



HAL
open science

Implicit Regularization in Deep Tensor Factorization

Paolo Milanese, Hachem Kadri, Stéphane Ayache, Thierry Artières

► **To cite this version:**

Paolo Milanese, Hachem Kadri, Stéphane Ayache, Thierry Artières. Implicit Regularization in Deep Tensor Factorization. International Joint Conference on Neural Networks (IJCNN), Jul 2021, Online, China. ⟨hal-03211964⟩

HAL Id: hal-03211964

<https://hal.science/hal-03211964v1>

Submitted on 1 May 2021

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

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



HAL Authorization

Implicit Regularization in Deep Tensor Factorization

Paolo Milanese¹, Hachem Kadri¹, Stéphane Ayache¹, Thierry Artières^{1,2}

¹Aix Marseille Univ, Université de Toulon, CNRS, LIS

²Ecole Centrale de Marseille

Marseille, France

firstname.lastname@lis-lab.fr

Abstract—Attempts of studying implicit regularization associated to gradient descent (GD) have identified matrix completion as a suitable test-bed. Late findings suggest that this phenomenon cannot be phrased as a minimization-norm problem, implying that a paradigm shift is required and that dynamics has to be taken into account. In the present work we address the more general setup of tensor completion by leveraging two popularized tensor factorization, namely Tucker and TensorTrain (TT). We track relevant quantities such as tensor nuclear norm, *effective rank*, generalized singular values and we introduce deep Tucker and TT unconstrained factorization to deal with the completion task. Experiments on both synthetic and real data show that gradient descent promotes solution with low-rank, and validate the conjecture saying that the phenomenon has to be addressed from a dynamical perspective.

Index Terms—tensor factorization, deep learning, Tucker decomposition, tensor-train, effective rank

I. INTRODUCTION

Deep Neural Networks (DNNs) are currently the state of the art in a variety of tasks ranging from computer vision to speech recognition and bioinformatics. Though their deployment is ubiquitous, a clear theoretical understanding still lies out of our reach; for example, currently employed DNNs are largely overparametrized and yet, they are able to generalize to unseen data [1]–[3]. Although a number of tricks (including small batch processing, early stopping of training and norms constraints over layers weights, just to name a few) are currently used by practitioners to enhance inference capacities, it has become clear that none of these devices is neither a sufficient nor a necessary condition to guarantee generalization of DNNs [2]. Moreover, experiments show that DNNs trained by back propagation with vanilla gradient descent and no explicit regularization do generalize reasonably well [4], suggesting that the optimization algorithm plays a role in driving the loss function towards a local minimum with good generalization properties. Common wisdom says that implicit regularization induced by gradient descent can be rephrased by means of some variational principle: for example, it is known that for linear models, gradient descent finds the solution with minimal Frobenius norm [2], thus one could expect this to hold true in more complex scenarios, possibly by replacing Frobenius norm with other

norms or quasi-norms architecture-depending. A recent line of research addresses this issue within the framework of matrix completion. To put it more clearly, let us consider a ground truth matrix $M \in \mathbb{R}^{n \times m}$ which is partially observed and let $\Omega \subset \{1, \dots, n\} \times \{1, \dots, m\}$ be the set indexing the observed entries. The completion task consists in finding the matrix W minimizing the loss

$$\ell(W) = \frac{1}{2} \sum_{(i,j) \in \Omega} (W_{ij} - M_{ij})^2. \quad (1)$$

The above problem being ill-posed, one needs some additional constraints and the usual strategy consists in seeking for the minimizer with minimal rank, motivated by the observation that in many practical applications the sought matrix only has few degrees of freedom. However, as the rank minimization is a NP-hard problem [5], one is led to minimize the nuclear norm (*i.e.* the sum of the singular values) which is a convex problem and formally equivalent [6]—provided that number of observed entries is large enough and under some “incoherence” assumption—to the rank minimization problem. A natural approach to solve this problem by means of gradient descent would be to factorize W as $W = W_1 W_2$, with $W_1 \in \mathbb{R}^{n \times \hat{r}}$, $W_2 \in \mathbb{R}^{\hat{r} \times m}$, with \hat{r} some estimate of the rank of M , plug it in (1) and run the gradient descent algorithm. Notice that this corresponds to train a 2-layers Linear Neural Network on the dataset $\{(i, j), M_{ij}\}_{(i,j) \in \Omega}$ with loss given by Eq. (1).

However, a somewhat surprising finding is that, even in the case where the shared dimension of W_1, W_2 is unconstrained, namely $\hat{r} = \min\{m, n\}$, the solution of (1) via gradient descent turns out to be a low-rank matrix. This empirical evidence together with a proof in a restrictive setup led Gunasekar *et al.* to conjecture the following:

Conjecture 1 ([7], roughly stated). *With small enough learning rate and initialization, gradient descent on full-dimensional matrix factorization converges to the solution with minimal nuclear norm.*

This conjecture supports the viewpoint of conventional learning theory, claiming that implicit regularization can be rephrased as a norm-minimization problem.

Arora *et al.* [8] pushed the analysis further by considering deep matrix factorization, which consists in parametrizing W

This work was funded by the French national research agency (grant number ANR16-CE23-0006 “Deep in France”).

in Eq. (1) as $W = W_L W_{L-1} \dots W_2 W_1$ for some $L \in \mathbb{N}$ and with $\{W_i\}_{i=1}^L$ to be such that no constraint on the rank is present. Notice that deep matrix factorization is a generalization of shallow matrix factorization setup investigated by [7], which corresponds to choose $L = 2$. On the one hand they empirically showed that depth enhances recovery performances and, on the other hand, they found that the proof supporting the conjecture raised in [7] carries over to the deep setup, meaning that the argument of [7] cannot distinguish the shallow setup ($L = 2$) from the deep one. As this fact contradict experimental evidence and having proved that gradient descent over deep matrix factorization do not minimize any quasi-norm, they were led to challenge Conjecture 1 and to raise the following one:

Conjecture 2 ([8], roughly stated). *Implicit regularization of gradient descent in deep matrix factorization cannot be formulated as a norm-minimization problem; instead one needs to take into account the trajectories of the gradient descent dynamics.*

In this respect they studied the gradient flow and they found out, both theoretically and empirically, that GD promotes sparsity of singular values of W . Lately, [9] exhibited a class of matrix completion problem for which fitting observed entries through gradient descent on deep matrix factorization leads all norms and quasi-norms to grow to infinity whereas other rank surrogates, such as effective rank (cf Section IV), converge towards their minimum. This can be considered as one more clue to the fact that implicit regularization needs to be understood as a dynamical phenomenon. To summarize, not only an understanding of implicit regularization is lacking, but actually the mathematical formalism one needs to address the problem is unclear. Though the conventional viewpoint—claiming that the problem can be phrased as some norm minimization—is becoming more and more untrustworthy and a shift of paradigm seems to be required, one cannot reject it yet.

In the present work we address implicit regularization of gradient descent in the more general setup of tensor completion. Tensors can be considered as multi-dimensional arrays and they have been attracting an increasing interest as many problems require the manipulation of quantities which are, by their very nature, multivariate. Although tensor completion could be viewed as a byproduct of matrix completion and solved by means of techniques developed for the latter, a more careful analysis shows that downgrading tensor completion to a matrix framework—e.g. via slicings or unfoldings—totally discard the multi-way nature of data leading to a degradation of performance. The first issue one has to face when dealing with tensors is that the rank is not uniquely defined; many definitions are possible, among which the canonical rank, Tucker rank and Tensor-Train are the most popular (cf Section II for more details). Once a suitable notion of rank has been established, a first approach to solve tensor completion consists in imposing a low-rank structure to the tensor to be recovered, by leveraging, e.g. the Canonical

Polyadic (CP) [10] or Tucker [11] decomposition; the main drawback of these methods is that an estimate of the rank is required, which is not always available. A second approach relies on the assumption that matrix nuclear norm is a convex surrogate of rank, and tries to redefine an appropriate tensor nuclear norm [12], [13], [14], which generally leads to much reliable results. A somewhat hybrid path consists in imposing norms constraints on factorized matrices instead than on tensor unfolding, in order to reduce computational burden [15] [16]. For a more comprehensive account on the topic of tensor factorization we refer the reader to [17].

The idea of exploring implicit regularization in deep learning with the view of tensor factorization has been suggested and partially studied in the recent work of [9], which only dealt with CP factorization. In the present paper, we go further by addressing other two tensor decompositions, namely Tucker and Tensor-Train, which have proved themselves to be more expressive and manageable than CP. We define deep Tucker and Tensor-Train *unconstrained* factorization and introduce the notion of tensor effective rank to characterize implicit regularization in these models. Our main contributions are:

- (i) We study tensor factorization in the overparameterized regime. More specifically, instead of constraining the rank of the tensor decomposition and reducing the number of parameters, we propose Tucker and Tensor-Train unconstrained factorization that have deep structure and representation flexibility.
- (ii) We generalize the notion of effective rank, introduced for matrices as a continuous surrogate of the matrix rank, to tensors and use it to characterize the implicit regularization in deep tensor factorization. To our knowledge, this is the first proposition on how to extend effective rank to the context of Tucker and Tensor-Train decompositions.
- (iii) We show experimentally that gradient descent on Tucker and Tensor-Train unconstrained factorization tends to produce solutions with low tensor rank. The rank depends on the type of factorization. As far as we are aware, this is the first time such behaviour has been analysed for Tucker and Tensor-Train tensors.
- (iv) Our empirical study provides insights into the behaviour of tensor singular values in deep tensor factorization, and supports the observation that the implicit regularization has to be addressed from a dynamical perspective.

II. TENSOR ALGEBRA

In this section we introduce tensors and the related multi-linear algebra. For a more comprehensive account on this subject we refer the reader to [18]. We will denote vector by boldface lowercase letters, e.g. $v \in \mathbb{R}^{I_1}$ whereas matrices and tensors will be denoted respectively by boldface uppercase and calligraphic letters, e.g. $M \in \mathbb{R}^{I_1 \times I_2}$ and $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$. The (i_1, i_2, \dots, i_N) -th entry of $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ will be denoted by $\mathcal{T}_{i_1, i_2, \dots, i_N}$ where $i_n = 1, 2, \dots, I_n$, for all

$n = [1, \dots, N]$. Given two tensors $\mathcal{T}, \mathcal{S} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ their scalar product writes

$$\langle \mathcal{T}, \mathcal{S} \rangle = \sum_{i_1} \sum_{i_2} \dots \sum_{i_N} \mathcal{S}_{i_1, i_2, \dots, i_N} \mathcal{T}_{i_1, i_2, \dots, i_N},$$

and the Frobenius norm of \mathcal{T} is defined as $\|\mathcal{T}\|_F = \sqrt{\langle \mathcal{T}, \mathcal{T} \rangle}$. Given a 3-rd order tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, its slices are the matrices obtained by fixing one index: as an example, $\mathbf{T}[i_2] \in \mathbb{R}^{I_1 \times I_3}$ is the slice obtained by fixing $I_2 = i_2$.

Given N vectors $\mathbf{v}_1 \in \mathbb{R}^{I_1}, \mathbf{v}_2 \in \mathbb{R}^{I_2}, \dots, \mathbf{v}_N \in \mathbb{R}^{I_N}$, their outer product is the tensor whose (i_1, \dots, i_N) -th entry writes $(\mathbf{v}_1 \circ \mathbf{v}_2 \circ \dots \circ \mathbf{v}_N)_{i_1, i_2, \dots, i_N} = (\mathbf{v}_1)_{i_1} (\mathbf{v}_2)_{i_2} \dots (\mathbf{v}_N)_{i_N}$ for all $i_n \in [1, \dots, I_n], n \in [1, \dots, N]$. An N -th order tensor \mathcal{T} is called a rank-1 tensor if it can be written as the outer product of N vectors, i.e. $\mathcal{T} = \mathbf{v}_1 \circ \mathbf{v}_2 \circ \dots \circ \mathbf{v}_N$. This lead us to the first definition of tensor rank.

Definition II.1. *The canonical rank of an arbitrary N -th order tensor \mathcal{T} is the minimal number of rank-1 tensors that yield \mathcal{T} in a linear combination.*

The canonical rank is the most straightforward extension of the matrix rank to the tensor realm; however its computation is an NP-hard problem [19], thus alternative and more maneuverable definitions of the tensor rank have been established. To define them let us define the (Tucker) n -unfolding of a tensor as follows: Given an N -th order tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, the matrix $\mathbf{T}_{(n)} \in \mathbb{R}^{I_n \times (I_{n+1} I_{n+2} \dots I_N I_1 I_2 \dots I_{n-1})}$ is called the (Tucker) unfolding of \mathcal{T} along the n -th mode. This bring us to define the n -rank and the Tucker rank.

Definition II.2. *The tensor n -rank of \mathcal{T} is defined as the rank of $\mathbf{T}_{(n)}$. The Tucker rank, also called multi-linear rank, is defined as the N -tuple whose i -th entry is the i -th rank of \mathcal{T} .*

Different n -ranks of a tensor are not necessarily the same and generally speaking they do not coincide with the canonical rank, though it can be found that any n -rank is bounded by the canonical rank. Tucker rank is intimately related to the higher-order singular value decomposition, on which our analysis relies. The first step to take is to define a proper way to multiply a tensor and a matrix.

Definition II.3. *The n -mode product of a tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ with a matrix $\mathbf{U} \in \mathbb{R}^{J \times I_n}$ is a tensor of size $(I_1 \times I_2 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N)$ denoted by $\mathcal{T} \times_n \mathbf{U}$ and its entries are defined as*

$$(\mathcal{T} \times_n \mathbf{U})_{i_1 \dots i_{n-1} j i_{n+1} \dots i_N} = \sum_{i_n=1}^{I_n} \mathcal{T}_{i_1 \dots i_{n-1} i_n i_{n+1} \dots i_N} U_{j i_n}.$$

The n -mode product satisfies to the following property: given $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ and two matrices $\mathbf{F} \in \mathbb{R}^{J \times I_n}, \mathbf{G} \in \mathbb{R}^{K \times J}$, one has

$$(\mathcal{T} \times_n \mathbf{F}) \times_n \mathbf{G} = \mathcal{T} \times_n (\mathbf{G} \cdot \mathbf{F}),$$

which generalizes to more than two matrices. We are in shape to write a Singular Value Decomposition (SVD) applying to

tensors, named Higher-Order Singular Value Decomposition (HOSVD).

Theorem 1 (HOSVD, [20]). *Any tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ can be written as the product*

$$\mathcal{T} = \mathcal{C} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)},$$

where

- 1) $\mathbf{U}^{(n)}$ is a unitary $(I_n \times I_n)$ -matrix, usually denoted as factor matrix;
- 2) \mathcal{C} is a $(I_1 \times I_2 \times \dots \times I_N)$ -tensor, usually referred to as core tensor; of which the sub-tensors $\mathcal{C}_{i_n=\alpha}$ obtained by fixing the n -th index to α are such that
 - (i) they are mutually orthogonal: given two sub-tensors $\mathcal{C}_{i_n=\alpha}$ and $\mathcal{C}_{i_n=\beta}$, they are orthogonal for all possible values of n, α, β subject to $\alpha \neq \beta$:

$$\langle \mathcal{C}_{i_n=\alpha}, \mathcal{C}_{i_n=\beta} \rangle = 0,$$

- (ii) they are ordered:

$$\|\mathcal{C}_{i_n=1}\|_F \geq \|\mathcal{C}_{i_n=2}\|_F \dots \geq \|\mathcal{C}_{i_n=I_n}\|_F \geq 0.$$

The analogy with matrix SVD is straightforward: the Frobenius norm of subtensors $\|\mathcal{C}_{i_n=i}\|_F := \sigma_i^n$ are the n -mode singular values and the column vectors of the matrices $\mathbf{U}^{(n)}$ are the n -mode singular vectors. As in the matrix case, where the number of non-zero singular values controls the rank, in the higher-order setup we have that, if R_n is equal to the highest index for which $\|\mathcal{C}_{i_n=R_n}\| > 0$ then one has that the n -rank of \mathcal{T} is equal to R_n . Tucker decomposition [21] with rank (R_1, R_2, \dots, R_N) can be obtained from HOSVD by considering only the first R_n column vectors of the matrix $\mathbf{U}^{(n)}$ for each $n \in [1, \dots, N]$.

Tensor-Train (TT) decomposition [22] writes \mathcal{T} as a product of third order tensors $\mathcal{T}_k \in \mathbb{R}^{R_{k-1} \times I_k \times R_k}$ called cores. The TT-rank of the decomposition is given by the tuple (R_0, R_1, \dots, R_N) where $R_0 = R_N = 1$. By denoting $\mathcal{T}_k[i_k]$ the slice obtained from \mathcal{T}_k by taking $I_k = i_k$, TT decomposition writes

$$\mathcal{T}_{i_1, i_2, \dots, i_N} = \mathbf{T}_1[i_1] \mathbf{T}_2[i_2] \dots \mathbf{T}_N[i_N]. \quad (2)$$

TT algorithm is naturally associated with the following unfolding: given an N -th order tensor $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, the matrix $\mathbf{T}_{[n]} \in \mathbb{R}^{(I_1 I_2 \dots I_n) \times (I_{n+1} I_{n+2} \dots I_N)}$ is called the TT unfolding of \mathcal{T} along the n -th mode. In order to define TT generalized singular values, we recall TT-SVD algorithm ([22], Algorithm 1). Given $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ and a TT-rank (R_0, R_1, \dots, R_N) , it starts by considering the unfolding $\mathbf{T}_{[1]}$ and its partial SVD denoted by $U_n S_n V_n^T$ truncated to the first R_1 singular values. Then, the first core \mathcal{T}_1 will be given by reshaping U_1 as a 3-rd order tensor with size (R_0, I_1, R_1) and, at the following iteration an SVD truncated to the first R_i singular values is applied to $S_1 V_1^T$. The algorithm stops after $N - 1$ iterations. For any $n \in [1, \dots, N - 1]$, we say that the diagonal entries of S_n are the n -mode singular values associated with TT decomposition.

III. DEEP TENSOR FACTORIZATION

In this section, we introduce Tucker and TT unconstrained factorization for deep tensor completion. Tensor completion by mean of tensor factorization is the generalization of Eq. (1). Let $\mathcal{M} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ the ground truth tensor we want to recover and let us denote by $\Omega \subset \{1, 2, \dots, I_1\} \times \{1, 2, \dots, I_2\} \times \dots \times \{1, 2, \dots, I_N\}$ the set which indexes the observed entries. We want to minimize the loss $\ell : \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N} \rightarrow \mathbb{R}_{\geq 0}$ defined as

$$\ell(\mathcal{W}) = \frac{1}{2} \sum_{(i_1, \dots, i_N) \in \Omega} (\mathcal{W}_{i_1, \dots, i_N} - \mathcal{M}_{i_1, \dots, i_N})^2. \quad (3)$$

To do so, we adopt a deep learning perspective and propose Tucker and TT “overparameterized” factorization where no assumptions are made with respect to the rank. This is crucial in order to analyze the effect of the implicit regularization on deep tensor factorization.

a) *Deep Tucker Unconstrained Factorization:* We write \mathcal{W} as the product of a core tensor times a set of factor matrices and inspired by deep matrix factorization [8], we include depth by writing

$$\begin{aligned} \mathcal{W} = \mathcal{D} \times_1 & \left(\mathbf{V}_1^{(1)} \cdot \mathbf{V}_2^{(1)} \dots \mathbf{V}_{k_1}^{(1)} \right) \\ & \times_2 \left(\mathbf{V}_1^{(2)} \cdot \mathbf{V}_2^{(2)} \dots \mathbf{V}_{k_2}^{(2)} \right) \\ & \times_3 \dots \times_N \left(\mathbf{V}_1^{(N)} \cdot \mathbf{V}_2^{(N)} \dots \mathbf{V}_{k_N}^{(N)} \right), \end{aligned} \quad (4)$$

where \mathcal{D} is a $(I_1 \times I_2 \times \dots \times I_N)$ -th tensor and the matrices $\mathbf{V}_i^{(n)}$ are in $I_n \times I_n$ for all $n \in \{1, \dots, N\}$ and $i \in \{1, \dots, k_n\}$ and we refer to Eq. (4) as Tucker Unconstrained Factorization (UF). A Tucker UF with k_n factor matrices along mode n , $n \in [1, \dots, N]$ will be said to have depth (k_1, k_2, \dots, k_N) and we adopt the convention that if $k_n = 0$ for some n , then the n -mode multiplication writes $\mathcal{D} \times_n \mathbb{1}_{I_n}$, where $\mathbb{1}_{I_n}$ is the $I_n \times I_n$ identity matrix. Thus the depth- $(0, 0, \dots, 0)$ Tucker UF writes

$$\mathcal{W} = \mathcal{D} \times_1 \mathbb{1}_{I_1} \times_2 \mathbb{1}_{I_2} \dots \times_N \mathbb{1}_{I_N} = \mathcal{D}.$$

Notice that Eq. (4) has no explicit constraints on any of the n -rank and that it is the most natural way to carry over the deep matrix factorization setup addressed in [8] to Tucker low-ranked tensors.

b) *Deep TT Unconstrained Factorization:* As far as Tensor-Train decomposition is concerned, in order not to introduce any explicit constraint on the TT-rank (recall the TT-SVD algorithm in the previous section) we just need to keep the entire SVD of each $\mathbf{T}_{[n]}$; thus the unconstrained Tensor Train factorization of \mathcal{W} writes:

$$\mathcal{W}_{i_1, i_2, \dots, i_N} = \mathbf{W}_1[i_1] \mathbf{W}_2[i_2] \dots \mathbf{W}_N[i_N], \quad (5)$$

such that, for each $n \in [1, \dots, N-1]$, $\mathcal{W}_n \in \mathbb{R}^{R_{n-1} \times I_n \times R_n}$, $R_n = \min \{I_1 I_2 \dots I_n, I_{n+1} \dots I_N\}$ and with $R_0 = R_N = 1$. We denote Eq. (5) as Tensor-Train UF. The notion of depth here is intrinsic to the TT decomposition, since the number of core tensors is fixed according

to the order of the original tensor. It is, however, possible to include other notions of depth in the parameterization given by Eq. (5). At least, two options can be explored: the first one is straightforward from Tucker UF and consists in factorizing each TT core by means of Eq. (4). The second strategy consists in writing each core \mathcal{W}_n , $n \in [1, \dots, N]$ by mean of a TT decomposition. We defer this to future work.

IV. TENSOR EFFECTIVE RANKS

When performing tensor completion, brute performance of the completion task is grasped by the reconstruction error, defined as the normalized Frobenius distance $\|\mathcal{W}^* - \mathcal{M}\|_F / \|\mathcal{M}\|_F$ between the solution \mathcal{W}^* of Eq. (3) and the ground-truth tensor \mathcal{M} .

As far as Tucker factorization is concerned, a suitable extension of the nuclear norm to the tensor realm has been proposed by [12], and writes

$$\|\mathcal{W}\|_*^{tuc} = \sum_{n=1}^N \alpha_n \|\mathbf{W}_{(n)}\|_* \quad \text{s.t.} \quad \sum_{n=1}^N \alpha_n = 1.$$

In the aforementioned paper, authors also established several algorithms based on the minimization of $\|\cdot\|_*^{tuc}$, the most effective of which is named High Accuracy Low Rank Tensor Completion (HALRTC) and takes advantage of Alternating Direction Method of Multipliers (ADMM) to perform the task. On the other hand, when dealing with TT factorization, [14] introduced a convenient notion of nuclear norm, which reads

$$\|\mathcal{W}\|_*^{TT} = \sum_{n=1}^{N-1} \alpha_n \|\mathbf{W}_{[n]}\|_* \quad \text{s.t.} \quad \sum_{n=1}^{N-1} \alpha_n = 1,$$

as well as an algorithm, named TT-SILRTC, to deal with the minimization of $\|\cdot\|_*^{TT}$ by means of Block Coordinate Descent method. These norms arise as convex approximations of the rank functions which are much easier to handle analytically and numerically.

In the matrix case, the notion of effective rank has also been proposed in [23] as a continuous surrogate of the matrix rank and it proved itself, as remarked in [8], to provide useful insights for the deep matrix factorization setup. The effective rank of a matrix is computed from the entropy of its singular value distribution. In this respect we propose a new definition of effective rank which allow us to cover Tucker and TT ranks. We start by recalling the definition of Shannon entropy: given a discrete probability distribution $\{p_i\}_{i=1}^n$, the Shannon entropy $H(\{p_i\}_{i=1}^n)$ is defined as $H(\{p_i\}_{i=1}^n) = -\sum_i p_i \cdot \ln p_i$, with the convention that $0 \cdot \ln 0 = 0$.

Definition IV.1 (Tucker effective rank). *Let $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ be a N -th order tensor and let $\mathcal{C} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \dots \times_N \mathbf{U}^{(N)}$ be the associated HOSVD. For $n \in [1, \dots, N]$, consider the n -mode singular values $\{\sigma_i^n\}_{i=1}^{I_n}$ and the associate probability distribution $\left\{ \rho_i^n = \sigma_i^n(\mathcal{T}) / \sum_{j=1}^{I_n} \sigma_j^n \right\}_{i=1}^{I_n}$. The Tucker effective rank along the n -th mode of \mathcal{T} is defined as $erank_n^{tuc}(\mathcal{T}) = \exp \{H(\rho_1^n(\mathcal{T}), \rho_2^n(\mathcal{T}), \dots, \rho_{I_n}^n(\mathcal{T}))\}$.*

Definition IV.2 (TT effective rank). Let $\mathcal{T} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ and let $\mathbf{U}_n \mathbf{S}_n \mathbf{V}_n^T$ be the SVD associated with the n -th unfolding $\mathbf{T}_{[n]}$ given by the TT-SVD algorithm, denote by $\{\sigma_i^n\}_{i=1}^{I_n}$ the singular values and consider the associate probability distribution $\left\{ \rho_i^n = \sigma_i^n(\mathcal{T}) / \sum_{j=1}^{I_n} \sigma_j^n \right\}_{i=1}^{I_n}$. The TT effective rank along the n -th mode of \mathcal{T} is defined as $\text{erank}_n^{TT}(\mathcal{T}) = \exp \left\{ H(\rho_1^n(\mathcal{T}), \rho_2^n(\mathcal{T}), \dots, \rho_{I_n}^n(\mathcal{T})) \right\}$.

It is easy to see that Proposition 1 in [23], which gives lower and upper bounds of matrix effective rank, can be straightforwardly extended to cover Tucker and TT effective ranks.

Property 1. Let R_n be the Tucker n -rank (respectively TT n -rank), we have that the Tucker effective rank (resp. TT effective rank) is bounded by $1 \leq \text{erank}_n^{\text{tuc}}(\mathcal{T}) \leq R_n$ (resp. $1 \leq \text{erank}_n^{TT}(\mathcal{T}) \leq R_n$).

As for the case of matrix completion, a crucial question, which may have important implications in understanding the implicit regularization in deep learning, is whether gradient descent in deep tensor factorization converges to solution with minimal nuclear norm, or whether there are other measures that are more relevant to characterize this phenomenon. The experiments described in the next section were designed to provide some clues to this question.

V. EXPERIMENTAL ANALYSIS

Our tensor completion experiments focus on the most interesting regime where few entries are observed. They are designed to better understand the implicit regularization in deep tensor factorization. We first evaluate the reconstruction performance of Tucker and TT (unconstrained) factorization. We then explore whether the implicit regularization induced by gradient-based optimization can be characterized by the minimization of tensor norms or of tensor ranks. Finally, we investigate the impact of the implicit regularization on the singular values dynamics of the learned tensor.

A. Experimental setup

Synthetic data. We generated a $30 \times 30 \times 30$ tensor with Tucker-rank equal to $(6, 6, 6)$ and entries sampled from a centered and reduced normal distribution. Recalling Eq. (4), entries of \mathcal{W} are sampled from a centered normal distribution with variance proportional to σ , $\sigma \in \{0.1, 0.05, 0.025\}$. We also generated a 4-th order tensor of size $12 \times 12 \times 12 \times 12$ with TT-rank equal to $(1, 6, 6, 6, 1)$ with entries sampled from a centered normal distribution with variance proportional to σ , $\sigma \in \{0.075, 0.05, 0.025\}$.

Real data. The comprehensive climate data set (CCDS) is a collection of climate records of North America [24]. The data set contains monthly observations of 17 variables such as carbon dioxide and temperature spanning from 1990 to 2001 across 125 observation locations thus giving a $138 \times 17 \times 125$ tensor.¹ The Meteo-UK data set is collected from the mete-

orological office in UK.² It contains monthly measurements of 5 variables in 16 stations across UK from 1960 to 2000 resulting in a $492 \times 5 \times 16$ tensor; measures were scaled to have zero mean and unitary variance.

Training. Minimization of the loss given by Eq. (3) is done by backpropagation using SGD on full batch; training is stopped when training loss is smaller than 10^{-6} or when $4 \cdot 10^6$ iteration are elapsed. Learning rate is set to 0.1. When measuring reconstruction error, nuclear norms and effective ranks, we took the average over three runs initialized with different seeds.

Baselines. All operations involving tensor decomposition are performed with TensorLy [25]. We compared Tucker UF to HALRTC algorithm [12] and to Tucker HOOI algorithm described in [11]. Both these baseline are available in the PyTen Toolbox [17]. As far as the TT decomposition is concerned, we compared TT UF with TT-SILRTC [14].

B. Reconstruction Error and Depth Effect

For the sake of readability, we point out that the label “measure” on the y -axis of all figures refers to the values taken by the observable specified in the plot’s title. Let us start by looking at the “reconstruction error” plots in Figures 1 and 3 concerning experiments on synthetic tensors with low Tucker rank and TT rank respectively. We can see that both deep Tucker and TT UF outperforms nuclear norm based methods, in particular for the regime where fewer entries are observed (left side of the plots). We also observe that by reducing the variance of the normal distribution used at initialization recovery capacity increases as well. In the case of Tucker UF, we get by Figure 1 (b) that depth plays a beneficial role on completion task. When considering the reconstruction error on CCDS (Figure 4(a)) and Meteo UK (Figure 5(a)) we see that the tendency is much less noticeable with all method giving comparable results.

C. Nuclear Norm vs. Effective Rank

Let us now analyze “Tucker nuclear norm” and “TT nuclear norm” plots in Figures 1 and 3(a). We notice that both Tucker and TT UF do not minimize nuclear norm as much as HALRTC and TT-SILRTC (i.e. nuclear norm-based method) do: HALRTC and TT-SILRTC curves lie indeed below Tucker and TT UF’s. The situation is reversed as soon as we look at the “Tucker effective rank” and “TT effective rank” plots: they show that Tucker and TT UFs minimize effective rank much more than HALRTC and TT-SILRTC methods. This fact gives us strong insight about regularization associated with gradient descent, as it pushes us to questioning Conjecture 1. Nuclear norms plots for real data in Figures 4(a) and 5(a) confirm the tendency we just described for synthetic data. In the case of synthetic data where the rank is known, we also consider Tucker “metric” (resp. TT “metric”) which we define as the Frobenius norm of the difference between the ground truth tensor and a low-rank Tucker (resp. TT) decomposition of the

¹The aggregated and processed data set can be found at <https://viterbi-web.usc.edu/liu32/data.html>.

²<http://www.metoffice.gov.uk/public/weather/climate-historic>



Fig. 1: Tensor completion experiments of a $30 \times 30 \times 30$ tensor with Tucker rank equal to $(6, 6, 6)$ by using Tucker UFs. Dashed lines represent ground truth values of the observable considered. Figure (a) compares depth $(1, 1, 1)$ Tucker UF with HALRTC and Tucker HOOI (where Tucker rank has been set to $(8, 8, 8)$) methods. In (b) we compare Tucker UFs with different depth from $(1, 1, 1)$ to $(3, 3, 3)$.

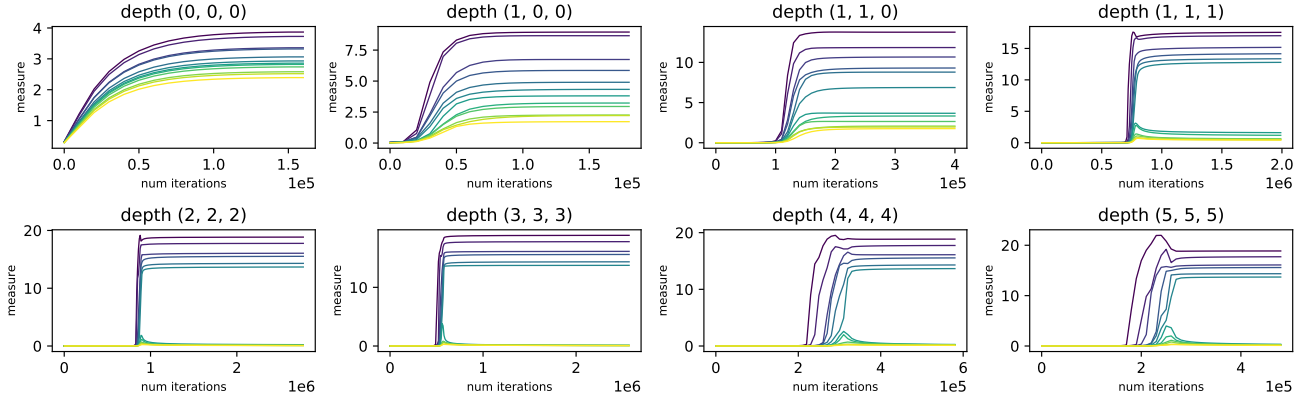


Fig. 2: Time evolution of singular values over iterations of SGD training process to perform tensor completion of a $30 \times 30 \times 30$ tensor with Tucker rank equal to $(6, 6, 6)$ by using Tucker UFs with different depth for Tucker factorization. In each experiment 2750 entries are observed. Plots represents the 12 largest 2-mode singular values.

tensor learned by solving the deep Tucker (resp. TT) UF.³ For these two metrics UFs outperform nuclear norms based methods (see Figures 1 and 3(a)).

D. Singular Value Dynamics

Let us now examine plots displaying time trajectories of generalized singular values (Figures 2, 3(b), 4(b) and 5(b)). In the case of Tucker UF we clearly see that the deeper the model, the sparser the solution given by SGD; moreover, for deep factorization, the transition from non-sparse to sparse solutions happens more abruptly. This observation is consistent with previous results in the matrix case [8]. Sparsity can also be observed for TT UF, though in this case depth comparison

³Note that here we abusively use the term “metric” since all the properties of a distance may not be satisfied.

was not possible. For the matrix case, Theorem 3 in [8] gives a theoretical proof of the behavior of these plots. Roughly stated, it says that

$$\dot{\sigma}(t) \propto (\sigma(t))^{2(1-1/L)} \cdot \gamma(t), \quad (6)$$

where $\gamma(t)$ is the reconstruction error at iteration t , and σ and L being a singular value and the depth of the matrix factorization and it can be interpreted by saying that the exponent controlling the sparsity of the solution is the depth L . For deep Tucker UF, Figure 2 shows that depth plays a key role in driving the learned tensor toward a sparser solution; though experiments did not allow to isolate the behaviour of a single mode as depth k_n along this mode varies, it is reasonable to suppose that, analogously to the matrix setup, k_n might be indeed the exponent which controls sparsity.




Fig. 3: Tensor completion experiments on a $12 \times 12 \times 12 \times 12$ tensor with Tensor-Train rank equal to $(1, 6, 6, 6, 1)$. Dashed lines represent ground truth values of the observable considered. Figure (a) compares TT UFs with TT-SILRTC whereas in (b) the 12 largest 2-mode singular values are depicted.

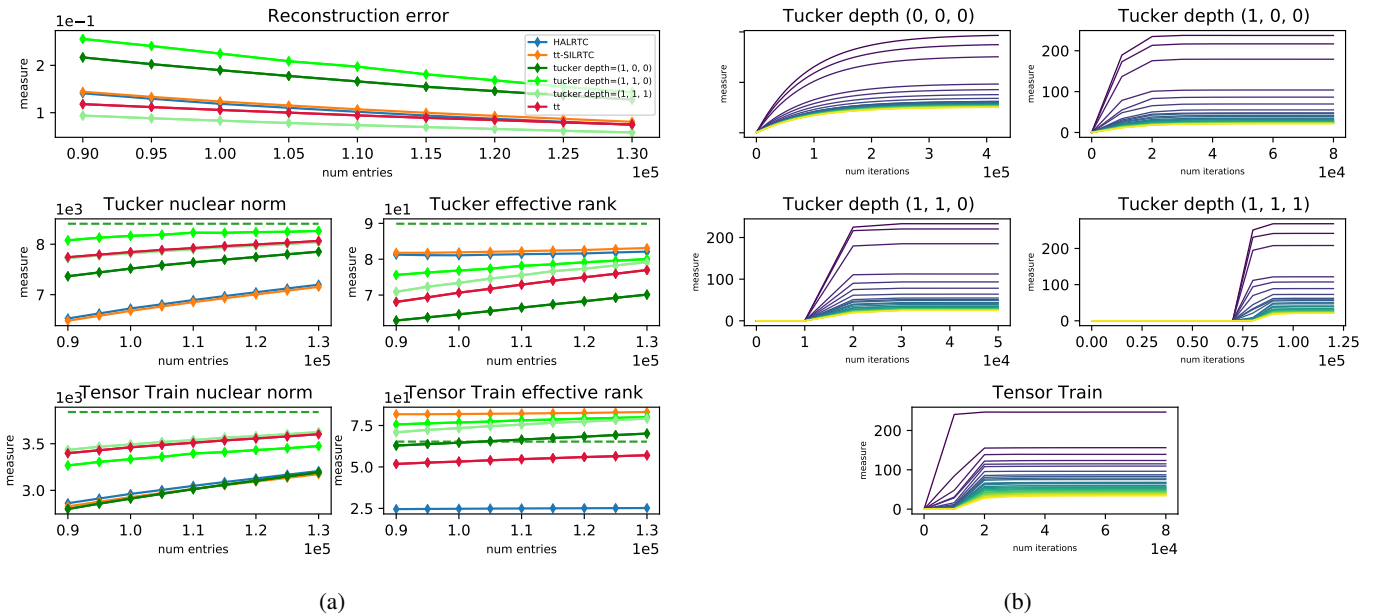


Fig. 4: Tensor completion experiments on the CCDS data set. (a) shows a comparison between Tucker and TT UFs together with HALRTC and TT-SILRTC. (b) shows time evolution of 2-mode singular values where 90000 entries were observed.

We designed an experiment to investigate if a similar behaviour as the one for the matrix case could be observed in the tensor case. Figure 6 shows the evolution of $\ln(\hat{\sigma}_i^n(t)) - \ln(\gamma(t))$ as a function of $\ln(\sigma_i^n(t))$ for different depth values. The linear behaviour of the curves gives an empirical argument in favour of the singular values dynamics provided in the case of matrix completion (see Eq. 6). It is also interesting to note that, similarly to the matrix case, the slope of the curve rises with the depth of Tucker Unconstrained Factorization.

More formally, let us consider deep tensor factorization and a (k_1, k_2, \dots, k_N) -Tucker UF. Analogously to the matrix case, one can imagine that there exists a function on the depths of all the tensor modes, f , so that the generalized singular values

σ_i^n , $\forall n \in [1, \dots, N]$, $\forall i \in [1, \dots, I_n]$, evolve according to

$$\hat{\sigma}_i^n(t) \propto (\sigma_i^n(t))^{f(k_1, \dots, k_N)} \cdot \gamma(t). \quad (7)$$

VI. CONCLUSION

In this paper we explore the implicit regularization induced by gradient descent optimization in deep tensor factorization using the notion of effective rank. Our study provides empirical evidence that tensor nuclear norms is not minimized along the training process, which however tends to produce low tensor rank solutions. In the case of Tucker UF we were able to verify the beneficial effects of depth on the reconstruction error and to formulate a conjecture which gives account of the

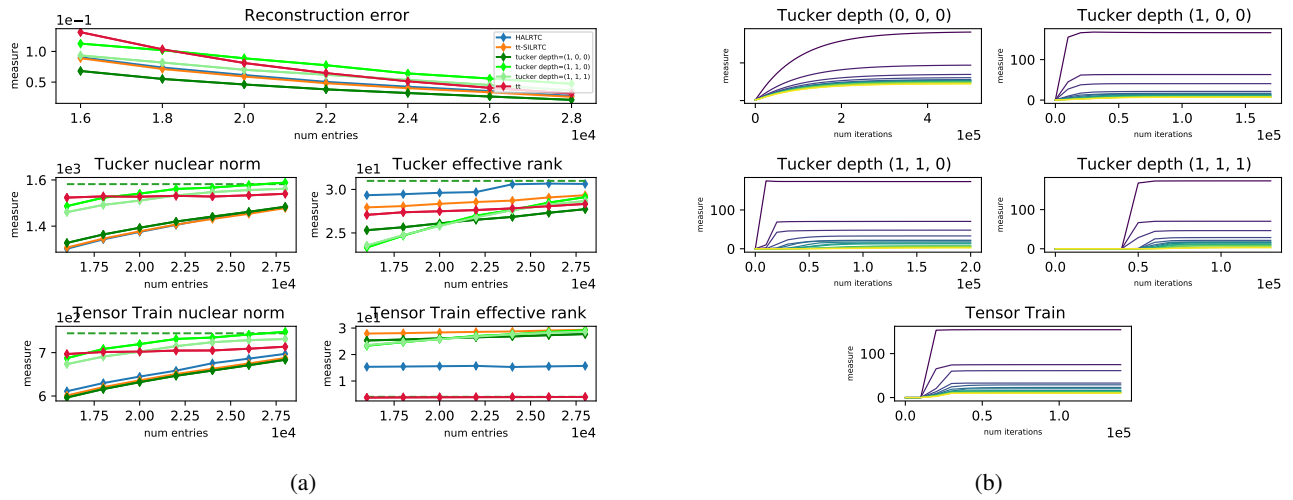


Fig. 5: Tensor completion experiments on the Meteo UK data set. (a) shows a comparison between Tucker and TT UFUs together with HALRTC and TT-SILRTC. (b) shows time evolution of 2-mode singular values singular values, where 20000 entries were observed.




Fig. 6: Behaviour of the singular value dynamics obtained from a $30 \times 30 \times 30$ tensor with Tucker rank equal to $(6, 6, 6)$ and for depths ranging from $(1, 1, 1)$ to $(4, 4, 4)$. Lower left part stand for the beginning of the learning (high loss, small singular value with small derivatives). Upper right part corresponds to convergence (final singular values with small derivatives, small loss).

behavior of generalized singular values. Future research will focus on investigating deeper the dependency of the function f that governs the dynamics of the singular values.

REFERENCES

- [1] K. Kawaguchi, L. P. Kaelbling, and Y. Bengio, “Generalization in deep learning,” *arXiv:1710.05468v6*, 2020.
- [2] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” in *ICLR*, 2017.
- [3] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, “Exploring generalization in deep learning,” in *NeurIPS*, 2017.
- [4] Ruder, S., “An overview of gradient descent optimization algorithms,” *arXiv:1609.04747v2*, 2017.
- [5] N. Gillis and F. Glineur, “Low-rank matrix approximation with weights or missing data is np-hard,” *SIAM J. Matrix Anal. Appl.*, vol. 32, no. 4, pp. 1149–1165, 2011.
- [6] E. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Found. Comput. Math.*, vol. 717, no. 9, 2009.
- [7] S. Gunasekar, B. E. Woodworth, S. Bhojanapalli, B. Neyshabur, and N. Srebro, “Implicit regularization in matrix factorization,” in *NeurIPS*, 2017.
- [8] S. Arora, N. Cohen, W. Hu, and Y. Luo, “Implicit regularization in deep matrix factorization,” in *NeurIPS*, 2019.
- [9] N. Razin and N. Cohen, “Implicit regularization in deep learning may not be explainable by norms,” in *NeurIPS*, 2020.
- [10] G. Tomasi and R. Bro, “PARAFAC and missing values,” *Chemometrics and Intelligent Laboratory Systems*, vol. 75, no. 2, pp. 163–180, 2005.
- [11] L. De Lathauwer, B. De Moor, and J. Vandewalle, “On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higher-order tensors,” *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 4, pp. 1324–1342, 2000.
- [12] J. Liu, P. Musialski, P. Wonka, and J. Ye, “Tensor completion for estimating missing values in visual data,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, pp. 208–220, 2013.
- [13] S. Gandy, B. Recht, and I. Yamada, “Tensor completion and low-n-rank tensor recovery via convex optimization,” *Inverse problems*, vol. 27, no. 2, p. 025010, 2011.
- [14] H. N. Phien, H. D. Tuan, J. A. Bengua, and M. N. Do, “Efficient tensor completion: Low-rank tensor train,” *arXiv:1601.01083*, 2016.
- [15] J. Ying, H. Lu, Q. Wei, J. Cai, D. Guo, J. Wu, Z. Chen, and X. Qu, “Hankel matrix nuclear norm regularized tensor completion for N -dimensional exponential signals,” *IEEE Trans. on Signal Process.*, 2017.
- [16] Y. Liu, F. Shang, H. Cheng, J. Cheng, and H. Tong, “Factor matrix trace norm minimization for low-rank tensor completion,” in *SDM*, 2014.
- [17] Q. Song, H. Ge, J. Caverlee, and X. Hu, “Tensor completion algorithms in big data analytics,” *ACM Trans. Knowl. Discov. Data*, vol. 13, 2019.
- [18] T. G. Kolda and B. W. Bader, “Tensor decompositions and applications,” *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009.
- [19] J. Haastad, “Tensor rank is NP-complete,” *Journal of Algorithms*, vol. 4, no. 11, pp. 644–654, 1990.
- [20] L. De Lathauwer, B. De Moor, and J. Vandewalle, “A multilinear singular value decomposition,” *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 4, pp. 1253–1278, 2000.
- [21] Tucker, L.R., “Some mathematical notes on three-mode factor analysis,” *Psychometrika*, vol. 31, no. 3, pp. 279–311, 1966.
- [22] I. V. Oseledets, “Tensor-Train decomposition,” *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2295–2317, 2011.
- [23] O. Roy and M. Vetterli, “The effective rank: a measure of effective dimensionality,” in *European Signal Processing Conference*, 2007.
- [24] A. C. Lozano, H. Li, A. Niculescu-Mizil, Y. Liu, C. Perlich, J. Hosking, and N. Abe, “Spatial-temporal causal modeling for climate change attribution,” in *KDD*, 2009.
- [25] J. Kossaifi, Y. Panagakis, A. Anandkumar, and M. Pantic, “TensorLy: Tensor learning in Python,” *The Journal of Machine Learning Research*, vol. 20, no. 1, pp. 925–930, 2019.