

---

# Implicit Regularization in Hierarchical Tensor Factorization and Deep Convolutional Neural Networks

---

Noam Razin<sup>†</sup> Asaf Maman<sup>†</sup> Nadav Cohen<sup>†</sup>

## Abstract

In the pursuit of explaining implicit regularization in deep learning, prominent focus was given to matrix and tensor factorizations, which correspond to simplified neural networks. It was shown that these models exhibit an implicit tendency towards low matrix and tensor ranks, respectively. Drawing closer to practical deep learning, the current paper theoretically analyzes the implicit regularization in hierarchical tensor factorization, a model equivalent to certain deep convolutional neural networks. Through a dynamical systems lens, we overcome challenges associated with hierarchy, and establish implicit regularization towards low hierarchical tensor rank. This translates to an implicit regularization towards locality for the associated convolutional networks. Inspired by our theory, we design explicit regularization discouraging locality, and demonstrate its ability to improve the performance of modern convolutional networks on non-local tasks, in defiance of conventional wisdom by which architectural changes are needed. Our work highlights the potential of enhancing neural networks via theoretical analysis of their implicit regularization.<sup>1</sup>

## 1 Introduction

One of the central mysteries in deep learning is the ability of neural networks to generalize when having far more learnable weights than training examples. The fact that this generalization takes place even in the absence of any explicit regularization (Zhang et al., 2017) has given rise to a common view by which gradient-based optimization induces an *implicit regularization* — a tendency to fit data with

<sup>†</sup>Blavatnik School of Computer Science, Tel Aviv University, Israel. Correspondence to: Noam Razin <noam-razin@mail.tau.ac.il>.

*Proceedings of the 39<sup>th</sup> International Conference on Machine Learning*, Baltimore, Maryland, USA, PMLR 162, 2022. Copyright 2022 by the author(s).

<sup>1</sup>Due to lack of space, a significant portion of the paper is deferred to the appendices. We refer the reader to Razin et al. (2022) for a self-contained version of the text.

predictors of low complexity (see, e.g., Neyshabur (2017)). Efforts to mathematically formalize this intuition have led to theoretical focus on *matrix* and *tensor factorizations*.<sup>2</sup>

Matrix factorization refers to minimizing a given loss (over matrices) by parameterizing the solution as a product of matrices, and optimizing the resulting objective via gradient descent. Tensor factorization is a generalization of this procedure to multi-dimensional arrays. There, a tensor is learned through gradient descent over a sum-of-outer-products parameterization. Various works analyzed the implicit regularization in matrix and tensor factorizations (Gunasekar et al., 2017; Li et al., 2018; Du et al., 2018; Arora et al., 2019; Razin & Cohen, 2020; Chou et al., 2020; Li et al., 2021a; Razin et al., 2021; Ge et al., 2021). Though initially conjectured to be equivalent to norm minimization (see Gunasekar et al. (2017)), recent studies (Arora et al., 2019; Razin & Cohen, 2020; Chou et al., 2020; Li et al., 2021a; Razin et al., 2021) suggest that this is not the case in general, and instead adopt a dynamical systems approach. They establish that gradient descent (with small learning rate and near-zero initialization) induces a momentum-like effect on the components of a factorization, leading them to move slowly when small and quickly when large. This implies a form of incremental learning that results in low matrix rank solutions for matrix factorization, and low tensor rank solutions for tensor factorization.

From a deep learning perspective, matrix factorization can be seen as a linear neural network (fully connected neural network with linear activation), and, in a similar vein, tensor factorization corresponds to a certain shallow (depth two) non-linear convolutional neural network (see Cohen et al. (2016b); Razin et al. (2021)). As theoretical surrogates for deep learning, the practical relevance of these models is limited. The former lacks non-linearity, while the latter misses depth — both crucial features of modern neural networks. A natural extension of matrix and tensor factorizations that accounts for both non-linearity and depth is *hierarchical tensor factorization*,<sup>3</sup> which corresponds to a

<sup>2</sup>The term “tensor factorization” refers throughout to the classic *CP factorization* (see Kolda & Bader (2009); Hackbusch (2012) for an introduction to tensor factorizations).

<sup>3</sup>The term “hierarchical tensor factorization” refers throughout to a variant of the Hierarchical Tucker factorization (Hackbusch & Kühn, 2009), presented in Section 3.

class of *deep non-linear* convolutional neural networks (Cohen et al., 2016b) that have demonstrated promising performance in practice (Cohen & Shashua, 2014; Cohen et al., 2016a; Sharir et al., 2016; Stoudenmire, 2018; Grant et al., 2018; Felser et al., 2021), and have been key to the study of expressiveness in deep learning (Cohen et al., 2016b; Cohen & Shashua, 2016; 2017; Cohen et al., 2017; 2018; Sharir & Shashua, 2018; Levine et al., 2018a;b; Balda et al., 2018; Khrukov et al., 2018; 2019; Levine et al., 2019).

In this paper, we conduct the first theoretical analysis of implicit regularization in hierarchical tensor factorization. As opposed to tensor factorization, which is a simple construct dating back to at least the early 20<sup>th</sup> century (Hitchcock, 1927), hierarchical tensor factorization was formally introduced only recently (Hackbusch & Kühn, 2009), and is much more elaborate. We circumvent the challenges brought forth by the added hierarchy through identification of *local components*, and characterization of their evolution under gradient descent (with small learning rate and near-zero initialization). The characterization reveals that they are subject to a momentum-like effect, identical to that in matrix and tensor factorizations. Accordingly, local components are learned incrementally, leading to solutions with low *hierarchical tensor rank* — a central concept in tensor analysis (Grasedyck, 2010; Grasedyck et al., 2013). Theoretical and empirical demonstrations validate our analysis.

For the deep convolutional networks corresponding to hierarchical tensor factorization, hierarchical tensor rank is known to measure the strength of dependencies modeled between spatially distant input regions (patches of pixels in the context of image classification) — see Cohen & Shashua (2017); Levine et al. (2018a;b). The established tendency towards low hierarchical tensor rank therefore implies a bias towards local (short-range) dependencies, in accordance with the fact that convolutional networks often struggle or completely fail to learn tasks entailing long-range dependencies (see, e.g., Wang et al. (2016); Linsley et al. (2018); Mlynarski et al. (2019); Hong et al. (2020); Kim et al. (2020)). However, while this failure is typically attributed solely to a limitation in expressive capability (*i.e.* to an inability of convolutional networks to represent functions modeling long-range dependencies — see Cohen & Shashua (2017); Linsley et al. (2018); Kim et al. (2020)), our analysis reveals that it also originates from implicit regularization. This suggests that the difficulty in learning long-range dependencies may be countered via *explicit* regularization, in contrast to conventional wisdom by which architectural modifications are needed. Through a series of controlled experiments we confirm this prospect, demonstrating that explicit regularization designed to promote high hierarchical tensor rank can significantly improve the performance of modern convolutional networks (*e.g.* ResNet18 and ResNet34 from He et al. (2016)) on tasks involving long-range dependencies.

Our results bring forth the possibility that deep learning ar-

chitectures considered suboptimal for certain tasks (*e.g.* convolutional networks for natural language processing tasks) may be greatly improved through a right choice of explicit regularization. Theoretical understanding of implicit regularization may be key to discovering such regularizers.

The remainder of the paper is organized as follows. Section 2 outlines existing dynamical characterizations of implicit regularization in matrix and tensor factorizations. Section 3 presents the hierarchical tensor factorization model, as well as its interpretation as a deep non-linear convolutional network. In Section 4 we characterize the dynamics of gradient descent over hierarchical tensor factorization, establishing that they lead to low hierarchical tensor rank. Section 5 explains why low hierarchical tensor rank means locality for a corresponding convolutional network. In Section 6 we demonstrate that the locality of modern convolutional networks can be countered using dedicated explicit regularization. Finally, Section 7 provides a summary. Due to lack of space we defer technical details behind hierarchical tensor factorization and its analysis to Appendices A and B respectively, and our review of related work to Appendix C.

## 2 Preliminaries: Matrix and Tensor Factorizations

To put our work into context, we overview known characterizations of implicit regularization in matrix and tensor factorizations.

Throughout the paper, when referring to a norm we mean the standard Frobenius (Euclidean) norm, denoted  $\|\cdot\|$ . For  $N \in \mathbb{N}$ , we let  $[N] := \{1, \dots, N\}$ . For vectors, matrices, or tensors, parenthesized superscripts denote elements in a collection, *e.g.*  $(\mathbf{w}^{(n)} \in \mathbb{R}^D)_{n=1}^N$ , while subscripts refer to entries, *e.g.*  $\mathbf{W}_{i,j} \in \mathbb{R}$  is the  $(i,j)$ 'th entry of  $\mathbf{W} \in \mathbb{R}^{D,D'}$ . An exception to this rule are the subscripts “M”, “T”, and “H”, which signify relation to matrix, tensor, and hierarchical tensor factorizations, respectively. A colon indicates all entries in an axis, *e.g.*  $\mathbf{W}_{i,:} \in \mathbb{R}^{D'}$  is the  $i$ 'th row and  $\mathbf{W}_{:,j} \in \mathbb{R}^D$  is the  $j$ 'th column of  $\mathbf{W}$ .

### 2.1 Matrix Factorization: Incremental Matrix Rank Learning

Consider the task of minimizing a differentiable and locally smooth<sup>4</sup> loss  $\mathcal{L}_M : \mathbb{R}^{D,D'} \rightarrow \mathbb{R}_{\geq 0}$ , where  $D, D' \in \mathbb{N}$ . For example,  $\mathcal{L}_M$  can be a matrix completion loss — mean squared error over observed entries from a ground truth matrix. Matrix factorization with hidden dimensions  $D_2, \dots, D_L \in \mathbb{N}$  refers to parameterizing the solution  $\mathbf{W}_M \in \mathbb{R}^{D,D'}$  as a product of  $L$  matrices, *i.e.* as  $\mathbf{W}_M = \mathbf{W}^{(1)} \dots \mathbf{W}^{(L)}$ , where  $\mathbf{W}^{(l)} \in \mathbb{R}^{D_l, D_{l+1}}$  for  $l = 1, \dots, L$ ,

<sup>4</sup>A differentiable function  $g : \mathbb{R}^D \rightarrow \mathbb{R}$  is *locally smooth* if for any compact subset  $\mathcal{B} \subset \mathbb{R}^D$  there exists  $\beta \in \mathbb{R}_{\geq 0}$  such that  $\|\nabla g(\mathbf{x}) - \nabla g(\mathbf{y})\| \leq \beta \cdot \|\mathbf{x} - \mathbf{y}\|$  for all  $\mathbf{x}, \mathbf{y} \in \mathcal{B}$ .




Figure 1: Dynamics of gradient descent over matrix, tensor, and hierarchical tensor factorizations — incremental learning leads to low matrix, tensor, and hierarchical tensor ranks, respectively. **Left:** top 10 singular values of the end matrix in a depth 3 matrix factorization when minimizing the mean squared error over observed entries from a matrix rank 5 ground truth (matrix completion loss). **Middle:** top 10 component norms of an order 3 tensor factorization when minimizing the mean squared error over observed entries from a tensor rank 5 ground truth (tensor completion loss). **Right:** top 10 local component norms at node  $\{1, 2, 3, 4\}$  of an order 4 hierarchical tensor factorization induced by a perfect binary mode tree, when minimizing the mean squared error over observed entries from a hierarchical tensor rank  $(5, 5, 5, 5, 5)$  ground truth (tensor completion loss). **All:** initial factorization weights were sampled independently from a zero-mean Gaussian distribution. Notice that, in accordance with existing analyses for matrix and tensor factorizations (Section 2) and our analysis for hierarchical tensor factorization (Section 4), the singular values, component norms, and local component norms move slowly when small and quickly when large, creating an incremental learning process that results in effectively low matrix, tensor, and hierarchical tensor rank solutions, respectively. In all factorizations this implicit regularization led to accurate reconstruction of the low rank ground truth (reconstruction errors were 0.001, 0.001, and 0.005, respectively). For further details such as loss definitions and factorization dimensions, as well as additional experiments for hierarchical tensor factorization, see Appendix G.

$D_1 := D$ , and  $D_{L+1} := D'$ , and minimizing the resulting objective  $\phi_{\mathbf{M}}(\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(L)}) := \mathcal{L}_{\mathbf{M}}(\mathbf{W}_{\mathbf{M}})$  using gradient descent. We call  $\mathbf{W}_{\mathbf{M}}$  the *end matrix* of the factorization. It is possible to explicitly constrain the values that  $\mathbf{W}_{\mathbf{M}}$  can take by limiting the hidden dimensions  $D_2, \dots, D_L$ . However, from an implicit regularization perspective, the case of interest is where the search space is unconstrained, thus we consider  $D_2, \dots, D_L \geq \min\{D, D'\}$ . Matrix factorization can be viewed as applying a linear neural network for minimizing  $\mathcal{L}_{\mathbf{M}}$ , and as such, serves a prominent theoretical model in deep learning (see *e.g.* Gunasekar et al. (2017); Du et al. (2018); Li et al. (2018); Arora et al. (2019); Gidel et al. (2019); Mulayoff & Michaeli (2020); Blanc et al. (2020); Gissin et al. (2020); Razin & Cohen (2020); Chou et al. (2020); Yun et al. (2021); Li et al. (2021a)).

Several characterizations of implicit regularization in matrix factorization have suggested that gradient descent, with small learning rate and near-zero initialization, induces a form of incremental matrix rank learning (Gidel et al., 2019; Gissin et al., 2020; Chou et al., 2020; Li et al., 2021a). Below we follow the presentation of Arora et al. (2019), which in line with other analyses, modeled small learning rate through the infinitesimal limit, *i.e.* via *gradient flow*:

$$\frac{d}{dt} \mathbf{W}^{(l)}(t) = -\frac{\partial}{\partial \mathbf{W}^{(l)}} \phi_{\mathbf{M}}(\mathbf{W}^{(1)}(t), \dots, \mathbf{W}^{(L)}(t))$$

for all  $t \geq 0$  and  $l \in [L]$ . Under gradient flow, the difference  $\mathbf{W}^{(l)}(t)^\top \mathbf{W}^{(l)}(t) - \mathbf{W}^{(l+1)}(t) \mathbf{W}^{(l+1)}(t)^\top$  remains constant through time for any  $l \in [L-1]$  (Arora et al., 2018). This implies that the *unbalancedness magnitude*, defined as  $\max_l \|\mathbf{W}^{(l)}(t)^\top \mathbf{W}^{(l)}(t) - \mathbf{W}^{(l+1)}(t) \mathbf{W}^{(l+1)}(t)^\top\|$ , does not change through time, thus becomes relatively small as optimization moves away from the origin, more so the closer initialization is to zero. Accordingly, it is common practice

to treat the case of unbalancedness magnitude zero as an idealization of standard near-zero initializations (see, *e.g.*, Saxe et al. (2014); Arora et al. (2018); Bartlett et al. (2018); Bah et al. (2019); Arora et al. (2019)).

With unbalancedness magnitude zero, the  $r$ 'th singular value of the end matrix  $\mathbf{W}_{\mathbf{M}}(t) = \mathbf{W}^{(1)}(t) \dots \mathbf{W}^{(L)}(t)$  ( $r \in [\min\{D, D'\}]$ ), denoted  $\sigma_{\mathbf{M}}^{(r)}(t) \in \mathbb{R}$ , evolves by (*cf.* Arora et al. (2019)):<sup>5</sup>

$$\frac{d}{dt} \sigma_{\mathbf{M}}^{(r)}(t) = \sigma_{\mathbf{M}}^{(r)}(t)^{2 - \frac{2}{L}} L \langle -\nabla \mathcal{L}_{\mathbf{M}}(\mathbf{W}_{\mathbf{M}}(t)), \mathcal{C}_{\mathbf{M}}^{(r)}(t) \rangle, \quad (1)$$

where  $\mathcal{C}_{\mathbf{M}}^{(r)}(t) := \mathbf{u}^{(r)}(t) \mathbf{v}^{(r)}(t)^\top \in \mathbb{R}^{D, D'}$  is the  $r$ 'th singular component of  $\mathbf{W}_{\mathbf{M}}(t)$ , meaning  $\mathbf{u}^{(r)}(t) \in \mathbb{R}^D$  and  $\mathbf{v}^{(r)}(t) \in \mathbb{R}^{D'}$  are, respectively, left and right singular vectors of  $\mathbf{W}_{\mathbf{M}}(t)$  corresponding to  $\sigma_{\mathbf{M}}^{(r)}(t)$ . As evident from Equation (1), two factors govern the evolution rate of a singular value  $\sigma_{\mathbf{M}}^{(r)}(t)$ . The first factor,  $\langle -\nabla \mathcal{L}_{\mathbf{M}}(\mathbf{W}_{\mathbf{M}}(t)), \mathcal{C}_{\mathbf{M}}^{(r)}(t) \rangle$ , is a projection of the singular component  $\mathcal{C}_{\mathbf{M}}^{(r)}(t)$  onto  $-\nabla \mathcal{L}_{\mathbf{M}}(\mathbf{W}_{\mathbf{M}}(t))$ , the direction of steepest descent with respect to the end matrix. The more the singular component is aligned with  $-\nabla \mathcal{L}_{\mathbf{M}}(\mathbf{W}_{\mathbf{M}}(t))$ , the faster the singular value grows. The second, more critical factor, is  $\sigma_{\mathbf{M}}^{(r)}(t)^{2 - \frac{2}{L}} L$ , which implies that the rate of change of the singular value is proportional to its size exponentiated by  $2 - 2/L$ . This brings rise to a momentum-like effect, which attenuates the movement of small singular values and accelerates the movement of large ones. We may thus expect that if the matrix factorization is initialized near the origin, singular values progress slowly at first, and

<sup>5</sup>The dynamical characterization of singular values in Equation (1) requires  $\mathcal{L}_{\mathbf{M}}$  to be analytic, a property met by standard loss functions such as the square and cross-entropy losses.

then, one after the other they reach a critical threshold and quickly rise, until convergence is attained. Such incremental learning phenomenon leads to low matrix rank solutions. It is demonstrated empirically in Figure 1 (left), which reproduces an experiment from Arora et al. (2019). We note that under certain technical conditions, the incremental matrix rank learning phenomenon can be used to prove exact matrix rank minimization (Li et al., 2021a).

## 2.2 Tensor Factorization: Incremental Tensor Rank Learning

A depth two matrix factorization boils down to parameterizing a sought-after solution as a sum of tensor (outer) products between column vectors of  $\mathbf{W}^{(1)}$  and row vectors of  $\mathbf{W}^{(2)}$ . Namely, since  $\mathbf{W}_M = \mathbf{W}^{(1)}\mathbf{W}^{(2)}$  we may write  $\mathbf{W}_M = \sum_{r=1}^R \mathbf{W}_{:,r}^{(1)} \otimes \mathbf{W}_{r,:}^{(2)}$ , where  $R$  is the dimension shared between  $\mathbf{W}^{(1)}$  and  $\mathbf{W}^{(2)}$ , and  $\otimes$  stands for the tensor product.<sup>6</sup> Note that the minimal number of summands  $R$  required for  $\mathbf{W}_M$  to express a given matrix  $\mathbf{W}$  is precisely the latter’s matrix rank.

By allowing each summand to be a tensor product of more than two vectors, we may transition from a factorization for matrices to a factorization for tensors. In tensor factorization, a sought-after solution  $\mathcal{W}_T \in \mathbb{R}^{D_1, \dots, D_N}$  — an order  $N \geq 3$  tensor with *modes* (axes) of *dimensions*  $D_1, \dots, D_N \in \mathbb{N}$  — is parameterized as  $\mathcal{W}_T = \sum_{r=1}^R \mathbf{W}_{:,r}^{(1)} \otimes \dots \otimes \mathbf{W}_{:,r}^{(N)}$ , where  $\mathbf{W}^{(n)} \in \mathbb{R}^{D_n, R}$  for  $n \in [N]$ . Each term  $\mathbf{W}_{:,r}^{(1)} \otimes \dots \otimes \mathbf{W}_{:,r}^{(N)}$  in this sum is called a *component*, and  $\mathcal{W}_T$  is referred to as the *end tensor* of the factorization. Given a differentiable and locally smooth loss  $\mathcal{L}_T : \mathbb{R}^{D_1, \dots, D_N} \rightarrow \mathbb{R}_{\geq 0}$ , e.g. mean squared error over observed entries from a ground truth tensor (i.e. a tensor completion loss), the goal is to minimize the objective  $\phi_T(\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)}) := \mathcal{L}_T(\mathcal{W}_T)$ . In analogy with matrix factorization, the minimal number of components  $R$  required for  $\mathcal{W}_T$  to express a given tensor  $\mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N}$  is defined to be the latter’s *tensor rank*, and the case of interest is when  $R$  is sufficiently large to not restrict tensor rank (i.e. to admit an unconstrained search space).

Similarly to how matrix factorization corresponds to a linear neural network, tensor factorization is known (see Cohen et al. (2016b); Razin et al. (2021)) to be equivalent to a certain shallow (depth two) non-linear convolutional network (with multiplicative non-linearity). By virtue of this equivalence, illustrated in Figure 2 (top), tensor factorization is considered closer to practical deep learning than matrix factorization.

As in matrix factorization, gradient flow over tensor factor-

<sup>6</sup>Given  $\mathcal{U} \in \mathbb{R}^{D_1, \dots, D_N}$ ,  $\mathcal{V} \in \mathbb{R}^{H_1, \dots, H_K}$ , their tensor product  $\mathcal{U} \otimes \mathcal{V} \in \mathbb{R}^{D_1, \dots, D_N, H_1, \dots, H_K}$  is defined by  $[\mathcal{U} \otimes \mathcal{V}]_{d_1, \dots, d_{N+K}} = \mathcal{U}_{d_1, \dots, d_N} \cdot \mathcal{V}_{d_{N+1}, \dots, d_{N+K}}$ . For vectors  $\mathbf{u} \in \mathbb{R}^D$ ,  $\mathbf{v} \in \mathbb{R}^{D'}$ , the tensor product  $\mathbf{u} \otimes \mathbf{v}$  is equal to  $\mathbf{u}\mathbf{v}^\top \in \mathbb{R}^{D, D'}$ .

ization induces invariants of optimization. In particular, the differences between squared norms of vectors in the same component, i.e.  $\|\mathbf{W}_{:,r}^{(n)}(t)\|^2 - \|\mathbf{W}_{:,r'}^{(n)}(t)\|^2$  for  $n, n' \in [N]$  and  $r \in [R]$ , are constant through time (Razin et al., 2021). This leads to the following definition of unbalancedness magnitude:  $\max_{n, n', r} \|\|\mathbf{W}_{:,r}^{(n)}(t)\|^2 - \|\mathbf{W}_{:,r'}^{(n')}(t)\|^2\|$ , which does not change during optimization, therefore remains small throughout if initialization is close to the origin. Under the idealized assumption of unbalancedness magnitude zero (corresponding to infinitesimally small initialization), the norm of the  $r$ ’th component in the factorization ( $r \in [R]$ ), i.e.  $\sigma_T^{(r)}(t) := \|\otimes_{n=1}^N \mathbf{W}_{:,r}^{(n)}(t)\|$ , evolves by (cf. Razin et al. (2021)):

$$\frac{d}{dt} \sigma_T^{(r)}(t) = \sigma_T^{(r)}(t)^{2 - \frac{2}{N}} N \langle -\nabla \mathcal{L}_T(\mathcal{W}_T(t)), \mathcal{C}_T^{(r)}(t) \rangle, \quad (2)$$

where  $\mathcal{C}_T^{(r)}(t) := \otimes_{n=1}^N \overline{\mathbf{W}}_{:,r}^{(n)}(t)$ , with  $\overline{\mathbf{W}}_{:,r}^{(n)}(t)$  defined as  $\mathbf{W}_{:,r}^{(n)}(t) / \|\mathbf{W}_{:,r}^{(n)}(t)\|$  for all  $n \in [N]$  (by convention, if  $\mathbf{W}_{:,r}^{(n)}(t) = 0$  then  $\overline{\mathbf{W}}_{:,r}^{(n)}(t) = 0$ ), denotes the  $r$ ’th normalized component. Comparing Equation (2) to Equation (1) reveals that the evolution rate of a component norm in tensor factorization is *structurally identical* to that of a singular value in matrix factorization. Specifically, it is determined by two factors, analogous to those in Equation (1): (i) a projection of the normalized component  $\mathcal{C}_T^{(r)}(t)$  onto  $-\nabla \mathcal{L}_T(\mathcal{W}_T(t))$ , which encourages growth of components that are aligned with the direction of steepest descent with respect to the end tensor; and (ii)  $\sigma_T^{(r)}(t)^{2 - \frac{2}{N}} N$ , which induces a momentum-like effect, leading component norms to move slower when small and faster when large. This suggests that, in analogy with matrix factorization, components tend to be learned incrementally, yielding a bias towards low tensor rank. Figure 1 (middle) demonstrates the phenomenon empirically, reproducing an experiment from Razin et al. (2021). Similarly to the case of matrix factorization, under certain technical conditions, the incremental tensor rank learning phenomenon can be used to prove exact tensor rank minimization (Razin et al., 2021).

## 3 Hierarchical Tensor Factorization

In this section we present the hierarchical tensor factorization model. Due to lack of space we defer the formal definitions to Appendix A, and provide here an overview introducing the central concepts.

As discussed in Section 2.2, tensor factorization produces an order  $N$  end tensor through a sum of components, each combining  $N$  vectors using the tensor product operator. It is customary to represent this computation through a shallow tree structure with  $N$  leaves, corresponding to the weight matrices  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)}$ , that are directly connected to the root, which computes the end tensor  $\mathcal{W}_T = \sum_{r=1}^R \mathbf{W}_{:,r}^{(1)} \otimes \dots \otimes \mathbf{W}_{:,r}^{(N)}$ .

Generalizing this scheme to an arbitrary tree gives rise to



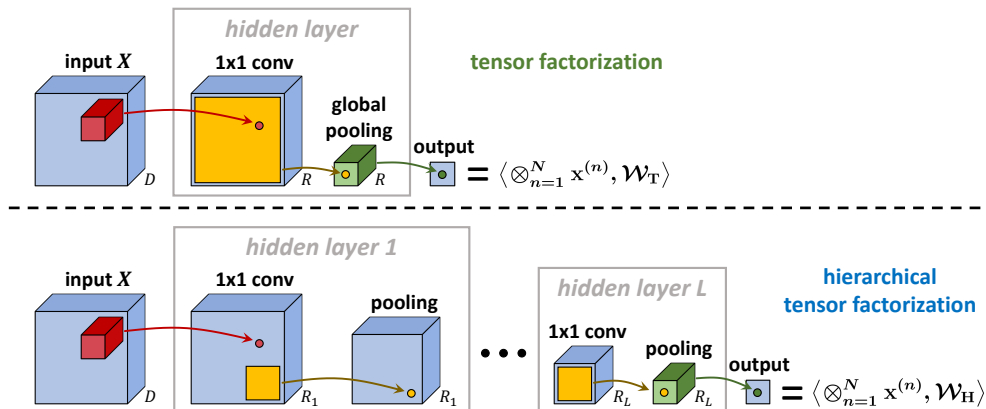


Figure 2: Tensor factorization corresponds to a class of shallow (depth two) non-linear convolutional networks, while hierarchical tensor factorization corresponds to a class of *deep* non-linear convolutional networks. These correspondences have been studied extensively (see references in Section 3). For completeness, we briefly describe them herein and provide a formal proof in Appendix D. **Top:** the shallow network equivalent to tensor factorization processes an input  $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) \in \mathbb{R}^{D_1} \times \dots \times \mathbb{R}^{D_N}$  (illustration assumes  $D_1 = \dots = D_N = D$  to avoid clutter) using a single hidden layer, which consists of: (i) locally connected linear operator with  $R$  channels, computing  $(\mathbf{W}^{(1)})^\top \mathbf{x}^{(1)}, \dots, (\mathbf{W}^{(N)})^\top \mathbf{x}^{(N)}$  with learnable weights  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)}$  (this operator is referred to as “1 × 1 conv” in appeal to the common case of weight sharing, *i.e.*  $\mathbf{W}^{(1)} = \dots = \mathbf{W}^{(N)}$ ); and (ii) channel-wise global product pooling (multiplicative non-linearity). Summing over the resulting activations then yields the scalar output  $\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \sum_{r=1}^R \otimes_{n=1}^N \mathbf{W}^{(n)}_{:,r} \rangle = \langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_T \rangle$ . Hence, functions realized by this class of networks are naturally represented via tensor factorization, where the number of components  $R$  and the weight matrices  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)}$  of the factorization correspond to the width and learnable weights of the network, respectively. **Bottom:** for a hierarchical tensor factorization induced by a perfect  $P$ -ary mode tree, the equivalent network is a deep variant of that associated with tensor factorization. It has  $L = \log_P N$  hidden layers instead of just one, with channel-wise product pooling operating over windows of size  $P$  as opposed to globally. After passing an input  $(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) \in \mathbb{R}^{D_1} \times \dots \times \mathbb{R}^{D_N}$  through all hidden layers, a final linear layer produces the network’s scalar output  $\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \rangle$ , where  $\mathcal{W}_H$  is the end tensor of the hierarchical tensor factorization, whose weight matrices are equal to the network’s learnable weights. Thus, functions realized by this class of networks are naturally represented via hierarchical tensor factorization. We note that, as shown in Cohen & Shashua (2016), by considering *generalized hierarchical tensor factorizations* it is possible to account for various non-linearities beyond multiplicative, *i.e.* for product pooling being converted to a different pooling operator (*e.g.* max or average), optionally preceded by a non-linear activation (*e.g.* rectified linear unit).

hierarchical tensor factorization. Given a tree, or formally, a *mode tree* of the hierarchical tensor factorization (Definition 2 in Appendix A), the scheme progresses from leaves to root. Each internal node combines tensors produced by its children to form higher-order tensors, until finally the root outputs an order  $N$  end tensor. Different mode trees bring about different hierarchical tensor factorizations, which are essentially a composition of many local tensor factorizations, each corresponding to a different location in the mode tree. We refer to the components of these local tensor factorizations as the *local components* (Definition 3 in Appendix A) of the hierarchical factorization — see Figure 3 for an illustration.

A mode tree of a hierarchical tensor factorization induces a notion of rank called *hierarchical tensor rank* (Definition 5 in Appendix A). The hierarchical tensor rank is a tuple whose entries correspond to locations in the mode tree. The value held by an entry is characterized by the number of local components at the corresponding location, similarly to how tensor rank is characterized by the number of components in a tensor factorization (see Section 2.2). Motivated by matrix and tensor ranks being implicitly minimized in matrix and tensor factorizations, respectively (Section 2), in Section 4 we explore the possibility of hierarchical tensor

rank being implicitly minimized in hierarchical tensor factorization. That is, we investigate the prospect of gradient flow with near-zero initialization over hierarchical tensor factorization learning solutions that can be represented with few local components at all locations of the mode tree.

**Equivalence to a class of deep non-linear convolutional networks** As discussed in Section 2, matrix factorization can be seen as a linear neural network, and, in a similar vein, tensor factorization corresponds to a certain shallow (depth two) non-linear convolutional network. A drawback of these models as theoretical surrogates for deep learning is that the former lacks non-linearity, while the latter misses depth. Hierarchical tensor factorization accounts for both of these limitations: for appropriate mode trees, it is known (see Cohen et al. (2016b)) to be equivalent to a class of deep non-linear convolutional networks (with multiplicative non-linearity). These networks have demonstrated promising performance in practice (Cohen & Shashua, 2014; Cohen et al., 2016a; Sharir et al., 2016; Stoudenmire, 2018; Grant et al., 2018; Felser et al., 2021), and their equivalence to hierarchical tensor factorization has been key to the study of expressiveness in deep learning (Cohen et al., 2016b; Cohen & Shashua, 2016; 2017; Cohen et al., 2017; 2018; Levine et al., 2018a;b; Balda et al., 2018; Khulkov et al., 2018;




Figure 3: Hierarchical tensor factorization consists of multiple local tensor factorizations. **Left:** tensor factorization represents an order  $N$  tensor as a sum of components, each combining  $N$  vectors through the tensor product operator. Accordingly, it is represented by a shallow tree where all leaves are directly connected to the root. **Right:** hierarchical tensor factorization adheres to an arbitrary tree structure (figure depicts a perfect binary tree), producing an order  $N$  tensor by iteratively combining multiple local tensor factorizations. The components of the local tensor factorizations constituting the hierarchical tensor factorization are defined to be its local components. For a formal description of hierarchical tensor factorization see Appendix A.

2019). The equivalence is illustrated in Figure 2 (bottom) and rigorously proven in Appendix D.

#### 4 Analysis: Incremental Hierarchical Tensor Rank Learning

In this section we theoretically analyze the implicit regularization in hierarchical tensor factorization. Our analysis extends known results for matrix and tensor factorizations outlined in Section 2. In particular, we show that the implicit regularization in hierarchical tensor factorization induces an incremental learning process that results in low hierarchical tensor rank, similarly to how matrix and tensor factorizations incrementally learn solutions with low matrix and tensor ranks, respectively. To facilitate this extension, while overcoming the challenges arising from the complexity of the hierarchical tensor factorization model, we characterize the evolution of the local components introduced in Section 3. Due to lack of space we defer the formal treatment to Appendix B, and provide here an overview covering the main results.

As discussed in Section 2, for both matrix and tensor factorizations, there exists an invariant of optimization whose deviation from zero is referred to as unbalancedness magnitude, and it is common to treat the case of unbalancedness magnitude zero as an idealization of standard near-zero initializations. With unbalancedness magnitude zero, singular values in a matrix factorization evolve by Equation (1), and component norms in a tensor factorization move per Equation (2). Equations (1) and (2) are structurally identical, and are interpreted as implying incremental learning of singular values and component norms, respectively, *i.e.* of matrix and tensor ranks, respectively. This interpretation was initially supported by experiments (such as those reported in Figure 1 (left and middle)), and later via proofs of exact matrix and tensor rank minimization under certain technical conditions.

In Appendix B.1 we show that in analogy with matrix and tensor factorizations, hierarchical tensor factorization entails an invariant of optimization (Lemma 2), which leads to a corresponding notion of unbalancedness magnitude (Defini-

tion 6). For the canonical case of unbalancedness magnitude zero (corresponding to standard near-zero initializations), we prove the following.

**Theorem** (informal version of Theorem 1 in Appendix B.1). *The norm of the  $r$ 'th local component associated with node  $\nu$  in the mode tree, denoted  $\sigma_H^{(\nu,r)}(t)$ , evolves by:*

$$\frac{d}{dt} \sigma_H^{(\nu,r)}(t) = \sigma_H^{(\nu,r)}(t)^{2 - \frac{2}{L_\nu}} L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), C_H^{(\nu,r)}(t) \rangle, \quad (3)$$

where  $L_\nu$  is the number of weight vectors in the local component, and  $C_H^{(\nu,r)}(t)$  is the direction it imposes on the end tensor  $\mathcal{W}_H(t)$  of the hierarchical tensor factorization.

Appendix E generalizes the above theorem by relieving the assumption of unbalancedness magnitude zero. Namely, it establishes that Equation (3) holds approximately when unbalancedness magnitude at initialization is small. Equation (3) is structurally identical to Equations (1) and (2), therefore the evolution rate of a local component norm in hierarchical tensor factorization mirrors the evolution rates of a singular value in matrix factorization and a component norm in tensor factorization. One is thus led to interpret Equation (3) as implying incremental learning of local component norms, *i.e.* of hierarchical tensor rank (see Section 3). We support this interpretation through experiments analogous to those typically conducted for supporting the interpretation of Equations (1) and (2) as implying incremental learning of matrix and tensor ranks, respectively — see Figure 1 (right) as well as Appendix G. Moreover, we consider technical conditions similar to those assumed for proving exact matrix and tensor rank minimization by matrix and tensor factorizations, respectively, and establish theoretical results aimed at facilitating a proof of exact hierarchical tensor rank minimization — see Appendix B.2. Completing the missing steps for deriving such a proof is regarded as a promising direction for future work.

We note in passing that hierarchical tensor rank is partially (not totally) ordered, and as such, represents a new type of complexity measure that may further our understanding of implicit regularization in deep learning. See Appendix B.3 for details.

## 5 Low Hierarchical Tensor Rank Implies Locality

In Section 4 we established that the implicit regularization in hierarchical tensor factorization favors solutions with low hierarchical tensor rank. A natural question that arises is what are the implications of this tendency for the class of deep convolutional networks equivalent to hierarchical tensor factorization (illustrated in Figure 2 (bottom)). It is known (Cohen & Shashua, 2017; Levine et al., 2018a;b) that for this class of networks, hierarchical tensor rank measures the strength of dependencies modeled between spatially distant input regions (patches of pixels in the context of image classification) — see brief explanation in Section 5.1 below, and formal derivation in Appendix F. An implicit regularization towards low hierarchical tensor rank thus implies a bias towards local (short-range) dependencies. While seemingly benign, this observation is shown in Section 6 to bring forth a practical method for improving performance of contemporary convolutional networks (e.g. ResNet18 and ResNet34 from He et al. (2016)) on tasks with long-range dependencies.

### 5.1 Locality via Separation Rank

Given a multivariate function  $f$  with scalar output, a popular measure of dependencies between a set of input variables and its complement is known as *separation rank*. The separation rank, formally presented in Definition 1 below, was originally introduced in Beylkin & Mohlenkamp (2002), and has since been employed for various applications (Harrison et al., 2003; Hackbusch, 2006; Beylkin et al., 2009), as well as analyses of expressiveness in deep learning (Cohen & Shashua, 2017; Cohen et al., 2017; Levine et al., 2018a;b; 2020; Wies et al., 2021; Levine et al., 2022). It is also prevalent in quantum mechanics, where it serves as a measure of entanglement (Levine et al., 2018b).

Consider the convolutional network equivalent to a hierarchical tensor factorization with mode tree  $\mathcal{T}$ . It turns out (see formal derivation in Appendix F) that for functions realized by this network, separation ranks measuring dependencies between distinct regions of the input are precisely equal to entries of the hierarchical tensor rank with respect to  $\mathcal{T}$  (recall that, as discussed in Section 3, the hierarchical tensor rank is a tuple). Thus, low hierarchical tensor rank implies that the separation ranks are low, which in turn means that dependencies modeled between distinct input regions are weak, *i.e.* that only local dependencies are prominent.

**Definition 1.** The *separation rank* of  $f : \times_{n=1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  with respect to  $I \subset [N]$ , denoted  $\text{sep}(f; I)$ , is the minimal  $R \in \mathbb{N} \cup \{0\}$  for which there exist  $g_1, \dots, g_R : \times_{i \in I} \mathbb{R}^{D_i} \rightarrow \mathbb{R}$  and  $\bar{g}_1, \dots, \bar{g}_R : \times_{j \in [N] \setminus I} \mathbb{R}^{D_j} \rightarrow \mathbb{R}$  such that:

$$f(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \sum_{r=1}^R g_r((\mathbf{x}^{(i)})_{i \in I}) \cdot \bar{g}_r((\mathbf{x}^{(j)})_{j \in [N] \setminus I}).$$

**Interpretation** The separation rank of  $f$  with respect to  $I$  is the minimal number of summands required to express  $f$ , where each summand is a product of two functions — one that operates over variables indexed by  $I$ , and another that operates over the remaining variables. If  $\text{sep}(f; I) = 1$ , the function is separable, meaning it does not model any interaction between the sets of variables. In a statistical setting, where  $f$  is a probability density function, this would mean that  $(\mathbf{x}^{(i)})_{i \in I}$  and  $(\mathbf{x}^{(j)})_{j \in [N] \setminus I}$  are statistically independent. The higher  $\text{sep}(f; I)$  is, the farther  $f$  is from separability, *i.e.* the stronger the dependencies it models between  $(\mathbf{x}^{(i)})_{i \in I}$  and  $(\mathbf{x}^{(j)})_{j \in [N] \setminus I}$ .

## 6 Countering Locality of Convolutional Networks via Regularization

Convolutional networks often struggle or completely fail to learn tasks that entail strong dependencies between spatially distant regions of the input (e.g. patches of pixels in image classification, or tokens in natural language processing tasks) — see Wang et al. (2016); Linsley et al. (2018); Mlynarski et al. (2019); Hong et al. (2020); Kim et al. (2020). Conventional wisdom attributes this failure to the local nature of convolutional network architectures, *i.e.* to their inability to express long-range dependencies (see Cohen & Shashua (2017); Linsley et al. (2018); Kim et al. (2020)). This suggests that addressing the problem requires modifying the architecture. Our theory reveals that there is also an implicit regularization at play, giving rise to the possibility of countering the locality of convolutional networks via *explicit* regularization, without modifying their architecture. In the current section we affirm this possibility, demonstrating that carefully designed regularization can greatly improve the performance of contemporary convolutional networks on tasks involving long-range dependencies. For brevity, we defer some implementation details and experiments to Appendix G.

We conducted a series of experiments, using the ubiquitous ResNet18 and ResNet34 convolutional networks (He et al., 2016), over two types of image classification datasets in which the distance between salient regions can be controlled. The first type, referred to as "IsSameClass," comprises datasets we constructed, where the goal is to predict whether two randomly sampled CIFAR10 (Krizhevsky, 2009) images are of the same class or not. Each classified input is a  $32 \times 224$  image filled with zeros, over which the CIFAR10 images are placed (symmetrically around the center) at a predetermined distance from each other (to comply with ResNets, inputs were padded to have size  $224 \times 224$ ). By increasing the predetermined distance between CIFAR10 images, we produce datasets that require stronger modeling of long-range dependencies. The second type of datasets we experiment with is taken from the Pathfinder challenge (Linsley et al., 2018; Kim et al., 2020; Tay et al., 2021) — a standard benchmark for modeling




Figure 4: Samples from IsSameClass and Pathfinder datasets. For further details on their creation process see Appendix G.2.2. **Left:** positive and negative samples from IsSameClass datasets with 0 and 160 pixels between images, respectively. The label is 1 if the two CIFAR10 images are of the same class, and 0 otherwise. For the sake of illustration, background is displayed as white instead of black, and padding is not shown (*i.e.* only the raw  $32 \times 224$  input is presented). **Right:** positive and negative samples from Pathfinder challenge (Linsley et al., 2018) datasets with connecting path lengths 6 and 9, respectively. A connecting path is one that joins the two circles, and if present, its length is measured in the number of dashes. The label of a sample is 1 if it includes a connecting path (*i.e.* if the two circles are connected), and 0 otherwise.

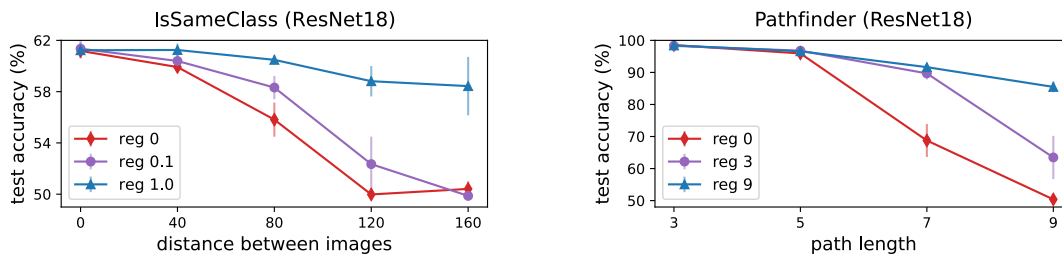


Figure 5: Dedicated explicit regularization can counter the locality of convolutional networks, significantly improving performance on tasks with long-range dependencies. Plots present test accuracies achieved by a randomly initialized ResNet18 over IsSameClass (left) and Pathfinder (right) datasets, with varying spatial distances between salient regions of the input (CIFAR10 images in IsSameClass and connected circles in Pathfinder — see Figure 4). For each dataset, the network was trained via stochastic gradient descent to minimize a regularized objective, consisting of the binary cross-entropy loss and the dedicated regularization described in Section 6.1. The legend specifies the regularization coefficients used. Markers and error bars report means and standard deviations, respectively, taken over five different runs for the corresponding combination of dataset and regularization coefficient. As expected, when increasing the (spatial) range of dependencies required to be modeled, the test accuracy obtained by an unregularized network (regularization coefficient zero) substantially deteriorates, reaching the vicinity of the trivial value 50%. Conventional wisdom attributes this failure to a limitation in the expressive capability of convolutional networks (*i.e.* to their inability to represent functions modeling long-range dependencies). However, as can be seen, applying the dedicated regularization significantly improved performance, without any architectural modification. Appendix G provides further implementation details, as well as additional experiments: (*i*) using ResNet34; and (*ii*) showing similar improvements when the baseline network (“reg 0”) is already regularized via standard techniques (weight decay or dropout).

long-range dependencies. In Pathfinder, each image contains two white circles and multiple dashed paths (curves) over a black background, and the goal is to predict whether the circles are connected by a path or not. The length of connecting paths is predetermined, allowing control over the spatial range of dependencies required to be modeled. Representative examples from IsSameClass and Pathfinder datasets are displayed in Figure 4.

Figure 5 shows (for ResNet18; similar results for ResNet34 are reported in Appendix G) that when fitting IsSameClass and Pathfinder datasets, increasing the strength of long-range dependencies (*i.e.* the distance between CIFAR10 images in IsSameClass, and the connecting path length in Pathfinder) leads to significant degradation in test accuracy, to the point of it being not much better than chance. This phenomenon complies with existing evidence from Cohen & Shashua (2017); Linsley et al. (2018); Kim et al. (2020) showing failure of convolutional networks in learning tasks with long-range dependencies. However, while

these works address the problem by modifying the architecture, we tackle it through explicit regularization (described in Section 6.1 below) designed to promote high separation ranks implying long-range dependencies between image regions (see Section 5). As evident in Figure 5, our regularization significantly improves test accuracy. This suggests that the tendency towards locality of modern convolutional networks may in large part be due to implicit regularization, not an inherent limitation of expressive power as often believed. Our findings showcase that deep learning architectures considered suboptimal for certain tasks may be greatly improved through a right choice of explicit regularization. Theoretical understanding of implicit regularization may be key to discovering such regularizers.

## 6.1 Explicit Regularization Promoting Long-Range Dependencies

We describe below the explicit regularization applied in our experiments to counter the locality of convolutional net-



works. We emphasize that this regularization is based on our theory, and merely serves as an example to how the performance of convolutional networks on tasks involving long-range dependencies can be improved without modifying their architecture. Further evaluation and improvement of our regularization are regarded as promising directions for future work.

Denote by  $f_{\Theta}(\mathbf{X})$  the output of a neural network, where  $\Theta$  stands for its learnable weights, and  $\mathbf{X} := (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$  represents an input image, with each  $\mathbf{x}^{(n)}$  standing for a pixel. Suppose we are given a subset of indices  $I \subset [N]$ , with complement  $J := [N] \setminus I$ , and we would like to encourage the network to learn a function  $f_{\Theta}$  that models strong dependence between  $\mathbf{X}_I := (\mathbf{x}^{(i)})_{i \in I}$  (pixels indexed by  $I$ ) and  $\mathbf{X}_J := (\mathbf{x}^{(j)})_{j \in J}$  (those indexed by  $J$ ). As discussed in Section 5.1, a standard measure of such dependence is the separation rank, provided in Definition 1. If the separation rank of  $f_{\Theta}$  with respect to  $I$  is one, meaning no dependence between  $\mathbf{X}_I$  and  $\mathbf{X}_J$  is modeled, then we may write  $f_{\Theta}(\mathbf{X}) = g(\mathbf{X}_I) \cdot \bar{g}(\mathbf{X}_J)$  for some functions  $g$  and  $\bar{g}$ . This implies that  $\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X}) = \bar{g}(\mathbf{X}_J) \cdot \nabla g(\mathbf{X}_I)$ , meaning that a change in  $\mathbf{X}_J$  (with  $\mathbf{X}_I$  held fixed) does not affect the direction of  $\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X})$ , only its magnitude (and possibly its sign). This observation suggests that, in order to learn a function  $f_{\Theta}$  modeling strong dependence between  $\mathbf{X}_I$  and  $\mathbf{X}_J$ , one may add a regularization term that promotes a change in the direction of  $\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X})$  whenever  $\mathbf{X}_J$  is altered (with  $\mathbf{X}_I$  held fixed).

The regularization applied in our experiments is of the type outlined above, with  $I$  and  $J$  chosen to promote long-range dependencies. Namely, at each iteration of stochastic gradient descent we randomly choose disjoint subsets of indices  $I$  and  $J$  corresponding to contiguous (distinct) image regions. Then, for each image  $\mathbf{X}$  in the iteration’s batch, we let  $\mathbf{X}'$  be the result of replacing the pixels in  $\mathbf{X}$  indexed by  $J$  with alternative values taken from a different image in the training set. Finally, we compute  $|\langle \nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X}), \nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X}') \rangle| \cdot \|\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X})\|^{-1} \|\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X}')\|^{-1}$  — (absolute value of) cosine of the angle between  $\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X})$  and  $\nabla_{\mathbf{X}_I} f_{\Theta}(\mathbf{X}')$  — average it across the batch, multiply the average by a constant coefficient, and add the result to the minimized objective.<sup>7</sup> For further details see Appendix G.2.2.

## 7 Summary

Incremental matrix rank learning in matrix factorization (Arora et al., 2019; Gidel et al., 2019; Gissin et al., 2020; Chou et al., 2020; Li et al., 2021a) and incremental tensor rank learning in tensor factorization (Razin et al., 2021; Ge et al., 2021) were important discoveries on the

<sup>7</sup>Each artificially generated image  $\mathbf{X}'$  is used only to compute the regularization term, not as an additional training instance incurring its own loss. Our proposed regularization is therefore fundamentally different from data augmentation.

path to explaining implicit regularization in deep learning. The current paper takes an additional step along this path, establishing incremental hierarchical tensor rank learning in hierarchical tensor factorization. It circumvents the complexity of the hierarchical tensor factorization model by introducing the notion of local components, and theoretically analyzing their evolution throughout optimization. Experiments validate the theory.

While matrix factorization corresponds to linear neural networks and tensor factorization to certain shallow (depth two) non-linear convolutional neural networks (Cohen et al., 2016b; Razin et al., 2021), hierarchical tensor factorization is equivalent to a class of *deep* non-linear convolutional neural networks (Cohen et al., 2016b). It therefore jointly accounts for both non-linearity and depth — two critical aspects of deep learning. For the class of convolutional networks equivalent to hierarchical tensor factorization, low hierarchical tensor rank translates to weak modeling of dependencies between spatially distant input regions (Cohen & Shashua, 2017; Levine et al., 2018a;b). Our theory thus suggests an implicit regularization towards locality in convolutional networks. While the locality of convolutional networks is widely accepted, it is typically ascribed to expressive properties determined by their architecture (see, e.g., Cohen & Shashua (2017); Linsley et al. (2018); Kim et al. (2020)). The fact that implicit regularization also plays a role indicates that it might be possible to counter this locality via explicit regularization. We verify this prospect empirically, demonstrating that explicit regularization designed to promote high hierarchical tensor rank vastly improves the performance of modern convolutional networks (ResNet18 and ResNet34 from He et al. (2016)) on tasks with long-range dependencies.

Taken together, the theory and experiments presented in this paper bring forth the possibility that deep learning architectures considered suboptimal for certain tasks (e.g. convolutional networks for natural language processing tasks) may be greatly improved through a right choice of explicit regularization. Theoretical understanding of implicit regularization may be key to discovering such regularizers.

## Acknowledgments

This work was supported by a Google Research Scholar Award, a Google Research Gift, the Yandex Initiative in Machine Learning, the Israel Science Foundation (grant 1780/21), Len Blavatnik and the Blavatnik Family Foundation, and Amnon and Anat Shashua. NR is supported by the Apple Scholars in AI/ML and the Tel Aviv University Center for AI and Data Science (TAD) PhD fellowships.

## References

Arora, S., Cohen, N., and Hazan, E. On the optimization of deep networks: Implicit acceleration by overparameterization. In

- International Conference on Machine Learning (ICML)*, pp. 244–253, 2018.
- Arora, S., Cohen, N., Hu, W., and Luo, Y. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 7413–7424, 2019.
- Azulay, S., Moroshko, E., Nacson, M. S., Woodworth, B., Srebro, N., Globerson, A., and Soudry, D. On the implicit bias of initialization shape: Beyond infinitesimal mirror descent. *International Conference on Machine Learning (ICML)*, 2021.
- Bah, B., Rauhut, H., Terstiege, U., and Westdickenberg, M. Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. *arXiv preprint arXiv:1910.05505*, 2019.
- Balda, E. R., Behboodi, A., and Mathar, R. A tensor analysis on dense connectivity via convolutional arithmetic circuits. *Preprint*, 2018.
- Bartlett, P., Helmbold, D., and Long, P. Gradient descent with identity initialization efficiently learns positive definite linear transformations. In *International Conference on Machine Learning (ICML)*, pp. 520–529, 2018.
- Beylkin, G. and Mohlenkamp, M. J. Numerical operator calculus in higher dimensions. *Proceedings of the National Academy of Sciences*, 99(16):10246–10251, 2002.
- Beylkin, G., Garcke, J., and Mohlenkamp, M. J. Multivariate regression and machine learning with sums of separable functions. *SIAM Journal on Scientific Computing*, 31(3):1840–1857, 2009.
- Blanc, G., Gupta, N., Valiant, G., and Valiant, P. Implicit regularization for deep neural networks driven by an ornstein-uhlenbeck like process. In *Conference on Learning Theory (COLT)*, 2020.
- Chizat, L. and Bach, F. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory (COLT)*, pp. 1305–1338, 2020.
- Chou, H.-H., Gieshoff, C., Maly, J., and Rauhut, H. Gradient descent for deep matrix factorization: Dynamics and implicit bias towards low rank. *arXiv preprint arXiv:2011.13772*, 2020.
- Chou, H.-H., Maly, J., and Rauhut, H. More is less: Inducing sparsity via overparameterization. *arXiv preprint arXiv:2112.11027*, 2021.
- Cohen, N. and Shashua, A. Simnets: A generalization of convolutional networks. *Advances in Neural Information Processing Systems (NeurIPS), Deep Learning Workshop*, 2014.
- Cohen, N. and Shashua, A. Convolutional rectifier networks as generalized tensor decompositions. *International Conference on Machine Learning (ICML)*, 2016.
- Cohen, N. and Shashua, A. Inductive bias of deep convolutional networks through pooling geometry. *International Conference on Learning Representations (ICLR)*, 2017.
- Cohen, N., Sharir, O., and Shashua, A. Deep simnets. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016a.
- Cohen, N., Sharir, O., and Shashua, A. On the expressive power of deep learning: A tensor analysis. *Conference On Learning Theory (COLT)*, 2016b.
- Cohen, N., Sharir, O., Levine, Y., Tamari, R., Yakira, D., and Shashua, A. Analysis and design of convolutional networks via hierarchical tensor decompositions. *Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI) Special Issue on Deep Learning Theory*, 2017.
- Cohen, N., Tamari, R., and Shashua, A. Boosting dilated convolutional networks with mixed tensor decompositions. *International Conference on Learning Representations (ICLR)*, 2018.
- Da Silva, C. and Herrmann, F. J. Optimization on the hierarchical tucker manifold—applications to tensor completion. *Linear Algebra and its Applications*, 481:131–173, 2015.
- Du, S. S., Hu, W., and Lee, J. D. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 384–395, 2018.
- Eftekhari, A. and Zygalakis, K. Limitations of implicit bias in matrix sensing: Initialization rank matters. *arXiv preprint arXiv:2008.12091*, 2021.
- Ergen, T. and Pilanci, M. Implicit convex regularizers of cnn architectures: Convex optimization of two-and three-layer networks in polynomial time. *International Conference on Learning Representations (ICLR)*, 2021.
- Felser, T., Trenti, M., Sestini, L., Gianelle, A., Zuliani, D., Lucchesi, D., and Montangelo, S. Quantum-inspired machine learning on high-energy physics data. *npj Quantum Information*, 7(1):1–8, 2021.
- Ge, R., Ren, Y., Wang, X., and Zhou, M. Understanding deflation process in over-parametrized tensor decomposition. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Gidel, G., Bach, F., and Lacoste-Julien, S. Implicit regularization of discrete gradient dynamics in linear neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3196–3206, 2019.
- Gissin, D., Shalev-Shwartz, S., and Daniely, A. The implicit bias of depth: How incremental learning drives generalization. *International Conference on Learning Representations (ICLR)*, 2020.
- Grant, E., Benedetti, M., Cao, S., Hallam, A., Lockhart, J., Stojevic, V., Green, A. G., and Severini, S. Hierarchical quantum classifiers. *npj Quantum Information*, 4(1):1–8, 2018.
- Grasedyck, L. Hierarchical singular value decomposition of tensors. *SIAM Journal on Matrix Analysis and Applications*, 31(4):2029–2054, 2010.
- Grasedyck, L., Kressner, D., and Tobler, C. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.
- Gunasekar, S., Woodworth, B. E., Bhojanapalli, S., Neyshabur, B., and Srebro, N. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6151–6159, 2017.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 9461–9471, 2018.

- Hackbusch, W. On the efficient evaluation of coalescence integrals in population balance models. *Computing*, 78(2):145–159, 2006.
- Hackbusch, W. *Tensor spaces and numerical tensor calculus*, volume 42. Springer, 2012.
- Hackbusch, W. and Kühn, S. A new scheme for the tensor representation. *Journal of Fourier analysis and applications*, 15(5): 706–722, 2009.
- HaoChen, J. Z., Wei, C., Lee, J., and Ma, T. Shape matters: Understanding the implicit bias of the noise covariance. In *Conference on Learning Theory (COLT)*, 2021.
- Harrison, R. J., Fann, G. I., Yanai, T., and Beylkin, G. Multiresolution quantum chemistry in multiwavelet bases. In *International Conference on Computational Science*, pp. 103–110. Springer, 2003.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 770–778, 2016.
- Hitchcock, F. L. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4): 164–189, 1927.
- Hong, D., Gao, L., Yao, J., Zhang, B., Plaza, A., and Chanussot, J. Graph convolutional networks for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 2020.
- Hu, W., Xiao, L., Adlam, B., and Pennington, J. The surprising simplicity of the early-time learning dynamics of neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Jagadeesan, M., Razenshteyn, I., and Gunasekar, S. Inductive bias of multi-channel linear convolutional networks with bounded weight norm. *arXiv preprint arXiv:2102.12238*, 2021.
- Ji, Z. and Telgarsky, M. Gradient descent aligns the layers of deep linear networks. *International Conference on Learning Representations (ICLR)*, 2019a.
- Ji, Z. and Telgarsky, M. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory (COLT)*, pp. 1772–1798, 2019b.
- Ji, Z. and Telgarsky, M. Directional convergence and alignment in deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Kargas, N. and Sidiropoulos, N. D. Nonlinear system identification via tensor completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 4420–4427, 2020.
- Kargas, N. and Sidiropoulos, N. D. Supervised learning and canonical decomposition of multivariate functions. *IEEE Transactions on Signal Processing*, 69:1097–1107, 2021.
- Khrulkov, V., Novikov, A., and Oseledets, I. Expressive power of recurrent neural networks. *International Conference on Learning Representations (ICLR)*, 2018.
- Khrulkov, V., Hrinchuk, O., and Oseledets, I. Generalized tensor models for recurrent neural networks. *International Conference on Learning Representations (ICLR)*, 2019.
- Kim, J., Linsley, D., Thakkar, K., and Serre, T. Disentangling neural mechanisms for perceptual grouping. *International Conference on Learning Representations (ICLR)*, 2020.
- Kohn, K., Merkh, T., Montúfar, G., and Trager, M. Geometry of linear convolutional networks. *arXiv preprint arXiv:2108.01538*, 2021.
- Kolda, T. G. Multilinear operators for higher-order decompositions. Technical report, 2006.
- Kolda, T. G. and Bader, B. W. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Levine, Y., Sharir, O., and Shashua, A. Benefits of depth for long-term memory of recurrent networks. *International Conference on Learning Representations (ICLR) Workshop*, 2018a.
- Levine, Y., Yakira, D., Cohen, N., and Shashua, A. Deep learning and quantum entanglement: Fundamental connections with implications to network design. *International Conference on Learning Representations (ICLR)*, 2018b.
- Levine, Y., Sharir, O., Cohen, N., and Shashua, A. Quantum entanglement in deep learning architectures. *To appear in Physical Review Letters*, 2019.
- Levine, Y., Wies, N., Sharir, O., Bata, H., and Shashua, A. Limits to depth efficiencies of self-attention. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Levine, Y., Wies, N., Jannai, D., Navon, D., Hoshen, Y., and Shashua, A. The inductive bias of in-context learning: Rethinking pretraining example design. *International Conference on Learning Representations (ICLR)*, 2022.
- Li, Y., Ma, T., and Zhang, H. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Proceedings of the 31st Conference On Learning Theory (COLT)*, pp. 2–47, 2018.
- Li, Z., Luo, Y., and Lyu, K. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. *International Conference on Learning Representations (ICLR)*, 2021a.
- Li, Z., Wang, T., and Arora, S. What happens after sgd reaches zero loss?—a mathematical framework. *arXiv preprint arXiv:2110.06914*, 2021b.
- Linsley, D., Kim, J., Veerabadran, V., Windolf, C., and Serre, T. Learning long-range spatial dependencies with horizontal gated recurrent units. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Lyu, K. and Li, J. Gradient descent maximizes the margin of homogeneous neural networks. *International Conference on Learning Representations (ICLR)*, 2020.
- Lyu, K., Li, Z., Wang, R., and Arora, S. Gradient descent on two-layer nets: Margin maximization and simplicity bias. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Milanesi, P., Kadri, H., Ayache, S., and Artières, T. Implicit regularization in deep tensor factorization. In *International Joint Conference on Neural Networks (IJCNN)*, 2021.

- Min, H., Tarmoun, S., Vidal, R., and Mallada, E. On the explicit role of initialization on the convergence and implicit bias of overparameterized linear networks. *International Conference on Machine Learning (ICML)*, 2021.
- Mlynarski, P., Delingette, H., Criminisi, A., and Ayache, N. 3d convolutional neural networks for tumor segmentation using long-range 2d context. *Computerized Medical Imaging and Graphics*, 73:60–72, 2019.
- Moroshko, E., Gunasekar, S., Woodworth, B., Lee, J. D., Srebro, N., and Soudry, D. Implicit bias in deep linear classification: Initialization scale vs training accuracy. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Mulayoff, R. and Michaeli, T. Unique properties of wide minima in deep networks. In *International Conference on Machine Learning (ICML)*, 2020.
- Mulayoff, R., Michaeli, T., and Soudry, D. The implicit bias of minima stability: A view from function space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Nacson, M. S., Gunasekar, S., Lee, J., Srebro, N., and Soudry, D. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In *International Conference on Machine Learning (ICML)*, 2019a.
- Nacson, M. S., Lee, J., Gunasekar, S., Savarese, P. H. P., Srebro, N., and Soudry, D. Convergence of gradient descent on separable data. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, 2019b.
- Neyshabur, B. Implicit regularization in deep learning. *PhD thesis*, 2017.
- Oymak, S. and Soltanolkotabi, M. Overparameterized nonlinear learning: Gradient descent takes the shortest path? In *International Conference on Machine Learning (ICML)*, pp. 4951–4960, 2019.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- Pesme, S., Pillaud-Vivien, L., and Flammarion, N. Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Rauhut, H., Schneider, R., and Stojanac, Ž. Low rank tensor recovery via iterative hard thresholding. *Linear Algebra and its Applications*, 523:220–262, 2017.
- Razin, N. and Cohen, N. Implicit regularization in deep learning may not be explainable by norms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Razin, N., Maman, A., and Cohen, N. Implicit regularization in tensor factorization. *International Conference on Machine Learning (ICML)*, 2021.
- Razin, N., Maman, A., and Cohen, N. Implicit regularization in hierarchical tensor factorization and deep convolutional neural networks. *arXiv preprint arXiv:2201.11729*, 2022.
- Sarussi, R., Brutzkus, A., and Globerson, A. Towards understanding learning in neural networks with linear teachers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *International Conference on Learning Representations (ICLR)*, 2014.
- Shachaf, G., Brutzkus, A., and Globerson, A. A theoretical analysis of fine-tuning with linear teachers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Sharir, O. and Shashua, A. On the expressive power of overlapping architectures of deep learning. *International Conference on Learning Representations (ICLR)*, 2018.
- Sharir, O., Tamari, R., Cohen, N., and Shashua, A. Tensorial mixture models. *arXiv preprint*, 2016.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Steinlechner, M. Riemannian optimization for high-dimensional tensor completion. *SIAM Journal on Scientific Computing*, 38(5):S461–S484, 2016.
- Stoudenmire, E. M. Learning relevant features of data with multi-scale tensor networks. *Quantum Science and Technology*, 3(3):034003, 2018.
- Tay, Y., Dehghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D. Long range arena: A benchmark for efficient transformers. *International Conference on Learning Representations (ICLR)*, 2021.
- Teschl, G. *Ordinary differential equations and dynamical systems*, volume 140. American Mathematical Soc., 2012.
- Vardi, G. and Shamir, O. Implicit regularization in relu networks with the square loss. In *Conference on Learning Theory (COLT)*, 2021.
- Vardi, G., Shamir, O., and Srebro, N. On margin maximization in linear and relu networks. *arXiv preprint arXiv:2110.02732*, 2021.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., and Van Gool, L. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pp. 20–36. Springer, 2016.
- Wies, N., Levine, Y., Jannai, D., and Shashua, A. Which transformer architecture fits my data? a vocabulary bottleneck in self-attention. *International Conference on Machine Learning (ICML)*, 2021.
- Woodworth, B., Gunasekar, S., Lee, J. D., Moroshko, E., Savarese, P., Golan, I., Soudry, D., and Srebro, N. Kernel and rich regimes in overparameterized models. In *Conference on Learning Theory (COLT)*, pp. 3635–3673, 2020.
- Yun, C., Krishnan, S., and Mobahi, H. A unifying view on implicit bias in training linear neural networks. *International Conference on Learning Representations (ICLR)*, 2021.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations (ICLR)*, 2017.



## A Formal Presentation of Hierarchical Tensor Factorization

In this appendix we formally present the hierarchical tensor factorization model, which was informally overviewed in Section 3.

The structure of a hierarchical tensor factorization is determined by a mode tree.

**Definition 2.** Let  $N \in \mathbb{N}$ . A *mode tree*  $\mathcal{T}$  over  $[N]$  is a rooted tree in which:

- every node is labeled by a subset of  $[N]$ ;
- there are exactly  $N$  leaves, labeled  $\{1\}, \dots, \{N\}$ ; and
- the label of an interior (non-leaf) node is the union of the labels of its children.

We identify nodes with their labels, *i.e.* with the corresponding subsets of  $[N]$ , and accordingly treat  $\mathcal{T}$  as a subset of  $2^{[N]}$ . Furthermore, we denote the set of all interior nodes by  $\text{int}(\mathcal{T}) \subset \mathcal{T}$ , the parent of a non-root node  $\nu \in \mathcal{T} \setminus \{[N]\}$  by  $Pa(\nu) \in \mathcal{T}$ , and the children of  $\nu \in \text{int}(\mathcal{T})$  by  $C(\nu) \subset \mathcal{T}$ . When enumerating over children of a node, *i.e.* over  $C(\nu)$  for  $\nu \in \text{int}(\mathcal{T})$ , an arbitrary fixed ordering is assumed.

One may consider various mode trees, each leading to a different hierarchical tensor factorization. Notable choices include: (i) a shallow tree (comprising only leaves and root), which reduces the hierarchical tensor factorization to a tensor factorization (Section 2.2); and (ii) a perfect binary tree (applicable if  $N$  is a power of 2) whose corresponding hierarchical tensor factorization is perhaps the most extensively studied. Figure 6(a) illustrates these two choices.

Along with a mode tree  $\mathcal{T}$ , what defines a hierarchical tensor factorization are the number of local components at each interior node, denoted  $(R_\nu \in \mathbb{N})_{\nu \in \text{int}(\mathcal{T})}$ . The induced hierarchical tensor factorization is parameterized by weight matrices  $(\mathbf{W}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}})_{\nu \in \mathcal{T}}$ , where  $R_{Pa([N])} := 1$  and  $R_{\{1\}} := D_1, \dots, R_{\{N\}} := D_N$ . It creates the *end tensor*  $\mathcal{W}_H \in \mathbb{R}^{D_1, \dots, D_N}$  by constructing intermediate tensors of increasing order while traversing  $\mathcal{T}$  from leaves to root as follows:

for all  $\nu \in \{\{1\}, \dots, \{N\}\}$  and  $r \in [R_{Pa(\nu)}]$ :

$$\underbrace{\mathcal{W}^{(\nu, r)}}_{\text{order } 1} := \mathbf{W}_{:,r}^{(\nu)},$$

for all  $\nu \in \text{int}(\mathcal{T}) \setminus \{[N]\}$  and  $r \in [R_{Pa(\nu)}]$  (traverse interior nodes of  $\mathcal{T}$  from leaves to root, non-inclusive):

$$\underbrace{\mathcal{W}^{(\nu, r)}}_{\text{order } |\nu|} := \pi_\nu \left( \sum_{r'=1}^{R_\nu} \mathbf{W}_{r',r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right), \quad (4)$$

$$\underbrace{\mathcal{W}_H}_{\text{order } N} := \pi_{[N]} \left( \sum_{r'=1}^{R_{[N]}} \mathbf{W}_{r',1}^{([N])} \left[ \otimes_{\nu_c \in C([N])} \mathcal{W}^{(\nu_c, r')} \right] \right),$$

where  $\pi_\nu$ , for  $\nu \in \mathcal{T}$ , is a *mode permutation* operator which arranges the modes (axes) of its input such that they comply with an ascending order of  $\nu$ .<sup>8</sup>

Hierarchical tensor factorization can be viewed as a composition of multiple local tensor factorizations, one for each interior node in the mode tree. The local tensor factorization for  $\nu \in \text{int}(\mathcal{T})$  comprises  $R_\nu$  components, referred to as the local components at node  $\nu$  of the hierarchical tensor factorization — see Figure 6(b) for an illustration, and Definition 3 below.

**Definition 3.** For  $\nu \in \text{int}(\mathcal{T})$  and  $r \in [R_\nu]$ , the  $(\nu, r)$ 'th *local component* of the hierarchical tensor factorization is  $\mathbf{W}_{r,:}^{(\nu)} \otimes (\otimes_{\nu_c \in C(\nu)} \mathbf{W}_{:,r}^{(\nu_c)})$ . We use  $\text{LC}(\nu, r)$  to denote the set comprising  $\mathbf{W}_{r,:}^{(\nu)}$  and  $(\mathbf{W}_{:,r}^{(\nu_c)})_{\nu_c \in C(\nu)}$ , and  $\sigma_H^{(\nu, r)} := \|\otimes_{\mathbf{w} \in \text{LC}(\nu, r)} \mathbf{w}\|$  to denote the norm of the  $(\nu, r)$ 'th local component.

Mode trees of hierarchical tensor factorizations give rise to the notion of hierarchical tensor rank (*cf.* Grasedyck (2010)), which is based on matrix ranks of specific *matricizations* of a tensor (*cf.* Section 3.4 in Kolda (2006)).

<sup>8</sup>For  $\nu \in \text{int}(\mathcal{T})$ , denote its  $K := |C(\nu)|$  children by  $\nu_1, \dots, \nu_K$ , and the elements of  $\nu_k$  by  $j_1^k < \dots < j_{|\nu_k|}^k$  for  $k \in [K]$ . Let  $h : [|\nu|] \rightarrow [|\nu|]$  be the permutation sorting the tuple  $(j_1^1, \dots, j_{|\nu_1|}^1, \dots, j_1^K, \dots, j_{|\nu_K|}^K)$  in ascending order. Then, the mode permutation operator for  $\nu$  is defined by:  $\pi_\nu(\mathcal{W})_{d_1, \dots, d_{|\nu|}} = \mathcal{W}_{d_{h(1)}, \dots, d_{h(|\nu|)}}$ , where  $\mathcal{W}$  is an order  $|\nu|$  tensor.

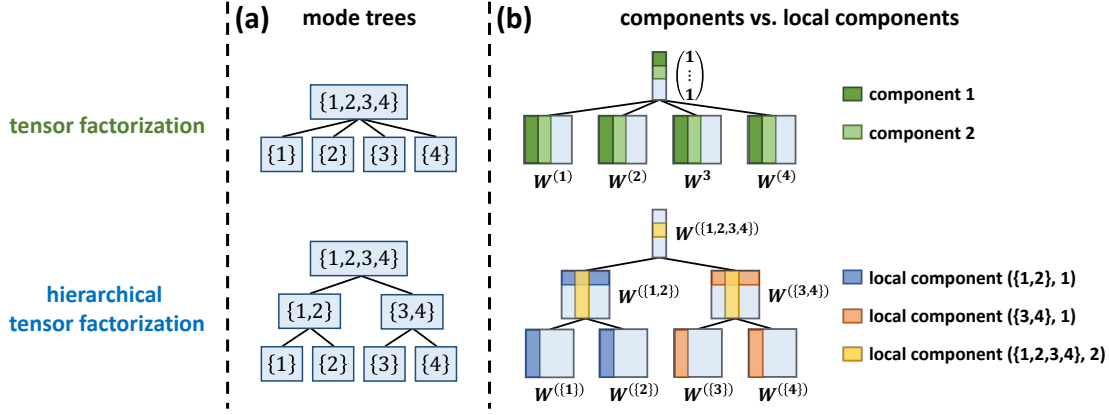


Figure 6: (a) Exemplar mode trees (Definition 2) for order  $N = 4$  hierarchical tensor factorization. Top corresponds to the degenerate case of tensor factorization, while bottom represents the most common choice (perfect binary tree). (b) Components of a tensor factorization (top) vs. local components (Definition 3) of a hierarchical tensor factorization (bottom). The  $r$ 'th component of a tensor factorization can be seen as the tensor product between a linear coefficient, which is set to 1, and the  $r$ 'th columns of  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(N)}$ . The local components of a hierarchical tensor factorization are the components of the local tensor factorizations forming it. For example, the  $r$ 'th local component at node  $\{1, 2\}$  in the hierarchical tensor factorization illustrated above is the tensor product between the  $r$ 'th row of  $\mathbf{W}^{(\{1,2\})}$  and the  $r$ 'th columns of its children's weight matrices  $\mathbf{W}^{(\{1\})}$  and  $\mathbf{W}^{(\{2\})}$ .

**Definition 4.** The *matricization* of  $\mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N}$  with respect to  $I \subset [N]$ , denoted  $\llbracket \mathcal{W}; I \rrbracket \in \mathbb{R}^{\prod_{i \in I} D_i, \prod_{j \in [N] \setminus I} D_j}$ , is its arrangement as a matrix where rows correspond to modes indexed by  $I$  and columns correspond to the remaining modes.<sup>9</sup>

**Definition 5.** The *hierarchical tensor rank* of  $\mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N}$  with respect to mode tree  $\mathcal{T}$  is the tuple comprising the matrix ranks of  $\mathcal{W}$ 's matricizations according to all nodes in  $\mathcal{T}$  except for the root, *i.e.*  $(\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$ . The order of entries in the tuple does not matter as long as it is consistent.

Unless stated otherwise, when referring to the hierarchical tensor rank of a hierarchical tensor factorization's end tensor, the rank is with respect to the mode tree of the factorization. Hierarchical tensor rank differs markedly from tensor rank. Specifically, even when the hierarchical tensor rank is low, *i.e.* the matrix ranks of matricizations according to all nodes in the mode tree are low, the tensor rank is typically extremely high (exponential in the order of the tensor — see Cohen et al. (2016a)).

Lemma 1 below states that the number of local components in a hierarchical tensor factorization controls the hierarchical tensor rank of its end tensor. More precisely,  $R_\nu$  — the number of local components at  $\nu \in \text{int}(\mathcal{T})$  — upper bounds the matrix rank of matricizations according to the children of  $\nu$ .

**Lemma 1** (adaptation of Theorem 7 in Cohen et al. (2018)). *For any interior node  $\nu \in \text{int}(\mathcal{T})$  and child  $\nu_c \in C(\nu)$ , it holds that  $\text{rank } \llbracket \mathcal{W}_H; \nu_c \rrbracket \leq R_\nu$ .*

*Proof.* Deferred to Appendix H.3. □

We may explicitly restrict the hierarchical tensor rank of end tensors  $\mathcal{W}_H$  (Equation (4)) by limiting  $(R_\nu)_{\nu \in \text{int}(\mathcal{T})}$ . However, since our interest lies in the implicit regularization of gradient descent, *i.e.* in the types of end tensors it will find without explicit constraints, we treat the case where  $(R_\nu)_{\nu \in \text{int}(\mathcal{T})}$  can be arbitrarily large.

Given a differentiable and locally smooth loss  $\mathcal{L}_H : \mathbb{R}^{D_1, \dots, D_N} \rightarrow \mathbb{R}_{\geq 0}$ , we consider parameterizing the solution as a hierarchical tensor factorization (Equation (4)), and optimizing the resulting (non-convex) objective:

$$\phi_H((\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}}) := \mathcal{L}_H(\mathcal{W}_H). \quad (5)$$

In line with analyses of implicit regularization in matrix and tensor factorizations (see Section 2), we model small learning rate for gradient descent via gradient flow:

$$\frac{d}{dt} \mathbf{W}^{(\nu)}(t) = -\frac{\partial}{\partial \mathbf{W}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T}}) \quad (6)$$

<sup>9</sup>Denoting the elements in  $I$  by  $i_1 < \dots < i_{|I|}$  and those in  $[N] \setminus I$  by  $j_1 < \dots < j_{N-|I|}$ , the matricization  $\llbracket \mathcal{W}; I \rrbracket$  holds the entries of  $\mathcal{W}$  such that  $\mathcal{W}_{d_1, \dots, d_N}$  is placed in row index  $1 + \sum_{l=1}^{|I|} (d_{i_l} - 1) \prod_{l'=1}^{l-1} D_{i_{l'}}$  and column index  $1 + \sum_{l=1}^{N-|I|} (d_{j_l} - 1) \prod_{l'=1}^{l-1} D_{j_{l'}}$ .

for all  $t \geq 0$  and  $\nu \in \mathcal{T}$ , where  $(\mathbf{W}^{(\nu)}(t))_{\nu \in \mathcal{T}}$  denote the weight matrices at time  $t$  of optimization.

Over matrix and tensor factorizations, gradient flow initialized near zero is known to minimize matrix and tensor ranks, respectively (Section 2). In particular, it leads to solutions that can be represented using few components. A natural question that arises is whether a similar phenomenon takes place in hierarchical tensor factorization: does gradient flow with small initialization learn solutions that can be represented with few local components at all locations of the mode tree? That is, does it learn solutions of low hierarchical tensor rank? In Section 4 we answer this question affirmatively.

## B Formal Analysis: Incremental Hierarchical Tensor Rank Learning

In this appendix we formally deliver our theoretical analysis of implicit regularization in hierarchical tensor factorization, which was informally overviewed in Section 4.

### B.1 Evolution of Local Component Norms

Lemma 2 below establishes an invariant of optimization: the differences between squared norms of weight vectors in the same local component are constant through time.

**Lemma 2.** *For all  $\nu \in \text{int}(\mathcal{T})$ ,  $r \in [R_\nu]$ , and  $\mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)$ :*

$$\|\mathbf{w}(t)\|^2 - \|\mathbf{w}'(t)\|^2 = \|\mathbf{w}(0)\|^2 - \|\mathbf{w}'(0)\|^2, \quad t \geq 0.$$

*Proof sketch (proof in Appendix H.4).* A straightforward derivation shows that  $\frac{d}{dt}\|\mathbf{w}(t)\|^2 = \frac{d}{dt}\|\mathbf{w}'(t)\|^2$  for all  $t \geq 0$ . Integrating both sides with respect to time completes the proof.  $\square$

The above invariant leads to the following definition of unbalancedness magnitude.

**Definition 6.** The *unbalancedness magnitude* of a hierarchical tensor factorization (Equation (4)) is:

$$\max_{\nu \in \text{int}(\mathcal{T}), r \in [R_\nu], \mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)} \left| \|\mathbf{w}\|^2 - \|\mathbf{w}'\|^2 \right|.$$

Lemma 2 implies that the unbalancedness magnitude remains constant throughout optimization. In the common regime of near-zero initialization, it will start off small, and stay small throughout. The closer initialization is to zero, the smaller the unbalancedness magnitude is. In accordance with analyses for matrix and tensor factorizations (see Section 2), we treat the case of unbalancedness magnitude zero as an idealization of standard near-zero initializations. Theorem 1 analyzes this case, characterizing the dynamics for norms of local components.

**Theorem 1.** *Assume unbalancedness magnitude zero at initialization. Let  $\mathcal{W}_H(t)$  denote the end tensor (Equation (4)) and  $(\sigma_H^{(\nu, r)}(t))_{\nu \in \text{int}(\mathcal{T}), r \in [R_\nu]}$  denote the norms of local components (Definition 3) at time  $t \geq 0$  of optimization. Then, for any  $\nu \in \text{int}(\mathcal{T})$  and  $r \in [R_\nu]$ :*

$$\frac{d}{dt}\sigma_H^{(\nu, r)}(t) = \sigma_H^{(\nu, r)}(t)^{2 - \frac{2}{L_\nu}} L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, r)}(t) \rangle, \quad (7)$$

where  $L_\nu := |C(\nu)| + 1$  is the number of weight vectors in a local component at node  $\nu$ , and  $\mathcal{C}_H^{(\nu, r)}(t) \in \mathbb{R}^{D_1, \dots, D_N}$  is the end tensor obtained by normalizing the  $r$ 'th local component at node  $\nu$  and setting all other local components at node  $\nu$  to zero, i.e. by replacing in Equation (4)  $\mathcal{W}^{(\nu, r')}$  with  $\pi_\nu((\sigma_H^{(\nu, r)})^{-1} \mathbf{W}_{r, r'}^{(\nu)} [\otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')}] )$  for all  $r' \in [R_{Pa(\nu)}]$ . By convention,  $\mathcal{C}_H^{(\nu, r)}(t) = 0$  if  $\sigma_H^{(\nu, r)}(t) = 0$ .

*Proof sketch (proof in Appendix H.5).* If  $\sigma_H^{(\nu, r)}(t)$  is zero at some  $t \geq 0$ , then we show that it must be identically zero through time, leading both sides of Equation (7) to be equal to zero. Otherwise, differentiating the local component's norm with respect to time, we obtain:

$$\frac{d}{dt}\sigma_H^{(\nu, r)}(t) = \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, r)}(t) \rangle \cdot \sum_{\mathbf{w} \in \text{LC}(\nu, r)} \prod_{\mathbf{w}' \in \text{LC}(\nu, r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2.$$

Since the unbalancedness magnitude is zero at initialization, Lemma 2 implies that  $\|\mathbf{w}(t)\|^2 = \|\mathbf{w}'(t)\|^2 = \sigma_H^{(\nu, r)}(t)^{2/L_\nu}$  for all  $\mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)$ , which together with the expression above for  $\frac{d}{dt}\sigma_H^{(\nu, r)}(t)$  establishes Equation (7).  $\square$

As can be seen from Equation (7), the evolution of local component norms in a hierarchical tensor factorization is structurally identical to the evolution of singular values in matrix factorization (Equation (1)) and component norms in tensor factorization (Equation (2)). Specifically, it is dictated by two factors: a projection term,  $\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, r)}(t) \rangle$ , and a self-dependence term,  $\sigma_H^{(\nu, r)}(t)^{2-\frac{2}{L_\nu}} L_\nu$ . Analogous to a singular component  $\mathcal{C}_M^{(r)}(t)$  in matrix factorization and a normalized component  $\mathcal{C}_T^{(r)}(t)$  in tensor factorization,  $\mathcal{C}_H^{(\nu, r)}(t)$  is the direction that the  $(\nu, r)$ 'th local component imposes on  $\mathcal{W}_H(t)$ .<sup>10</sup> The projection of  $\mathcal{C}_H^{(\nu, r)}(t)$  onto  $-\nabla \mathcal{L}_H(\mathcal{W}_H(t))$  therefore promotes growth of local components that align  $\mathcal{W}_H(t)$  with  $-\nabla \mathcal{L}_H(\mathcal{W}_H(t))$ , the direction of steepest descent. More critical is the self-dependence term,  $\sigma_H^{(\nu, r)}(t)^{2-\frac{2}{L_\nu}} L_\nu$ , which induces a momentum-like effect that attenuates the movement of small local components and accelerates the movement of large ones. It suggests that, in analogy with matrix and tensor factorizations, local components tend to be learned incrementally, yielding a bias towards low hierarchical tensor rank. This prospect is affirmed empirically in Figure 1 (right) as well as Appendix G, and is supported theoretically in Appendix B.2.

**Evolution of local component norms under arbitrary initialization** Theorem 1 can be extended to account for arbitrary initialization, *i.e.* for initialization with unbalancedness magnitude different from zero. For conciseness we defer this extension to Appendix E, while noting that if initialization has small unbalancedness magnitude — as is the case with any near-zero initialization — then local component norms approximately evolve per Equation (7), *i.e.* the result of Theorem 1 approximately holds.

## B.2 Implicit Hierarchical Tensor Rank Minimization

As discussed in Section 2, under certain technical conditions, the incremental matrix and tensor rank learning phenomena, induced by the implicit regularization in matrix and tensor factorizations, can be used to prove exact matrix and tensor rank minimization, respectively. Below we consider similar technical conditions, and provide theoretical results aimed at facilitating an analogous proof for hierarchical tensor factorization, *i.e.* a proof that its implicit regularization leads to exact hierarchical tensor rank minimization. We begin by illustrating how, under said conditions, the incremental hierarchical tensor rank learning phenomenon established in Theorem 1 leads to solutions with many small local components (Appendix B.2.1). We then show that this implies proximity to low hierarchical tensor rank (Appendix B.2.2). Throughout the above, the main step missing in order to derive a complete proof of exact hierarchical tensor rank minimization, is confirmation that a certain alignment inequality (Equation (8)) holds throughout optimization. We regard this as an important direction for future work.

### B.2.1 ILLUSTRATIVE DEMONSTRATION OF SMALL LOCAL COMPONENTS

Below we qualitatively demonstrate how the dynamical characterization derived in Appendix B.1 implies that the implicit regularization in hierarchical tensor factorization can lead to solutions with small local components. Under the setting and notation of Theorem 1, consider an initialization  $(\mathbf{U}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}})_{\nu \in \mathcal{T}}$  for the weight matrices of the hierarchical tensor factorization, scaled by  $\alpha \in \mathbb{R}_{>0}$ . That is,  $\mathbf{W}^{(\nu)}(0) = \alpha \cdot \mathbf{U}^{(\nu)}$  for all  $\nu \in \mathcal{T}$ . Focusing on some interior node  $\nu \in \text{int}(\mathcal{T})$ , let  $r, \bar{r} \in [R_\nu]$ , and assume for simplicity that  $\nu$  is not degenerate, in the sense that it has more than one child. Suppose also that at initialization the norm of the  $(\nu, r)$ 'th local component is greater than the norm of the  $(\nu, \bar{r})$ 'th local component, *i.e.*  $\sigma_H^{(\nu, r)}(0) > \sigma_H^{(\nu, \bar{r})}(0)$ , and that  $\mathcal{C}_H^{(\nu, r)}(t)$  is at least as aligned as  $\mathcal{C}_H^{(\nu, \bar{r})}(t)$  with the direction of steepest descent up to a time  $T > 0$ , *i.e.* for all  $t \in [0, T]$ :

$$\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, r)}(t) \rangle \geq \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, \bar{r})}(t) \rangle. \quad (8)$$

Then, by Theorem 1 for all  $t \in [0, T]$ :<sup>11</sup>

$$\sigma_H^{(\nu, \bar{r})}(t)^{-2+\frac{2}{L_\nu}} \frac{d}{dt} \sigma_H^{(\nu, \bar{r})}(t) \leq \sigma_H^{(\nu, r)}(t)^{-2+\frac{2}{L_\nu}} \frac{d}{dt} \sigma_H^{(\nu, r)}(t).$$

Integrating both sides with respect to time, we may upper bound  $\sigma_H^{(\nu, \bar{r})}(t)$  with a function of  $\sigma_H^{(\nu, r)}(t)$ :

$$\sigma_H^{(\nu, \bar{r})}(t) \leq \left[ \sigma_H^{(\nu, r)}(t)^{-\frac{L_\nu-2}{L_\nu}} + \alpha^{-(L_\nu-2)} \cdot \text{const} \right]^{-\frac{L_\nu}{L_\nu-2}}, \quad (9)$$

<sup>10</sup>Indeed, just as in matrix factorization  $\mathbf{W}_M = \sum_r \sigma_M^{(r)} \cdot \mathcal{C}_M^{(r)}$ , and in tensor factorization  $\mathcal{W}_T = \sum_r \sigma_T^{(r)} \cdot \mathcal{C}_T^{(r)}$ , the end tensor of a hierarchical tensor factorization decomposes as  $\mathcal{W}_H = \sum_{r=1}^{R_\nu} \sigma_H^{(\nu, r)} \cdot \mathcal{C}_H^{(\nu, r)}$  (implied by Lemmas 5 and 12 in Appendix H).

<sup>11</sup>A local component cannot reach the origin unless it was initialized there (implied by Lemma 14 in Appendix H). Accordingly, we disregard the trivial case where  $\sigma_H^{(\nu, \bar{r})}(t) = 0$  for some  $t \in [0, T]$ .



where  $const$  stands for a positive value that does not depend on  $t$  and  $\alpha$ . Equation (9) reveals a gap between  $\sigma_H^{(\nu,r)}(t)$  and  $\sigma_H^{(\nu,\bar{r})}(t)$  that is more significant the smaller the initialization scale  $\alpha$  is. In particular, regardless of how large  $\sigma_H^{(\nu,r)}(t)$  is,  $\sigma_H^{(\nu,\bar{r})}(t)$  is upper bounded by a value that approaches zero as  $\alpha \rightarrow 0$ . Hence, initializing near zero produces solutions with small local components.

### B.2.2 SMALL LOCAL COMPONENTS IMPLY PROXIMITY TO LOW HIERARCHICAL TENSOR RANK

The following proposition establishes that small local components in a hierarchical tensor factorization imply that its end tensor can be well approximated with low hierarchical tensor rank.

**Proposition 1.** *Consider an assignment for the weight matrices  $(\mathbf{W}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}})_{\nu \in \mathcal{T}}$  of a hierarchical tensor factorization, and let  $B := \max_{\nu \in \mathcal{T}} \|\mathbf{W}^{(\nu)}\|$ . Assume without loss of generality that at each  $\nu \in \text{int}(\mathcal{T})$ , local components are ordered by their norms, i.e.  $\sigma_H^{(\nu,1)} \geq \dots \geq \sigma_H^{(\nu,R_\nu)}$ . Then, for any  $\epsilon \geq 0$  and  $(R'_\nu \in \mathbb{N})_{\nu \in \text{int}(\mathcal{T})}$ , if  $\sum_{r=R'_\nu+1}^{R_\nu} \sigma_H^{(\nu,r)} \leq \epsilon \cdot (|\mathcal{T}| - N)^{-1} B^{|C(\nu)|+1-|\mathcal{T}|}$  for all  $\nu \in \text{int}(\mathcal{T})$ , it holds that:*

$$\inf_{\substack{\mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N} \text{ s.t.} \\ \forall \nu \in \mathcal{T} \setminus \{[N]\}: \text{rank}[\mathcal{W}; \nu] \leq R'_{Pa(\nu)}}} \|\mathcal{W}_H - \mathcal{W}\| \leq \epsilon,$$

i.e.  $\mathcal{W}_H$  is within  $\epsilon$ -distance from the set of tensors whose hierarchical tensor rank is no greater (element-wise) than  $(R'_{Pa(\nu)})_{\nu \in \mathcal{T} \setminus \{[N]\}}$ .

*Proof sketch (proof in Appendix H.6).* Let  $\overline{\mathcal{W}}_{HT}^S$  be the end tensor obtained after pruning all local components indexed by  $\mathcal{S} := \{(\nu, r) : \nu \in \text{int}(\mathcal{T}), r \in \{R'_\nu + 1, \dots, R_\nu\}\}$ , i.e. after setting to zero the  $r$ 'th row of  $\mathbf{W}^{(\nu)}$  and the  $r$ 'th column of  $\mathbf{W}^{(\nu_c)}$  for all  $(\nu, r) \in \mathcal{S}$  and  $\nu_c \in C(\nu)$ . The desired result follows by showing that  $\text{rank}[\overline{\mathcal{W}}_{HT}^S; \nu] \leq R'_{Pa(\nu)}$  for all  $\nu \in \mathcal{T} \setminus \{[N]\}$ , and upper bounding  $\|\mathcal{W}_H - \overline{\mathcal{W}}_{HT}^S\|$  by  $\epsilon$ .  $\square$

### B.3 Partially Ordered Complexity Measure

Existing attempts to explain implicit regularization in deep learning typically argue for reduction of some complexity measure that is *totally ordered* (meaning that within any two values for this measure, there must be one smaller than or equal to the other), for example a norm (Gunasekar et al., 2017; Soudry et al., 2018; Li et al., 2018; Woodworth et al., 2020; Lyu et al., 2021). Recent evidence suggests that obtaining a complete explanation through such complexity measures may not be possible (Razin & Cohen, 2020; Vardi & Shamir, 2021). Hierarchical tensor rank (which we have shown to be implicitly reduced by a class of deep non-linear convolutional networks) represents a new type of complexity measure, in the sense that it is *partially ordered*. Specifically, while it entails a standard (product) partial order —  $(r_1, \dots, r_K) \leq (r'_1, \dots, r'_K)$  if and only if  $r_i \leq r'_i$  for all  $i \in [K]$  — it does not admit a natural total order. Indeed, Proposition 2 below shows that there exist simple learning problems in which, among the data-fitting solutions, there are multiple minimal hierarchical tensor ranks, none smaller than or equal to the other. We believe the notion of a partially ordered complexity measure may pave the way to furthering our understanding of implicit regularization in deep learning.

**Proposition 2.** *For every order  $N \in \mathbb{N}_{\geq 3}$  and mode dimensions  $D_1, \dots, D_N \in \mathbb{N}_{\geq 2}$ , there exists a tensor completion problem (i.e. a loss  $\mathcal{L}(\mathcal{W}) = \frac{1}{|\Omega|} \sum_{(d_1, \dots, d_N) \in \Omega} (\mathcal{W}_{d_1, \dots, d_N} - \mathcal{W}_{d_1, \dots, d_N}^*)^2$  with ground truth  $\mathcal{W}^* \in \mathbb{R}^{D_1, \dots, D_N}$  and set of observed entries  $\Omega \subset [D_1] \times \dots \times [D_N]$ ) in which, for every mode tree  $\mathcal{T}$  over  $[N]$  (Definition 2), the set of hierarchical tensor ranks for tensors fitting the observations includes multiple minimal elements (under the standard product partial order), none smaller than or equal to the other. That is, the set  $\mathcal{R}_{\mathcal{T}} := \{(\text{rank}[\mathcal{W}; \nu])_{\nu \in \mathcal{T} \setminus \{[N]\}} : \mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N}, \mathcal{L}(\mathcal{W}) = 0\}$  includes elements  $(R_\nu)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  and  $(R'_\nu)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  for which the following hold: (i) there exists no  $(R''_\nu)_{\nu \in \mathcal{T} \setminus \{[N]\}} \in \mathcal{R}_{\mathcal{T}} \setminus \{(R_\nu)_{\nu \in \mathcal{T} \setminus \{[N]\}}, (R'_\nu)_{\nu \in \mathcal{T} \setminus \{[N]\}}\}$  satisfying  $(R''_\nu)_\nu \leq (R_\nu)_\nu$  or  $(R''_\nu)_\nu \leq (R'_\nu)_\nu$ ; and (ii) neither  $(R_\nu)_\nu \leq (R'_\nu)_\nu$  nor  $(R'_\nu)_\nu \leq (R_\nu)_\nu$ .*

*Proof sketch (proof in Appendix H.7).* We construct a tensor completion problem and two solutions  $\mathcal{W}$  and  $\mathcal{W}'$  (tensors fitting observed entries) such that, for every mode tree  $\mathcal{T}$ , the hierarchical tensor ranks with respect to  $\mathcal{T}$  of  $\mathcal{W}$  and  $\mathcal{W}'$  are two different minimal elements of  $\mathcal{R}_{\mathcal{T}}$ .  $\square$

## C Related Work

A large and growing body of literature has theoretically investigated the implicit regularization brought forth by gradient-based optimization. Works along this line have treated various models, including: linear predictors (Soudry et al., 2018;

Gunasekar et al., 2018; Nacson et al., 2019b; Ji & Telgarsky, 2019b; Shachaf et al., 2021); polynomially parameterized linear models with a single output (Ji & Telgarsky, 2019a; Woodworth et al., 2020; Moroshko et al., 2020; Azulay et al., 2021; HaoChen et al., 2021; Pesme et al., 2021; Li et al., 2021b; Chou et al., 2021); shallow non-linear neural networks (Hu et al., 2020; Vardi & Shamir, 2021; Sarussi et al., 2021; Mulayoff et al., 2021; Lyu et al., 2021); homogeneous networks (Lyu & Li, 2020; Vardi et al., 2021); and ultra-wide networks (Oymak & Soltanolkotabi, 2019; Chizat & Bach, 2020). Arguably the most widely analyzed model is matrix factorization, whose study was extended to tensor factorization (Gunasekar et al., 2017; Du et al., 2018; Li et al., 2018; Arora et al., 2019; Gidel et al., 2019; Mulayoff & Michaeli, 2020; Blanc et al., 2020; Gissin et al., 2020; Razin & Cohen, 2020; Chou et al., 2020; Eftekhari & Zygalkis, 2021; Yun et al., 2021; Min et al., 2021; Li et al., 2021a; Razin et al., 2021; Milanese et al., 2021; Ge et al., 2021). Our work generalizes existing results for matrix and tensor factorizations (see Section 2) to hierarchical tensor factorization — a considerably richer and more complex model.

Hierarchical tensor factorization was originally introduced in Hackbusch & Kühn (2009). By virtue of its equivalence to different types of (non-linear) neural networks, it has been paramount to the study of expressiveness in deep learning (Cohen et al., 2016b; Sharir et al., 2016; Cohen & Shashua, 2016; 2017; Cohen et al., 2017; Sharir & Shashua, 2018; Cohen et al., 2018; Levine et al., 2018a;b; Balda et al., 2018; Khruikov et al., 2018; 2019; Levine et al., 2019). It is also used in different contexts, for example recovery of low (hierarchical tensor) rank tensors (Da Silva & Herrmann, 2015; Steinlechner, 2016; Rauhut et al., 2017; Kargas & Sidiropoulos, 2020; 2021). To the best of our knowledge, this paper is the first to study implicit regularization of gradient-based optimization over hierarchical tensor factorization.

With regards to convolutional networks, theoretical investigations of their implicit regularization are scarce. Existing works in this category treat linear (Gunasekar et al., 2018; Jagadeesan et al., 2021; Kohn et al., 2021) and homogeneous (Nacson et al., 2019a; Lyu & Li, 2020; Ji & Telgarsky, 2020) models.<sup>12</sup> None of these works have pointed out an implicit regularization towards local dependencies, as our theory does (Sections 4 and 5). Although the locality of convolutional networks is widely accepted, it is typically ascribed to expressive properties determined by their architecture (see, *e.g.*, Cohen & Shashua (2017); Linsley et al. (2018); Kim et al. (2020)). Our work is the first to indicate that it also originates from implicit regularization. As we demonstrate in Section 6, this observation can have far reaching implications to the performance of convolutional networks in practice.

## D Hierarchical Tensor Factorization as Deep Non-Linear Convolutional Network

In this appendix, we formally state and prove a known correspondence between hierarchical tensor factorization and certain deep non-linear convolutional networks (*cf.* Cohen et al. (2016b)). For conciseness, we assume the tensor order  $N$  is a power of  $P \in \mathbb{N}_{\geq 2}$  and the mode dimensions  $D_1, \dots, D_N$  are equal, and focus on the factorization induced by a perfect  $P$ -ary mode tree (Definition 2) that combines nodes with adjacent indices.

Let  $L := \log_P N$  denote the height of the mode tree, and associate each of its nodes with a respective location  $(l, n)$ , where  $l \in [L+1]$  is the level in the tree (numbered from leaves to root in ascending order), and  $n \in [N/P^{l-1}]$  is the index inside the level (see Figure 7 for an illustration). Adapting Equation (4) to the current setting, the end tensor is computed as follows:

for all  $n \in [N]$  and  $r \in [R_1]$ :

$$\underbrace{\mathcal{W}^{(1,n,r)}}_{\text{order } 1} := \mathbf{W}_{:,r}^{(1,n)},$$

for all  $l \in \{2, \dots, L\}$ ,  $n \in [N/P^{l-1}]$ , and  $r \in [R_l]$  (traverse interior nodes of  $\mathcal{T}$  from leaves to root, non-inclusive):

$$\underbrace{\mathcal{W}^{(l,n,r)}}_{\text{order } P^{l-1}} := \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \left[ \otimes_{p=(n-1) \cdot P+1}^{n \cdot P} \mathcal{W}^{(l-1,p,r')} \right],$$

$$\underbrace{\mathcal{W}_H}_{\text{order } N} := \sum_{r'=1}^{R_L} \mathbf{W}_{r',1}^{(L+1,1)} \left[ \otimes_{p=1}^P \mathcal{W}^{(L,p,r')} \right],$$

(10)




Figure 7: Perfect  $P$ -ary mode tree that combines adjacent indices, for order  $N = 4$  and  $P = 2$ .

<sup>12</sup>There have also been works studying implicit effects of explicit regularizers for convolutional networks (Ergen & Pilanci, 2021), but these are outside the scope of our paper.

where  $(\mathbf{W}^{(l,n)} \in \mathbb{R}^{R_{l-1}, R_l})_{l \in [L+1], n \in [N/P^{l-1}]}$  are the factorization’s weight matrices,  $R_{L+1} = 1$ , and  $R_0 := D_1 = \dots = D_N$ .

The deep non-linear convolutional network corresponding to the above factorization (illustrated in Figure 2 (right)) has  $L$  hidden layers, the  $l$ ’th one comprising a locally connected linear operator with  $R_l$  channels followed by channel-wise product pooling with window size  $P$  (multiplicative non-linearity). Denoting by  $(\mathbf{h}^{(l-1,1)}, \dots, \mathbf{h}^{(l-1, N/P^{l-1})}) \in \mathbb{R}^{R_{l-1}} \times \dots \times \mathbb{R}^{R_{l-1}}$  the output of the  $l-1$ ’th hidden layer, where  $(\mathbf{h}^{(0,1)}, \dots, \mathbf{h}^{(0,N)}) := (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$  is the network’s input, the locally connected operator of the  $l$ ’th layer computes  $(\mathbf{W}^{(l,n)})^\top \mathbf{h}^{(l-1,n)}$  for each index  $n \in [N/P^{l-1}]$ . We refer to this operator as “ $1 \times 1$  conv” in appeal to the case of weight sharing, where  $\mathbf{W}^{(l,1)} = \dots = \mathbf{W}^{(l, N/P^{l-1})}$ . Following the locally connected operator, for each  $n \in [N/P^l]$  and  $r \in [R_l]$ , the pooling operator computes  $\prod_{p=(n-1) \cdot P+1}^{n \cdot P} [(\mathbf{W}^{(l,p)})^\top \mathbf{h}^{(l-1,p)}]_r$ , thereby producing  $(\mathbf{h}^{(l,1)}, \dots, \mathbf{h}^{(l, N/P^l)})$ . After passing the input through all hidden layers, a final linear layer, whose weights are  $\mathbf{W}^{(L+1,1)}$ , yields the scalar output of the network  $(\mathbf{W}^{(L+1,1)})^\top \mathbf{h}^{(L,1)}$ . Notice that the weight matrices of the hierarchical tensor factorization are exactly the learnable weights of the network, and  $R_{l-1}$  — the number of local components (Definition 3) at nodes in level  $l$  of the factorization — is the width of the network’s  $l-1$ ’th hidden layer.

The above formulation of the network supports not only sequential inputs (*e.g.* audio and text), but also inputs arranged as multi-dimensional arrays (*e.g.* two-dimensional images). The choice of how to assign the indices  $1, \dots, N$  to input elements determines the geometry of pooling windows throughout the network (Cohen & Shashua, 2017).

Proposition 3 below implies that we may view solution of a prediction task using the deep convolutional network described above as a hierarchical tensor factorization problem, and vice versa. For example, solving tensor completion and certain sensing problems using hierarchical tensor factorization amounts to applying the corresponding network to a regression task.

**Proposition 3** (adapted from Cohen et al. (2016b)). *Let  $f_\Theta : \times_{n=1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  be the function realized by the deep non-linear convolutional network described above, where  $\Theta := (\mathbf{W}^{(l,n)})_{l \in [L+1], n \in [N/P^{l-1}]}$  stands for the network’s weights. Denote by  $\mathcal{W}_H$  the end tensor of the hierarchical tensor factorization specified in Equation (10). Then, for all  $\mathbf{x}^{(1)} \in \mathbb{R}^{D_1}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^{D_N}$ :*

$$f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \rangle.$$

*Proof sketch (proof in Appendix H.8).* By induction over the layers of the network, we show that the output of the  $l$ ’th convolutional layer (linear output layer for  $l = L+1$ ) at index  $n$  and channel  $r$  is  $\langle \otimes_{p=(n-1) \cdot P+1}^{n \cdot P} \mathbf{x}^{(p)}, \mathcal{W}^{(l,n,r)} \rangle$ , where  $\mathcal{W}^{(L+1,1,1)} := \mathcal{W}_H$ , and all other  $\mathcal{W}^{(l,n,r)}$  are the intermediate tensors formed when computing  $\mathcal{W}_H$  according to Equation (10). Since  $f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$  is the output of the  $L+1$ ’th layer at index 1 and channel 1, applying the inductive claim for  $l = L+1, n = 1$ , and  $r = 1$  concludes the proof.  $\square$

We conclude this appendix by noting that in the special case where  $P = N$ , if the weight matrix of the root node holds ones, the hierarchical tensor factorization reduces to a tensor factorization, and the corresponding convolutional network has a single hidden layer (with global product pooling) followed by a final summation layer. We thus obtain the equivalence between tensor factorization and a shallow non-linear convolutional network as a corollary of Proposition 3.

## E Evolution of Local Component Norms Under Arbitrary Initialization

Theorem 1 in Appendix B.1 characterizes the evolution of local component norms in a hierarchical tensor factorization, under the assumption of unbalancedness magnitude zero at initialization. Theorem 2 below extends the characterization to account for arbitrary initialization. It establishes that if the unbalancedness magnitude at initialization is small — as is the case under any near-zero initialization — local component norms approximately evolve per Theorem 1.

**Theorem 2.** *With the context and notations of Theorem 1, assume unbalancedness magnitude  $\epsilon \geq 0$  at initialization. Then, for any  $\nu \in \text{int}(\mathcal{T})$ ,  $r \in [R_\nu]$ , and time  $t \geq 0$  at which  $\sigma_H^{(\nu,r)}(t) > 0$ :<sup>13</sup>*

<sup>13</sup>Since norms are not differentiable at the origin, when  $\sigma_H^{(\nu,r)}(t)$  is equal to zero it may not be differentiable with respect to time.

- If  $\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle \geq 0$ , then:

$$\begin{aligned} \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &\leq \left( \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon \right)^{L_\nu-1} \cdot L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle, \\ \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &\geq \frac{\sigma_H^{(\nu,r)}(t)^2}{\sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon} \cdot L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle; \end{aligned} \quad (11)$$

- otherwise, if  $\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle < 0$ , then:

$$\begin{aligned} \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &\geq \left( \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon \right)^{L_\nu-1} \cdot L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle, \\ \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &\leq \frac{\sigma_H^{(\nu,r)}(t)^2}{\sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon} \cdot L_\nu \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle. \end{aligned} \quad (12)$$

*Proof sketch (proof in Appendix H.9).* The proof follows a line similar to that of Theorem 1, except that here conservation of unbalancedness magnitude leads to  $\|\mathbf{w}(t)\|^2 \leq \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon$  for all  $\mathbf{w} \in \text{LC}(\nu, r)$ . Applying this inequality to  $\frac{d}{dt} \sigma_H^{(\nu,r)}(t) = \langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \rangle \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \prod_{\mathbf{w}' \in \text{LC}(\nu,r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2$  yields Equations (11) and (12).  $\square$

## F Hierarchical Tensor Rank as Measure of Long-Range Dependencies

Section 5 discusses the known fact by which the hierarchical tensor rank (Definition 5) of a hierarchical tensor factorization measures the strength of long-range dependencies modeled by the equivalent convolutional network (see Cohen & Shashua (2017); Levine et al. (2018a;b)). For the convenience of the reader, the current appendix formally explains this fact.

Consider a hierarchical tensor factorization with mode tree  $\mathcal{T}$  (Definition 2), weight matrices  $\Theta := (\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}}$ , and an equivalent convolutional network realizing a parametric input-output function  $f_\Theta$ . As claimed in Section 3 (and formally justified in Appendix D), the function realized by the convolutional network takes the form  $f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \rangle$ , where  $\mathcal{W}_H$  stands for the end tensor of the factorization (Equation (4)). Proposition 4 below establishes that for any subset of indices  $I \subset [N]$ , the matrix rank of  $\mathcal{W}_H$ 's matricization according to  $I$  is equal to the separation rank (Definition 1) of  $f_\Theta$  with respect to  $I$ , i.e.  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket = \text{sep}(f_\Theta; I)$ . In particular, the hierarchical tensor rank of  $\mathcal{W}_H$  with respect to  $\mathcal{T}$  —  $(\text{rank} \llbracket \mathcal{W}_H; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  — amounts to  $(\text{sep}(f_\Theta; \nu))_{\nu \in \mathcal{T} \setminus \{[N]\}}$ . In the canonical case where nodes in  $\mathcal{T}$  hold adjacent indices, the separation ranks of  $f_\Theta$  with respect to them measure the dependencies modeled between distinct areas of the input, i.e. the non-local (long-range) dependencies.

**Proposition 4** (adaptation of Claim 1 in Cohen & Shashua (2017)). *Consider a hierarchical tensor factorization with mode tree  $\mathcal{T}$  (Definition 2) and weight matrices  $\Theta := (\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}}$ , and denote its end tensor by  $\mathcal{W}_H$  (Equation (4)). Let  $f_\Theta : \times_{n=1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  be defined by  $f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) := \langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \rangle$ . Then, for all  $I \subset [N]$ :*

$$\text{rank} \llbracket \mathcal{W}_H; I \rrbracket = \text{sep}(f_\Theta; I).$$

*Proof sketch (proof in Appendix H.10).* To prove that  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \geq \text{sep}(f_\Theta; I)$ , we derive a representation of  $f_\Theta$  as a sum of  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket$  terms, each being a product between a function that operates over  $(\mathbf{x}^{(i)})_{i \in I}$  and another that operates over the remaining input variables. For the converse,  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \leq \text{sep}(f_\Theta; I)$ , we prove that for any grid tensor  $\mathcal{W}$  of a function  $f$ , i.e. tensor holding the outputs of  $f$  over a grid of inputs, it holds that  $\text{rank} \llbracket \mathcal{W}; I \rrbracket \leq \text{sep}(f; I)$ . We conclude by showing that  $\mathcal{W}_H$  is a grid tensor of  $f_\Theta$ .  $\square$

## G Further Experiments and Implementation Details

### G.1 Further Experiments

Figures 8, 9, and 10 supplement Figure 1 by including, respectively: (i) plots of additional local component norms and singular values during optimization in the experiment presented by Figure 1 (right); (ii) experiments with tensor sensing loss; and (iii) experiments with different hierarchical tensor factorization orders and mode trees, as well as different ground truth hierarchical tensor ranks. Figure 11 portrays an experiment identical to that of Figure 5, but with ResNet34 in place of



ResNet18. Figures 12 and 13 extend Figures 5 and 11, respectively, by presenting results obtained with baseline networks that are already regularized using standard techniques (weight decay and dropout).

## G.2 Implementation Details

In this subappendix we provide implementation details omitted from our experimental reports (Figure 1, Section 6, and Appendix G.1). Source code for reproducing our results and figures, based on the PyTorch framework (Paszke et al., 2017), can be found at [https://github.com/asafmaman101/imp\\_reg\\_htf](https://github.com/asafmaman101/imp_reg_htf). All experiments were run on a single Nvidia RTX 2080 Ti GPU.

### G.2.1 INCREMENTAL HIERARCHICAL TENSOR RANK LEARNING (FIGURES 1, 8, 9, AND 10)

**Figure 1 (left):** the minimized matrix completion loss was  $\mathcal{L}_M(\mathbf{W}_M) = \frac{1}{|\Omega|} \sum_{(i,j) \in \Omega} ((\mathbf{W}_M)_{i,j} - \mathbf{W}_{i,j}^*)^2$ , where  $\Omega$  denotes a set of 2048 observed entries chosen uniformly at random (without repetition) from a matrix rank 5 ground truth  $\mathbf{W}^* \in \mathbb{R}^{64,64}$ . We generated  $\mathbf{W}^*$  by computing  $\mathbf{W}^{*(1)}\mathbf{W}^{*(2)}$ , with each entry of  $\mathbf{W}^{*(1)} \in \mathbb{R}^{64,5}$  and  $\mathbf{W}^{*(2)} \in \mathbb{R}^{5,64}$  drawn independently from the standard normal distribution, and subsequently normalizing the result to be of Frobenius norm 64 (square root of its number of entries). Reconstruction error with respect to  $\mathbf{W}^*$  is based on normalized Frobenius distance, *i.e.* for a solution  $\mathbf{W}_M$  it is  $\|\mathbf{W}_M - \mathbf{W}^*\| / \|\mathbf{W}^*\|$ . The matrix factorization applied to the task was of depth 3 and had hidden dimensions 64 between its layers so that its rank was unconstrained. Standard deviation for initialization was set to 0.001.

**Figure 1 (middle):** the minimized tensor completion loss was  $\mathcal{L}_T(\mathcal{W}_T) = \frac{1}{|\Omega|} \sum_{(d_1,d_2,d_3) \in \Omega} ((\mathcal{W}_T)_{d_1,d_2,d_3} - \mathcal{W}_{d_1,d_2,d_3}^*)^2$ , where  $\Omega$  denotes a set of 2048 observed entries chosen uniformly at random (without repetition) from a tensor rank 5 ground truth  $\mathcal{W}^* \in \mathbb{R}^{16,16,16}$ . We generated  $\mathcal{W}^*$  by computing  $\sum_{r=1}^5 \mathbf{W}_{:,r}^{*(1)} \otimes \mathbf{W}_{:,r}^{*(2)} \otimes \mathbf{W}_{:,r}^{*(3)}$ , with each entry of  $\mathbf{W}^{*(1)}$ ,  $\mathbf{W}^{*(2)}$ , and  $\mathbf{W}^{*(3)} \in \mathbb{R}^{16,5}$  drawn independently from the standard normal distribution, and subsequently normalizing the result to be of Frobenius norm 64 (square root of its number of entries). Reconstruction error with respect to  $\mathcal{W}^*$  is based on normalized Frobenius distance, *i.e.* for a solution  $\mathcal{W}_T$  it is  $\|\mathcal{W}_T - \mathcal{W}^*\| / \|\mathcal{W}^*\|$ . The tensor factorization applied to the task had  $R = 256$  components so that its tensor rank was unconstrained.<sup>14</sup> Standard deviation for initialization was set to 0.001.

**Figure 1 (right):** the minimized tensor completion loss was  $\mathcal{L}_H(\mathcal{W}_H) = \frac{1}{|\Omega|} \sum_{(d_1,\dots,d_4) \in \Omega} ((\mathcal{W}_H)_{d_1,\dots,d_4} - \mathcal{W}_{d_1,\dots,d_4}^*)^2$ , where  $\Omega$  denotes a set of 2048 observed entries chosen uniformly at random (without repetition) from a hierarchical tensor rank  $(5, 5, 5, 5, 5, 5)$  ground truth  $\mathcal{W}^* \in \mathbb{R}^{8,8,8,8,8}$ . We generated  $\mathcal{W}^*$  according to Equation (4) using a perfect binary mode tree  $\mathcal{T}$  over  $[4]$  and weight matrices  $(\mathbf{W}^{*(\nu)})_{\nu \in \mathcal{T}}$ , where  $\mathbf{W}^{*(\nu)} \in \mathbb{R}^{8,5}$  for  $\nu \in \{\{1\}, \dots, \{4\}\}$ ,  $\mathbf{W}^{*(\nu)} \in \mathbb{R}^{5,5}$  for  $\nu \in \text{int}(\mathcal{T}) \setminus \{\{4\}\}$ , and  $\mathbf{W}^{*(\{4\})} \in \mathbb{R}^{5,1}$ . We sampled the entries of  $(\mathbf{W}^{*(\nu)})_{\nu \in \mathcal{T}}$  independently from the standard normal distribution, and subsequently normalized the ground truth to be of Frobenius norm 64 (square root of its number of entries). Reconstruction error with respect to  $\mathcal{W}^*$  is based on normalized Frobenius distance, *i.e.* for a solution  $\mathcal{W}_H$  it is  $\|\mathcal{W}_H - \mathcal{W}^*\| / \|\mathcal{W}^*\|$ . The hierarchical tensor factorization applied to the task had 512 local components at all interior nodes due to computational and memory considerations (increasing the number of local components had no substantial impact on the dynamics). Standard deviation for initialization was set to 0.01.

**Figure 8:** plots correspond to the same experiment presented in Figure 1 (right).

**Figure 9:** implementation details are identical to those of Figure 1 (right), except that the following tensor sensing loss was minimized:  $\mathcal{L}_H(\mathcal{W}_H) = \sum_{i=1}^{2048} (\langle \otimes_{n=1}^4 \mathbf{x}^{(i,n)}, \mathcal{W}_H \rangle - \langle \otimes_{n=1}^4 \mathbf{x}^{(i,n)}, \mathcal{W}^* \rangle)^2$ , where the entries of  $(\mathbf{x}^{(i,1)}, \dots, \mathbf{x}^{(i,4)}) \in \mathbb{R}^8 \times \dots \times \mathbb{R}^8$  were sampled independently from a zero-mean Gaussian distribution with standard deviation  $4096^{-1/8}$  (ensures each measurement tensor  $\otimes_{n=1}^4 \mathbf{x}^{(i,n)}$  has expected square Frobenius norm 1).

**Figure 10:** implementation details are identical to those of Figure 1 (right), except that: (i) the ground truth tensor was of order 9 with modes of dimension 3, Frobenius norm  $\sqrt{19683}$  (square root of its number of entries), hierarchical tensor rank  $(2, \dots, 2)$ , and was generated according to a perfect ternary mode tree; (ii) reconstruction was based on 9840 entries chosen uniformly at random; (iii) the hierarchical tensor factorization applied to the task had 100 local components at all interior nodes; and (iv) standard deviation for initialization was set to 0.1.

**All:** using sample sizes smaller than those specified above led to similar results, up until a point where solutions found had fewer non-zero singular values, components, or local components (at all nodes) than the ground truths. Gradient descent was

<sup>14</sup>For any  $D_1, \dots, D_N \in \mathbb{N}$ , setting  $R = (\prod_{n=1}^N D_n) / \max\{D_n\}_{n=1}^N$  suffices for expressing all tensors in  $\mathbb{R}^{D_1, \dots, D_N}$  (Lemma 3.41 in Hackbusch (2012)).




Figure 8: Dynamics of gradient descent over order four hierarchical tensor factorization with a perfect binary mode tree (on tensor completion task) — incremental learning leads to low hierarchical tensor rank. For the hierarchical tensor factorization experiment in Figure 1 (right), plots present the evolution of additional quantities during optimization. **Left and second to left:** top 10 local component norms at nodes  $\{1, 2\}$  and  $\{1, 2, 3, 4\}$  (respectively) in the mode tree (the latter also appears in Figure 1 (right)). **Second to right and right:** top 10 singular values of the end tensor’s matricizations according to  $\{1, 2\}$  and  $\{1, 3\}$  (respectively). The former corresponds to a node in the mode tree, meaning its rank is part of the end tensor’s hierarchical tensor rank, whereas the latter does not. **All:** notice that, in line with our analysis (Section 4), local component norms move slower when small and faster when large, creating an incremental process that leads to low hierarchical tensor rank solutions. Moreover, the singular values of the end tensor’s matricizations according to nodes in the mode tree exhibit a similar behavior, whereas those of matricizations according to index sets outside the mode tree do not. The rank of a matricization lower bounds the (non-hierarchical) tensor rank (Remark 6.21 in Hackbusch (2012)). Thus, while the hierarchical tensor rank of the obtained solution is low, its tensor rank is high. For further implementation details, such as loss definition and factorization dimensions, see Appendix G.2.




Figure 9: Dynamics of gradient descent over order four hierarchical tensor factorization with a perfect binary mode tree (on tensor sensing task) — incremental learning leads to low hierarchical tensor rank. This figure is identical to Figure 8, except that the minimized mean squared error was based on random linear measurements (instead of randomly chosen entries). For further implementation details, such as loss definition and factorization dimensions, see Appendix G.2.




Figure 10: Dynamics of gradient descent over order nine hierarchical tensor factorization with a perfect ternary mode tree — incremental learning leads to low hierarchical tensor rank. This figure is identical to Figure 8, except that: (i) the hierarchical tensor factorization employed had order nine and complied with a perfect ternary mode tree; and (ii) the ground truth tensor was of hierarchical tensor rank  $(2, \dots, 2)$  (Definition 5). For further implementation details, such as loss definition and factorization dimensions, see Appendix G.2.

initialized randomly by sampling each weight in the factorization independently from a zero-mean Gaussian distribution, and was run until the loss remained under  $5 \cdot 10^{-5}$  for 100 iterations in a row. For each figure, experiments were carried out with initialization standard deviations 0.1, 0.05, 0.01, 0.005, 0.001, and 0.0005. Reported are representative runs striking a balance between the potency of the incremental learning effect and run time. Reducing standard deviations further did not yield a significant change in the dynamics, yet resulted in longer optimization times due to vanishing gradients around the origin.

To facilitate more efficient experimentation, we employed an adaptive learning rate scheme, where at each iteration a base learning rate is divided by the square root of an exponential moving average of squared gradient norms. That is,




Figure 11: Dedicated explicit regularization can counter the locality of convolutional networks, significantly improving performance on tasks with long-range dependencies. This figure is identical to Figure 5, except that: (i) experiments were carried out using a randomly initialized ResNet34 (as opposed to ResNet18); and (ii) it includes evaluation over a Pathfinder dataset with path length 11, since up until path length 9 an unregularized network still obtained non-trivial performance. For further details see Appendix G.2.2.




Figure 12: Dedicated explicit regularization can counter the locality of convolutional networks (regularized via standard techniques), significantly improving performance on tasks with long-range dependencies. This figure is identical to Figure 5, except that instead of applying our regularizer (Section 6.1) to a baseline unregularized network, the baseline networks here were regularized using either weight decay or dropout, and are compared to the results obtained when applying our regularization in addition to them. Figure 5 shows that the test accuracy obtained by an unregularized network substantially deteriorates when increasing the (spatial) range of dependencies required to be modeled. From the plots above it is evident that, even when employing standard regularization techniques such as weight decay or dropout, a similar degradation in performance occurs. As was the case for unregularized networks, applying our dedicated regularization, in addition to these techniques, significantly improved performance. In particular, for the combination of our regularization and dropout, the test accuracy was high across all datasets. For further details such as regularization hyperparameters, see Appendix G.2.2.




Figure 13: Dedicated explicit regularization can counter the locality of convolutional networks (regularized via standard techniques), significantly improving performance on tasks with long-range dependencies. This figure is identical to Figure 12, except that: (i) experiments were carried out using a randomly initialized ResNet34 (as opposed to ResNet18); and (ii) it includes evaluation over a Pathfinder dataset with path length 11, since up until path length 9 networks regularized using weight decay or dropout still obtained non-trivial performance. For further details such as regularization hyperparameters, see Appendix G.2.2.

for base learning rate  $\eta = 10^{-2}$  and weighted average coefficient  $\beta = 0.99$ , at iteration  $t$  the learning rate was set to  $\eta_t = \eta / (\sqrt{\gamma_t / (1 - \beta^t)} + 10^{-6})$ , where  $\gamma_t = \beta \cdot \gamma_{t-1} + (1 - \beta) \cdot \|\partial/\partial\Theta\phi(\Theta(t))\|^2$ ,  $\gamma_0 = 0$ ,  $\phi$  stands for any one of  $\phi_M$ ,  $\phi_T$ , or  $\phi_H$ , and  $\Theta$  denotes the corresponding factorization’s weights. We emphasize that only the learning rate (step size) is affected by this scheme, not the direction of movement. Comparisons between the scheme and optimization with a fixed learning rate showed no significant difference in terms of the dynamics, while run times of the former were considerably shorter.

### G.2.2 COUNTERING LOCALITY OF CONVOLUTIONAL NETWORKS VIA REGULARIZATION (FIGURES 5, 11, 12, AND 13)

In all experiments, we randomly initialized the ResNet18 and ResNet34 networks according to the default PyTorch (Paszke et al., 2017) implementation. The (regularized) binary cross-entropy loss was minimized via stochastic gradient descent with

Table 1: Hyperparameters for the regularizations employed in the experiments of Figures 12 and 13. For every model and dataset type combination, table reports the weight decay coefficient and dropout probability used when applied individually, as well as when combined with our regularization (described in Section 6.1), whose coefficients are also specified. These hyperparameters were tuned on the datasets with largest spatial range between salient regions of the input. That is, for each model separately, their values on IsSameClass datasets were set to those achieving the best test accuracy over a dataset with 160 pixels between CIFAR10 images. Similarly, their values on Pathfinder datasets were set to those achieving the best test accuracy over a dataset with connecting path length 9 for ResNet18 and path length 11 for ResNet34. For further details see the captions of Figures 12 and 13, as well as Appendix G.2.2.

	ResNet18		ResNet34	
	IsSameClass	Pathfinder	IsSameClass	Pathfinder
Weight Decay	0.001	0.01	0.01	0.001
Dropout	0.6	0.5	0.3	0.2
Ours & Weight Decay	1 & 0.001	0.1 & 0.01	1 & 0.0001	0.1 & 0.001
Ours & Dropout	1 & 0.5	0.1 & 0.4	1 & 0.5	0.5 & 0.3

learning rate 0.01, momentum coefficient 0.9, and batch size 64 (for ResNet34 we used a batch size of 32 and accumulated gradients over two batches due to GPU memory considerations). Optimization proceeded until perfect training accuracy was attained for 20 consecutive epochs or 150 epochs elapsed (runs without regularization always reached perfect training accuracy). For each dataset and model combination, runs were carried out using the regularization described in Section 6.1 with coefficients 0, 0.1, 0.5, 1, 3, 6, 9, and 10. Values lower than those reported in Figures 5 and 11 had no noticeable impact, whereas higher values typically did not allow fitting the training data. Table 1 specifies the hyperparameters used for the different regularizations in the experiments of Figures 12 and 13. Dropout layers shared the same probability hyperparameter, and were inserted before blocks expanding the number of channels, *i.e.* before the first convolutional layers with 128, 256, and 512 output channels (the default ResNet18 and ResNet34 implementations do not include dropout).

At each stochastic gradient descent iteration, the subset of indices  $I$  and  $J$  used for computing the regularized objective were sampled as follows. For IsSameClass datasets, we set  $I$  to be the indices marking either the left or right CIFAR10 image uniformly at random, and then let  $J$  be the indices corresponding to the remaining CIFAR10 image. For Pathfinder datasets,  $I$  and  $J$  were set to non-overlapping  $2 \times 2$  patches chosen uniformly across the input. In order to prevent additional computational overhead, alternative values for pixels indexed by  $J$  were taken from other images in the batch (as opposed to from the whole training set). Specifically, we used a permutation without fixed points to shuffle the pixel patches indexed by  $J$  across the batch.

IsSameClass datasets consisted of 5000 training and 10000 test samples. Each sample was generated by first drawing uniformly at random a label from  $\{0, 1\}$  and an image from CIFAR10. Then, depending on the chosen label, another image was sampled either from the same class (for label 1) or from all other classes (for label 0). Lastly, the CIFAR10 images were placed at a predetermined horizontal distance from each other around the center of a  $224 \times 224$  image filled with zeros. For example, when the horizontal distance is 0, the CIFAR10 images are adjacent, and when it is 160, they reside in opposite borders of the  $224 \times 224$  input. Pathfinder datasets consisted of 10000 training and 10000 test samples. Given a path length, the corresponding dataset was generated according to the protocol of Linsley et al. (2018), with hyperparameters: circle radius 3, paddle length 5, paddle thickness 2, inner paddle margin 3, and continuity 1.8. See Linsley et al. (2018) for additional information regarding the data generation process. As to be expected, when running a subset of all experiments using larger training set sizes (for both IsSameClass and Pathfinder datasets), we observed improved generalization across the board. Nevertheless, the addition of training samples did not alleviate the degradation in test accuracy observed for larger horizontal distances and path lengths, nor did it affect the beneficial impact of our regularization. That is, the trends observed in Figures 5, 11, 12, and 13 remained intact up to a certain shift upwards.

## H Deferred Proofs

### H.1 Additional Notation

Before delving into the proofs, we introduce the following notation.

**General** A colon is used to indicate a range of entries in a mode, *e.g.*  $\mathbf{W}_{i,:} \in \mathbb{R}^{D'}$  and  $\mathbf{W}_{:,j} \in \mathbb{R}^D$  are the  $i$ 'th row and  $j$ 'th column of  $\mathbf{W} \in \mathbb{R}^{D,D'}$ , respectively, and  $\mathbf{W}_{:,j} \in \mathbb{R}^{i,j}$  is the sub-matrix of  $\mathbf{W}$  consisting of its first  $i$  rows and  $j$  columns. For  $\mathcal{W} \in \mathbb{R}^{D_1, \dots, D_N}$ , we let  $\text{vec}(\mathcal{W}) \in \mathbb{R}^{\prod_{n=1}^N D_n}$  be its arrangement as a vector. The tensor and Kronecker products are denoted by  $\otimes$  and  $\odot$ , respectively.



**Hierarchical tensor factorization** For a mode tree  $\mathcal{T}$  over  $[N]$  (Definition 2), we denote the set of nodes in the subtree of  $\mathcal{T}$  whose root is  $\nu \in \mathcal{T}$  by  $\mathcal{T}(\nu) \subset \mathcal{T}$ . The sets of left and right siblings of  $\nu \in \mathcal{T}$  are denoted by  $\overleftarrow{S}(\nu)$  and  $\overrightarrow{S}(\nu)$ , respectively. For  $\nu \in \mathcal{T}$ , we let  $\mathcal{W}^{(\nu, \cdot)}$  be the tensor obtained by stacking  $(\mathcal{W}^{(\nu, r)})_{r=1}^{R_{Pa(\nu)}}$  into a single tensor, i.e.  $\mathcal{W}_{:, \dots, :, r}^{(\nu, \cdot)} = \mathcal{W}^{(\nu, r)}$  for all  $r \in [R_{Pa(\nu)}]$ . Given weight matrices  $(\mathbf{W}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}})_{\nu \in \mathcal{T}}$ , the function mapping them to the end tensor they produce according to Equation (4) is denoted by  $\mathcal{H}((\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}})$ . For  $\nu \in \mathcal{T}$ , with slight abuse of notation we let  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, \cdot)})$  be the function mapping  $(\mathcal{W}^{(\nu, r)})_{r=1}^{R_{Pa(\nu)}}$  and weight matrices outside of  $\mathcal{T}(\nu)$  to the end tensor they produce. Lastly, we denote by  $\overleftarrow{S}(\nu)$  and  $\overrightarrow{S}(\nu)$  the sets of left and right siblings of  $\nu \in \mathcal{T}$ , respectively.

## H.2 Useful Lemmas

### H.2.1 TECHNICAL

**Lemma 3.** For any  $\mathcal{U} \in \mathbb{R}^{D_1, \dots, D_N}$ ,  $\mathcal{V} \in \mathbb{R}^{H_1, \dots, H_K}$ , and  $I \subset [N + K]$ :

$$\llbracket \mathcal{U} \otimes \mathcal{V}; I \rrbracket = \llbracket \mathcal{U}; I \cap [N] \rrbracket \odot \llbracket \mathcal{V}; I - N \cap [K] \rrbracket,$$

where  $I - N := \{i - N : i \in I\}$ .

*Proof.* The identity follows directly from the definitions of the tensor and Kronecker products.  $\square$

**Lemma 4.** For any  $\mathbf{U} \in \mathbb{R}^{D_1, D_2}$ ,  $\mathbf{V} \in \mathbb{R}^{D_2, D_3}$ , and  $\mathbf{w} \in \mathbb{R}^{D_4}$ , the following holds:

$$(\mathbf{UV}) \odot \mathbf{w}^\top = \mathbf{U}(\mathbf{V} \odot \mathbf{w}^\top) \quad , \quad \mathbf{w}^\top \odot (\mathbf{UV}) = \mathbf{U}(\mathbf{w}^\top \odot \mathbf{V}).$$

*Proof.* According to the mixed-product property of the Kronecker product, for any matrices  $\mathbf{A}, \mathbf{A}', \mathbf{B}, \mathbf{B}'$  for which  $\mathbf{AA}'$  and  $\mathbf{BB}'$  are defined, it holds that  $(\mathbf{AA}') \odot (\mathbf{BB}') = (\mathbf{A} \odot \mathbf{B})(\mathbf{A}' \odot \mathbf{B}')$ . Thus:

$$(\mathbf{UV}) \odot \mathbf{w}^\top = (\mathbf{UV}) \odot (\mathbf{1} \cdot \mathbf{w}^\top) = (\mathbf{U} \odot \mathbf{1})(\mathbf{V} \odot \mathbf{w}^\top) = \mathbf{U}(\mathbf{V} \odot \mathbf{w}^\top),$$

where  $\mathbf{1}$  is treated as the 1-by-1 identity matrix. Similarly:

$$\mathbf{w}^\top \odot (\mathbf{UV}) = (\mathbf{1} \cdot \mathbf{w}^\top) \odot (\mathbf{UV}) = (\mathbf{1} \odot \mathbf{U})(\mathbf{w}^\top \odot \mathbf{V}) = \mathbf{U}(\mathbf{w}^\top \odot \mathbf{V}).$$

$\square$

### H.2.2 HIERARCHICAL TENSOR FACTORIZATION

Suppose that use a hierarchical tensor factorization with mode tree  $\mathcal{T}$ , weight matrices  $(\mathbf{W}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}})_{\nu \in \mathcal{T}}$ , and end tensor  $\mathcal{W}_H \in \mathbb{R}^{D_1, \dots, D_N}$  (Equations (4)) to minimize  $\phi_H$  (Equation (5)) via gradient flow (Equation (6)). Under this setting, we prove the following technical lemmas.

**Lemma 5.** The functions  $\mathcal{H}((\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}})$  and  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, \cdot)})$ , for  $\nu \in \mathcal{T}$ , defined in Appendix H.1, are multilinear.

*Proof.* We begin by proving that  $\mathcal{H}((\mathbf{W}^{(\nu)})_{\nu \in \mathcal{T}})$  is multilinear. Fix  $\nu \in \mathcal{T}$ , and let  $\mathbf{W}^{(\nu)}, \mathbf{U}^{(\nu)} \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}}$ , and  $\alpha > 0$ .

**Homogeneity** Denote by  $(\mathcal{U}_\alpha^{(\nu', r)})_{\nu' \in \mathcal{T}, r \in [R_{Pa(\nu')}]}$  the intermediate tensors produced when computing the end tensor  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \alpha \cdot \mathbf{W}^{(\nu)})$  according to Equation (4) (there denoted  $(\mathcal{W}^{(\nu', r)})_{\nu', r}$ ). If  $\nu$  is a leaf node, then  $\mathcal{U}_\alpha^{(\nu, r)} = \alpha \cdot \mathbf{W}_{:, r}^{(\nu)} = \alpha \cdot \mathcal{W}^{(\nu, r)}$  for all  $r \in [R_{Pa(\nu)}]$ . Otherwise, if  $\nu$  is an interior node, a straightforward computation leads to the same conclusion, i.e. for all  $r \in [R_{Pa(\nu)}]$ :

$$\begin{aligned} \mathcal{U}_\alpha^{(\nu, r)} &= \pi_\nu \left( \sum_{r'=1}^{R_\nu} \alpha \cdot \mathbf{W}_{r', r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) \\ &= \alpha \cdot \pi_\nu \left( \sum_{r'=1}^{R_\nu} \mathbf{W}_{r', r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) \\ &= \alpha \cdot \mathcal{W}^{(\nu, r)}, \end{aligned}$$

where the second equality is by the linearity of  $\pi_\nu$  (recall it is merely a reordering of the tensor entries). Moving on to the parent of  $\nu$ , multilinearity of the tensor product implies that for all  $r \in [R_{Pa(Pa(\nu))}]$ :

$$\begin{aligned} \mathcal{U}_\alpha^{(Pa(\nu),r)} &= \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \otimes \mathcal{U}_\alpha^{(\nu, r')} \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \right] \right) \\ &= \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \otimes (\alpha \cdot \mathcal{W}^{(\nu, r')}) \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \right] \right) \\ &= \alpha \cdot \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \otimes_{\nu_c \in C(Pa(\nu))} \mathcal{W}^{(\nu_c, r')} \right] \right) \\ &= \alpha \cdot \mathcal{W}^{(Pa(\nu), r)}. \end{aligned}$$

An inductive claim over the path from  $\nu$  to the root  $[N]$  therefore yields:

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \alpha \cdot \mathbf{W}^{(\nu)}) = \alpha \cdot \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}). \quad (13)$$

**Additivity** We let  $(\mathcal{U}^{(\nu', r)})_{\nu' \in \mathcal{T}, r \in [R_{Pa(\nu')}]}$  and  $(\mathcal{U}_+^{(\nu', r)})_{\nu' \in \mathcal{T}, r \in [R_{Pa(\nu')}]}$  denote the intermediate tensors produced when computing  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \mathbf{U}^{(\nu)})$  and  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \mathbf{W}^{(\nu)} + \mathbf{U}^{(\nu)})$  according to Equation (4), respectively. If  $\nu$  is a leaf node, we have that  $\mathcal{U}_+^{(\nu, r)} = \mathbf{W}_{:,r}^{(\nu)} + \mathbf{U}_{:,r}^{(\nu)} = \mathcal{W}^{(\nu, r)} + \mathcal{U}^{(\nu, r)}$  for all  $r \in [R_{Pa(\nu)}]$ . Otherwise, if  $\nu$  is an interior node, we arrive at the same conclusion, *i.e.* for all  $r \in [R_{Pa(\nu)}]$ :

$$\begin{aligned} \mathcal{U}_+^{(\nu, r)} &= \pi_\nu \left( \sum_{r'=1}^{R_\nu} (\mathbf{W}_{r',r}^{(\nu)} + \mathbf{U}_{r',r}^{(\nu)}) \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) \\ &= \pi_\nu \left( \sum_{r'=1}^{R_\nu} \mathbf{W}_{r',r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) + \pi_\nu \left( \sum_{r'=1}^{R_\nu} \mathbf{U}_{r',r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) \\ &= \mathcal{W}^{(\nu, r)} + \mathcal{U}^{(\nu, r)}, \end{aligned}$$

where the second equality is by the linearity of  $\pi_\nu$ . Then, for any  $r \in [R_{Pa(Pa(\nu))}]$ :

$$\begin{aligned} \mathcal{U}_+^{(Pa(\nu), r)} &= \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \otimes \mathcal{U}_+^{(\nu, r')} \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \right] \right) \\ &= \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \otimes (\mathcal{W}^{(\nu, r')} + \mathcal{U}^{(\nu, r')}) \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \right] \right) \\ &= \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \otimes_{\nu_c \in C(Pa(\nu))} \mathcal{W}^{(\nu_c, r')} \right] \right) + \\ &\quad \pi_{Pa(\nu)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r',r}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \otimes \mathcal{U}^{(\nu, r')} \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c, r')} \right) \right] \right) \\ &= \mathcal{W}^{(Pa(\nu), r)} + \mathcal{U}^{(Pa(\nu), r)}, \end{aligned}$$

where the penultimate equality is by multilinearity of the tensor product as well as linearity of  $\pi_{Pa(\nu)}$ . An induction over the path from  $\nu$  to the root thus leads to:

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \mathbf{W}^{(\nu)} + \mathbf{U}^{(\nu)}) = \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}) + \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \mathbf{U}^{(\nu)}). \quad (14)$$

Equations (13) and (14) establish that  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}})$  is multilinear. The proof for  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, \cdot)})$  follows by analogous derivations.  $\square$

**Lemma 6.** Suppose there exists  $\nu \in \mathcal{T}$  such that  $\mathcal{W}^{(\nu, r)} = 0$  for all  $r \in [R_{Pa(\nu)}]$ , where  $\mathcal{W}^{(\nu, r)}$  is as defined in Equation (4). Then,  $\mathcal{W}_H = 0$ .

*Proof.* Since  $\mathcal{W}^{(\nu,r)} = 0$  for all  $r \in [R_{Pa(\nu)}]$ , for any  $r' \in [R_{Pa(Pa(\nu))}]$  we have that:

$$\begin{aligned}\mathcal{W}^{(Pa(\nu),r')} &= \pi_{Pa(\nu)} \left( \sum_{r=1}^{R_{Pa(\nu)}} \mathbf{W}_{r,r'}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c,r)} \right) \otimes \mathcal{W}^{(\nu,r)} \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c,r)} \right) \right] \right) \\ &= \pi_{Pa(\nu)} \left( \sum_{r=1}^{R_{Pa(\nu)}} \mathbf{W}_{r,r'}^{(Pa(\nu))} \left[ \left( \otimes_{\nu_c \in \overleftarrow{S}(\nu)} \mathcal{W}^{(\nu_c,r)} \right) \otimes 0 \otimes \left( \otimes_{\nu_c \in \overrightarrow{S}(\nu)} \mathcal{W}^{(\nu_c,r)} \right) \right] \right) \\ &= 0.\end{aligned}$$

Thus, the claim readily follows by an induction up the path from  $\nu$  to the root  $[N]$ .  $\square$

**Lemma 7.** For any  $\nu \in \text{int}(\mathcal{T})$  and  $r \in [R_{Pa(\nu)}]$ :

$$\left\| \mathcal{W}^{(\nu,r)} \right\| \leq \left\| \mathbf{W}_{:,r}^{(\nu)} \right\| \cdot \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, :)} \right\|,$$

where  $\mathcal{W}^{(\nu_c, :)}$ , for  $\nu_c \in C(\nu)$ , is the tensor obtained by stacking  $(\mathcal{W}^{(\nu_c, r')})_{r'=1}^{R_\nu}$  into a single tensor, i.e.  $\mathcal{W}_{:, \dots, :, r'}^{(\nu_c, :)} = \mathcal{W}^{(\nu_c, r')}$  for all  $r' \in [R_\nu]$ .

*Proof.* By the definition of  $\mathcal{W}^{(\nu,r)}$  (Equation (4)) we have that:

$$\begin{aligned}\left\| \mathcal{W}^{(\nu,r)} \right\| &= \left\| \pi_\nu \left( \sum_{r'=1}^{R_\nu} \mathbf{W}_{r',r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right) \right\| \\ &= \left\| \sum_{r'=1}^{R_\nu} \mathbf{W}_{r',r}^{(\nu)} \left[ \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right] \right\|,\end{aligned}$$

where the second equality is due to the fact that  $\pi_\nu$  merely reorders entries of a tensor, and therefore does not alter its Frobenius norm. Vectorizing each  $\otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')}$ , we may write  $\left\| \mathcal{W}^{(\nu,r)} \right\|$  as the Frobenius norm of a matrix-vector product:

$$\left\| \mathcal{W}^{(\nu,r)} \right\| = \left\| \left( \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, 1)} \right), \dots, \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, R_\nu)} \right) \right) \mathbf{W}_{:,r}^{(\nu)} \right\|.$$

Hence, sub-multiplicativity of the Frobenius norm gives:

$$\left\| \mathcal{W}^{(\nu,r)} \right\| \leq \left\| \mathbf{W}_{:,r}^{(\nu)} \right\| \cdot \left\| \left( \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, 1)} \right), \dots, \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, R_\nu)} \right) \right) \right\|. \quad (15)$$

Notice that:

$$\begin{aligned}\left\| \left( \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, 1)} \right), \dots, \text{vec} \left( \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, R_\nu)} \right) \right) \right\|^2 &= \sum_{r'=1}^{R_\nu} \left\| \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r')} \right\|^2 \\ &= \sum_{r'=1}^{R_\nu} \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, r')} \right\|^2 \\ &\leq \prod_{\nu_c \in C(\nu)} \left( \sum_{r'=1}^{R_\nu} \left\| \mathcal{W}^{(\nu_c, r')} \right\|^2 \right) \\ &= \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, :)} \right\|^2,\end{aligned}$$

where the second transition is by the fact that the norm of a tensor product is equal to the product of the norms, and the inequality is due to  $\prod_{\nu_c \in C(\nu)} \left( \sum_{r'=1}^{R_\nu} \left\| \mathcal{W}^{(\nu_c, r')} \right\|^2 \right)$  being a sum of non-negative elements which includes  $\sum_{r'=1}^{R_\nu} \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, r')} \right\|^2$ . Taking the square root of both sides in the equation above and plugging it into Equation (15) completes the proof.  $\square$

**Lemma 8.** For any  $\nu \in \text{int}(\mathcal{T})$ :

$$\left\| \mathcal{W}^{(\nu, :)} \right\| \leq \left\| \mathbf{W}^{(\nu)} \right\| \cdot \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, :)} \right\|,$$

where  $\mathcal{W}^{(\nu_c, :)}$ , for  $\nu_c \in C(\nu)$ , is the tensor obtained by stacking  $(\mathcal{W}^{(\nu_c, r)})_{r=1}^{R_\nu}$  into a single tensor, i.e.  $\mathcal{W}_{:, \dots, :, r}^{(\nu_c, :)} = \mathcal{W}^{(\nu_c, r)}$  for all  $r \in [R_\nu]$ .

*Proof.* We may explicitly write  $\|\mathcal{W}^{(\nu, \cdot)}\|^2$  as follows:

$$\|\mathcal{W}^{(\nu, \cdot)}\|^2 = \sum_{r=1}^{R_{Pa(\nu)}} \|\mathcal{W}^{(\nu, r)}\|^2. \quad (16)$$

For each  $r \in [R_{Pa(\nu)}]$ , by Lemma 7 we know that:

$$\|\mathcal{W}^{(\nu, r)}\|^2 \leq \|\mathbf{W}_{:,r}^{(\nu)}\|^2 \cdot \prod_{\nu_c \in C(\nu)} \|\mathcal{W}^{(\nu_c, \cdot)}\|^2.$$

Thus, going back to Equation (16) we arrive at:

$$\|\mathcal{W}^{(\nu, \cdot)}\|^2 \leq \sum_{r=1}^{R_{Pa(\nu)}} \|\mathbf{W}_{:,r}^{(\nu)}\|^2 \cdot \prod_{\nu_c \in C(\nu)} \|\mathcal{W}^{(\nu_c, \cdot)}\|^2 = \|\mathbf{W}^{(\nu)}\|^2 \cdot \prod_{\nu_c \in C(\nu)} \|\mathcal{W}^{(\nu_c, \cdot)}\|^2.$$

Taking the square root of both sides concludes the proof.  $\square$

**Lemma 9.** For any  $\nu \in \mathcal{T}$  and  $\Delta \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}}$ :

$$\left\langle \frac{\partial}{\partial \mathbf{W}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \Delta \right\rangle = \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta) \right\rangle.$$

*Proof.* We treat  $(\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}$  as fixed, and with slight abuse of notation consider:

$$\phi_H^{(\nu)}(\mathbf{W}^{(\nu)}) := \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}).$$

For  $\Delta \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}}$ , by multilinearity of  $\mathcal{H}$  (Lemma 5) we have that:

$$\begin{aligned} \phi_H^{(\nu)}(\mathbf{W}^{(\nu)} + \Delta) &= \mathcal{L}_H\left(\mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \mathbf{W}^{(\nu)} + \Delta\right)\right) \\ &= \mathcal{L}_H\left(\mathcal{W}_H + \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta\right)\right). \end{aligned}$$

According to the first order Taylor approximation of  $\mathcal{L}_H$  we may write:

$$\begin{aligned} \phi_H^{(\nu)}(\mathbf{W}^{(\nu)} + \Delta) &= \mathcal{L}_H(\mathcal{W}_H) + \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta\right) \right\rangle + o(\|\Delta\|) \\ &= \phi_H^{(\nu)}(\mathbf{W}^{(\nu)}) + \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta\right) \right\rangle + o(\|\Delta\|). \end{aligned}$$

The term  $\left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta\right) \right\rangle$  is a linear function of  $\Delta$ . Therefore, uniqueness of the linear approximation of  $\phi_H^{(\nu)}$  at  $\mathbf{W}^{(\nu)}$  implies:

$$\left\langle \frac{d}{d\mathbf{W}^{(\nu)}} \phi_H^{(\nu)}(\mathbf{W}^{(\nu)}), \Delta \right\rangle = \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \Delta\right) \right\rangle.$$

Noticing that  $\frac{\partial}{\partial \mathbf{W}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}) = \frac{d}{d\mathbf{W}^{(\nu)}} \phi_H^{(\nu)}(\mathbf{W}^{(\nu)})$  completes the proof.  $\square$

**Lemma 10.** For any  $\nu \in \text{int}(\mathcal{T})$ ,  $r \in [R_\nu]$ , and  $\Delta \in \mathbb{R}^{R_{Pa(\nu)}}$ :

$$\left\langle \frac{\partial}{\partial \mathbf{W}_{r,:}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \Delta^\top \right\rangle = \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\Delta^\top)\right) \right\rangle, \quad (17)$$

where  $\text{PadR}_r(\Delta^\top) \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}}$  is the matrix whose  $r$ 'th row is  $\Delta^\top$ , and all the rest are zero. Furthermore, for any  $\nu_c \in C(\nu)$  and  $\Delta \in \mathbb{R}^{R_{\nu_c}}$ :

$$\left\langle \frac{\partial}{\partial \mathbf{W}_{:,r}^{(\nu_c)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \Delta \right\rangle = \left\langle \nabla \mathcal{L}_H(\mathcal{W}_H), \mathcal{H}\left((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\Delta)\right) \right\rangle, \quad (18)$$

where  $\text{PadC}_r(\Delta) \in \mathbb{R}^{R_{\nu_c}, R_\nu}$  is the matrix whose  $r$ 'th column is  $\Delta$ , and all the rest are zero.



*Proof.* Equations (17) and (18) are direct implications of Lemma 9 since:

$$\left\langle \frac{\partial}{\partial \mathbf{W}_{r,:}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \Delta^\top \right\rangle = \left\langle \frac{\partial}{\partial \mathbf{W}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \text{PadR}_r(\Delta^\top) \right\rangle,$$

and

$$\left\langle \frac{\partial}{\partial \mathbf{W}_{:,r}^{(\nu_c)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \Delta \right\rangle = \left\langle \frac{\partial}{\partial \mathbf{W}^{(\nu_c)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}), \text{PadC}_r(\Delta) \right\rangle.$$

□

**Lemma 11.** *Let  $\nu \in \text{int}(\mathcal{T})$ ,  $\nu_c \in C(\nu)$ , and  $r \in [R_\nu]$ . If both  $\mathbf{W}_{r,:}^{(\nu)} = 0$  and  $\mathbf{W}_{:,r}^{(\nu_c)} = 0$ , then:*

$$\frac{\partial}{\partial \mathbf{W}_{r,:}^{(\nu)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}) = 0, \quad (19)$$

and

$$\frac{\partial}{\partial \mathbf{W}_{:,r}^{(\nu_c)}} \phi_H((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T}}) = 0. \quad (20)$$

*Proof.* We show that  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\Delta^\top)) = 0$  for all  $\Delta \in \mathbb{R}^{Pa(\nu)}$ . Equation (19) then follows from Equation (17) in Lemma 10. Fix some  $\Delta \in \mathbb{R}^{Pa(\nu)}$  and let  $(\mathcal{U}^{(\nu', r')})_{\nu' \in \mathcal{T}, r' \in [R_{Pa(\nu')}]}$  be the intermediate tensors produced when computing  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\Delta^\top))$  according to Equation (4) (there denoted  $(\mathcal{W}^{(\nu', r')})_{\nu', r'}$ ). For any  $\bar{r} \in [R_{Pa(\nu)}]$  we have that:

$$\mathcal{U}^{(\nu, \bar{r})} = \pi_\nu \left( \sum_{r'=1}^{R_\nu} \text{PadR}_r(\Delta^\top)_{r', \bar{r}} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r')} \right] \right) = \pi_\nu \left( \Delta_{\bar{r}} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r')} \right] \right).$$

The fact that  $\mathbf{W}_{r,:}^{(\nu)} = 0$  implies that  $\mathcal{W}^{(\nu_c, r)} := \pi_{\nu_c} \left( \sum_{r'=1}^{R_{\nu_c}} \mathbf{W}_{r', r}^{(\nu_c)} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r')} \right] \right) = 0$ , and so for every  $r' \in [R_{Pa(\nu)}]$ :

$$\mathcal{U}^{(\nu, \bar{r})} = \pi_\nu \left( \Delta_{\bar{r}} \left[ \left( \otimes_{\nu' \in \overleftarrow{S}(\nu_c)} \mathcal{W}^{(\nu', r)} \right) \otimes 0 \otimes \left( \otimes_{\nu' \in \overrightarrow{S}(\nu_c)} \mathcal{W}^{(\nu', r)} \right) \right] \right) = 0.$$

Lemma 6 then gives  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\Delta^\top)) = 0$ , completing this part of the proof.

Next, we show that  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\Delta)) = 0$  for all  $\Delta \in \mathbb{R}^{\nu_c}$ . Equation (18) in Lemma 10 then yields Equation (20). Fix some  $\Delta \in \mathbb{R}^{\nu_c}$  and let  $(\mathcal{V}^{(\nu', r')})_{\nu' \in \mathcal{T}, r' \in [R_{Pa(\nu')}]}$  be the intermediate tensors produced when computing  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\Delta))$  according to Equation (4). For any  $\bar{r} \in [R_\nu] \setminus \{r\}$ :

$$\mathcal{U}^{(\nu_c, \bar{r})} = \pi_{\nu_c} \left( \sum_{r'=1}^{R_{\nu_c}} \text{PadC}_r(\Delta)_{r', \bar{r}} \left[ \otimes_{\nu' \in C(\nu_c)} \mathcal{W}^{(\nu', r')} \right] \right) = \pi_{\nu_c} \left( \sum_{r'=1}^{R_{\nu_c}} 0 \cdot \left[ \otimes_{\nu' \in C(\nu_c)} \mathcal{W}^{(\nu', r')} \right] \right) = 0.$$

Thus, for any  $\hat{r} \in [R_{Pa(\nu)}]$  we may write  $\mathcal{U}^{(\nu, \hat{r})} = \pi_\nu \left( \mathbf{W}_{r, \hat{r}}^{(\nu)} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{U}^{(\nu', r)} \right] \right)$ . Since  $\mathbf{W}_{r,:}^{(\nu)} = 0$ , we get that  $\mathcal{U}^{(\nu, \hat{r})} = 0$  for all  $\hat{r} \in [R_{Pa(\nu)}]$ , which by Lemma 6 leads to  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\Delta)) = 0$ . □

**Lemma 12.** *For any  $\nu \in \text{int}(\mathcal{T})$ ,  $\nu_c \in C(\nu)$ , and  $r \in [R_\nu]$ , the following hold:*

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)})) = \sigma_H^{(\nu, r)} \cdot \mathcal{C}_H^{(\nu, r)}, \quad (21)$$

and

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\mathbf{W}_{:,r}^{(\nu_c)})) = \sigma_H^{(\nu, r)} \cdot \mathcal{C}_H^{(\nu, r)}, \quad (22)$$

where  $\text{PadR}_r(\Delta^\top) \in \mathbb{R}^{R_\nu, R_{Pa(\nu)}}$  is the matrix whose  $r$ 'th row is  $\Delta^\top$ , and all the rest are zero,  $\text{PadC}_r(\Delta) \in \mathbb{R}^{R_{\nu_c}, R_\nu}$  is the matrix whose  $r$ 'th column is  $\Delta$ , and all the rest are zero, and  $\mathcal{C}_H^{(\nu, r)}$  is as defined in Theorem 1.

*Proof.* Starting with Equation (21), let  $(\mathbf{U}^{(\nu', r')})_{\nu' \in \mathcal{T}, r' \in [R_{Pa(\nu')}]}$  be the intermediate tensors formed when computing  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)}))$  according to Equation (4) (there denoted  $(\mathcal{W}^{(\nu', r')})_{\nu', r'}$ ). Clearly, for any  $\nu' \in \mathcal{T}(\nu) \setminus \{\nu\}$  — a node in the subtree of  $\nu$  which is not  $\nu$  — it holds that  $\mathcal{U}^{(\nu', r')} = \mathcal{W}^{(\nu', r')}$  for all  $r' \in [R_{Pa(\nu')}]$ . Thus, for all  $r' \in [R_{Pa(\nu)}]$  we have that:

$$\mathcal{U}^{(\nu, r')} = \pi_\nu \left( \sum_{\bar{r}=1}^{R_\nu} \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)})_{\bar{r}, r'} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', \bar{r})} \right] \right) = \pi_\nu \left( \mathbf{W}_{r, r'}^{(\nu)} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r)} \right] \right). \quad (23)$$

If  $\sigma_H^{(\nu, r)} = \|\mathbf{W}_{r,:}^{(\nu)} \otimes (\otimes_{\nu' \in C(\nu)} \mathbf{W}_{:,r}^{(\nu')})\| = \|\mathbf{W}_{r,:}^{(\nu)}\| \prod_{\nu' \in C(\nu)} \|\mathbf{W}_{:,r}^{(\nu')}\| = 0$ , then either  $\mathbf{W}_{r,:}^{(\nu)} = 0$  or  $\mathbf{W}_{:,r}^{(\nu')} = 0$  for some  $\nu' \in C(\nu)$ . We claim that in both cases  $\mathcal{U}^{(\nu, r')} = 0$  for all  $r' \in [R_{Pa(\nu)}]$ . Indeed, if  $\mathbf{W}_{r,:}^{(\nu)} = 0$  this immediately follows from Equation (23). On the other hand, if  $\mathbf{W}_{:,r}^{(\nu')} = 0$  for some  $\nu' \in C(\nu)$ , then  $\mathcal{W}^{(\nu', r)} = 0$ , which combined with Equation (23) also implies that  $\mathcal{U}^{(\nu, r')} = 0$  for all  $r' \in [R_{Pa(\nu)}]$ . Hence, Lemma 6 establishes Equation (22) for the case of  $\sigma_H^{(\nu, r)} = 0$ :

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)})) = 0 = \sigma_H^{(\nu, r)} \cdot \mathcal{C}_H^{(\nu, r)}.$$

Now, suppose that  $\sigma_H^{(\nu, r)} \neq 0$  and let  $\mathcal{U}^{(\nu, :)}$  be the tensor obtained by stacking  $(\mathcal{U}^{(\nu, r')})_{r'=1}^{R_{Pa(\nu)}}$  into a single tensor, i.e.  $\mathcal{U}_{:, \dots, :, r'}^{(\nu, :)} = \mathcal{U}^{(\nu, r')}$  for all  $r' \in [R_{Pa(\nu)}]$ . Multilinearity of  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{U}^{(\nu, :)})$  (Lemma 5) leads to:

$$\begin{aligned} \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)})) &= \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{U}^{(\nu, :)}) \\ &= \sigma_H^{(\nu, r)} \cdot \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, (\sigma_H^{(\nu, r)})^{-1} \mathcal{U}^{(\nu, :)}). \end{aligned} \quad (24)$$

From Equation (23) we know that  $(\sigma_H^{(\nu, r)})^{-1} \mathcal{U}_{:, \dots, :, r'}^{(\nu, :)} = \pi_\nu((\sigma_H^{(\nu, r)})^{-1} \mathbf{W}_{r, r'}^{(\nu)} [\otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r)}])$  for all  $r' \in [R_{Pa(\nu)}]$ . Thus, by the definition of  $\mathcal{C}_H^{(\nu, r)}$  we may conclude that:

$$\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, (\sigma_H^{(\nu, r)})^{-1} \mathcal{U}^{(\nu, :)}) = \mathcal{C}_H^{(\nu, r)}. \quad (25)$$

Combining Equations (24) and (25) yields Equation (22), completing this part of the proof.

Turning our attention to Equation (22), let  $(\mathcal{V}^{(\nu', r')})_{\nu' \in \mathcal{T}, r' \in [R_{Pa(\nu')}]}$  be the intermediate tensors produced when computing  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\mathbf{W}_{:,r}^{(\nu_c)}))$  according to Equation (4) (there denoted  $(\mathcal{W}^{(\nu', r')})_{\nu', r'}$ ). Clearly, for any  $\nu' \in \mathcal{T}(\nu) \setminus \{\nu, \nu_c\}$  — a node in the subtree of  $\nu$  which is not  $\nu$  nor  $\nu_c$  — it holds that  $\mathcal{V}^{(\nu', r')} = \mathcal{W}^{(\nu', r')}$  for all  $r' \in [R_{Pa(\nu')}]$ . Thus,  $\mathcal{V}^{(\nu_c, r)} = \pi_{\nu_c}(\sum_{\bar{r}=1}^{R_{\nu_c}} \text{PadC}_r(\mathbf{W}_{:,r}^{(\nu_c)})_{\bar{r}, r} [\otimes_{\nu' \in C(\nu_c)} \mathcal{W}^{(\nu', \bar{r})}]) = \mathcal{W}^{(\nu_c, r)}$ , whereas for any  $r' \in [R_\nu] \setminus \{r\}$ :

$$\begin{aligned} \mathcal{V}^{(\nu_c, r')} &= \pi_{\nu_c} \left( \sum_{\bar{r}=1}^{R_{\nu_c}} \text{PadC}_r(\mathbf{W}_{:,r}^{(\nu_c)})_{\bar{r}, r'} \left[ \otimes_{\nu' \in C(\nu_c)} \mathcal{W}^{(\nu', \bar{r})} \right] \right) \\ &= \pi_\nu \left( \sum_{\bar{r}=1}^{R_{\nu_c}} 0 \cdot \left[ \otimes_{\nu' \in C(\nu_c)} \mathcal{W}^{(\nu', \bar{r})} \right] \right) \\ &= 0. \end{aligned}$$

Putting it all together, for any  $r' \in [R_{Pa(\nu)}]$  we may write:

$$\mathcal{V}^{(\nu, r')} = \pi_\nu \left( \sum_{\bar{r}=1}^{R_\nu} \mathbf{W}_{\bar{r}, r'}^{(\nu)} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{U}^{(\nu', \bar{r})} \right] \right) = \pi_\nu \left( \mathbf{W}_{r, r'}^{(\nu)} \left[ \otimes_{\nu' \in C(\nu)} \mathcal{W}^{(\nu', r)} \right] \right).$$

From this point, following steps analogous to those used for proving Equation (21) based on Equation (23) yields Equation (22).  $\square$

**Lemma 13.** For any  $\nu \in \text{int}(\mathcal{T})$ ,  $\nu_c \in C(\nu)$ , and  $r \in [R_\nu]$ :

$$\frac{d}{dt} \left\| \mathbf{W}_{r,:}^{(\nu)}(t) \right\|^2 = 2\sigma_H^{(\nu, r)}(t) \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu, r)}(t) \right\rangle = \frac{d}{dt} \left\| \mathbf{W}_{:,r}^{(\nu_c)}(t) \right\|^2,$$

where  $\mathcal{C}_H^{(\nu, r)}(t)$  is as defined in Theorem 1.

*Proof.* Differentiating  $\|\mathbf{W}_{:,r}^{(\nu_c)}(t)\|^2$  with respect to time we get:

$$\frac{d}{dt} \|\mathbf{W}_{:,r}^{(\nu_c)}(t)\|^2 = 2 \left\langle \mathbf{W}_{:,r}^{(\nu_c)}(t), \frac{d}{dt} \mathbf{W}_{:,r}^{(\nu_c)}(t) \right\rangle = -2 \left\langle \mathbf{W}_{:,r}^{(\nu_c)}(t), \frac{\partial}{\partial \mathbf{W}_{:,r}^{(\nu_c)}} \phi_{\mathcal{H}}((\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T}}) \right\rangle.$$

By Equation (18) from Lemma 10 we have that:

$$\frac{d}{dt} \|\mathbf{W}_{:,r}^{(\nu_c)}(t)\|^2 = -2 \left\langle \nabla \mathcal{L}_{\mathcal{H}}(\mathcal{W}_{\mathcal{H}}(t)), \mathcal{H}\left((\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T} \setminus \{\nu_c\}}, \text{PadC}_r(\mathbf{W}_{:,r}^{(\nu_c)}(t))\right) \right\rangle.$$

Then, applying Equation (22) from Lemma 12 concludes:

$$\frac{d}{dt} \|\mathbf{W}_{:,r}^{(\nu_c)}(t)\|^2 = 2\sigma_{\mathcal{H}}^{(\nu,r)}(t) \left\langle -\nabla \mathcal{L}_{\mathcal{H}}(\mathcal{W}_{\mathcal{H}}(t)), \mathcal{C}_{\mathcal{H}}^{(\nu,r)}(t) \right\rangle.$$

A similar argument yields the desired result for  $\|\mathbf{W}_{r,:}^{(\nu)}(t)\|^2$ . Differentiating with respect to time we obtain:

$$\frac{d}{dt} \|\mathbf{W}_{r,:}^{(\nu)}(t)\|^2 = 2 \left\langle \mathbf{W}_{r,:}^{(\nu)}(t), \frac{d}{dt} \mathbf{W}_{r,:}^{(\nu)}(t) \right\rangle = -2 \left\langle \mathbf{W}_{r,:}^{(\nu)}(t), \frac{\partial}{\partial \mathbf{W}_{r,:}^{(\nu)}} \phi_{\mathcal{H}}((\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T}}) \right\rangle.$$

By Equation (17) from Lemma 10 we may write:

$$\frac{d}{dt} \|\mathbf{W}_{r,:}^{(\nu)}(t)\|^2 = -2 \left\langle \nabla \mathcal{L}_{\mathcal{H}}(\mathcal{W}_{\mathcal{H}}(t)), \mathcal{H}\left((\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T} \setminus \{\nu\}}, \text{PadR}_r(\mathbf{W}_{r,:}^{(\nu)}(t))\right) \right\rangle.$$

Lastly, applying Equation (21) from Lemma 12 completes the proof:

$$\frac{d}{dt} \|\mathbf{W}_{r,:}^{(\nu)}(t)\|^2 = 2\sigma_{\mathcal{H}}^{(\nu,r)}(t) \left\langle -\nabla \mathcal{L}_{\mathcal{H}}(\mathcal{W}_{\mathcal{H}}(t)), \mathcal{C}_{\mathcal{H}}^{(\nu,r)}(t) \right\rangle.$$

□

**Lemma 14.** *Let  $\nu \in \text{int}(\mathcal{T})$  and  $r \in [R_{\nu}]$ . If there exists a time  $t_0 \geq 0$  at which  $\mathbf{w}(t_0) = 0$  for all  $\mathbf{w} \in \text{LC}(\nu, r)$ , then:*

$$\mathbf{w}(t) = 0 \quad , t \geq 0, \mathbf{w} \in \text{LC}(\nu, r),$$

*i.e.  $\mathbf{w}(t)$  is identically zero for all  $\mathbf{w} \in \text{LC}(\nu, r)$ .*

*Proof.* Standard existence and uniqueness theorems (e.g. Theorem 2.2 in Teschl (2012)) imply that the system of differential equations governing gradient flow over  $\phi_{\mathcal{H}}$  (Equation (6)) has a unique solution that passes through  $(\mathbf{W}^{(\nu')}(t_0))_{\nu' \in \mathcal{T}}$  at time  $t_0$ . It therefore suffices to show that there exist  $(\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}$  satisfying Equation (6) such that  $\bar{\mathbf{W}}^{(\nu')}(t_0) = \mathbf{W}^{(\nu')}(t_0)$  for all  $\nu' \in \mathcal{T}$ , for which  $\bar{\mathbf{W}}_{r,:}^{(\nu)}(t)$  and  $(\bar{\mathbf{W}}_{:,r}^{(\nu_c)}(t))_{\nu_c \in C(\nu)}$  are zero for all  $t \geq 0$  (recall that  $\text{LC}(\nu, r)$  consists of  $\mathbf{W}_{r,:}^{(\nu)}$  and  $(\mathbf{W}_{:,r}^{(\nu_c)})_{\nu_c \in C(\nu)}$ ).

We denote by  $\Theta_{\nu,r}(t)$  all factorization weights at time  $t \geq 0$ , except for those in  $\text{LC}(\nu, r)$ , i.e.:

$$\Theta_{\nu,r}(t) := (\mathbf{W}^{(\nu')}(t))_{\nu' \in \mathcal{T} \setminus (\{\nu\} \cup C(\nu))} \cup (\mathbf{W}_{:,r'}^{(\nu_c)}(t))_{\nu_c \in C(\nu), r' \in [R_{\nu}] \setminus \{r\}} \cup (\mathbf{W}_{r',:}^{(\nu)}(t))_{r' \in [R_{\nu}] \setminus \{r\}}.$$

We construct  $(\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}$  as follows. First, let  $\bar{\mathbf{W}}_{r,:}^{(\nu)}(t) := 0$  and  $\bar{\mathbf{W}}_{:,r}^{(\nu_c)}(t) := 0$  for all  $\nu_c \in C(\nu)$  and  $t \geq 0$ . Then, considering  $\mathbf{W}_{r,:}^{(\nu)}(t)$  and  $(\mathbf{W}_{:,r}^{(\nu_c)}(t))_{\nu_c \in C(\nu)}$  as fixed to zero, we denote by  $\bar{\phi}_{HT}(\Theta_{\nu,r}(t))$  the induced objective over all other weights, and let

$$\bar{\Theta}_{\nu,r}(t) := (\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T} \setminus (\{\nu\} \cup C(\nu))} \cup (\bar{\mathbf{W}}_{:,r'}^{(\nu_c)}(t))_{\nu_c \in C(\nu), r' \in [R_{\nu}] \setminus \{r\}} \cup (\bar{\mathbf{W}}_{r',:}^{(\nu)}(t))_{r' \in [R_{\nu}] \setminus \{r\}}$$

be a gradient flow path over  $\bar{\phi}_{HT}$  satisfying  $\bar{\Theta}_{\nu,r}(t_0) = \Theta_{\nu,r}(t_0)$ . By definition, it holds that  $\bar{\mathbf{W}}^{(\nu')}(t_0) = \mathbf{W}^{(\nu')}(t_0)$  for all  $\nu' \in \mathcal{T}$ . Thus, it remains to show that  $(\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}$  obey the differential equations defining gradient flow over  $\phi_{\mathcal{H}}$

(Equation (6)). To see it is so, notice that since  $\bar{\mathbf{W}}_{r,:}^{(\nu)}(t)$  and  $(\bar{\mathbf{W}}_{:,r}^{(\nu_c)}(t))_{\nu_c \in C(\nu)}$  are identically zero, by the definition of  $\bar{\phi}_{HT}$  we have that:

$$\frac{d}{dt} \bar{\Theta}_{\nu,r}(t) = -\frac{d}{d\Theta_{\nu,r}} \bar{\phi}_{HT}(\bar{\Theta}_{\nu,r}(t)) = -\frac{\partial}{\partial \Theta_{\nu,r}} \phi_H((\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}). \quad (26)$$

Furthermore, by Lemma 11 we obtain:

$$\frac{d}{dt} \bar{\mathbf{W}}_{r,:}^{(\nu)}(t) = 0 = -\frac{\partial}{\partial \mathbf{W}_{r,:}^{(\nu)}} \phi_H((\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}), \quad (27)$$

and for all  $\nu_c \in C(\nu)$ :

$$\frac{d}{dt} \bar{\mathbf{W}}_{:,r}^{(\nu_c)}(t) = 0 = -\frac{\partial}{\partial \mathbf{W}_{:,r}^{(\nu_c)}} \phi_H((\bar{\mathbf{W}}^{(\nu')}(t))_{\nu' \in \mathcal{T}}). \quad (28)$$

Combining Equations (26), (27), and (28), completes the proof:

$$\frac{d}{dt} \bar{\mathbf{W}}^{(\nu')}(t) = -\frac{\partial}{\partial \mathbf{W}^{(\nu')}} \phi_H((\bar{\mathbf{W}}^{(\bar{\nu})}(t))_{\bar{\nu} \in \mathcal{T}}), \quad \nu' \in \mathcal{T}.$$

□

### H.3 Proof of Lemma 1

The proof follows a line similar to that of Theorem 7 in Cohen et al. (2018), extending from binary to arbitrary trees its upper bound on  $(\text{rank} \llbracket \mathcal{W}_H; \nu \rrbracket)_{\nu \in \mathcal{T}}$ .

Towards deriving a matricized form of Equation (4), we define the notion of *index set reduction*. For  $\nu \in \mathcal{T}$ , whose elements are denoted by  $i_1 < \dots < i_{|\nu|}$ , the reduction of  $\nu$  onto  $\nu' \in \mathcal{T}$  is defined by:

$$\nu|_{\nu'} := \{n \in \llbracket \nu' \rrbracket : i_n \in \nu \cap \nu'\}.$$

Now, fix  $\nu \in \text{int}(\mathcal{T})$  and  $\nu_c \in C(\nu)$ . By Lemma 3 and the linearity of the matricization operator, we may write the computation of  $\llbracket \mathcal{W}_H; \nu_c \rrbracket$  based on Equation (4) as follows:

For  $\bar{\nu} \in \{\{1\}, \dots, \{N\}\}$  (traverses leaves of  $\mathcal{T}$ ):

$$\mathcal{W}^{(\bar{\nu},r)} := \mathbf{W}_{:,r}^{(\bar{\nu})}, \quad r \in [R_{Pa(\bar{\nu})}],$$

for  $\bar{\nu} \in \text{int}(\mathcal{T}) \setminus \{\{N\}\}$  (traverses interior nodes of  $\mathcal{T}$  from leaves to root, non-inclusive):

$$\llbracket \mathcal{W}^{(\bar{\nu},r)}; \nu_c|_{\bar{\nu}} \rrbracket := \mathbf{Q}^{(\bar{\nu})} \left( \sum_{r'=1}^{R_{\bar{\nu}}} \mathbf{W}_{r',r}^{(\bar{\nu})} \left[ \odot_{\nu' \in C(\bar{\nu})} \llbracket \mathcal{W}^{(\nu',r')} ; \nu_c|_{\nu'} \rrbracket \right] \right) \bar{\mathbf{Q}}^{(\bar{\nu})}, \quad r \in [R_{Pa(\bar{\nu})}],$$

$$\llbracket \mathcal{W}_H; \nu_c \rrbracket = \mathbf{Q}^{([N])} \left( \sum_{r'=1}^{R_{[N]}} \mathbf{W}_{r',1}^{([N])} \left[ \odot_{\nu' \in C([N])} \llbracket \mathcal{W}^{(\nu',r')} ; \nu_c|_{\nu'} \rrbracket \right] \right) \bar{\mathbf{Q}}^{([N])},$$

where  $\mathbf{Q}^{(\bar{\nu})}$  and  $\bar{\mathbf{Q}}^{(\bar{\nu})}$ , for  $\bar{\nu} \in \text{int}(\mathcal{T})$ , are permutation matrices rearranging the rows and columns, respectively, to accord with an ascending order of  $\bar{\nu}$ , i.e. they fulfill the role of  $\pi_{\bar{\nu}}$  in Equation (4). For  $r \in [R_{Pa(\nu)}]$ , let us focus on  $\llbracket \mathcal{W}^{(\nu,r)}; \nu_c|_{\nu} \rrbracket$ . Since  $\nu_c|_{\nu_c} = \llbracket \nu_c \rrbracket$  and  $\nu_c|_{\nu'} = \emptyset$  for all  $\nu' \in C(\nu) \setminus \{\nu_c\}$ , we have that:

$$\begin{aligned} & \llbracket \mathcal{W}^{(\nu,r)}; \nu_c|_{\nu} \rrbracket \\ &= \mathbf{Q}^{(\nu)} \left( \sum_{r'=1}^{R_{\nu}} \mathbf{W}_{r',r}^{(\nu)} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu_c)} \llbracket \mathcal{W}^{(\nu',r')} ; \emptyset \rrbracket \right) \odot \llbracket \mathcal{W}^{(\nu_c,r')} ; \llbracket \nu_c \rrbracket \rrbracket \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu_c)} \llbracket \mathcal{W}^{(\nu',r')} ; \emptyset \rrbracket \right) \right] \right) \bar{\mathbf{Q}}^{(\nu)}. \end{aligned}$$

Notice that  $\llbracket \mathcal{W}^{(\nu',r')} ; \emptyset \rrbracket$  is a row vector, whereas  $\llbracket \mathcal{W}^{(\nu_c,r')} ; \llbracket \nu_c \rrbracket \rrbracket$  is a column vector, for each  $\nu' \in C(\nu) \setminus \{\nu_c\}$  and  $r' \in [R_{\nu}]$ . Commutativity of the Kronecker product between a row and column vectors therefore leads to:

$$\begin{aligned} \llbracket \mathcal{W}^{(\nu,r)}; \nu_c|_{\nu} \rrbracket &= \mathbf{Q}^{(\nu)} \left( \sum_{r'=1}^{R_{\nu}} \mathbf{W}_{r',r}^{(\nu)} \left[ \llbracket \mathcal{W}^{(\nu_c,r')} ; \llbracket \nu_c \rrbracket \rrbracket \odot \left( \odot_{\nu' \in C(\nu) \setminus \{\nu_c\}} \llbracket \mathcal{W}^{(\nu',r')} ; \emptyset \rrbracket \right) \right] \right) \bar{\mathbf{Q}}^{(\nu)} \\ &= \mathbf{Q}^{(\nu)} \left( \sum_{r'=1}^{R_{\nu}} \mathbf{W}_{r',r}^{(\nu)} \left[ \llbracket \mathcal{W}^{(\nu_c,r')} ; \llbracket \nu_c \rrbracket \rrbracket \left( \odot_{\nu' \in C(\nu) \setminus \{\nu_c\}} \llbracket \mathcal{W}^{(\nu',r')} ; \emptyset \rrbracket \right) \right] \right) \bar{\mathbf{Q}}^{(\nu)}, \end{aligned}$$



where the second equality is by the fact that for any column vector  $\mathbf{u}$  and row vector  $\mathbf{v}$  it holds that  $\mathbf{u} \odot \mathbf{v} = \mathbf{u}\mathbf{v}$ . Defining  $\mathbf{B}^{(\nu_c)}$  to be the matrix whose column vectors are  $\llbracket \mathcal{W}^{(\nu_c, 1)}; \llbracket \nu_c \rrbracket \rrbracket, \dots, \llbracket \mathcal{W}^{(\nu_c, R_\nu)}; \llbracket \nu_c \rrbracket \rrbracket$ , we can express the term between  $\mathbf{Q}^{(\nu)}$  and  $\bar{\mathbf{Q}}^{(\nu)}$  in the equation above as a  $\mathbf{B}^{(\nu_c)}\mathbf{A}^{(\nu, r)}$ , where  $\mathbf{A}^{(\nu, r)}$  is defined to be the matrix whose rows are  $\mathbf{W}_{1, r}^{(\nu)}(\odot_{\nu' \in C(\nu) \setminus \{\nu_c\}} \llbracket \mathcal{W}^{(\nu', 1)}; \emptyset \rrbracket), \dots, \mathbf{W}_{R_\nu, r}^{(\nu)}(\odot_{\nu' \in C(\nu) \setminus \{\nu_c\}} \llbracket \mathcal{W}^{(\nu', R_\nu)}; \emptyset \rrbracket)$ . That is:

$$\llbracket \mathcal{W}^{(\nu, r)}; \nu_c | \nu \rrbracket = \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \mathbf{A}^{(\nu, r)} \bar{\mathbf{Q}}^{(\nu)}. \quad (29)$$

The proof proceeds by propagating  $\mathbf{B}^{(\nu_c)}$  and the left permutation matrices up the tree, until reaching a representation of  $\llbracket \mathcal{W}_H; \nu_c \rrbracket$  as a product of matrices that includes  $\mathbf{B}^{(\nu_c)}$ . Since  $\mathbf{B}^{(\nu_c)}$  has  $R_\nu$  columns, this will imply that the rank of  $\llbracket \mathcal{W}_H; \nu_c \rrbracket$  is at most  $R_\nu$ , as required.

We begin with the propagation step from  $\nu$  to  $Pa(\nu)$ . For  $r \in [R_{Pa(Pa(\nu))}]$ , we examine:

$$\begin{aligned} & (\mathbf{Q}^{(Pa(\nu))})^{-1} \llbracket \mathcal{W}^{(Pa(\nu), r)}; \nu_c | Pa(\nu) \rrbracket (\bar{\mathbf{Q}}^{(Pa(\nu))})^{-1} \\ &= \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \nu_c | \nu' \rrbracket \right) \odot \llbracket \mathcal{W}^{(\nu, r)}; \nu_c | \nu \rrbracket \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \nu_c | \nu' \rrbracket \right) \right]. \end{aligned}$$

Plugging in Equation (29) while noticing that  $\nu_c | \nu' = \emptyset$  for any  $\nu'$  which is not an ancestor of  $\nu_c$ , we arrive at:

$$\sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right]. \quad (30)$$

Let  $r' \in [R_{Pa(\nu)}]$ . Since  $\llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket$  is a row vector for any  $\nu' \in C(Pa(\nu)) \setminus \{\nu\}$ , so are  $\odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket$  and  $\odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket$ . Applying Lemma 4 twice we therefore have that:

$$\begin{aligned} & \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \\ &= \left( \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \right] \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \\ &= \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right]. \end{aligned}$$

Going back to Equation (30), we obtain:

$$\begin{aligned} & \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right] \\ &= \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left( \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right] \right) \\ &= \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \left( \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right] \right). \quad (31) \end{aligned}$$

For brevity, we denote the matrix multiplying  $\mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)}$  from the right in the equation above by  $\mathbf{A}^{(Pa(\nu), r)}$ , i.e.:

$$\mathbf{A}^{(Pa(\nu), r)} := \sum_{r'=1}^{R_{Pa(\nu)}} \mathbf{W}_{r', r}^{(Pa(\nu))} \left[ \left( \odot_{\nu' \in \overleftarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \odot \left( \mathbf{A}^{(\nu, r')} \bar{\mathbf{Q}}^{(\nu)} \right) \odot \left( \odot_{\nu' \in \overrightarrow{S}(\nu)} \llbracket \mathcal{W}^{(\nu', r')}; \emptyset \rrbracket \right) \right].$$

Recalling that the expression in Equation (31) is of  $(\mathbf{Q}^{(Pa(\nu))})^{-1} \llbracket \mathcal{W}^{(Pa(\nu), r)}; \nu_c | Pa(\nu) \rrbracket (\bar{\mathbf{Q}}^{(Pa(\nu))})^{-1}$  completes the propagation step:

$$\llbracket \mathcal{W}^{(Pa(\nu), r)}; \nu_c | Pa(\nu) \rrbracket = \mathbf{Q}^{(Pa(\nu))} \mathbf{Q}^{(\nu)} \mathbf{B}^{(\nu_c)} \mathbf{A}^{(Pa(\nu), r)} \bar{\mathbf{Q}}^{(Pa(\nu))}.$$

Continuing this process, we propagate  $\mathbf{B}^{(\nu_c)}$ , along with the left permutation matrices, upwards in the tree until reaching the root. This brings forth the following representation of  $\llbracket \mathcal{W}_H; \nu_c \rrbracket$ :

$$\llbracket \mathcal{W}_H; \nu_c \rrbracket = \mathbf{Q}^{([N])} \mathbf{Q} \mathbf{B}^{(\nu_c)} \mathbf{A}^{([N])} \bar{\mathbf{Q}}^{([N])},$$

for appropriate  $\mathbf{Q}$  and  $\mathbf{A}^{([N])}$  encompassing the propagated permutation matrices and the ‘‘remainder’’ of the decomposition, respectively. Since  $\mathbf{B}^{(\nu_c)}$  has  $R_\nu$  columns, we may conclude:

$$\text{rank} \llbracket \mathcal{W}_H; \nu_c \rrbracket \leq \text{rank} \mathbf{B}^{(\nu_c)} \leq R_\nu.$$

□

#### H.4 Proof of Lemma 2

For any  $\nu \in \text{int}(\mathcal{T})$ ,  $r \in [R_\nu]$ , and  $\mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)$ , Lemma 13 implies that:

$$\frac{d}{dt} \|\mathbf{w}(t)\|^2 = 2\sigma_{\text{H}}^{(\nu, r)}(t) \left\langle -\nabla \mathcal{L}_{\text{H}}(\mathcal{W}_{\text{H}}(t)), \mathcal{C}_{\text{H}}^{(\nu, r)}(t) \right\rangle = \frac{d}{dt} \|\mathbf{w}'(t)\|^2,$$

where  $\mathcal{C}_{\text{H}}^{(\nu, r)}(t)$  is as defined in Theorem 1. For  $t \geq 0$ , integrating both sides with respect to time leads to:

$$\|\mathbf{w}(t)\|^2 - \|\mathbf{w}(0)\|^2 = \|\mathbf{w}'(t)\|^2 - \|\mathbf{w}'(0)\|^2.$$

Rearranging the equality above yields the desired result.  $\square$

#### H.5 Proof of Theorem 1

Let  $t \geq 0$ .

First, suppose that  $\sigma_{\text{H}}^{(\nu, r)}(t) = 0$ . Since the unbalancedness magnitude at initialization is zero, from Lemma 2 we know that  $\|\mathbf{w}(t)\| = \|\mathbf{w}'(t)\|$  for any  $\mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)$ . Hence, the fact that  $\sigma_{\text{H}}^{(\nu, r)}(t) = 0$  implies that  $\|\mathbf{w}(t)\| = 0$  for all  $\mathbf{w} \in \text{LC}(\nu, r)$ . Lemma 14 then establishes that  $\sigma_{\text{H}}^{(\nu, r)}(t)$  is identically zero through time, in which case both sides of Equation (7) are equal to zero.

We now move to the case where  $\sigma_{\text{H}}^{(\nu, r)}(t) > 0$ . Since  $\sigma_{\text{H}}^{(\nu, r)}(t) = \|\otimes_{\mathbf{w} \in \text{LC}(\nu, r)} \mathbf{w}(t)\| = \prod_{\mathbf{w} \in \text{LC}(\nu, r)} \|\mathbf{w}(t)\|$  (the norm of a tensor product is equal to the product of the norms), by the product rule of differentiation we have that:

$$\frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)^2 = \sum_{\mathbf{w} \in \text{LC}(\nu, r)} \frac{d}{dt} \|\mathbf{w}(t)\|^2 \cdot \prod_{\mathbf{w}' \in \text{LC}(\nu, r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2.$$

Applying Lemma 13 then leads to:

$$\begin{aligned} \frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)^2 &= \sum_{\mathbf{w} \in \text{LC}(\nu, r)} 2\sigma_{\text{H}}^{(\nu, r)}(t) \left\langle -\nabla \mathcal{L}_{\text{H}}(\mathcal{W}_{\text{H}}(t)), \mathcal{C}_{\text{H}}^{(\nu, r)}(t) \right\rangle \cdot \prod_{\mathbf{w}' \in \text{LC}(\nu, r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2 \\ &= 2\sigma_{\text{H}}^{(\nu, r)}(t) \left\langle -\nabla \mathcal{L}_{\text{H}}(\mathcal{W}_{\text{H}}(t)), \mathcal{C}_{\text{H}}^{(\nu, r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu, r)} \prod_{\mathbf{w}' \in \text{LC}(\nu, r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2. \end{aligned}$$

From the chain rule we know that  $\frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)^2 = 2\sigma_{\text{H}}^{(\nu, r)}(t) \cdot \frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)$  (note that  $\frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)$  surely exists because  $\sigma_{\text{H}}^{(\nu, r)}(t) > 0$ ). Thus:

$$\begin{aligned} \frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t) &= \frac{1}{2} \sigma_{\text{H}}^{(\nu, r)}(t)^{-1} \frac{d}{dt} \sigma_{\text{H}}^{(\nu, r)}(t)^2 \\ &= \left\langle -\nabla \mathcal{L}_{\text{H}}(\mathcal{W}_{\text{H}}(t)), \mathcal{C}_{\text{H}}^{(\nu, r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu, r)} \prod_{\mathbf{w}' \in \text{LC}(\nu, r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2. \end{aligned} \quad (32)$$

According to Lemma 2, the unbalancedness magnitude remains zero through time, and so  $\|\mathbf{w}(t)\| = \|\mathbf{w}'(t)\|$  for any  $\mathbf{w}, \mathbf{w}' \in \text{LC}(\nu, r)$ . Recalling that  $L_\nu := C(\nu) + 1$  is the number of weight vectors in a local component at  $\nu$ , this implies that for each  $\mathbf{w} \in \text{LC}(\nu, r)$ :

$$\|\mathbf{w}(t)\|^2 = \|\mathbf{w}(t)\|^{L_\nu \cdot \frac{2}{L_\nu}} = \left( \prod_{\mathbf{w}' \in \text{LC}(\nu, r)} \|\mathbf{w}'(t)\| \right)^{\frac{2}{L_\nu}} = \sigma_{\text{H}}^{(\nu, r)}(t)^{\frac{2}{L_\nu}}. \quad (33)$$

Plugging Equation (33) into Equation (32) completes the proof.  $\square$

#### H.6 Proof of Proposition 1

We begin by establishing the following key lemma, which upper bounds the distance between the end tensor  $\mathcal{W}_{\text{H}}$  and the one obtained after setting a local component to zero.

**Lemma 15.** *Let  $\nu \in \text{int}(\mathcal{T})$  and  $r \in [R_\nu]$ . Denote by  $\overline{\mathcal{W}}_{\text{HT}}^{(\nu, r)}$  the end tensor obtained by pruning the  $(\nu, r)$ 'th local component, i.e. by setting the  $r$ 'th row of  $\mathbf{W}^{(\nu)}$  and the  $r$ 'th columns of  $(\mathbf{W}^{(\nu_c)})_{\nu_c \in \mathcal{C}(\nu)}$  to zero. Then:*

$$\|\mathcal{W}_{\text{H}} - \overline{\mathcal{W}}_{\text{HT}}^{(\nu, r)}\| \leq \sigma_{\text{H}}^{(\nu, r)} \cdot \prod_{\nu' \in \mathcal{T} \setminus (\{\nu\} \cup \mathcal{C}(\nu))} \|\mathbf{W}^{(\nu')}\|.$$

*Proof.* Let  $(\bar{\mathbf{W}}^{(\nu')})_{\nu' \in \mathcal{T}}$  be the weight matrices corresponding to  $\bar{\mathcal{W}}_{HT}^{(\nu,r)}$ , i.e.  $\bar{\mathbf{W}}^{(\nu)}$  is the weight matrix obtained by setting the  $r$ 'th row of  $\mathbf{W}^{(\nu)}$  to zero,  $(\bar{\mathbf{W}}^{(\nu_c)})_{\nu_c \in C(\nu)}$  are the weight matrices obtained by setting the  $r$ 'th columns of  $(\mathbf{W}^{(\nu_c)})_{\nu_c \in C(\nu)}$  to zero, and  $\bar{\mathbf{W}}^{(\nu')} = \mathbf{W}^{(\nu')}$  for all  $\nu' \in \mathcal{T} \setminus (\{\nu\} \cup C(\nu))$ . Accordingly, we denote by  $(\bar{\mathcal{W}}^{(\nu',r')})_{\nu' \in \mathcal{T}, r' \in [R_{Pa(\nu')}]}$  the intermediate tensors produced when computing  $\bar{\mathcal{W}}_{HT}^{(\nu,r)}$  according to Equation (4) (there denoted  $(\mathcal{W}^{(\nu',r')})_{\nu', r'}$ ).

By definition,  $\mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, :)}) = \mathcal{W}_H$  and  $\mathcal{H}((\bar{\mathbf{W}}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \bar{\mathcal{W}}^{(\nu, :)}) = \bar{\mathcal{W}}_{HT}^{(\nu,r)}$ . Since  $\mathcal{H}$  is multilinear (Lemma 5) and  $\bar{\mathbf{W}}^{(\nu')} = \mathbf{W}^{(\nu')}$  for all  $\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)$ , we have that:

$$\begin{aligned} \left\| \mathcal{W}_H - \bar{\mathcal{W}}_{HT}^{(\nu,r)} \right\| &= \left\| \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, :)}) - \mathcal{H}((\bar{\mathbf{W}}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \bar{\mathcal{W}}^{(\nu, :)}) \right\| \\ &= \left\| \mathcal{H}((\mathbf{W}^{(\nu')})_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)}, \mathcal{W}^{(\nu, :)} - \bar{\mathcal{W}}^{(\nu, :)}) \right\|. \end{aligned}$$

Heading from the root downwards, subsequent applications of Lemma 8 over all nodes in the mode tree, except those belonging to the sub-tree whose root is  $\nu$ , then yield:

$$\left\| \mathcal{W}_H - \bar{\mathcal{W}}_{HT}^{(\nu,r)} \right\| \leq \left\| \mathcal{W}^{(\nu, :)} - \bar{\mathcal{W}}^{(\nu, :)} \right\| \cdot \prod_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)} \left\| \mathbf{W}^{(\nu')} \right\|. \quad (34)$$

Notice that for any  $r' \in [R_{Pa(\nu)}]$ :

$$\begin{aligned} \left( \mathcal{W}^{(\nu, :)} - \bar{\mathcal{W}}^{(\nu, :)} \right)_{:, \dots, :, r'} &= \sum_{\bar{r} \in [R_\nu]} \mathbf{W}_{\bar{r}, r'}^{(\nu)} \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, \bar{r})} - \sum_{\bar{r} \in [R_\nu] \setminus \{r\}} \mathbf{W}_{\bar{r}, r'}^{(\nu)} \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, \bar{r})} \\ &= \mathbf{W}_{r, r'}^{(\nu)} \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r)}. \end{aligned}$$

Thus, a straightforward computation shows:

$$\begin{aligned} \left\| \mathcal{W}^{(\nu, :)} - \bar{\mathcal{W}}^{(\nu, :)} \right\|^2 &= \sum_{r'=1}^{R_{Pa(\nu')}} \left\| \mathbf{W}_{r, r'}^{(\nu)} \otimes_{\nu_c \in C(\nu)} \mathcal{W}^{(\nu_c, r)} \right\|^2 \\ &= \sum_{r'=1}^{R_{Pa(\nu')}} \left( \mathbf{W}_{r, r'}^{(\nu)} \right)^2 \cdot \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, r)} \right\|^2 \\ &= \left\| \mathbf{W}_{r, :}^{(\nu)} \right\|^2 \cdot \prod_{\nu_c \in C(\nu)} \left\| \mathcal{W}^{(\nu_c, r)} \right\|^2, \end{aligned}$$

where the second equality is by the fact that the norm of a tensor product is equal to the product of the norms. From Lemma 7 we get that  $\left\| \mathcal{W}^{(\nu_c, r)} \right\| \leq \left\| \mathbf{W}_{:, r}^{(\nu_c)} \right\| \cdot \prod_{\nu' \in C(\nu_c)} \left\| \mathcal{W}^{(\nu', :)} \right\|$  for all  $\nu_c \in C(\nu)$ , which leads to:

$$\begin{aligned} \left\| \mathcal{W}^{(\nu, :)} - \bar{\mathcal{W}}^{(\nu, :)} \right\|^2 &\leq \left\| \mathbf{W}_{r, :}^{(\nu)} \right\|^2 \cdot \prod_{\nu_c \in C(\nu)} \left( \left\| \mathbf{W}_{:, r}^{(\nu_c)} \right\|^2 \cdot \prod_{\nu' \in C(\nu_c)} \left\| \mathcal{W}^{(\nu', :)} \right\|^2 \right) \\ &= \left( \sigma_H^{(\nu, r)} \right)^2 \cdot \prod_{\nu_c \in C(\nu), \nu' \in C(\nu_c)} \left\| \mathcal{W}^{(\nu', :)} \right\|^2. \end{aligned}$$

Taking the square root of both sides and plugging the inequality above into Equation (34), we arrive at:

$$\left\| \mathcal{W}_H - \bar{\mathcal{W}}_{HT}^{(\nu,r)} \right\| \leq \sigma_H^{(\nu, r)} \cdot \prod_{\nu_c \in C(\nu), \nu' \in C(\nu_c)} \left\| \mathcal{W}^{(\nu', :)} \right\| \cdot \prod_{\nu' \in \mathcal{T} \setminus \mathcal{T}(\nu)} \left\| \mathbf{W}^{(\nu')} \right\|.$$

Applying Lemma 8 iteratively over the sub-trees whose roots are  $C(\nu_c)$  gives:

$$\prod_{\nu_c \in C(\nu), \nu' \in C(\nu_c)} \left\| \mathcal{W}^{(\nu', :)} \right\| \leq \prod_{\nu' \in \mathcal{T}(\nu) \setminus (\{\nu\} \cup C(\nu))} \left\| \mathbf{W}^{(\nu')} \right\|,$$

concluding the proof.  $\square$

With Lemma 15 in hand, we are now in a position to prove Proposition 1. Let

$$\mathcal{S} := \{(\nu, r) : \nu \in \text{int}(\mathcal{T}), r \in \{R_\nu^l + 1, \dots, R_\nu\}\},$$

and denote by  $\overline{\mathcal{W}}_{HT}^S$  the end tensor obtained by pruning all local components in  $\mathcal{S}$ , *i.e.* by setting to zero the  $r$ 'th row of  $\mathbf{W}^{(\nu)}$  and the  $r$ 'th column of  $\mathbf{W}^{(\nu_c)}$  for all  $(\nu, r) \in \mathcal{S}$  and  $\nu_c \in C(\nu)$ . As can be seen from Equation (4), we may equivalently discard these weight vectors instead of setting them to zero. Doing so, we arrive at a representation of  $\overline{\mathcal{W}}_{HT}^S$  as the end tensor of  $(\overline{\mathbf{W}}^{(\nu)} \in \mathbb{R}^{R'_\nu \times R'_{Pa(\nu)}})_{\nu \in \mathcal{T}}$ , where  $R'_{Pa([N])} = 1$ ,  $R'_{\{n\}} = D_n$  for  $n \in [N]$ , and  $\overline{\mathbf{W}}^{(\nu)} = \mathbf{W}^{(\nu)}_{:R'_\nu, :R'_{Pa(\nu)}}$  for all  $\nu \in \mathcal{T}$ . Hence, Lemma 1 implies that for any  $\nu \in \mathcal{T}$  the rank of  $\llbracket \overline{\mathcal{W}}_{HT}^S; \nu \rrbracket$  is at most  $R'_{Pa(\nu)}$ . This means that it suffices to show that:

$$\|\mathcal{W}_H - \overline{\mathcal{W}}_{HT}^S\| \leq \epsilon. \quad (35)$$

For  $i \in \llbracket |\mathcal{S}| \rrbracket$ , let  $\mathcal{S}_i \subset \mathcal{S}$  be the set comprising the first  $i$  local components in  $\mathcal{S}$  according to an arbitrary order. Adding and subtracting  $\overline{\mathcal{W}}_{HT}^{\mathcal{S}_i}$  for all  $i \in \llbracket |\mathcal{S}| - 1 \rrbracket$ , and applying the triangle inequality, we have:

$$\|\mathcal{W}_H - \overline{\mathcal{W}}_{HT}^S\| \leq \sum_{i=0}^{|\mathcal{S}|-1} \left\| \overline{\mathcal{W}}_{HT}^{\mathcal{S}_i} - \overline{\mathcal{W}}_{HT}^{\mathcal{S}_{i+1}} \right\|,$$

where  $\overline{\mathcal{W}}_{HT}^{\mathcal{S}_0} := \mathcal{W}_H$ . Upper bounding each term in the sum according to Lemma 15, while noticing that pruning a local component can only decrease the norms of weight matrices and other local components in the factorization, we obtain:

$$\begin{aligned} \|\mathcal{W}_H - \overline{\mathcal{W}}_{HT}^S\| &\leq \sum_{\nu \in \text{int}(\mathcal{T})} \sum_{r=R'_\nu+1}^{R_\nu} \sigma_H^{(\nu, r)} \cdot \prod_{\nu' \in \mathcal{T} \setminus (\{\nu\} \cup C(\nu))} \|\mathbf{W}^{(\nu')}\| \\ &\leq \sum_{\nu \in \text{int}(\mathcal{T})} B^{|\mathcal{T}|-1-|C(\nu)|} \cdot \sum_{r=R'_\nu+1}^{R_\nu} \sigma_H^{(\nu, r)}, \end{aligned}$$

where the latter inequality is by recalling that  $B = \max_{\nu \in \mathcal{T}} \|\mathbf{W}^{(\nu)}\|$ . Since for all  $\nu \in \text{int}(\mathcal{T})$  we have that  $\sum_{r=R'_\nu+1}^{R_\nu} \sigma_H^{(\nu, r)} \leq \epsilon \cdot (|\mathcal{T}| - N)^{-1} B^{|C(\nu)|+1-|\mathcal{T}|}$ , Equation (35) readily follows.  $\square$

## H.7 Proof of Proposition 2

Consider the tensor completion problem defined by the set of observed entries

$$\Omega = \{(1, \dots, 1, 1, 1, 1), (1, \dots, 1, 1, 1, 2), (1, \dots, 1, 2, 2, 1), (1, \dots, 1, 2, 2, 2)\}$$

and ground truth  $\mathcal{W}^* \in \mathbb{R}^{D_1, \dots, D_N}$ , whose values at those locations are:

$$\mathcal{W}_{1, \dots, 1, 1, :2, :2}^* = \begin{bmatrix} 1 & 0 \\ ? & ? \end{bmatrix}, \quad \mathcal{W}_{1, \dots, 1, 2, :2, :2}^* = \begin{bmatrix} ? & ? \\ 0 & 1 \end{bmatrix}, \quad (36)$$

where ? stands for an unobserved entry. We define two solutions for the tensor completion problem,  $\mathcal{W}$  and  $\mathcal{W}'$  in  $\mathbb{R}^{D_1, \dots, D_N}$ , as follows:

$$\mathcal{W}_{1, \dots, 1, 1, :2, :2} := \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}, \quad \mathcal{W}_{1, \dots, 1, 2, :2, :2} := \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}, \quad \mathcal{W}'_{1, \dots, 1, 1, :2, :2} := \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathcal{W}'_{1, \dots, 1, 2, :2, :2} := \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

and the remaining entries of  $\mathcal{W}$  and  $\mathcal{W}'$  hold zero. Clearly,  $\mathcal{L}(\mathcal{W}) = \mathcal{L}(\mathcal{W}') = 0$ .

Fix a mode tree  $\mathcal{T}$  over  $[N]$ . Since  $\mathcal{W}$  and  $\mathcal{W}'$  fit the observed entries their hierarchical tensor ranks with respect to  $\mathcal{T}$ ,  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  and  $(\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$ , are in  $\mathcal{R}_{\mathcal{T}}$ . We prove that neither  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  nor  $(\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  (with respect to the standard product partial order), by examining the matrix ranks of the matricizations of  $\mathcal{W}$  and  $\mathcal{W}'$  according to  $\{N-2\} \in \mathcal{T}$  and  $\{N-1\} \in \mathcal{T}$  (recall that any mode tree has leaves  $\{1\}, \dots, \{N\}$ ). For  $\{N-2\}$ , we have that  $\text{rank} \llbracket \mathcal{W}; \{N-2\} \rrbracket = 2$  whereas  $\text{rank} \llbracket \mathcal{W}'; \{N-2\} \rrbracket = 1$ . To see it is so, notice that:

$$\llbracket \mathcal{W}; \{N-2\} \rrbracket_{:2, :4} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \quad \llbracket \mathcal{W}'; \{N-2\} \rrbracket_{:2, :4} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix},$$

and all other entries of  $\llbracket \mathcal{W}; \{N-2\} \rrbracket$  and  $\llbracket \mathcal{W}'; \{N-2\} \rrbracket$  hold zero. This means that  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  does not hold. On the other hand, for  $\{N-1\}$  we have that  $\text{rank} \llbracket \mathcal{W}; \{N-1\} \rrbracket = 1$  while  $\text{rank} \llbracket \mathcal{W}'; \{N-1\} \rrbracket = 2$ , because:

$$\llbracket \mathcal{W}; \{N-1\} \rrbracket_{:2, :4} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix}, \quad \llbracket \mathcal{W}'; \{N-1\} \rrbracket_{:2, :4} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix},$$

and the remaining entries of  $\llbracket \mathcal{W}; \{N-1\} \rrbracket$  and  $\llbracket \mathcal{W}'; \{N-1\} \rrbracket$  hold zero. This implies that  $(\text{rank } \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  does not hold, and so the hierarchical tensor ranks of  $\mathcal{W}$  and  $\mathcal{W}'$  are incomparable, *i.e.* neither is smaller than or equal to the other.

It remains to show that there exists no  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}} \in \mathcal{R}_{\mathcal{T}} \setminus \{(\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}, (\text{rank } \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}\}$  satisfying  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  or  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank } \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$ . Assume by way of contradiction that there exists such  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}}$ , and let  $\mathcal{W}'' \in \mathbb{R}^{D_1, \dots, D_N}$  be a solution of this hierarchical tensor rank. We now prove that  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  entails a contradiction. Since  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}}$  is not equal to the hierarchical tensor rank of  $\mathcal{W}$ , there exists  $\nu \in \mathcal{T} \setminus \{[N]\}$  for which  $\text{rank } \llbracket \mathcal{W}''; \nu \rrbracket = R''_{\nu} < \text{rank } \llbracket \mathcal{W}; \nu \rrbracket$ . Let us examine the possible cases:

- If  $\nu$  does not contain  $N-2$ ,  $N-1$ , and  $N$ , then  $\text{rank } \llbracket \mathcal{W}; \nu \rrbracket = 1$  as all rows but the first of this matricization are zero. In this case  $\llbracket \mathcal{W}''; \nu \rrbracket = R''_{\nu} = 0$ , implying that  $\mathcal{W}''$  is the zero tensor, in contradiction to it fitting the (non-zero) observed entries from Equation (36).
- If  $\nu$  contains  $N$  but not  $N-2$  and  $N-1$ , then  $\text{rank } \llbracket \mathcal{W}; \nu \rrbracket = 2$  since:

$$\llbracket \mathcal{W}; \nu \rrbracket_{:,2:,4} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

and all other entries of  $\llbracket \mathcal{W}; \nu \rrbracket$  hold zero. In this case  $\llbracket \mathcal{W}''; \nu \rrbracket = R''_{\nu} < 2$ . However, the fact that  $\mathcal{W}''$  fits the observed entries from Equation (36) leads to a contradiction, as  $\llbracket \mathcal{W}''; \nu \rrbracket$  must contain at least two linearly independent columns. To see it is so, notice that:

$$\llbracket \mathcal{W}''; \nu \rrbracket_{:,2:,4} = \begin{bmatrix} 1 & ? & ? & 0 \\ 0 & ? & ? & 1 \end{bmatrix},$$

where recall that ? stands for an unobserved entry.

- If  $\nu$  contains  $N-1$  but not  $N-2$  and  $N$ , then  $\text{rank } \llbracket \mathcal{W}; \nu \rrbracket = 1$  since:

$$\llbracket \mathcal{W}; \nu \rrbracket_{:,2:,4} = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix},$$

and all other entries of  $\llbracket \mathcal{W}; \nu \rrbracket$  hold zero. In this case  $\llbracket \mathcal{W}''; \nu \rrbracket = R''_{\nu} = 0$ , which means that  $\mathcal{W}''$  is the zero tensor, in contradiction to it fitting the (non-zero) observed entries from Equation (36).

- If  $\nu$  contains  $N-2$  but not  $N-1$  and  $N$ , then  $\text{rank } \llbracket \mathcal{W}; \nu \rrbracket = 2$  since:

$$\llbracket \mathcal{W}; \nu \rrbracket_{:,2:,4} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix},$$

and all other entries of  $\llbracket \mathcal{W}; \nu \rrbracket$  hold zero. In this case  $\llbracket \mathcal{W}''; \nu \rrbracket = R''_{\nu} < 2$ . Noticing that  $\llbracket \mathcal{W}''; \{N-2\} \rrbracket_{:,2:,4} = \llbracket \mathcal{W}''; \nu \rrbracket_{:,2:,4}$ , and that entries of  $\llbracket \mathcal{W}''; \{N-2\} \rrbracket$  outside its top 2-by-4 submatrix hold zero, we get that  $\llbracket \mathcal{W}''; \{N-2\} \rrbracket = \llbracket \mathcal{W}''; \nu \rrbracket < 2$ . Furthermore, from the assumption that  $(R''_{\nu})_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank } \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  and the previous three cases, we know that  $R''_{\{n\}} = \llbracket \mathcal{W}; \{n\} \rrbracket = 1$  for all  $n \in [N-3]$ ,  $R''_{\{N\}} = \llbracket \mathcal{W}; \{N\} \rrbracket = 2$ , and  $R''_{\{N-1\}} = \llbracket \mathcal{W}; \{N-1\} \rrbracket = 1$ . Any tensor  $\mathcal{V} \in \mathbb{R}^{D_1, \dots, D_N}$  that satisfies  $\text{rank } \llbracket \mathcal{V}; \{n\} \rrbracket \leq R''_{\{n\}} \in \mathbb{N}$  for all  $n \in [N]$  can be represented as:

$$\mathcal{V} = \sum_{r_1=1}^{R_{\{1\}}} \dots \sum_{r_N=1}^{R_{\{N\}}} \mathcal{C}_{r_1, \dots, r_N} \otimes_{n=1}^N \mathbf{U}_{:,r_n}^{(n)},$$

where  $\mathcal{C} \in \mathbb{R}^{R_{\{1\}}, \dots, R_{\{N\}}}$  and  $(\mathbf{U}^{(n)} \in \mathbb{R}^{D_n, R_{\{n\}}})_{n=1}^N$  (see, *e.g.*, Section 4 in Kolda & Bader (2009)). Thus, there exist  $c_1, c_2 \in \mathbb{R}$ ,  $(\mathbf{U}^{(n)} \in \mathbb{R}^{D_n, 1})_{n=1}^{N-1}$ , and  $\mathbf{U}^{(N)} \in \mathbb{R}^{D_N, 2}$  such that:

$$\mathcal{W}'' = c_1 \cdot (\otimes_{n=1}^{N-1} \mathbf{U}_{:,1}^{(n)}) \otimes \mathbf{U}_{:,1}^{(N)} + c_2 \cdot (\otimes_{n=1}^{N-1} \mathbf{U}_{:,1}^{(n)}) \otimes \mathbf{U}_{:,2}^{(N)}.$$

By multilinearity of the tensor product, we may write:  $\mathcal{W}'' = (\otimes_{n=1}^{N-1} \mathbf{U}_{:,1}^{(n)}) \otimes (c_1 \cdot \mathbf{U}_{:,1}^{(N)} + c_2 \cdot \mathbf{U}_{:,2}^{(N)})$ , and so  $\mathcal{W}''$  has tensor rank one (it can be represented as a single non-zero tensor product between vectors). Since the tensor rank of a given tensor upper bounds the ranks of its matricizations (Remark 6.21 in Hackbusch (2012)),  $R''_{\{n\}} = \text{rank } \llbracket \mathcal{W}''; \{n\} \rrbracket = 1$  for all  $n \in [N]$  (the matrix ranks of these matricizations cannot be zero as  $\mathcal{W}''$  is not the zero tensor). Hence, we have arrived at a contradiction —  $2 = R''_{\{N\}} \leq 1$ .



- Contradictions in the remaining cases, where  $\nu$  contains  $N - 2, N - 1$ , and  $N$ , or any two of them, readily follow from the previous cases due to the fact that  $\llbracket \mathcal{V}; \nu \rrbracket = \llbracket \mathcal{V}; [N] \setminus \nu \rrbracket^\top$  for any tensor  $\mathcal{V} \in \mathbb{R}^{D_1, \dots, D_N}$ , and that the matrix rank of a matrix is equal to the matrix rank of its transpose. In particular, for any such  $\nu$ , it holds that  $\text{rank} \llbracket \mathcal{W}''; \nu \rrbracket = \text{rank} \llbracket \mathcal{W}''; [N] \setminus \nu \rrbracket$  and  $\text{rank} \llbracket \mathcal{W}; \nu \rrbracket = \text{rank} \llbracket \mathcal{W}; [N] \setminus \nu \rrbracket$ . Therefore, if  $\text{rank} \llbracket \mathcal{W}''; \nu \rrbracket = R_\nu'' < \text{rank} \llbracket \mathcal{W}; \nu \rrbracket$ , then  $\text{rank} \llbracket \mathcal{W}''; [N] \setminus \nu \rrbracket < \text{rank} \llbracket \mathcal{W}; [N] \setminus \nu \rrbracket$ . Since  $\nu$  contains  $N - 2, N - 1$ , and  $N$ , or any two of them, its complement  $[N] \setminus \nu$  contains none or just one of them. Each of these scenarios was already covered in previous cases, which imply that  $\text{rank} \llbracket \mathcal{W}''; [N] \setminus \nu \rrbracket < \text{rank} \llbracket \mathcal{W}; [N] \setminus \nu \rrbracket$  entails a contradiction.

In all cases, we have established that the existence of  $(R_\nu'')_{\nu \in \mathcal{T} \setminus \{[N]\}} \in \mathcal{R}_{\mathcal{T}}$ , different from  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  and  $(\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$ , satisfying  $(R_\nu'')_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  leads to a contradiction. The claim for  $\mathcal{W}'$ , i.e. that there exists no such  $(R_\nu'')_{\nu \in \mathcal{T} \setminus \{[N]\}}$  satisfying  $(R_\nu'')_{\nu \in \mathcal{T} \setminus \{[N]\}} \leq (\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$ , is proven analogously. Combined with the previous part of the proof, in which we established that neither  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  nor  $(\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  is smaller than or equal to the other, we conclude that  $(\text{rank} \llbracket \mathcal{W}; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  and  $(\text{rank} \llbracket \mathcal{W}'; \nu \rrbracket)_{\nu \in \mathcal{T} \setminus \{[N]\}}$  are two different minimal elements of  $\mathcal{R}_{\mathcal{T}}$ .  $\square$

### H.8 Proof of Proposition 3

For  $l \in [L]$ , the output of the  $l$ 'th convolutional layer at index  $n \in [N/P^{l-1}]$  and channel  $r \in [R_l]$  depends solely on inputs  $\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})}$ . Hence, we denote it by  $\text{conv}_{l,n,r}(\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})})$ . We may view the output linear layer as a  $1 \times 1$  convolutional layer with a single output channel. Accordingly, let  $\text{conv}_{L+1,1,1}(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) := f_\theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)})$  and  $\mathcal{W}^{(L+1,1,1)} := \mathcal{W}_H$ .

We show by induction over the layer  $l \in [L+1]$  that for any  $n \in [N/P^{l-1}]$  and  $r \in [R_l]$ :

$$\text{conv}_{l,n,r}(\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})}) = \left\langle \otimes_{p=(n-1) \cdot P^{l-1} + 1}^{n \cdot P^{l-1}} \mathbf{x}^{(p)}, \mathcal{W}^{(l,n,r)} \right\rangle. \quad (37)$$

For  $l = 1$ , let  $n \in [N]$  and  $r \in [R_1]$ . From the definition of  $\mathcal{W}^{(1,n,r)}$  (Equation (10)) we can see that:

$$\text{conv}_{1,n,r}(\mathbf{x}^{(n)}) = \left\langle \mathbf{x}^{(n)}, \mathbf{W}_{:,r}^{(1,n)} \right\rangle = \left\langle \mathbf{x}^{(n)}, \mathcal{W}^{(1,n,r)} \right\rangle.$$

Now, assuming that the inductive claim holds for  $l-1 \geq 1$ , we prove that it holds for  $l$ . Fix some  $n \in [N/P^{l-1}]$  and  $r \in [R_l]$ . The  $l$ 'th convolutional layer is applied to the output of the  $l-1$ 'th hidden layer, denoted  $(\mathbf{h}^{(l-1,1)}, \dots, \mathbf{h}^{(l-1, N/P^{l-1})}) \in \mathbb{R}^{R_{l-1}} \times \dots \times \mathbb{R}^{R_{l-1}}$ . Each  $\mathbf{h}^{(l-1,n)}$ , for  $n \in [N/P^{l-1}]$ , is a result of the product pooling operation (with window size  $P$ ) applied to the output of the  $l-1$ 'th convolutional layer. Thus:

$$\begin{aligned} \text{conv}_{l,n,r}(\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})}) &= \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \cdot \mathbf{h}_{r'}^{(l-1,n)} \\ &= \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \cdot \prod_{p=(n-1) \cdot P + 1}^{n \cdot P} \text{conv}_{l-1,p,r'}(\mathbf{x}^{((p-1) \cdot P^{l-2} + 1)}, \dots, \mathbf{x}^{(p \cdot P^{l-2})}). \end{aligned}$$

The inductive assumption for  $l-1$  then implies that:

$$\text{conv}_{l,n,r}(\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})}) = \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \cdot \prod_{p=(n-1) \cdot P + 1}^{n \cdot P} \left\langle \otimes_{n'=(p-1) \cdot P^{l-2} + 1}^{p \cdot P^{l-2}} \mathbf{x}^{(n')}, \mathcal{W}^{(l-1,p,r')} \right\rangle.$$

For any tensors  $\mathcal{A}, \mathcal{A}', \mathcal{B}, \mathcal{B}'$  such that  $\mathcal{A}$  is of the same dimensions as  $\mathcal{A}'$  and  $\mathcal{B}$  is of the same dimensions as  $\mathcal{B}'$ , it holds that  $\langle \mathcal{A} \otimes \mathcal{B}, \mathcal{A}' \otimes \mathcal{B}' \rangle = \langle \mathcal{A}, \mathcal{A}' \rangle \cdot \langle \mathcal{B}, \mathcal{B}' \rangle$ . We may therefore write:

$$\begin{aligned} &\text{conv}_{l,n,r}(\mathbf{x}^{((n-1) \cdot P^{l-1} + 1)}, \dots, \mathbf{x}^{(n \cdot P^{l-1})}) \\ &= \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \cdot \left\langle \otimes_{p=(n-1) \cdot P + 1}^{n \cdot P} \left( \otimes_{n'=(p-1) \cdot P^{l-2} + 1}^{p \cdot P^{l-2}} \mathbf{x}^{(n')} \right), \otimes_{p=(n-1) \cdot P + 1}^{n \cdot P} \mathcal{W}^{(l-1,p,r')} \right\rangle \\ &= \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \cdot \left\langle \otimes_{p=(n-1) \cdot P^{l-1} + 1}^{n \cdot P^{l-1}} \mathbf{x}^{(p)}, \otimes_{p=(n-1) \cdot P + 1}^{n \cdot P} \mathcal{W}^{(l-1,p,r')} \right\rangle \\ &= \left\langle \otimes_{p=(n-1) \cdot P^{l-1} + 1}^{n \cdot P^{l-1}} \mathbf{x}^{(p)}, \sum_{r'=1}^{R_{l-1}} \mathbf{W}_{r',r}^{(l,n)} \left[ \otimes_{p=(n-1) \cdot P + 1}^{n \cdot P} \mathcal{W}^{(l-1,p,r')} \right] \right\rangle. \end{aligned}$$

Noticing that  $\mathcal{W}^{(l,n,r)} = \sum_{r'=1}^{R_l-1} \mathbf{W}_{r',r}^{(l,n)} \left[ \otimes_{p=(n-1) \cdot P+1}^{n \cdot P} \mathcal{W}^{(l-1,p,r')} \right]$  (Equation (10)) establishes Equation (37).

Applying the inductive claim for  $l = L + 1, n = 1$ , and  $r = 1$ , while recalling that  $L = \log_P N$ , yields:

$$f_\theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \text{conv}_{L+1,1,1}(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}^{(L+1,1,1)} \right\rangle = \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \right\rangle.$$

□

## H.9 Proof of Theorem 2

Let  $t \geq 0$  be a time at which  $\sigma_H^{(\nu,r)}(t) := \|\otimes_{\mathbf{w} \in \text{LC}(\nu,r)} \mathbf{w}(t)\| = \prod_{\mathbf{w} \in \text{LC}(\nu,r)} \|\mathbf{w}(t)\| > 0$ . We differentiate  $\sigma_H^{(\nu,r)}(t)^2$  with respect to time as done in the proof of Theorem 1 (Appendix H.5). From the product rule and Lemma 13 we get that:

$$\frac{d}{dt} \sigma_H^{(\nu,r)}(t)^2 = 2\sigma_H^{(\nu,r)}(t) \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \prod_{\mathbf{w}' \in \text{LC}(\nu,r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2.$$

Since according to the chain rule  $\frac{d}{dt} \sigma_H^{(\nu,r)}(t)^2 = 2\sigma_H^{(\nu,r)}(t) \cdot \frac{d}{dt} \sigma_H^{(\nu,r)}(t)$ , the equation above leads to:

$$\frac{d}{dt} \sigma_H^{(\nu,r)}(t) = \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \prod_{\mathbf{w}' \in \text{LC}(\nu,r) \setminus \{\mathbf{w}\}} \|\mathbf{w}'(t)\|^2. \quad (38)$$

By Lemma 2, the unbalancedness magnitude is constant through time, and so it remains equal to  $\epsilon$  — its value at initialization. Hence, for any  $\mathbf{w} \in \text{LC}(\nu, r)$ :

$$\|\mathbf{w}(t)\|^2 \leq \min_{\mathbf{w}' \in \text{LC}(\nu,r)} \|\mathbf{w}'(t)\|^2 + \epsilon = \left( \min_{\mathbf{w}' \in \text{LC}(\nu,r)} \|\mathbf{w}'(t)\| \right)^{L_\nu \cdot \frac{2}{L_\nu}} + \epsilon \leq \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon. \quad (39)$$

If  $\left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \geq 0$ , applying the inequality above to each  $\|\mathbf{w}'(t)\|^2$  in Equation (38) yields the upper bound from Equation (11):

$$\begin{aligned} \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &\leq \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \prod_{\mathbf{w}' \in \text{LC}(\nu,r) \setminus \{\mathbf{w}\}} \left( \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon \right) \\ &= \left( \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon \right)^{L_\nu-1} \cdot L_\nu \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle. \end{aligned}$$

To prove the lower bound from Equation (11), we multiply and divide each summand on the right hand side of Equation (38) by the corresponding  $\|\mathbf{w}(t)\|^2$  (non-zero because  $\sigma_H^{(\nu,r)}(t) > 0$ ), *i.e.*:

$$\begin{aligned} \frac{d}{dt} \sigma_H^{(\nu,r)}(t) &= \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \|\mathbf{w}(t)\|^{-2} \cdot \prod_{\mathbf{w}' \in \text{LC}(\nu,r)} \|\mathbf{w}'(t)\|^2 \\ &= \left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle \sigma_H^{(\nu,r)}(t) \cdot \sum_{\mathbf{w} \in \text{LC}(\nu,r)} \|\mathbf{w}(t)\|^{-2}. \end{aligned}$$

By Equation (39) we know that  $\|\mathbf{w}(t)\|^{-2} \geq \left( \sigma_H^{(\nu,r)}(t)^{\frac{2}{L_\nu}} + \epsilon \right)^{-1}$ . Thus, applying this inequality to the equation above establishes the desired lower bound.

If  $\left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle < 0$ , the upper and lower bounds in Equation (12) readily follow by similar derivations, where the difference in the direction of inequalities is due to the negativity of  $\left\langle -\nabla \mathcal{L}_H(\mathcal{W}_H(t)), \mathcal{C}_H^{(\nu,r)}(t) \right\rangle$ . □

## H.10 Proof of Proposition 4

We partition the proof into two parts: the first shows that  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \geq \text{sep}(f_\Theta; I)$ , and the second establishes the converse.

**Proof of lower bound** ( $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \geq \text{sep}(f_\Theta; I)$ ) Denote  $R := \text{rank} \llbracket \mathcal{W}_H; I \rrbracket$ , and assume without loss of generality that  $I = \llbracket I \rrbracket$ . Since  $\llbracket \mathcal{W}_H; I \rrbracket$  is a rank  $R$  matrix, there exist  $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(R)} \in \mathbb{R}^{\prod_{n=1}^{|I|} D_n}$  and  $\bar{\mathbf{v}}^{(1)}, \dots, \bar{\mathbf{v}}^{(R)} \in \mathbb{R}^{\prod_{n=|I|+1}^N D_n}$  such that:

$$\llbracket \mathcal{W}_H; I \rrbracket = \sum_{r=1}^R \mathbf{v}^{(r)} (\bar{\mathbf{v}}^{(r)})^\top.$$

For each  $r \in [R]$ , let  $\mathcal{V}^{(r)} \in \mathbb{R}^{D_1, \dots, D_{|I|}}$  be the tensor whose arrangement as a column vector is equal to  $\mathbf{v}^{(r)}$ , i.e.  $\llbracket \mathcal{V}^{(r)}; I \rrbracket = \mathbf{v}^{(r)}$ . Similarly, for every  $r \in [R]$  let  $\bar{\mathcal{V}}^{(r)} \in \mathbb{R}^{D_{|I|+1}, \dots, D_N}$  be the tensor whose arrangement as a row vector is equal to  $(\bar{\mathbf{v}}^{(r)})^\top$ , i.e.  $\llbracket \bar{\mathcal{V}}^{(r)}; \emptyset \rrbracket = (\bar{\mathbf{v}}^{(r)})^\top$ . Then:

$$\begin{aligned} \llbracket \mathcal{W}_H; I \rrbracket &= \sum_{r=1}^R \mathbf{v}^{(r)} (\bar{\mathbf{v}}^{(r)})^\top \\ &= \sum_{r=1}^R \llbracket \mathcal{V}^{(r)}; I \rrbracket \odot \llbracket \bar{\mathcal{V}}^{(r)}; \emptyset \rrbracket \\ &= \sum_{r=1}^R \llbracket \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)}; I \rrbracket \\ &= \llbracket \sum_{r=1}^R \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)}; I \rrbracket, \end{aligned}$$

where the third equality makes use of Lemma 3, and the last equality is by linearity of the matricization operator. Since matricizations merely reorder the entries of tensors, the equation above implies that  $\mathcal{W}_H = \sum_{r=1}^R \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)}$ . We therefore have that:

$$\begin{aligned} f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) &= \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{W}_H \right\rangle \\ &= \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \sum_{r=1}^R \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)} \right\rangle \\ &= \sum_{r=1}^R \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)} \right\rangle. \end{aligned}$$

For any  $\mathcal{A}, \mathcal{A}' \in \mathbb{R}^{D_1, \dots, D_{|I|}}$  and  $\mathcal{B}, \mathcal{B}' \in \mathbb{R}^{D_{|I|+1}, \dots, D_N}$  it holds that  $\langle \mathcal{A} \otimes \mathcal{B}, \mathcal{A}' \otimes \mathcal{B}' \rangle = \langle \mathcal{A}, \mathcal{A}' \rangle \cdot \langle \mathcal{B}, \mathcal{B}' \rangle$ . Thus:

$$f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \sum_{r=1}^R \left\langle \otimes_{n=1}^N \mathbf{x}^{(n)}, \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)} \right\rangle = \sum_{r=1}^R \left\langle \otimes_{n=1}^{|I|} \mathbf{x}^{(n)}, \mathcal{V}^{(r)} \right\rangle \cdot \left\langle \otimes_{n=|I|+1}^N \mathbf{x}^{(n)}, \bar{\mathcal{V}}^{(r)} \right\rangle.$$

By defining  $g_r : \times_{n=1}^{|I|} \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  and  $\bar{g}_r : \times_{n=|I|+1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$ , for  $r \in [R]$ , as:

$$g_r(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(|I|)}) = \left\langle \otimes_{n=1}^{|I|} \mathbf{x}^{(n)}, \mathcal{V}^{(r)} \right\rangle, \quad \bar{g}_r(\mathbf{x}^{(|I|+1)}, \dots, \mathbf{x}^{(N)}) = \left\langle \otimes_{n=|I|+1}^N \mathbf{x}^{(n)}, \bar{\mathcal{V}}^{(r)} \right\rangle,$$

we arrive at the following representation of  $f_\Theta$  as a sum, where each summand is a product of two functions — one that operates over inputs indexed by  $I$  and another that operates over inputs indexed by  $[N] \setminus I$ :

$$f_\Theta(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \sum_{r=1}^R g_r(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(|I|)}) \cdot \bar{g}_r(\mathbf{x}^{(|I|+1)}, \dots, \mathbf{x}^{(N)}).$$

Since the separation rank of  $f_\Theta$  is the minimal number of summands required to express it in such a manner, we conclude that  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket = R \geq \text{sep}(f_\Theta; I)$ .

**Proof of upper bound** ( $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \leq \text{sep}(f_\Theta; I)$ ) Towards proving the upper bound, we establish the following lemma.

**Lemma 16.** *Given  $f : \times_{n=1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  and any  $(\mathbf{x}^{(1, h_1)} \in \mathbb{R}^{D_1})_{h_1=1}^{H_1}, \dots, (\mathbf{x}^{(N, h_N)} \in \mathbb{R}^{D_N})_{h_N=1}^{H_N}$ , let  $\mathcal{W} \in \mathbb{R}^{H_1, \dots, H_N}$  be the tensor defined by  $\mathcal{W}_{h_1, \dots, h_N} := f(\mathbf{x}^{(1, h_1)}, \dots, \mathbf{x}^{(N, h_N)})$  for all  $(h_1, \dots, h_N) \in [H_1] \times \dots \times [H_N]$ . Then, for any  $I \subset [N]$ :*

$$\text{rank} \llbracket \mathcal{W}; I \rrbracket \leq \text{sep}(f; I).$$

*In words, for any tensor holding the outputs of  $f$  over a grid of inputs, the rank of its matricization according to  $I$  is upper bounded by the separation rank of  $f$  with respect to  $I$ .*

*Proof.* If  $\text{sep}(f; I)$  is  $\infty$  or zero, i.e.  $f$  cannot be represented as a finite sum of separable functions (with respect to  $I$ ) or is identically zero, then the claim is trivial. Otherwise, denote  $R := \text{sep}(f; I)$ , and assume without loss of generality that  $I = [|I|]$ . Let  $g_1, \dots, g_R : \times_{n=1}^{|I|} \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  and  $\bar{g}_1, \dots, \bar{g}_R : \times_{n=|I|+1}^N \mathbb{R}^{D_n} \rightarrow \mathbb{R}$  such that:

$$f(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}) = \sum_{r=1}^R g_r(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(|I|)}) \cdot \bar{g}_r(\mathbf{x}^{(|I|+1)}, \dots, \mathbf{x}^{(N)}). \quad (40)$$

We define  $(\mathcal{V}^{(r)} \in \mathbb{R}^{D_1, \dots, D_{|I|}})_{r=1}^R$  to be the tensors holding the outputs of  $(g_r)_{r=1}^R$  over the grid of inputs

$$(\mathbf{x}^{(1, h_1)})_{h_1=1}^{H_1}, \dots, (\mathbf{x}^{(|I|, h_{|I|})})_{h_{|I|=1}}^{H_{|I|}},$$

*i.e.* for all  $h_1, \dots, h_{|I|} \in [H_1] \times \dots \times [H_{|I|}]$  and  $r \in [R]$  it holds that  $\mathcal{V}_{h_1, \dots, h_{|I|}}^{(r)} = g_r(\mathbf{x}^{(1, h_1)}, \dots, \mathbf{x}^{(|I|, h_{|I|})})$ . Similarly, we let  $(\bar{\mathcal{V}}^{(r)} \in \mathbb{R}^{D_{|I|+1}, \dots, D_N})_{r=1}^R$  be the tensors holding the outputs of  $(\bar{g}_r)_{r=1}^R$  over their respective grid of inputs, *i.e.* for all  $h_{|I|+1}, \dots, h_N \in [H_{|I|+1}] \times \dots \times [H_N]$  and  $r \in [R]$  it holds that  $\bar{\mathcal{V}}_{h_{|I|+1}, \dots, h_N}^{(r)} = \bar{g}_r(\mathbf{x}^{(|I|+1, h_{|I|+1})}, \dots, \mathbf{x}^{(N, h_N)})$ .

By Equation (40) and the definitions of  $\mathcal{W}$ ,  $(\mathcal{V}^{(r)})_{r=1}^R$ , and  $(\bar{\mathcal{V}}^{(r)})_{r=1}^R$ , we have that for any  $h_1, \dots, h_N \in [H_1] \times \dots \times [H_N]$ :

$$\begin{aligned} \mathcal{W}_{h_1, \dots, h_N} &= f(\mathbf{x}^{(1, h_1)}, \dots, \mathbf{x}^{(N, h_N)}) \\ &= \sum_{r=1}^R g_r(\mathbf{x}^{(1, h_1)}, \dots, \mathbf{x}^{(|I|, h_{|I|})}) \cdot \bar{g}_r(\mathbf{x}^{(|I|+1, h_{|I|+1})}, \dots, \mathbf{x}^{(N, h_N)}) \\ &= \sum_{r=1}^R \mathcal{V}_{h_1, \dots, h_{|I|}}^{(r)} \cdot \bar{\mathcal{V}}_{h_{|I|+1}, \dots, h_N}^{(r)}, \end{aligned}$$

which means that  $\mathcal{W} = \sum_{r=1}^R \mathcal{V}^{(r)} \otimes \bar{\mathcal{V}}^{(r)}$ . From the linearity of the matricization operator and Lemma 3 we then get that  $\llbracket \mathcal{W}; I \rrbracket = \sum_{r=1}^R \llbracket \mathcal{V}^{(r)}; I \rrbracket \odot \llbracket \bar{\mathcal{V}}^{(r)}; \emptyset \rrbracket$ . Since  $\llbracket \mathcal{V}^{(r)}; I \rrbracket$  is a column vector and  $\llbracket \bar{\mathcal{V}}^{(r)}; \emptyset \rrbracket$  is a row vector for all  $r \in [R]$ , we have arrived at a representation of  $\llbracket \mathcal{W}; I \rrbracket$  as a sum of  $R$  tensor products between two vectors. A tensor product of two vectors is a rank one matrix, and so, due to the sub-additivity of rank we conclude:  $\text{rank} \llbracket \mathcal{W}; I \rrbracket \leq R = \text{sep}(f; I)$ .  $\square$

Now, consider the grid of inputs defined by the standard bases of  $\mathbb{R}^{D_1}, \dots, \mathbb{R}^{D_N}$ , *i.e.* by:

$$(\mathbf{e}^{(1, d_1)} \in \mathbb{R}^{D_1})_{d_1=1}^{D_1}, \dots, (\mathbf{e}^{(N, d_N)} \in \mathbb{R}^{D_N})_{d_N=1}^{D_N},$$

where  $\mathbf{e}^{(n, d_n)}$  is the vector holding one at its  $d_n$ 'th entry and zero elsewhere for  $n \in [N]$  and  $d_n \in [D_n]$ . With Lemma 16 in hand,  $\text{rank} \llbracket \mathcal{W}_H; I \rrbracket \leq \text{sep}(f_\Theta; I)$  follows by showing that  $\mathcal{W}_H$  is the tensor holding the outputs of  $f_\Theta$  over this grid of inputs. Indeed, for all  $d_1, \dots, d_N \in [D_1] \times \dots \times [D_N]$ :

$$f_\Theta(\mathbf{e}^{(1, d_1)}, \dots, \mathbf{e}^{(N, d_N)}) = \langle \otimes_{n=1}^N \mathbf{e}^{(n, d_n)}, \mathcal{W}_H \rangle = (\mathcal{W}_H)_{d_1, \dots, d_N}.$$

$\square$