $l \in [2, L - 1]$, the coordinates of h_t^l (resp. x_t^l) all converge to the same deterministic constant,*
- (ii) *for any $l \in [2, L - 1]$, the coordinates of $m \nabla_{x_t^l} f_t(\xi_t)$ (resp. $m \nabla_{h_t^l} f_t(\xi_t)$) all converge to the same deterministic constant,*
- (iii) *for any $l \in [3, L - 1]$, the entries of $(W^l(t) - W^l(0))$ all converge to the same deterministic constant.*

The rigorous version of this theorem, and its proof, formalized within the framework of the Tensor Program, are presented in Appendix K.2.

5.2 Not Scaling the Bias Terms

In this section, we consider a version of IPs where we remove the pre-factor $1/m$ for the bias terms of layers $l \geq 2$. We thus consider the following computations in the forward pass:

$$\begin{aligned} h_t^1 &= w^1(t)\xi + b^1(t), \\ h_t^l &= \left(m^{-1} w^l(t) x_t^{l-1} \right) + b^l(t), \quad l \in [2, L] \\ f_t(\xi) &= \left(m^{-1} (w^{L+1}(t))^\top x_t^L \right) + b^{L+1}(t), \end{aligned} \tag{5.2}$$

which in other terms simply means that $B^l(t) = b^l(t)$ for $l \in [1, L + 1]$. We use the same initialization for the bias terms as in IPs: $b^l(0) = v^l$ for $l \in [1, L + 1]$, where the entries of v^l are i.i.d. following $\mathcal{N}(0, 1)$. We call *IP-bias* the modified version of the integrable parameterization described by Equations (5.2).

Gaussian first forward pass. For the first forward pass we have that the pre-activation of the first layer h_0^1 is the same as in IPs at initialization and thus has i.i.d. Gaussian coordinates. On the other hand, $h_0^l \simeq b^l(0) = v^l \sim \mathcal{N}(0, 1)$ as $m \rightarrow \infty$, so that the coordinates of the pre-activations of all the intermediate layers now behave as standard Gaussians in the large-width limit. Note that in contrast to IP-non-centered, the coordinates of h_0^l **do not** depend on the input data for $l \geq 2$ in the large-width limit. Similarly, we have $f_0(\xi) \simeq v^{L+1} \sim \mathcal{N}(0, 1)$ (which does not depend on the input ξ) as $m \rightarrow \infty$.

First parameter updates. The first backward pass still vanishes as in integrable parameterizations because of the multiplications by $(W^l(0))^\top = m^{-1/2}(m^{-1/2}U^l)$. Indeed, we have $\nabla_{x^L} f_0(\xi) = W^{L+1}(0) = m^{-1}U^{L+1}$, and $\nabla_{h^L} f_0(\xi) = m^{-1}U^{L+1} \odot \sigma'(h_0^L)$, so that the coordinates of $\nabla_{x^L} f_0(\xi)$ and $\nabla_{h^L} f_0(\xi)$ are in $\Theta(m^{-1})$. For $l \in [1, L-1]$, we have that $m\nabla_{x^l} f_0(\xi) = m^{-1/2}(m^{-1/2}(U^{l+1})^\top)(m\nabla_{h^{l+1}} f_0(\xi))$, and an easy induction shows that the coordinates of $\nabla_{x^l} f_0(\xi)$ and $\nabla_{h^l} f_0(\xi)$ are in $\Theta(m^{-1}m^{-(L-l)/2})$ for any $l \in [1, L]$. Note that as in the forward pass, the backward pass at $t = 0$ also does not depend on the first training input ξ_0 except for $\nabla_{h^1} f_0(\xi_0)$. We get the following formulas for the first weight and bias updates at $t = 0$:

$$\begin{aligned}
\Delta W^1(1) &= -\eta m^{-c_1} \partial_2 \ell(y_0, f_0(\xi_0)) (\nabla_{h^1} f_0(\xi_0)) \xi_0^\top, \\
\Delta B^1(1) &= -\eta m^{-c_1} \partial_2 \ell(y_0, f_0(\xi_0)) \nabla_{h^1} f_0(\xi_0), \\
\Delta W^l(1) &= -\eta m^{-(2+c_l)} \partial_2 \ell(y_0, f_0(\xi_0)) \frac{(m \nabla_{h^l} f_0(\xi_0)) x_0^{l-1}}{m}, \quad l \in [2, L], \\
\Delta B^l(1) &= -\eta m^{-c_l} \partial_2 \ell(y_0, f_0(\xi_0)) \nabla_{h^l} f_0(\xi_0), \quad l \in [2, L], \\
\Delta W^{L+1}(1) &= -\eta m^{-(1+c_{L+1})} \partial_2 \ell(y_0, f_0(\xi_0)) x_0^L / m, \\
\Delta B^{L+1}(1) &= -\eta m^{-c_{L+1}} \partial_2 \ell(y_0, f_0(\xi_0)).
\end{aligned} \tag{5.3}$$

Initial learning rates. Because the backward pass vanishes in the infinite-width limit, the learning rate exponents c_l still need to be chosen carefully in order to escape the initial stationary point. However, the two following points stand out: (1) because the first forward pass does not vanish as in the Naive-IP, the choice of c_l does not require any homogeneity property, and needs not be as large (in absolute value) as for IP-LLR (see the values in the case $p = 1$ in the comment after [Theorem 3.2](#)); (2) Because we removed the pre-factor m^{-1} from the bias terms, $\Delta W^l(1)$ and $\Delta B^l(1)$ do not have compatible magnitudes, which suggests setting a separate learning rate exponent ϵ_l for the bias terms, different from c_l for layers $l \in [2, L+1]$, in order to have non-trivial updates for both the weights and the bias terms. In light of the previous comment and of the update formulas of [Equations \(5.3\)](#), we set, at $t = 0$, the learning rate exponents for the weights to

$$\begin{aligned}
c_1 &= -1 - (L-1)/2 = -(L+1)/2, \\
c_l &= -2 - (L-l)/2 = -(L-l+4)/2, \\
c_{L+1} &= -1,
\end{aligned} \tag{5.4}$$

and for the bias terms to

$$\begin{aligned}
\epsilon_1 &= c_1 = -(L+1)/2, \\
\epsilon_l &= -(L-l+2)/2, \\
\epsilon_{L+1} &= 0.
\end{aligned} \tag{5.5}$$

One may compare the learning rates exponents for the weights with those of IP-LLR with a degree of homogeneity $p = 1$, which are $c_1 = c_{L+1} = -(L+1)/2$, and $c_l = -(L+2)/2$, where the absolute value of the exponent does not decrease with the layer l for intermediate

layers. Even when the learning rates are appropriately scaled as in Equations (5.4) and (5.5), $\Delta W^l(1)$ does not depend on the first training input for $l \in [2, L+1]$. We thus get the following informal theorem, whose formal version within the framework of the Tensor Program is given in Appendix K.1.

Theorem 5.2 (Informal). *Consider the IP-bias as in Equations (5.2), with the initial learning rates as in Equations (5.4) and (5.5). Then, for any input $\xi \in \mathbb{R}^d$ to the network, $h_0^l(\xi), x_0^l(\xi)$ for $l \geq 2$, and $f_0(\xi)$ do not depend on ξ in the limit $m \rightarrow \infty$. In addition, $\Delta W^l(1)$ also does not depend on the first training input ξ_0 in the infinite-width limit $m \rightarrow \infty$ for $l \in [3, L+1]$.*

Learning rates at step $t \geq 1$. Repeating the calculations of the forward pass with the updates of Equation (5.3) and with the learning rates for the weights and bias terms as described in Equations (5.4) and (5.5), we readily get that the coordinates of the second forward pass are in $\Theta(1)$. Then, it is direct to see that the choice of the Naive-IP learning rate exponents $c_1 = c_{L+1} = -1$, and $c_l = -2$ for $l \in [2, L]$ for the weights and $\epsilon_l = -1$ for $l \in [1, L]$ and $\epsilon_{L+1} = 0$ for the bias terms yields non-vanishing and non-exploding updates for the weights and the bias terms at $t = 1$.

Degeneracy at time $t \geq 1$. It follows that the same choice of learning rate exponents as at $t = 1$ also induce non-vanishing and non-exploding updates in the limit $m \rightarrow \infty$ at later time steps $t \geq 1$. With this choice of learning rates we thus get, for any $t \geq 1$, and for $l \in [2, L]$,

$$h_t^l \simeq v^l + \sum_{s=1}^{t-1} \Delta W^l(s) x_t^{l-1}.$$

With the choice of learning rates prescribed above for $t \geq 1$, the products $\Delta W^l(s) x_t^{l-1}$ are finite and their numerical value strongly depends on the values of η and $\partial_2 \ell(y_s, f_s(\xi_s))$. Typically, their product is rather small (*e.g.*, $\leq 10^{-2}$), and this means that the initial bias term v^l has the dominant contribution to h_t^l . Therefore, in addition to Theorem 5.2, it can be also be argued that the forward pass in the intermediate layers only weakly depends on the training data and on the input to the network at time steps t .

6 Numerical Experiments

In this section we investigate numerically the behavior of the models previously introduced in this work, namely Naive-IP, IP-LLR, IP-bias, IP-non-centered and μ P. In contrast to the theoretical analysis carried out in Sections 3, 4, and 5, we examine the performance of the models on a multi-class classification task (instead of a single output prediction) and we train them using mini-batch SGD (instead of single-sample SGD). In addition to these two points, we adopt the following slight modifications compared to our theoretical setting.

Standard deviation of initial weights. In our numerical experiments, we allow the initial Gaussian weight matrices U^l and vectors v^l to have entries drawn from $\mathcal{N}(0, \delta_l^2)$ where δ_l can be different from 1 for $l \in [1, L]$, but is independent of m . As hinted in Remark 4.2 and explained more in detail in Remarks F.2 and F.3, this is to avoid issues (vanishing or explosion of the forward/backward pass) with the depth L . The choices of the standard deviation of the Gaussian depend on the activation function and are summarized in Table 1.

activation	ReLU	GeLU	ELU	tanh
init. std	$\sqrt{2}$	2	1	1

Table 1: Standard deviation δ_l of the initial Gaussian entries of layers $l \in [1, L]$ for different choices of activation functions.

Re-scaling the standard deviation of the first layer. All the models we consider have $a_1 = 0$ so that, as mentioned in Section 5.2, the coordinates of h_0^1 follow $\mathcal{N}(0, \|\xi\|^2 + 1)$ and the variance is equal to $\sum_{k=1}^d \xi_k^2 + 1$. To avoid having too large a variance when the (fixed) dimension d is large, we re-scale the standard deviation of the first layer’s weights and bias term at initialization by dividing it by $\sqrt{d+1}$, that is we use the Gaussian law $\mathcal{N}(0, \delta_1^2/(d+1))$ to initialize the entries of $w^1(0)$ and $b^1(0)$.

Calibrating the initial base learning rates for IP-LLR. As discussed in Section 4.2, IP-LLR basically amounts to training with μP but forgetting the initialization in the intermediate layers for the first update. We thus roughly have $W^l(1) \simeq \Delta W^l(1)$ for any $l \in [2, L]$, and the base learning rate η directly influences the magnitude of $\Delta W^l(1)$ and thus that of h_1^l . Typical values for the learning rates, the initial loss derivative $\partial_2 \ell(y_0, 0)$, and the averaged inner products involved in the second forward pass are rather small (*e.g.*, $\leq 10^{-1}$), and this will cause the pre-activations of the second forward pass to be of small magnitude, and this effect compounds quickly with depth as the pre-activations of layer $(l-1)$ are then multiplied by $\Delta W^l(1)$. This will in turn lead to very small values for the second weight updates $\Delta W^l(2)$ and can considerably slow down learning in practice. To overcome this issue, we simply calibrate the initial values of the base learning rates η_l (but cap them at a value of 500 to avoid too large initial updates) of layers $l \in [2, L]$ at $t = 0$, so that the magnitude of the pre-activation of the intermediate layers in the second forward pass is equal to 1 on average over the second training batch.

Note that this calibration results in base learning rates η_l which **do not** depend on m (they do depend on L however) in the large-width limit as the coordinates of h_1^l have non-zero and finite values for large m . In contrast, this is not possible with the Naive-IP as the coordinates of h_1^l converge to zero as fast as some power of m , which would result in the base learning rate η_l depending on m which is prohibited (by definition of the base learning rate).

All the points above can be handled within the framework of the Tensor Program, but they would unnecessarily over-complicate the analysis and the formulas, which is why we used a simpler setting in our theoretical analysis.

6.1 Experimental Setup

We evaluate the performance of the different models on two datasets: MNIST², containing 60,000 training samples and 10,000 test samples, and CIFAR-10³, containing 50,000 training samples and 10,000 test samples. Both datasets consist in a 10-class image classification task. Since we consider only fully-connected networks, we use gray-scale images which we also flatten for both datasets, which means the input dimension is $d = 28 \times 28 = 784$ for MNIST and $d = 32 \times 32 = 1024$ for CIFAR-10.

We train for 600 SGD steps on MNIST and 1200 steps on CIFAR-10 using a base learning rate $\eta = 0.01$, a batch-size $B = 512$, and the cross-entropy loss, which satisfies Assumption 1.

²<http://yann.lecun.com/exdb/mnist/>

³<http://www.cs.toronto.edu/~kriz/cifar.html>




Figure 1: Loss *vs.* number of optimization (SGD) steps on MNIST for different activation functions.

For each experiment, we run $N_{\text{trials}} = 5$ trials with different random initializations. The hyperparameters are summarized in Table 2.

L	m	d_{MNIST}	d_{CIFAR}	ℓ	η	B	N_{trials}
6	1024	784	1024	cross-ent.	0.01	512	5

Table 2: Hyperparameters for training models.

6.2 Naive-IP is Trivial but Large Initial Learning Rates Induce Learning

In this section we compare the numerical performance of Naive-IP and IP-LLR on MNIST for different activation functions. Essentially, the results we present corroborate Proposition 3.1 and Theorem 4.1, except that numerical evidence tends to show that those results hold with less restrictive assumptions on the activation function than what we consider in the theoretical part, as already hinted in Point 2 of Remark 4.2.

As observed in Figure 1, while the loss (averaged over a batch) stays at its initial value for Naive-IP, we observe a decrease for IP-LLR whose strength depends on the choice of activation function. Similarly, Figure 2 depicts the evolution of the mean absolute output during training, that is, we plot for any step t the quantity $(1/B) \sum_{i=1}^B (1/10) \sum_{k=1}^{10} |f_{k,t}(\xi_t^{(i)})|$, where $\xi_t^{(i)}$ is the i -th sample in the batch at time t and for any class label $k \in [1, 10]$, $f_{k,t}(\xi)$ is the k -th entry of the output of the model (logits for class k) on input ξ . We also observe here that there is no change in the output for the Naive-IP which stays equal to 0 during the course of training, whereas for IP-LLR, the mean absolute output value increases from its initial value, equal to 0, to some positive quantity whose value depends on the activation function. The solid line in both plots denotes the mean of the metric of interest over multiple (5) random trials while the shaded area represents a 95% confidence interval around the mean. There is no shaded area for Naive-IP since the output of the network is equal to the deterministic constant 0 at any time step for large m , as stated in Proposition 3.1.

Finally, we show in Table 3 the test accuracy (averaged over 5 random runs) *at the end of training* for the Naive-IP and IP-LLR for different activation functions. The Naive-IP has the same test accuracy of 0.098 independently of the activation function, which is roughly equal to that of random guessing which would yield an accuracy of 0.10 as there are 10 classes. In contrast, IP-LLR has higher-than-chance test accuracy for every choice of activation function, and while ReLU appears to perform poorly, all other activations perform relatively well with ELU and GeLU achieving an error lower than 5%.




Figure 2: Mean absolute output *vs.* number of optimization (SGD) steps on MNIST for different activation functions.

model \ activation	ReLU	GeLU	ELU	tanh
Naive-IP	0.098	0.098	0.098	0.098
IP-LLR	0.113	0.956	0.964	0.932

Table 3: Test accuracies on MNIST for various activation functions.

6.3 IP-LLR *vs.* μP

We compare the numerical performance of IP-LLR and μP on both MNIST and CIFAR-10, and investigate the reasons behind the differences observed between different models and different non-linearities.

As observed in Tables 4 and 5, the performance, as measured by the accuracy on the test set, is consistent across activation functions for μP whereas the gaps are larger for IP-LLR. However, the best test accuracy for μP and IP-LLR are comparable: the former achieves 0.975 test accuracy on MNIST and 0.419 test accuracy on CIFAR-10 with $\sigma = \text{GeLU}$ while the latter achieves 0.964 test accuracy on MNIST and 0.383 test accuracy on CIFAR-10 with $\sigma = \text{ELU}$.

Performance and rank collapse. The consistency of μP across activation functions and the lack of consistency for IP-LLR can be explained by (or at least correlated with) the diversity, measured in terms of rank, of the (pre-)activations at different layers on large batches of samples. Indeed, as shown in (Daneshmand et al., 2020), the rank of the family of pre-activations (considered over large batches) has a dramatic impact on the observed performance of models. In fact, the authors argue that this might be the reason behind the empirical success of batch normalization: it allows the rank of these families of pre-activations to remain large even when the number of hidden layers L is large, whereas they show there is a collapse in the rank without the batch-normalization operation, which coincides with poor accuracy. This problem is exacerbated in IP-LLR because the contribution of the initial weight matrices (which are full-rank) vanishes after the first gradient step, thereby lowering considerably the rank of the family of pre-activations. Two effects are then at play: (1) the choice of the activation function σ can induce large differences in the rank of the family of vectors $(W^l(1)\sigma(h))_{h \in \mathcal{S}}$, where \mathcal{S} is a large set of vectors; (2) the impact of the activation function on (1) is compounding with depth and can lead to dramatically small rank (equal to 1 in the worst case) towards the last layers of the network.

In Figures 3 we plot the rank (the y -axis is in log-scale) of the families $(h_1^l(\xi))_{\xi \in \mathcal{S}}$

model \ activation	ReLU	GeLU	ELU	tanh
IP-LLR	0.113	0.956	0.964	0.932
μ P	0.954	0.975	0.928	0.905

Table 4: Test accuracies on MNIST for various activation functions.

model \ activation	ReLU	GeLU	ELU	tanh
IP-LLR	0.100	0.329	0.383	0.284
μ P	0.407	0.419	0.356	0.304

Table 5: Test accuracies on CIFAR-10 for various activation functions.

and $(x_1^l(\xi))_{\xi \in \mathcal{S}}$ for $l \in [1, L]$, where \mathcal{S} is the set comprised of the first 5,000 training inputs of MNIST. The numerical “rank” is computed as in (Daneshmand et al., 2020) with `torch.matrix_rank()` which regards singular values below $\sigma_{\max} \times m \times 10^{-7}$ as zero. We observe that for IP-LLR, the rank of those families with $\sigma = \text{ReLU}$ is one order of magnitude smaller than for other activation functions after layer $l = 4$ and even collapses to 1 in the last layers, which might explain its poor performance, whereas for μ P all activation functions induce comparable ranks which remain at least on the order of 10^2 at any layer. We believe the latter fact is due to the non-vanishing contributions of the initial Gaussian matrices which are full-rank (with probability 1). In contrast, it would seem like IP-LLR is much more sensitive to the choice of activation function and we identify the vanishing of the contribution of the initial weights for intermediate layers as a probable cause for this effect.

Whether the difference between ReLU and other activation functions for IP-LLR is actually due to the difference between the homogeneity property with $\sigma = \text{ReLU}$ and the effective linearization property for other activation functions (as highlighted in Remark 4.2) or to other inherent characteristics of the activation functions is still an open question and we leave it for future work.

6.4 Learning is Degenerate for IP-bias and IP-non-centered

In this section we show numerically that IP-non-centered and IP-bias (see Sections 5.1 and 5.2 respectively) are able to escape the initial stationary point but that the resulting dynamics do not seem effective as observed through the evolution of the training loss.

Figure 4 shows that both models are indeed able to escape the initial stationary point as the magnitude of the output evolves non-trivially during training but in contrast Figure 5, depicting the training losses on MNIST and CIFAR-10 for both models, shows that learning is very slow for those models and that the dynamics are not effective in reducing the training loss.

Additionally, as summarized in Table 6, the slow decrease of the training loss translates into poor test accuracy at the end of training comparatively with IP-LLR and μ P, even with the best choice of activation function.

7 Conclusion

Recent research has shown that the parameterization of a neural network has a dramatic impact on its training dynamics, and therefore, on the type of functions that it is able to learn. Until now, the parameterizations used by practitioners have been restricted to




Figure 3: Ranks (log-scale) of the families of pre-activations ($h_1^l(\xi)$) and activations ($x_1^l(\xi)$) at time $t = 1$ on MNIST *vs.* layer l for different activation functions.

dataset \ model	IP-LLR	μP	IP-bias	IP-non-centered
MNIST	0.964	0.975	0.113	0.209
CIFAR-10	0.383	0.419	0.100	0.154

Table 6: Test accuracies (averaged over 5 random runs) at the end of training on MNIST and CIFAR-10. For each model, we show the maximum (averaged) accuracy over all activation functions. For each model, the activation function which performs best is the same for both datasets and the pairing model \rightarrow activation is the following: IP-LLR \rightarrow ELU, $\mu P \rightarrow$ GeLU, IP-bias \rightarrow GeLU, IP-non-centered \rightarrow ELU.

standard schemes which rely on the analysis of the the first forward and backward passes. In the present work, pushing the analysis beyond the first gradient step (which is made possible by the Tensor Program framework), we have studied how to train neural networks with parameterizations that enjoy radically different behaviors, such as forgetting the contribution of the initial weights after the first weight update.

The parameterizations we have analyzed, which we refer to as *integrable parameterizations*, have been previously described with tools from the *mean-field* literature, and we have deepened our understanding of these models with a different perspective. Indeed, we have shown that these parameterizations are trivial for deep networks with centered i.i.d. initialization and a constant learning rate: they are stuck at initialization. This observation led us to explore various ways to escape this initial stationary point and initiate learning. Among those methods, we found that the only one that does not lead to a degenerate behaviour is to use large learning rates for the first gradient step. We proved that in the infinite-width limit the resulting dynamic is equivalent to a modification of μP where the initial weights are removed after the first gradient step. Importantly, the random fluctuations around the limit—which are ignored in the mean-field description—turn out to actually be essential for our analysis, since it is by amplifying them that we are able to escape the stationary point.

Extending our theoretical results to a more general class of activation functions requires more thorough technical work and is left as an open problem. Also, analyzing rigorously the impact of the presence or absence of the initial weight matrices on the learning behavior appears to be an interesting avenue for future research. Finally, understanding the generalization properties of IP-LLR and μP remains an important open question but is beyond the scope of this paper.




Figure 4: Mean absolute value of the output during training.




Figure 5: Loss *vs.* number of SGD steps.

Acknowledgements

Karl Hajjar and Christophe Giraud receive respectively full and partial support from the Agence Nationale de la Recherche (ANR), reference ANR-19-CHIA-0021-01 “BiSCottE”.

Appendix

A Notations

We introduce here some additional notations that will come in handy in the text and equations presented in the Appendix.

Hat matrices. We define the following matrices and output weight vector (see Definition 2.1 for the definitions of the matrices U^l):

$$\begin{cases} \widehat{W}^1 = U^1 \\ \widehat{W}^l = m^{-1/2}U^l, \quad l \in [2, L+1]. \end{cases} \quad (\text{A.1})$$

The pre-factor in $m^{-1/2}$ is the natural re-scaling of the i.i.d. Gaussian matrices when their input dimension grows to infinity due to the central limit theorem (CLT).

Omegas. For any ac-parameterization, we define $\omega_1 := m^{-a_1}$, and for any $l \in [2, L+1]$, $\omega_l := m^{1/2-a_l}$. To avoid blow-up or vanishing in the first layer, all the parameterizations we study have $\omega_1 = 1$. This is the case for integrable parameterizations, the NTK parameterization and for μP . For integrable parameterizations we also have $\omega_l = m^{-1/2}$ for $l \in [2, L+1]$, but for μP , $\omega_l = 1$ if $l \in [2, L]$ and $\omega_{L+1} = m^{-1/2}$ (see Section B.3 for a detailed description of μP).

Those ω_l naturally appear in the calculations as the magnitudes of the first forward pass of an ac-parameterization of a neural network. The term m^{-a_l} comes from the scaling pre-factor of the effective weights, and the added $m^{1/2}$ appears when expressing the computation in function of the naturally scaled \widehat{W}^l : $W^l(0) = \omega_l \widehat{W}^l$.

Scalar limits. For any scalar ω which depends on m , we denote by $\overset{\circ}{\omega}$ the almost sure limit (when it exists) of this scalar as $m \rightarrow \infty$.

Gradients. We define for any t and l ,

$$\begin{cases} dh_t^l := \nabla_{h_t^l} f_t(\xi_t) \\ dx_t^l := \nabla_{x_t^l} f_t(\xi_t) \\ dw^l(t) := \nabla_{w^l(t)} f_t(\xi_t) \\ db^l(t) := \nabla_{b^l(t)} f_t(\xi_t) \\ \chi_t := \partial_2 \ell(y_t, f_t(\xi_t)). \end{cases}$$

The equations of backpropagation give:

$$\begin{aligned}
dx_t^L &= W^{L+1}(t) \\
dw^L(t) &= m^{-a_{L+1}} x_t^L \\
dh_t^l &= dx_t^l \odot \sigma'(h_t^l) \\
dx_t^{l-1} &= (W^l(t))^\top dh_t^l \\
dw^l(t) &= m^{-a_l} dh_t^l (x_t^{l-1})^\top, \\
db^l(t) &= m^{-a_l} dh_t^l.
\end{aligned}$$

As noted in Definition 2.1 Remark 2.1, one has for $l \in [1, L]$,

$$\Delta w^l(t) = -\eta m^{-c_l} \chi_t dw^l(t) = -\eta m^{-(a_l+c_l)} \chi_t dh_t^l (x_t^{l-1})^\top, \quad (\text{A.2})$$

$$\Delta W^l(t) = m^{-a_l} \Delta w^l(t) = -\eta m^{-(2a_l+c_l)} \chi_t dh_t^l (x_t^{l-1})^\top, \quad (\text{A.3})$$

$$\Delta B^l(t) = m^{-a_l} \Delta b^l(t) = -\eta m^{-(2a_l+c_l)} \chi_t dh_t^l, \quad (\text{A.4})$$

and for $l = L + 1$

$$\Delta w^{L+1}(t) = -\eta m^{-c_{L+1}} \chi_t dw^{L+1}(t) = -\eta m^{-(a_{L+1}+c_{L+1})} \chi_t x_t^L, \quad (\text{A.5})$$

$$\Delta W^{L+1}(t) = m^{-a_{L+1}} \Delta w^{L+1}(t) = -\eta m^{-(2a_{L+1}+c_{L+1})} \chi_t x_t^L, \quad (\text{A.6})$$

$$\Delta B^{L+1}(t) = m^{-a_{L+1}} \Delta b^{L+1}(t) = -\eta m^{-(2a_{L+1}+c_{L+1})} \chi_t. \quad (\text{A.7})$$

Z variables. As described in Section B.2, the variables Z with a superscript will be used to denote the random variable whose law describes the evolution of all coordinates of a given vector of the forward or backward pass at a given layer in the limit $m \rightarrow \infty$.

Tilde variables. For $z \in \{h_t^l, x_t^l, dh_t^l, dx_t^l\}$, we will use \tilde{z} to denote a variable “without scale”, *i.e.*, such that $Z^{\tilde{z}}$ has positive and finite variance (see Definition F.1). When we do so, we always have $z = \lambda \tilde{z}$ for some scalar λ (which might depend on m). The tilde variables of the backward pass for $t \geq 1$ might have different expressions in different contexts or in different proofs, but we still use the same notation every time as the exact definition should always be clear from the context.

B An overview of the Tensor Program technique

The Tensor Program technique, first introduced by in Yang (2019), was initially developed to better understand the behavior at initialization of networks whose weights are initialized i.i.d. with standard Gaussians as the number of units in each layer grows to infinity. Since the output of a hidden unit in layer $l \geq 2$ is given by $\sum_{q=1}^m W_{pq}^l(0) x_{0,q}^{l-1}$, the magnitude of the weights need to be downscaled by some negative power of m to avoid blow-up as $m \rightarrow \infty$. Scalings which have naturally appeared in the literature are $m^{-1/2}$ and m^{-1} , and lead to different types of limits.

Using a first version of the Tensor Program (referred to as NETSOR), it is shown in (Yang, 2019) that the output at initialization of a neural network of **any architecture** (fully-connected, recurrent, convolutional, with normalization, attention, ...) whose weights are initialized with $W^l(0) = m^{-1/2} U^l$ for $l \geq 2$ (*i.e.*, $a_l = 0$ and $b_l = 1/2$ for $l \geq 2$ in the ac-parameterization) is a Gaussian process in the infinite-width limit.

Going further, and in the light of the recent literature on the neural tangent kernel, Yang (2020a) studies the first backward pass of networks initialized as above in the limit where

$m \rightarrow \infty$ and has shown that the neural tangent kernel at initialization, defined as $K(\xi, \bar{\xi}) := \langle \nabla_{\theta} f_0(\theta(0); \xi), \nabla_{\theta} f_0(\theta(0); \bar{\xi}) \rangle$ converges to a deterministic limit for any architecture.

Finally, and most importantly for our work, the Tensor Program is extended in (Yang, 2020b) to cover the forward and backward passes of networks of any architecture **at any time step** and not just at initialization. The crucial step taken in (Yang, 2020b) is to be able to describe the evolution of quantities where both a weight matrix W^l and its transpose $(W^l)^{\top}$ are involved. (Yang and Hu, 2021) then applies the results and theorems of (Yang, 2020b) in the particular context of ac-parameterizations (or rather abc-parameterizations as defined by Yang and Hu, 2021) to describe the infinite-width limits of neural networks with different parameterizations.

B.1 Intuition behind the technique

To explain the intuition behind the Tensor Program technique and how it comes into play for neural networks, let us first look at the forward pass of a fully-connected network with L hidden layers after t steps of SGD. Assume single samples $(\xi_0, y_0), \dots, (\xi_{t-1}, y_{t-1})$ are used at each step for simplicity. Consider a neural network in any ac-parameterization and an input ξ to the network. Using Equation (A.3) for the updates, the forward pass of the network at time t is given by:

$$\begin{aligned} h_t^1 &= W^1(0)\xi - \eta m^{-(2a_1+c_1)} \sum_{s=0}^{t-1} \chi_s (\xi_s^{\top} \xi) dh_s^1 \\ h_t^l &= W^l(0)x_t^{l-1} - \eta m^{-(2a_l+c_l)} \sum_{s=0}^{t-1} \chi_s \left((x_s^{l-1})^{\top} x_t^{l-1} \right) dh_s^l \quad l \in [2, L] \\ f_t(\xi) &= (W^{L+1}(0))^{\top} x_t^L - \eta m^{-(2a_{L+1}+c_{L+1})} \sum_{s=0}^{t-1} \chi_s (x_s^L)^{\top} x_t^L. \end{aligned}$$

To understand what happens in the forward pass, one thus needs to understand the behavior of the multiplication by i.i.d. Gaussian matrices, that of vectors dh_s^l of the backward pass as well as that of the inner products $(x_s^{l-1})^{\top} x_t^{l-1}$. As $m \rightarrow \infty$, the sums defining the matrix multiplications and inner products involve an infinity of terms and one must therefore understand how those quantities scale in the limit.

Before we dive into the matrix multiplications, let us look more precisely at what the vectors dh_s^l look like. We have:

$$\begin{aligned} dh_s^l &= dx_s^l \odot \sigma'(h_s^l) \\ dx_s^l &= (W^{l+1}(0))^{\top} dh_s^{l+1} - \eta m^{-(2a_l+c_l)} \sum_{u=0}^s \chi_u \left((dh_u^{l+1})^{\top} dh_s^{l+1} \right) x_u^l \quad l \in [2, L]. \end{aligned}$$

We observe that inner products appear again, and that in contrast with the forward pass, it is now the multiplication by the transpose of i.i.d. Gaussian matrices which appears.

We already see that two main quantities appear in the calculations: The initial i.i.d. Gaussian matrices, and vectors which are generated either (i) through the multiplication of another vector with a Gaussian matrix or its transpose, or (ii) through some form of non-linearity involving other vectors as well as the activation function σ and/or its derivative σ' . Before trying to understand how the inner products behave, let us first dive into the multiplication by i.i.d. Gaussian matrices.

B.1.1 Multiplication by i.i.d. Gaussian matrices

The multiplication of a random vector by an i.i.d. Gaussian matrix can happen in two different scenarios: (i) the input vector is independent of the Gaussian weights, and (ii) the input vector is correlated with the Gaussian weights, which, in the case of neural networks, will translate into saying that the transpose of the weight matrix is used somewhere to compute the input vector.

Independent input vector. Consider a list $(x_q)_{q \in \mathbb{N}^*}$ of i.i.d. random variables with finite first and second moments, independent of U^l , and consider multiplying this vector by the i.i.d. Gaussian matrix U^l . At any finite-width m the p -th entry of $U^l x$ is given by

$$\sum_{q=1}^m U_{pq}^l x_q \underset{m \rightarrow \infty}{\simeq} m^{1/2} \mathcal{N}(0, \mathbb{E}[x_1^2])$$

The terms $(U_{pq}^l x_q)_{q \geq 1}$ are i.i.d. with mean 0 and finite variance $\mathbb{E}[x_1^2]$ because x_q is independent of U_{pq}^l . Therefore, by a central limit argument, the sum will behave like $m^{1/2} \mathcal{N}(0, \mathbb{E}[x_1^2])$ for large m . It is thus natural to scale the sum by $m^{-1/2}$, or equivalently to consider $\widehat{W}^l = m^{-1/2} U^l$ (as defined in Equation A.1) for matrix multiplications.

With the above result in mind, we take a look at the first forward pass at initialization of a network where all the weight matrices are initialized as $W^l(0) = \widehat{W}^l$ (i.e., $a_l = 0$, $a_l = 1/2$, $l \in [2, L + 1]$). We consider an input $\xi \in \mathbb{R}^d$ to the network and compute the pre-activations of each layer recursively. For the first layer, we get that for any $p \in [m]$,

$$\begin{aligned} h_{0,p}^1 &= (\widehat{W}^1 \xi)_p = (U^1 \xi)_p \\ &= \sum_{q=1}^d \xi_q U_{pq}^1 \sim \mathcal{N}(0, \|\xi\|^2) \end{aligned}$$

Since the $(U_{pq}^1)_q$ are i.i.d. standard Gaussians, the linear combination above is also a Gaussian with mean 0 and variance $\sum_q \xi_q^2 = \|\xi\|^2$. Note that since the lists $(U_{pq}^1)_q$ are independent for different p , the vector h_0^1 has i.i.d. coordinates all distributed as $\mathcal{N}(0, \|\xi\|^2)$. We also note that adding a bias term initialized as $\mathcal{N}(0, 1)$ would simply change the variance to $\|\xi\|^2 + 1$.

Then for the second layer we get that for any $p \in [m]$:

$$h_{0,p}^2 = \frac{1}{\sqrt{m}} \sum_{q=1}^m U_{pq}^2 \sigma(h_{0,q}^1) \xrightarrow{m \rightarrow \infty} \mathcal{N}(0, \mathbb{E}[\sigma(h_{0,1}^1)^2])$$

The terms $(U_{pq}^2 \sigma(h_{0,q}^1))_q$ are i.i.d. with mean zero, and by a central limit argument, we have that the coordinates of h_0^2 converge in law towards $\mathcal{N}(0, \mathbb{E}[\sigma(h_{0,1}^1)^2])$ where $\mathbb{E}[\sigma(h_{0,1}^1)^2]$ is simply $\mathbb{E}[\sigma(Z)^2]$ with $Z \sim \mathcal{N}(0, \|\xi\|^2)$. Those coordinates are also independent (and Gaussian at any finite width m) **conditionally** on h_0^1 because the lists $(U_{pq}^2)_q$ are independent (Gaussians) for different p . The different coordinates of h_0^2 are identically distributed at any finite width m and remain so in the limit. They are not strictly speaking independent at finite width but the intuition is that they become so in the limit $m \rightarrow \infty$ as they also become Gaussian, and that is how they should be thought of in the context of the Tensor Program.

Repeating the calculations above at every layer, we can intuitively describe the forward pass in the infinite-width limit by describing the law of a single random variable Z_l for each

layer (whose law is the common law of all the coordinates of the pre-activations h_0^l), and by the hand-wavy calculations above, we get the following recursion for the variables Z :

$$\begin{aligned} Z_1 &\sim \mathcal{N}(0, \|\xi\|^2) \\ Z_{l+1} &\sim \mathcal{N}(0, \mathbb{E}[\sigma(Z_l)^2]), \quad l \in [1, L] \end{aligned}$$

Having discussed the case where the input vectors are not correlated with the weight matrix, we now move on to the case where there is some correlation between the two.

Correlated input vector. As the simplest form of correlation, we consider a vector $x = (\widehat{W}^l)^\top z$ where $(z_q)_{q \in \mathbb{N}^*}$ is a list of i.i.d. random variables independent of \widehat{W}^l with finite first and second moments, and we consider the result of the multiplication $h = \widehat{W}^l x$. For any $p \in [m]$, we have

$$\begin{aligned} h_p &= \sum_{q=1}^m \sum_{r=1}^m \widehat{W}_{pq}^l \widehat{W}_{rq}^l z_r \\ &= \left[\frac{1}{m} \sum_{q=1}^m \left(U_{pq}^l \right)^2 \right] z_p + \frac{1}{\sqrt{m}} \sum_{r \neq p} z_r \left(\frac{1}{\sqrt{m}} \sum_{q=1}^m U_{pq}^l U_{rq}^l \right) \end{aligned}$$

By the law of large numbers, the first term will converge almost surely to z_p as $m \rightarrow \infty$. For the second term, the intuition is that for any $r \neq p$ the terms $(1/\sqrt{m}) \sum_q U_{pq}^l U_{rq}^l$ become distributed as independent Gaussians as m becomes large by a central limit argument. Then, by another central limit argument, intuitively, the sum over $r \neq p$ should also become distributed as $\mathcal{N}(0, \mathbb{E}[z_1^2])$. In the limit $m \rightarrow \infty$, we thus expect the coordinates of $h = \widehat{W}^l (\widehat{W}^l)^\top z$ to be the sum of two terms: a first term distributed as z_1 where the correlation between the entries of \widehat{W}^l and $(\widehat{W}^l)^\top$ comes into play, and a second term distributed as $\mathcal{N}(0, \mathbb{E}[z_1^2])$ which is purely Gaussian and where the correlation between the entries of \widehat{W}^l and $(\widehat{W}^l)^\top$ has no effect.

The aim of the Tensor Program series (Yang, 2019, 2020a,b) is to formalize those intuitions into theorems and rigorous calculations. Of course, the calculations become more complex when we introduce non-linearities and consider later steps in training than the initialization, but what the Tensor Program shows is that the intuitions above still hold.

To summarize, the intuition is that in the large-width limit, the coordinates of pre-activation vectors become i.i.d. and we thus only need to track the law of a single real-valued random variable. Therefore, any average of some function of the coordinates should converge to an expectation in the limit $m \rightarrow \infty$ by a law of large number argument. Finally, any multiplication by \widehat{W}^l yields two terms where one is purely Gaussian and the other depends on the expression of the vector that is multiplied by \widehat{W}^l in function of $(\widehat{W}^l)^\top$.

B.2 Mathematical formalism

The mathematical formalism of the Tensor Program goes beyond neural network computations and describes the evolution of any computational systems (with some restrictions) in the limit $m \rightarrow \infty$. The computational system is comprised of different vectors whose dimensions are equal to m which can be generated from a set of initial vectors in various ways. The Tensor Program is defined by the sequence of mathematical operations which produce the vectors from previously generated vectors. The operations are the same at any

given width m , only the size of the vectors and matrices involved change with m , and the aim of the Tensor Program is to provide the tools (formalism and theorems) to be able to describe the behavior of the system in the limit $m \rightarrow \infty$. As described in the intuitions of the previous section B.1, the coordinates of vectors in the program are roughly i.i.d. as $m \rightarrow \infty$ and variables Z are introduced to describe the common law of the coordinates in the limit $m \rightarrow \infty$.

Initial vectors. Consider a set $\mathcal{V} := \{v^1, \dots, v^N\} \in (\mathbb{R}^m)^N$ of *initial vectors* such that:

- (i) the coordinates $(v_p)_{p \in [m]}$ are i.i.d. for any $v \in \mathcal{V}$ and any m . We call Z^v a real-valued random variable whose law is the same as that of all the coordinates.
- (ii) The joint law of $Z^{\mathcal{V}} := (Z^{v^1}, \dots, Z^{v^N})$ is a Gaussian $\mathcal{N}(\mu_{\text{init}}, \Sigma_{\text{init}})$ for any m (the variables Z^v do not actually depend on m , but this is simply to say that at any width m and for any $p \in [m]$, the law of (v_p^1, \dots, v_p^N) is the same N -dimensional Gaussian).

Initial scalars. Similarly, we define a list of initial scalars $\theta_1, \dots, \theta_M$ which can depend on m and for which the only requirement is that each θ_r converges almost surely to some finite limit $\overset{\circ}{\theta}_r$ as $m \rightarrow \infty$.

Initial Gaussian matrices. Consider a set $\mathcal{W} := \{\widehat{W}^1, \dots, \widehat{W}^P\} \in (\mathbb{R}^{m \times m})^P$, such that $\widehat{W}_{pq}^r \sim \mathcal{N}(0, 1/m)$ i.i.d. over p, q for any r , and the $(\widehat{W}^r)_{r \in [P]}$ are independent of each other and independent of the vectors in \mathcal{V} . Since we consider a more general setting than neural networks, we do not index those matrices by l and can have $P \neq L$ but for neural networks, those initial matrices will always be the initialization of the weight matrices of the intermediate layers $l \in [2, L]$, appropriately scaled.

Generation of new vectors/scalars. Given previously generated vectors v^1, \dots, v^k , previously generated scalars $\theta_1, \dots, \theta_r$, and a non-linearity $\psi(\cdot; \cdot) : \mathbb{R}^k \times \mathbb{R}^r \rightarrow \mathbb{R}$, we can, in the following ways, generate:

MatMul a vector $z = \widehat{W}v$ for any $v \in \{v^1, \dots, v^k\}$ and $\widehat{W} \in \mathcal{W}$.

NonLin a vector $z = \psi(v^1, \dots, v^k; \theta_1, \dots, \theta_r)$ where ψ is taken element-wise, *i.e.*, $z_p = \psi(v_p^1, \dots, v_p^k; \theta_1, \dots, \theta_r)$ for any $p \in [m]$ and for any m .

Moment a scalar $\omega = \frac{1}{m} \sum_{p=1}^m \psi(v_p^1, \dots, v_p^k; \theta_1, \dots, \theta_r) \in \mathbb{R}$.

The non-linearity used does not have to actually depend on all the previous vectors and/or scalars, but we present the operations this way for simplicity.

Given those operations, the Tensor Program framework allows to seamlessly describe the infinite-width limit of the computational system defining a given Tensor Program by tracking recursively the laws of the variables Z whose law represents the common law of the coordinates of a given vector. Indeed, every vector z in the program (initial or generated using previous vectors in the program) will roughly have i.i.d. coordinates in the limit $m \rightarrow \infty$, and the Tensor Program associates a real-valued random variable Z^z to the vector z . Then, associated with the operations on vectors and scalars above are the following operations on the corresponding variables Z which come as their natural counterparts in the infinite-width limit to track the evolution of the laws of the variables Z :

ZInit For initial vectors $v \in \mathcal{V}$, define $\dot{Z}^v = 0$ and $\widehat{Z}^v = Z^v$. The purpose of those notations will become clear in the **ZMatMul** section.

ZMoment Given a scalar $\omega = (1/m) \sum_{p=1}^m \psi(z_p^1, \dots, z_p^k; \theta_1, \dots, \theta_r)$, define

$$\mathring{\omega} = \mathbb{E} \left[\psi(Z^{z^1}, \dots, Z^{z^k}; \mathring{\theta}_1, \dots, \mathring{\theta}_r) \right] \quad (\text{B.1})$$

ZNonLin Given $z = \psi(z^1, \dots, z^k; \theta_1, \dots, \theta_r)$, define:

$$Z^z = \psi(Z^{z^1}, \dots, Z^{z^k}; \mathring{\theta}_1, \dots, \mathring{\theta}_r) \quad (\text{B.2})$$

ZMatMul Given $z = \widehat{W}v$ for a previous vector v and $\widehat{W} \in \mathcal{W}$, $Z^z = \widehat{Z}^z + \dot{Z}^z$ is the sum of two terms:

ZHat $\widehat{Z}^z \sim \mathcal{N} \left(0, \mathbb{E} \left[(Z^v)^2 \right] \right)$ is a purely Gaussian term. Additionally, if we let $\mathcal{W}_{\widehat{W}}$ be the set of all vectors in the program of the form $\widehat{W}u$ for some u in the program, the vector $Z^{\mathcal{W}_{\widehat{W}}} = (Z^h)_{h \in \mathcal{W}_{\widehat{W}}}$ is defined to be jointly Gaussian with covariance matrix given by:

$$\text{cov}(Z^{W^x}, Z^{W^y}) = \mathbb{E}[Z^x Z^y]$$

Moreover, the vector $Z^{\mathcal{W}_{\widehat{W}}}$ is defined to be mutually independent of the list of Z^u for u in $\{\widehat{Z}^v : v \in \mathcal{V} \cup_{W \in \mathcal{W} \cup \mathcal{W}^\top, W \neq \widehat{W}} \mathcal{W}_W\}$ where $\mathcal{W}^\top := \{\widehat{W}^\top : \widehat{W} \in \mathcal{W}\}$, and \mathcal{W}_W is the set of vectors in the program of the form Wu for some vector u in the program.

ZDot \dot{Z}^z comes from the potential interactions (correlations) between \widehat{W} and \widehat{W}^\top in the computation of z . One can always unwind the expression of Z^v and express it in function of the $\widehat{Z}^{\widehat{W}^\top y}$ for some x in the program, that is we can always write Z^v as $Z^v = \phi(\widehat{Z}^{\widehat{W}^\top y^1}, \dots, \widehat{Z}^{\widehat{W}^\top y^k}, \widehat{Z}^{x^1}, \dots, \widehat{Z}^{x^r}; \mathring{\theta}_1, \dots, \mathring{\theta}_s)$ with x^1, \dots, x^r such that \widehat{W}^\top is never used in the computation of those vectors. Then, define:

$$\dot{Z}^z = \sum_{j=1}^k \mathbb{E} \left[\frac{\partial Z^v}{\partial \widehat{Z}^{\widehat{W}^\top y_j}} \right] Z^{y_j} \quad (\text{B.3})$$

where $\partial Z^v / \partial \widehat{Z}^{\widehat{W}^\top y_j}$ is simply defined as the j -th partial derivative of ϕ above when expressing Z^v as required for \dot{Z} . As noted in (Yang and Hu, 2021), if ϕ is not everywhere differentiable, one can leverage Stein's lemma to replace the formula in Equation (B.3) by a linear algebra formula.

Now that we have introduced the necessary concepts and described the content of a Tensor Program, we can move on to present the main theorem derived in (Yang and Hu, 2021) which connects the mathematical operations used at finite-width with the infinite-width limit of the computational system defining a Tensor Program. The “*master theorem*” formulated in (Yang and Hu, 2021) is surprisingly simple (although the proof is much more intricate) yet very powerful, and goes as follows (see Yang and Hu, 2021, Theorem 7.4):

Theorem B.1 (Master Theorem). *Given a Tensor Program, for any vectors x^1, \dots, x^k and scalars $\theta_1, \dots, \theta_r$ in the program, and for any pseudo-Lipschitz non-linearity ψ (see Definition 3.1, page 9), one has that:*

$$\frac{1}{m} \sum_{p=1}^m \psi(x_p^1, \dots, x_p^k; \theta_1, \dots, \theta_r) \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E} \left[\psi \left(Z^{x^1}, \dots, Z^{x^k}; \mathring{\theta}_1, \dots, \mathring{\theta}_r \right) \right]$$

Remark B.1.

1. The theorem essentially states that even though the coordinates of vectors in the program are not rigorously i.i.d., they appear so from the perspective of the average by a suitable non-linearity so that a law of large number type of result holds. Note that for neural networks, even though the coordinates of the (pre-)activations follow the same law when using i.i.d. initialization for the weights, it is not *a priori* clear that we can consider them as independent copies, and thus that we can summarize the computations using a single real-valued variable, but the master theorem shows that from the perspective of averaging, this is in fact the case in the infinite-width limit.
2. In (Yang and Hu, 2021), different versions of the Tensor Program are presented in the sense that different classes of non-linearities are allowed. These differences induce minor subtleties in the master theorem and in the proofs. However, most of the results in the main text of the paper require that the non-linearities be pseudo-Lipschitz (which is the stronger assumption), both in NonLin and in the master theorem. The Assumption 2 on the activation function σ and its derivative σ' ensures that any quantity appearing in the forward or backward computation of a neural network can be expressed as pseudo-Lipschitz non-linearity.
3. What the Tensor Program and its master theorem show is that to understand the behavior of the computational system in the infinite-width limit, one simply needs to track the operations on the variables Z which mimic the recursive operations in the computational system. Then, quantities which involve sum over coordinates such as inner products between the vectors in the program (which occur in the forward and backward passes of a neural network, as well as in the computation of the neural tangent kernel), or norm computations are easily described, when properly re-normalized, through expectations involving the corresponding variables Z . The main difficulty is that it is actually hard (computationally and in the mathematical formulation) to track the correlations between different Z because, as explained in (Yang and Hu, 2021), of the necessary unwinding in the definition of \dot{Z} , so that the computational graph associated with the operations on the variables Z is hard to implement in practice.

B.3 The maximal update parameterization μP

We close this section by presenting briefly the maximal update parameterization considered in (Yang and Hu, 2021). To quantify the learning abilities of a given parameterization, Yang and Hu (2021) introduce the notions of *feature learning* and *feature kernel evolution* at a given layer $l \in [1, L]$, which we recall below. Both these definitions concern the large-width limit of the networks:

Definition B.1 (Feature Learning). An ac-parameterization is said to admit **feature learning** at the l -th layer if the quantity $\Delta x_t^l(\xi) := x_t^l(\xi) - x_0^l(\xi)$ is such that there exists a training routine for which, almost surely, there exists a constant $C > 0$ such that $\|\Delta x_t^l(\xi)\|^2/m \geq C$ for large enough m .

Definition B.2 (Kernel Evolution). An ac-parameterization is said to **evolve the feature kernel** at the l -th layer if the quantity $\Delta F_t^l(\xi, \bar{\xi}) := [x_t^l(\xi)^\top x_t^l(\bar{\xi}) - x_0^l(\xi)^\top x_0^l(\bar{\xi})]/m$ is such that there exists a training routine for which, almost surely, there exists a constant $C > 0$ such that for large enough m , $\Delta F_t^l(\xi, \bar{\xi}) \geq C$.

(Yang and Hu, 2021) goes about categorizing whether different ac-parameterizations admit feature learning or not. One of the striking result presented is that there is essentially a

dichotomy (depending on the values of $(a_l, c_l)_{l \in [L+1]}$) among ac-parameterizations: an ac-parameterization either admits feature learning (and evolves the feature kernel) or is in the kernel regime, meaning that the quantities in definitions B.1 and B.2 converge to 0 almost surely so that in the infinite width limit, the evolution of the prediction function f_t is deterministic and depends only on the previous prediction function f_{t-1} and the loss at time $(t-1)$ through a (deterministic) kernel $K(\xi, \bar{\xi}) = \lim_{m \rightarrow \infty} (x_0^L(\xi))^T x_0^L(\bar{\xi})/m$ (or a rescaled version thereof).

The categorization result proved in (Yang and Hu, 2021) holds for a certain class of ac-parameterizations which are deemed *stable* and *non-trivial*. Stable refers to the fact that the pre-activations and output (h_0^l and $f_0(\xi)$ respectively) at initialization do not blow-up as $m \rightarrow \infty$ at any layer. As already hinted in Section B.1, this corresponds to having $a_1 = 0$ and $a_l \geq 1/2$ for $l \in [2, L+1]$. Non-trivial refers to the fact that the pre-activations of all layers do not converge to 0 almost surely as $m \rightarrow \infty$ at initialization. This corresponds to having $a_1 \leq 0$ and $a_l \leq 1/2$ for $l \in [2, L]$. It is mentioned in (Yang and Hu, 2021) that those parameterizations for which the pre-activations of the intermediate layers converge to 0 almost surely should stay at their initialization throughout the course of training, and we actually prove in Section 3, using the Tensor Program technique, that this is the case when $L \geq 3$ in the setting where $a_1 = 0$ and $a_l = 1$ for $l \in [2, L+1]$ (i.e., integrable parameterizations) unless one uses large (polynomial in m) initial learning rates, a scenario which is not covered in (Yang and Hu, 2021). We show that in this case, integrable parameterizations are only trivial at initialization (the pre-activations of all layers except the first one converge to 0 in the infinite-width limit) and are actually in a feature learning regime at all layers after the first gradient step ($t \geq 1$).

The maximal update parameterization μP introduced in (Yang and Hu, 2021) is the result of the analysis of the values of a_l , and c_l for which the parameterization admits feature learning at every layer, and maximally so in the sense that if we were to reduce the value of a_l then the Δx_t^l introduced in Definition B.1 or the pre-activations h_t^l would blow-up as $m \rightarrow \infty$. In essence, μP corresponds to the values of a_l , and c_l for which Δx_t^l is as large as possible (with regards to its dependency on m) at every layer without creating any instabilities (pre-activations or updates blowing-up) in the limit $m \rightarrow \infty$. A quick analysis of the updates at $t = 0$ shows that the choice $a_1 = 0$, $a_l = 1/2$ for $l \in [2, L]$, and $a_{L+1} = 1$ associated with $c_l = -1$ for all $l \in [L+1]$ achieves this, and it is rigorously shown in (Yang and Hu, 2021) that this choice of ac-parameterization induces an update such that, $\|\Delta W^l(t)x_t^{l-1}\|^2/m = \Theta(1)$. We thus adopt the following definition for μP which is the same as in (Yang and Hu, 2021, Definition 5.1) but re-parameterized to remove the redundant b in the abc-parameterization:

Definition B.3 (μP). The maximal update parameterization μP is defined by the following choice of parameterization:

$$\begin{aligned} a_1 &= 0, & c_1 &= -1, \\ a_l &= 1/2, & c_l &= -1, \quad l \in [2, L], \\ a_{L+1} &= 1, & c_{L+1} &= -1. \end{aligned}$$

C Useful preliminary results

We show in this section a couple of useful results which will prove helpful in the proofs.

C.1 Positive finite moments of pseudo-Lipschitz functions of Gaussians

Lemma C.1 (Positive finite moments with polynomially bounded non-linearities). *Let ϕ be a polynomially bounded non-linearity which is not almost everywhere 0, and let $Z \sim \mathcal{N}(0, v^2)$ with $v^2 < \infty$. Then, for any $p \in \mathbb{R}_+$:*

(i) $0 \leq \mathbb{E}[|\phi(Z)|^p] < \infty$,

(ii) if in addition $v^2 > 0$, $0 < \mathbb{E}[|\phi(Z)|^p] < \infty$.

Proof. If $v^2 = 0$, and then $\phi(Z) = \phi(0)$ almost surely, so that $\mathbb{E}[|\phi(Z)|^p] = |\phi(0)|^p < \infty$.

Now, assume $v^2 > 0$. Since ϕ is bounded by a polynomial of some degree $r > 0$, $|\phi(z)| \leq C(1 + |z|^r)$ for some $C > 0$. Then, $|\phi(z)|^p = \exp(p \ln(|\phi(z)|)) \leq C^p(1 + |z|^r)^p$. Since $v^2 > 0$, we have

$$\begin{aligned} \mathbb{E}[|\phi(Z)|^p] &= \frac{1}{\sqrt{2\pi v^2}} \int_{\mathbb{R}} |\phi(z)|^p e^{-z^2/2v^2} dz \\ &\leq \frac{1}{\sqrt{2\pi v^2}} \int_{\mathbb{R}} C^p(1 + |z|^r)^p e^{-z^2/2v^2} dz < \infty. \end{aligned}$$

Finally, since ϕ is not almost everywhere 0, neither is $|\phi|^p$ which shows the integral in the first equality above is not 0, and gives $\mathbb{E}[|\phi(Z)|^p] > 0$. \square

C.2 The Z dots are 0 in the first forward-backward pass

Lemma C.2 ($\dot{Z} = 0$ in the first forward-backward pass). *Consider an ac-parameterization of an L -hidden layer fully-connected neural network with $a_1 \geq 0$ and $a_l \geq 1/2$ for $l \in [2, L+1]$, and with a non-linearity satisfying Assumption 2. Then for any $l \geq 2$, $\dot{Z} \widehat{W}^l x_0^{l-1} = 0$, and for any $l \in [1, L]$, $\dot{Z} (\widehat{W}^l)^\top dh_0^l = 0$.*

Remark C.1. This lemma applies to the NTK, μP , and integrable parameterizations (in particular IP-LLR) as well as HP and HPZ.

Proof. Consider any ac-parameterization of a fully-connected neural network which has $a_1 \geq 0$ and $a_l \geq 1/2$ for $l \in [2, L+1]$, and with a non-linearity satisfying Assumption 2. Define $\omega_1 = m^{-a_1}$ and $\omega_l = m^{-(a_l-1/2)}$ for $l \geq 2$, and the initial scalar $\alpha_{L+1} := m^{-a_{L+1}}$. The conditions on the a_l guarantee that the ω_l converge almost surely to either 0 or 1 and α_{L+1} converges almost surely to 0, which allows applying the rules of the Tensor Program.

For any $l \in [2, L]$, since the computation of x_0^{l-1} , and thus of $Z^{x_0^{l-1}}$ do not involve $(\widehat{W}^l)^\top$, $\dot{Z} \widehat{W}^l x_0^{l-1} = 0$ as per the ZDot rule of the Tensor Program. In addition, $Z^{h_0^l} = \omega_1 (\widehat{Z}^{\widehat{W}^1 \xi} + \widehat{Z}^{v^1})$ and by definition, $\widehat{Z}^{\widehat{W}^1 \xi} \sim \mathcal{N}(0, \|\xi\|^2)$ and $\widehat{Z}^{v^1} \sim \mathcal{N}(0, 1)$ are independent Gaussians, which shows that $Z^{h_0^l} \sim \mathcal{N}(0, \omega_1^2 (\|\xi\|^2 + 1))$ whose variance is finite because $\omega_1^2 \in \{0, 1\}$. By Lemma C.1, this also shows that $\mathbb{E}[(Z^{x_0^1})^2] < \infty$. Let $l \in [2, L]$ and assume that $\mathbb{E}[(Z^{h_0^{l-1}})^2] < \infty$ and $\mathbb{E}[(Z^{x_0^{l-1}})^2] < \infty$. We have $h_0^l = \omega_l \widehat{W}^l x_0^{l-1} + m^{-a_l} v^l$. Since m^{-2a_l} converges to 0 almost surely, we can consider it as an initial scalar in the program, which gives by ZNonLin $Z^{h_0^l} = \omega_l \widehat{Z}^{\widehat{W}^l x_0^{l-1}} + 0 \times \widehat{Z}^{v^l}$. $\widehat{Z}^{v^l} \sim \mathcal{N}(0, 1)$ by definition since v^l is an initial vector in the program, so that $Z^{h_0^l} = \omega_l \widehat{Z}^{\widehat{W}^l x_0^{l-1}} \sim \mathcal{N}(0, \omega_l^2 \mathbb{E}[Z^{x_0^{l-1}}]^2)$ whose variance is finite by the induction hypothesis and because $\omega_l \in \{0, 1\}$. Then by Lemma C.1, we also get that $\mathbb{E}[(Z^{x_0^l})^2] < \infty$, which concludes the induction.

Let us now deal with the first backward pass for any ac-parameterization. The result will essentially boil down to having the expectation of the derivatives defining the \dot{Z} being 0 because the weight matrices are initialized with 0 mean and because of an independence argument. We have $dx_0^L = W^{L+1}(0) = m^{-\alpha_{L+1}}U^{L+1}$, and $dh_0^L = dx_0^L \odot \sigma'(Z^{h_0^L})$. By ZNonLin we thus have

$$\begin{aligned} Z^{dx_0^L} &= \overset{\circ}{\alpha}_{L+1} Z^{U^{L+1}}, \\ Z^{dh_0^L} &= \overset{\circ}{\alpha}_{L+1} Z^{U^{L+1}} \sigma'(Z^{h_0^L}). \end{aligned}$$

Now let $l \in [1, L]$. $dx_0^{l-1} = (\widehat{W}^l)^\top dh_0^l$ gives

$$Z^{(\widehat{W}^l)^\top dh_0^l} = \widehat{Z}^{(\widehat{W}^l)^\top dh_0^l} + \dot{Z}^{(\widehat{W}^l)^\top dh_0^l},$$

and to understand what $\dot{Z}^{(\widehat{W}^l)^\top dh_0^l}$ is, we need to expand the expression of $Z^{dh_0^l}$ in function of variables which were generated with \widehat{W}^l . So far, the only variable where \widehat{W}^l was used is $h_0^l = \omega_l \widehat{W}^l x_0^{l-1}$ (with the convention that $x_0^0 = \xi_0$). We thus need to expand the expression of $Z^{dh_0^l}$ in function of $\widehat{Z}^{\widehat{W}^l x_0^{l-1}}$. We have, for $l = L$

$$\begin{aligned} Z^{dh_0^L} &= \overset{\circ}{\alpha}_{L+1} Z^{U^{L+1}} \sigma'(\overset{\circ}{\omega}_L Z^{\widehat{W}^L x_0^{L-1}}) \\ &= \overset{\circ}{\alpha}_{L+1} \widehat{Z}^{U^{L+1}} \sigma'(\overset{\circ}{\omega}_L \widehat{Z}^{\widehat{W}^L x_0^{L-1}}), \end{aligned}$$

where the last equality stems from the fact that $Z^{\widehat{W}^L x_0^{L-1}} = \widehat{Z}^{\widehat{W}^L x_0^{L-1}}$ in the first forward pass, and the fact that U^{L+1} is an initial vector in the program which gives by definition $\widehat{Z}^{U^{L+1}} = Z^{U^{L+1}}$. We can formally write this as

$$Z^{dh_0^L} = \Psi(\widehat{Z}^{\widehat{W}^L x_0^{L-1}}, \widehat{Z}^{U^{L+1}}; \overset{\circ}{\alpha}_{L+1}, \overset{\circ}{\omega}_L),$$

where $\Psi(z_1, z_2; \theta_1, \theta_2) := \theta_1 z_2 \sigma'(\theta_2 z_1)$ is a pseudo-Lipschitz function because σ' is, and we have

$$\frac{\partial \Psi}{\partial z_1}(z_1, z_2; \theta_1, \theta_2) = \theta_1 \theta_2 z_2 \sigma''(\theta_2 z_1).$$

We get that by definition

$$\begin{aligned} \dot{Z}^{(\widehat{W}^L)^\top dh_0^L} &= \mathbb{E} \left[\frac{\partial Z^{dh_0^L}}{\partial \widehat{Z}^{\widehat{W}^L x_0^{L-1}}} \right] Z^{x_0^{L-1}} \\ &= \mathbb{E} \left[\frac{\partial \Psi}{\partial z_1}(\widehat{Z}^{\widehat{W}^L x_0^{L-1}}, \widehat{Z}^{U^{L+1}}; \overset{\circ}{\alpha}_{L+1}, \overset{\circ}{\omega}_L) \right] Z^{x_0^{L-1}} \\ &= \overset{\circ}{\alpha}_{L+1} \overset{\circ}{\omega}_L \mathbb{E}[Z^{U^{L+1}} \sigma''(\overset{\circ}{\omega}_L \widehat{Z}^{\widehat{W}^L x_0^{L-1}})] Z^{x_0^{L-1}} \\ &= \overset{\circ}{\alpha}_{L+1} \overset{\circ}{\omega}_L \underbrace{\mathbb{E}[\widehat{Z}^{U^{L+1}}]}_0 \underbrace{\mathbb{E}[\sigma''(\overset{\circ}{\omega}_L \widehat{Z}^{\widehat{W}^L x_0^{L-1}})]}_{< \infty} \underbrace{Z^{x_0^{L-1}}}_{< \infty \text{ a.s.}}, \end{aligned}$$

where the last equality stems from the fact that by ZHat, $\widehat{Z}^{\widehat{W}^L x_0^{L-1}}$ is independent of $\widehat{Z}^{U^{L+1}}$ because U^{L+1} is an initial vector in the program. The fact that the second expectation finite is because $\overset{\circ}{\omega}_L \in \{0, 1\}$, σ'' is polynomially bounded, and $\widehat{Z}^{\widehat{W}^L x_0^{L-1}}$ is a Gaussian with mean 0 and finite variance since $\mathbb{E}[(Z^{x_0^{L-1}})^2] < \infty$. This gives $\dot{Z}^{(\widehat{W}^L)^\top dh_0^L} = 0$.

Now suppose $l \in [1, L - 1]$ and assume $\dot{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1} = 0$ which gives $Z(\widehat{W}^{l+1})^\top dh_0^{l+1} = \widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}$. We have

$$\begin{aligned} Z^{dh_0^l} &= Z^{dx_0^l} \sigma'(Z^{h_0^l}) \\ &= \mathring{\omega}_{l+1} Z(\widehat{W}^{l+1})^\top dh_0^{l+1} \sigma'(\mathring{\omega}_l Z \widehat{W}^{l+1} x_0^{l-1}) \\ &= \mathring{\omega}_{l+1} \widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1} \sigma'(\mathring{\omega}_l \widehat{Z} \widehat{W}^{l+1} x_0^{l-1}) \end{aligned}$$

where we have used that previous \dot{Z} are 0 to replace the Z with \widehat{Z} . We can once more formally write this as

$$Z^{dh_0^l} = \Psi(\widehat{Z} \widehat{W}^{l+1} x_0^{l-1}, \widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}; \mathring{\omega}_{l+1}, \mathring{\omega}_l)$$

with exactly the same Ψ as for $l = L$. We get that by definition

$$\begin{aligned} \dot{Z}(\widehat{W}^l)^\top dh_0^l &= \mathbb{E} \left[\frac{\partial Z^{dh_0^l}}{\partial \widehat{Z} \widehat{W}^l x_0^{l-1}} \right] Z^{x_0^{l-1}} \\ &= \mathbb{E} \left[\frac{\partial \Psi}{\partial z_1}(\widehat{Z} \widehat{W}^l x_0^{l-1}, \widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}; \mathring{\omega}_{l+1}, \mathring{\omega}_l) \right] Z^{x_0^{l-1}} \\ &= \mathring{\omega}_{l+1} \mathring{\omega}_l \mathbb{E}[\widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1} \sigma''(\mathring{\omega}_l \widehat{Z} \widehat{W}^l x_0^{l-1})] Z^{x_0^{l-1}} \\ &= \mathring{\omega}_{l+1} \mathring{\omega}_l \underbrace{\mathbb{E}[\widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}]}_0 \underbrace{\mathbb{E}[\sigma''(\mathring{\omega}_l \widehat{Z} \widehat{W}^l x_0^{l-1})]}_{< \infty} \underbrace{Z^{x_0^{l-1}}}_{< \infty a.s.} \\ &= 0 \end{aligned}$$

Where the first expectation is 0 because by definition $\widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}$ is a Gaussian with 0 mean and an easy induction (from $l = L$ to $l = 1$) shows that, as for the forward pass, $\mathbb{E}[(Z^{dx_0^l})^2] < \infty$ and $\mathbb{E}[(Z^{dh_0^l})^2] < \infty$, which implies that $\widehat{Z}(\widehat{W}^{l+1})^\top dh_0^{l+1}$ has finite variance. The second expectation is finite because $\mathring{\omega}_l \in \{0, 1\}$, $\widehat{Z} \widehat{W}^l x_0^{l-1}$ is a Gaussian with 0 mean by definition and finite variance, and because σ'' is polynomially bounded since σ' is pseudo-Lipschitz. \square

C.3 Gaussian output in the infinite-width limit

Lemma C.3 (Gaussian output). *For every $m \in \mathbb{N}^*$, let x^m and w^m be independent random vectors in \mathbb{R}^m such that*

$$\begin{cases} \frac{1}{m} \|x^m\|^2 \xrightarrow[m \rightarrow \infty]{a.s.} \sigma_\infty^2 \\ w_j^m \sim \mathcal{N}(0, 1/m) \text{ i.i.d. over } j = 1, \dots, m. \end{cases}$$

Then

$$(w^m)^\top x^m \xrightarrow[m \rightarrow \infty]{law} \mathcal{N}(0, \sigma_\infty^2)$$

Proof. Consider two sequences of independent vectors of growing dimension $(w^m)_m$ and $(x^m)_m$ as in Lemma C.3. Conditionally on x^m , the random variable $(w^m)^\top x^m$ follows a Gaussian $\mathcal{N}(0, \|x^m\|^2/m)$ distribution. Since $\|x^m\|^2/m$ converges to σ_∞^2 almost surely, the conditional distribution of $(w^m)^\top x^m$ given x^m converges to a Gaussian $\mathcal{N}(0, \sigma_\infty^2)$ distribution. The lemma follows. \square

C.4 Convergence of the coordinates to the limiting distribution Z

Lemma C.4 (Convergence to the limit distribution). *For any vector h in the Tensor Program we have for any $\alpha \in \mathbb{N}^*$,*

$$h_\alpha \xrightarrow[m \rightarrow \infty]{\text{law}} Z^h$$

Remark C.2.

1. Let h^1, \dots, h^k be k vectors in the program, let $\theta_1, \dots, \theta_p$ be p scalars in the program, and let $\phi : \mathbb{R}^{k+p} \rightarrow \mathbb{R}$ be a pseudo-Lipschitz function. Then applying the previous Lemma C.4 to $h := \phi(h^1, \dots, h^k; \theta_1, \dots, \theta_p)$ (which is in the program by `NonLin`), shows that for any α , $\phi(h_\alpha^1, \dots, h_\alpha^k; \theta_1, \dots, \theta_p)$ converges in law to $Z^h = \phi(Z^{h^1}, \dots, Z^{h^k}; \overset{\circ}{\theta}_1, \dots, \overset{\circ}{\theta}_p)$.
2. A stronger form of convergence can occur depending on the parameterization we look at and the context. Indeed, if for example Z^h turns out to be a constant, then we already get convergence in probability instead of in law. If in addition the convergence is “fast enough”, it can occur almost surely.

Proof. Let h be a vector in the program, and consider the corresponding random variable Z^h . All we need is to prove that for any $\alpha \in \mathbb{N}^*$ and any bounded 1-Lipschitz function ϕ , we have $\mathbb{E}[\phi(h_\alpha)] \rightarrow \mathbb{E}[\phi(Z^h)]$, as m goes to infinity. We first observe that the Master Theorem B.1 ensures the convergence

$$\frac{1}{m} \sum_{\beta=1}^m \phi(h_\beta) \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[\phi(Z^h)].$$

Secondly, for any m , the distribution of h_1, \dots, h_m is exchangeable by symmetry, so that we get

$$\mathbb{E}[\phi(h_\alpha)] = \mathbb{E} \left[\frac{1}{m} \sum_{\beta=1}^m \phi(h_\beta) \right] \xrightarrow[m \rightarrow \infty]{} \mathbb{E}[\phi(Z^h)],$$

where the convergence is obtained by dominated convergence, which concludes the proof. \square

D Proof of the triviality of IPs: Proposition 3.1

Proof. Fix a time $t \geq 0$ and an input $\xi \in \mathbb{R}^d$ for the whole proof. We first show that the coordinates of the (pre-)activations of any layer $l \geq 2$ converge to 0 almost surely at initialization. To that end, we prove that the corresponding Z 's are equal to 0. Then we show a similar result for the backward pass, and finally conclude the proof by an induction.

D.1 Proof at $t = 0$

D.1.1 First forward pass

Tensor program setup: We consider a Tensor Program as defined in

$$\begin{cases} \widehat{W}^{l+1} = U^{L+1}, \\ U^1 \xi_0, \dots, U^1 \xi_t, U^1 \xi, \\ v^1, \dots, v^L, \end{cases}$$

and the initial scalars

$$\begin{cases} \chi_0, \dots, \chi_t, \\ \omega := m^{-1/2}, \nu := m^{-1}, \tau := m^{-2}, \\ m^{-1}v^{L+1}; \end{cases}$$

and with initial weight matrices

$$\widehat{W}^2, \dots, \widehat{W}^L.$$

Recall that the \widehat{W}^l are defined in Equation (A.1) of Appendix A. Note that for any $m \in \mathbb{N}^*$ and $j \in [m]$, we have

$$\left(U_j^{L+1}, (U^1 \xi_0)_j, \dots, (U^1 \xi_t)_j, (U^1 \xi)_j, v_j^1, \dots, v_j^L \right) \sim \mathcal{N} \left(0, \begin{pmatrix} 1 & 0 & 0 \\ 0 & M & 0 \\ 0 & 0 & I_L \end{pmatrix} \right),$$

where $M := \text{Gram}(\xi_0, \dots, \xi_t, \xi) = (\xi_r^\top \xi_s)_{0 \leq r, s \leq t+1}$ and I_L is the identity matrix of size $L \times L$ where we have set $\xi_{t+1} := \xi$.

Convergence of the initial scalars: ω, ν, τ as well as $m^{-1}v^{L+1}$ all converge almost surely towards 0. For the χ_s we will show below in the proof that they all converge to constants almost surely, thereby meeting the requirements of the Tensor Program. It is important to note that there is no circular logic to prove the χ_s converge almost surely. Indeed, each time we apply the master theorem to prove the convergence of $f_s(\xi_s)$ to a constant almost surely and thus that of χ_s , we apply it to a restricted Tensor Program where only the scalars $(\chi_r)_{0 \leq r < s}$ appear (and there is no such scalar needed to prove the convergence of χ_0 as shown below) which will already have been proved to converge almost surely.

1st forward pass: We drop the dependency of the forward and backward passes on ξ for brevity. $h_0^1 = U^1 \xi + v^1$ is the sum of two initial vectors in the program and has iid Gaussian coordinates $\mathcal{N}(0, \|\xi\|^2 + 1)$. By definition, $\widehat{Z}^{h_0^1} = \widehat{Z}^{U^1 \xi} + \widehat{Z}^{v^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$ since the two Gaussians appearing in the sum are independent. By **NonLin**, we have that since $x_0^1 = \sigma(h_0^1)$, $Z^{x_0^1} = \sigma(Z^{h_0^1})$. Note that $\mathbb{E}[\sigma(Z^{h_0^1})^2] < \infty$ since $Z^{h_0^1}$ is Gaussian with finite variance and σ is pseudo-Lipschitz and thus polynomially bounded.

Since $L \geq 2$, we can write $h_0^2 = m^{-1/2} \widehat{W}^2 x_0^1 + m^{-1} v^2$ (otherwise there is no h_0^2 and we simply have $f_0(\xi_0) = m^{-1} (U^2)^\top x_0^1$), which implies by **NonLin** that $Z^{h_0^2} = \overset{\circ}{\omega} Z^{\widehat{W}^2 x_0^1} + \overset{\circ}{\nu} Z^{v^2}$ with $\overset{\circ}{\omega} = \overset{\circ}{\nu} = 0$ and

$$Z^{\widehat{W}^2 x_0^1} = \widehat{Z}^{\widehat{W}^2 x_0^1} + \dot{Z}^{\widehat{W}^2 x_0^1}.$$

$\dot{Z}^{\widehat{W}^2 x_0^1} = 0$ by Lemma C.2, and $\widehat{Z}^{\widehat{W}^2 x_0^1} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_0^1})^2])$ and $0 \leq \mathbb{E}[(Z^{x_0^1})^2] < \infty$. We thus have $Z^{h_0^2} = \overset{\circ}{\omega} \widehat{Z}^{\widehat{W}^2(0)x_0^1} = 0$. Similarly, we also get that $\overset{\circ}{\nu} Z^{v^2} = 0$. We then have by **ZNonLin** $Z^{x_0^2} = \sigma(Z^{h_0^2}) = \sigma(0) = 0$.

Let $l \in [2, L-1]$ and assume $Z^{h_0^l} = 0$. Then, $Z^{x_0^l} = \sigma(Z^{h_0^l}) = 0$, and since $h_0^{l+1} = \omega \widehat{W}^{l+1} x_0^l + \nu v^{l+1}$, by **ZNonLin**, $Z^{h_0^{l+1}} = \overset{\circ}{\omega} Z^{\widehat{W}^{l+1} x_0^l} + \overset{\circ}{\nu} Z^{v^{l+1}}$ where by **ZMatMul**,

$$Z^{\widehat{W}^{l+1} x_0^l} = \widehat{Z}^{\widehat{W}^{l+1} x_0^l} + \dot{Z}^{\widehat{W}^{l+1} x_0^l},$$

and $\dot{Z}\widehat{W}^{l+1}x_0^l = 0$ by Lemma C.2. By ZHat, $\widehat{Z}\widehat{W}^{l+1}x_0^l \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_0^l})^2])$, and since $\dot{\omega} = 0$, $\dot{\omega}\widehat{Z}\widehat{W}^{l+1}x_0^l = 0$. Similarly, $\dot{v}Z^{v^{l+1}} = 0$. Then, by ZNonLin $Zx_0^{l+1} = \sigma(Z^{h_0^{l+1}}) = \sigma(0) = 0$, which concludes the induction.

We thus have only to deal with the last layer $L + 1$ to finish the first forward pass. We have $f_0(\xi) = m^{-1}((U^{L+1}(0))^\top x_0^L + v^{L+1}) = (1/m) \sum_{i=1}^m U_i^{L+1} x_{0,i}^L + m^{-1}v^{L+1}$. Since U^{L+1} and x_0^L are vectors in the program, $(1/m) \sum_{i=1}^m U_i^{L+1} x_{0,i}^L$ is a scalar in the program by the **Moment** rule, and it therefore converges almost surely to $\mathbb{E}[Z^{U^{L+1}} Z^{x_0^L}]$ by the Master Theorem. Now because U^{L+1} is an initial vector in the program, by definition, $Z^{U^{L+1}} = \widehat{Z}^{U^{L+1}} \sim \mathcal{N}(0, 1)$ is independent of $Z^{x_0^L}$. We thus get $\mathbb{E}[Z^{U^{L+1}} Z^{x_0^L}] = \mathbb{E}[Z^{U^{L+1}}] \mathbb{E}[Z^{x_0^L}] = 0$. On the other hand, $m^{-1}v^{L+1}$ is an initial scalar in the program which converges to 0 almost surely, so that $f_0(\xi)$ converges almost surely to 0.

D.1.2 First backward pass

1st backward pass: We can apply the previous reasoning of the forward pass with ξ_0 instead of ξ and we get that $f_0(\xi_0) \rightarrow 0$ almost surely. Therefore, since $\chi_0 = \partial_2 \ell(y_0, f_0(\xi_0))$ and $\partial_2 \ell(y_0, \cdot)$ is continuous by assumption, $\chi_0 \rightarrow \partial_2 \ell(y_0, 0) =: \dot{\chi}_0$ almost surely. We have $dx_0^L = m^{-1}U^{L+1}$ which makes it a vector in the program by **NonLin**, and $Z^{dx_0^L} = \dot{v}Z^{U^{L+1}}$. Since $Z^{U^{L+1}} \sim \mathcal{N}(0, 1)$ has finite variance and $\dot{v} = 0$, we have $Z^{dx_0^L} = 0$. $dh_0^L = dx_0^L \odot \sigma'(h_0^L)$ implies by ZNonLin $Z^{dh_0^L} = Z^{dx_0^L} \sigma'(Z^{h_0^L}) = 0 \times \sigma'(0) = 0$.

One has:

$$Z^{mdx_0^{L-1}} = \dot{\omega}(\widehat{Z}(\widehat{W}^L)^\top(mdh_0^L)) + \dot{Z}(\widehat{W}^L)^\top(mdh_0^L),$$

where $mdh_0^L = U^{L+1} \odot \sigma'(h_0^L)$. By Lemma C.2, $\dot{Z}(\widehat{W}^L)^\top(mdh_0^L) = 0$ (essentially, \widehat{W}^L never appears in the computation of dh_0^L), and by ZHat, $\widehat{Z}(\widehat{W}^L)^\top(mdh_0^L) \sim \mathcal{N}(0, \mathbb{E}[(Z^{mdh_0^L})^2])$, and by independence of $Z^{U^{L+1}}$ and $Z^{h_0^L}$,

$$\mathbb{E}[(Z^{mdh_0^L})^2] = \mathbb{E}[(Z^{U^{L+1}})^2] \mathbb{E}[(\sigma'(Z^{h_0^L}))^2] = \sigma'(0)^2$$

which is finite. Since $\dot{\omega} = 0$ we get $Z^{mdx_0^{L-1}} = 0$. $dh_0^{L-1} = dx_0^{L-1} \odot \sigma'(h_0^{L-1})$ implies by ZNonLin $Z^{mdh_0^{L-1}} = Z^{mdx_0^{L-1}} \sigma'(Z^{h_0^{L-1}}) = 0 \times \sigma'(Z^{h_0^{L-1}}) = 0$.

Let $l \in [2, L]$ (which is non-empty since $L \geq 2$) and assume $Z^{mdx_0^l} = Z^{mdh_0^l} = 0$. $mdx_0^{l-1} = \omega(\widehat{W}^l)^\top(mdh_0^l)$ implies by ZMatMul

$$Z^{mdx_0^{l-1}} = \dot{\omega}(\widehat{Z}(\widehat{W}^l)^\top(mdh_0^l)) + \dot{Z}(\widehat{W}^l)^\top(mdh_0^l).$$

By Lemma C.2, $\dot{Z}(\widehat{W}^l)^\top(mdh_0^l) = 0$, and by ZHat, $\widehat{Z}(\widehat{W}^l)^\top(mdh_0^l) \sim \mathcal{N}(0, \mathbb{E}[(Z^{mdh_0^l})^2])$. By the assumption above, $\mathbb{E}[(Z^{mdh_0^l})^2] = 0$, and since $\dot{\omega} = 0$ we get $Z^{mdx_0^{l-1}} = 0$. $dh_0^{l-1} = dx_0^{l-1} \odot \sigma'(h_0^{l-1})$ implies by ZNonLin $Z^{mdh_0^{l-1}} = Z^{mdx_0^{l-1}} \sigma'(Z^{h_0^{l-1}}) = 0 \times \sigma'(Z^{h_0^{l-1}})$. $Z^{h_0^{l-1}}$ is not 0 if $l = 2$, but since it is Gaussian with finite variance, and σ' is pseudo-Lipschitz by assumption, $\sigma'(Z^{h_0^{l-1}})$ is finite almost surely, and $Z^{mdh_0^{l-1}} = 0$ almost surely, which concludes the induction.

D.2 Induction step

Induction: Since we proved the result of the theorem for $t = 0$ in the first forward pass, we might as well assume $t \geq 1$. Let $s \in [0, t - 1]$ be an integer. In all that follows, for

any $r \in [0, s]$, for $z \in \{h_r^l, x_r^l, dh_r^l, dx_r^l\}$, we use z to denote $z(\xi_r)$. We make the following induction hypothesis: for any $r \in [0, s]$

$$\begin{cases} Z^{h_r^1} = Z^{U^1 \xi_r + v^1} \sim \mathcal{N}(0, \|\xi_r\|^2 + 1) \\ Z^{h_r^l} = 0 \text{ almost surely, } l \in [2, L] \\ f_r(\xi_r), f_r(\xi) \rightarrow 0 \text{ almost surely} \\ \chi_r \rightarrow \overset{\circ}{\chi}_r := \partial_2 \ell(y_r, 0) \text{ almost surely} \\ Z^{mdx_r^l} = Z^{mdh_r^l} = 0 \text{ almost surely, } l \in [1, L-1], \\ Z^{mdx_r^L} = U^{L+1}. \end{cases}$$

The aim is then to prove the same claims for $r = s+1$. Let us first start with the expressions of $\Delta W^l(s+1)$ and $\Delta B^l(s+1)$. We will use Equation (A.3) and the fact that $c_l + 2 \geq 0$ if $l \in [2, L]$, and $c_l + 1 \geq 0$ for $l = 1$, and $l = L+1$. We have by Equations (A.3) and (A.6)

$$\begin{aligned} \Delta W^1(s+1) &= -\eta m^{-c_1} \sum_{r=0}^s \chi_r dh_r^1 \xi_r^\top, \\ \Delta W^l(s+1) &= -\eta m^{-(2+c_l)} \sum_{r=0}^s \chi_r dh_r^l (x_r^{l-1})^\top, \quad l \in [2, L], \\ \Delta W^{L+1}(s+1) &= -\eta m^{-(1+c_{L+1})} \sum_{r=0}^s \chi_r x_r^L / m, \end{aligned}$$

and by Equations (A.4) and (A.7)

$$\begin{aligned} \Delta B^1(s+1) &= -\eta m^{-c_1} \sum_{r=0}^s \chi_r dh_r^1, \\ \Delta B^l(s+1) &= -\eta m^{-(2+c_l)} \sum_{r=0}^s \chi_r dh_r^l, \quad l \in [2, L], \\ \Delta B^{L+1}(s+1) &= -\eta m^{-(1+c_{L+1})} \sum_{r=0}^s \chi_r / m. \end{aligned}$$

In the following, we use for $z \in \{h_{s+1}^l, x_{s+1}^l, dh_{s+1}^l, dx_{s+1}^l\}$, we use z to denote $z(\xi)$ (and not $z(\xi_{s+1})$ for now). Using that in the Naive-IP, $c_1 = c_{L+1} = -1$, and $c_l = -2$ for $l \in [2, L]$, we have

$$\begin{aligned} \Delta W^1(s+1)\xi + \Delta B^1(s+1) &= -\eta \sum_{r=0}^s (\xi_s^\top \xi + 1) \chi_r (mdh_r^1), \\ \Delta W^l(s+1)x_{s+1}^{l-1} + \Delta B^l(s+1) &= -\eta \sum_{r=0}^s \chi_r \frac{((x_r^{l-1})^\top x_{s+1}^{l-1}) + 1}{m} (mdh_r^l), \quad l \in [2, L], \\ (\Delta W^{L+1}(s+1))^\top x_{s+1}^L + \Delta B^{L+1}(s+1) &= -\eta \sum_{r=0}^s \chi_r \frac{(x_r^L)^\top x_{s+1}^L + 1}{m}. \end{aligned}$$

To prove the claims above for $r = s+1$, we will first induct from $l = 1$ to $l = L$ for the forward pass and then induct from $l = L$ to $l = 1$ for the backward pass.

D.2.1 Forward pass at step $s + 1$

Forward pass at step $(s + 1)$: $h_{s+1}^1 = U^1\xi + v^1 + \Delta W^1(s+1)\xi + \Delta b^1(s+1)$ and by **ZNonLin**

$$\begin{aligned} Z^{h_{s+1}^1} &= \widehat{Z}^{U^1\xi} - \eta \sum_{r=0}^s (\xi_s^\top \xi + 1) \overset{\circ}{\chi}_r \underbrace{Z^{mdh_r^1}}_{0 \text{ a.s.}} \\ &= \widehat{Z}^{U^1\xi} = Z^{h_0^1(\xi)} \quad \text{almost surely.} \end{aligned}$$

Note that the scalars $(\chi_r)_{0 \leq r \leq s}$ are now valid scalars in the program by the induction hypothesis which allows applying the Tensor Program rules with those scalars as well as the master theorem. This gives $Z^{h_{s+1}^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$, and we then have $Z^{x_{s+1}^1} = \sigma(\widehat{Z}^{h_0^1(\xi)}) = Z^{x_0^1(\xi)}$ for which we have already proven $\mathbb{E}[(Z^{x_0^1(\xi)})^2] < \infty$.

$$h_{s+1}^2 = \omega \widehat{W}^2 x_{s+1}^1 + \tau v^2 - \eta \sum_{r=0}^s \chi_r \frac{(x_r^1)^\top x_{s+1}^1 + 1}{m} (mdh_r^2).$$

Because x_{s+1}^1 is a vector in the program, by **ZMatMul**

$$Z^{\widehat{W}^2 x_{s+1}^1} = \widehat{Z}^{\widehat{W}^2 x_{s+1}^1} + \dot{Z}^{\widehat{W}^2 x_{s+1}^1},$$

and because $Z^{x_{s+1}^1} = \sigma(\widehat{Z}^{U^1\xi + v^1})$ is only a function of the initial vectors $U^1\xi$ and v^1 , and not of any vector computed used $(\widehat{W}^2)^\top$, $\dot{Z}^{\widehat{W}^2 x_{s+1}^1} = 0$ by **ZDot**, and $\widehat{Z}^{\widehat{W}^2 x_{s+1}^1} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_{s+1}^1})^2])$ is a Gaussian with finite variance by **ZHat**. $(x_r^1)^\top x_{s+1}^1/m$ is a valid scalar in the program by the moment rule, and by the Master theorem,

$$((x_r^1)^\top x_{s+1}^1 + 1)/m \xrightarrow[m \rightarrow \infty]{\text{a.s.}} \mathbb{E}[Z^{x_r^1} Z^{x_{s+1}^1}] = \mathbb{E}[\sigma(Z^{U^1\xi_r + v^1})\sigma(Z^{U^1\xi + v^1})],$$

and because $U^1\xi_r$, v^1 and $U^1\xi$ are initial vectors in the program, $(Z^{U^1\xi_r + v^1}, Z^{U^1\xi + v^1})$ is jointly Gaussian by definition with finite covariance matrix

$$\begin{pmatrix} \|\xi_r\|^2 + 1 & \xi_r^\top \xi + 1 \\ \xi_r^\top \xi + 1 & \|\xi\|^2 + 1 \end{pmatrix},$$

which ensures the expectation above is finite because σ is polynomially bounded since it is pseudo-Lipschitz. We thus have

$$\begin{aligned} Z^{h_{s+1}^2} &= 0 \times \underbrace{\widehat{Z}^{\widehat{W}^2 x_{s+1}^1}}_{< \infty} + 0 \times Z^{v^2} - \eta \sum_{r=0}^s \underbrace{\overset{\circ}{\chi}_r}_{< \infty} \underbrace{\mathbb{E}[Z^{x_r^1} Z^{x_{s+1}^1}]}_{< \infty} \underbrace{Z^{mdh_r^2}}_0 \\ Z^{h_{s+1}^2} &= 0. \end{aligned}$$

We then get $Z^{x_{s+1}^2} = \sigma(0) = 0$ and thus $\mathbb{E}[(Z^{x_{s+1}^2})^2] = 0$.

Let $l \in [2, L - 1]$ and assume $Z^{h_{s+1}^l} = 0$.

$$h_{s+1}^{l+1} = \omega \widehat{W}^{l+1} + \tau v^{l+1} + x_{s+1}^l - \eta \sum_{r=0}^s \chi_r \frac{(x_r^l)^\top x_{s+1}^l + 1}{m} (mdh_r^{l+1}).$$

Now, since x_{s+1}^l is a vector in the program, $((x_r^l)^\top x_{s+1}^l + 1)/m$ is a scalar in the program by the **Moment** operation, which converges almost surely, by the Master Theorem, to

$$\mathbb{E}[Z^{x_r^l} Z^{x_{s+1}^l}] = \mathbb{E}[\sigma(Z^{h_r^l})\sigma(Z^{h_{s+1}^l})] = \sigma(0)^2 = 0.$$

By ZNonLin,

$$Z^{h_{s+1}^{l+1}} = \overset{\circ}{\omega} Z^{\widehat{W}^{l+1} x_{s+1}^l} + \overset{\circ}{\tau} Z^{v^{l+1}} - \eta \sum_{r=0}^s \underbrace{\overset{\circ}{\chi}_r}_{< \infty} \underbrace{\mathbb{E}[Z^{x_r^l} Z^{x_{s+1}^l}]}_{< \infty} \underbrace{Z^{mdh_r^{l+1}}}_0.$$

On the other hand,

$$Z^{\widehat{W}^{l+1} x_{s+1}^l} = \widehat{Z}^{\widehat{W}^{l+1} x_{s+1}^l} + \dot{Z}^{\widehat{W}^{l+1} x_{s+1}^l},$$

and since $Z^{x_{s+1}^l} = \sigma(Z^{h_{s+1}^l}) = \sigma(0) = 0$ is a constant almost surely, the derivatives defining \dot{Z} are equal to 0 (its expression as a function of the previous \widehat{Z} is a constant because any \widehat{Z} gets multiplied by 0) so that $\dot{Z}^{\widehat{W}^{l+1} x_{s+1}^l} = 0$, and $\widehat{Z}^{\widehat{W}^{l+1} x_{s+1}^l} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_{s+1}^l})^2]) = 0$. With $\overset{\circ}{\omega}_{l+1} = 0$ and $\overset{\circ}{\tau} = 0$, this yields $Z^{h_{s+1}^{l+1}} = 0$, and therefore $\mathbb{E}[(Z^{x_{s+1}^{l+1}})^2] = \mathbb{E}[\sigma(Z^{h_{s+1}^{l+1}})^2] = \sigma(0)^2 = 0$.

We now deal with the last layer $l = L + 1$ in the forward pass.

$$f_{s+1}(\xi) = m^{-1} (U^{L+1})^\top x_{s+1}^L - \eta \sum_{r=0}^s \chi_r \frac{(x_r^L)^\top x_{s+1}^L + 1}{m}.$$

Since $U^{L+1}, x_{s+1}^L, x_r^L$ are vectors in the program, by the Master Theorem, we have:

$$m^{-1} (U^{L+1})^\top x_{s+1}^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_{s+1}^L}] = \sigma(0) \underbrace{\mathbb{E}[Z^{U^{L+1}}]}_0 = 0,$$

and

$$\frac{(x_r^L)^\top x_{s+1}^L + 1}{m} \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{x_r^L} Z^{x_{s+1}^L}] = \sigma(0)^2 = 0.$$

We thus get

$$\sum_{r=0}^s \chi_r \frac{(x_r^L)^\top x_{s+1}^L + 1}{m} \xrightarrow[m \rightarrow \infty]{a.s.} \sum_{r=0}^s \underbrace{\overset{\circ}{\chi}_r}_{< \infty} \times 0 = 0.$$

This shows that

$$f_{s+1}(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} 0.$$

Doing the exact same reasoning as above with ξ_{s+1} instead of ξ for $r = s + 1$ gives us the first 3 claims of the induction hypothesis for $r = s + 1$.

D.2.2 Backward pass at step $s + 1$

Backward pass at step $(s + 1)$: the fourth claim $\chi_{s+1} \rightarrow \overset{\circ}{\chi}_{s+1} = \partial_2 \ell(y_{s+1}, 0)$ is a consequence of the fact that $f_{s+1}(\xi_{s+1}) \rightarrow 0$ almost surely, combined with the facts that $\chi_{s+1} = \partial_2 \ell(y_{s+1}, f_{s+1}(\xi_{s+1}))$ and that $\partial_2 \ell(y_{s+1}, \cdot)$ is continuous by assumption. In all the rest of this proof, for $z \in \{h_{s+1}^l, x_{s+1}^l, dh_{s+1}^l, dx_{s+1}^l\}$ we now use z to denote $z(\xi_{s+1})$ and not $z(\xi)$ anymore.

$mdx_{s+1}^L = w^{L+1}(s+1) = U^{L+1} - \eta \sum_{r=0}^s \chi_r x_r^L$ yields by ZNonLin

$$\begin{aligned} Z^{mdx_{s+1}^L} &= Z^{U^{L+1}} - \eta \sum_{r=0}^s \overset{\circ}{\chi}_r \underbrace{Z^{x_r^L}}_0 \\ Z^{dx_{s+1}^L} &= Z^{U^{L+1}}. \end{aligned}$$

We thus have $Z^{dx_{s+1}^L} = \overset{\circ}{\tau} Z^{mdx_{s+1}^L} = 0$, and $Z^{dh_{s+1}^L} = Z^{dx_{s+1}^L} \sigma'(Z^{h_{s+1}^L}) = 0 \times \sigma'(0) = 0$ almost surely.

One has:

$$mdx_{s+1}^{L-1} = \omega(\widehat{W}^L)^\top (mdh_{s+1}^L) - \eta \sum_{r=0}^s \chi_r \frac{(mdh_r^L)^\top mdh_{s+1}^L}{m} x_r^{L-1},$$

so that

$$Z^{mdx_{s+1}^{L-1}} = \overset{\circ}{\omega} Z^{(\widehat{W}^L)^\top (mdh_{s+1}^L)} - \eta \sum_{r=0}^s \overset{\circ}{\chi}_r \mathbb{E}[Z^{mdh_r^L} Z^{mdh_{s+1}^L}] Z^{x_r^{L-1}}.$$

Now, we have $\mathbb{E}[Z^{mdh_r^L} Z^{mdh_{s+1}^L}] = \mathbb{E}[(Z^{U^{L+1}})^2] \sigma'(0)^2 = \sigma'(0)^2$ which is finite. On the other hand, because $Z^{mdh_{s+1}^L} = Z^{U^{L+1}}$ does not depend on $Z^{\widehat{W}^L}$, we get that $\dot{Z}^{(\widehat{W}^L)^\top (mdh_{s+1}^L)} = 0$ and $\widehat{Z}^{(\widehat{W}^L)^\top (mdh_{s+1}^L)} \sim \mathcal{N}(0, 1)$ so that $\overset{\circ}{\omega} Z^{(\widehat{W}^L)^\top (mdh_{s+1}^L)} = 0$. It follows that $Z^{mdx_{s+1}^{L-1}} = 0$, and since $Z^{mdh_{s+1}^{L-1}} = Z^{mdx_{s+1}^{L-1}} \sigma'(Z^{h_{s+1}^{L-1}})$ we also get $Z^{mdh_{s+1}^{L-1}} = 0$.

Let $l \in [2, L]$ and assume $Z^{mdx_{s+1}^l} = Z^{mdh_{s+1}^l} = 0$. Then

$$mdx_{s+1}^{l-1} = \omega(\widehat{W}^l)^\top (mdh_{s+1}^l) - \eta \sum_{r=0}^s \chi_r \frac{(mdh_r^l)^\top mdh_{s+1}^l}{m} x_r^{l-1}.$$

Since $(mdh_r^l)^\top$ and mdh_{s+1}^l are vectors in the program, $(mdh_r^l)^\top mdh_{s+1}^l / m$ is a scalar in the program which converges almost surely, by the Master Theorem, to $\mathbb{E}[Z^{mdh_r^l} Z^{mdh_{s+1}^l}] = 0$. On the other hand $\dot{Z}^{(\widehat{W}^l)^\top (mdh_{s+1}^l)} = 0$ because $Z^{mdh_{s+1}^l}$ is a constant (its expression in function of the previous \widehat{Z} is constant equal to 0), and $\widehat{Z}^{(\widehat{W}^l)^\top (mdh_{s+1}^l)} \sim \mathcal{N}(0, \mathbb{E}[(Z^{mdh_0^l})^2])$ is almost surely 0 because $\mathbb{E}[(Z^{mdh_0^l})^2] = 0$. By ZNonLin we have

$$\begin{aligned} Z^{dx_{s+1}^{l-1}} &= \underbrace{\overset{\circ}{\omega}}_0 \underbrace{\widehat{Z}^{(\widehat{W}^l)^\top (mdh_{s+1}^l)}}_0 - \eta \sum_{r=0}^s \underbrace{\overset{\circ}{\chi}_r}_{< \infty} \underbrace{\mathbb{E}[Z^{dh_r^l} Z^{dh_{s+1}^l}]}_0 \underbrace{Z^{x_r^{l-1}}}_0 \\ Z^{dx_{s+1}^{l-1}} &= 0. \end{aligned}$$

Finally, $Z^{dh_{s+1}^{l-1}} = Z^{dx_{s+1}^{l-1}} \sigma'(Z^{h_{s+1}^{l-1}})$ yields $Z^{dh_{s+1}^{l-1}} = 0$ because $Z^{h_{s+1}^{l-1}} = 0$. This proves the last claim of the induction hypothesis for $r = s + 1$ and thus concludes the induction and therefore the proof. \square

E Preliminaries on positively homogeneous functions

In this section we give a description of activation functions σ satisfying Assumption 3. The fact that σ is positively p -homogeneous translates as

$$\sigma(z) = \begin{cases} \alpha z^p & \text{if } z \geq 0 \\ \beta |z|^p & \text{if } z < 0. \end{cases}$$

Additionally, one has

$$\sigma'(z) = \begin{cases} \alpha p z^{p-1} & \text{if } z \geq 0 \\ -\beta p |z|^{p-1} & \text{if } z < 0, \end{cases}$$

so that σ' is positively $(p-1)$ -homogeneous with $\sigma'(0) = 0$. Since $p \geq 2$, both σ and σ' are continuous and σ' is differentiable everywhere except at 0 if $p = 2$. It is immediate to check that both σ and σ' are pseudo-Lipschitz and that σ , σ' and σ'' are also polynomially bounded functions. The non-negativity assumption on σ gives $\alpha \geq 0, \beta \geq 0$, the fact that σ is not identically 0 leads to $\alpha > 0$ or $\beta > 0$, and finally the fact that σ has faster growth on the positive part of the real line yields $\alpha > \beta \geq 0$. One notices that the faster growth assumption is stronger than the assumption that σ is not identically zero, and the latter could thus be gotten rid of. The conditions on α and β can thus simply be summarized as

$$\alpha > \beta \geq 0 \tag{E.1}$$

With these conditions, we have that $\sigma(z) > 0$ for $z > 0$, and $\sigma'(z)z \geq 0$ for $z \neq 0$, that is $\text{sign}(\sigma'(z)) = \text{sign}(z)$.

F Preliminaries for [Theorem 3.2](#) and [Theorem 4.1](#)

In all this section since we assume positive homogeneity of the activation function, we also consider parameterizations with no bias terms except at the first layer.

F.1 Tilde variables

Definition F.1 (Scaleless variables at initialization). Let $\xi \in \mathbb{R}^d$ be an input vector. Independently of any parameterization, we consider the following variables “*without scale*” at initialization :

$$\begin{cases} \tilde{h}_0^1(\xi) := U^1 \xi + v^1 \\ \tilde{x}_0^1(\xi) := \sigma(\tilde{h}_0^1(\xi)) \end{cases} \quad \forall l \in [2, L], \begin{cases} \tilde{h}_0^l(\xi) := \widehat{W}^l \tilde{x}_0^{l-1}(\xi) \\ \tilde{x}_0^l(\xi) := \sigma(\tilde{h}_0^l(\xi)) \end{cases}$$

and define $\tilde{f}_0(\xi) := (\widehat{W}^{L+1})^\top \tilde{x}_0^L$, as well as

$$\begin{cases} d\tilde{x}_0^L(\xi) := U^{L+1} \\ d\tilde{h}_0^L(\xi) := d\tilde{x}_0^L(\xi) \odot \sigma'(\tilde{h}_0^L(\xi)) \end{cases} \quad \forall l \in [L-1], \begin{cases} d\tilde{x}_0^l(\xi) := (\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}(\xi) \\ d\tilde{h}_0^l(\xi) := d\tilde{x}_0^l(\xi) \odot \sigma'(\tilde{h}_0^l(\xi)) \end{cases}$$

where the \widehat{W}^l are defined in Equation [\(A.1\)](#).

Remark F.1. The tilde variables are independent of the choice of parameterization because, independently of the parameterization, $\widehat{W}_{pq}^l = m^{-1/2} U_{pq}^l \sim \mathcal{N}(0, 1/m)$ for $l \in [2, L+1]$ and $\widehat{W}_{pq}^1 = U_{pq}^1 \sim \mathcal{N}(0, 1)$. Those variables essentially reproduce the computations that take place in the forward (without any bias terms except at the first layer) and backward passes of any ac-parameterization but the magnitudes (the multiplying scalars ω_l) have been set to 1, essentially removing the additional scales which lead to explosion or vanishing as $m \rightarrow \infty$. The tilde variables of the forward pass at initialization correspond to the NTK parameterization. However this is not the case for the backward pass as the backward pass of NTK vanishes at initialization whereas the corresponding tilde variables have positive (> 0) variance as shown in Lemma [F.1](#) below.

Lemma F.1 (Scaleless variables have positive and finite second moment). *Let $\xi \in \mathbb{R}^d$ be an input vector, and consider a non-linearity σ satisfying Assumption 2. Then, dropping the dependency of the tilde variables on ξ , one has that for any $l \in [1, L]$, and for any $z \in \{\tilde{h}_0^l, \tilde{x}_0^l, d\tilde{h}_0^l, d\tilde{x}_0^l\}$, the second moment is positive and finite: $0 < \mathbb{E}[(Z^z)^2] < \infty$. More precisely, one has:*

$$\begin{aligned}
Z^{\tilde{h}_0^1} &\sim \mathcal{N}(0, \|\xi\|^2 + 1), & 0 < \mathbb{E}[(Z^{\tilde{x}_0^1})^2] < \infty \\
Z^{\tilde{h}_0^l} &\sim \mathcal{N}(0, V_{h,l}^2), & 0 < V_{h,l}^2 := \mathbb{E}[(Z^{\tilde{x}_0^{l-1}})^2] < \infty, & l \in [2, L], \\
0 < \mathbb{E}[(Z^{\tilde{x}_0^l})^2] < \infty, & & & l \in [2, L], \\
\tilde{f}_0(\xi) &\xrightarrow[m \rightarrow \infty]{law} \mathcal{N}(0, V_f^2), & 0 < V_f^2 := \mathbb{E}[(Z^{\tilde{x}_0^L})^2] < \infty, \\
Z^{d\tilde{x}_0^L} &\sim \mathcal{N}(0, 1) \\
Z^{d\tilde{x}_0^l} &\sim \mathcal{N}(0, V_{dx,l}^2), & 0 < V_{dx,l}^2 := \mathbb{E}[(Z^{d\tilde{h}_0^{l+1}})^2] < \infty, & l \in [1, L-1], \\
0 < \mathbb{E}[(Z^{d\tilde{h}_0^l})^2] < \infty, & & & l \in [1, L].
\end{aligned}$$

Remark F.2. As shown in Appendix M, those expectations, as well as the means (first and second moment) are tractable with $\sigma = \text{ReLU}$ and have simple expressions (for the first forward and backward passes). As shown in Appendices M.3 and M.5, the recursive formulas for the variances of the forward and backward passes can be unrolled, and to avoid explosion or vanishing with the depth L , one must initialize the i.i.d. Gaussian entries with a standard deviation of $\sqrt{2}$ to preserve the norm of the input signal.

Proof. Let $\xi \in \mathbb{R}^d$ be an input vector. We omit the dependency of the forward and backward passes on ξ for simplicity. We first induct from $l = 1$ to $l = L$ for the forward pass and then from $l = L$ to $l = 1$ for the backward pass. $\tilde{h}_0^1 = U^1 \xi + v^1$ is the sum of two initial vectors in the program, which follows two independent Gaussian laws by definition: $Z^{U^1 \xi} \sim \mathcal{N}(0, \|\xi\|^2)$, and $Z^{v^1} \sim \mathcal{N}(0, 1)$ independently of $Z^{U^1 \xi}$. We thus have $Z^{\tilde{h}_0^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$, which shows its variance is finite and > 0 , and by Lemma C.1, $0 < \mathbb{E}[(Z^{\tilde{x}_0^1})^2] < \infty$ since $Z^{\tilde{x}_0^1} = \sigma(Z^{\tilde{h}_0^1})$.

Now let $l \in [1, L-1]$ and assume $Z^{\tilde{h}_0^l} \sim \mathcal{N}(0, V_{h,l}^2)$ with $0 < V_{h,l}^2 < \infty$, and $0 < \mathbb{E}[(Z^{\tilde{x}_0^l})^2] < \infty$. By ZMatMul, $Z^{\tilde{h}_0^{l+1}} = Z^{\widehat{W}^{l+1} x_0^l}$ which is equal to $\widehat{Z}^{\widehat{W}^{l+1} x_0^l}$ by Lemma C.2. now by definition, $\widehat{Z}^{\widehat{W}^{l+1} x_0^l} \sim \mathcal{N}(0, \mathbb{E}[(Z^{\tilde{x}_0^l})^2])$, and the variance is > 0 and finite by the induction hypothesis, so that $0 < \mathbb{E}[(Z^{\tilde{h}_0^{l+1}})^2] < \infty$. Now by Lemma C.1 again, since $Z^{\tilde{x}_0^{l+1}} = \sigma(Z^{\tilde{h}_0^{l+1}})$, we also get that $0 < \mathbb{E}[(Z^{\tilde{x}_0^{l+1}})^2] < \infty$ which concludes the induction for the first L layers of the forward pass.

$\tilde{f}_0(\xi) = (\widehat{W}^{L+1})^\top \tilde{x}_0^L$ and $\widehat{W}_j^{L+1} \sim \mathcal{N}(0, 1/m)$ for every m , and by the Master Theorem, since $\|\tilde{x}_0^L\|^2/m$ is a scalar in the program defined by the moment operation, it converges almost surely to $\mathbb{E}[(Z^{\tilde{x}_0^L})^2]$. Finally, since \tilde{x}_0^L is computed using only the \widehat{W}^l for $l \leq L$, \tilde{x}_0^L is independent of \widehat{W}^{L+1} . By Lemma C.3, $\tilde{f}_0(\xi)$ converges in law towards $\mathcal{N}(0, \mathbb{E}[(Z^{\tilde{x}_0^L})^2])$, and $0 < \mathbb{E}[(Z^{\tilde{x}_0^L})^2] < \infty$ by the previous induction.

$Z^{d\tilde{x}_0^L} = Z^{U^{L+1}}$ and since U^{L+1} is an initial vector in the program whose coordinates are iid following $\mathcal{N}(0, 1)$, we have by definition $Z^{U^{L+1}} \sim \mathcal{N}(0, 1)$. $Z^{d\tilde{h}_0^L} = Z^{d\tilde{x}_0^L} \sigma'(Z^{\tilde{h}_0^L}) = \widehat{Z}^{U^{L+1}} \sigma'(\widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}})$. Now by definition in ZHat, $\widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}$ is independent of $\widehat{Z}^{U^{L+1}}$ since U^{L+1}

is an initial vector in the program. This yields

$$\begin{aligned}\mathbb{E}[(Z^{d\tilde{h}_0^L})^2] &= \mathbb{E}\left[(\widehat{Z}^{U^{L+1}})^2\right] \mathbb{E}\left[\sigma'(Z^{\tilde{h}_0^L})^2\right] \\ &= 1 \times \mathbb{E}\left[\sigma'(Z^{\tilde{h}_0^L})^2\right].\end{aligned}$$

By assumption, σ' is pseudo-Lipschitz and thus polynomially bounded, and is not almost everywhere 0. By the induction above, $Z^{\tilde{h}_0^L} \sim \mathcal{N}(0, \mathbb{E}[(Z^{\tilde{x}_0^{L-1}})^2])$ with $0 < \mathbb{E}[(Z^{\tilde{x}_0^{L-1}})^2] < \infty$. By Lemma C.1 we thus have $0 < \mathbb{E}[\sigma'(Z^{\tilde{h}_0^L})^2] < \infty$, which shows $0 < \mathbb{E}[(Z^{d\tilde{h}_0^L})^2] < \infty$.

Now let $l \in [2, L]$ and assume $Z^{d\tilde{x}_0^l} \sim \mathcal{N}(0, V_{dx,l}^2)$ with $0 < V_{dx,l}^2 < \infty$, and assume $0 < \mathbb{E}[(Z^{d\tilde{h}_0^l})^2] < \infty$. $Z^{d\tilde{x}_0^{l-1}} = Z^{(\widehat{W}^l)^\top d\tilde{h}_0^l}$ and $Z^{(\widehat{W}^l)^\top d\tilde{h}_0^l} = \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l}$ by Lemma C.2. By definition, $\widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l} \sim \mathcal{N}(0, \mathbb{E}[(Z^{d\tilde{h}_0^l})^2])$, so that $\mathbb{E}[(Z^{d\tilde{x}_0^{l-1}})^2] = \mathbb{E}[(Z^{d\tilde{h}_0^l})^2]$ and thus $0 < \mathbb{E}[(Z^{d\tilde{x}_0^{l-1}})^2] < \infty$ by the induction hypothesis. We have

$$Z^{d\tilde{h}_0^{l-1}} = Z^{d\tilde{x}_0^{l-1}} \sigma'(Z^{\tilde{h}_0^{l-1}}) = \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l} \sigma'(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}})$$

if $l \geq 3$, and

$$Z^{d\tilde{h}_0^1} = Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}) = \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \sigma'(\widehat{Z}^{\widehat{W}^1 \xi_0 + v^1})$$

if $l = 2$. In any case, the random variable inside σ' is independent of the other variable in the product. We thus get

$$\mathbb{E}[(Z^{d\tilde{h}_0^{l-1}})^2] = \underbrace{\mathbb{E}\left[(Z^{d\tilde{x}_0^{l-1}})^2\right]}_{>0, <\infty} \underbrace{\mathbb{E}\left[\sigma'(Z^{\tilde{h}_0^{l-1}})^2\right]}_{>0, <\infty}$$

where the bounds on the second expectation are obtained using Lemma C.1. This concludes the induction for the backward pass and thus the proof. \square

F.2 Expression of the forward and backward passes of ac-parameterizations in function of the tilde variables with homogeneity

Lemma F.2 (Forward pass with homogeneity at $t = 0$). *Consider any ac-parameterization of an L -hidden layer neural network with a p -homogeneous activation function, and $p \geq 1$. Let $\xi \in \mathbb{R}^d$ be an input to the network. Then, omitting the dependency of the forward pass and the tilde variables on ξ , one has:*

$$h_0^l = \gamma_{f,l} \tilde{h}_0^l, \quad l \in [1, L], \quad (\text{F.1})$$

$$x_0^l = (\gamma_{f,l})^p \tilde{x}_0^l, \quad l \in [1, L], \quad (\text{F.2})$$

$$f_0(\xi) = \gamma_{f,L+1} \tilde{f}_0(\xi), \quad (\text{F.3})$$

where, for any $l \in [1, L+1]$

$$\gamma_{f,l} := \left(\prod_{k=1}^l \omega_k^{p^{l-k}} \right).$$

Remark F.3.

1. $(\gamma_{f,l})^p = \left(\prod_{k=1}^l \omega_k^{p^{l-k+1}} \right)$.

2. When $p = 1$, $\gamma_{f,l}$ and $(\gamma_{f,l})^p$ simply reduce to $\omega_l \dots \omega_1$.
3. For integrable parameterizations, for any $l \in [1, L + 1]$, $\gamma_{f,l} = m^{-\sum_{k=0}^{l-2} p^k/2}$. The latter term is 1 when $l = 1$, and otherwise $m^{-(l-1)/2}$ if $p = 1$ and $m^{-(p^{l-1}-1)/2(p-1)}$ if $p > 1$. For μP , $\gamma_{f,l} = 1$ for any $l \in [1, L]$ because $\omega_l = 1$ for μP if $l \in [1, L]$.
4. Instead of homogeneity, assume σ is differentiable, has non-zero derivative in 0 and $\sigma(0) = 0$. Also assume that $\omega_1 = 1$ (i.e., $a_1 = 0$) and $\omega_l \rightarrow 0$ (i.e., $a_l > 1/2$) for $l \in [2, L]$, which is the case in integrable parameterizations. Then, we have $h_0^1 = \tilde{h}_0^1$, and $h_0^2 = \omega_2 \tilde{h}_0^2$, so that $x_0^2 = \sigma(\omega_2 \tilde{h}_0^2)$ and as $m \rightarrow \infty$, $x_0^2 \simeq \omega_2 \sigma'(0) \tilde{h}_0^2$. Then similarly, we have for $h_0^3 \simeq \omega_3 \omega_2 \sigma'(0) \widehat{W}^3 \tilde{h}_0^2$ and $x_0^3 \simeq \omega_3 \omega_2 \sigma'(0)^2 \widehat{W}^3 \tilde{h}_0^2$. An easy induction then gives $h_0^l = \sigma'(0)^{l-2} (\omega_l \dots \omega_2) \widehat{W}^l \dots \widehat{W}^2 \tilde{h}_0^2$. This thus resembles the case of a $p = 1$ positively homogeneous function, except that the first forward pass is effectively linearized after layer 1, but the magnitude of the forward pass at different layers is also well understood in this case so that the learning rates for the first update can be chosen appropriately (e.g., for integrable parameterizations). In particular, the initial learning rates of IP-LLR for $p = 1$ will also produce non-trivial weight updates at $t = 0$ in this setting, which will in turn induce learning. Finally, setting the initial standard deviations of the weight matrices equal to $|\sigma'(0)|^{-1}$ instead of 1 for the intermediate layers avoids problems with the depth L .

Proof. $h_0^1 = m^{-a_1}$ implies that $h_0^1 = \omega_1 (U^1 \xi + v^1) = \omega_1 \tilde{h}_0^1$, which entails $x_0^1 = \omega_1^p \tilde{x}_0^1$ because σ is positively p -homogeneous and $\omega_1 \geq 0$. Now let $l \in [1, L - 1]$ and assume $h_0^l = (\prod_{k=1}^l \omega_k^{p^{l-k}}) \tilde{h}_0^l$, and $x_0^l = (\prod_{k=1}^l \omega_k^{p^{l-k+1}}) \tilde{x}_0^l$. Then

$$\begin{aligned}
h_0^{l+1} &= \omega_{l+1} \widehat{W}^{l+1}(0) x_0^l \\
&= \omega_{l+1} \left(\prod_{k=1}^l \omega_k^{p^{l-k+1}} \right) \widehat{W}^{l+1}(0) \tilde{x}_0^l \\
&= \left(\prod_{k=1}^{l+1} \omega_k^{p^{l+1-k}} \right) \tilde{h}_0^{l+1}
\end{aligned}$$

Since σ is positively homogeneous, we have

$$\begin{aligned}
x_0^{l+1} &= \sigma(h_0^{l+1}) \\
&= \left(\prod_{k=1}^{l+1} \omega_k^{p^{l+1-k}} \right)^p \sigma(\tilde{h}_0^{l+1}) \\
&= \left(\prod_{k=1}^{l+1} \omega_k^{p^{l+2-k}} \right) \sigma(\tilde{h}_0^{l+1})
\end{aligned}$$

This concludes the induction and gives the result for any $l \in [1, L]$. To conclude, we compute the expression of $f_0(\xi) = \omega_{L+1} (\widehat{W}^{L+1}(0))^\top x_0^L = \omega_{L+1} (\prod_{k=1}^L \omega_k^{p^{L-k+1}}) (\widehat{W}^{L+1}(0))^\top \tilde{x}_0^L = (\prod_{k=1}^{L+1} \omega_k^{p^{L+1-k}}) \tilde{f}_0(\xi)$. \square

Lemma F.3 (Backward pass with homogeneity at $t = 0$). *Consider any ac-parameterization of an L -hidden layer neural network with a positively p -homogeneous activation function, and $p \geq 1$. Let $\xi_0 \in \mathbb{R}^d$ be the first training input. Then, omitting the dependency of the forward*

and backward passes, as well as that of the tilde variables on ξ_0 , one has for any $l \in [1, L]$:

$$dx_0^l = m^{-a_{L+1}} \gamma_{b,l} \left(\prod_{k=l+1}^L \gamma_{f,k} \right)^{p-1} d\tilde{x}_0^l, \quad (\text{F.4})$$

$$dh_0^l = m^{-a_{L+1}} \gamma_{b,l} \left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} d\tilde{h}_0^l, \quad (\text{F.5})$$

where, for any $l \in [1, L]$,

$$\gamma_{b,l} = \prod_{k=l+1}^L \omega_k.$$

Remark F.4.

1. By swapping the products, one has that

$$\prod_{k=l+1}^L \gamma_{f,k} = \prod_{k=1}^L \omega_k^{\sum_{r=\max(k,l+1)}^L p^{r-k}}.$$

2. When $p = 1$, $\left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} = 1$ for any $l \in [1, L+1]$.

3. For integrable parameterizations, $\gamma_{b,l} = m^{-(L-l)/2}$ for any $l \in [1, L]$. For μP , $\gamma_{b,l} = 1$ for any $l \in [1, L]$.

4. For $l = L$, $\gamma_{b,L} = 1$, $\prod_{k=l+1}^L \gamma_{f,k} = 1$, $\prod_{k=l}^L \gamma_{f,k} = \gamma_{f,L}$.

Proof. $dx_0^L = W^L(0) = m^{-a_{L+1}} U^{L+1} = m^{-a_{L+1}} d\tilde{x}_0^L$,

$$\begin{aligned} dh_0^L &= dx_0^L \odot \sigma'(h_0^L) \\ &= m^{-a_{L+1}} d\tilde{x}_0^L \odot \sigma'(\gamma_{f,L} \tilde{h}_0^L) \\ &= m^{-a_{L+1}} (\gamma_{f,L})^{p-1} d\tilde{x}_0^L \odot \sigma'(\tilde{h}_0^L) \end{aligned}$$

where the second equality stems from Lemma F.2 and the last equality stems from $\omega_L \dots \omega_1 > 0$ and the positive $(p-1)$ -homogeneity of σ' . Let $l \in [2, L]$ and assume that dx_0^l satisfies Equation (F.4) and dh_0^l satisfies Equation (F.5). Then

$$\begin{aligned} dx_0^{l-1} &= \omega_l (\widehat{W}^l)^\top dh_0^l \\ &= m^{-a_{L+1}} \omega_l \gamma_{b,l} \left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} (\widehat{W}^l)^\top d\tilde{h}_0^l \\ &= m^{-a_{L+1}} \gamma_{b,l-1} \left(\prod_{k=(l-1)+1}^L \gamma_{f,k} \right)^{p-1} d\tilde{x}_0^{l-1}, \end{aligned}$$

and

$$\begin{aligned}
dh_0^{l-1} &= dx_0^l \odot \sigma'(h_0^l) \\
&= m^{-a_{L+1}} \gamma_{b,l-1} \left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} d\tilde{x}_0^{l-1} \odot \sigma'(\gamma_{f,l} \tilde{h}_0^{l-1}) \\
&= m^{-a_{L+1}} \gamma_{b,l-1} \left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} (\gamma_{f,l-1})^{p-1} d\tilde{x}_0^l \odot \sigma'(\tilde{h}_0^l) \\
&= m^{-a_{L+1}} \gamma_{b,l-1} \left(\prod_{k=l-1}^L \gamma_{f,k} \right)^{p-1} d\tilde{h}_0^{l-1},
\end{aligned}$$

where we have used Lemma F.2 in the second equality, the positive $(p-1)$ -homogeneity of σ' combined with $\omega_l \dots \omega_1 > 0$ in the third equality and the definition of $d\tilde{h}_0^{l-1}$ in the last. This thus concludes the proof by induction. \square

Lemma F.4 (Weight updates with homogeneity at $t = 0$). *Consider any ac-parameterization of an L -hidden layer neural network with a positively p -homogeneous activation function, and $p \geq 1$. Let $\xi_0 \in \mathbb{R}^d$ be the first training input. Then, omitting the dependency of the forward and backward passes, as well as that of the tilde variables on ξ_0 , one has:*

$$\begin{aligned}
\Delta W^1(1) &= -\eta \chi_0 m^{-(a_{L+1}+2a_1+c_1)} \omega_1^{p^L-1} \left(\prod_{k=2}^L \omega_k^{p^{L-k+1}} \right) d\tilde{h}_0^1 \xi_0^\top, \\
\Delta B^1(1) &= -\eta \chi_0 m^{-(a_{L+1}+2a_1+c_1)} \omega_1^{p^L-1} \left(\prod_{k=2}^L \omega_k^{p^{L-k+1}} \right) d\tilde{h}_0^1, \\
\Delta W^l(1) &= -\eta \chi_0 m^{-(a_{L+1}+2a_l+c_l-1)} \left(\prod_{k=1}^L \omega_k^{p^{L-k+1}} \right) \omega_l^{-1} \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m}, \quad l \in [2, L], \\
\Delta W^{L+1}(1) &= -\eta \chi_0 m^{-(2a_{L+1}+c_{L+1}-1)} \left(\prod_{k=1}^L \omega_k^{p^{L-k+1}} \right) \tilde{x}_0^L / m.
\end{aligned}$$

Remark F.5. For $p = 1$, we have

$$\begin{aligned}
\omega_1^{p^L-1} \left(\prod_{k=2}^L \omega_k^{p^{L-k+1}} \right) &= \omega_1 \dots \omega_L \\
\left(\prod_{k=1}^L \omega_k^{p^{L-k+1}} \right) \omega_l^{-1} &= \omega_1 \dots \omega_{l-1} \omega_{l+1} \dots \omega_L \\
\prod_{k=1}^L \omega_k^{p^{L-k+1}} &= \omega_1 \dots \omega_L
\end{aligned}$$

Proof. Before we begin with the proof, we start with a first basic result which will be used repeatedly in the proof. Let $N \in \mathbb{N}^*$. By Equation (A.3), we have

$$(p-1) \sum_{r=0}^N p^r = \sum_{r=1}^{N+1} p^r - \sum_{r=0}^N p^r = p^{N+1} - 1.$$

Now that this is established, let us look at the update for the first layer. We have

$$\begin{aligned}\Delta W^1(1) &= -\eta m^{-(2a_1+c_1)} \chi_0 d\tilde{h}_0^1 \xi_0^\top \\ &= -\eta m^{-(2a_1+c_1+a_{L+1})} \gamma_{b,1} \left(\prod_{k=1}^L \gamma_{f,k} \right)^{p-1} \chi_0 d\tilde{h}_0^1 \xi_0^\top,\end{aligned}$$

where we have used Lemmas F.2 and F.3 in the second equality. Now, we have

$$\gamma_{b,1} = \prod_{k=2}^L \omega_k,$$

and by the first point in Remark F.4, we have (with $l = 1$)

$$\begin{aligned}\left(\prod_{k=1}^L \gamma_{f,k} \right)^{p-1} &= \prod_{k=1}^L \omega_k^{(p-1) \sum_{r=k}^L p^{r-k}} \\ &= \prod_{k=1}^L \omega_k^{(p-1) \sum_{r=0}^{L-k} p^r} \\ &= \prod_{k=1}^L \omega_k^{p^{L-k+1} - 1}.\end{aligned}$$

It follows that

$$\gamma_{b,1} \left(\prod_{k=1}^L \gamma_{f,k} \right)^{p-1} = \omega_1^{p^L - 1} \prod_{k=2}^L \omega_k^{p^{L-k+1}}$$

The formula for $\Delta B^1(1)$ follows from the expression of dh_0^1 in function of $d\tilde{h}_0^1$ and from Equation (A.4).

Let $l \in [2, L]$

$$\begin{aligned}\Delta W^l(1) &= -\eta m^{-(2a_l+c_l)} \chi_0 d\tilde{h}_0^l (x_0^{l-1})^\top \\ &= -\eta m^{-(2a_l+c_l+a_{L+1}-1)} \chi_0 \gamma_{b,l} (\gamma_{f,l-1})^p \left(\prod_{k=l}^L \gamma_{f,k} \right)^{p-1} \frac{d\tilde{h}_0^l (x_0^{l-1})^\top}{m}\end{aligned}$$

Now, we have

$$\gamma_{b,l} = \prod_{k=l+1}^L \omega_k.$$

In addition, by the first point of Remark F.3, we have

$$(\gamma_{f,l-1})^p = \left(\prod_{k=1}^{l-1} \omega_k^{p^{l-k}} \right),$$

and by the first point in Remark F.4

$$\begin{aligned}
\left(\prod_{k=l}^L \gamma_{f,k}\right)^{p-1} &= \left(\prod_{k=1}^{l-1} \omega_k^{(p-1)\sum_{r=l}^L p^{r-k}}\right) \times \left(\prod_{k=l}^L \omega_k^{(p-1)\sum_{r=k}^L p^{r-k}}\right) \\
&= \left(\prod_{k=1}^{l-1} \omega_k^{(p-1)p^{l-k}\sum_{r=l}^L p^{r-l}}\right) \times \left(\prod_{k=l}^L \omega_k^{(p-1)\sum_{r=0}^{L-k} p^r}\right) \\
&= \left(\prod_{k=1}^{l-1} \omega_k^{(p-1)p^{l-k}\sum_{r=0}^{L-l} p^r}\right) \times \left(\prod_{k=l}^L \omega_k^{(p-1)\sum_{r=0}^{L-k} p^r}\right).
\end{aligned}$$

Let us now look, for each $k \in [1, L]$, at the power of ω_k which appears in the product $\gamma_{b,l}(\gamma_{f,l-1})^p \left(\prod_{k=l}^L \gamma_{f,k}\right)^{p-1}$. If $k \in [1, l-1]$, the exponent for ω_k is equal to

$$\begin{aligned}
p^{l-k} + (p-1)p^{l-k}\sum_{r=0}^{L-l} p^r &= p^{l-k} \left((p-1)\sum_{r=0}^{L-l} p^r + 1 \right) \\
&= p^{l-k} \left(p^{L-l+1} - 1 + 1 \right) \\
&= p^{L-k+1}.
\end{aligned}$$

If $k = l$, the exponent for ω_l is equal to

$$(p-1)\sum_{r=0}^{L-l} p^r = p^{L-l+1} - 1.$$

If $k \in [l+1, L]$, the exponent for ω_k is equal to

$$1 + (p-1)\sum_{r=0}^{L-k} p^r = 1 + p^{L-k+1} - 1 = p^{L-k+1}.$$

Thus, for every $k \neq l$, the exponent for ω_k is equal to p^{L-k+1} , and for $k = l$, the exponent for ω_l is equal to $p^{L-l+1} - 1$. It follows that

$$\gamma_{b,l}(\gamma_{f,l-1})^p \left(\prod_{k=l}^L \gamma_{f,k}\right)^{p-1} = \left(\prod_{k=1}^L \omega_k^{p^{L-k+1}}\right) \omega_l^{-1}.$$

Finally,

$$\begin{aligned}
\Delta W^{L+1}(1) &= -\eta m^{-(2a_{L+1}+c_{L+1})} \chi_0 x_0^L \\
&= -\eta m^{-(2a_{L+1}+c_{L+1}-1)} \chi_0 (\gamma_{f,L})^p \tilde{x}_0^L / m,
\end{aligned}$$

where we have used Lemma F.2 in the second equality. From the first point of Remark F.3, we get that

$$(\gamma_{f,L})^p = \prod_{k=1}^L \omega_k^{p^{L-k+1}},$$

which concludes the proof. □

Corollary F.4.1 (Weight updates of IP with homogeneity at $t = 0$). *Consider an integrable parameterization of an L -hidden layer neural network with no bias terms except at the first layer, and a positively p -homogeneous activation function, and $p \geq 1$. Let $\xi_0 \in \mathbb{R}^d$ be the first training input. Then, omitting the dependency of the forward and backward passes, as well as that of the tilde variables on ξ_0 , one has:*

$$\begin{aligned}\Delta W^1(1) &= -\eta\chi_0 m^{-(c_1-\gamma_1(p))} d\tilde{h}_0^1 \xi_0^\top, \\ \Delta B^1(1) &= -\eta\chi_0 m^{-(c_1-\gamma_1(p))} d\tilde{h}_0^1, \\ \Delta W^l(1) &= -\eta\chi_0 m^{-(c_l-\gamma_l(p))} \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m}, \quad l \in [2, L], \\ \Delta W^{L+1}(1) &= -\eta\chi_0 m^{-(c_{L+1}-\gamma_{L+1}(p))} \tilde{x}_0^L / m,\end{aligned}$$

where the $\gamma_l(p)$ are given in Definition 3.2.

Proof. For integrable parameterizations, $\omega_1 = 1$, $\omega_l = m^{-1/2}$ for $l \in [2, L]$, and $a_{L+1} = 1$. For the first layer, we have $a_{L+1} + 2a_1 + c_1 = 1$. On the other hand,

$$\begin{aligned}\omega_1^{p^{L-1}} \left(\prod_{k=2}^L \omega_k^{p^{L-k+1}} \right) &= \prod_{k=2}^L m^{-p^{L-k+1}/2} \\ &= m^{-\sum_{k=2}^L p^{L-k+1}/2} \\ &= m^{-\sum_{k=1}^{L-1} p^k/2} \\ &= m^{-1/2(\sum_{k=0}^{L-1} p^k - 1)},\end{aligned}$$

so that

$$\begin{aligned}m^{-(a_{L+1}+2a_1+c_1)} \omega_1^{p^{L-1}} \left(\prod_{k=2}^L \omega_k^{p^{L-k+1}} \right) &= m^{-c_1} m^{-1/2(\sum_{k=0}^{L-1} p^k - 1)} m^{-1} \\ &= m^{-c_1} m^{-1/2(\sum_{k=0}^{L-1} p^k + 1)} \\ &= m^{-c_1} m^{\gamma_1(p)},\end{aligned}$$

by Definition 3.2, which gives the result for the first layer's update ($\Delta W^1(1)$ and $\Delta B^1(1)$). Let $l \in [2, L]$. $a_{L+1} + 2a_l - 1 = 1 + 2 - 1 = 2$. On the other hand,

$$\begin{aligned}\left(\prod_{k=1}^L \omega_k^{p^{L-k+1}} \right) \omega_l^{-1} &= m^{-1/2(\sum_{k=0}^{L-1} p^k - 1)} m^{1/2} \\ &= m^{-1/2 \sum_{k=0}^{L-1} p^k + 1},\end{aligned}$$

so that

$$\begin{aligned}m^{-(a_{L+1}+2a_l+c_l)} \left(\prod_{k=1}^L \omega_k^{p^{L-k+1}} \right) \omega_l^{-1} &= m^{-c_l} m^{-1/2 \sum_{k=0}^{L-1} p^k + 1} m^{-2} \\ &= m^{-c_l} m^{-1/2 \sum_{k=0}^{L-1} p^k - 1} \\ &= m^{-c_l} m^{\gamma_l(p)},\end{aligned}$$

by Definition 3.2, which proves the result for the updates of the intermediate layers. Finally, we have $2a_{L+1} - 1 = 2 - 1 = 1$, and on the other hand, because $\omega_1 = 1$, as in the first update, we find

$$\prod_{k=1}^L \omega_k^{p^{L-k+1}} = m^{-1/2(\sum_{k=0}^{L-1} p^k - 1)},$$

so that

$$\begin{aligned}
m^{-(2a_{L+1}+c_{L+1}-1)} \prod_{k=1}^L \omega_k^{p^{L-k+1}} &= m^{-c_{L+1}} m^{-1/2(\sum_{k=0}^{L-1} p^k - 1)} m^{-1} \\
&= m^{-c_{L+1}} m^{-1/2(\sum_{k=0}^{L-1} p^k + 1)} \\
&= m^{-c_{L+1}} m^{\gamma_{L+1}(p)},
\end{aligned}$$

by Definition 3.2, which gives the result for the last layer's update and therefore concludes the proof. \square

Corollary F.4.2 (Weight updates of IP-LLR at $t = 0$). *Consider an IP-LLR parameterization of an L -hidden layer neural network with a p -homogeneous activation function, and $p \geq 1$. Let $\xi_0 \in \mathbb{R}^d$ be the first training input. Then, omitting the dependency of the forward and backward passes of IP-LLR, as well as that of the tilde variables on ξ_0 , one has:*

$$\begin{aligned}
\Delta W^1(1) &= -\eta \chi_0 d\tilde{h}_0^1 \xi_0^\top, \\
\Delta B^1(1) &= -\eta \chi_0 d\tilde{h}_0^1, \\
\Delta W^l(1) &= -\eta \chi_0 \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m}, \quad l \in [2, L], \\
\Delta W^{L+1}(1) &= -\eta \chi_0 \tilde{x}_0^L / m.
\end{aligned}$$

Proof. This is a simple consequence of Corollary F.4.1 and the fact that for IP-LLR $c_l = \gamma_l(p)$ at $t = 0$ by definition (see Definition 4.1) for any $l \in [1, L + 1]$. \square

Lemma F.5 (Weight updates of μP at $t = 0$). *Consider the μP parameterization given in Definition B.3 with a differentiable activation function σ . Let $\xi_0 \in \mathbb{R}^d$ be the first training input. Then, omitting the dependency of the forward and backward passes of μP , as well as that of the tilde variables on ξ_0 , one has:*

$$\begin{aligned}
\Delta W^1(1) &= -\eta \chi_0 d\tilde{h}_0^1 \xi_0^\top \\
\Delta B^1(1) &= -\eta \chi_0 d\tilde{h}_0^1 \\
\Delta W^l(1) &= -\eta \chi_0 \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m}, \quad l \in [2, L] \\
\Delta W^{L+1}(1) &= -\eta \chi_0 \tilde{x}_0^L / m
\end{aligned}$$

Remark F.6.

1. Although the formulas are identical with those for IP-LLR when the activation function is positively p -homogeneous, this **does not** mean that the weight updates are exactly equal. Indeed, although the tilde variables do not depend on the choice of parameterization and will thus be the same in μP as in IP-LLR, the variable χ_0 which appears in the formulas is parameterization-dependent as it depends on $f_0(\xi)$ which itself depends on the choice of parameterization.
2. There is no strong assumption on the activation function here (*e.g.*, homogeneity) as μP is designed to have such updates which induce feature learning at all layers.
3. Note that the coordinates of $\Delta W^l(1)$ are in $\Theta(m^{-1})$ whereas that of $W^l(0)$ are in $\Theta(m^{-1/2})$ for $l \in [2, L]$, so that paradoxically, even though μP is designed to produce “maximal updates” (in a certain sense), we have that $\Delta W_{jq}^l(1)/W_{jq}^l(0) = \Theta(m^{-1/2}) \rightarrow 0$ as $m \rightarrow \infty$: the relative displacement of the weights is zero in the infinite-width

limit. More generally, we have that for μP $(W_{jq}^l(t) - W_{jq}^l(0))/W_{jq}^l(0) \rightarrow 0$ as $m \rightarrow \infty$ if $t \geq 1$, which means that weights of the intermediate layers do not move away from their initialization in the infinite-width limit for μP , even if the (pre-)activations of every layer are maximally updated. This is in stark contrast with IP-LLR for which both $W^l(0)$ and $\Delta W^l(1)$ are in $\Theta(m^{-1})$ for the intermediate layers $l \in [2, L]$: the weights do move relatively to their initialization in the infinite-width limit.

Proof. μP is designed so that its forward pass has $h_0^l = \tilde{h}_0^l$ for any $l \in [1, L]$. Indeed, the choice of pre-factors for the weights with μP lead to the same recursive equations for the forward pass as the tilde variables, except for $f_0(\xi)$ which is equal to $m^{-1/2}\tilde{f}_0(\xi)$. For the backward pass, one has that for μP , $dx_0^L = W^{L+1}(0) = m^{-1}U^{L+1} = m^{-1}d\tilde{x}_0^L$. We then have

$$\begin{aligned} dh_0^L &= dx_0^L \odot \sigma'(h_0^L) \\ &= m^{-1}d\tilde{x}_0^L \odot \sigma'(\tilde{h}_0^L) \\ &= m^{-1}d\tilde{h}_0^L. \end{aligned}$$

Let $l \in [1, L-1]$, and assume that $dx_0^{l+1} = m^{-1}d\tilde{x}_0^{l+1}$ and $dh_0^{l+1} = m^{-1}d\tilde{h}_0^{l+1}$. Then, we have

$$\begin{aligned} dx_0^l &= (W^{l+1}(0))^\top dh_0^{l+1} \\ &= m^{-1}(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1} \\ &= m^{-1}d\tilde{x}_0^l. \end{aligned}$$

Similarly, we have

$$\begin{aligned} dh_0^l &= dx_0^l \odot \sigma'(h_0^l) \\ &= m^{-1}d\tilde{x}_0^l \odot \sigma'(\tilde{h}_0^l) \\ &= m^{-1}d\tilde{h}_0^l, \end{aligned}$$

which proves by induction that for any $l \in [1, L]$, $dx_0^l = m^{-1}d\tilde{x}_0^l$ and $dh_0^l = m^{-1}d\tilde{h}_0^l$ for μP . Recall that for μP , $a_1 = 0$, $a_l = 1/2$ for $l \in [2, L]$ and $a_{L+1} = 1$, and $c_l = -1$ for any $l \in [1, L+1]$. Now by Equations (A.3) and (A.4), the first weight updates give:

$$\begin{aligned} \Delta W^1(1) &= -\eta\chi_0 m^{-c_1} m^{-1} d\tilde{h}_0^1 \xi_0^\top \\ &= -\eta\chi_0 d\tilde{h}_0^1 \xi_0^\top, \end{aligned}$$

and

$$\begin{aligned} \Delta B^1(1) &= -\eta\chi_0 m^{-c_1} m^{-1} d\tilde{h}_0^1 \\ &= -\eta\chi_0 d\tilde{h}_0^1, \end{aligned}$$

For $l \in [2, L]$, we have

$$\begin{aligned} \Delta W^l(1) &= -\eta\chi_0 m^{-(1+c_l)} m^{-1} d\tilde{h}_0^l (x_0^{l-1})^\top \\ &= -\eta\chi_0 \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m}. \end{aligned}$$

Finally,

$$\begin{aligned} \Delta W^{L+1}(1) &= -\eta\chi_0 m^{-(2+c_{L+1})} x_0^L \\ &= -\eta\chi_0 \tilde{x}_0^L / m, \end{aligned}$$

which concludes the proof. \square

G Dynamics of the infinite-width limit of IP-LLR

Lemma G.1 (IP-LLR is zero at initialization). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, for any input vector $\xi \in \mathbb{R}^d$, one has that $\tilde{h}_0^l(\xi), \tilde{x}_0^l(\xi), d\tilde{x}_0^l, d\tilde{h}_0^l$ are vectors in the Tensor Program program for any $l \in [2, L]$, and additionally:*

$$\begin{aligned} f_0(\xi) &\xrightarrow[m \rightarrow \infty]{a.s.} 0 \\ \chi_0 &\xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{\chi}_0 := \partial_2 \ell(y_0, 0) \end{aligned}$$

Remark G.1. The result on the almost sure convergence of χ_0 ensures that the latter is a valid initial scalar in the Tensor Program defining the computations associated with the IP-LLR parameterization.

Proof. Because σ and σ' are pseudo-Lipschitz (since $p \geq 2$, see Appendix E), the tilde variables of the first forward and backward passes ($\tilde{h}_0^l, \tilde{x}_0^l, d\tilde{x}_0^l, d\tilde{x}_1^l$) are vectors in the program given Definition F.1 by the ZNonLin and ZMatMul rules. Additionally, by Lemma F.2,

$$\begin{aligned} f_0(\xi_0) &= m^{-\sum_{k=0}^{L-1} p^k/2} m^{-1/2} (U^{L+1})^\top \tilde{x}_0^L \\ &= m^{1/2} m^{-\sum_{k=0}^{L-1} p^k/2} m^{-1} (U^{L+1})^\top \tilde{x}_0^L \\ &= m^{-\sum_{k=1}^{L-1} p^k/2} m^{-1} (U^{L+1})^\top \tilde{x}_0^L \end{aligned}$$

Now, $m^{-1} (U^{L+1})^\top \tilde{x}_0^L \rightarrow \mathbb{E}[Z^{U^{L+1}} Z^{\tilde{x}_0^L}]$ almost surely by the master theorem, and $Z^{\tilde{x}_0^L} = \sigma(Z^{\widehat{W}^L \tilde{x}_0^{L-1}})$. By the Lemma C.2, $Z^{\widehat{W}^L \tilde{h}_0^{L-1}} = \widehat{Z}^{\widehat{W}^L \tilde{h}_0^{L-1}}$, and by the ZHat rule, the latter variable is independent of $Z^{U^{L+1}}$ since $Z^{U^{L+1}}$ is an initial vector in the program. This gives $\mathbb{E}[Z^{U^{L+1}} Z^{\tilde{x}_0^L}] = \mathbb{E}[Z^{U^{L+1}}] \mathbb{E}[Z^{\tilde{x}_0^L}] = 0 \times \mathbb{E}[Z^{\tilde{x}_0^L}]$. By Lemma F.1, and Lemma C.1, $\mathbb{E}[Z^{\tilde{x}_0^L}] < \infty$ because σ is polynomially bounded. We thus get $\mathbb{E}[Z^{U^{L+1}} Z^{\tilde{x}_0^L}] = 0$, and since $m^{-\sum_{k=1}^{L-1} p^k/2} \in (0, 1]$, $f_0(\xi_0) \rightarrow 0$ almost surely. Recall that by definition (see Appendix A) $\chi_0 = \partial_2 \ell(y_0, f_0(\xi_0))$. Since $f_0(\xi_0) \rightarrow 0$ almost surely, and since $\partial_2 \ell(y_0, \cdot)$ is continuous by assumption, we have that $\chi_0 \rightarrow \partial_2 \ell(y_0, 0) =: \overset{\circ}{\chi}_0$, which concludes the proof. \square

Definition G.1 (Tilde variables in the backward pass after initialization). For any ac-parameterization with $a_{L+1} = 1$, define for any $t \geq 1$,

$$\begin{aligned} d\tilde{x}_t^L &= m dx_t^L, \quad d\tilde{h}_t^L = d\tilde{x}_t^L \odot \sigma'(h_t^L), \\ d\tilde{x}_t^l &= (W^{l+1}(t))^\top d\tilde{h}_t^{l+1}, \quad l \in [1, L-1], \\ d\tilde{h}_t^l &= d\tilde{x}_t^l \odot \sigma'(h_t^l), \quad l \in [1, L-1]. \end{aligned}$$

Remark G.2.

1. One could in general define $d\tilde{x}_t^l$ to be equal to $m^{a_{L+1}} dx_t^l$ but since all the ac-parameterizations we study in this paper, *i.e.*, integrable parameterizations, μP , or hybrid versions thereof have $a_{L+1} = 1$, we limit the formulas to this case. The tilde variables are the right quantity to look at because of the term $m^{-a_{L+1}}$ which appears in the gradient *w.r.t* to x_t^L and then propagate to all the other variables of the backward pass by the equations of backpropagation.
2. Recall that in the definition above, it is implicitly assumed that the computations of the forward and backward passes at any time step s are done with the input $\xi = \xi_s$.

Lemma G.2 (Relationship between tilde and non-tilde variables). *For any ac-parameterization with $a_{L+1} = 1$, for any $t \geq 1$, and for any ξ , dropping the dependency of the forward and backward passes on ξ at time t , one has:*

$$\forall l \in [1, L+1], \quad dx_t^l = m^{-1}d\tilde{x}_t^l, \quad dh_t^l = m^{-1}d\tilde{h}_t^l.$$

Proof. $dx_t^L = m^{-1}d\tilde{x}_t^L$. $dh_t^L = dx_t^L \odot \sigma'(h_t^L) = m^{-1}d\tilde{x}_t^L \odot \sigma'(h_t^L) = m^{-1}d\tilde{h}_t^L$. Now let $l \in [2, L]$ and assume $dx_t^l = m^{-1}d\tilde{x}_t^l$, $dh_t^l = m^{-1}d\tilde{h}_t^l$. Then $dx_t^{l-1} = (W^l(t))^\top dh_t^l = m^{-1}(W^l(t))^\top d\tilde{h}_t^l = m^{-1}d\tilde{x}_t^{l-1}$, and $dh_t^{l-1} = dx_t^{l-1} \sigma(h_t^{l-1}) = m^{-1}d\tilde{x}_t^{l-1} \sigma(h_t^{l-1}) = m^{-1}d\tilde{h}_t^{l-1}$ which concludes the proof by induction. \square

Lemma G.3 (Weight updates for IP-LLR at any time step). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 1$, and let $t \geq 1$. Then, dropping the dependency of the forward and backward passes on ξ_t at time t , one has:*

$$\begin{aligned} \Delta W^{L+1}(t+1) &= -\eta \chi_t x_t^L / m, \\ \Delta W^l(t+1) &= -\eta \chi_t \frac{d\tilde{h}_t^l(x_t^{l-1})^\top}{m}, \quad l \in [2, L], \\ \Delta W^1(t+1) &= -\eta \chi_t d\tilde{h}_t^1 \xi_t^\top, \\ \Delta B^1(t+1) &= -\eta \chi_t d\tilde{h}_t^1. \end{aligned}$$

Proof. Using Equation (A.3), we have $\Delta W^{L+1}(t) = -\eta \chi_t m^{-(2a_{L+1}+c_{L+1})} x_t^L = -\eta \chi_t x_t^L / m$ because $2a_{L+1} + c_{L+1} = 2 - 1 = 1$ in IP-LLR since $t \geq 1$. For $l \in [2, L]$

$$\begin{aligned} \Delta W^l(t) &= -\eta \chi_t m^{-(2a_l+c_l)} dh_t^l(x^{l-1})^\top \\ &= -\eta \chi_t \frac{d\tilde{h}_t^l(x_t^{l-1})^\top}{m}, \end{aligned}$$

by Lemma G.2 and because $2a_l + c_l = 2 - 2 = 0$ for $t \geq 1$ in IP-LLR. $\Delta W^1(t) = -\eta \chi_t m^{-(2a_1+c_1)} dh_t^1 \xi_t^\top = -\eta \chi_t d\tilde{h}_t^1 \xi_t^\top$ by Lemma G.2 and because $2a_1 + c_1 = 0 - 1 = -1$ for $t \geq 1$ in IP-LLR. Finally, by Equation (A.4), we have $\Delta B^1(t) = -\eta \chi_t m^{-(2a_1+c_1)} dh_t^1 = -\eta \chi_t d\tilde{h}_t^1$ by Lemma G.2 and because $2a_1 + c_1 = -1$. \square

Theorem G.4 (Weights in IP-LLR at time t). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 1$. Then, for any $t \geq 1$, one has:*

- (i) $W^1(t) = U^1 - \eta \chi_0 d\tilde{h}_0^1 \xi_0^\top - \eta \left(\sum_{s=1}^{t-1} \chi_s d\tilde{h}_s^1 \xi_s^\top \right)$,
- (ii) $B^1(t) = v^1 - \eta \chi_0 d\tilde{h}_0^1 - \eta \left(\sum_{s=1}^{t-1} \chi_s d\tilde{h}_s^1 \right)$,
- (iii) $W^l(t) = \omega_l \widehat{W}^l - \eta \chi_0 \frac{d\tilde{h}_0^l(\tilde{x}_0^{l-1})^\top}{m} - \eta \left(\sum_{s=1}^{t-1} \chi_s \frac{d\tilde{h}_s^l(x_s^{l-1})^\top}{m} \right)$, $l \in [2, L]$,
- (iv) $W^{L+1}(t) = U^{L+1}/m - \eta \chi_0 \tilde{x}_0^L / m - \eta \left(\sum_{s=1}^{t-1} \chi_s x_s^L / m \right)$.

Proof. We have already seen the formulas are correct for $t = 1$ by Corollary F.4.2. Then, by Lemma G.3, an easy induction immediately yields the result. \square

Lemma G.5 (Backward pass of IP-LLR at time t). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 1$. Then, for any $t \geq 1$, dropping the dependency of the forward pass at time t on ξ_t , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

- (i) $d\tilde{x}_t^L = w^{L+1}(t) = U^{L+1} - \eta\chi_0\tilde{x}_0^L - \eta\sum_{s=1}^{t-1}\chi_s x_s^L$,
- (ii) $d\tilde{x}_t^{l-1} = \omega_l(\widehat{W}^l)^\top d\tilde{h}_t^l - \eta\chi_0\frac{(d\tilde{h}_0^l)^\top d\tilde{h}_t^l}{m}\tilde{x}_0^{l-1} - \eta\sum_{s=1}^{t-1}\chi_s\frac{(d\tilde{h}_s^l)^\top d\tilde{h}_t^l}{m}x_s^{l-1}$, $l \in [2, L]$.

Proof. By definition, we have

$$\begin{aligned} d\tilde{x}_t^L &= m dx_t^L \\ &= mW^{L+1}(t) \\ &= U^{L+1} - \eta\chi_0\tilde{x}_0^L - \eta\sum_{s=1}^{t-1}\chi_s x_s^L \end{aligned}$$

where the last equality stems from [Theorem G.4](#).

Let $l \in [2, L]$, we have:

$$\begin{aligned} d\tilde{x}_t^{l-1} &= (W^l(t))^\top d\tilde{h}_t^l \\ &= \omega_l(\widehat{W}^l)^\top d\tilde{h}_t^l - \eta\chi_0\frac{(d\tilde{h}_0^l)^\top d\tilde{h}_t^l}{m}\tilde{x}_0^{l-1} - \eta\sum_{s=1}^{t-1}\chi_s\frac{(d\tilde{h}_s^l)^\top d\tilde{h}_t^l}{m}x_s^{l-1} \end{aligned}$$

where the second equality stems from [Theorem G.4](#). \square

Lemma G.6 (Z for the forward pass of IP-LLR at time $t = 1$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Let $\xi \in \mathbb{R}^d$ be an input to the network. Then, for any $l \in [1, L]$, $h_1^l(\xi)$, $x_1^l(\xi)$, $d\tilde{x}_1^l$, $d\tilde{h}_1^l$ are vectors in the program, $f_1(\xi)$ is a scalar in the program, and χ_1 is a valid initial scalar in the program. Additionally, dropping the dependency of the forward pass at time $t = 1$ on ξ , and of the first forward and backward passes on ξ_0 , one has:*

- (i) $Z^{h_1^1} = Z^{W^1(1)\xi + B^1(1)} = Z^{U^1\xi} + Z^{v^1} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^1}$,
- (ii) $Z^{h_1^l} = Z^{W^l(1)x_1^{l-1}} = \overset{\circ}{\omega}_l Z^{\widehat{W}^l x_1^{l-1}} - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}] Z^{d\tilde{h}_0^l}$, $l \in [2, L]$,
- (iii) $f_1(\xi) = (W^{L+1}(1))^\top x_1^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_1^L}] - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}]$.

Proof. By [Theorem G.4](#), with $t = 1$, one has that $h_1^1 = U^1\xi + v^1 - \eta\chi_0(\xi_0^\top\xi + 1)d\tilde{h}_0^1$. By [Lemma G.1](#), $d\tilde{h}_0^1$ is a vector in the Tensor Program and χ_0 is a valid initial scalar in the program which has an almost sure limit $\overset{\circ}{\chi}_0 := \partial_2 \ell(y_0, 0)$ as $m \rightarrow \infty$. In addition, $U^1\xi$ and v^1 are initial vectors in the program, which thus shows that h_1^1 is a vector in the program by the **NonLin** operation. This also gives that $x_1^1 = \sigma(h_1^1)$ is a vector in the program since σ is pseudo-Lipschitz (see [Appendix E](#)). Moreover, by **ZNonLin**, we have $Z^{h_1^1} = Z^{U^1\xi} + Z^{v^1} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^1}$. Let $l \in [2, L]$ and assume that h_1^{l-1}, x_1^{l-1} are vectors in the program. Then, by [Theorem G.4](#) with $t = 1$, we get

$$h_1^l = \omega_l \widehat{W}^l x_1^{l-1} - \eta\chi_0 \frac{(\tilde{x}_0^{l-1})^\top x_1^{l-1}}{m} d\tilde{h}_0^l.$$

$(\tilde{x}_0^{l-1})^\top x_1^{l-1}/m$ is a scalar in the program by the **Moment** operation, and thus by the **MatMul** and **NonLin** operations, h_1^l is a vector in the program and thus so is $x_1^l = \sigma(h_1^l)$, which proves by induction that this is the case for any $l \in [2, L]$. By **ZNonLin** we thus have

$$Z^{h_1^l} = \overset{\circ}{\omega}_l Z^{\widehat{W}^l x_1^{l-1}} - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}] Z^{d\tilde{h}_0^l}.$$

We then have by [Theorem G.4](#) with $t = 1$,

$$f_1(\xi) = m^{-1}(U^{L+1})^\top x_1^L - \eta\chi_0 \frac{(\tilde{x}_0^L)^\top x_1^L}{m}$$

$U^{L+1} - \eta\chi_0\tilde{x}_0^L$ is a vector in the program by the **NonLin** operation, and the quantity $m^{-1}(U^{L+1} - \eta\chi_0\tilde{x}_0^L)^\top x_1^L$ is thus a scalar in the program by the **Moment** operation, and by the master theorem, we get $f_1(\xi) \rightarrow \mathbb{E}[Z^{U^{L+1}} Z^{x_1^L}] - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}]$ almost surely, since both expectations are finite by [Lemma L.1](#). Since we did the previous reasoning with an arbitrary ξ , we also get that $h_1^l(\xi_1), x_1^l(\xi_1)$ are vectors in the program for any $l \in [1, L]$ and that the formulas in (i), (ii), and (iii) hold when the input is ξ_1 . In particular, $f_1(\xi_1)$ converges to a finite almost sure limit $\overset{\circ}{f}_1(\xi_1)$, and thus the continuity of $\partial_2 \ell(y_1, \cdot)$ ensures the almost sure convergence of χ_1 towards $\overset{\circ}{\chi}_1 := \partial_2 \ell(y_1, \overset{\circ}{f}_1(\xi_1))$, which means χ_1 is a valid initial scalar in the Tensor Program. Then, dropping the dependency of the second forward pass (at $t = 1$) on ξ_1 , we get by [Lemma G.5](#) with $t = 1$:

$$d\tilde{x}_1^L = U^{L+1} - \eta\chi_0\tilde{x}_0^L$$

which is a vector in the program by **NonLin**. Then $d\tilde{h}_1^L = d\tilde{x}_1^L \odot \sigma'(h_1^L)$ is also a vector in the program since σ' is pseudo-Lipschitz. Let $l \in [2, L-1]$ and assume that $d\tilde{x}_1^{l+1}$ and $d\tilde{h}_1^{l+1}$ are vectors in the program. Then by [Lemma G.5](#) with $t = 1$, we have

$$d\tilde{x}_1^l = \omega_{l+1}(\widehat{W}^{l+1})^\top d\tilde{h}_1^{l+1} - \eta\chi_0 \frac{(d\tilde{h}_0^{l+1})^\top d\tilde{h}_1^{l+1}}{m} \tilde{x}_0^l$$

$(d\tilde{h}_0^{l+1})^\top d\tilde{h}_1^{l+1}/m$ is a scalar in the program by the **Moment** operation and by **MatMul** and **NonLin** we thus get that $d\tilde{x}_1^l$ is a vector in the program. Then $d\tilde{h}_1^l = d\tilde{x}_1^l \odot \sigma'(h_1^l)$ is also a vector in the program since σ' is pseudo-Lipschitz, which concludes the induction and with it the proof. \square

Theorem G.7 (Z for the forward pass of IP-LLR at time t). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Let $\xi \in \mathbb{R}^d$ be an input to the network. Then, for any $l \in [1, L]$, $h_s^l(\xi), x_s^l(\xi), d\tilde{x}_s^l, d\tilde{h}_s^l$ are vectors in the program, $f_s(\xi)$ is a scalar in the program, and χ_s is a valid initial scalar in the program. Additionally, dropping the dependency of the forward pass at time t on ξ , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

$$(i) \quad Z^{h_t^1} = Z^{W^1(t)\xi + B^1(t)} = Z^{U^1\xi + V^1} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^1} - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s(\xi_s^\top\xi + 1)Z^{d\tilde{h}_s^1} \right),$$

(ii) for any $l \in [2, L]$,

$$Z^{h_t^l} = Z^{W^l(t)x_t^{l-1}} = \overset{\circ}{\omega}_l Z^{\widehat{W}^l x_t^{l-1}} - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_0^l} - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_s^l} \right),$$

$$(iii) \quad f_t(\xi) = (W^{L+1}(t))^\top x_t^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_t^L}] - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_t^L}] - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_t^L}] \right).$$

Proof. We prove that the vectors and scalars in the claim of the theorem are part of the program by induction. Then the formulas of (i), (ii), and (iii) are a simple consequence of the **ZNonLin** operation. The case $t = 1$ has been treated in [Lemma G.6](#). Let $t \geq 1$ and assume that the vectors and scalars in the claim of the theorem are part of the program for any $s \in [1, t]$. By [Theorem G.4](#), one has that

$$\begin{aligned} h_{t+1}^1 &= W^1(t+1)\xi + B^1(t+1) \\ &= U^1\xi + V^1 - \eta\chi_0(\xi_0^\top\xi + 1)d\tilde{h}_0^1 - \eta \left(\sum_{s=1}^t \chi_s(\xi_s^\top\xi + 1)d\tilde{h}_s^1 \right) \end{aligned}$$

By the induction hypothesis and **NonLin**, we thus get that h_{t+1}^1 is a vector in the program and thus so is $x_{t+1}^1 = \sigma(h_{t+1}^1)$ since σ is polynomially bounded. Let $l \in [2, L]$ and assume that $h_{t+1}^{l-1}, x_{t+1}^{l-1}$ are vectors in the program. Then, by **Theorem G.4**, we get

$$h_{t+1}^l = \omega_l \widehat{W}^l x_{t+1}^{l-1} - \eta \chi_0 \frac{(\tilde{x}_0^{l-1})^\top x_{t+1}^{l-1}}{m} d\tilde{h}_0^l - \eta \left(\sum_{s=1}^t \chi_s \frac{(x_s^{l-1})^\top x_{t+1}^{l-1}}{m} d\tilde{h}_s^l \right).$$

For any $s \in [1, t]$, $(x_s^{l-1})^\top x_{t+1}^{l-1}/m$ and $(\tilde{x}_0^{l-1})^\top x_{t+1}^{l-1}/m$ are scalars in the program by the induction hypothesis and the **Moment** operation. Thus by the **MatMul** and **NonLin** operations, h_{t+1}^l is a vector in the program and thus so is $x_{t+1}^l = \sigma(h_{t+1}^l)$, which proves by induction that this is the case for any $l \in [2, L]$. We then have by **Theorem G.4**,

$$f_{t+1}(\xi) = m^{-1} \left(U^{L+1} - \eta \chi_0 \tilde{x}_0^L - \eta \sum_{s=1}^t \chi_s x_s^L \right)^\top x_{t+1}^L$$

$U^{L+1} - \eta \chi_0 \tilde{x}_0^L - \eta \sum_{s=1}^t \chi_s x_s^L$ is a vector in the program by the induction hypothesis and the **NonLin** operation. Then, by the **Moment** operation, $f_{t+1}(\xi)$ is a scalar in the program since x_{t+1}^L is also a vector in the program, and by the master theorem, we have

$$f_{t+1}(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_{t+1}^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_{t+1}^L}] - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_{t+1}^L}].$$

The limit is finite by Lemma **L.1** since by an easy induction any Z which appears is a polynomially bounded function of a Gaussian vector with finite covariance matrix. Since we did the previous reasoning with an arbitrary ξ , we also get that $h_{t+1}^l(\xi_{t+1}), x_{t+1}^l(\xi_{t+1})$ are vectors in the program for any $l \in [1, L]$. In particular, $f_{t+1}(\xi_{t+1})$ converges to an almost sure limit $\overset{\circ}{f}_{t+1}(\xi_{t+1})$, and thus the continuity of $\partial_2 \ell(y_{t+1}, \cdot)$ ensures the almost sure convergence of χ_{t+1} towards $\overset{\circ}{\chi}_{t+1} := \partial_2 \ell(y_{t+1}, \overset{\circ}{f}_{t+1}(\xi_{t+1}))$, which means χ_{t+1} is a valid initial scalar in the Tensor Program. Then, dropping the dependency of the forward pass at $t+1$ on ξ_{t+1} , we get by Lemma **G.5**:

$$d\tilde{x}_{t+1}^L = U^{L+1} - \eta \chi_0 \tilde{x}_0^L - \eta \sum_{s=1}^t \chi_s x_s^L$$

which is a vector in the program by **NonLin**. Then $d\tilde{h}_{t+1}^L = d\tilde{x}_{t+1}^L \odot \sigma'(h_{t+1}^L)$ is also a vector in the program since σ' is pseudo-Lipschitz. Let $l \in [2, L-1]$ and assume that $d\tilde{x}_{t+1}^{l+1}$ and $d\tilde{h}_{t+1}^{l+1}$ are vectors in the program. Then by Lemma **G.5**, we have

$$d\tilde{x}_{t+1}^l = \omega_{l+1} (\widehat{W}^{l+1})^\top d\tilde{h}_{t+1}^{l+1} - \eta \chi_0 \frac{(d\tilde{h}_0^{l+1})^\top d\tilde{h}_{t+1}^{l+1}}{m} \tilde{x}_0^l - \eta \sum_{s=1}^t \chi_s \frac{(d\tilde{h}_s^{l+1})^\top d\tilde{h}_{t+1}^{l+1}}{m} x_s^l$$

$(d\tilde{h}_s^{l+1})^\top d\tilde{h}_{t+1}^{l+1}/m$ is a scalar in the program for any $s \in [0, t]$ by the **Moment** operation and by **MatMul** and **NonLin** we thus get that $d\tilde{x}_{t+1}^l$ is a vector in the program. Then $d\tilde{h}_{t+1}^l = d\tilde{x}_{t+1}^l \odot \sigma'(h_{t+1}^l)$ is also a vector in the program since σ' is pseudo-Lipschitz, which concludes the induction. Then we get the claims of (i), (ii) and (iii) simply by applying the **ZNonLin** rule to the formulas derived above for the pre-activations h_{t+1}^l . \square

Corollary G.7.1 (Z for the forward pass of IP-LLR at time t). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, for any $t \geq 1$, and for any input $\xi \in \mathbb{R}^d$, dropping the dependency of the forward pass at time t on ξ , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

$$(i) \quad Z^{h_t^1} = Z^{W^1(t)\xi + B^1(t)} = Z^{U^1\xi} + Z^{v^1} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{h}_0^1} - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s (\xi_s^\top \xi + 1) Z^{d\tilde{h}_s^1} \right)$$

(ii) for any $l \in [2, L]$,

$$Z^{h_t^l} = Z^{W^l(t)x_t^{l-1}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_0^l} - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_s^l} \right),$$

$$(iii) \quad f_t(\xi) = (W^{L+1}(t))^\top x_t^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_t^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_t^L}] - \eta \left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_t^L}] \right).$$

Proof. The formulas are readily obtained by [Theorem G.7](#) coupled with the fact that we have $\overset{\circ}{\omega}_l Z^{\widehat{W}^l x_t^{l-1}} = 0$ for any $l \in [2, L]$, and $t \geq 1$, which stems from [Theorem L.9](#). \square

Remark G.3. Note that there is no circular logic here since only [Theorem G.7](#) is used to prove the results of [Appendix L.1](#) (and in particular [Theorem L.9](#)), so that using [Theorem L.9](#) for [Corollary G.7.1](#) does not lead to any issue.

Theorem G.8 (*Zs of backward pass of IP-LLR at time t*). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, for any $t \geq 1$, dropping the dependency of the forward pass at time t on ξ_t , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

$$(i) \quad Z^{d\tilde{x}_t^L} = Z^{w^{L+1}(t)} = Z^{U^{L+1}} - \eta \overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L} - \eta \sum_{s=1}^{t-1} \overset{\circ}{\chi}_s Z^{x_s^L},$$

$$(ii) \quad Z^{d\tilde{x}_t^{l-1}} = \overset{\circ}{\omega}_l Z^{(\widehat{W}^l)^\top d\tilde{h}_t^l} - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_t^l}] Z^{\tilde{x}_0^{l-1}} - \eta \sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{d\tilde{h}_s^l} Z^{d\tilde{h}_t^l}] Z^{x_s^{l-1}}, \quad l \in [2, L].$$

Proof. We have already proved in [Theorem G.7](#) that for any $s \in [1, t]$ the vectors of the forward (h_s^l, x_s^l for $l \in [1, L]$) and the backward pass ($d\tilde{x}_s^l, d\tilde{h}_s^l$ for $l \in [1, L]$) at time s are part of the program and similarly at $t = 0$ by [Lemma G.1](#). Then, claims (i) and (ii) readily follow from applying the `ZNonLin` rule to the formulas of [Lemma G.5](#). \square

Corollary G.8.1 (*Zs of backward pass of IP-LLR at time t*). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, for any $t \geq 1$, dropping the dependency of the forward pass at time t on ξ_t , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

$$(i) \quad Z^{d\tilde{x}_t^L} = Z^{w^{L+1}(t)} = Z^{U^{L+1}} - \eta \overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L} - \eta \sum_{s=1}^{t-1} \overset{\circ}{\chi}_s Z^{x_s^L}$$

$$(ii) \quad Z^{d\tilde{x}_t^{l-1}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_t^l}] Z^{\tilde{x}_0^{l-1}} - \eta \sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{d\tilde{h}_s^l} Z^{d\tilde{h}_t^l}] Z^{x_s^{l-1}}, \quad l \in [2, L].$$

Proof. The formulas are readily obtained by [Theorem G.8](#) and the fact that $Z^{(\widehat{W}^l)^\top d\tilde{h}_t^l} = 0$ for any $l \in [2, L]$ and $t \geq 1$, which stems from [Theorem L.9](#). \square

Remark G.4. Note that a similar statement can be made as in [Remark G.3](#) regarding circular logic since only [Theorem G.8](#) is used to prove the results of [Appendix L.1](#).

G.1 Second forward pass of IP-LLR ($t = 1$)

In this section, we prove that for IP-LLR, we have $0 < \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}] < \infty$ for any $l \in [1, L]$ under the assumption that $\overset{\circ}{\chi}_0 := \lim_{m \rightarrow \infty} \chi_0 \neq 0$. To obtain those results, we use the formulas from [Corollary G.7.1](#) for $t = 1$, which are obtained using the main result from [Appendix L](#), namely [Theorem L.9](#). We choose to put [Appendix L](#) towards the end of the [Appendix](#) section as its main result is quite intuitive: any multiplication by matrices with pre-factors in m^{-1} result in a vector whose coordinates (the corresponding Z) converge to 0 almost surely at any time step. The proof however requires a long and cumbersome induction and we thus leave it for the later stages of the [Appendix](#) so as not to break the narrative of the [Appendix](#).

The finiteness of the expectations $\mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_1^1}]$ is a simple consequence of Lemma L.1, but the fact that they are > 0 requires more work as we will see below. Since we work with IP-LLR, recall that we consider a bias term at the first layer only.

Lemma G.9 (1st layer of forward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with an activation function σ satisfying Assumption 3. Let ξ be an input to the network, and assume $\overset{\circ}{\chi}_0 \neq 0$. Then, dropping the dependency of the first forward-backward pass on ξ_0 , and that of the second forward pass on ξ , one has:*

$$(i) \quad Z^{h_1^1} = Z^{\tilde{h}_0^1(\xi)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{h}_0^1} = Z^{\tilde{h}_0^1(\xi)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}),$$

$$(ii) \quad (Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi)}, Z^{d\tilde{x}_0^1}) \sim \mathcal{N} \left(0, \begin{pmatrix} \|\xi_0\|^2 + 1 & \xi_0^\top \xi + 1 & 0 \\ \xi_0^\top \xi + 1 & \|\xi\|^2 + 1 & 0 \\ 0 & 0 & \mathbb{E}[(Z^{d\tilde{h}_0^2})^2] \end{pmatrix} \right),$$

$$(iii) \quad 0 < \mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_1^1}] < \infty.$$

Proof. We have by Corollary G.7.1 at time $t = 1$

$$\begin{aligned} Z^{x_1^1} &= \sigma \left(Z^{h_1^1(\xi)} \right) \\ &= \sigma \left(Z^{\tilde{h}_0^1(\xi)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{h}_0^1} \right). \end{aligned}$$

Moreover, since $d\tilde{h}_0^1 = d\tilde{x}_0^1 \odot \sigma'(Z^{\tilde{h}_0^1})$, since all the vectors are part of the Tensor Program, by ZNonLin we have $Z^{d\tilde{h}_0^1} = Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1})$, so that

$$Z^{x_1^1} = \sigma \left(Z^{\tilde{h}_0^1(\xi)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}) \right).$$

Finally, we have

$$Z^{\tilde{x}_0^1} = \sigma(Z^{\tilde{h}_0^1}).$$

From the rules of ZInit and ZHat, we have that

$$(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi)}, Z^{d\tilde{x}_0^1}) \sim \mathcal{N} \left(0, \begin{pmatrix} S & 0 \\ 0 & \mathbb{E}[(Z^{d\tilde{h}_0^2})^2] \end{pmatrix} \right),$$

with

$$S := \begin{pmatrix} \|\xi_0\|^2 + 1 & \xi_0^\top \xi + 1 \\ \xi_0^\top \xi + 1 & \|\xi\|^2 + 1 \end{pmatrix}.$$

By Lemma F.1, $\mathbb{E}[(Z^{d\tilde{h}_0^2})^2] < \infty$, so that the covariance matrix is finite and thus $Z^{\tilde{x}_0^1} Z^{x_1^1}$ is a polynomially bounded function of a Gaussian vector which shows that the expectation is finite by Lemma L.1. It is also non-negative since σ is non-negative. To prove that it is positive, one needs only prove that the integrand is not almost everywhere 0. By Lemma F.1, $\mathbb{E}[(Z^{d\tilde{h}_0^2})^2] > 0$ so that the covariance matrix is invertible if and only if S is invertible. We have

$$\det(S) = (\|\xi_0\|^2 \|\xi\|^2 - (\xi_0^\top \xi)^2) + \|\xi_0 - \xi\|^2,$$

which is the sum of two non-negative terms by Cauchy-Schwarz's inequality, and is thus 0 if and only if both terms are zero. The first term is zero only when ξ and ξ_0 are proportional,

and if in addition the second term is zero than $\xi = \xi_0$. The distribution of the Gaussian vector appearing in $Z^{\tilde{x}_0^1} Z^{x_1^1}$ thus depends on whether or not ξ_0 and ξ are equal.

Case when $\xi = \xi_0$. Then, calling $\lambda := -\eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1)$, we have

$$\mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_1^1}] = \int \sigma(z) \sigma(z - \lambda u \sigma'(z)) p_z(z) p_u(u) dz du,$$

where p_z and p_u are the densities of the two Gaussians $\mathcal{N}(0, \|\xi_0\|^2 + 1)$ and $\mathcal{N}(0, \mathbb{E}[(Z^{d\tilde{h}_0^2})^2])$ respectively, which are not degenerate, so that $p_z(z) > 0$ for any z and similarly for $p_u(u)$. Since $Z^{d\tilde{x}_0^1}$ and $-Z^{d\tilde{x}_0^1}$ have the same distribution and since it is independent of $Z^{\tilde{h}_0^1}$, we can assume $\lambda \geq 0$ W.L.O.G (if $\lambda \leq 0$ we can always do the change of variable $u \leftarrow -u$ in the integral above since $p_u(-u) = p_u(u)$). Consider the point $(z^*, u^*) := (1, -1)$, at which the integrand in the integral above is > 0 , because σ and σ' are > 0 on the positive part of the real line (see Appendix E) and $\lambda \geq 0$. The integral is then positive, because the integrand is a continuous function, since σ and σ' are continuous (see again Appendix E).

Case when $\xi \neq \xi_0$. Then, we have

$$\mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_1^1}] = \int \sigma(u) \sigma(v - \lambda z \sigma'(u)) p_{u,v}(u, v) p_z(z) dudvdz,$$

where $p_{u,v}$ and p_z are the densities of non-degenerate Gaussians and are thus well-defined and positive everywhere. Again, we can assume $\lambda \geq 0$ W.L.O.G. We consider the point $(u^*, v^*, z^*) = (1, 1, -1)$ at which the integrand is > 0 since σ and σ' are positive on the positive part of the real line. Hence, the integral is > 0 because the integrand is a continuous function, since σ and σ' are continuous, which concludes the proof. \square

Lemma G.10 (Intermediate layer of forward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with an activation function σ satisfying Assumption 3. Let ξ be an input to the network, let $l \in [2, L]$, and assume $\overset{\circ}{\chi}_0 \neq 0$. Then, dropping the dependency of the first forward-backward pass on ξ_0 , and that of the second forward pass on ξ , one has:*

- (i) $Z^{h_1^l} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}] Z^{d\tilde{h}_0^l} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}] \widehat{Z}^{d\tilde{x}_0^l} \sigma'(\widehat{Z}^{\tilde{h}_0^l})$,
- (ii) $Z^{\tilde{h}_0^l}$ and $Z^{d\tilde{x}_0^l}$ are independent,
- (iii) $0 < \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}] < \infty$.

Proof. We prove the result by induction on l , the case of $l = 1$ has already been dealt with in Lemma G.9. Let $l \in [1, L-1]$, and assume $0 < \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}] < \infty$. Calling $\lambda := -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}]$, we have $\lambda \neq 0$ by assumption and by the induction hypothesis. Then, by Corollary G.7.1 with $t = 1$, we have

$$Z^{h_1^{l+1}} = -\lambda Z^{d\tilde{h}_0^{l+1}}.$$

Moreover, $d\tilde{h}_0^{l+1} = d\tilde{x}_0^{l+1} \odot \sigma'(\tilde{h}_0^{l+1})$, and since all the vectors are part of the Tensor Program, we have by ZNonLin $Z^{d\tilde{h}_0^{l+1}} = Z^{d\tilde{x}_0^{l+1}} \sigma'(Z^{\tilde{h}_0^{l+1}})$. On the other hand, by Lemma C.2, we have $Z^{d\tilde{h}_0^{l+1}} = \widehat{Z}^{d\tilde{h}_0^{l+1}}$ and $Z^{\tilde{h}_0^{l+1}} = \widehat{Z}^{\tilde{h}_0^{l+1}}$, and finally by the ZHat rule, since $\tilde{h}_0^{l+1} = \widehat{W}^{l+1} \tilde{x}_1^l$ and $d\tilde{x}_0^{l+1} = U^{L+1}$ if $l = L - 1$ and $(\widehat{W}^{l+2})^\top d\tilde{h}^{l+2}$ otherwise, we get that $\widehat{Z}^{\tilde{h}_0^{l+1}}$ and $\widehat{Z}^{d\tilde{h}_0^{l+1}}$ are independent. In addition, we have

$$\mathbb{E}[Z^{\tilde{x}_0^{l+1}} Z^{x_1^{l+1}}] = \mathbb{E}[\sigma(Z^{\tilde{h}_0^{l+1}}) \sigma(-\lambda Z^{d\tilde{x}_0^{l+1}} \sigma'(Z^{\tilde{h}_0^{l+1}}))].$$

The expectation is non-negative because σ is and it is finite by Lemma L.1 because the integrand is a polynomially bounded function of the Gaussian vector $(Z^{\tilde{h}_0^{l+1}}, Z^{d\tilde{x}_0^{l+1}})$ (and thus of Z_0 , see Definition L.1). Using the positive p -homogeneity of σ and the fact that $\text{sign}(\sigma'(z)) = \text{sign}(z)$ (see Appendix E), and calling $\epsilon = \text{sign}(\lambda) \in \{-1, 1\}$, we have

$$\begin{aligned} \mathbb{E}[Z^{\tilde{x}_0^{l+1}} Z^{x_1^{l+1}}] &= \mathbb{E} \left[\mathbb{E} \left[Z^{\tilde{x}_0^{l+1}} Z^{x_1^{l+1}} \middle| Z^{\tilde{h}_0^{l+1}} \right] \right] \\ &= |\lambda|^p \mathbb{E} \left[|\sigma(Z^{\tilde{h}_0^{l+1}})| |\sigma'(Z^{\tilde{h}_0^{l+1}})|^p \mathbb{E} \left[\sigma(-\epsilon \text{sign}(Z^{\tilde{h}_0^{l+1}}) Z^{d\tilde{x}_0^{l+1}}) \middle| Z^{\tilde{h}_0^{l+1}} \right] \right], \end{aligned}$$

Now since $\epsilon \text{sign}(Z^{\tilde{h}_0^{l+1}}) \in \{-1, 1\}$, $Z^{d\tilde{x}_0^{l+1}}$ and $\epsilon \text{sign}(Z^{\tilde{h}_0^{l+1}}) Z^{d\tilde{x}_0^{l+1}}$ have the same distribution conditionally on $Z^{\tilde{h}_0^{l+1}}$, so that

$$\begin{aligned} \mathbb{E} \left[\sigma(-\epsilon \text{sign}(Z^{\tilde{h}_0^{l+1}}) Z^{d\tilde{x}_0^{l+1}}) \middle| Z^{\tilde{h}_0^{l+1}} \right] &= \mathbb{E} \left[\sigma(Z^{d\tilde{x}_0^{l+1}}) \middle| Z^{\tilde{h}_0^{l+1}} \right] \\ &= \mathbb{E} \left[\sigma(Z^{d\tilde{x}_0^{l+1}}) \right]. \end{aligned}$$

We thus get

$$\mathbb{E}[Z^{\tilde{x}_0^{l+1}} Z^{x_1^{l+1}}] = |\lambda|^p \mathbb{E} \left[|\sigma(Z^{\tilde{h}_0^{l+1}})| |\sigma'(Z^{\tilde{h}_0^{l+1}})|^p \right] \mathbb{E}[\sigma(Z^{d\tilde{x}_0^{l+1}})],$$

and both expectations are positive because they are non-negative and their integrands are > 0 on the positive part of the real line and the Gaussians involved have non-zero density on this subset of \mathbb{R} as they are not degenerate by Lemma F.1. This proves $\mathbb{E}[Z^{\tilde{x}_0^{l+1}} Z^{x_1^{l+1}}] > 0$ and concludes the proof by induction. \square

Lemma G.11 (Last layer of forward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with an activation function σ satisfying Assumption 3. Let ξ be an input to the network, and assume $\overset{\circ}{\chi}_0 \neq 0$. Then, dropping the dependency of the first forward-backward pass on ξ_0 , and that of the second forward pass on ξ , one has:*

- (i) $f_1(\xi) = (W^{L+1}(1))^\top x_1^L \xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{f}_1(\xi) := \mathbb{E}[Z^{U^{L+1}} Z^{x_1^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}]$,
- (ii) $Z^{U^{L+1}}$ and $Z^{\tilde{h}_0^L}$ are independent,
- (iii) $0 < \overset{\circ}{f}_1(\xi) < \infty$.

Proof. Claim (i) comes from Lemma G.6, in which we have already proved that the limit $\overset{\circ}{f}_1(\xi)$ is finite as a result of Lemma L.1 and the fact that the integrands are polynomially bounded functions of the Gaussian vector $(Z^{\tilde{h}_0^L}, Z^{U^{L+1}})$ which has finite (and diagonal as we will see shortly) covariance matrix. In addition, by Lemma C.2, we have $Z^{\tilde{h}_0^L} = \widehat{Z}^{\tilde{h}_0^L}$ and by definition in ZInit $Z^{U^{L+1}} = \widehat{Z}^{U^{L+1}}$. Finally, by the ZHat rule, the latter two random variables are independent since $\tilde{h}_0^L = \widehat{W}^L \tilde{x}_0^{L-1}$. Let $\epsilon := \text{sign}(\overset{\circ}{\chi}_0)$ and $\lambda_l := \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}]$ for $l \in \{L-1, L\}$. We have $\lambda_{L-1}, \lambda_L > 0$ by Lemma G.10, and using again the fact that $\text{sign}(\sigma'(z)) = \text{sign}(z)$ and the positive p -homogeneity of σ , we have

$$\begin{aligned} \mathbb{E}[Z^{U^{L+1}} Z^{x_1^L}] &= \mathbb{E} \left[\mathbb{E} \left[Z^{U^{L+1}} Z^{x_1^L} \middle| Z^{\tilde{h}_0^L} \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[Z^{U^{L+1}} \sigma(-\eta \overset{\circ}{\chi}_0 \lambda_{L-1} Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L})) \middle| Z^{\tilde{h}_0^L} \right] \right] \\ &= |\eta \lambda_{L-1} \overset{\circ}{\chi}_0|^p \mathbb{E} \left[|\sigma'(Z^{\tilde{h}_0^L})|^p \mathbb{E} \left[Z^{U^{L+1}} \sigma(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \right]. \end{aligned}$$

Since $Z^{U^{L+1}}$ and $-Z^{U^{L+1}}$ have the same distribution, and it is independent of $Z^{\tilde{h}_0^L}$, and since $\epsilon \text{sign}(Z^{\tilde{h}_0^L}) \in \{-1, 1\}$, we have

$$\begin{aligned} \mathbb{E} \left[-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}} \sigma(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] &= \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \\ &= \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \right], \end{aligned}$$

so that

$$\begin{aligned} \mathbb{E} \left[|\sigma'(Z^{\tilde{h}_0^L})|^p \mathbb{E} \left[Z^{U^{L+1}} \sigma(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \right] &= \\ -\epsilon \mathbb{E} \left[\text{sign}(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^p \right] \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \right]. \end{aligned}$$

We thus get

$$\mathbb{E}[Z^{U^{L+1}} Z^{x_1^L}] = -\epsilon |\eta \lambda_{L-1} \overset{\circ}{\chi}_0|^p \mathbb{E} \left[\text{sign}(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^p \right] \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \right]$$

We now prove that both expectations are positive. This is where the assumption that $\alpha > \beta$ (see Appendix E) appears to be crucial. We start with the first one. Since $Z^{\tilde{h}_0^L}$ has a zero-mean Gaussian distribution with positive variance (by Lemma F.1), its density p_z is positive everywhere and symmetric, and we have

$$\begin{aligned} \mathbb{E} \left[\text{sign}(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^p \right] &= \int_{z=0}^{+\infty} (\alpha p)^p z^{p(p-1)} p_z(z) dz + \int_{z=-\infty}^0 -(\beta p)^p (-z)^{p(p-1)} p_z(z) dz \\ &= (\alpha p)^p \int_{z=0}^{+\infty} z^{p(p-1)} p_z(z) dz - (\beta p)^p \int_{z=0}^{+\infty} z^{p(p-1)} dz \\ &= (\alpha^p - \beta^p) p^p \int_{z=0}^{+\infty} z^{p(p-1)} p_z(z) dz. \end{aligned}$$

The second equality stems from the change of variable $z \leftarrow -z$ in the second integral and from the symmetry of p_z with respect to $z = 0$. The last integral is > 0 because its integrand is > 0 on the corresponding domain, and $\alpha^p - \beta^p > 0$ since $\alpha > \beta$ by assumption and $p > 0$. For the second expectation, we get with a similar reasoning that

$$\begin{aligned} \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \right] &= \int_{u=0}^{+\infty} u \alpha u^p p_u(u) du + \int_{u=-\infty}^0 u \beta (-u)^p p_u(u) du \\ &= (\alpha - \beta) \int_{u=0}^{+\infty} u^{p+1} p_u(u) du, \end{aligned}$$

which shows the expectation is > 0 .

We now look at the second term in $\overset{\circ}{f}_1(\xi)$: $-\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}] = -\epsilon \eta |\overset{\circ}{\chi}_0| \lambda_L$. Summing this up with the first term, we get

$$\overset{\circ}{f}_1(\xi) = -\epsilon \underbrace{\left[|\eta \lambda_{L-1} \overset{\circ}{\chi}_0|^p \mathbb{E} \left[\text{sign}(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^p \right] \mathbb{E} \left[Z^{U^{L+1}} \sigma(Z^{U^{L+1}}) \right] + \eta |\overset{\circ}{\chi}_0| \lambda_L \right]}_{>0}$$

which concludes the proof. \square

Theorem G.12 (Non-trivial learning of IP-LLR at $t = 1$). *Consider an IP-LLR parameterization of an L -hidden layer neural network with an activation function σ satisfying*

Assumption 3. Let $\xi \in \mathbb{R}^d$ be an input to the network, and assume $\xi_0, \xi, \overset{\circ}{\chi}_0 \neq 0$. Then, one has:

$$(i) \quad f_0(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} 0$$

$$(ii) \quad f_1(\xi) \xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{f}_1(\xi) \neq 0$$

Proof. Claim (i) has already been proved in Lemma G.1, and claim (ii) has been proved in Lemma G.11 above. \square

Remark G.5. Note that since only quantities of the first ($t = 0$) forward and backward passes and second ($t = 1$) forward pass appear in Lemmas G.9, G.10, G.11, and Theorem G.12 we only need to assume we have an integrable parameterization with $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ at $t = 0$.

H Proof that no constant learning rate is possible: Theorem 3.2

In this section we prove the result of Theorem 3.2 by splitting the proof in two steps. First we show in Lemma H.1 that to have stable and non-vanishing updates for integrable parameterizations at $t = 1$, one must use the learning rate exponents $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ at $t = 0$. Then we show some preliminary results on the second backward pass (at $t = 1$) for integrable parameterizations when $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ at $t = 0$, and some other preliminary results on the third forward pass (at $t = 2$) when additionally one uses $c_1 = -1$, $c_l = -2$ for $l \in [2, L]$ and $c_{L+1} = -1$ at $t = 1$. Then we show in Lemma H.4, using those preliminary results, that assuming we have $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ at $t = 0$, to have stable and non-vanishing updates at $t = 2$ for integrable parameterizations, one must use the learning rate exponents $c_1 = -1$, $c_l = -2$ for $l \in [2, L]$ and $c_{L+1} = -1$ at $t = 1$.

H.1 Proof of the first implication for the learning rates at $t = 0$

Lemma H.1 (Learning rates for stable learning with IP at $t = 0$). *Consider an L -hidden layer fully-connected neural network with $L \geq 3$ in the integrable parameterization, and with no bias terms, except for the first layer. Assume that the activation function σ satisfies Assumption 3, and that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) \neq 0$. Assume further that $\xi_0^\top \xi_1 \neq 0$. Finally assume that Equation (3.1) holds:*

$$\begin{cases} \frac{1}{m} \|\Delta W^l(1) x_1^{l-1}\|^2 = \Theta(1), & l \in [1, L] \\ (\Delta W^{L+1}(1))^\top x_1^L = \Theta(1) \end{cases}$$

Then, one necessarily has that at $t = 0$, $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ (see Definition 3.2).

Proof. With the notations introduced in Appendix A, the assumptions on the limit of the loss terms at $t = 0$ imply $\overset{\circ}{\chi}_0 \neq 0$. Let us consider the updates at $t = 0$. By Corollary F.4.1, we have

$$\Delta W^1(1) \xi_1 = -m^{-(c_1 - \gamma_1(p))} \eta \chi_0 (\xi_0^\top \xi) d\tilde{h}_0^1,$$

so that

$$\frac{1}{m} \|\Delta W^1(1) \xi_1\|^2 = m^{-2(c_1 - \gamma_1(p))} [\eta \chi_0 (\xi_0^\top \xi)]^2 \frac{1}{m} \sum_{q=1}^m \left(d\tilde{h}_{0,q}^1 \right)^2.$$

From the master theorem, we get that $\sum_{q=1}^m (d\tilde{h}_{0,q}^1)^2/m$ converges almost surely towards $\mathbb{E}[(Z^{d\tilde{h}_0^1})^2]$ which is > 0 and finite by Lemma F.1. On the other hand, $[\eta\chi_0(\xi_0^\top \xi_1)]^2$ converges almost surely to $[\eta\chi_0^\circ(\xi_0^\top \xi_1)]^2$, which is > 0 by assumption, and finite.

If $c_1 > \gamma_1(p)$, then $c_1 - \gamma_1(p) > 0$, and $\|\Delta W^1(1)\xi_1\|^2/m \rightarrow 0$ almost surely, which is impossible since by assumption, almost surely, there exists $A > 0$ such that for large enough m , $A \leq \|\Delta W^1(1)\xi_1\|^2/m$.

If $c_1 < \gamma_1(p)$, then $c_1 - \gamma_1(p) < 0$, and $\|\Delta W^1(1)\xi_1\|^2/m \rightarrow \infty$ almost surely, which is impossible since by assumption, almost surely, there exists $B > 0$ such that for large enough m , $\|\Delta W^1(1)\xi_1\|^2/m \leq B$.

We thus have that $c_1 = \gamma_1(p)$. Let $l \in [1, L-1]$ and assume that $c_k = \gamma_k(p)$ for $k \in [1, l]$. Then by Lemmas G.9 and G.10, we have $0 < \mathbb{E}[Z^{\tilde{x}_0^k} Z^{x_1^k}] < \infty$ for any $k \in [1, l]$. We have

$$\frac{1}{m} \|\Delta W^{l+1}(1)x_1^l\|^2 = m^{-2(c_{l+1} - \gamma_{l+1}(p))} \left[\eta\chi_0 \frac{(\tilde{x}_0^l)^\top x_1^l}{m} \right]^2 \frac{1}{m} \sum_{q=1}^m (d\tilde{h}_{0,q}^{l+1})^2.$$

From the master theorem, we get that $\sum_{q=1}^m (d\tilde{h}_{0,q}^{l+1})^2/m$ converges almost surely towards $\mathbb{E}[(Z^{d\tilde{h}_0^{l+1}})^2]$ which is > 0 and finite by Lemma F.1. On the other hand, $[\eta\chi_0(\tilde{x}_0^l)^\top x_1^l/m]^2$ converges almost surely to $[\eta\chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}]]^2$, which is > 0 and finite.

If $c_{l+1} > \gamma_{l+1}(p)$, then $c_{l+1} - \gamma_{l+1}(p) > 0$, and $\|\Delta W^{l+1}(1)x_1^l\|^2/m \rightarrow 0$ almost surely, which is impossible since by assumption, almost surely, there exists $A > 0$ such that for large enough m , $A \leq \|\Delta W^{l+1}(1)x_1^l\|^2/m$.

If $c_{l+1} < \gamma_{l+1}(p)$, then $c_{l+1} - \gamma_{l+1}(p) < 0$, and $\|\Delta W^{l+1}(1)x_1^l\|^2/m \rightarrow \infty$ almost surely, which is impossible since by assumption, almost surely, there exists $B > 0$ such that for large enough m , $\|\Delta W^{l+1}(1)x_1^l\|^2/m \leq B$.

Therefore, we have $c_{l+1} = \gamma_{l+1}(p)$. By induction, we thus get that $c_l = \gamma_l(p)$ for any $l \in [1, L]$, which means in particular that $0 < \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}] < \infty$ by Lemma G.10. Finally, we have

$$(\Delta W^{L+1}(1))^\top x_1^L = -m^{-(c_{L+1} - \gamma_{L+1}(p))} \eta\chi_0 \frac{(\tilde{x}_0^L)^\top x_1^L}{m}.$$

The term $\eta\chi_0(\tilde{x}_0^L)^\top x_1^L/m$ converges almost surely towards $\eta\chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_1^L}]$, whose absolute value is > 0 and finite. Therefore, if $c_{L+1} > \gamma_{L+1}(p)$ then $c_{L+1} - \gamma_{L+1}(p) > 0$ so that $(\Delta W^{L+1}(1))^\top x_1^L \rightarrow 0$ almost surely, which is impossible since by assumption, almost surely, there exists $A > 0$ such that for large enough m , $A \leq |(\Delta W^{L+1}(1))^\top x_1^L|$. If $c_{L+1} < \gamma_{L+1}(p)$ then $c_{L+1} - \gamma_{L+1}(p) < 0$ so that $(\Delta W^{L+1}(1))^\top x_1^L \rightarrow \infty$ almost surely, which is impossible since by assumption, almost surely, there exists $B > 0$ such that for large enough m , $|(\Delta W^{L+1}(1))^\top x_1^L| \leq B$. Thus, we must have $c_{L+1} = \gamma_{L+1}(p)$, which concludes the proof for the first part. \square

H.2 Preliminaries on the second backward pass ($t = 1$)

Before we move on to the proof of the second part of the claim of Theorem 3.2, we stop and prove some preliminary results on the second backward pass (at $t = 1$) which will come in

handy later on. Similarly to what we did for $\mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_1^l}]$, we wish to prove that the quantity $0 < \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] < \infty$ for any $l \in [2, L]$.

Lemma H.2 (Backward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization of an L hidden-layer network, and assume that the activation function σ satisfies Assumption 3, and that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) \neq 0$. Then, one has that for any $l \in [2, L]$,*

$$0 < \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] < \infty$$

Remark H.1. Note that since only quantities of the first ($t = 0$) and second ($t = 1$) forward and backward passes appear, we only need to assume we have an integrable parameterization with $c_l = \gamma_l(p)$ for any $l \in [1, L + 1]$ at $t = 0$.

Proof. We start with $l = L$, and then induct over l from $l = L$ to $l = 2$, and we recall that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) =: \dot{\chi}_0$ by definition (see Appendix A), which is thus $\neq 0$ by assumption.

The case $l = L$. By Corollary G.8.1, we have $Z^{d\tilde{h}_0^L} = Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L})$ and $Z^{d\tilde{h}_1^L} = (Z^{U^{L+1}} - \eta \dot{\chi}_0 \sigma(Z^{\tilde{h}_0^L})) \sigma'(Z^{h_1^L})$. We thus have

$$\mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}] = \underbrace{\mathbb{E}[(Z^{U^{L+1}})^2 \sigma'(Z^{\tilde{h}_0^L}) \sigma'(Z^{h_1^L})]}_{:=A} + \eta |\dot{\chi}_0| \underbrace{\mathbb{E}[-\epsilon Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L}) \sigma(Z^{\tilde{h}_0^L}) \sigma'(Z^{h_1^L})]}_{:=B},$$

with $\epsilon := \text{sign}(\dot{\chi}_0)$ and we deal with both terms separately. First, by Corollary G.7.1 we re-write $Z^{h_1^L}$ as

$$Z^{h_1^L} = -\eta |\dot{\chi}_0| \epsilon \lambda Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L}),$$

where $\lambda := \mathbb{E}[Z^{\tilde{x}_0^{L-1}} Z^{x_1^{L-1}}] > 0$ by Lemma G.10. Using the fact that $\text{sign}(\sigma'(z)) = \text{sign}(z)$ and the positive $(p - 1)$ -homogeneity of σ' , we have

$$\sigma'(Z^{h_1^L}) = (\eta |\dot{\chi}_0| \lambda)^{p-1} |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}).$$

The first term in $\mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}]$ is thus equal to

$$\begin{aligned} A &= (\eta |\dot{\chi}_0| \lambda)^{p-1} \mathbb{E} \left[\mathbb{E} \left[(Z^{U^{L+1}})^2 \sigma'(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \right] \\ &= (\eta |\dot{\chi}_0| \lambda)^{p-1} \mathbb{E} \left[\sigma'(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \mathbb{E} \left[(Z^{U^{L+1}})^2 \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \right] \\ &= (\eta |\dot{\chi}_0| \lambda)^{p-1} \mathbb{E} \left[\sigma'(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \right] \mathbb{E} \left[(Z^{U^{L+1}})^2 \sigma'(Z^{U^{L+1}}) \right]. \end{aligned}$$

The third equality stems from the fact that $-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}$ and $Z^{U^{L+1}}$ have the same distribution conditionally on $Z^{\tilde{h}_0^L}$, and from the fact that $(Z^{U^{L+1}})^2 = (-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}})^2$. We now show that both expectations are > 0 . Calling p_z the density of the Gaussian $Z^{\tilde{h}_0^L}$ which is symmetric and positive everywhere since $Z^{\tilde{h}_0^L}$ is not degenerate, the first term is equal to

$$\begin{aligned} \mathbb{E} \left[\sigma'(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \right] &= (\alpha p)^p \int_{z=0}^{+\infty} z^{p(p-1)} p_z(z) dz - (\beta p)^p \int_{z=-\infty}^0 (-z)^{p(p-1)} p_z(z) dz \\ &= (\alpha^p - \beta^p) p^p \int_{z=0}^{+\infty} z^{p(p-1)} p_z(z) dz, \end{aligned}$$

where we have used the change of variable $z \leftarrow -z$ in the second equality, and the last quantity is > 0 since $\alpha > \beta$. With similar calculations, we get

$$\mathbb{E} \left[(Z^{U^{L+1}})^2 \sigma'(Z^{U^{L+1}}) \right] = (\alpha - \beta) p \int_{u=0}^{+\infty} u^{p+1} p_u(u) du > 0,$$

where p_u is the density of the standard Gaussian $Z^{U^{L+1}}$. This thus shows that $A > 0$.

We now turn to the second term B . We have:

$$\begin{aligned} B &= (\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} \times \mathbb{E} \left[\sigma'(Z^{\tilde{h}_0^L}) \sigma(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \text{sign}(Z^{\tilde{h}_0^L}) \times \right. \\ &\quad \left. \mathbb{E} \left[(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^L}) Z^{U^{L+1}}) \middle| Z^{\tilde{h}_0^L} \right] \right] \\ &= (\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} \times \mathbb{E} \left[\sigma'(Z^{\tilde{h}_0^L}) \sigma(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^{p-1} \text{sign}(Z^{\tilde{h}_0^L}) \right] \mathbb{E} \left[Z^{U^{L+1}} \sigma'(Z^{U^{L+1}}) \right] \\ &= (\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} \times \mathbb{E} \left[\sigma(Z^{\tilde{h}_0^L}) |\sigma'(Z^{\tilde{h}_0^L})|^p \right] \mathbb{E} \left[Z^{U^{L+1}} \sigma'(Z^{U^{L+1}}) \right]. \end{aligned}$$

We now prove again that both expectations are > 0 . The first integrand is non-negative everywhere and positive on the positive part of the real line where the Gaussian $Z^{\tilde{h}_0^L}$ has non-zero density, which shows the first expectation is > 0 . The same argument holds for the second expectation since $Z^{U^{L+1}}$ and $\sigma'(Z^{U^{L+1}})$ are of the same sign, which also leads to a positive expectation, which finally gives $B > 0$, thereby concluding the proof.

The case $l \in [2, L-1]$.

Let $l \in [2, L-1]$ and assume $0 < \nu := \mathbb{E}[Z^{d\tilde{h}_0^{l+1}} Z^{d\tilde{h}_1^{l+1}}] < \infty$. Calling $\epsilon := \text{sign}(\overset{\circ}{\chi}_0)$, on the one hand, we have by Corollary G.8.1

$$Z^{d\tilde{x}_1^l} = -\eta|\overset{\circ}{\chi}_0|\nu\epsilon\sigma(Z^{\tilde{h}_0^l}),$$

and on the other hand, with $\lambda := \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}]$, which is > 0 by Lemmas G.10 and G.9 (if $l = 2$)

$$\sigma'(Z^{h_1^l}) = (\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} |\sigma'(Z^{\tilde{h}_0^l})|^{p-1} \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^l}) Z^{d\tilde{x}_0^l}).$$

Recalling that $Z^{d\tilde{h}_0^l} = Z^{d\tilde{x}_0^l} \sigma'(Z^{\tilde{h}_0^l})$ and $Z^{d\tilde{h}_1^l} = Z^{d\tilde{x}_1^l} \sigma'(Z^{h_1^l})$, this leads to

$$\begin{aligned} \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] &= \eta|\overset{\circ}{\chi}_0|\nu(\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} \mathbb{E} \left[(-\epsilon \text{sign}(Z^{\tilde{h}_0^l}) Z^{d\tilde{x}_0^l}) \sigma'(-\epsilon \text{sign}(Z^{\tilde{h}_0^l}) Z^{d\tilde{x}_0^l}) \right. \\ &\quad \left. \text{sign}(Z^{\tilde{h}_0^l}) \sigma'(Z^{\tilde{h}_0^l}) |\sigma'(Z^{\tilde{h}_0^l})|^{p-1} \sigma(Z^{\tilde{h}_0^l}) \right], \end{aligned}$$

which, by conditioning on $Z^{\tilde{h}_0^l}$ and since $-\epsilon \text{sign}(Z^{\tilde{h}_0^l}) Z^{d\tilde{x}_0^l}$ and $Z^{d\tilde{x}_0^l}$ have the same distribution conditionally on $Z^{\tilde{h}_0^l}$, and since $\text{sign}(\sigma'(z)) = \text{sign}(z)$, gives

$$\mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] = \eta|\overset{\circ}{\chi}_0|\nu(\eta|\overset{\circ}{\chi}_0|\lambda)^{p-1} \mathbb{E} \left[Z^{d\tilde{x}_0^l} \sigma'(Z^{d\tilde{x}_0^l}) \right] \mathbb{E} \left[|\sigma'(Z^{\tilde{h}_0^l})|^p \sigma(Z^{\tilde{h}_0^l}) \right].$$

The term in front of the expectations is positive by assumption, and both expectations are positive because their integrands are both non-negative and positive on the positive part of the real line where the Gaussians $Z^{d\tilde{x}_0^l}$ and $Z^{\tilde{h}_0^l}$ have non-zero density. The expectations are also finite by Lemma L.1 because their integrands are polynomially bounded functions of some Gaussian vector with finite covariance variance matrix. By induction, we thus get that $0 < \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] < \infty$ for any $l \in [2, L]$, which concludes the proof. \square

H.3 Preliminaries on the third forward pass ($t = 2$)

In this section we wish to prove that similarly to the second forward pass, the quantities the quantities $\mathbb{E}[Z^{x_1^l} Z^{x_2^l}]$ and $\mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}]$ (which appear in the third forward pass at $t = 2$) are > 0 for any $l \in [1, L]$ when using the IP-LLR learning rates at $t = 0$ and $t = 1$. We assume here that the training samples ξ_0, ξ_1, ξ_2 are all distinct, which is probably not necessary for the result to hold but simplifies somewhat some parts of the proof and is in any case a very natural assumption.

Lemma H.3 (Forward pass of IP-LLR at $t = 2$). *Consider the IP-LLR parameterization of an L hidden-layer network, and assume that the activation function σ satisfies Assumption 3, and that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) \neq 0$ and $\lim_{m \rightarrow \infty} \partial_2 \ell(y_1, f_1(\xi_1)) \neq 0$. Assume further that the first three training samples ξ_0, ξ_1, ξ_2 are all distinct. Then, one has that for any $l \in [1, L]$,*

$$\begin{aligned} 0 &< \mathbb{E}[Z^{x_1^l} Z^{x_2^l}] < \infty \\ 0 &< \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}] < \infty \end{aligned}$$

Proof. We start with the case $l = 1$ and then induct over l from $l = 1$ to $l = L$ for both expectations simultaneously as the derivations are very similar.

The case $l = 1$.

Let us first unwind the expressions of $Z^{h_1^1}$ and $Z^{h_2^1}$. We have

$$Z^{h_1^1} = Z^{\tilde{h}_0^1(\xi_1)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_1 + 1) Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}),$$

and

$$Z^{h_2^1} = Z^{\tilde{h}_0^1(\xi_2)} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_2 + 1) Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}) - \eta \overset{\circ}{\chi}_1 (\xi_1^\top \xi_2 + 1) Z^{d\tilde{x}_1^1} \sigma'(Z^{h_1^1}).$$

The case of $\mathbb{E}[Z^{x_1^1} Z^{x_2^1}]$.

Recalling that $Z^{d\tilde{x}_1^1} = -\eta \overset{\circ}{\chi}_0 \nu \sigma(Z^{\tilde{h}_0^1})$ where $\nu := \mathbb{E}[Z^{d\tilde{h}_0^2} Z^{d\tilde{h}_1^2}]$. With the assumption that ξ_0, ξ_1, ξ_2 are all distinct, the vector $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)}, Z^{\tilde{h}_0^1(\xi_2)}, Z^{d\tilde{x}_0^1})$ has a non-degenerate Gaussian distribution, and we thus get

$$\begin{aligned} \mathbb{E}[Z^{x_1^1} Z^{x_2^1}] &= \int \sigma(u_1 - \mu_0 z \sigma'(u_0)) \sigma(u_2 - \mu_1 z \sigma'(u_0) + \mu_2 \sigma(u_0) \sigma'(u_1 - \mu_0 z \sigma'(u_0))) \times \\ &\quad q(u_0, u_1, u_2) p_z(z) d(u_0, u_1, u_2) dz, \end{aligned}$$

where $\mu_0 := \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_1 + 1)$, $\mu_1 := \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_2 + 1)$ and $\mu_2 := \eta^2 \overset{\circ}{\chi}_0 \overset{\circ}{\chi}_1 \nu (\xi_1^\top \xi_2 + 1)$, and q and p_z are the densities of non-degenerate Gaussians and are thus positive everywhere. Now the integrand is non-negative everywhere and we wish to show that it is positive at some given point of \mathbb{R}^4 , and it is also a polynomially bounded function of $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)}, Z^{\tilde{h}_0^1(\xi_2)}, Z^{d\tilde{x}_0^1})$ which shows that the expectation is finite. Since $Z^{d\tilde{x}_0^1}$ and $-Z^{d\tilde{x}_0^1}$ have the same distribution and it is independent of $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)}, Z^{\tilde{h}_0^1(\xi_2)})$, we can assume that $\mu_0 \geq 0$ W.L.O.G. Consider the point $(u_0^*, u_1^*, u_2^*, z^*)$ defined as $u_0^* = u_1^* = 1$, $z^* = -1$ and

$$u_2^* := |\mu_1| \sigma'(1) + |\mu_2| \sigma(1) \sigma'(1 + \mu_0 \sigma'(1)) + 1.$$

We show below that the integrand is > 0 at $(u_0^*, u_1^*, u_2^*, z^*)$. Since it is also a continuous function of (u_0, u_1, u_2, z) , we get that the expectation is positive.

Let us now show that the integrand is > 0 at $(u_0^*, u_1^*, u_2^*, z^*)$. We have

$$u_1^* - \mu_0 z^* \sigma'(u_0^*) = 1 + \mu_0 \sigma'(1) \geq 1 > 0,$$

and

$$-\mu_1 z^* \sigma'(u_0^*) = \mu_1 \sigma'(1) \geq -|\mu_1| \sigma'(1),$$

and finally

$$\mu_2 \sigma(u_0^*) \sigma'(u_1^* - \mu_0 z^* \sigma'(u_0^*)) = \mu_2 \sigma(1) \sigma'(1 + \mu_0 \sigma'(1)) \geq -|\mu_2| \sigma(1) \sigma'(1 + \mu_0 \sigma'(1)).$$

With the choice for u_2^* , one has that $u_2^* - \mu_1 z \sigma'(u_0^*) + \mu_2 \sigma(u_0^*) \sigma'(u_1^* - \mu_0 z \sigma'(u_0^*)) \geq 1 > 0$, which concludes the proof because σ is positive on the positive part of the real line.

The case of $\mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_2^1}]$.

We have

$$\begin{aligned} \mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_2^1}] &= \int \sigma(u_0) \sigma(u_2 - \mu_1 z \sigma'(u_0) + \mu_2 \sigma(u_0) \sigma'(u_1 - \mu_0 z \sigma'(u_0))) \times \\ &\quad q(u_0, u_1, u_2) p_z(z) d(u_0, u_1, u_2) dz, \end{aligned}$$

As for the case of $\mathbb{E}[Z^{x_1^1} Z^{x_2^1}]$, we show that the integrand is > 0 at the same point $(u_0^*, u_1^*, u_2^*, z^*)$ as above, and since it is also a continuous function of (u_0, u_1, u_2, z) , we get that the expectation is positive. It is also finite by Lemma L.1 because its integrand is a polynomially bounded function of $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)}, Z^{\tilde{h}_0^1(\xi_2)}, Z^{d\tilde{x}_0^1})$.

The case $l \in [2, L - 1]$.

Let $l \in [2, L - 1]$ and assume $\tau := \mathbb{E}[Z^{x_1^{l-1}} Z^{x_2^{l-1}}] > 0$ and $\rho := \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_2^{l-1}}] > 0$. Calling $\lambda := \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_1^{l-1}}]$ which is > 0 by Lemma G.10, and $\nu := \mathbb{E}[Z^{d\tilde{h}_0^{l+1}} Z^{d\tilde{h}_1^{l+1}}]$ which is also > 0 by Lemma H.2, we have

$$Z^{h_1^l} = -\eta \overset{\circ}{\chi}_0 \lambda Z^{d\tilde{x}_0^l} \sigma'(Z^{\tilde{h}_0^l}),$$

and

$$Z^{h_2^l} = -\eta \overset{\circ}{\chi}_0 \rho Z^{d\tilde{x}_0^l} \sigma'(Z^{\tilde{h}_0^l}) - \eta \overset{\circ}{\chi}_1 \tau Z^{d\tilde{x}_1^l} \sigma'(Z^{h_1^l}).$$

Finally recall that $Z^{d\tilde{x}_1^l} = -\eta \overset{\circ}{\chi}_0 \nu \sigma(Z^{\tilde{h}_0^l})$, and let us call $\mu_0 := \eta \overset{\circ}{\chi}_0 \lambda$, $\mu_1 := \eta \overset{\circ}{\chi}_0 \rho$ and $\mu_2 := \eta^2 \overset{\circ}{\chi}_0 \overset{\circ}{\chi}_1 \tau \nu$. μ_0 is $\neq 0$ because of the assumption on $\overset{\circ}{\chi}_0$. Since $Z^{d\tilde{x}_0^l}$ and $-Z^{d\tilde{x}_0^l}$ have the same distribution and it is independent of $Z^{\tilde{h}_0^l}$, we can assume $\mu_0 > 0$ W.L.O.G. Note then that since μ_1 is of the same sign as μ_0 ($\lambda \rho > 0$), this also implies $\mu_1 > 0$, and μ_2 has the sign of $\overset{\circ}{\chi}_1$. By assumption, $\overset{\circ}{\chi}_1 \neq 0$, and by the induction hypothesis and Lemma H.2 we have $\mu_2 \neq 0$.

The case of $\mathbb{E}[Z^{x_1^l} Z^{x_2^l}]$.

We have

$$\mathbb{E}[Z^{x_1^l} Z^{x_2^l}] = \int \sigma(-\mu_0 z \sigma'(u)) \sigma(-\mu_1 z \sigma'(u) + \mu_2 \sigma(u) \sigma'(-\mu_0 z \sigma'(u))) p_u(u) p_z(z) du dz,$$

where p_u and p_z are the densities of non-degenerate Gaussians ($Z^{\tilde{h}_0^l}$ and $Z^{d\tilde{x}_0^l}$ respectively) and are thus positive everywhere. Now the integrand is non-negative everywhere and we wish

to show that it is positive at some given point of \mathbb{R}^2 . The integrand is also a polynomially bounded function of $(Z^{\tilde{h}_0^l}, Z^{d\tilde{x}_0^l})$ which shows that the expectation is finite by Lemma L.1. Let $z^* = -1$ and $u > 0$. Then, $-\mu_0 z^* \sigma'(u) = \mu_0 \sigma'(u) > 0$ so that $\sigma(-\mu_0 z^* \sigma'(u)) > 0$. On the other hand, $-\mu_1 z^* \sigma'(u) = \mu_1 \alpha p u^{p-1}$, and

$$\begin{aligned} \mu_2 \sigma(u) \sigma'(-\mu_0 z^* \sigma'(u)) &= \mu_2 \alpha u^p \alpha p (\mu_0 \alpha p)^{p-1} u^{(p-1)^2} \\ &\geq -(\alpha p) |\mu_2| \alpha (\mu_0 \alpha p)^{p-1} u^{p-1} u^{(p-1)^2+1}. \end{aligned}$$

This leads to

$$-\mu_1 z^* \sigma'(u) + \mu_2 \sigma(u) \sigma'(-\mu_0 z^* \sigma'(u)) \geq \alpha p u^{p-1} \left[\mu_1 - |\mu_2| \alpha (\mu_0 \alpha p)^{p-1} u^{(p-1)^2+1} \right].$$

The quantity in the bracket is > 0 as soon as

$$u < \left[\frac{\mu_1}{|\mu_2| \alpha (\mu_0 \alpha p)^{p-1}} \right]^{\frac{1}{(p-1)^2+1}} =: \varepsilon$$

Calling $u^* := \varepsilon/2$, we thus get that the integrand is > 0 at (u^*, z^*) , and since it is a continuous function of (u, z) , the integral is positive.

The case of $\mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}]$.

We have

$$\mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}] = \int \sigma(u) \sigma(-\mu_1 z \sigma'(u) + \mu_2 \sigma(u) \sigma'(-\mu_0 z \sigma'(u))) p_u(u) p_z(z) du dz,$$

The integrand is non-negative everywhere and with $z^* = -1$ and $u^* = \varepsilon/2$ as above, one shows that the integrand is > 0 at (u^*, z^*) which in turn implies that the expectation is positive. It is also finite for the same reasons as $\mathbb{E}[Z^{x_1^l} Z^{x_2^l}]$. This now concludes the induction over $l \in [1, L-1]$ which thus shows that $\mathbb{E}[Z^{x_1^l} Z^{x_2^l}]$ and $\mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}]$ are > 0 and finite for any $l \in [1, L-1]$. Those expectations are also finite as their integrands are polynomially bounded functions of Gaussian vectors which have finite covariance matrices.

The case $l = L$.

Let $\tau := \mathbb{E}[Z^{x_1^{L-1}} Z^{x_2^{L-1}}] > 0$ and $\rho := \mathbb{E}[Z^{\tilde{x}_0^{L-1}} Z^{x_2^{L-1}}] > 0$ by the previous induction. Calling $\lambda := \mathbb{E}[Z^{\tilde{x}_0^{L-1}} Z^{x_1^{L-1}}]$ which is > 0 by Lemma G.10, we have

$$Z^{h_1^L} = -\eta \overset{\circ}{\chi}_0 \lambda Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L}),$$

and

$$Z^{h_2^L} = -\eta \overset{\circ}{\chi}_0 \rho Z^{U^{L+1}} \sigma'(Z^{\tilde{h}_0^L}) - \eta \overset{\circ}{\chi}_1 \tau Z^{d\tilde{x}_1^L} \sigma'(Z^{h_1^L}).$$

Finally recall that $Z^{d\tilde{x}_1^L} = Z^{U^{L+1}} - \eta \overset{\circ}{\chi}_0 \sigma'(Z^{\tilde{h}_0^L})$, and let us call $\mu_0 := \eta \overset{\circ}{\chi}_0 \lambda$, $\mu_1 := \eta \overset{\circ}{\chi}_0 \rho$, $\mu_2 := \eta \overset{\circ}{\chi}_1 \tau$, and finally $\mu_3 := \eta^2 \overset{\circ}{\chi}_0 \overset{\circ}{\chi}_1 \tau$. Since $Z^{d\tilde{x}_0^L}$ and $-Z^{d\tilde{x}_0^L}$ have the same distribution and it is independent of $Z^{\tilde{h}_0^L}$, we can assume $\mu_0 > 0$ W.L.O.G. Note then that since μ_1 is of the same sign as μ_0 , this also implies $\mu_1 > 0$, and μ_2 has the sign of $\overset{\circ}{\chi}_1$. In addition, with the assumptions and previous results, we have $\mu_2 \neq 0$ and $\mu_3 \neq 0$.

The case of $\mathbb{E}[Z^{x_1^L} Z^{x_2^L}]$.

We have

$$\begin{aligned} \mathbb{E}[Z^{x_1^L} Z^{x_2^L}] &= \int \sigma(-\mu_0 z \sigma'(u)) \sigma(-\mu_1 z \sigma'(u) + (-\mu_2 z + \mu_3 \sigma(u)) \sigma'(-\mu_0 z \sigma'(u))) \times \\ &\quad p_u(u) p_z(z) du dz, \end{aligned}$$

where p_u and p_z are the densities of non-degenerate Gaussians ($Z^{\tilde{h}_0^L}$ and $Z^{U^{L+1}}$ respectively) and are thus positive everywhere. Now the integrand is non-negative everywhere and we wish to show that it is positive at some point of \mathbb{R}^2 . The integrand is also a polynomially bounded function of $(Z^{\tilde{h}_0^L}, Z^{U^{L+1}})$ and the expectation is thus finite by Lemma L.1. We first take a closer look at the second term inside σ . Let $z \leq 0, u \geq 0$. We have

$$-\mu_1 z \sigma'(u) = \mu_1 |z| \alpha p u^{p-1},$$

as well as

$$-\mu_2 z + \mu_3 \sigma(u) = -\mu_2 z + \mu_3 \alpha u^p,$$

and

$$\sigma'(-\mu_0 z \sigma'(u)) = \alpha p (\mu_0 \alpha p)^{p-1} |z|^{p-1} u^{(p-1)^2}.$$

We thus get that

$$\begin{aligned} -\mu_1 z \sigma'(u) + (-\mu_2 z + \mu_3 \sigma(u)) \sigma'(-\mu_0 z \sigma'(u)) = \\ \alpha p |z| u^{p-1} \underbrace{\left[\mu_1 + (-\mu_2 z + \mu_3 \alpha |u|^p) (\mu_0 \alpha p)^{p-1} |z|^{p-2} |u|^{(p-1)(p-2)} \right]}_{F(u,z)} \end{aligned}$$

Because $p - 2 \geq 0$, the function F is continuous over \mathbb{R}^2 , and we have $F(0, 0) = \mu_1 > 0$. Therefore, there exists $u^* > 0$ and $z^* < 0$ such that $F(u^*, z^*) > 0$. With such a pair (u^*, z^*) we get that the integrand is > 0 at (u^*, z^*) , and since it is a continuous function of (u, z) , it follows that the expectation is positive.

The case of $\mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_2^L}]$.

A similar argument to the case of $\mathbb{E}[Z^{x_1^L} Z^{x_2^L}]$ applies and we get that the expectation is positive, which concludes the proof. \square

H.4 Proof of the second implication

Lemma H.4 (Learning rates for stable learning with IP at $t = 1$). *Consider an L -hidden layer fully-connected neural network with $L \geq 3$ in the integrable parameterization, and with no bias terms, except at the first layer. Assume that the activation function σ satisfies Assumption 3, and that $\lim_{m \rightarrow \infty} \partial_2 \ell(y_0, f_0(\xi_0)) \neq 0$ and $\lim_{m \rightarrow \infty} \partial_2 \ell(y_1, f_0(\xi_1)) \neq 0$. Assume further that $\xi_1^\top \xi_2 \neq 0$, that the first three training samples ξ_0, ξ_1, ξ_2 are all distinct, and that at $t = 0$ (i.e., to compute $\Delta W^l(1)$) $c_l = \gamma_l(p)$ (see Definition 3.2) for any $l \in [1, L + 1]$. Finally assume that Equation (3.2) holds:*

$$\begin{cases} \frac{1}{m} \|\Delta W^l(2) x_2^{l-1}\|^2 = \Theta(1), & l \in [1, L] \\ (\Delta W^{L+1}(2))^\top x_2^L = \Theta(1) \end{cases}$$

Then, one necessarily has that at $t = 1$, $c_1 = c_{l+1} = -1$ and $c_l = -2$ for any $l \in [2, L]$.

Proof. We first treat the case $l = 1$ and then induct over l from $l = 2$ to $l = L$ and conclude by the case $l = L + 1$. Note that because of the assumptions, Lemma H.2 holds and the claim of Lemma H.3 will hold at layer l as soon as we show $c_1 = -1$ and $c_k = -2$ for $k \in [2, L]$.

The case $l = 1$.

We have

$$\Delta W^1(2)\xi_2 = -\eta m^{-(1+c_1)}\chi_1(\xi_1^\top \xi_2)d\tilde{x}_1^1 \odot \sigma'(h_1^1),$$

so that

$$\frac{1}{m}\|\Delta W^1(2)\xi_2\|^2 = m^{-2(1+c_1)}(\eta\chi_1(\xi_1^\top \xi_2))^2 \frac{1}{m}\|d\tilde{x}_1^1 \odot \sigma'(h_1^1)\|^2.$$

Recall that $d\tilde{x}_1^1 = -\eta\chi_0((d\tilde{h}_0^2)^\top d\tilde{h}_1^2)/m\sigma(\tilde{h}_0^1)$, so that by the Master Theorem,

$$\frac{1}{m}\|d\tilde{x}_1^1 \odot \sigma'(h_1^1)\|^2 \xrightarrow[m \rightarrow \infty]{a.s.} (\eta\overset{\circ}{\chi}_0\nu)^2 \mathbb{E}[\sigma(Z^{\tilde{h}_0^1})^2 \sigma'(Z^{h_1^1})^2],$$

where $\nu := \mathbb{E}[Z^{d\tilde{h}_0^2} Z^{d\tilde{h}_1^2}] > 0$ by Lemma H.2. The term in front of the expectation is > 0 with the assumptions. On the other hand, the term $(\eta\chi_1(\xi_1^\top \xi_2))^2$ converges almost surely towards $(\eta\overset{\circ}{\chi}_1(\xi_1^\top \xi_2))^2$ which is also > 0 with the assumptions. We show below that the expectation is > 0 , which proves that c_1 must be equal to 1 since by assumption $\frac{1}{m}\|\Delta W^1(2)\xi_2\|^2 = \Theta(1)$. Recall that

$$Z^{h_1^1} = Z^{\tilde{h}_0^1(\xi_1)} - \eta\overset{\circ}{\chi}_0(\xi_0^\top \xi_1 + 1)Z^{d\tilde{x}_0^1} \sigma'(Z^{\tilde{h}_0^1}).$$

The integrand in the expectation is non-negative, and it simply remains to show that is not almost surely zero. Because $Z^{d\tilde{x}_0^1}$ and $-Z^{d\tilde{x}_0^1}$ have the same distribution, and since it is independent of $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)})$, we can assume W.L.O.G. that $\mu := \eta\overset{\circ}{\chi}_0(\xi_0^\top \xi_1 + 1) \geq 0$. As usual, the vector $(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)})$ has a Gaussian distribution which is degenerate only if $\xi_1 = \xi_0$, which is precluded by the assumptions. Note that in any case, the expectation is finite by Lemma L.1 since its integrand is a polynomially bounded function of a Gaussian vector with finite covariance matrix. Since $\xi \neq \xi_0$ by assumption, we have

$$\mathbb{E}[\sigma(Z^{\tilde{h}_0^1})^2 \sigma'(Z^{h_1^1})^2] = \int \sigma(u)^2 \sigma'(v - \mu z \sigma'(u))^2 p_{u,v}(u, v) p_z(z) du dv dz,$$

where $p_{u,v}$ and p_z are the densities of non-degenerate Gaussians ($(Z^{\tilde{h}_0^1}, Z^{\tilde{h}_0^1(\xi_1)})$ and $Z^{d\tilde{x}_0^1}$ respectively) and are thus well-defined and positive everywhere. Again, one sees that at point $(u^*, v^*, z^*) = (1, 1, -1)$ the integrand is > 0 , and since it is a continuous function, this proves that the expectation is positive. It is also finite by Lemma L.1 since the integrand is a polynomially bounded function of a Gaussian vector with finite covariance matrix.

The case $l \in [1, L - 1]$

Let $l \in [2, L - 1]$. We have already shown that $c_1 = -1$. Assume now that $c_k = -2$ for $k \in [2, l - 1]$ (note that if $l = 2$ this means no additional assumption). Then we have

$$\Delta W^l(2)x_2^{l-1} = -\eta m^{-(2+c_l)}\chi_1 \frac{(x_1^{l-1})^\top x_2^{l-1}}{m} d\tilde{x}_1^l \odot \sigma'(h_1^l),$$

so that

$$\frac{1}{m}\|\Delta W^l(2)x_2^{l-1}\|^2 = m^{-2(2+c_l)} \left(\eta\chi_1 \frac{(x_1^{l-1})^\top x_2^{l-1}}{m} \right)^2 \frac{1}{m}\|d\tilde{x}_1^l \odot \sigma'(h_1^l)\|^2.$$

In addition, we have $d\tilde{x}_1^l = -\eta\chi_0((d\tilde{h}_0^{l+1})^\top d\tilde{h}_1^{l+1})/m\sigma(\tilde{h}_0^l)$, so that by the Master Theorem,

$$\frac{1}{m}\|d\tilde{x}_1^l \odot \sigma'(h_1^l)\|^2 \xrightarrow[m \rightarrow \infty]{a.s.} (\eta\overset{\circ}{\chi}_0\nu)^2 \mathbb{E}[\sigma(Z^{\tilde{h}_0^l})^2 \sigma'(Z^{h_1^l})^2],$$

where $\nu := \mathbb{E}[Z^{d\tilde{h}_0^{l+1}} Z^{d\tilde{h}_1^{l+1}}]$ is such that $0 < \nu < \infty$ by Lemma H.2. Recall that

$$Z^{h_1^l} = -\eta\overset{\circ}{\chi}_0\lambda Z^{d\tilde{x}_0^l}\sigma'(Z^{\tilde{h}_0^l}),$$

with $\lambda := \mathbb{E}[Z^{x_0^{l-1}} Z^{x_1^{l-1}}]$ such that $0 < \lambda < \infty$ by Lemmas G.9 and G.10, which leads to

$$\mathbb{E}[\sigma(Z^{\tilde{h}_0^l})^2\sigma'(Z^{h_1^l})^2] = \int \sigma(u)^2\sigma'(-\mu z\sigma'(u))^2 p_u(u)p_z(z) dudz,$$

where $\mu := \eta\overset{\circ}{\chi}_0\lambda$ which is $\neq 0$ with the assumptions, and p_u and p_z are the densities of two non-degenerate Gaussians ($Z^{\tilde{h}_0^l}$ and $Z^{d\tilde{x}_0^l}$ respectively) and are thus positive everywhere. Since $Z^{d\tilde{x}_0^l}$ and $-Z^{d\tilde{x}_0^l}$ have the same distribution and it is independent of $Z^{\tilde{h}_0^l}$ we can assume $\mu > 0$ W.L.O.G. Then, we see that at point $(u^*, z^*) = (1, -1)$ the integrand is > 0 , and since it is a continuous function, this proves that the expectation is positive. It is also finite by Lemma L.1 since the integrand is a polynomially bounded function of a Gaussian vector with finite covariance matrix. The term $(\eta\overset{\circ}{\chi}_0\nu)^2$ in front of the expectation is > 0 and finite with the assumptions. Finally the term $(\eta\chi_1((x_1^{l-1})^\top x_2^{l-1})/m)^2$ converges almost surely towards $(\eta\overset{\circ}{\chi}_1\tau)^2$ by the Master Theorem, where $\tau := \mathbb{E}[Z^{x_1^{l-1}} Z^{x_2^{l-1}}]$ is > 0 and finite by Lemma H.3, which shows that $(\eta\overset{\circ}{\chi}_1\tau)^2$ is > 0 and finite with the assumptions. Since $\|d\tilde{x}_1^l \odot \sigma'(h_1^l)\|^2/m = \Theta(1)$ by assumption, then c_l must be equal to -2 otherwise $\|d\tilde{x}_1^l \odot \sigma'(h_1^l)\|^2/m$ would either converge towards 0 or diverge towards ∞ almost surely.

The case $l = L$. We have already proved that at $t = 1$, $c_1 = -1$ and $c_l = -2$ for $l \in [2, L]$. We have

$$|(\Delta W^{L+1}(2))^\top x_2^L| = \eta m^{-(1+c_{L+1})} |\chi_1| \left| \frac{(x_1^L)^\top x_2^L}{m} \right|.$$

By the Master Theorem,

$$\left| \frac{(x_1^L)^\top x_2^L}{m} \right| \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{x_1^L} Z^{x_2^L}]$$

which is > 0 and finite by Lemma H.3. On the other hand, $\eta|\chi_1|$ converges almost surely towards $\eta|\overset{\circ}{\chi}_1|$ which is also > 0 and finite. This shows that since $|(\Delta W^{L+1}(2))^\top x_2^L| = \Theta(1)$ then we must have $c_{L+1} = -1$ to avoid vanishing towards 0 or explosion towards $+\infty$ as $m \rightarrow \infty$, which concludes the proof. \square

I Proof of the non-triviality of IP-LLR: Theorem 4.1

Proof. Claims (i) and (ii) of Theorem 4.1 have already been shown in Theorem G.12. Claim (iii) simply stems from Corollary G.7.1 with $t = 2$ and the fact that all the variables Z which appear are polynomially bounded functions of the vector Z_0 (see Definition L.1) by a simple induction. \square

J Proof of the equivalence between IP-LLR and μ P: Proposition 4.1 and Theorem 4.2

In this section, we present the proofs of the equivalence between IP-LLR and hybrid versions of μ P both at finite-width and in the large-width limit. Because we need to use the homogeneity property, we consider a positively p -homogeneous activation function σ and no bias terms except at the first layer for all the parameterizations we consider. We assume $p \geq 1$ for the finite-width case, which includes ReLU, and $p \geq 2$ in the infinite-width case as we use the Tensor Program framework for the proof and thus require some smoothness.

J.1 Finite-width equivalence: Proposition 4.1

We start with a preliminary Lemma showing the equivalence at $t = 1$ and then do the proof of Proposition 4.1 by induction.

J.1.1 Equivalence at $t = 1$

Lemma J.1 (First weight updates of HP). *Consider the IP-LLR and HP parameterizations with a positively p -homogeneous activation function, and $p \geq 1$, and no bias terms except at the first layer, and let us sub/super-script the variables of each models with IP and HP respectively. Assume the first training sample (ξ_0, y_0) and the loss ℓ are the same for both parameterizations. Assume further that $\chi_0^{HP} \neq 0$, and simply denote by η the base learning rate of the IP-LLR parameterization. Finally consider for HP the initial learning rate: $\eta_{HP}(0) = (\chi_0^{IP}/\chi_0^{HP})\eta$, and let $\xi \in \mathbb{R}^d$ be an input to both networks. Then, dropping the dependency of the weights at $t = 1$ on η and η_{HP} , one has:*

$$\begin{aligned} \forall l \in [1, L + 1], \quad W_{HP}^l(1) &= W_{IP}^l(1) \\ B_{HP}^1(1) &= B_{IP}^1(1) \\ f_1^{HP}(\xi) &= f_1^{IP}(\xi) \end{aligned}$$

Proof. By definition (see Section 4.2.1), we have

$$\begin{aligned} W_{HP}^1(1) &= W_{IP}^1(0) + \Delta W_{\mu P}^1(1) \\ B_{HP}^1(1) &= B_{IP}^1(0) + \Delta B_{\mu P}^1(1) \\ W_{HP}^l(1) &= W_{IP}^l(0) + \Delta W_{\mu P}^l(1), \quad l \in [2, L] \\ W_{HP}^{L+1}(1) &= W_{IP}^{L+1}(0) + \Delta W_{\mu P}^{L+1}(1) \end{aligned}$$

Using Corollaries F.4.2, and Lemma F.5, and the fact that $\eta_{HP}(0)\chi_0^{\mu P} = \eta\chi_0^{IP}$ we have:

$$\begin{aligned} \Delta W_{\mu P}^1(1) &= -\eta_{HP}(0)\chi_0^{\mu P} d\tilde{h}_0^1 \xi_0^\top \\ &= -\eta\chi_0^{IP} d\tilde{h}_0^1 \xi_0^\top \\ &= \Delta W_{IP}^1(1), \end{aligned}$$

$$\begin{aligned} \Delta B_{\mu P}^1(1) &= -\eta_{HP}(0)\chi_0^{\mu P} d\tilde{h}_0^1 \\ &= -\eta\chi_0^{IP} d\tilde{h}_0^1 \\ &= \Delta B_{IP}^1(1), \end{aligned}$$

and, for $l \in [2, L]$

$$\begin{aligned} \Delta W_{\mu P}^l(1) &= -\eta_{HP}\chi_0^{\mu P} \frac{d\tilde{h}_0^l(\tilde{x}_0^{l-1})^\top}{m} \\ &= -\eta\chi_0^{IP} \frac{d\tilde{h}_0^l(\tilde{x}_0^{l-1})^\top}{m} \\ &= \Delta W_{IP}^l(1), \end{aligned}$$

and finally

$$\begin{aligned} \Delta W_{\mu P}^{L+1}(1) &= -\eta_{HP}\chi_0^{\mu P} \tilde{x}_0^L/m \\ &= -\eta\chi_0^{IP} \tilde{x}_0^L/m \\ &= \Delta W_{IP}^{L+1}(1), \end{aligned}$$

where the $\Delta W_{IP}^l(1)$ and $\Delta B^1(1)$ are computed with the base learning rate η . We then get $W_{HP}^l(1) = W_{IP}^l(1)$ for all l , and it follows that for any input ξ , $f_1^{HP}(\xi) = f_1^{IP}(\xi)$. \square

J.1.2 Proof of Proposition 4.1

Proof. We first show by induction that the effective weight matrices and the effective biases of the first layer are the same for both parameterizations at any time step ≥ 1 , which will then immediately yield the result. We have already shown in Lemma J.1 that with the choice of initial learning rate for HP, $W_{\text{HP}}^l(1) = W_{\text{IP}}^l(1)$ for all $l \in [1, L+1]$, and $B_{\text{HP}}^l(1) = B_{\text{IP}}^l(1)$ as well as $f_1^{\text{HP}}(\xi) = f_1^{\text{IP}}(\xi)$.

Now let $s \geq 1$, and assume that for all $l \in [1, L+1]$, $W_{\text{HP}}^l(s) = W_{\text{IP}}^l(s)$, and $B_{\text{HP}}^l(s) = B_{\text{IP}}^l(s)$. We want to show that this also holds true for the next time step $s+1$. An easy induction shows that since the effective weights of all layers are equal, and since by assumption the s -th training sample (ξ_s, y_s) is the same for both parameterization, we get that for any $l \in [1, L+1]$, $x_{s,\text{HP}}^l = x_{s,\text{IP}}^l$, $h_{s,\text{HP}}^l = h_{s,\text{IP}}^l$, as well as $f_s^{\text{HP}}(\xi_s) = f_s^{\text{IP}}(\xi_s)$, and therefore $\chi_s^{\text{HP}} = \chi_s^{\text{IP}}$ since by assumption both parameterization use the same loss. This in turn will give by another easy induction that for any $l \in [1, L+1]$, $dx_{s,\text{HP}}^l = dx_{s,\text{IP}}^l$, $dh_{s,\text{HP}}^l = dh_{s,\text{IP}}^l$. Now, by Equation (A.3) we have, on the one hand (recall that $s+1 \geq 2$ so that the base learning for both models for the $(s+1)$ -th SGD step is η)

$$\Delta W_{\text{HP}}^1(s+1) = -\eta m^{-(2a_1^{\mu\text{P}} + c_1^{\mu\text{P}})} dh_{s,\text{HP}}^1 \xi_s^\top$$

and for $l \in [2, L]$

$$\Delta W_{\text{HP}}^l(s) = -\eta m^{-(2a_l^{\mu\text{P}} + c_l^{\mu\text{P}})} dh_{s,\text{HP}}^l x_{s,\text{HP}}^{l-1}$$

and finally

$$\Delta W_{\text{HP}}^{L+1}(s) = -\eta m^{-(2a_{L+1}^{\mu\text{P}} + c_{L+1}^{\mu\text{P}})} x_{s,\text{HP}}^L$$

On the other hand, we have

$$\Delta W_{\text{IP}}^1(s) = -\eta m^{-(2a_1^{\text{IP}} + c_1^{\text{IP}})} dh_{s,\text{IP}}^1 \xi_s^\top$$

and for $l \in [2, L]$

$$\Delta W_{\text{IP}}^l(s) = -\eta m^{-(2a_l^{\text{IP}} + c_l^{\text{IP}})} dh_{s,\text{IP}}^l x_{s,\text{IP}}^{l-1}$$

and finally

$$\Delta W_{\text{IP}}^{L+1}(s) = -\eta m^{-(2a_{L+1}^{\text{IP}} + c_{L+1}^{\text{IP}})} x_{s,\text{IP}}^L$$

To see that the quantities are equal, we only need to observe that since $s+1 \geq 1$

$$\begin{aligned} 2a_1^{\mu\text{P}} + c_1^{\mu\text{P}} &= -1 = 2a_1^{\text{IP}} + c_1^{\text{IP}} \\ 2a_l^{\mu\text{P}} + c_l^{\mu\text{P}} &= 0 = 2a_l^{\text{IP}} + c_l^{\text{IP}} \\ 2a_{L+1}^{\mu\text{P}} + c_{L+1}^{\mu\text{P}} &= 1 = 2a_{L+1}^{\text{IP}} + c_{L+1}^{\text{IP}} \end{aligned}$$

(recall that for $s \geq 1$, $c_{L+1}^{\text{IP}} = c_{L+1}^{\text{IP}} = -1$, and $c_l^{\text{IP}} = -2$ for $l \in [2, L]$). We thus find $\Delta W_{\text{HP}}^l(s) = \Delta W_{\text{IP}}^l(s)$ for all l , and since $W_{\text{HP}}^l(s) = W_{\text{IP}}^l(s)$ by assumption, we get $W_{\text{HP}}^l(s+1) = W_{\text{IP}}^l(s+1)$ for all l which concludes the induction.

The effective weights being equal in both parameterizations for all time steps ≥ 1 , we get that at time step $t \geq 1$, for any input $\xi \in \mathbb{R}^d$, the outputs $f_t^{\text{HP}}(\xi)$ and $f_t^{\text{IP}}(\xi)$ are the same, which concludes the proof. \square

J.2 Infinite-width equivalence: [Theorem 4.2](#)

In this section we prove [Theorem 4.2](#) which states the equivalence between IP-LLR (see [Definition 4.1](#)) and HPZ (see [Section 4.2.2](#)). We start by a couple of preliminary results on the dynamics of HPZ, then proceed to prove the main induction step over t , and finally conclude by putting the results together to prove the theorem.

J.2.1 Preliminary results

Lemma J.2 (μP is zero at initialization). *Consider the μP parameterization with an activation function satisfying [Assumption 2](#) and a loss function ℓ satisfying [Assumption 1](#), and no bias terms except at the first layer. Let $\xi \in \mathbb{R}^d$ be an input to the network. One has:*

$$\begin{aligned} f_0(\xi) &\xrightarrow[m \rightarrow \infty]{a.s.} 0 \\ \chi_0 &\xrightarrow[m \rightarrow \infty]{a.s.} \overset{\circ}{\chi}_0 := \partial_2 \ell(y_0, 0) \end{aligned}$$

Remark J.1. The result on the almost sure convergence of χ_0 ensures that the latter is a valid initial scalar in the Tensor Program defining the computations associated with μP (and thus and HPZ). Also note that the limit of χ_0 is the same as for IP-LLR (see [Lemma G.1](#)).

Proof. μP is designed so that $h_0^l = \tilde{h}_0^l$ and $x_0^l = \tilde{x}_0^l$ for any $l \in [2, L]$, and as already proved in [Lemma G.1](#), the tilde variables are vectors in the Tensor Program. Since $f_0(\xi) = m^{-1}(U^{L+1})^\top \tilde{x}_0^L$ we get by the master theorem that $f_0(\xi)$ converges almost surely towards $\mathbb{E}[Z^{U^{L+1}} Z^{\tilde{x}_0^L}]$. By [Lemma L.2](#), $Z^{\tilde{h}_0^L} = \widehat{Z}^{\tilde{h}_0^L}$ and $Z^{U^{L+1}} = \widehat{Z}^{U^{L+1}}$ by definition, and by the ZHat rule, $\widehat{Z}^{\tilde{h}_0^L}$ and $\widehat{Z}^{U^{L+1}}$ are independent, and since $\mathbb{E}[Z^{U^{L+1}}] = 0$ and $\mathbb{E}[(Z^{\tilde{x}_0^L})^2] < \infty$ we get that $f_0(\xi)$ converges almost surely towards 0. The result on the limit of χ_0 is then simply a consequence of the fact that $\partial_2 \ell(y_0, \cdot)$ is continuous by assumption. \square

Lemma J.3 (Weight updates for μP at any time step). *Consider the μP parameterization with a differentiable activation function σ and no bias terms except at the first layer, and let $t \geq 1$. Then, dropping the dependency of the forward and backward passes on ξ_t at time t , one has:*

$$\begin{aligned} \Delta W^{L+1}(t+1) &= -\eta \chi_t x_t^L / m, \\ \Delta W^l(t+1) &= -\eta \chi_t \frac{d\tilde{h}_t^l (x_t^{l-1})^\top}{m}, \quad l \in [2, L], \\ \Delta W^1(t+1) &= -\eta \chi_t d\tilde{h}_t^1 \xi_t^\top, \\ \Delta B^1(t+1) &= -\eta \chi_t d\tilde{h}_t^1. \end{aligned}$$

Remark J.2. Because HPZ and μP have the same parameterization for $t \geq 1$ (see [Section 4.2.2](#)), the formulas above for the updates are the same for HPZ, the only difference is that, at finite width, the x_t^l and $d\tilde{h}_t^l$ differ from HPZ to μP because $W_{\text{HPZ}}^l(t) = W_{\mu\text{P}}^l(t) - W_{\mu\text{P}}^l(0)$. Note that the formulas are also exactly the same as for IP-LLR (see [Lemma G.3](#)) but again the quantities x_t^l and $d\tilde{h}_t^l$ differ for μP and IP-LLR because of the initial weight contribution in $W^l(t)$ which is different for the intermediate layers of both parameterizations.

Proof. By [Equation \(A.6\)](#), we have

$$\begin{aligned} \Delta W^{L+1}(t+1) &= -\eta m^{-(2a_{L+1} + c_{L+1})} \chi_t x_t^L \\ &= -\eta \chi_t x_t^L / m, \end{aligned}$$

because $2a_{L+1} + c_{L+1} = 2 - 1 = 1$ for μP . For $l \in [2, L]$, we have by Equation (A.3)

$$\begin{aligned}\Delta W^l(t+1) &= -\eta\chi_t m^{-(2a_l+c_l)} dh_t^l (x_t^{l-1})^\top \\ &= -\eta\chi_t \frac{d\tilde{h}_t^l (x_t^{l-1})^\top}{m},\end{aligned}$$

because $dh_t^l = m^{-1}d\tilde{h}_t^l$ and $2a_l + c_l = 1 - 1 = 0$ for μP . Finally, for $l = 1$ we have again by Equation (A.3)

$$\begin{aligned}\Delta W^1(t+1) &= -\eta\chi_t m^{-(2a_1+c_1)} dh_t^1 \xi_t^\top \\ &= -\eta\chi_t d\tilde{h}_t^1 \xi_t^\top,\end{aligned}$$

because $2a_1 + c_1 = -1$ for μP and $dh_t^1 = d\tilde{h}_t^1$. A similar argument holds for $\Delta B^1(t+1)$, which concludes the proof. \square

Theorem J.4 (Weights in HPZ at time t). *Consider the HPZ parameterization with a differentiable activation function σ and no bias terms except at the first layer. Then, for any $t \geq 1$, one has:*

- (i) $W^1(t) = U^1 - \eta\chi_0 d\tilde{h}_0^1 \xi_0^\top - \eta \left(\sum_{s=1}^{t-1} \chi_s d\tilde{h}_s^1 \xi_s^\top \right)$,
- (ii) $B^1(t) = v^1 - \eta\chi_0 d\tilde{h}_0^1 - \eta \left(\sum_{s=1}^{t-1} \chi_s d\tilde{h}_s^1 \right)$,
- (iii) $W^l(t) = -\eta\chi_0 \frac{d\tilde{h}_0^l (\tilde{x}_0^{l-1})^\top}{m} - \eta \left(\sum_{s=1}^{t-1} \chi_s \frac{d\tilde{h}_s^l (x_s^{l-1})^\top}{m} \right)$, $l \in [2, L]$,
- (iv) $W^{L+1}(t) = U^{L+1}/m - \eta\chi_0 \tilde{x}_0^L/m - \eta \left(\sum_{s=1}^{t-1} \chi_s x_s^L/m \right)$.

Proof. The formulas are correct at $t = 1$ by definition of HPZ and by Lemma F.5 which gives the first weight updates for μP . Then, an easy induction using Lemma J.3 yields the result. \square

Lemma J.5 (Backward pass of HPZ at time t). *Consider the HPZ parameterization with a differentiable activation function σ and no bias terms except at the first layer. Then, for any $t \geq 1$, dropping the dependency of the forward pass at time t on ξ_t , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

- (i) $d\tilde{x}_t^L = w^{L+1}(t) = U^{L+1} - \eta\chi_0 \tilde{x}_0^L - \eta \sum_{s=1}^{t-1} \chi_s x_s^L$,
- (ii) $d\tilde{x}_t^{l-1} = -\eta\chi_0 \frac{(d\tilde{h}_0^l)^\top d\tilde{h}_t^l}{m} \tilde{x}_0^{l-1} - \eta \sum_{s=1}^{t-1} \chi_s \frac{(d\tilde{h}_s^l)^\top d\tilde{h}_t^l}{m} x_s^{l-1}$, $l \in [2, L]$.

Proof. By definition, we have

$$\begin{aligned}d\tilde{x}_t^L &= m dx_t^L \\ &= m W^{L+1}(t) \\ &= U^{L+1} - \eta\chi_0 \tilde{x}_0^L - \eta \sum_{s=1}^{t-1} \chi_s x_s^L\end{aligned}$$

where the last equality stems from Theorem J.4.

Let $l \in [2, L]$, we have:

$$\begin{aligned}d\tilde{x}_t^{l-1} &= (W^l(t))^\top d\tilde{h}_t^l \\ &= -\eta\chi_0 \frac{(d\tilde{h}_0^l)^\top d\tilde{h}_t^l}{m} \tilde{x}_0^{l-1} - \eta \sum_{s=1}^{t-1} \chi_s \frac{(d\tilde{h}_s^l)^\top d\tilde{h}_t^l}{m} x_s^{l-1}\end{aligned}$$

where the second equality stems from Theorem J.4. \square

Lemma J.6 (Z for the forward pass of HPZ at time $t = 1$). *Consider the HPZ parameterization with an activation function σ satisfying Assumption 2 and no bias terms except at the first layer. Let $\xi \in \mathbb{R}^d$ be an input to the network. Then, for any $l \in [1, L]$, $h_1^l(\xi), x_1^l(\xi), d\tilde{x}_1^l, d\tilde{h}_1^l$ are vectors in the program, $f_1(\xi)$ is a scalar in the program, and χ_1 is a valid initial scalar in the program. Additionally, dropping the dependency of the forward pass at time $t = 1$ on ξ , and of the first forward and backward passes on ξ_0 , one has:*

- (i) $Z^{h_1^1} = Z^{W^1(1)\xi + B^1(1)} = Z^{U^1\xi} + Z^{v^1} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^1}$,
- (ii) $Z^{h_1^l} = Z^{W^l(1)x_1^{l-1}} = -\eta\overset{\circ}{\chi}_0\mathbb{E}[Z^{\tilde{x}_0^{l-1}}Z^{x_1^{l-1}}]Z^{d\tilde{h}_0^l}$, $l \in [2, L]$,
- (iii) $f_1(\xi) = (W^{L+1}(1))^\top x_1^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}}Z^{x_1^L}] - \eta\overset{\circ}{\chi}_0\mathbb{E}[Z^{\tilde{x}_0^L}Z^{x_1^L}]$.

Proof. By Theorem J.4, with $t = 1$, one has that $h_1^1 = U^1\xi + v^1 - \eta\chi_0(\xi_0^\top\xi + 1)d\tilde{h}_0^1$. By Lemma G.1, $d\tilde{h}_0^1$ is a vector in the Tensor Program (recall that the tilde variables at initialization do not depend on the choice of parameterization) and by Lemma J.2 χ_0 is a valid initial scalar in the program which has an almost sure limit $\overset{\circ}{\chi}_0 := \partial_2\ell(y_0, 0)$ as $m \rightarrow \infty$ (see Remark J.1). In addition, $U^1\xi$ and v^1 are initial vectors in the program, which thus shows that h_1^1 is a vector in the program by the NonLin operation. This also gives that $x_1^1 = \sigma(h_1^1)$ is a vector in the program since σ is pseudo-Lipschitz by assumption. Moreover, by ZNonLin, we have $Z^{h_1^1} = Z^{U^1\xi} + Z^{v^1} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^1}$. Let $l \in [2, L]$ and assume that h_1^{l-1}, x_1^{l-1} are vectors in the program. Then, by Theorem J.4 with $t = 1$, we get

$$h_1^l = -\eta\chi_0 \frac{(\tilde{x}_0^{l-1})^\top x_1^{l-1}}{m} d\tilde{h}_0^l.$$

$(\tilde{x}_0^{l-1})^\top x_1^{l-1}/m$ is a scalar in the program by the Moment operation, and thus by the MatMul and NonLin operations, h_1^l is a vector in the program and thus so is $x_1^l = \sigma(h_1^l)$, which proves by induction that this is the case for any $l \in [2, L]$. By ZNonLin we thus have

$$Z^{h_1^l} = -\eta\overset{\circ}{\chi}_0\mathbb{E}[Z^{\tilde{x}_0^{l-1}}Z^{x_1^{l-1}}]Z^{d\tilde{h}_0^l}.$$

We then have by Theorem J.4 with $t = 1$,

$$f_1(\xi) = m^{-1}(U^{L+1})^\top x_1^L - \eta\chi_0 \frac{(\tilde{x}_0^L)^\top x_1^L}{m}$$

$U^{L+1} - \eta\chi_0\tilde{x}_0^L$ is a vector in the program by the NonLin operation, and the quantity $m^{-1}(U^{L+1} - \eta\chi_0\tilde{x}_0^L)^\top x_1^L$ is thus a scalar in the program by the Moment operation, and by the master theorem, we get $f_1(\xi) \rightarrow \mathbb{E}[Z^{U^{L+1}}Z^{x_1^L}] - \eta\overset{\circ}{\chi}_0\mathbb{E}[Z^{\tilde{x}_0^L}Z^{x_1^L}]$ almost surely, since both expectations are finite by Lemma L.1. Since we did the previous reasoning with an arbitrary ξ , we also get that $h_1^l(\xi_1), x_1^l(\xi_1)$ are vectors in the program for any $l \in [1, L]$ and that the formulas in (i), (ii), and (iii) hold when the input is ξ_1 . In particular, $f_1(\xi_1)$ converges to a finite almost sure limit $\overset{\circ}{f}_1(\xi_1)$, and thus the continuity of $\partial_2\ell(y_1, \cdot)$ ensures the almost sure convergence of χ_1 towards $\overset{\circ}{\chi}_1 := \partial_2\ell(y_1, \overset{\circ}{f}_1(\xi_1))$, which means χ_1 is a valid initial scalar in the Tensor Program. Then, dropping the dependency of the second forward pass (at $t = 1$) on ξ_1 , we get by Theorem J.5 with $t = 1$:

$$d\tilde{x}_1^L = U^{L+1} - \eta\chi_0\tilde{x}_0^L$$

which is a vector in the program by NonLin. Then $d\tilde{h}_1^L = d\tilde{x}_1^L \odot \sigma'(h_1^L)$ is also a vector in the program by NonLin since σ' is pseudo-Lipschitz. Let $l \in [2, L-1]$ and assume that $d\tilde{x}_1^{l+1}$ and $d\tilde{h}_1^{l+1}$ are vectors in the program. Then by Theorem J.5 with $t = 1$, we have

$$d\tilde{x}_1^l = -\eta\chi_0 \frac{(d\tilde{h}_0^{l+1})^\top d\tilde{h}_1^{l+1}}{m} \tilde{x}_0^l$$

$(d\tilde{h}_0^{l+1})^\top d\tilde{h}_1^{l+1}/m$ is a scalar in the program by the **Moment** operation and by **MatMul** and **NonLin** we thus get that $d\tilde{x}_1^l$ is a vector in the program. Then $d\tilde{h}_1^l = d\tilde{x}_1^l \odot \sigma'(h_1^l)$ is also a vector in the program since σ' is pseudo-Lipschitz, which concludes the induction and with it the proof. \square

Lemma J.7 (*Zs of HPZ and IP-LLR are equal at $t = 1$*). *Consider the HPZ and IP-LLR parameterization with an activation function σ satisfying Assumption 3, and no bias terms except at the first layer, and let us sub/super-script the variables of each models with HPZ and IP respectively. Let $\xi \in \mathbb{R}^d$ be an input to the networks, and assume that HPZ and IP-LLR share the same training samples (ξ_0, y_0) and (ξ_1, y_1) at $t = 0$ and $t = 1$, the same loss function ℓ satisfying Assumption 1, and the same base learning rate η . Then dropping the dependency of the first forward and backward passes on ξ_0 and that of the second forward passes on ξ , we have:*

- (i) $Z^{h^l, \text{HPZ}} = Z^{h^l, \text{IP}}, \quad Z^{x^l, \text{HPZ}} = Z^{x^l, \text{IP}}, \quad l \in [1, L],$
- (ii) $\lim_{m \rightarrow \infty} f_1^{\text{HPZ}}(\xi) = \lim_{m \rightarrow \infty} f_1^{\text{IP}}(\xi),$
- (iii) $\overset{\circ}{\chi}_1^{\text{HPZ}} = \overset{\circ}{\chi}_1^{\text{IP}},$
- (iv) $Z^{d\tilde{x}_1^l, \text{HPZ}} = Z^{d\tilde{x}_1^l, \text{IP}}, \quad Z^{d\tilde{h}_1^l, \text{HPZ}} = Z^{d\tilde{h}_1^l, \text{IP}}, \quad l \in [1, L].$

Proof. By Lemmas G.1 and G.1 we have $\overset{\circ}{\chi}_0^{\text{IP}} = \overset{\circ}{\chi}_0^{\mu\text{P}} = \overset{\circ}{\chi}_0^{\text{HPZ}} = \partial_2 \ell(y_0, 0)$, which we simply call $\overset{\circ}{\chi}_0$ in the remainder of this proof for simplicity. By Corollary G.6 and Lemma J.6 we have

$$Z^{h^1, \text{IP}} = Z^{U^1 \xi} + Z^{v^1} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{h}_0^1},$$

and

$$Z^{h^1, \text{HPZ}} = Z^{U^1 \xi} + Z^{v^1} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi + 1) Z^{d\tilde{h}_0^1},$$

and since the tilde variables are computed independently of any parameterization, we have $Z^{h^1, \text{HPZ}} = Z^{h^1, \text{IP}}$. Because IP and HPZ share the same activation function we also get $Z^{x^1, \text{HPZ}} = Z^{x^1, \text{IP}}$. Now let $l \in [2, L]$ and assume $Z^{h^{l-1}, \text{HPZ}} = Z^{h^{l-1}, \text{IP}}$ as well as $Z^{x^{l-1}, \text{HPZ}} = Z^{x^{l-1}, \text{IP}}$. By Corollary G.6 and Lemma J.6 we have

$$Z^{h^l, \text{IP}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_{1, \text{IP}}^{l-1}}] Z^{d\tilde{h}_0^l},$$

and

$$Z^{h^l, \text{HPZ}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_{1, \text{HPZ}}^{l-1}}] Z^{d\tilde{h}_0^l},$$

which shows $Z^{h^l, \text{IP}} = Z^{h^l, \text{HPZ}}$ since the tilde variables are independent of any choice of parameterization. Since the activation function σ is the same for both models we also get $Z^{x^l, \text{IP}} = Z^{x^l, \text{HPZ}}$ which concludes the induction. For the output of the networks, we have by Corollary G.6 and Lemma J.6

$$f_1^{\text{IP}}(\xi) \xrightarrow[m \rightarrow \infty]{\text{a.s.}} \mathbb{E}[Z^{U^{L+1}} Z^{x_{1, \text{IP}}^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_{1, \text{IP}}^L}],$$

and

$$f_1^{\text{HPZ}}(\xi) \xrightarrow[m \rightarrow \infty]{\text{a.s.}} \mathbb{E}[Z^{U^{L+1}} Z^{x_{1, \text{HPZ}}^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_{1, \text{HPZ}}^L}],$$

and since $Z^{x_{1,\text{IP}}^L} = Z^{x_{1,\text{HPZ}}^L}$ by the previous induction and the tilde variables are independent of the parameterization, we get $\lim_{m \rightarrow \infty} f_1^{\text{HPZ}}(\xi) = \lim_{m \rightarrow \infty} f_1^{\text{IP}}(\xi) =: \overset{\circ}{f}_1(\xi)$. Since $\chi_1^{\text{IP}} = \partial_2(y_1, f_1^{\text{IP}}(\xi))$ and $\chi_1^{\text{HPZ}} = \partial_2(y_1, f_1^{\text{HPZ}}(\xi))$, by continuity of $\partial_2 \ell(y_1, \cdot)$ we get that $\overset{\circ}{\chi}_1^{\text{HPZ}} = \partial_2 \ell(y_1, \overset{\circ}{f}_1(\xi)) = \overset{\circ}{\chi}_1^{\text{IP}}$.

For the backward pass, we have by Lemma J.5 that $d\tilde{x}_{1,\text{HPZ}}^L = U^{L+1} - \eta\chi_{0,\text{HPZ}}\tilde{x}_0^L$ which gives by NonLin $Z^{d\tilde{x}_{1,\text{HPZ}}^L} = ZU^{L+1} - \eta\overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L}$ which is also equal to $Z^{d\tilde{x}_{1,\text{IP}}^L}$ by Lemma G.5 since the tilde variables are independent of the choice of parameterization. Then, we also get $Z^{d\tilde{h}_{1,\text{HPZ}}^L} = Z^{d\tilde{x}_{1,\text{HPZ}}^L} \sigma'(Z^{h_{1,\text{HPZ}}^L})$ and $Z^{d\tilde{h}_{1,\text{IP}}^L} = Z^{d\tilde{x}_{1,\text{HPZ}}^L} \sigma'(Z^{h_{1,\text{IP}}^L})$ which shows $Z^{d\tilde{h}_{1,\text{HPZ}}^L} = Z^{d\tilde{h}_{1,\text{IP}}^L}$. Let $l \in [1, L-1]$ and assume $Z^{d\tilde{x}_{1,\text{HPZ}}^{l+1}} = Z^{d\tilde{x}_{1,\text{IP}}^{l+1}}$ as well as $Z^{d\tilde{h}_{1,\text{HPZ}}^{l+1}} = Z^{d\tilde{h}_{1,\text{IP}}^{l+1}}$. By Lemma J.5, we have

$$d\tilde{x}_{1,\text{HPZ}}^l = -\eta\chi_0 \frac{(d\tilde{h}_0^{l+1})^\top d\tilde{h}_{1,\text{HPZ}}^{l+1}}{m} \tilde{x}_0^l$$

which gives by the master theorem and the ZNonLin

$$Z^{d\tilde{x}_{1,\text{HPZ}}^l} = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^{l+1}} Z^{d\tilde{h}_{1,\text{HPZ}}^{l+1}}] Z^{\tilde{x}_0^l}$$

which is the same expression as $Z^{d\tilde{x}_{1,\text{IP}}^l}$ by Lemma G.5. It then follows that $Z^{d\tilde{h}_{1,\text{HPZ}}^l} = Z^{d\tilde{h}_{1,\text{IP}}^l}$, which concludes the induction and with it the proof. \square

Theorem J.8 (Z for the forward pass of HPZ at time t). *Consider the HPZ parameterization with an activation function σ satisfying Assumption 2 and no bias terms except at the first layer. Let $\xi \in \mathbb{R}^d$ be an input to the network. Then, for any $l \in [1, L]$, $h_s^l(\xi)$, $x_s^l(\xi)$, $d\tilde{x}_s^l$, $d\tilde{h}_s^l$ are vectors in the program, $f_s(\xi)$ is a scalar in the program, and χ_s is a valid initial scalar in the program. Additionally, dropping the dependency of the forward pass at time t on ξ , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

- (i) $Z^{h_t^l} = Z^{W^l(t)\xi + B^l(t)} = Z^{U^l\xi + Z^{v^l}} - \eta\overset{\circ}{\chi}_0(\xi_0^\top\xi + 1)Z^{d\tilde{h}_0^l} - \eta\left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s(\xi_s^\top\xi + 1)Z^{d\tilde{h}_s^l}\right)$,
- (ii) $Z^{h_t^l} = Z^{W^l(t)x_t^{l-1}} = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_0^l} - \eta\left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_s^l}\right)$, $l \in [2, L]$,
- (iii) $f_t(\xi) = (W^{L+1}(t))^\top x_t^L \xrightarrow[m \rightarrow \infty]{a.s.} \mathbb{E}[Z^{U^{L+1}} Z^{x_t^L}] - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^L} Z^{x_t^L}] - \eta\left(\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_t^L}]\right)$.

Proof. The proof is exactly the same as for Theorem G.7 except that whenever a multiplication by $W^l(0)$ appears with $l \in [2, L]$, it is now replaced by 0, but the reasoning and all the arguments are the same, which in summary uses an induction over t as well as the master theorem and the ZNonLin rule from the Tensor Program. \square

Theorem J.9 (Zs of backward pass of HPZ at time t). *Consider the HPZ parameterization with an activation function σ satisfying Assumption 2 and no bias terms except at the first layer. Then, for any $t \geq 1$, dropping the dependency of the forward pass at time t on ξ_t , and of the previous forward and backward passes on the corresponding ξ_s , one has:*

- (i) $Z^{d\tilde{x}_t^L} = Z^{W^{L+1}(t)} = Z^{U^{L+1}} - \eta\overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L} - \eta\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s Z^{x_s^L}$,
- (ii) $Z^{d\tilde{x}_t^{l-1}} = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_t^l}] Z^{\tilde{x}_0^{l-1}} - \eta\sum_{s=1}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{d\tilde{h}_s^l} Z^{d\tilde{h}_t^l}] Z^{x_s^{l-1}}$, $l \in [2, L]$.

Proof. As for Theorem J.8, the proof follows exactly the same pattern as for Theorem J.9 except that whenever a multiplication by $W^l(0)$ appears with $l \in [2, L]$, it is now replaced by 0. \square

J.2.2 Induction on t

Lemma J.10 (Induction step on the Z s of the forward pass). *Consider the IP-LLR and HPZ parameterizations with an activation function σ satisfying Assumption 3 and no bias terms except at the first layer, and let us sub/super-script the variables of each models with IP and HP respectively. Let $s \geq 1$, $\xi \in \mathbb{R}^d$ be an input to the networks, and assume that the training routine (see Definition 2.2) is the same for both models with a loss satisfying Assumption 1. Assume further that, dropping the dependency of the forward and backward passes at time $t = r$ on ξ_r , for all $r \in [1, s]$, we have:*

- (i) $Z^{h^l_{\text{HPZ},r}} = Z^{h^l_{\text{IP},r}}, \quad Z^{x^l_{\text{HPZ},r}} = Z^{x^l_{\text{IP},r}}, \quad l \in [1, L],$
- (ii) $\lim_{m \rightarrow \infty} f_r^{\text{HPZ}}(\xi) = \lim_{m \rightarrow \infty} f_r^{\text{IP}}(\xi),$
- (iii) $\overset{\circ}{\chi}_r^{\text{HPZ}} = \overset{\circ}{\chi}_r^{\text{IP}},$
- (iv) $Z^{d\tilde{h}^l_{\text{HPZ},r}} = Z^{d\tilde{h}^l_{\text{IP},r}}, \quad Z^{d\tilde{x}^l_{\text{HPZ},r}} = Z^{d\tilde{x}^l_{\text{IP},r}}, \quad l \in [1, L].$

Then, dropping the dependency of the forward pass at time $t = s + 1$ on ξ , one has:

- (v) $Z^{h^l_{\text{HPZ},s+1}} = Z^{h^l_{\text{IP},s+1}}, \quad Z^{x^l_{\text{HPZ},s+1}} = Z^{x^l_{\text{IP},s+1}}, \quad l \in [1, L],$
- (vi) $\lim_{m \rightarrow \infty} f_{s+1}^{\text{HPZ}}(\xi) = \lim_{m \rightarrow \infty} f_{s+1}^{\text{IP}}(\xi),$
- (vii) $\overset{\circ}{\chi}_{s+1}^{\text{HPZ}} = \overset{\circ}{\chi}_{s+1}^{\text{IP}},$
- (viii) $Z^{d\tilde{h}^l_{\text{HPZ},s+1}} = Z^{d\tilde{h}^l_{\text{IP},s+1}}, \quad Z^{d\tilde{x}^l_{\text{HPZ},s+1}} = Z^{d\tilde{x}^l_{\text{IP},s+1}}, \quad l \in [1, L].$

Proof. Since by assumption, for any $r \in [1, s]$, the Z s of the forward and backward passes are equal for both parameterizations, we drop the dependency of those quantities on the model, and for $z \in \{h_r^l, x_r^l, d\tilde{h}_r^l, d\tilde{x}_r^l\}$, we simply call $Z^{z_{\text{HPZ}}} = Z^{z_{\text{IP}}} = Z^z$. Similarly we simply call $\overset{\circ}{\chi}_r^{\text{HPZ}} = \overset{\circ}{\chi}_r^{\text{IP}} = \overset{\circ}{\chi}_r$. We have G.7

$$\begin{aligned} Z^{h^1_{\text{HPZ},s+1}} &= Z^{U^1 \xi_{s+1}} - \eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_{s+1} + 1) Z^{d\tilde{h}_0^1} - \eta \sum_{r=1}^s \overset{\circ}{\chi}_r (\xi_r^\top \xi_{s+1} + 1) Z^{d\tilde{h}_r^1} \\ &= Z^{h^1_{\text{IP},s+1}} \end{aligned}$$

where the first equality stems from Theorem J.8 and the second one from Theorem G.7. Since both parameterizations use the same linearity σ , we get $Z^{x^1_{\text{HPZ},s+1}} = \sigma(Z^{x^1_{\text{HPZ},s+1}}) = \sigma(Z^{h^1_{\text{IP},s+1}}) = Z^{x^1_{\text{IP},s+1}}$.

Let $l \in [2, L]$ and assume $Z^{h^{l-1}_{\text{HPZ},s+1}} = Z^{h^{l-1}_{\text{IP},s+1}}, Z^{x^{l-1}_{\text{HPZ},s+1}} = Z^{x^{l-1}_{\text{IP},s+1}}$. By Theorem J.8, we have

$$\begin{aligned} Z^{h^l_{\text{HPZ},s+1}} &= -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x^{l-1}_{\text{HPZ},s+1}}] Z^{d\tilde{h}_0^l} - \eta \sum_{r=1}^s \overset{\circ}{\chi}_r \mathbb{E}[Z^{x_r^{l-1}} Z^{x^{l-1}_{\text{HPZ},s+1}}] Z^{d\tilde{h}_r^l} \\ &= -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x^{l-1}_{\text{IP},s+1}}] Z^{d\tilde{h}_0^l} - \eta \sum_{r=1}^s \overset{\circ}{\chi}_r \mathbb{E}[Z^{x_r^{l-1}} Z^{x^{l-1}_{\text{IP},s+1}}] Z^{d\tilde{h}_r^l} \\ &= Z^{h^l_{\text{HPZ},s+1}} \end{aligned}$$

where the last equality stems from Theorem G.7. Since both parameterizations used the same non-linearity σ , we get $Z^{x^{l+1}_{\text{HPZ},s+1}} = Z^{x^{l+1}_{\text{IP},s+1}}$.

By induction, we thus get that for any $l \in [1, L]$, $Z^{h_{\text{HPZ},s+1}^l} = Z^{h_{\text{IP},s+1}^l}$, and $Z^{x_{\text{HPZ},s+1}^l} = Z^{x_{\text{IP},s+1}^l}$, which proves (v). We can thus drop the dependency of h_{s+1}^l and x_{s+1}^l on the model HPZ or IP. Now, we thus have by [Theorem J.8](#)

$$\begin{aligned} \lim_{m \rightarrow \infty} f_{s+1}^{\text{HPZ}}(\xi) &= \mathbb{E}[Z^{U^{L+1}} Z^{x_{s+1}^L}] - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\bar{x}_0^L} Z^{x_{s+1}^L}] - \eta \left(\sum_{r=1}^s \overset{\circ}{\chi}_r \mathbb{E}[Z^{x_r^L} Z^{x_{s+1}^L}] \right) \\ &= \lim_{m \rightarrow \infty} f_{s+1}^{\text{IP}}(\xi) \end{aligned}$$

where the last equality stems from [Theorem G.7](#), which proves (vi). Then (vi) combined with the continuity of $\partial_2 \ell(y_{s+1}, \cdot)$ proves (vii), and we can thus imply denote $\overset{\circ}{\chi}_{s+1}^{\text{IP}} = \overset{\circ}{\chi}_{s+1}^{\text{HPZ}} = \overset{\circ}{\chi}_{s+1}$. By [Theorems J.9](#) and [G.8](#) we get $Z^{d\bar{x}_{s+1, \text{HPZ}}^L} = Z^{d\bar{x}_{s+1, \text{IP}}^L}$, from which it follows that $Z^{dh_{s+1, \text{HPZ}}^L} = Z^{dh_{s+1, \text{IP}}^L}$ by (v) and since both models share the same activation function. Finally, given the previous result, with (i), (iii), (iv), (v) and (vii), an easy induction gives (viii) with the formulas of [Theorems J.9](#) and [G.8](#), which concludes the proof. \square

J.2.3 Proof of [Theorem 4.2](#)

Proof. The claim has already been proved at $t = 0$ by [Lemmas G.1](#) and [J.2](#), and at $t = 1$ by [Lemma J.7](#). Then, by [Lemma J.10](#), we get the result at any time step $t \geq 1$ by induction. \square

K Formal versions of the results for the alternative methods to escape the initial stationary point

K.1 Formal version of [Theorem 5.2](#)

Theorem K.1 (Formal). *Consider the IP-bias parameterization as in [Equations \(5.2\)](#), with the initial learning rates as in [Equations \(5.4\)](#) and [\(5.5\)](#). Assume the activation function σ satisfies [Assumption 2](#) and the loss ℓ satisfies [Assumption 1](#). Then, for any input $\xi \in \mathbb{R}^d$ to the network, $Z^{h_0^l(\xi)}$, $Z^{x_0^l(\xi)}$ for $l \geq 2$, and $\lim_{m \rightarrow \infty} f_0(\xi)$ do not depend on ξ . In addition, for any vector x in the program such that Z^x does not depend on on the first training input ξ_0 , $Z^{\Delta W^{l(1)}x}$ for $l \in [3, L]$, and $\lim_{m \rightarrow \infty} (\Delta W^{L+1}(1))^\top x$ do not depend on ξ_0 .*

Proof. We have $h_0^1 = U^1 \xi + v^1$ so that $Z^{h_0^1} = \widehat{Z}^{U^1 \xi} + \widehat{Z}^{v^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$, and $Z^{x_0^1} = \sigma(Z^{h_0^1})$. At the second layer $l = 2$, we have $h_0^2 = m^{-1/2} \widehat{W}^2 x_0^1 + v^2$ so that by [ZNonLin](#) $Z^{h_0^2} = 0 \times \widehat{Z}^{\widehat{W}^2 x_0^1} + Z^{v^2}$, and $\widehat{Z}^{\widehat{W}^2 x_0^1} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_0^1})^2])$. Because σ is pseudo-Lipschitz, it is also polynomially bounded, and the variance of the Gaussian is finite by [Lemma C.1](#), so that $Z^{h_0^2} = Z^{v^2} \sim \mathcal{N}(0, 1)$ which does not depend on ξ . Therefore, $Z^{x_0^2} = \sigma(Z^{h_0^2})$ also does not depend on ξ . Let $l \in [3, L]$ and assume that $Z^{h_0^{l-1}} = Z^{v^{l-1}}$ and $Z^{x_0^{l-1}} = \sigma(Z^{v^{l-1}})$. Then, we have $Z^{h_0^l} = 0 \times \widehat{Z}^{\widehat{W}^l x_0^{l-1}} + Z^{v^l}$, and $\widehat{Z}^{\widehat{W}^l x_0^{l-1}} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_0^{l-1}})^2])$, and the variance is again finite by the same arguments as for $l = 2$. We thus get $Z^{h_0^l} = Z^{v^l}$ and $Z^{x_0^l} = \sigma(Z^{h_0^l}) = \sigma(Z^{v^l})$ which concludes the induction and shows that $Z^{h_0^l}$ and $Z^{x_0^l}$ do not depend on ξ for all intermediate layers l .

For the output of the network, we have by the master theorem that $m^{-1}(U^{L+1})^\top x_0^L$ converges almost surely to $\mathbb{E}[Z^{U^{L+1}} \sigma(Z^{v^L})] = 0$ since $Z^{U^{L+1}}$ has mean 0 and is independent of Z^{v^L} . Since $f_0(\xi) = m^{-1}(U^{L+1})^\top x_0^L + v^{L+1}$ where $v^{L+1} \sim \mathcal{N}(0, 1)$, we have that $f_0(\xi)$ converges almost surely to the Gaussian variable v^{L+1} which does not depend on ξ . For the backward pass, recall that we call $d\bar{x}_0^l := m^{-1} m^{-(L-l)/2} \nabla_{x^l} f_0(\xi_0)$ and $d\bar{h}_0^l := m^{-1} m^{-(L-l)/2} \nabla_{h^l} f_0(\xi_0)$.

Then, we have $d\tilde{x}_0^L = U^{L+1}$, $d\tilde{h}_0^L = U^{L+1} \odot \sigma'(h_0^L)$, and a simple induction shows that for any $l \in [1, L-1]$, $d\tilde{x}_0^l = (\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}$, $d\tilde{h}_0^l = d\tilde{x}_0^l \odot \sigma'(h_0^l)$. We thus have $Z^{d\tilde{x}_0^l} = Z^{U^{L+1}}$ and $Z^{d\tilde{h}_0^l} = Z^{U^{L+1}} \sigma'(Z^{v^L})$ which does not depend on the first training input ξ_0 . With the recursive formulas above, and since $Z^{h_0^l} = Z^{v^l}$ for $l \in [2, L]$, it is clear that $Z^{d\tilde{x}_0^l}$ and $Z^{d\tilde{h}_0^l}$ do not depend on ξ_0 for $l \in [2, L]$.

Finally, let x be a vector in the program for which Z^x does not depend on ξ , and let $l \in [3, L]$. Then, by design, with the initial learning rates of Equation (5.4) for the weights with IP-bias, we have

$$\Delta W^l(1)x = -\eta\chi_0 \frac{(x_0^{l-1})^\top x}{m} d\tilde{h}_0^l,$$

so that by ZNonLin

$$Z^{\Delta W^l(1)x} = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{x_0^{l-1}} Z^x] Z^{d\tilde{h}_0^l},$$

where $\overset{\circ}{\chi}_0 := \partial_2 \ell(y_0, v^{L+1})$. Since v^{L+1} , $Z^{x_0^{l-1}}$, Z^x and $Z^{d\tilde{h}_0^l}$ do not depend on ξ_0 ($l-1$ and l are both in $[2, L]$), $Z^{\Delta W^l(1)x}$ also does not depend on the first training input ξ_0 . To conclude, we have by the master theorem that $(\Delta W^{L+1}(1))^\top x$ converges almost surely towards $-\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{x_0^L} Z^x]$ which does not depend on ξ_0 since this is the case for v^{L+1} , $Z^{x_0^L}$ and Z^x , which concludes the proof. \square

K.2 Formal version of Theorem 5.1

Theorem K.2 (Formal). *Consider IP-non-centered with the Naive-IP learning rates at every time step. Assume the activation function σ satisfies Assumption 2 and the loss ℓ satisfies Assumption 1, and let $t \geq 0$ and $\xi \in \mathbb{R}^d$ be an input to the network. Then, calling $d\tilde{x}_s^l := m\nabla_{x^l} f_s(\xi_s)$ and $d\tilde{h}_s^l := m\nabla_{h^l} f_s(\xi_s)$, one has that:*

(i) for any $l \in [2, L-1]$, Z^{h^l} and Z^{x^l} are deterministic constants,

(ii) for any $l \in [2, L-1]$, $Z^{d\tilde{x}_t^l}$ and $Z^{d\tilde{h}_t^l}$ deterministic constants,

(iii) for any $l \in [3, L-1]$, and for any vector x in the program, we have that

$$Z^{(W^l(t+1) - W^l(0))x} = \left(-\eta \sum_{s=0}^t \overset{\circ}{\chi}_s Z^{d\tilde{h}_s^l} Z^{x^{l-1}} \right) \mathbb{E}[Z^x].$$

Remark K.1. Point (iii) highlights the fact that in the infinite-width limit the (random) matrix operator $(w^l(t) - w^l(0))$ acts on a vector x as if all the entries of the matrix operator were equal to a single deterministic constant $\left(-\eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s Z^{d\tilde{h}_s^l} Z^{x^{l-1}} \right)$, because then the averages over the coordinates of x involved in $(W^l(t) - W^l(0))x$ would simply yield $\mathbb{E}[Z^x]$ by the master theorem of the Tensor Program.

The proof Theorem K.2 can be found in Appendix K.2.2. The proof is done by inducting over t , and we present the case $t = 0$ and the induction step first in Appendix K.2.1.

K.2.1 Preliminaries

Lemma K.3 (First forward-backward pass and weight updates). *Claims (i), (ii) and (iii) of Theorem K.2 hold at $t = 0$.*

Proof. h_0^1 and $x_0^1 = \sigma(h_0^1)$ are vectors in the program by the **MatMul** and **NonLin** rules since σ is pseudo-Lipschitz by assumption, and $Z^{h_0^1} = Z^{U^1\xi} + Z^{v^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$, and finally $Z^{x_0^1} = \sigma(Z^{h_0^1})$. Now, we have (recall that as defined in Section 5.1 J is the matrix full of ones)

$$h_0^2 = m^{-1/2}\widehat{W}^2x_0^1 + m^{-1}v^2 + u_2m^{-1}Jx_0^1.$$

$m^{-1/2}\widehat{W}^2x_0^1 + m^{-1}v^2$ is a valid vector in the program by **MatMul** and **NonLin** because the initial scalars $m^{-1/2}$ and m^{-1} converge to 0 almost surely, and $Z^{m^{-1/2}\widehat{W}^2x_0^1 + m^{-1}v^2} = 0 \times \widehat{Z}^{\widehat{W}^2x_0^1} + 0 \times \widehat{Z}^{v^2}$. By the **ZHat** rule we get that $\widehat{Z}^{\widehat{W}^2x_0^1} \sim \mathcal{N}(0, \mathbb{E}[(Z^{x_0^1})^2])$, with finite variance by Lemma C.1 since σ is pseudo-Lipschitz and thus polynomially bounded, and $\widehat{Z}^{v^2} \sim \mathcal{N}(0, 1)$. We thus get $Z^{m^{-1/2}(\widehat{W}^2x_0^1 + v^2)} = 0$. On the other hand, $\theta := (1/m)\sum_{q=1}^m x_{0,q}^1$ is a valid scalar in the program by the **Moment** rule and it converges almost surely to $\overset{\circ}{\theta} = \mathbb{E}[Z^{x_0^1}]$ by the master theorem. The coordinates of $u_2m^{-1}Jx_0^1$ are thus all equal to $u_2\theta$, and the vector $u_2m^{-1}Jx_0^1$ is thus equal to $\psi(x_0^1; \theta)$ coordinate-wise where the function $\psi(\cdot; \cdot) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is pseudo-Lipschitz and depends **only** on the second variable with $\psi(x; \alpha) = u_2\alpha$. By the **NonLin** rule $u_2m^{-1}Jx_0^1$ is thus a vector in the program and by **ZNonLin** we thus get $Z^{u_2m^{-1}Jx_0^1} = \psi(Z^{x_0^1}; \overset{\circ}{\theta}) = u_2\mathbb{E}[Z^{x_0^1}]$. We thus finally get

$$Z^{h_0^2} = u_2\mathbb{E}[Z^{x_0^1}],$$

which is a (finite) deterministic constant. Then the same statement holds for $Z^{x_0^2} = \sigma(u_2\mathbb{E}[Z^{x_0^1}])$. Let $l \in [3, L]$ and assume that h_0^{l-1} and x_0^{l-1} are vectors in the program and that $Z^{h_0^{l-1}}$ and $Z^{x_0^{l-1}}$ are deterministic constants. Then, we have

$$h_0^l = m^{-1/2}\widehat{W}^lx_0^{l-1} + m^{-1}v^l + u_lm^{-1}Jx_0^{l-1}$$

As for the case $l = 2$, we get that $m^{-1/2}\widehat{W}^lx_0^{l-1} + m^{-1}v^l$ is a vector in the program with $Z^{m^{-1/2}\widehat{W}^lx_0^{l-1} + m^{-1}v^l} = 0$, and $u_lm^{-1}Jx_0^{l-1} = \psi(x_0^{l-1}; \theta)$ is a vector in the program with $\psi(z; \alpha) = u_l\alpha$ (recall that ψ is taken coordinate-wise) depending only on the second variable and $\theta := (1/m)\sum_{q=1}^m x_{0,q}^{l-1}$ is a valid scalar in the program by the **Moment** rule, which, by the master theorem, converges almost surely towards $\overset{\circ}{\theta} = \mathbb{E}[Z^{x_0^{l-1}}] = Z^{x_0^{l-1}}$ since the latter is a deterministic constant by the induction hypothesis. By **NonLin** h_0^l is a vector in the program and by **ZNonLin** $Z^{h_0^l} = \psi(Z^{x_0^{l-1}}; \overset{\circ}{\theta}) = u_lZ^{x_0^{l-1}}$ which is a deterministic constant. The same claim holds for $Z^{x_0^l} = \sigma(u_lZ^{x_0^{l-1}})$, which concludes the induction for the forward pass. For the backward pass we get $d\tilde{x}_0^L = w^{L+1}(0) = U^{L+1} + u_{L+1}\mathbf{1}$ so that by **ZNonLin** $Z^{d\tilde{x}_0^L} = Z^{U^{L+1}} + u_{L+1} \sim \mathcal{N}(u_{L+1}, 1)$ since u_{L+1} is a valid initial scalar in the program as it converges almost surely to u_{L+1} . We then have $Z^{d\tilde{h}_0^L} = Z^{d\tilde{x}_0^L}\sigma'(Z^{h_0^L})$. Note that both $Z^{d\tilde{x}_0^L}$ and $Z^{d\tilde{h}_0^L}$ are not deterministic constants because U^{L+1} is Gaussian with variance 1. We then have:

$$d\tilde{x}_0^{L-1} = m^{-1/2}(\widehat{W}^L)^\top d\tilde{h}_0^L + u_Lm^{-1}J^\top d\tilde{h}_0^L$$

As usual the first term $m^{-1/2}(\widehat{W}^L)^\top d\tilde{h}_0^L$ is a vector in the program by **MatMul** and **NonLin** and $Z^{m^{-1/2}(\widehat{W}^L)^\top d\tilde{h}_0^L} = 0$. For the second term, since $J^\top = J$, $m^{-1}J^\top d\tilde{h}_0^L$ is also a vector in the program and $Z^{m^{-1}J^\top d\tilde{h}_0^L} = u_L\mathbb{E}[Z^{d\tilde{h}_0^L}]$. We thus get that $d\tilde{x}_0^{L-1}$ is a vector in the program with $Z^{d\tilde{x}_0^{L-1}} = u_L\mathbb{E}[Z^{d\tilde{h}_0^L}]$ which is a deterministic constant. Then, $d\tilde{h}_0^{L-1}$ is also a vector in the program and by **ZNonLin** $Z^{d\tilde{h}_0^{L-1}} = Z^{d\tilde{x}_0^{L-1}}\sigma'(Z^{h_0^L})$ is a deterministic constant. Repeating

the reasoning above at any layer $l \in [2, L - 1]$, an easy induction (as in the forward pass) shows that $d\tilde{x}_0^l$ and $d\tilde{h}_0^l$ are vectors in the program and that $Z^{d\tilde{x}_0^l}$ and $Z^{d\tilde{h}_0^l}$ are deterministic constants. Note that $Z^{d\tilde{x}_0^1} = u_2 \mathbb{E}[Z^{d\tilde{h}_0^2}] = u_2 Z^{d\tilde{h}_0^2}$ is also a deterministic constant but that $Z^{d\tilde{h}_0^1} = Z^{d\tilde{x}_0^1} \sigma'(Z^{h_0^1})$ is not because $Z^{h_0^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$. Let $l \in [3, L - 1]$, and let x be a vector in the program. With the Naive-IP learning rates, we have

$$\Delta W^l(1) = -\eta \chi_0 \frac{d\tilde{h}_0^l(x_0^{l-1})^\top}{m}$$

Since $l \in [3, L - 1]$, $Z^{d\tilde{h}_0^l}$ is a deterministic constant, and since $l - 1 \in [2, L - 2]$, $Z^{x_0^{l-1}}$ is also a deterministic constant. By `ZNonLin` and `ZMoment` we get

$$\begin{aligned} Z^{\Delta W^l(1)x} &= -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{x_0^{l-1}} Z^x] Z^{d\tilde{h}_0^l} \\ &= -\eta \overset{\circ}{\chi}_0 Z^{d\tilde{h}_0^l} Z^{x_0^{l-1}} \mathbb{E}[Z^x] \end{aligned}$$

which concludes the proof. Note that χ_0 is a valid initial scalar in the program because $f_0(\xi_0) = m^{-1}(U^{L+1})^\top x_0^L + u_{L+1} m^{-1} \mathbf{1}^\top x_0^L$ converges almost surely, by the master theorem, to $\mathbb{E}[Z^{U^{L+1}} Z^{x_0^L}] + u_{L+1} \mathbb{E}[Z^{x_0^L}] = u_{L+1} Z^{x_0^L}$ since $Z^{x_0^L}$ is a deterministic constant and $Z^{U^{L+1}} \sim \mathcal{N}(0, 1)$ has mean zero. Since $\partial_2 \ell(y_0, \cdot)$ is continuous by assumption, χ_0 converges almost surely towards $\overset{\circ}{\chi}_0 := \partial_2 \ell(y_0, u_{L+1} Z^{x_0^L})$. \square

Lemma K.4 (Induction step at time $t \geq 1$). *Let $t \geq 1$ and assume claims (i), (ii) and (iii) of [Theorem K.2](#) hold at all time steps $s \in [0, t - 1]$. Then claims (i), (ii) and (iii) also hold at time step t .*

Proof. With the Naive-IP learning rate exponents, we get that for any $t \geq 1$,

$$\begin{aligned} W^1(t) &= U^1 - \eta \sum_{s=0}^{t-1} \chi_s d\tilde{h}_s^1 \xi_s^\top, \\ B^1(t) &= v^1 - \eta \sum_{s=0}^{t-1} \chi_s d\tilde{h}_s^1, \\ W^l(t) &= m^{-1}(U^l + u_l J) - \eta \sum_{s=0}^{t-1} \chi_s \frac{d\tilde{h}_s^l(x_s^{l-1})^\top}{m}, & l \in [2, L], \\ B^l(t) &= m^{-1}v^l - \eta m^{-1} \sum_{s=0}^{t-1} \chi_s d\tilde{h}_s^l, & l \in [2, L], \\ W^{L+1}(t) &= m^{-1}(U^{L+1} + u_{L+1}) - \eta \sum_{s=0}^{t-1} \chi_s \frac{x_s^L}{m}, \\ B^{L+1}(t) &= m^{-1}v^{L+1} - \eta m^{-1} \sum_{s=0}^{t-1} \chi_s. \end{aligned}$$

By a simple induction, all the h_s^l, x_s^l and $d\tilde{x}_s^l, d\tilde{h}_s^l$ are part of and the scalars χ_s are valid scalars in the program which have a constant almost sure limit, and by `ZNonLin` we get:

$$Z^{h_t^1} = Z^{U^1 \xi} + Z^{v^1} - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s (\xi_s^\top \xi + 1) Z^{d\tilde{h}_s^1}$$

and $Z^{x_t^1} = \sigma(Z^{h_t^1})$ is not a deterministic constant because $Z^{U^1\xi} + Z^{v^1} \sim \mathcal{N}(0, \|\xi\|^2 + 1)$. Let $l \in [2, L - 1]$. We have

$$\begin{aligned} Z^{h_t^l} &= 0 \times Z^{\widehat{W}^l x_t^{l-1}} + 0 \times Z^{v^l} + u_l \mathbb{E}[Z^{x_t^{l-1}}] - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_s^l} \\ &= u_l \mathbb{E}[Z^{x_t^{l-1}}] - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] Z^{d\tilde{h}_s^l}, \end{aligned}$$

which is a deterministic constant with the assumption on the $Z^{d\tilde{h}_s^l}$ since $l \in [2, L - 1]$. Note that if $l \in [3, L - 1]$, we even have that the expectations simplify and we get $Z^{h_t^l} = (u_l - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s Z^{x_s^{l-1}} Z^{d\tilde{h}_s^l}) Z^{x_t^{l-1}}$. In any case, $Z^{x_t^l} = \sigma(Z^{h_t^l})$ is also a deterministic constant. For the output of the network, we have

$$f_t(\xi) = \frac{(U^{L+1})^\top x_t^L}{m} + u_{L+1} \frac{\mathbf{1}^\top x_t^L}{m} + m^{-1}(v^{L+1} - \eta \sum_{s=0}^{t-1} \chi_s) - \eta \sum_{s=0}^{t-1} \chi_s \frac{(x_s^L)^\top x_t^L}{m}$$

so that even if the x_s^L are not deterministic, $f_t(\xi)$ still converges almost surely, by the master theorem, to $\mathbb{E}[(Z^{U^{L+1}} + u_{L+1}) Z^{x_t^L}] - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_t^L}]$, and since $\partial_2 \ell(y_t, \cdot)$ is continuous by assumption, χ_t converges almost surely towards the constant $\partial_2 \ell(y_t, \mathbb{E}[(Z^{U^{L+1}} + u_{L+1}) Z^{x_t^L}] - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^L} Z^{x_t^L}])$. For the backward pass, we get:

$$Z^{d\tilde{x}_t^L} = Z^{w^{L+1}(t)} = Z^{U^{L+1}} + u_{L+1} - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s Z^{x_s^L}$$

and $Z^{d\tilde{h}_t^L} = Z^{d\tilde{x}_t^L} \sigma'(Z^{h_t^L})$. Let $l \in [2, L - 1]$, we have

$$\begin{aligned} d\tilde{x}_t^l &= (W^{l+1}(t))^\top d\tilde{h}_t^{l+1} \\ &= m^{-1/2} (\widehat{W}^{l+1})^\top d\tilde{h}_t^{l+1} + m^{-1} u_{l+1} J d\tilde{h}_t^{l+1} - \eta \sum_{s=0}^{t-1} \chi_s \frac{(d\tilde{h}_s^{l+1})^\top d\tilde{h}_t^{l+1}}{m} x_s^l, \end{aligned}$$

so that by ZNonLin we get

$$Z^{d\tilde{x}_t^l} = u_{l+1} \mathbb{E}[Z^{d\tilde{h}_t^{l+1}}] - \eta \sum_{s=0}^{t-1} \overset{\circ}{\chi}_s \mathbb{E}[Z^{d\tilde{h}_s^{l+1}} Z^{d\tilde{h}_t^{l+1}}] Z^{x_s^l},$$

and since $l \in [2, L - 1]$, $Z^{x_s^l}$ is a deterministic constant and thus so is $Z^{d\tilde{x}_t^l}$. Then, $Z^{d\tilde{h}_t^l} = Z^{d\tilde{x}_t^l} \sigma'(Z^{h_t^l})$ and since $l \in [2, L - 1]$, $Z^{h_t^l}$ is a deterministic constant. Finally, let $l \in [3, L - 1]$, and let x be a vector in the program. We have

$$(W^l(t+1) - W^l(0))x = -\eta \sum_{s=0}^t \chi_s \frac{(x_s^{l-1})^\top x}{m} d\tilde{h}_s^l,$$

and by ZNonLin

$$\begin{aligned} Z^{(W^l(t+1) - W^l(0))x} &= -\eta \sum_{s=0}^t \overset{\circ}{\chi}_s \mathbb{E}[Z^{x_s^{l-1}} Z^x] Z^{d\tilde{h}_s^l} \\ &= \left(-\eta \sum_{s=0}^t \overset{\circ}{\chi}_s Z^{d\tilde{h}_s^l} Z^{x_s^{l-1}} \right) \mathbb{E}[Z^x], \end{aligned}$$

where the last equality stems from the fact that since $l \in [3, L - 1]$, $l - 1 \in [2, L - 2]$ and $Z^{x_s^{l-1}}$ is a deterministic constant for any $s \in [0, t]$. Since $l \in [2, L - 1]$, $Z^{d\tilde{h}_s^l}$ is also a deterministic constant, so that $-\eta \sum_{s=0}^t \overset{\circ}{\chi}_s Z^{d\tilde{h}_s^l} Z^{x_s^{l-1}}$ is a deterministic constant, which concludes the proof. \square

K.2.2 Proof of Theorem K.2

Proof. The result simply comes by induction over t using Lemmas K.3 and K.4. \square

L The variables associated with the initial weights vanish in IP-LLR

In this section we wish to study more precisely the evolution and the expression of the variables Z in the dynamics of IP-LLR at any time step t . To this end, we will show that the Z s of all the forward and backward variables in IP-LLR are functions only of the $\widehat{Z}^{\widehat{W}^l \bar{x}_0^{l-1}}$ and $\widehat{Z}^{(\widehat{W}^l)^\top d\bar{h}_0^l}$, as well as the initial vectors $U^1 \xi_0, \dots, U^1 \xi_t, v^1, U^{L+1}$. We will thus write

$$Z^z = \psi \left(\left(\widehat{Z}^{\widehat{W}^l \bar{x}_0^{l-1}} \right)_l, \left(\widehat{Z}^{(\widehat{W}^k)^\top d\bar{h}_0^k} \right)_k, (U^1 \xi_s)_s, v^1, U^{L+1} \right)$$

to **generically** denote that the variable Z^z is a function **only** of the variables which appear in the arguments: $\widehat{Z}^{\widehat{W}^l \bar{x}_0^{l-1}}$, $\widehat{Z}^{(\widehat{W}^k)^\top d\bar{h}_0^k}$, $U^1 \xi_s$, v^1 , and U^{L+1} , (where multiple values of l , k and s might actually appear in the argument). This function ψ (we will sometimes also use ϕ) will of course depend on the z under consideration, and we might denote it by ψ^z (or ϕ^z), but most of the time we will omit this dependency and simply use the symbol ψ for different variables to express that the variable Z^z is a function of the arguments of ψ only.

We will see that the function ψ appearing will always be polynomially bounded by some form of composition or product of polynomially bounded functions, which will allow us to prove that the corresponding Z^z is finite almost surely since its arguments, considered as a vector, follow a Gaussian distribution with finite variance (and thus finite moments of any order). Note that in the proofs, we will use extensively (without explicitly saying so) that if ϕ and ψ are polynomially bounded then $\phi \times \psi$ is also polynomially bounded, and if φ is a polynomially bounded function of a single variable then $\varphi \circ \psi$ is also polynomially bounded. We introduce the following definition and lemma which we will use extensively in the proof by induction:

Definition L.1 (Vector of initial vectors and first forward-backward). Let $t \geq 1$. Then, dropping the dependency on t , we define the random vector:

$$Z_0 = Z_{0,t} := \left(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_t}, Z^{v^1}, \widehat{Z}^{U^{L+1}}, \right. \\ \left. \widehat{Z}^{\widehat{W}^2 \bar{x}_0^1}, \dots, \widehat{Z}^{\widehat{W}^L \bar{x}_0^{L-1}}, \right. \\ \left. \widehat{Z}^{(\widehat{W}^2)^\top d\bar{h}_0^2}, \dots, \widehat{Z}^{(\widehat{W}^L)^\top d\bar{h}_0^L} \right)$$

Remark L.1.

1. Note that any function of $Z_{0,s}$ will also be a function of $Z_{0,t}$ for $t \geq s$, which is also why we suppress the dependency on t as we can always take the largest possible t when we make a specific claim which involves Z_0 .
2. Also note that by the ZDot rule of the Tensor Program, for any vector z in the Tensor Program such that Z^z is a function only of Z_0 , then for any $l \in [2, L]$:

$$\begin{cases} \dot{Z}^{\widehat{W}^l z} = \mathbb{E} \left[\frac{\partial Z^z}{\partial \widehat{Z}^{(\widehat{W}^l)^\top d\bar{h}_0^l}} \right] Z^{d\bar{h}_0^l} \\ \dot{Z}^{(\widehat{W}^l)^\top z} = \mathbb{E} \left[\frac{\partial Z^z}{\partial \widehat{Z}^{\widehat{W}^l \bar{x}_0^{l-1}}} \right] Z^{\bar{x}_0^{l-1}} \end{cases}$$

Lemma L.1 (Distribution of Z_0 and moments). *One has*

$$(i) \ Z_0 \sim \mathcal{N} \left(0, \begin{pmatrix} S & 0 & 0 \\ 0 & D_f & 0 \\ 0 & 0 & D_b \end{pmatrix} \right) \text{ with}$$

$$S := \begin{pmatrix} \Sigma & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \in \mathbb{R}^{(t+3) \times (t+3)}, \quad \Sigma_{rs} = \xi_r^\top \xi_s,$$

$$D_f := \begin{bmatrix} \mathbb{E}[(Z^{\tilde{x}_0^1})^2] & & \\ & \ddots & \\ & & \mathbb{E}[(Z^{\tilde{x}_0^{L-1}})^2] \end{bmatrix} \in \mathbb{R}^{(L-1) \times (L-1)},$$

$$D_b := \begin{bmatrix} \mathbb{E}[(Z^{d\tilde{h}_0^2})^2] & & \\ & \ddots & \\ & & \mathbb{E}[(Z^{d\tilde{h}_0^L})^2] \end{bmatrix} \in \mathbb{R}^{(L-1) \times (L-1)}.$$

(ii) $|\mathbb{E}[\psi(Z_0)]| < \infty$, and $|\psi(Z_0)| < \infty$ almost surely for any polynomially bounded function $\psi : \mathbb{R}^{t+2L} \rightarrow \mathbb{R}$.

Remark L.2. Note that the lemma stays valid even if ψ does not depend on the whole list of variables inside Z_0 but only on a couple of them, which will be the case in the Tensor Program. Point (ii) will be used repeatedly in different proofs to show that the expectations appearing in the forward and backward passes are finite.

Proof. Claim (i) simply comes from the definition of the initial vectors $U^1 \xi_0, \dots, U^1 \xi_t, U^{L+1}$ and from the ZHat rule in a Tensor Program. Claim (ii) then follows because all entries in the covariance matrix are finite by Lemma F.1, and since ψ is polynomially bounded and the moments of a Gaussian with finite variance are finite, $|\mathbb{E}[\psi(Z_0)]| \leq \mathbb{E}[|\psi(Z_0)|] < \infty$ and thus $|\psi(Z_0)| < \infty$ almost surely. \square

Note that by Lemmas F.2 and F.3, the first forward and backward passes of IP-LLR easily express in function of the entries of Z_0 . Let us now take care of the forward and backward passes at $t = 1$. As the dynamics evolve with time, the expression of the forward and backward passes of IP-LLR in function of Z_0 (or rather of some of the entries of Z_0) get more intricate. They are still easy to develop explicitly for $t = 1$ but we choose to simply express what variables appear in the expression of the forward and backward passes instead of giving the expression explicitly.

Lemma L.2 (Multiplications by initial weight matrices vanish with polynomially bounded variables). *Consider the IP-LLR parameterization and let z be a vector in the program such that $Z^z = \psi(Z_0)$ with ψ polynomially bounded. Then, one has that for any $l \in [2, L]$:*

(i) if $\frac{\partial Z^z}{\partial \widehat{W}^l} = \phi(Z_0)$ with ϕ polynomially bounded, then $Z^{W^{(0)}z} = 0$.

(ii) if $\frac{\partial Z^z}{\partial \widehat{W}^{l-1}} = \phi(Z_0)$ with ϕ polynomially bounded, then $Z^{(W^{(0)})^\top z} = 0$.

Proof. Let $l \in [2, L]$. We simply write

$$Z^{W^{(0)}z} = \overset{\circ}{\omega}_l \widehat{Z}^{\widehat{W}^l z} + \overset{\circ}{\omega}_l \dot{\widehat{Z}}^{\widehat{W}^l z}$$

where $\widehat{Z}^{W^{(0)}z} \sim \mathcal{N}(0, \mathbb{E}[(Z^z)^2])$ and the variance is finite by Lemma L.1 because $(Z^z)^2$ is a polynomially bounded function of Z_0 since Z^z is. This shows that $|\widehat{Z}^{W^{(0)}z}| < \infty$ almost

surely and thus that $\mathring{\omega}_l \widehat{Z}^{\widehat{W}^l z} = 0$ since $\mathring{\omega}_l = 0$ in IPs. On the other hand,

$$\dot{Z}^{\widehat{W}^l z} = \mathbb{E} \left[\frac{\partial Z^z}{\partial \widehat{Z}^{\widehat{W}^l \top} d\tilde{h}_0^l} \right] Z^{d\tilde{h}_0^l}$$

and the expectation is finite by Lemma L.1 since $\partial Z^z / \partial \widehat{Z}^{\widehat{W}^l \top} = \phi(Z_0)$ with ϕ polynomially bounded, and

$$Z^{d\tilde{h}_0^l} = \begin{cases} \widehat{Z}^{\widehat{W}^{l+1} \top} d\tilde{h}_{l+1_0} \sigma'(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}) & \text{if } l \in [2, L-1] \\ \widehat{Z}^{U^{L+1}} \sigma'(\widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}) & \text{if } l = L \end{cases}$$

In any case, $Z^{d\tilde{h}_0^l}$ is a polynomially bounded function of Z_0 and is thus finite almost surely, which entails $\mathring{\omega}_l \widehat{Z}^{\widehat{W}^l z} = 0$, and therefore $\widehat{Z}^{W^l(0)z} = 0$ which gives (i).

The same reasoning with $W^l(0)^\top$ gives (ii) if $\partial Z^z / \partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}} = \phi(Z_0)$ with ϕ polynomially bounded. □

L.0.1 The case $t = 1$

Lemma L.3 (Z_0 in the forward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, dropping the dependency of the forward pass on ξ_1 , one has:*

$$(i) \quad Z^{h_1^1} = \psi \left(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{\widehat{W}^2 \top} d\tilde{h}_0^2 \right)$$

$$(ii) \quad Z^{h_1^l} = \psi \left(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}, \widehat{Z}^{\widehat{W}^{l+1} \top} d\tilde{h}_0^{l+1} \right), \quad l \in [2, L-1]$$

$$(iii) \quad Z^{h_1^L} = \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right)$$

and

$$(iv) \quad \frac{\partial Z^{h_1^{l-1}}}{\widehat{Z}^{\widehat{W}^l \top} d\tilde{h}_0^l} = \psi(Z_0), \quad l \in [2, L]$$

$$(v) \quad \frac{\partial Z^{h_1^l}}{\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}} = \psi(Z_0), \quad l \in [2, L]$$

and **all** the different ψ that appear are polynomially bounded.

Remark L.3.

1. Recall that we simply use ψ or ϕ to mean that the variable **is a function of** the arguments of ψ (or ϕ) only, and that the different ψ and ϕ which appear in the different claims (i) to (v) are **not** actually the same.
2. For the partial derivatives we chose not to make a precise statement on which variables exactly appear in the expression as this will not matter and would only over-complicate things for close to none added-value.
3. Note that with the claims above, one can prove that $\mathring{\omega}_l Z^{\widehat{W}^l x_1^{l-1}} = 0$ because both of the terms \widehat{Z} and \dot{Z} defining $Z^{\widehat{W}^l x_1^{l-1}}$ are polynomially bounded functions of Gaussians which has finite covariance matrices, and $\mathring{\omega}_l = 0$ in IPs.

Proof. Using [Theorem G.7](#) with $\xi = \xi_1$ and $t = 1$ we have claim (i) because, first $\widehat{Z}^{d\tilde{x}_0^1} = \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2}$, and second σ' is polynomially bounded (see [Appendix E](#)). Claim (ii) also stems from [Theorem G.7](#) since $\widehat{Z}^{d\tilde{x}_0^l} = \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}}$, $\widehat{Z}^{\tilde{h}_0^l} = \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}$, and σ' is polynomially bounded. Finally, claim (iii) also stems from [Theorem G.7](#) since $\widehat{Z}^{d\tilde{x}_0^L} = Z^{U^{L+1}}$, $\widehat{Z}^{\tilde{h}_0^L} = \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}$, and σ' is polynomially bounded.

From [Theorem G.7](#), we get:

$$\frac{\partial Z^{h_1^1}}{\partial \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2}} = -\eta \overset{\circ}{\chi}_0 (\xi_0^\top \xi_1) \sigma'(\widehat{Z}^{U^1 \xi_0})$$

For $l \in [3, L]$, from [Theorem G.7](#), we get

$$\frac{\partial Z^{h_1^{l-1}}}{Z^{(\widehat{W}^l)^\top d\tilde{h}_0^l}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-2}} Z^{\tilde{x}_1^{l-2}}] \sigma'(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}})$$

which immediately gives claim (iv) since σ' is polynomially bounded and with claim (ii) and [Lemma L.1](#), we also have $|\mathbb{E}[Z^{\tilde{x}_0^{l-2}} Z^{\tilde{x}_1^{l-2}}]| < \infty$. Similarly, for $l \in [2, L]$, from [Theorem G.7](#), we get

$$\frac{\partial Z^{h_1^l}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^l}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{\tilde{x}_1^{l-1}}] Z^{d\tilde{x}_0^l} \sigma''(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}})$$

and $Z^{d\tilde{x}_0^{l-1}} = \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}}$ if $l \in [2, L-1]$, and $Z^{d\tilde{x}_0^{L-1}} = \widehat{Z}^{U^{L+1}}$ if $l = L$. Since the expectation is finite by claim (ii) and [Lemma L.1](#), and since σ'' is polynomially bounded, we get claim (v). \square

Lemma L.4 (Z_0 in the backward pass of IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, dropping the dependency of the forward and backward passes on ξ_1 , one has:*

$$(i) \quad \begin{aligned} Z^{d\tilde{x}_1^L} &= \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right), \\ Z^{d\tilde{h}_1^L} &= \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right) \end{aligned}$$

$$(ii) \quad \begin{aligned} Z^{d\tilde{x}_1^{l-1}} &= \psi \left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}} \right), \\ Z^{d\tilde{h}_1^{l-1}} &= \psi \left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}, \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l} \right), \quad l \in [3, L] \end{aligned}$$

$$(iii) \quad \begin{aligned} Z^{d\tilde{x}_1^1} &= \psi \left(\widehat{Z}^{U^1 \xi_0} \right), \\ Z^{d\tilde{h}_1^1} &= \psi \left(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right) \end{aligned}$$

and

$$(iv) \quad \frac{\partial Z^{d\tilde{h}_1^l}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}} = \psi(Z_0), \quad l \in [1, L]$$

$$(v) \quad \frac{\partial Z^{d\tilde{h}_1^l}}{\partial \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}}} = \psi(Z_0), \quad l \in [1, L-1]$$

and **all** the different ψ that appear are polynomially bounded.

Proof. For the backward pass, we have by definition of the tilde variables for $t \geq 1$, $d\tilde{x}_1^L = w^{L+1}(1) = U^{L+1} - \eta\chi_0\tilde{x}_1^L$ by Lemma F.4.2, and thus

$$Z^{d\tilde{x}_1^L} = \widehat{Z}^{U^{L+1}} - \eta\overset{\circ}{\chi}_0\sigma(\widehat{Z}^{\widehat{W}^L\tilde{x}_0^{L-1}})$$

Then,

$$Z^{d\tilde{h}_1^L} = Z^{d\tilde{x}_1^L}\sigma'(Z^{h_1^L})$$

which gives claim (i) since σ and σ' are polynomially bounded, and $Z^{h_1^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L\tilde{x}_0^{L-1}})$ and ψ is polynomially bounded by Lemma L.3.

For $l = L - 1$, we have

$$\begin{aligned} d\tilde{x}_1^{L-1} &= (W^L(1))^\top d\tilde{h}_1^L \\ &= \omega_L(\widehat{W}^L)^\top d\tilde{h}_1^L - \eta\chi_0 \frac{(d\tilde{h}_0^L)^\top d\tilde{h}_1^L}{m} \tilde{x}_0^{L-1} \end{aligned}$$

which gives

$$Z^{d\tilde{x}_1^{L-1}} = \overset{\circ}{\omega}_L Z^{(\widehat{W}^L)^\top d\tilde{h}_1^L} - \eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}] Z^{\tilde{x}_0^{L-1}}$$

Now, by the previous expression of $Z^{d\tilde{h}_1^L}$ and by Lemma L.3, we get

$$\frac{\partial Z^{d\tilde{h}_1^L}}{\partial \widehat{W}^L \tilde{x}_0^{L-1}} = -\eta\overset{\circ}{\chi}_0 \sigma'(\widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}) \sigma'(Z^{h_1^L}) + Z^{d\tilde{x}_1^L} \frac{\partial Z^{h_1^L}}{\partial \widehat{W}^L \tilde{x}_0^{L-1}} \sigma''(Z^{h_1^L})$$

and by claim (i) and Lemma L.3 we get

$$\frac{\partial Z^{d\tilde{h}_1^L}}{\partial \widehat{W}^L \tilde{x}_0^{L-1}} = \psi(Z_0)$$

with ψ polynomially bounded since σ' and σ'' are polynomially bounded. Therefore, by Lemma L.2, we get $\overset{\circ}{\omega}_L Z^{(\widehat{W}^L)^\top d\tilde{h}_1^L} = 0$.

We thus simply get

$$Z^{d\tilde{x}_1^{L-1}} = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}] Z^{\tilde{x}_0^{L-1}}$$

$Z^{d\tilde{h}_0^L}$ and $Z^{d\tilde{h}_1^L}$ are polynomially bounded functions of Z_0 and thus so is $Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}$, and by Lemma L.1, $|\mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}]| < \infty$. Since $Z^{d\tilde{x}_1^{L-1}} = \psi(Z^{\tilde{x}_0^{L-1}})$ with ψ polynomially bounded, we thus get $Z^{d\tilde{x}_1^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{L-1}\tilde{x}_0^{L-2}})$ and ψ is polynomially bounded (indeed: $\psi(z) = -\eta\overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}] \sigma(z)$).

We have

$$Z^{d\tilde{h}_1^{L-1}} = Z^{d\tilde{x}_1^{L-1}} \sigma'(Z^{h_1^{L-1}})$$

and since by Lemma L.3, $Z^{h_1^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{L-1}\tilde{x}_0^{L-2}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$ with ψ polynomially bounded, by the previous result for $Z^{d\tilde{x}_1^{L-1}}$ and since σ' is polynomially bounded we get

$$Z^{d\tilde{h}_1^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{L-1}\tilde{x}_0^{L-2}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$$

with ψ polynomially bounded.

We have

$$\begin{aligned} \frac{\partial Z^{d\tilde{h}_1^{L-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}} &= \frac{\partial Z^{d\tilde{x}_1^{L-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}} \sigma'(Z^{h_1^{L-1}}) + Z^{d\tilde{x}_1^{L-1}} \frac{\partial Z^{h_1^{L-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}} \sigma''(Z^{h_1^{L-1}}) \\ &= -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}] \sigma'(\widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}) \sigma'(Z^{h_1^{L-1}}) + \\ &\quad Z^{d\tilde{x}_1^{L-1}} \frac{\partial Z^{h_1^{L-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}} \sigma''(Z^{h_1^{L-1}}) \end{aligned}$$

By Lemma L.3 and since σ' and σ'' are polynomially bounded, and we have already proven that $Z^{d\tilde{x}_1^{L-1}} = \psi(Z_0)$ with ψ polynomially bounded, as well as $|\mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_1^L}]| < \infty$, we get

$$\frac{\partial Z^{d\tilde{h}_1^{L-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{L-2}}} = \psi(Z_0)$$

with ψ polynomially bounded.

Similarly, we have

$$\begin{aligned} \frac{\partial Z^{d\tilde{h}_1^{L-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}} &= \frac{\partial Z^{d\tilde{x}_1^{L-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}} \sigma'(Z^{h_1^{L-1}}) + Z^{d\tilde{x}_1^{L-1}} \frac{\partial Z^{h_1^{L-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}} \sigma''(Z^{h_1^{L-1}}) \\ &= Z^{d\tilde{x}_1^{L-1}} \frac{\partial Z^{h_1^{L-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}} \sigma''(Z^{h_1^{L-1}}) \end{aligned}$$

Now, we have shown above that $Z^{d\tilde{x}_1^{L-1}} = \psi(Z_0)$ with ψ polynomially bounded, and by Lemma L.3 we have that both $\partial Z^{h_1^{L-1}} / \partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}$ and $Z^{h_1^{L-1}}$ are polynomially bounded functions of Z_0 , which gives

$$\frac{\partial Z^{d\tilde{h}_1^{L-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}} = \psi(Z_0)$$

with ψ polynomially bounded.

Let $l \in [2, L-1]$ and assume claims (ii), (iv) and (v) are true for layer l . We have

$$Z^{d\tilde{x}_1^{l-1}} = \overset{\circ}{\omega}_l Z^{(\widehat{W}^l)^\top d\tilde{h}_1^l} - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] Z^{\tilde{x}_0^{l-1}}$$

Since by the induction hypothesis $\partial Z^{d\tilde{h}_1^l} / \partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}} = \psi(Z_0)$ with ψ polynomially bounded, and $Z^{d\tilde{h}_1^l}$ is a polynomially bounded function of Z_0 , by Lemma L.2 we get $\overset{\circ}{\omega}_l Z^{(\widehat{W}^l)^\top d\tilde{h}_1^l} = 0$. Then, we simply get

$$Z^{d\tilde{x}_1^{l-1}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] Z^{\tilde{x}_0^{l-1}}$$

Again here, since both $Z^{d\tilde{h}_0^l}$ and $Z^{d\tilde{h}_1^l}$ are polynomially bounded functions of Z_0 , then so is $Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}$, which shows by Lemma L.1 that $|\mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}]| < \infty$. If $l \geq 3$, since $Z^{\tilde{x}_0^{l-1}} = \sigma(\widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{l-2}})$ and σ is polynomially bounded, we get that

$$Z^{d\tilde{x}_1^{l-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1}\tilde{x}_0^{l-2}})$$

If $l = 2$, since $Z^{\tilde{x}_0^1} = \sigma(Z^{U^1 \xi_0})$ and σ is polynomially bounded we get:

$$Z^{d\tilde{x}_1^{l-1}} = \psi(Z^{U^1 \xi_0})$$

with ψ polynomially bounded.

We then have

$$Z^{d\tilde{h}_1^{l-1}} = Z^{d\tilde{x}_1^{l-1}} \sigma'(Z^{h_1^{l-1}})$$

and thus

$$\begin{aligned} \frac{\partial Z^{d\tilde{h}_1^{l-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}} &= \frac{\partial Z^{d\tilde{x}_1^{l-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1}(0) \tilde{x}_0^{l-2}}} \sigma'(Z^{h_1^{l-1}}) + Z^{d\tilde{x}_1^{l-1}} \frac{\partial Z^{h_1^{l-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}} \sigma''(Z^{h_1^{l-1}}) \\ &= -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}] \sigma'(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}) \sigma'(Z^{h_1^{l-1}}) + \\ &\quad Z^{d\tilde{x}_1^{l-1}} \frac{\partial Z^{h_1^{l-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}} \sigma''(Z^{h_1^{l-1}}) \end{aligned}$$

By Lemma L.3 as well as the previous result on $Z^{d\tilde{x}_1^{l-1}}$, and since $|\mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_1^l}]| < \infty$, and σ' and σ'' are polynomially bounded, we get that

$$\frac{\partial Z^{d\tilde{h}_1^{l-1}}}{\partial \widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}} = \psi(Z_0)$$

with ψ polynomially bounded.

Similarly

$$\begin{aligned} \frac{\partial Z^{d\tilde{h}_1^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^l}} &= \frac{\partial Z^{d\tilde{x}_1^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^l}} \sigma'(Z^{h_1^{l-1}}) + Z^{d\tilde{x}_1^{l-1}} \frac{\partial Z^{h_1^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^l}} \sigma''(Z^{h_1^{l-1}}) \\ &= Z^{d\tilde{x}_1^{l-1}} \frac{\partial Z^{h_1^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^l}} \sigma''(Z^{h_1^{l-1}}) \end{aligned}$$

and the three quantities in the product are polynomially bounded functions of Z_0 (shown above for the first term and by Lemma L.3 for the two other terms). We thus get

$$\frac{\partial Z^{d\tilde{h}_1^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^l}} = \psi(Z_0)$$

with ψ polynomially bounded. This concludes the induction and thus proves claims (ii), (iii), (iv) and (v) by induction. \square

Corollary L.4.1 (Multiplications by the initial weight matrices vanish in IP-LLR at $t = 1$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then for any $l \in [2, L]$, one has:*

$$\begin{cases} Z^{W^l(0)x_1^{l-1}} = \overset{\circ}{\omega}_l Z^{\widehat{W}^l x_1^{l-1}} = 0 \\ Z^{(W^l(0))^\top d\tilde{h}_1^l} = \overset{\circ}{\omega}_l Z^{(\widehat{W}^l)^\top d\tilde{h}_1^l} = 0 \end{cases}$$

Proof. Those results are actually hidden in the proof of Lemma L.4 and come from Lemma L.2. \square

L.0.2 The case $t = 2$

Lemma L.5 (Z_0 in the forward pass of IP-LLR at $t = 2$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, dropping the dependency of the forward pass on ξ_2 , one has:*

$$(i) \quad Z^{h_2^1} = \psi \left(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{U^1 \xi_2}, \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2} \right)$$

$$(ii) \quad Z^{h_2^l} = \psi \left(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}, \widehat{Z}^{(\widehat{W}^{l+1})^\top d \tilde{h}_0^{l+1}} \right), \quad l \in [2, L-1]$$

$$(iii) \quad Z^{h_2^L} = \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right)$$

and

$$(iv) \quad \frac{\partial Z^{h_2^{l-1}}}{\widehat{Z}^{(\widehat{W}^l)^\top d \tilde{h}_0^l}} = \psi(Z_0), \quad l \in [2, L]$$

$$(v) \quad \frac{\partial Z^{h_2^l}}{\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}} = \psi(Z_0), \quad l \in [2, L]$$

and

$$(vi) \quad Z^{W^{l(0)} x_2^{l-1}} = 0, \quad l \in [2, L]$$

and **all** the different ψ that appear are polynomially bounded.

Proof. We have

$$h_2^1 = U^1 \xi_2 - \eta \chi_0(\xi_0^\top \xi_2) d \tilde{h}_0^1 - \eta \chi_1(\xi_1^\top \xi_2) d \tilde{h}_1^1$$

which gives

$$Z^{h_2^1} = Z^{U^1 \xi_2} - \eta \overset{\circ}{\chi}_0(\xi_0^\top \xi_2) Z^{d \tilde{h}_0^1} - \eta \overset{\circ}{\chi}_1(\xi_1^\top \xi_2) Z^{d \tilde{h}_1^1}$$

By Lemma L.4 $Z^{d \tilde{h}_1^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2})$ and we also have $Z^{d \tilde{h}_0^1} = \psi(Z^{U^1 \xi_0}, \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2})$ where the different ψ are polynomially bounded, which gives claim (i).

We have

$$\frac{\partial Z^{h_2^1}}{\partial \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2}} = -\eta \overset{\circ}{\chi}_0(\xi_0^\top \xi_2) \frac{\partial Z^{d \tilde{h}_0^1}}{\partial \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2}} - \eta \overset{\circ}{\chi}_1(\xi_1^\top \xi_2) \frac{\partial Z^{d \tilde{h}_1^1}}{\partial \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2}}$$

with

$$\frac{\partial Z^{d \tilde{h}_0^1}}{\partial \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2}} = \sigma'(Z^{U^1 \xi_0})$$

which is a polynomially bounded function of Z_0 and so is $\partial Z^{d \tilde{h}_1^1} / \partial \widehat{Z}^{(\widehat{W}^2)^\top d \tilde{h}_0^2}$ by Lemma L.4. We thus get claim (iv) for $l = 2$.

We have

$$Z^{h_2^2} = \overset{\circ}{\omega}_2 Z^{\widehat{W}^2 x_2^1} - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_2^1}] Z^{d \tilde{h}_0^2} - \eta \overset{\circ}{\chi}_1 \mathbb{E}[Z^{x_1^1} Z^{x_2^1}] Z^{d \tilde{h}_1^2}$$

Now $Z^{x_2^1} = \sigma(Z^{h_2^1})$ is a polynomially bounded function of Z_0 because $Z^{h_2^1}$ is and σ is polynomially bounded. Secondly, we have

$$\frac{\partial Z^{x_2^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2} = \frac{\partial Z^{h_2^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2} \sigma'(Z^{h_2^1})$$

which is a polynomially bounded function of Z_0 by the previous results. By Lemma L.2 we get that $\omega_2^\circ \widehat{Z}^{\widehat{W}^2 x_2^1} = 0$ which gives claim (vi) for $l = 2$. In addition, this yields

$$Z^{h_2^2} = -\eta \chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_2^1}] Z^{d\tilde{h}_0^2} - \eta \chi_1^\circ \mathbb{E}[Z^{x_1^1} Z^{x_2^1}] Z^{d\tilde{h}_1^2}$$

which gives claim (ii) for $l = 2$ by the results for the backward passes at time $t = 0$ and $t = 1$ and because the expectations are finite since the integrands are polynomially bounded functions of Z_0 , as they are products of such variables by the induction hypothesis. Additionally, we have

$$\frac{\partial Z^{h_2^2}}{\partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3} = -\eta \chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^1} Z^{x_2^1}] \frac{\partial Z^{d\tilde{h}_0^2}}{\partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3} - \eta \chi_1^\circ \mathbb{E}[Z^{x_1^1} Z^{x_2^1}] \frac{\partial Z^{d\tilde{h}_1^2}}{\partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3}$$

and we have

$$\frac{\partial Z^{d\tilde{h}_0^2}}{\partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3} = \sigma'(\widehat{Z}^{\widehat{W}^2 \tilde{x}_0^1})$$

and $\partial Z^{d\tilde{h}_1^2} / \partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3$ is a polynomially bounded function of Z_0 by Lemma L.4. Once again, since the expectations are finite, we thus get that

$$\frac{\partial Z^{h_2^2}}{\partial \widehat{Z}(\widehat{W}^3(0))^\top d\tilde{h}_0^3} = \psi(Z_0)$$

with ψ polynomially bounded. A similar reasoning would prove that

$$\frac{\partial Z^{h_2^2}}{\partial \widehat{Z}^{\widehat{W}^3(0) \tilde{x}_0^2}} = \psi(Z_0)$$

with ψ polynomially bounded because

$$\frac{\partial Z^{d\tilde{h}_0^1}}{\partial \widehat{Z}^{\widehat{W}^3(0) \tilde{x}_0^2}} = \widehat{Z}^{(\widehat{W}^3(0))^\top d\tilde{h}_0^3} \sigma''(\widehat{Z}^{\widehat{W}^2 \tilde{x}_0^1})$$

and $\partial Z^{d\tilde{h}_1^1} / \partial \widehat{Z}^{\widehat{W}^3(0) \tilde{x}_0^2} = \psi(Z_0)$ with ψ polynomially bounded by Lemma L.4.

Let $l \in [2, L - 1]$ and assume claims (ii), (iv), (v), and (vi) for layer l . Then, we have:

$$Z^{h_2^{l+1}} = \omega_{l+1}^\circ \widehat{Z}^{\widehat{W}^{l+1} x_2^l} - \eta \chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}] Z^{d\tilde{h}_0^{l+1}} - \eta \chi_1^\circ \mathbb{E}[Z^{x_1^l} Z^{x_2^l}] Z^{d\tilde{h}_1^{l+1}}$$

Now $Z^{x_2^l} = \sigma(Z^{h_2^l})$ is a polynomially bounded function of Z_0 because $Z^{h_2^l}$ is by the induction hypothesis and σ is polynomially bounded. Secondly, we have

$$\frac{\partial Z^{x_2^l}}{\partial \widehat{Z}(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}} = \frac{\partial Z^{h_2^l}}{\partial \widehat{Z}(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}} \sigma'(Z^{h_2^l})$$

which is a polynomially bounded function of Z_0 by the induction hypothesis. By Lemma L.2 we get that $\omega_{l+1}^\circ Z^{\widehat{W}^{l+1}x_2^l} = 0$ which gives claim (vi) for layer $l + 1$. In addition, this yields

$$Z^{h_2^{l+1}} = -\eta\chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}] Z^{d\tilde{h}_0^{l+1}} - \eta\chi_1^\circ \mathbb{E}[Z^{x_1^l} Z^{x_2^l}] Z^{d\tilde{h}_1^{l+1}}$$

which gives claim (ii) for layer $l + 1$ by the results for the backward passes at time $t = 0$ and $t = 1$ and because the expectations are finite since the integrands are polynomially bounded functions of Z_0 , as they are products of such variables. The only thing that one has to be careful with is that if $l + 1 = L$, then $Z^{\tilde{h}_0^{l+1}} = \psi(Z^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ and $Z^{\tilde{h}_1^{l+1}} = \psi(Z^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ with both ψ polynomially bounded, which gives claim (iii). Otherwise, if $l + 1 \leq L - 1$, $Z^{\tilde{h}_0^{l+1}} = \psi(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}, \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}^{l+2}})$ and $Z^{\tilde{h}_1^{l+1}} = \psi(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}, \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}^{l+2}})$ with both ψ polynomially bounded, which gives claim (ii) for layer $l + 1$.

Now, if $l + 1 \leq L - 1$,

$$\frac{\partial Z^{h_2^{l+1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}} = -\eta\chi_0^\circ \mathbb{E}[Z^{\tilde{x}_0^l} Z^{x_2^l}] \frac{\partial Z^{d\tilde{h}_0^{l+1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}} - \eta\chi_1^\circ \mathbb{E}[Z^{x_1^l} Z^{x_2^l}] \frac{\partial Z^{d\tilde{h}_1^{l+1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}}$$

and we have

$$\frac{\partial Z^{d\tilde{h}_0^{l+1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}} = \sigma'(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l})$$

and $\partial Z^{d\tilde{h}_1^{l+1}} / \partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}$ is a polynomially bounded function of Z_0 by Lemma L.4. Once again, since the expectations are finite, we thus get that

$$\frac{\partial Z^{h_2^{l+1}}}{\partial \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}}} = \psi(Z_0)$$

with ψ polynomially bounded, which proves claim (iv) for layer $l + 1$. A similar reasoning would prove that

$$\frac{\partial Z^{h_2^{l+1}}}{\partial \widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}} = \psi(Z_0)$$

with ψ polynomially bounded because

$$\frac{\partial Z^{d\tilde{h}_0^{l+1}}}{\partial \widehat{Z}^{\widehat{W}^{l+2}(0) \tilde{x}_0^l}} = \begin{cases} \widehat{Z}^{(\widehat{W}^{l+2}(0))^\top d\tilde{h}_0^{l+2}} \sigma''(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}) & \text{if } l + 1 \leq L - 1 \\ \widehat{Z}^{U^{L+1}} \sigma''(\widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}) & \text{if } l + 1 = L \end{cases}$$

and $\partial Z^{d\tilde{h}_1^{l+1}} / \partial \widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l} = \psi(Z_0)$ with ψ polynomially bounded by Lemma L.4. This proves claim (v) and thus concludes the induction and with it the proof. \square

Lemma L.6 (Z_0 in the backward pass of IP-LLR at $t = 2$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, dropping the dependency of the forward and backward passes on ξ_2 , one has:*

- (i) $Z^{d\tilde{x}_2^L} = \psi\left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}\right),$
 $Z^{d\tilde{h}_2^L} = \psi\left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}\right)$
- (ii) $Z^{d\tilde{x}_2^{l-1}} = \psi\left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}^{l-2}}, \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l}\right),$
 $Z^{d\tilde{h}_2^{l-1}} = \psi\left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}^{l-2}}, \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l}\right), l \in [3, L]$

$$(iii) \quad \begin{aligned} Z^{d\tilde{x}_2^1} &= \psi \left(\widehat{Z}^{U^1\xi_0}, \widehat{Z}^{U^1\xi_1}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right), \\ Z^{d\tilde{h}_2^1} &= \psi \left(\widehat{Z}^{U^1\xi_0}, \widehat{Z}^{U^1\xi_1}, \widehat{Z}^{U^1\xi_2}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right) \end{aligned}$$

and

$$(iv) \quad \frac{\partial Z^{d\tilde{h}_2^l}}{\partial \widehat{Z}^{W^l \tilde{x}_0^{l-1}}} = \psi(Z_0), \quad l \in [2, L]$$

$$(v) \quad \frac{\partial Z^{d\tilde{h}_2^{l-1}}}{\partial \widehat{Z}^{(W^L)^\top d\tilde{h}_0^l}} = \psi(Z_0), \quad l \in [2, L]$$

and

$$(vi) \quad Z^{(W^l(0))^\top d\tilde{h}_2^l} = 0, \quad l \in [2, L]$$

and **all** the different ψ that appear are polynomially bounded.

Proof. We have:

$$Z^{d\tilde{x}_2^L} = \widehat{Z}^{U^{L+1}} - \eta \overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L} - \eta \overset{\circ}{\chi}_1 Z^{x_1^L}$$

where $Z^{\tilde{x}_0^L} = \sigma(\widehat{Z}^{W^L \tilde{x}_0^{L-1}})$ and $Z^{x_1^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{W^L \tilde{x}_0^{L-1}})$ with ψ polynomially bounded by Lemma L.3. Combining all this gives

$$Z^{d\tilde{x}_2^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{W^L \tilde{x}_0^{L-1}})$$

with ψ polynomially bounded since σ is also polynomially bounded. Then

$$Z^{d\tilde{h}_2^L} = Z^{d\tilde{x}_2^L} \sigma'(Z^{h_2^L})$$

and since $Z^{h_2^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{W^L \tilde{x}_0^{L-1}})$ with ψ polynomially bounded by Lemma L.5, we get

$$Z^{d\tilde{h}_2^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{W^L \tilde{x}_0^{L-1}})$$

with ψ polynomially bounded since σ' is also polynomially bounded. This thus proves claim (i). Now, we have

$$\frac{\partial Z^{d\tilde{h}_2^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} = \frac{\partial Z^{d\tilde{x}_2^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} \sigma'(Z^{h_2^L}) + Z^{d\tilde{x}_2^L} \frac{\partial Z^{h_2^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} \sigma''(Z^{h_2^L})$$

where $Z^{h_2^L}$, $\partial Z^{h_2^L} / \partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}$, and $Z^{d\tilde{x}_2^L}$ are polynomially bounded functions of Z_0 by the previous result and by Lemma L.5. We have

$$\frac{\partial Z^{d\tilde{x}_2^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} = -\eta \overset{\circ}{\chi}_0 \sigma'(\widehat{Z}^{W^L \tilde{x}_0^{L-1}}) - \eta \overset{\circ}{\chi}_1 \frac{\partial Z^{h_1^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} \sigma'(Z^{h_1^L})$$

which is a polynomially bounded function of Z_0 since σ' is polynomially bounded and by Lemma L.3. We thus get

$$\frac{\partial Z^{d\tilde{h}_2^L}}{\partial \widehat{Z}^{W^L \tilde{x}_0^{L-1}}} = \psi(Z_0)$$

with ψ polynomially bounded since σ' and σ'' are polynomially bounded. This proves (iv) for $l = L$.

We have:

$$Z^{d\tilde{x}_2^{L-1}} = \mathring{\omega}_L \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_2^L - \eta\mathring{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_2^L}] Z^{\tilde{x}_0^{L-1}} - \eta\mathring{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^L} Z^{d\tilde{h}_2^L}] Z^{x_1^{L-1}}$$

From the previous step we have that both $Z^{d\tilde{h}_2^L}$ and $\partial Z^{d\tilde{h}_2^L} / \partial \widehat{Z}(\widehat{W}^l)^\top d\tilde{x}_0^{L-1}$ are polynomially bounded functions of Z_0 . By Lemma L.2, this first shows that $\mathring{\omega}_L \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_2^L = 0$, and thus gives (vi) for $l = L$, leading to:

$$Z^{d\tilde{x}_2^{L-1}} = -\eta\mathring{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_2^L}] Z^{\tilde{x}_0^{L-1}} - \eta\mathring{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^L} Z^{d\tilde{h}_2^L}] Z^{x_1^{L-1}}$$

Now $Z^{\tilde{x}_0^{L-1}} = \sigma(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2})$ and by Lemma L.3, we also have that $Z^{x_1^{L-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^L)$ with ψ polynomially bounded. As always the expectations are finite by Lemma L.1 because the integrands are polynomially bounded functions of Z_0 as products of such variables. Since σ is also polynomially bounded, this gives

$$Z^{d\tilde{x}_2^{L-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^L)$$

Then, we have

$$Z^{d\tilde{h}_2^{L-1}} = Z^{d\tilde{x}_2^{L-1}} \sigma'(Z^{h_2^{L-1}})$$

and since $Z^{h_2^{L-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^L)$ with ψ polynomially bounded by Lemma L.5, we get

$$Z^{d\tilde{h}_2^{L-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^L)$$

with ψ polynomially bounded since σ' is also polynomially bounded. This thus proves claim (ii) for $l = L - 1$. Now, let $Z \in \{\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{L-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^L\}$. We have

$$\frac{\partial Z^{d\tilde{h}_2^{L-1}}}{\partial Z} = \frac{\partial Z^{d\tilde{x}_2^{L-1}}}{\partial Z} \sigma'(Z^{h_2^{L-1}}) + Z^{d\tilde{x}_2^{L-1}} \frac{\partial Z^{h_2^{L-1}}}{\partial Z} \sigma''(Z^{h_2^{L-1}})$$

where $Z^{h_2^{L-1}}$, $\partial Z^{h_2^{L-1}} / \partial Z$, and $Z^{d\tilde{x}_2^{L-1}}$ are polynomially bounded functions of Z_0 by the previous result and by Lemma L.5. We have

$$\frac{\partial Z^{d\tilde{x}_2^{L-1}}}{\partial Z} = -\eta\mathring{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^L} Z^{d\tilde{h}_2^L}] \frac{\partial Z^{\tilde{h}_0^{L-1}}}{\partial Z} \sigma'(Z^{\tilde{h}_0^{L-1}}) - \eta\mathring{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^L} Z^{d\tilde{h}_2^L}] \frac{\partial Z^{h_1^{L-1}}}{\partial Z} \sigma'(Z^{h_1^{L-1}})$$

which is a polynomially bounded function of Z_0 since σ' is polynomially bounded and by Lemma L.3.

For both possible values of Z , the expression of $\partial Z^{\tilde{h}_0^{L-1}} / \partial Z$ is easy to obtain and is a polynomially bounded function of Z_0 (this has actually already been shown for the proofs at time $t = 1$), and $Z^{h_1^{L-1}} / \partial Z = \psi(Z_0)$ with ψ polynomially bounded by Lemma L.3. Since the expectations are finite and σ' is polynomially bounded, we get

$$\frac{\partial Z^{d\tilde{x}_2^{L-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded and thus

$$\frac{\partial Z^{d\tilde{h}_2^{L-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded. This proves (iv) and (v) for $l = L - 1$.

Let $l \in [2, L - 1]$, and assume claims (ii), (iv), (v), are true at layer l and claim (vi) is true at layer $l + 1$. We have:

$$Z^{d\tilde{x}_2^{l-1}} = \overset{\circ}{\omega}_l \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_2^l - \eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_2^l}] Z^{\tilde{x}_0^{l-1}} - \eta \overset{\circ}{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^l} Z^{d\tilde{h}_2^l}] Z^{x_1^{l-1}}$$

From the induction hypothesis we have that both $Z^{d\tilde{h}_2^l}$ and $\partial Z^{d\tilde{h}_2^l} / \partial \widehat{Z}(\widehat{W}^l)^\top d\tilde{x}_0^{l-1}$ are polynomially bounded functions of Z_0 . By Lemma L.2, this first shows that $\overset{\circ}{\omega}_l \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_2^l = 0$, and thus gives (vi) for layer l , leading to:

$$Z^{d\tilde{x}_2^{l-1}} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_2^l}] Z^{\tilde{x}_0^{l-1}} - \eta \overset{\circ}{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^l} Z^{d\tilde{h}_2^l}] Z^{x_1^{l-1}}$$

Now, if $l - 1 \geq 2$, $Z^{\tilde{x}_0^{l-1}} = \sigma(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2})$ and by Lemma L.3, we also have that $Z^{x_1^{l-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l)$ with ψ polynomially bounded. On the other hand, if $l - 1 = 1$, we have $Z^{\tilde{x}_0^{l-1}} = \sigma(\widehat{Z}^{U^1 \xi_0})$ and we also have that $Z^{x_1^{l-1}} = \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2)$ by Lemma L.3. As always the expectations are finite by Lemma L.1 because the integrands are polynomially bounded functions of Z_0 as products of such variables. Since σ is also polynomially bounded, this gives

$$Z^{d\tilde{x}_2^{l-1}} = \begin{cases} \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l) & \text{if } l - 1 \geq 2 \\ \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2) & \text{if } l - 1 = 1 \end{cases}$$

Since $Z^{d\tilde{h}_2^{l-1}} = Z^{d\tilde{x}_2^{l-1}} \sigma'(Z^{h_2^{l-1}})$, by Lemma L.3 we get

$$Z^{d\tilde{h}_2^{l-1}} = \begin{cases} \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l) & \text{if } l - 1 \geq 2 \\ \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{U^1 \xi_2}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2) & \text{if } l - 1 = 1 \end{cases}$$

This gives claim (ii) for layer $l - 1$ and claim (iii) for the case when $l - 1 = 1$. Now, let $Z \in \{\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l\}$. We have

$$\frac{\partial Z^{d\tilde{h}_2^{l-1}}}{\partial Z} = \frac{\partial Z^{d\tilde{x}_2^{l-1}}}{\partial Z} \sigma'(Z^{h_2^{l-1}}) + Z^{d\tilde{x}_2^{l-1}} \frac{\partial Z^{h_2^{l-1}}}{\partial Z} \sigma''(Z^{h_2^{l-1}})$$

where $Z^{h_2^{l-1}}$ and $Z^{d\tilde{x}_2^{l-1}}$ are polynomially bounded functions of Z_0 by the previous result and by Lemma L.5. Also by Lemma L.5, we have

$$\frac{\partial Z^{h_2^{l-1}}}{\partial Z} = \begin{cases} 0 & \text{if } l - 1 = 1 \text{ and } Z = \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2 \\ \psi(Z_0) & \text{otherwise} \end{cases}$$

with ψ polynomially bounded. In any case, $\partial Z^{h_2^{l-1}} / \partial Z$ is a polynomially bounded function of Z_0 . On the other hand, we have

$$\frac{\partial Z^{d\tilde{x}_2^{l-1}}}{\partial Z} = -\eta \overset{\circ}{\chi}_0 \mathbb{E}[Z^{d\tilde{h}_0^l} Z^{d\tilde{h}_2^l}] \frac{\partial Z^{\tilde{h}_0^{l-1}}}{\partial Z} \sigma'(Z^{\tilde{h}_0^{l-1}}) - \eta \overset{\circ}{\chi}_1 \mathbb{E}[Z^{d\tilde{h}_1^l} Z^{d\tilde{h}_2^l}] \frac{\partial Z^{x_1^{l-1}}}{\partial Z} \sigma'(Z^{x_1^{l-1}})$$

For both possible values of Z , $\partial Z^{\tilde{h}_0^{l-1}} / \partial Z$ has an easy expression and is a polynomially bounded function of Z_0 (essentially because σ and its derivatives are polynomially bounded). On the other hand, $\partial Z^{x_1^{l-1}} / \partial Z$ is a polynomially bounded function of Z_0 by Lemma L.4. σ' is polynomially bounded, and the expectations are finite by Lemma L.1 since the integrands are

polynomially bounded functions of Z_0 as they are products of such functions by Lemma L.4 and by the induction hypothesis. We thus get that

$$\frac{\partial Z^{d\tilde{x}_2^{l-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded. We thus have that:

$$\frac{\partial Z^{d\tilde{h}_2^{l-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded, which proves claims (iv) and (v) at layer $l - 1$. This thus concludes the induction, and with it the proof. \square

L.0.3 The case $t \geq 2$

We have now treated the base case $t = 2$ and are thus equipped to do the induction for $t \geq 2$. To make things easier we first introduce some equations. Let $t \geq 2$, we define the following assertions, where the different ψ appearing are assumed to be polynomially bounded:

Forward pass at time t :

$$(i) \quad Z^{h_t^1} = \psi \left(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_t}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right) \quad (\text{L.1})$$

For $l \in [2, L]$,

$$(i) \quad Z^{h_t^l} = \psi \left(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}, \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}} \right) \quad (\text{L.2})$$

$$(iii) \quad Z^{h_t^L} = \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right) \quad (\text{L.3})$$

For $l \in [2, L]$,

$$(iv) \quad \frac{\partial Z^{h_t^{l-1}}}{\widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l}} = \psi(Z_0) \quad (\text{L.4})$$

$$(v) \quad \frac{\partial Z^{h_t^l}}{\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}} = \psi(Z_0) \quad (\text{L.5})$$

$$(vi) \quad Z^{W^l(0)x_t^{l-1}} = 0 \quad (\text{L.6})$$

Backward pass at time t :

$$(i1) \quad Z^{d\tilde{x}_t^L} = \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right) \quad (\text{L.7})$$

$$(i2) \quad Z^{d\tilde{h}_t^L} = \psi \left(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}} \right) \quad (\text{L.8})$$

For $l \in [3, L]$,

$$(ii1) \quad Z^{d\tilde{x}_t^{l-1}} = \psi \left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}, \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l} \right) \quad (\text{L.9})$$

$$(ii2) \quad Z^{d\tilde{h}_t^{l-1}} = \psi \left(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{l-2}}, \widehat{Z}^{(\widehat{W}^l)^\top d\tilde{h}_0^l} \right) \quad (\text{L.10})$$

$$(iii1) \quad Z^{d\tilde{x}_t^1} = \psi \left(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_{t-1}}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right) \quad (\text{L.11})$$

$$(iii2) \quad Z^{d\tilde{h}_t^1} = \psi \left(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_t}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2} \right) \quad (\text{L.12})$$

For $l \in [2, L]$,

$$(iv) \quad \frac{\partial Z^{d\tilde{h}_t^l}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{l-1}}} = \psi(Z_0) \quad (\text{L.13})$$

$$(v) \quad \frac{\partial Z^{d\tilde{h}_t^{l-1}}}{\partial \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^l}} = \psi(Z_0) \quad (\text{L.14})$$

$$(vi) \quad Z^{(W^{l(0)})^\top d\tilde{h}_t^l} = 0 \quad (\text{L.15})$$

Note that we have proved in Appendix L.0.2 that all the assertions above hold for $t = 2$. Our goal is now to show by induction that they hold for any $t \geq 2$. For this we prove the following two lemmas. The proofs will essentially follow exactly the same pattern as for $t = 2$, the only difference is that the formulas will involve more terms, but since any finite sum of polynomially bounded functions is polynomially bounded, we will get the same results. Before proving the lemmas, we introduce the following quantities for $0 \leq s < t$:

For $l \in [2, L]$

$$\gamma_{s,t,l}^f := \begin{cases} \mathbb{E}[Z^{\tilde{x}_0^{l-1}} Z^{x_t^{l-1}}] & \text{if } s = 0 \\ \mathbb{E}[Z^{x_s^{l-1}} Z^{x_t^{l-1}}] & \text{otherwise} \end{cases} \quad (\text{L.16})$$

For $l \in [1, L - 1]$

$$\gamma_{s,t,l}^b := \mathbb{E}[Z^{d\tilde{h}_s^{l+1}} Z^{d\tilde{h}_t^{l+1}}] \quad (\text{L.17})$$

$\gamma_{s,t,l}^f$ (resp. $\gamma_{s,t,l}^b$) will appear when expressing the variables of the l -th layer at time t in the forward (resp. backward) pass. We will show in the proofs that as for $t = 1$ and $t = 2$, those expectations are finite by Lemma L.1 because the integrands are polynomially bounded functions of Z_0 as they are products of such variables.

Lemma L.7 (Induction step in IP-LLR, forward pass). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Let $t \geq 2$, and assume that all of the assertions of Equation (L.1) up until Equation (L.15) hold for every time step $s \in [2, t]$. Then, the assertions of the forward pass, i.e., from Equation (L.1) up until Equation (L.6), hold at time $t + 1$.*

Proof. We follow the proof of Lemma L.5. By Theorem G.7, we have

$$Z^{h_{t+1}^1} = Z^{U^1 \xi_{t+1}} - \eta \sum_{s=0}^t \overset{\circ}{\chi}_s (\xi_s^\top \xi_{t+1}) Z^{d\tilde{h}_s^1}$$

By Lemma L.4 $Z^{d\tilde{h}_1^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2})$ and by assumption we also have $Z^{d\tilde{h}_s^1} = \psi(Z^{U^1 \xi_0}, \dots, Z^{U^1 \xi_s}, \widehat{Z}^{(\widehat{W}^2)^\top d\tilde{h}_0^2})$ where the different ψ are polynomially bounded, which gives claim (i) at time $t + 1$.

We have

$$\frac{\partial Z^{h_{t+1}^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2} = -\eta \sum_{s=0}^t \overset{\circ}{\chi}_s(\xi_s^\top \xi_{t+1}) \frac{\partial Z^{d\bar{h}_s^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2}$$

with

$$\frac{\partial Z^{d\bar{h}_0^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2} = \sigma'(Z^{U^1 \xi_0})$$

which is a polynomially bounded function of Z_0 and so is $\partial Z^{d\bar{h}_1^1} / \partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2$ by Lemma L.4. In addition, by assumption, for $s \in [2, t]$, $\partial Z^{d\bar{h}_s^1} / \partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2 = \psi(Z_0)$ with ψ polynomially bounded. We thus get claim (iv) for $l = 2$ at time $t + 1$.

We have by Theorem G.7

$$Z^{h_{t+1}^2} = \overset{\circ}{\omega}_2 Z^{\widehat{W}^2 x_{t+1}^1} - \eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,2}^f Z^{d\bar{h}_s^2}$$

Now $Z^{x_{t+1}^1} = \sigma(Z^{h_{t+1}^1})$ is a polynomially bounded function of Z_0 because $Z^{h_{t+1}^1}$ is and σ is polynomially bounded. Secondly, we have

$$\frac{\partial Z^{x_{t+1}^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2} = \frac{\partial Z^{h_{t+1}^1}}{\partial \widehat{Z}(\widehat{W}^2)^\top d\bar{h}_0^2} \sigma'(Z^{h_{t+1}^1})$$

which is a polynomially bounded function of Z_0 by the previous results and because σ' is polynomially bounded. By Lemma L.2 we get that $\overset{\circ}{\omega}_2 Z^{\widehat{W}^2 x_{t+1}^1} = 0$ which gives claim (vi) for $l = 2$ at time $t + 1$. In addition, this yields

$$Z^{h_{t+1}^2} = -\eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,2}^f Z^{d\bar{h}_0^2} - \eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,2}^f Z^{d\bar{h}_s^2}$$

The expectations defining the γ^f are finite by Lemma L.1 since the integrands are polynomially bounded functions of Z_0 , as they are products of such variables by the previous result on $Z^{x_{t+1}^1}$ and by the assumption. This gives claim (ii) for $l = 2$ by the results for the backward pass at time $t = 0$ and $t = 1$ and by the assumptions. Let $Z \in \{\widehat{Z}^{\widehat{W}^2 \bar{x}_0^1}, \widehat{Z}(\widehat{W}^3)^\top d\bar{h}_0^3\}$. We have

$$\frac{\partial Z^{h_{t+1}^2}}{\partial Z} = -\eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,2}^f \frac{\partial Z^{d\bar{h}_s^2}}{\partial Z}$$

$\partial Z^{d\bar{h}_0^2} / \partial Z$ has a simple expression and is a polynomially bounded function of Z_0 . Additionally, by the results of the backward pass for $t = 1$, and by assumption, for $s \in [1, t]$, $\partial Z^{d\bar{h}_s^2} / \partial Z = \psi(Z_0)$ with ψ polynomially bounded. Since the γ^f are finite, we thus get

$$\frac{\partial Z^{h_{t+1}^2}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded. This gives claims (iv) and (v) at time $t + 1$.

Let $l \in [2, L - 1]$ and assume claims (ii), (iv), (v), and (vi) for layer l at time $t + 1$. Then, by [Theorem G.7](#) we have:

$$Z^{h_{t+1}^{l+1}} = \overset{\circ}{\omega}_{l+1} Z^{\widehat{W}^{l+1} x_{t+1}^l} - \eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,l+1}^f Z^{d\tilde{h}_s^{l+1}}$$

Now $Z^{x_{t+1}^l} = \sigma(Z^{h_{t+1}^l})$ is a polynomially bounded function of Z_0 because $Z^{h_{t+1}^l}$ is by the induction hypothesis and σ is polynomially bounded. Secondly, we have

$$\frac{\partial Z^{x_{t+1}^l}}{\partial \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}}} = \frac{\partial Z^{h_{t+1}^l}}{\partial \widehat{Z}^{(\widehat{W}^{l+1})^\top d\tilde{h}_0^{l+1}}} \sigma'(Z^{h_{t+1}^l})$$

which is a polynomially bounded function of Z_0 by the induction hypothesis. By [Lemma L.2](#) we get that $\overset{\circ}{\omega}_{l+1} Z^{\widehat{W}^{2l} x_{t+1}^l} = 0$ which gives claim (vi) for layer $l + 1$ at time $t + 1$. In addition, this yields

$$Z^{h_{t+1}^{l+1}} = -\eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,l+1}^f Z^{d\tilde{h}_s^{l+1}}$$

The expectations defining the γ^f are finite by [Lemma L.1](#) since the integrands are polynomially bounded functions of Z_0 , as they are products of such variables by the assumption and by the induction hypothesis. If $l+1 = L$, we have, for any $s \in [0, s]$, $Z^{d\tilde{h}_s^{l+1}} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ with ψ polynomially bounded, which shows

$$Z^{h_{t+1}^{l+1}} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$$

with ψ polynomially bounded, which gives claim (iii). If $l + 1 \leq L - 1$, for any $s \in [0, s]$, $Z^{d\tilde{h}_s^{l+1}} = \psi(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}, \widehat{Z}^{(\widehat{W}^{l+2})^\top d\tilde{h}_0^{l+2}})$ with ψ polynomially bounded, which shows

$$Z^{h_{t+1}^{l+1}} = \psi(\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}, \widehat{Z}^{(\widehat{W}^{l+2})^\top d\tilde{h}_0^{l+2}})$$

ψ polynomially bounded, which shows claim (ii) at layer $l + 1$ for time $t + 1$. Let $Z \in \{\widehat{Z}^{\widehat{W}^{l+1} \tilde{x}_0^l}, \widehat{Z}^{(\widehat{W}^{l+2})^\top d\tilde{h}_0^{l+2}}\}$. Note that the second value is only valid if $l + 1 \leq L - 1$. Whenever Z is well-defined, we have

$$\frac{\partial Z^{h_{t+1}^{l+1}}}{\partial Z} = -\eta \sum_{s=0}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,l+1}^f \frac{\partial Z^{d\tilde{h}_s^{l+1}}}{\partial Z}$$

For both possible values of Z , $\partial Z^{d\tilde{h}_0^{l+1}} / \partial Z$ has a simple expression and is a polynomially bounded function of Z_0 . $Z^{d\tilde{h}_1^{l+1}} / \partial Z$ is a polynomially bounded function of Z_0 by the results of the backward pass at time $t = 1$ ([Lemma L.4](#)), and finally for $s \in [2, t]$, $Z^{d\tilde{h}_s^{l+1}} / \partial Z = \psi(Z_0)$ with ψ polynomially bounded by assumption. Since the γ^f are finite, this gives

$$\frac{\partial Z^{h_{t+1}^{l+1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded. This proves claim (iv) and (v) for layer $l + 1$ at time $t + 1$, and thus concludes the induction on l and with it the proof. \square

Lemma L.8 (Induction step in IP-LLR, backward pass). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Let $t \geq 2$, and assume that all of the assertions of [Equation \(L.1\)](#) up until [Equation \(L.15\)](#) for every time step $s \in [2, t]$. Additionally assume that the assertions of the forward pass, i.e., from [Equation \(L.1\)](#) up until [Equation \(L.6\)](#), hold at time $t + 1$. Then, the assertions of the backward pass, i.e., from [Equation \(L.7\)](#) up until [Equation \(L.15\)](#), hold at time $t + 1$.*

Proof. We follow the proof of Lemma L.6. We have:

$$Z^{d\tilde{x}_{t+1}^L} = \widehat{Z}^{U^{L+1}} - \eta \overset{\circ}{\chi}_0 Z^{\tilde{x}_0^L} - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s Z^{x_t^L}$$

where $Z^{\tilde{x}_0^L} = \sigma(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$, $Z^{x_1^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ with ψ polynomially bounded by Lemma L.3 and for $s \in [2, t]$, $Z^{x_s^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ with ψ polynomially bounded by assumption. Combining all this gives

$$Z^{d\tilde{x}_{t+1}^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$$

with ψ polynomially bounded since σ is also polynomially bounded. Then

$$Z^{d\tilde{h}_{t+1}^L} = Z^{d\tilde{x}_{t+1}^L} \sigma'(Z^{h_{t+1}^L})$$

and since $Z^{h_{t+1}^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$ with ψ polynomially bounded by assumption, we get

$$Z^{d\tilde{h}_{t+1}^L} = \psi(\widehat{Z}^{U^{L+1}}, \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}})$$

with ψ polynomially bounded since σ' is also polynomially bounded. This thus proves claim (i). Now, we have

$$\frac{\partial Z^{d\tilde{h}_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} = \frac{\partial Z^{d\tilde{x}_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} \sigma'(Z^{h_{t+1}^L}) + Z^{d\tilde{x}_{t+1}^L} \frac{\partial Z^{h_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} \sigma''(Z^{h_{t+1}^L})$$

where $Z^{h_{t+1}^L}$, $\partial Z^{h_{t+1}^L} / \partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}$, and $Z^{d\tilde{x}_{t+1}^L}$ are polynomially bounded functions of Z_0 by assumption and by the previous result on $Z^{d\tilde{x}_{t+1}^L}$. Additionally, we have

$$\frac{\partial Z^{d\tilde{x}_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} = -\eta \overset{\circ}{\chi}_0 \sigma'(\widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}) - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \frac{\partial Z^{h_s^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} \sigma'(Z^{h_s^L})$$

σ' is polynomially bounded and by the results of the forward pass at $t = 1$ (Lemma L.3) $\partial Z^{h_1^L} / \partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}} = \psi(Z_0)$ with ψ polynomially bounded. In addition, by assumption, for any $s \in [2, t]$, $\partial Z^{h_s^L} / \partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}} = \psi(Z_0)$ with ψ polynomially bounded. This thus gives

$$\frac{\partial Z^{d\tilde{x}_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} = \psi(Z_0)$$

with ψ polynomially bounded, and thus

$$\frac{\partial Z^{d\tilde{h}_{t+1}^L}}{\partial \widehat{Z}^{\widehat{W}^l \tilde{x}_0^{L-1}}} = \psi(Z_0)$$

with ψ polynomially bounded since σ' and σ'' are polynomially bounded. This proves (iv) for $l = L$ at time $t + 1$.

We have:

$$Z^{d\tilde{x}_{t+1}^{L-1}} = \overset{\circ}{\omega}_L \widehat{Z}^{\widehat{W}^l} \overset{\circ}{\gamma} d\tilde{h}_{t+1}^L - \eta \overset{\circ}{\chi}_0 \overset{\circ}{\gamma}_{0,t+1,L-1}^b Z^{\tilde{x}_0^{L-1}} - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \overset{\circ}{\gamma}_{s,t+1,L-1}^b Z^{x_s^{L-1}}$$

From the previous step we have that both $Z^{d\tilde{h}_{t+1}^L}$ and $\partial Z^{d\tilde{h}_{t+1}^L} / \partial \widehat{Z}^{\widehat{W}^L \tilde{x}_0^{L-1}}$ are polynomially bounded functions of Z_0 . By Lemma L.2, this first shows that $\overset{\circ}{\omega}_L \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_{t+1}^L} = 0$, and thus gives (vi) for $l = L$, leading to:

$$Z^{d\tilde{x}_{t+1}^{L-1}} = -\eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,L}^b Z^{\tilde{x}_0^{L-1}} - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,L}^b Z^{x_s^{L-1}}$$

Now $Z^{\tilde{x}_0^{L-1}} = \sigma(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-2}})$ and we also have that $Z^{x_1^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-2}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$ with ψ polynomially bounded by Lemma L.3. In addition, we have for any $s \in [2, t]$, we get $Z^{x_s^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-2}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$ with ψ polynomially bounded by assumption since it is the case for $Z^{h_s^{L-1}}$ and σ is polynomially bounded. As always the expectations defining the γ^b are finite by Lemma L.1 because the integrands are polynomially bounded functions of Z_0 as products of such variables by the results for the backward pass at times $t = 0$ and $t = 1$, by the assumptions and by the previous result on $Z^{d\tilde{h}_{t+1}^L}$. Since σ is also polynomially bounded, this gives

$$Z^{d\tilde{x}_{t+1}^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-1}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$$

Then, we have

$$Z^{d\tilde{h}_{t+1}^{L-1}} = Z^{d\tilde{x}_{t+1}^{L-1}} \sigma'(Z^{h_{t+1}^{L-1}})$$

and since $Z^{h_{t+1}^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-1}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$ with ψ polynomially bounded by assumption

$$Z^{d\tilde{h}_2^{L-1}} = \psi(\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-1}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L})$$

with ψ polynomially bounded since σ' is also polynomially bounded. This thus proves claim (ii) for $l = L - 1$. Now, let $Z \in \{\widehat{Z}^{\widehat{W}^{l-1} \tilde{x}_0^{L-1}}, \widehat{Z}^{(\widehat{W}^L)^\top d\tilde{h}_0^L}\}$. We have

$$\frac{\partial Z^{d\tilde{h}_{t+1}^{L-1}}}{\partial Z} = \frac{\partial Z^{d\tilde{x}_{t+1}^{L-1}}}{\partial Z} \sigma'(Z^{h_{t+1}^{L-1}}) + Z^{d\tilde{x}_{t+1}^{L-1}} \frac{\partial Z^{h_{t+1}^{L-1}}}{\partial Z} \sigma''(Z^{h_{t+1}^{L-1}})$$

where $Z^{h_{t+1}^{L-1}}$, $\partial Z^{h_{t+1}^{L-1}} / \partial Z$, and $Z^{d\tilde{x}_{t+1}^{L-1}}$ are polynomially bounded functions of Z_0 by assumption and by the previous result. We have

$$\frac{\partial Z^{d\tilde{x}_{t+1}^{L-1}}}{\partial Z} = -\eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,L-1}^b \frac{\partial Z^{\tilde{h}_0^{L-1}}}{\partial Z} \sigma'(Z^{\tilde{h}_0^{L-1}}) - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,L-1}^b \frac{\partial Z^{h_s^{L-1}}}{\partial Z} \sigma'(Z^{h_s^{L-1}})$$

For both possible values of Z , $\partial Z^{\tilde{h}_0^{L-1}} / \partial Z$ has a simple expression and is a polynomially bounded function of Z_0 , as is \tilde{h}_0^{L-1} . In addition, $Z^{h_1^{L-1}}$ and $\partial Z^{h_1^{L-1}} / \partial Z$ are polynomially bounded functions of Z_0 by the results of the forward pass at $t = 1$, and finally, for $s \in [2, t]$, $Z^{h_s^{L-1}}$ and $\partial Z^{h_s^{L-1}} / \partial Z$ are polynomially bounded functions of Z_0 by assumption. Since the γ^b are finite and σ' is polynomially bounded, we get

$$\frac{\partial Z^{d\tilde{x}_{t+1}^{L-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded and thus

$$\frac{\partial Z^{d\tilde{h}_{t+1}^{L-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded since σ' and σ'' are polynomially bounded. This proves (iv) and (v) for $l = L - 1$.

Let $l \in [2, L - 1]$, and assume claims (ii), (iv), (v), are true at layer l and claim (vi) is true at layer $l + 1$. We have:

$$Z^{d\tilde{x}_{t+1}^{l-1}} = \overset{\circ}{\omega}_l \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_{t+1}^l - \eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,l-1}^b Z^{\tilde{x}_0^{l-1}} - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,l-1}^b Z^{x_s^{l-1}}$$

From the induction hypothesis we have that both $Z^{d\tilde{h}_{t+1}^l}$ and $\partial Z^{d\tilde{h}_{t+1}^l} / \partial \widehat{Z}(\widehat{W}^l)^\top d\tilde{x}_0^{l-1}$ are polynomially bounded functions of Z_0 . By Lemma L.2, this first shows that $\overset{\circ}{\omega}_l \widehat{Z}(\widehat{W}^l)^\top d\tilde{h}_{t+1}^l = 0$, and thus gives (vi) for layer l , leading to:

$$Z^{d\tilde{x}_{t+1}^{l-1}} = -\eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,l-1}^b Z^{\tilde{x}_0^{l-1}} - \eta \sum_{s=1}^t \overset{\circ}{\chi}_s \gamma_{s,t+1,l-1}^b Z^{x_s^{l-1}}$$

Now, if $l - 1 \geq 2$, $Z^{\tilde{x}_0^{l-1}} = \sigma(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2})$ and by Lemma L.3, we also have that $Z^{x_1^{l-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l)$ with ψ polynomially bounded because it is the case for $Z^{h_1^{l-1}}$ and σ is polynomially bounded. In addition, by assumption, we have for any $s \in [2, t]$, $Z^{x_s^{l-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l)$ with ψ polynomially bounded since it is the case for $Z^{h_s^{l-1}}$ and σ is polynomially bounded. As always the expectations defining the γ^b are finite by Lemma L.1 because the integrands are polynomially bounded functions of Z_0 as products of such variables by the results of the backward passes at times $t = 0$ and $t = 1$ and by the induction hypothesis. We thus get

$$Z^{d\tilde{x}_{t+1}^{l-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l)$$

with ψ polynomially bounded. On the other hand, if $l - 1 = 1$, we have $Z^{\tilde{x}_0^1} = \sigma(\widehat{Z}^{U^1 \xi_0})$ and by Lemma L.3, we have $Z^{x_1^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2)$ with ψ polynomially bounded. In addition, by assumption we have for $s \in [2, t]$, $Z^{x_s^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_s}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2)$ with ψ polynomially bounded. Since σ is also polynomially bounded, this gives

$$Z^{d\tilde{x}_{t+1}^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_t}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2)$$

ψ polynomially bounded. Since $Z^{d\tilde{h}_{t+1}^{l-1}} = Z^{d\tilde{x}_{t+1}^{l-1}} \sigma'(Z^{h_{t+1}^{l-1}})$, and by assumption $Z^{h_{t+1}^{l-1}} = \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l)$ if $l - 1 \geq 2$, and otherwise $Z^{h_{t+1}^1} = \psi(\widehat{Z}^{U^1 \xi_0}, \dots, \widehat{Z}^{U^1 \xi_{t+1}}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2)$, we get

$$Z^{d\tilde{h}_{t+1}^{l-1}} = \begin{cases} \psi(\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l) & \text{if } l - 1 \geq 2 \\ \psi(\widehat{Z}^{U^1 \xi_0}, \widehat{Z}^{U^1 \xi_1}, \dots, \widehat{Z}^{U^1 \xi_{t+1}}, \widehat{Z}(\widehat{W}^2)^\top d\tilde{h}_0^2) & \text{if } l - 1 = 1 \end{cases}$$

This gives claim (ii) for layer $l - 1$ and claim (iii) for the case when $l - 1 = 1$. Now, let $Z \in \{\widehat{Z}(\widehat{W}^{l-1})^\top d\tilde{x}_0^{l-2}, \widehat{Z}(\widehat{W}^L)^\top d\tilde{h}_0^l\}$. We have

$$\frac{\partial Z^{d\tilde{h}_{t+1}^{l-1}}}{\partial Z} = \frac{\partial Z^{d\tilde{x}_{t+1}^{l-1}}}{\partial Z} \sigma'(Z^{h_{t+1}^{l-1}}) + Z^{d\tilde{x}_{t+1}^{l-1}} \frac{\partial Z^{h_{t+1}^{l-1}}}{\partial Z} \sigma''(Z^{h_{t+1}^{l-1}})$$

where $Z^{h_{t+1}^{l-1}}$ and $Z^{d\tilde{x}_{t+1}^{l-1}}$ are polynomially bounded functions of Z_0 by assumption and by the previous result on $Z^{d\tilde{x}_{t+1}^{l-1}}$. Also by assumption, we have

$$\frac{\partial Z^{h_{t+1}^{l-1}}}{\partial Z} = \begin{cases} 0 & \text{if } l - 1 = 1 \text{ and } Z = \widehat{Z}(\widehat{W}^2)^\top d\tilde{x}_0^1 \\ \psi(Z_0) & \text{otherwise} \end{cases}$$

with ψ polynomially bounded. In any case, $\partial Z^{h_{t+1}^{l-1}}/\partial Z$ is a polynomially bounded function of Z_0 . On the other hand, we have

$$\frac{\partial Z^{d\tilde{x}_2^{l-1}}}{\partial Z} = -\eta \overset{\circ}{\chi}_0 \gamma_{0,t+1,l-1}^b \frac{\partial Z^{\tilde{h}_0^{l-1}}}{\partial Z} \sigma'(Z^{\tilde{h}_0^{l-1}}) - \eta \sum_{s=1}^t \overset{\circ}{\chi}_t \gamma_{s,t+1,l-1}^b \frac{\partial Z^{h_s^{l-1}}}{\partial Z} \sigma'(Z^{h_s^{l-1}})$$

For both possible values of Z , $\partial Z^{\tilde{h}_0^{l-1}}/\partial Z$ has an easy expression and is a polynomially bounded function of Z_0 (essentially because σ and its derivatives are polynomially bounded). On the other hand, $\partial Z^{h_s^{l-1}}/\partial Z$ is a polynomially bounded function of Z_0 by assumption. σ' is polynomially bounded, and the γ^b are finite. We thus get that

$$\frac{\partial Z^{d\tilde{x}_{t+1}^{l-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded, and thus:

$$\frac{\partial Z^{d\tilde{h}_{t+1}^{l-1}}}{\partial Z} = \psi(Z_0)$$

with ψ polynomially bounded, which proves claims (iv) and (v) at layer $l-1$ for time $t+1$. This thus concludes the induction on l , and with it the proof. \square

L.1 Main result

Theorem L.9 (Multiplications by the initial weight matrices vanish in IP-LLR for $t \geq 1$). *Consider the IP-LLR parameterization with a positively p -homogeneous activation function, and $p \geq 2$. Then, for any $t \geq 1$, and for any $l \in [2, L]$, one has:*

$$\begin{cases} Z^{W^l(0)x_t^{l-1}} = \overset{\circ}{\omega}_l Z^{\widehat{W}^l x_t^{l-1}} = 0 \\ Z^{(W^l(0))^\top d\tilde{h}_t^l} = \overset{\circ}{\omega}_l Z^{(\widehat{W}^l)^\top d\tilde{h}_t^l} = 0 \end{cases}$$

Proof. The result for $t = 1$ has essentially been proved already early on in Corollary L.4.1 (which stems from Lemmas L.3 and L.4). For $t = 2$, the result has been proved in Lemmas L.5 and L.6. Then we can prove the result for any $t \geq 2$ by induction using Lemmas L.7 and L.8. \square

M Expectations with ReLU

In all this section, we consider $Z \sim \mathcal{N}(0, \sigma^2)$, so that $Z = \sigma U$ where $U \sim \mathcal{N}(0, 1)$.

M.1 First moment

For $\phi(z) = \max(0, z)$ and $Z \sim \mathcal{N}(0, \sigma^2)$, we have

$$\mathbb{E}[\phi(Z)] = \mathbb{E}[\phi(\sigma U)] = \frac{\sigma}{\sqrt{2\pi}} \int_0^\infty u e^{-u^2/2} du = \frac{\sigma}{\sqrt{2\pi}}.$$

M.2 Second moment

For $\phi(z) = \max(0, z)$ and $Z \sim \mathcal{N}(0, \sigma^2)$, we have

$$\mathbb{E}[\phi(Z)^2] = \frac{1}{2} \mathbb{E}[Z^2] = \frac{\sigma^2}{2}.$$

M.3 First forward pass moments

We have, for any $l \in [1, L]$, with $\sigma_0 := \sqrt{|\xi_0|^2 + 1}$,

$$\begin{aligned}\mathbb{E}[\widehat{Z}^{\tilde{h}_0^l}] &= 0, & \mathbb{E}[(\widehat{Z}^{\tilde{h}_0^l})^2] &= \frac{\sigma_0^2}{2^{l-1}} \\ \mathbb{E}[Z^{\tilde{x}_0^l}] &= \frac{\sigma_0}{\sqrt{2^l \pi}}, & \mathbb{E}[(Z^{\tilde{x}_0^l})^2] &= \frac{\sigma_0^2}{2^l}\end{aligned}$$

M.4 First derivative moments

For $\phi(z) = \max(0, z)$, we have $\phi'(z) = \mathbb{1}_{z \geq 0}$ almost everywhere, so for $Z \sim \mathcal{N}(0, \sigma^2)$, we have

$$\mathbb{E}[\phi'(Z)] = \mathbb{P}(Z \geq 0) = 1/2.$$

Note that since $\phi'(z)^p = \phi'(z)$ for any $p > 0$, all the moments of $\phi'(Z)$ are equal to the first moment.

M.5 First backward pass moments

We have, for any $l \in [1, L]$, with,

$$\begin{aligned}\mathbb{E}[\widehat{Z}^{d\tilde{x}_0^l}] &= 0, & \mathbb{E}[(\widehat{Z}^{d\tilde{x}_0^l})^2] &= \frac{1}{2^{L-l}} \\ \mathbb{E}[Z^{d\tilde{h}_0^l}] &= 0, & \mathbb{E}[(Z^{d\tilde{h}_0^l})^2] &= \frac{1}{2^{L-l+1}}\end{aligned}$$

References

- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 6158–6169, 2019.
- Dyego Araújo, Roberto I. Oliveira, and Daniel Yukimura. A mean-field limit for certain deep neural networks, 2019.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Andrew Barron. Barron, a.e.: Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39:930–945, 1993. doi: 10.1109/18.256500.
- Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.
- Yoshua Bengio, Nicolas Roux, Pascal Vincent, Olivier Delalleau, and Patrice Marcotte. Convex neural networks. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2006. URL <https://proceedings.neurips.cc/paper/2005/file/0fc170ecbb8ff1afb2c6de48ea5343e7-Paper.pdf>.

- Erwin Bolthausen. An iterative construction of solutions of the TAP equations for the Sherrington–Kirkpatrick model. *Communications in Mathematical Physics*, 325(1):333–366, 2014.
- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3040–3050, 2018.
- Lénaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/ae614c557843b1df326cb29c57225459-Paper.pdf>.
- Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.07289>.
- Hadi Daneshmand, Jonas Kohler, Francis Bach, Thomas Hofmann, and Aurelien Lucchi. Batch normalization provably avoids ranks collapse for randomly initialised deep networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18387–18398. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/d5ade38a2c9f6f073d69e1bc6b6e64c1-Paper.pdf>.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019.
- Cong Fang, Jason D. Lee, Pengkun Yang, and Tong Zhang. Modeling from features: a mean-field framework for over-parameterized deep neural networks, 2020.
- Mario Geiger, Arthur Jacot, Stefano Spigler, Franck Gabriel, Levent Sagun, Stephane d’Ascoli, Giulio Biroli, Clement Hongler, and Matthieu Wyart. Scaling description of generalization with number of parameters in deep learning. *Journal Of Statistical Mechanics-Theory And Experiment*, 2020(ARTICLE):023401, 2020a.
- Mario Geiger, Stefano Spigler, Arthur Jacot, and Matthieu Wyart. Disentangling feature and lazy training in deep neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(11):113301, 2020b.
- Mario Geiger, Stefano Spigler, Arthur Jacot, and Matthieu Wyart. Disentangling feature and lazy training in deep neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(11):113301, 2020c.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of

- Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <http://proceedings.mlr.press/v9/glorot10a.html>.
- Boris Hanin and Mihai Nica. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations*, 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*, pages 1026–1034, 2015.
- Dan Hendrycks and Kevin Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *CoRR*, abs/1606.08415, 2016. URL <http://arxiv.org/abs/1606.08415>.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *CoRR*, abs/1806.07572, 2018. URL <http://arxiv.org/abs/1806.07572>.
- Arthur Jacot, Franck Gabriel, François Ged, and Clément Hongler. Order and chaos: Ntk views on dnn normalization, checkerboard and boundary artifacts. *arXiv preprint arXiv:1907.05715*, 2019.
- Vera Kurková and Marcello Sanguineti. Bounds on rates of variable-basis and neural-network approximation. *Information Theory, IEEE Transactions on*, 47:2659 – 2665, 10 2001. doi: 10.1109/18.945285.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33): E7665–E7671, 2018.
- Hrushikesh Mhaskar. On the tractability of multivariate integration and approximation by neural networks. *J. Complexity*, 20:561–590, 08 2004. doi: 10.1016/j.jco.2003.11.004.
- Radford M Neal. *BAYESIAN LEARNING FOR NEURAL NETWORKS*. PhD thesis, University of Toronto, 1995.
- Radford M Neal. Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer, 1996.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multilayer neural networks. *CoRR*, abs/2001.11443, 2020. URL <https://arxiv.org/abs/2001.11443>.
- Atsushi Nitanda and Taiji Suzuki. Stochastic particle gradient descent for infinite ensembles. *arXiv preprint arXiv:1712.05438*, 2017.
- Huy Tuan Pham and Phan-Minh Nguyen. A note on the global convergence of multilayer neural networks in the mean field regime. *CoRR*, abs/2006.09355, 2020. URL <https://arxiv.org/abs/2006.09355>.
- Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, 2018.

- Grant M. Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of neural networks: An interacting particle system approach, 2019.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 2021.
- E. Weinan and Stephan Wojtowytsch. On the banach spaces associated with multi-layer relu networks: Function representation, approximation theory and gradient descent dynamics. *ArXiv*, abs/2007.15623, 2020.
- Greg Yang. Wide feedforward or recurrent neural networks of any architecture are gaussian processes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/5e69fda38cda2060819766569fd93aa5-Paper.pdf>.
- Greg Yang. Tensor programs ii: Neural tangent kernel for any architecture. *ArXiv*, abs/2006.14548, 2020a.
- Greg Yang. Tensor programs III: neural matrix laws. *CoRR*, abs/2009.10685, 2020b. URL <https://arxiv.org/abs/2009.10685>.
- Greg Yang and Edward J. Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11727–11737. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21c.html>.