Exact Learning Dynamics Of Deep Linear Networks With Prior Knowledge

Anonymous Author(s)
Affiliation
Address
email

Abstract

Learning in deep neural networks is known to depend critically on the knowledge embedded in the initial network weights. However, few theoretical results have precisely linked prior knowledge to learning dynamics. Here we derive exact solutions to the dynamics of learning with rich prior knowledge in deep linear networks by generalising Fukumizu's matrix Riccati solution [1]. While simple, deep linear networks retain a non-convex loss landscape and nonlinear learning dynamics that depend in detail on the initial weights of the network. We obtain explicit expressions for the evolving network function, hidden representational similarity, and neural tangent kernel over training for a broad class of initialisations and tasks. We characterise a class of task-independent initialisations that radically alters learning dynamics from slow step-like to fast exponential trajectories while converging to identical representational similarity, dissociating learning trajectories from the structure of internal representations. We discuss the implications of this finding for neural network weight initialisation schemes, continual learning and learning of structured knowledge. Finally, we characterise how network weights dynamically align with task structure, rigorously justifying why previous solutions successfully described learning from small weights without incorporating their fine-scale structure. Taken together, our results provide a mathematical toolkit for understanding the impact of prior knowledge on deep learning.

1 Introduction

A hallmark of human learning is our exquisite sensitivity to prior knowledge: what we already know affects how we subsequently learn [2]. For instance, having learned about the attributes of nine animals, we may learn about the tenth more quickly [3, 4, 5, 6]. In machine learning, the impact of prior knowledge on learning is evident in a range of paradigms including reversal learning [7], transfer learning [8, 9, 10, 11], continual learning [12, 13, 14], curriculum learning [15], and meta learning [16]. One form of prior knowledge in deep networks is the initial network state, which is known to strongly impact learning dynamics [17, 18, 19]. Even random initial weights of different variance can yield qualitative shifts in network behaviour between the *lazy* and *rich* regimes [20], imparting distinct inductive biases on the learning process. More broadly, rich representations such as those obtained through pretraining provide empirically fertile inductive biases for subsequent fine-tuning [21]. Yet while the importance of prior knowledge to learning is clear, our theoretical understanding remains limited, and fundamental questions remain about the implicit inductive biases of neural networks trained from structured initial weights. A better understanding of the impact of initialisation on gradient-based learning may lead to improved pretraining schemes and illuminate pathologies like catastrophic interference in continual learning [22].

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

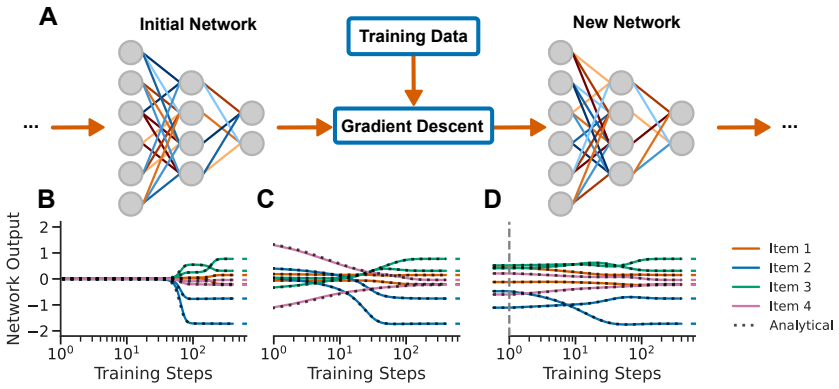

Figure 1: Learning with prior knowledge. **A** In our setting, a deep linear network is trained from a particular initialisation using gradient descent. **B**-**D** Network output for an example task over learning when starting from **B** small random weights, **C** large random weights, and **D** the weights of a previously learned task. Solid lines indicate simulations, dotted lines indicate the analytical solutions we derive in this work. The dynamics depend in detail on the initialisation.

Here, we address this gap by deriving exact solutions to the dynamics of learning in deep linear networks as a function of network initialisation, revealing an intricate and systematic dependence. We consider the setting depicted in Fig. 1A, where a network is trained with standard gradient descent from a potentially complex initialisation. When trained on the same task, different initialisations can radically change the trajectory of the network's output (Fig. 1B-D). Our approach, based on a matrix Riccati formalism [1], provides explicit analytical expressions for the network output over time (Fig. 1B-D dotted). While simple, deep linear networks have a non-convex loss landscape and have been shown to recapitulate several features of nonlinear deep networks while retaining mathematical tractability.

1.1 Related work

Our work builds on analyses of deep linear networks [23, 1, 17, 24, 10, 25, 26] which have shown that this simple model nevertheless has intricate fixed point structure and nonlinear learning dynamics reminiscent of phenomena seen in nonlinear networks. A variety of work has analysed convergence [27, 28], generalisation [10, 29, 30], and the implicit bias of gradient descent [31, 32, 33, 34]. This work mostly considers the *tabula rasa* case of small initial random weights, for which exact solutions are known [17]. By contrast our formalism describes dynamics from a much larger class of initial conditions and can describe alignment dynamics which do not occur in the *tabula rasa* setting. Most directly, our results build from the matrix Riccati formulation proposed by [1]. We extend and refine this result to obtain the dynamics of over- and under-complete networks; to obtain numerically stable forms of the matrix equations; and to more explicitly reveal the impact of initialisation.

A line of theoretical research has considered online learning dynamics in teacher-student settings [35, 36, 37], deriving ordinary differential equations for the average learning dynamics even in nonlinear networks. However, solving these equations requires numerical integration. By contrast, our approach provides explicit analytical solutions for the more restricted case of deep linear networks.

Other approaches for analysing deep network dynamics include the Neural Tangent Kernel (NTK) [38, 39, 40], which can describe nonlinear networks but not representation learning dynamics, and the mean field approach [41, 42, 43], which yields a description of representation learning dynamics in wide networks in terms of a PDE but does not allow explicit solutions. Our work is similar in seeking a subset of more tractable models that are amenable to analysis, but we focus on the impact of initialisation on representation learning dynamics and explicit solutions.

A large body of work has investigated the effect of different random initialisations on learning in deep networks. The role of initialisation in the vanishing gradient problem and proposals for better initialisation schemes have been illuminated by several works drawing on the central limit theorem [44, 17, 45, 18, 46], reviewed in [47, 19, 48]. These approaches typically guarantee that gradients do not vanish at the start of learning, but do not analytically describe the resulting learning

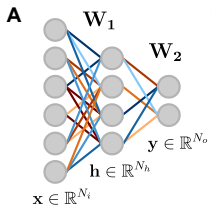

B

$$\mathbf{Q}\mathbf{Q}^T(t) = \begin{bmatrix} \mathbf{W}_1^T\mathbf{W}_1(t) & \mathbf{W}_1^T\mathbf{W}_2^T(t) \\ \mathbf{W}_2\mathbf{W}_1(t) & \mathbf{W}_2\mathbf{W}_2^T(t) \end{bmatrix}$$

C

$\hat{\mathbf{Y}}$: $\mathbf{W}_2\mathbf{W}_1(t)\mathbf{X}$

NTK: $N_o\,\mathbf{X}^T\mathbf{W}_1^T\mathbf{W}_1(t)\mathbf{X} + \mathrm{Tr}(\mathbf{W}_2\mathbf{W}_2^T(t))\mathbf{X}^T\mathbf{X}$

RSAs: $\mathbf{W}_1^T\mathbf{W}_1(t)\mathbf{X}$ and $\mathbf{W}_2\mathbf{W}_2^T(t)\mathbf{Y}$

Figure 2: Neural network architecture and the $\mathbf{Q}\mathbf{Q}^T(t)$ matrix. **A** Schematic representation of a two layer linear neural network with N_i inputs, N_h hidden nodes and N_o outputs. $\mathbf{W}_1$ and $\mathbf{W}_2$ are the first and second layer weights respectively. **B** The four quadrants of the $\mathbf{Q}\mathbf{Q}^T(t)$ matrix fully determine the temporal dynamics of **C:** The network function (orange), the neural tangent kernel (green and blue) and the representational similarity matrices of the input (green) and output (blue).

trajectories. Influential work has shown that network initialisation variance mediates a transition from *rich* representation learning to *lazy* NTK dynamics [20], which we analyse in our framework.

2 Preliminaries and setting

Consider a supervised learning task in which input vectors $\mathbf{x}_n \in \mathbb{R}^{N_i}$ from a set of P training pairs $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1\ldots P}$ have to be associated with their target output vectors $\mathbf{y}_n \in \mathbb{R}^{N_o}$. We learn this task with a two-layer linear network model (Fig. 2A) that produces the output prediction

$$\hat{\mathbf{y}}_n = \mathbf{W}_2\mathbf{W}_1\mathbf{x}_n, \tag{1}$$

with weight matrices $\mathbf{W}_1 \in \mathbb{R}^{N_h \times N_i}$ and $\mathbf{W}_2 \in \mathbb{R}^{N_i \times N_h}$ where N_h is the number of hidden units. The network's weights are optimised using full batch gradient descent with learning rate η (or respectively time constant $\tau = {}^1\!/\!_\eta$) on the mean squared error loss

$$\mathcal{L}(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{2}\left\langle ||\hat{\mathbf{y}} - \mathbf{y}||^2 \right\rangle, \tag{2}$$

where $\langle\cdot\rangle$ denotes the average over the dataset. The input and input-output correlation matrices of the dataset are

$$\tilde{\mathbf{\Sigma}}^{xx} = \frac{1}{P}\sum_{n=1}^{P}\mathbf{x}_n\mathbf{x}_n^T \in \mathbb{R}^{N_i \times N_i} \quad \text{and} \quad \tilde{\mathbf{\Sigma}}^{yx} = \frac{1}{P}\sum_{n=1}^{P}\mathbf{y}_n\mathbf{x}_n^T \in \mathbb{R}^{N_o \times N_i}. \tag{3}$$

Finally, the gradient optimisation starts from an initialisation $\mathbf{W}_2(0), \mathbf{W}_1(0)$. Our goal is to understand the full time trajectory of the network's output and internal representations as a function of this initialisation and the task statistics.

Our starting point is the seminal work of Fukumizu [1], which showed that the gradient flow dynamics could be written as a matrix Riccati equation with known solution. In particular, defining

$$\mathbf{Q} = \begin{bmatrix} \mathbf{W}_1^T \\ \mathbf{W}_2 \end{bmatrix} \quad \text{and} \quad \mathbf{F} = \begin{bmatrix} 0 & (\tilde{\mathbf{\Sigma}}^{yx})^T \\ \tilde{\mathbf{\Sigma}}^{yx} & 0 \end{bmatrix}, \tag{4}$$

Fukumizu showed that the continuous time dynamics of the matrix $\mathbf{Q}\mathbf{Q}^T$ is

$$\mathbf{Q}\mathbf{Q}^T(t) = e^{\mathbf{F}\frac{t}{\tau}}\mathbf{Q}(0)\left[\mathbf{I} + \frac{1}{2}\mathbf{Q}(0)^T\left(e^{\mathbf{F}\frac{t}{\tau}}\mathbf{F}^{-1}e^{\mathbf{F}\frac{t}{\tau}} - \mathbf{F}^{-1}\right)\mathbf{Q}(0)\right]^{-1}\mathbf{Q}(0)^T e^{\mathbf{F}\frac{t}{\tau}}, \tag{5}$$

with initial state $\mathbf{Q}(0)$, if the following four assumptions hold [1]:

Assumption 2.1 *The dimensions of the input and target vectors are identical, that is $N_i = N_o$.*

Assumption 2.2 *The input data is whitened, that is $\tilde{\mathbf{\Sigma}}^{xx} = \mathbf{I}$.*

Assumption 2.3 *The network's weight matrices are zero-balanced at the beginning of training, that is $\mathbf{W}_1(0)\mathbf{W}_1(0)^T = \mathbf{W}_2(0)^T\mathbf{W}_2(0)$. If this condition holds at initialisation, it will persist throughout training [17, 24].*

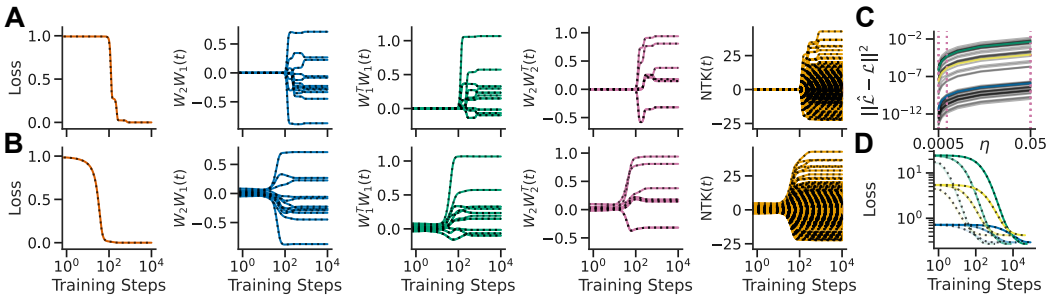

Figure 3: Exact learning dynamics **A** The temporal dynamics of the numerical simulation (coloured lines) of the loss, network function, input RSA, output RSA and NTK (columns 1-5 respectively) are exactly matched by the analytical solution (black dotted lines) for small initial weight values and **B** large initial weight values. **C** Each line shows the deviation of the analytical loss $\hat{\mathcal{L}}$ from the numerical loss $\mathcal{L}$ for one of $n = 50$ randomly sampled networks and targets across a range of learning rates $\eta \in [0.05, 0.0005]$. The deviation is low even for large learning rates and mutually decreases with the learning rate. **D** Numerical and analytical learning curves for three example networks (coloured lines in C) and three learning rates (pink dotted lines in C).

Assumption 2.4 *The input-output correlation matrix $\tilde{\mathbf{\Sigma}}^{yx}$ and the initial state of the network function have full rank, that is* $\text{rank}(\tilde{\mathbf{\Sigma}}^{xy}) = \text{rank}(\mathbf{W}_1(0)\mathbf{W}_2(0)) = N_i = N_o$.

For completeness, we include a derivation of this solution in Appendix A.

Rather than tracking the weights' dynamics directly, this approach tracks several key statistics collected in the matrix $\mathbf{QQ}^T$. In particular, this matrix can be separated into four quadrants with intuitive meaning (Fig. 2B): the off-diagonal blocks contain the network function, while the on-diagonal blocks contain the correlation structure in each weight matrix separately. These quantities permit calculation of the network's predictions, but also aspects of its internal representation including the representational similarity analysis (RSA) of the weights [49] and the network's neural tangent kernel [38, 39, 40] (Fig. 2C). Hence the solution in Eqn. (5) describes important aspects of network behaviour.

However, in this form, the solution has several limitations. First, it relies on general matrix exponentials and inverses, which are a barrier to explicit understanding. Second, when evaluated numerically, it is often unstable. And third, the equation is only valid for equal input and output dimensions. In the following section we overcome these limitations.

Implementation and simulation All simulation details are in Appendix G. Code to replicate all simulations and plots are in the supplementary material and available online under a *GPLv3* license [1]. Reproduction of all simulation takes <6 hours on a single AMD Ryzen 5950x.

3 Exact learning dynamics with prior knowledge

In this section we derive an exact and numerically stable solution for $\mathbf{QQ}^T$ that better reveals the learning dynamics, convergence behaviour and generalisation properties of two-layer linear networks with prior knowledge. First, we weaken the assumption of equal input and output dimensions 2.1:

Assumption 3.1 *The input dimension is unequal to the output dimension, that is $N_i \neq N_o$.*

Assumption 3.2 *The input-output correlation and the initial state of the network have sufficient rank, that is* $\text{rank}(\tilde{\mathbf{\Sigma}}^{xy}) = \text{rank}(\mathbf{W}_1(0)\mathbf{W}_2(0)) = \min(N_i, N_o)$.

Assumption 3.3 *The network is not bottlenecked, that is $N_h \geq \min(N_i, N_o)$.*

Next, to place the solution in a more explicit form, we make use of the singular value decomposition. Let the singular value decomposition of the initial network function and the input-output correlation

[1] github.com/anonymized

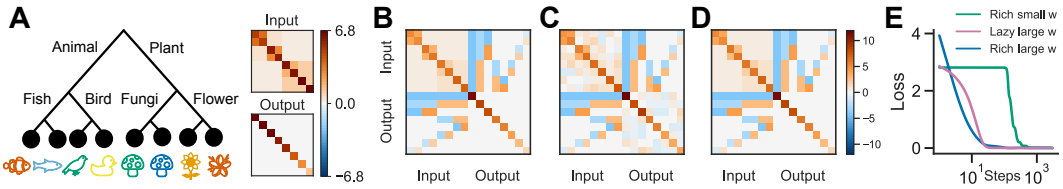

Figure 4: **A** Semantic learning task and its respective input and output RSA. Rows and columns in the RSA are identically ordered as the order of items in the hierarchical tree. **B** Final $\mathbf{QQ}^T$ matrices after training converged when initialised from random small weights, **C** random large weights (note how the upper left and lower right quadrant differ from the task's RSA) and **D** large zero-balanced weights. **E** Learning curves for the three different initialisations as in B (green), C (pink) and D (blue). While both large weight initialisations lead to exponentially learning curves, the small weight initialisation leads to a step like decay of the loss.

of the task be

$$\mathrm{SVD}(\mathbf{W}_2(0)\mathbf{W}_1(0)) = \mathbf{USV}^T \quad \text{and} \quad \mathrm{SVD}(\tilde{\mathbf{\Sigma}}^{yx}) = \tilde{\mathbf{U}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T. \tag{6}$$

We note that if $N_i < N_o$, the left and right singular vectors are of shape $N_o \times N_i$ and $N_i \times N_i$ respectively and we define $\mathbf{U}_\perp \in \mathbb{R}^{N_o \times (N_o - N_i)}$ as the matrix containing orthogonal column vectors that complete the basis, i.e., make $[\mathbf{U}\ \mathbf{U}_\perp]$ orthonormal. Conversely, we define $\mathbf{V}_\perp \in \mathbb{R}^{N_i \times (N_i - N_o)}$ similarly for the case of $N_i > N_o$.

Theorem 3.1 *Define* $\mathbf{B} = \mathbf{U}^T\tilde{\mathbf{U}} + \mathbf{V}^T\tilde{\mathbf{V}}$ *and* $\mathbf{C} = \mathbf{U}^T\tilde{\mathbf{U}} - \mathbf{V}^T\tilde{\mathbf{V}}$. *Under the assumptions of whitened inputs, 2.2, zero-balanced weights 2.3, unequal input-output dimensions 3.1, sufficient rank 3.2, no bottlenecks 3.3, and* $\mathbf{B}$ *nonsingular, the temporal dynamics of* $\mathbf{QQ}^T$ *are*

$$\begin{aligned}\mathbf{QQ}^T(t) = \mathbf{Z}\Big[&4e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\mathbf{B}^{-1}\mathbf{A}^{-1}\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}} + \left(\mathbf{I} - e^{-2\tilde{\mathbf{S}}\frac{t}{\tau}}\right)\tilde{\mathbf{S}}^{-1} \\ &- e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\mathbf{B}^{-1}\mathbf{C}\left(e^{-2\tilde{\mathbf{S}}\frac{t}{\tau}} - \mathbf{I}\right)\tilde{\mathbf{S}}^{-1}\mathbf{C}^T\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}} \\ &+ 4\frac{t}{\tau}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\mathbf{B}^{-1}(\mathbf{V}^T\mathbf{V}_\perp{\mathbf{V}_\perp}^T\mathbf{V} + \mathbf{U}^T\mathbf{U}_\perp{\mathbf{U}_\perp}^T\mathbf{U})\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\Big]^{-1}\mathbf{Z}^\mathbf{T}\end{aligned} \tag{7}$$

with

$$\mathbf{Z} = \begin{bmatrix}\tilde{\mathbf{V}}\left(\mathbf{I} - e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\mathbf{C}^T\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\right) + 2\mathbf{V}_\perp{\mathbf{V}_\perp}^T\mathbf{V}\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}} \\ \tilde{\mathbf{U}}\left(\mathbf{I} + e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\mathbf{C}^T\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\right) + 2\mathbf{U}_\perp{\mathbf{U}_\perp}^T\mathbf{U}\mathbf{B}^{-T}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\end{bmatrix}. \tag{8}$$

For a proof of Theorem 3.1 please refer to Appendix B.

With this solution we can calculate the exact temporal dynamics of the loss, network function, RSAs and NTK, as we verify in Fig. 3A, B. Because it contains only negative exponentials, it is numerically stable and provides high precision across a wide range of learning rates and network architectures (Fig. 3C, D).

Importantly, the solution depends on the inverse of $\mathbf{B}$. As we prove in Appendix B, $\mathbf{B}$ is nonsingular except in certain reversal learning cases. We note that a solution for the weights $\mathbf{W}_1(t)$ and $\mathbf{W}_2(t)$, i.e., $\mathbf{Q}(t)$, can be derived up to a time varying orthogonal transformation as demonstrated in Appendix B. Further, as time dependent variables only occur in matrix exponentials of diagonal matrices of negative sign, the temporal limit is well defined.

Theorem 3.2 *Under the assumptions of Theorem 3.1, the network function converges to the global minimum* $\tilde{\mathbf{U}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T$ *and acquires a rich task-specific internal representation, that is* $\mathbf{W}_1^T\mathbf{W}_1 = \tilde{\mathbf{V}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T$ *and* $\mathbf{W}_2\mathbf{W}_2^T = \tilde{\mathbf{U}}\tilde{\mathbf{S}}\tilde{\mathbf{U}}^T$.

The proof of Theorem 3.2 is in Appendix B. We now turn to several implications of these results.

4 Rich and lazy learning regimes and generalisation

Recent results have shown that large deep networks can operate in qualitatively distinct regimes depending on their weight initialisations [20, 50], the so called *rich* and *lazy* regimes. In the *rich*

regime, learning dynamics can be highly nonlinear and lead to task-specific solutions thought to lead to favourable generalisation properties [20, 51, 50]. By contrast, the *lazy* regime exhibits simple exponential learning dynamics and exploits high-dimensional nonlinear projections of the data produced by the initial random weights, leading to task-agnostic representations that attain zero training error but possibly lower generalisation performance [38, 39, 40]. Traditionally, the *rich* and *lazy* learning regimes have been respectively linked to low and high variance initial weights (in relation to the network layer size).

To illustrate these phenomena, we consider a semantic learning task in which a set of living things have to be correctly linked to their position in a hierarchical structure (Fig. 4A). The representational similarity of the input of the task ($\mathbf{X}\mathbf{X}^T = \tilde{\mathbf{V}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T$) reveals its inherent structure (Fig. 4A). For example, the representations of the two fishes are most similar to each other, less similar to birds and least similar to plants. Likewise, the representational similarity of the task's target values ($\mathbf{Y}\mathbf{Y}^T = \tilde{\mathbf{U}}\tilde{\mathbf{S}}\tilde{\mathbf{U}}^T$) reveals the primary groups among which items are organised. As a consequence, one can for example predict from an object being a fish that it is an animal and from an object being a plant that it is not a bird. Reflecting these structural relationships in internal representations can allow the *rich* regime to generalise in ways the *lazy* regime cannot. Crucially, $\mathbf{Q}\mathbf{Q}^T(t)$ contains the temporal dynamics of the weights' representational similarity and therefore can be used to study if a network finds a *rich* or *lazy* solution.

When training a two layer network from random small initial weights, the weights' input and output RSA (Fig. 4B, upper left and lower right quadrant) are identical to the task's structure. However, when training from large initial weights, the RSA reveals that the network has converged to a *lazy* solution (Fig. 4C). We emphasise that the network function in both cases is identical (Fig. 4B, C, lower left quadrant). And while their final loss is identical too, their learning dynamics evolve slow and step-wise in the case of small initial weights and fast and exponentially in the case of large initial weights (Fig. 4E), as predicted by previous work [20].

However, from Theorem 3.2 it directly follows that in our setup, the set of zero-balanced initial weights 2.3 is guaranteed to find a *rich* solution in which the weights' RSA is identical to the task's RSA, i.e., $\mathbf{W}_1^T\mathbf{W}_1 = \tilde{\mathbf{V}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T = \mathbf{X}\mathbf{X}^T$ and $\mathbf{W}_2\mathbf{W}_2^T = \tilde{\mathbf{V}}\tilde{\mathbf{S}}\tilde{\mathbf{V}}^T = \mathbf{Y}\mathbf{Y}^T$. Therefore, as zero-balanced weights may be large, there exist initial states that lead to *rich* solutions while evolving as rapid exponential learning curves (Fig. 4D, E). Crucially, these initialisations are task-agnostic, in the sense that they are independent of the task structure (cf. [52]). This finding applies to any learning task with well defined input-output correlation. For additional simulations see Appendix C. Hence our dynamics can describe the change in dynamics from step-like to exponential with increasing weight scale, and separate this dynamical phenomenon from the structure of internal representations.

5 Exact alignment dynamics

The learning dynamics of deep linear networks depend on the exact initial values of the synaptic weights. Previous solutions [17] for the exact temporal learning dynamics relied on full alignment between initial state and the target task, i.e. that $\mathbf{U} = \tilde{\mathbf{U}}$ and $\mathbf{V} = \tilde{\mathbf{V}}$. In other words, the initial state of the network was assumed to contain part of the task's structure before the onset of training. However, this assumption is violated in most real-world scenarios. As a consequence, most previous work has relied on the empirical observation that learning from *tabula rasa* small initial weights occurs in two phases: First, the network's singular vectors rapidly align with the task's singular vectors and then subsequently, the singular values are learned. Because this alignment process is fast, when training begins from small initial weights, the learning dynamics are still approximately described by the temporal learning dynamics of the singular values. This dynamic has been called a "silent alignment" process [26]. Here we leverage our matrix Riccati approach to provide exact solutions to the dynamics of this alignment process. Further, we lift the assumption of small random weights and characterise the transition from stage-like to exponential learning trajectories as initialisation variance increases, as seen in Section 4. We begin by deriving an alternate form that eases analysis of alignment in particular.

Theorem 5.1 *Assume that $N_i = N_o = N_h$. Let the weight matrices of a two layer linear network be initialised by $\mathbf{W}_1 = \mathbf{A}(0)\tilde{\mathbf{V}}^T$ and $\mathbf{W}_2 = \tilde{\mathbf{U}}\mathbf{A}(0)^T$, where $\mathbf{A}(0) \in \mathbb{R}^{N_h \times N_i}$ is an arbitrary, invertible matrix. Then, under assumptions 2.1-2.4, the temporal dynamics of $\mathbf{Q}\mathbf{Q}^T$ are fully*

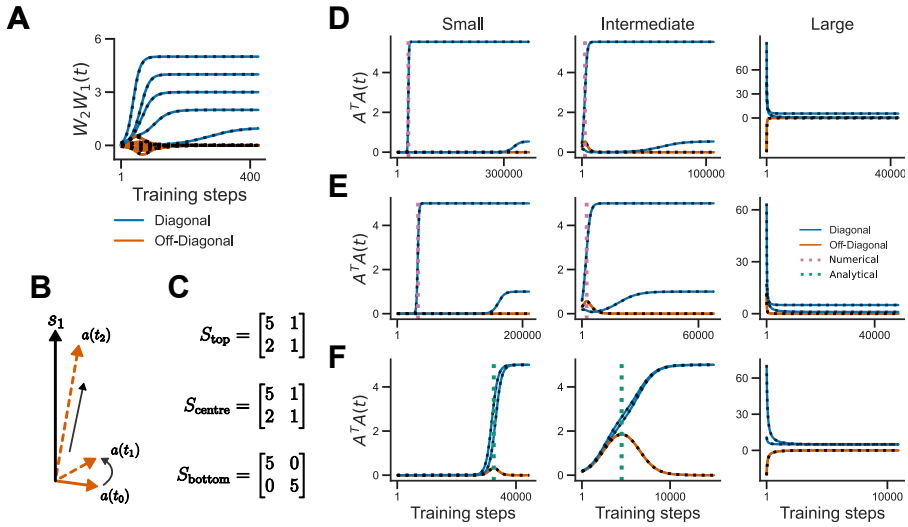

Figure 5: Alignment dynamics. **A** Analytical (black dotted lines) and numerical (colored lines) of the temporal dynamics of the on- and off-diagonal elements of $\mathbf{A}^T\mathbf{A}$ in blue and red, respectively. **B** Schematic representation of the alignment process. **C** Three target matrices with dense, unequal diagonal, and equal diagonal structure. **D-F** Alignment dynamics for the top (D), middle (E), and bottom (F) tasks depicted in panel C. Row F contains analytical predictions for the time of the peak of the off-diagonal (dashed green). The network is initialised as defined in D with small, intermediate and large variance.

determined by

$$\mathbf{A}^T\mathbf{A}(t) = \left[e^{-\tilde{\mathbf{S}}\frac{t}{\tau}}\left(\mathbf{A}(0)^T\mathbf{A}(0)\right)^{-1}e^{-\tilde{\mathbf{S}}\frac{t}{\tau}} + (\mathbf{I} - e^{-2\tilde{\mathbf{S}}\frac{t}{\tau}})\tilde{\mathbf{S}}^{-1}\right]^{-1}. \tag{9}$$

We say that a network function is *aligned* with a task if the network's and the task's left and right singular vectors (eq. 6) are identical, that is, $\tilde{\mathbf{U}} = \mathbf{U}(t)$ and $\tilde{\mathbf{V}} = \mathbf{V}(t)$.

Theorem 5.2 *The network function is aligned if and only if the off-diagonal values of $\mathbf{A}^T\mathbf{A}(t)$ are zero.*

For proofs of Theorems 5.1 and 5.2, please refer to Appendix D. We remark that the form of the initial weights is less general than that in Theorem 3.1, and in particular implies $\mathbf{UV} = \tilde{\mathbf{U}}\tilde{\mathbf{V}}$. As shown in Fig. 5A, using Eq. 9 we can exactly describe the temporal dynamics of the on- and off-diagonal values of $\mathbf{A}^T\mathbf{A}(t)$. Intuitively, off-diagonal values in $\mathbf{A}^T\mathbf{A}$ represent counterproductive cross-coupling between different singular vector modes, and eventually converge to zero despite transient peaks. From Theorem 5.2 it follows that the temporal dynamics of the alignment process depend on the trajectories of the off-diagonal entries in $\mathbf{A}^T\mathbf{A}$.

However, the outer inverse in equation 9 renders it difficult to study high dimensional networks. Therefore, we first focus on small networks with input and output dimension $N_i = 2$ and $N_o = 2$, for which a lengthy but explicit analytical solution is given in Appendix D. In this setting, the structure of the weight initialisation and task are encoded in the matrices

$$\mathbf{A}(0)^T\mathbf{A}(0) = \begin{bmatrix} a_1(0) & b(0) \\ b(0) & a_2(0) \end{bmatrix} \quad \text{and} \quad \tilde{\mathbf{S}} = \begin{bmatrix} s_1 & 0 \\ 0 & s_2 \end{bmatrix}, \tag{10}$$

where the parameters $a_1(0)$ and $a_2(0)$ represent the component of the initialisation that is aligned with the task, and $b(0)$ represents cross-coupling, such that taking $b(0) = 0$ recovers previously known solutions for the aligned case [17]. We use this setting to demonstrate two features of the learning dynamics.

Alignment dynamics. First, we track alignment by considering the dynamics of the off-diagonal element $b(t)$ (Fig. 5D-F red lines). At convergence, the off-diagonal element shrinks to zero as shown in Appendix D. However, strikingly, it can exhibit non-monotonic trajectories with transient

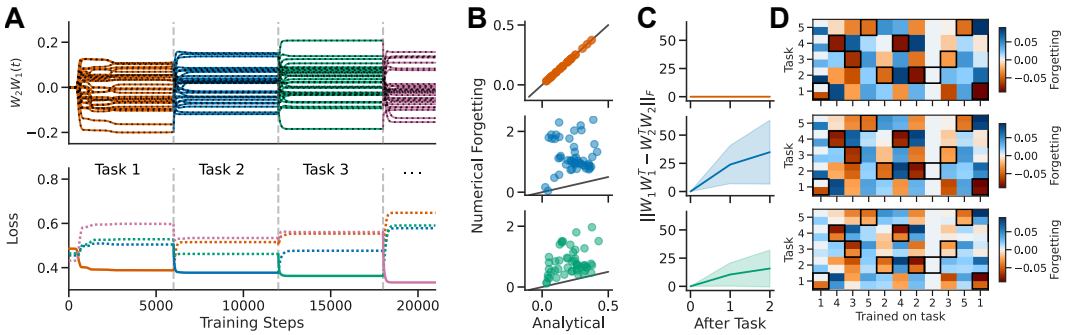

Figure 6: Continual learning. **A** Starting from random initial weights, a network is trained on a sequence of tasks (colored lines show simulation and black dotted lines analytical results). As the network function is optimised on the current task (top row), the loss of the other tasks increases. **B** Comparison of the numerical and analytical amount of catastrophic forgetting on a first task after training on a second task in for $n = 50$ linear (red), $\tanh$ (blue) and ReLU (green) networks. **C** Weight alignment before the onset of training and after training on a first and second task for $n = 50$ networks in linear (red), $\tanh$ (blue) and ReLU (green) networks. Shaded error show $\pm\,\mathrm{std}$. **D** Amount of catastrophic forgetting for each of 5 tasks when trained on a random sequence of tasks for a linear (top), $\tanh$ (center) and ReLU (bottom) network. The top half of each square shows the pre-computed analytical amount of forgetting and the bottom half the actual numerical value. The current task highlighted with a frame.

peaks or troughs partway through the learning process. In particular, in Appendix D we derive the time of the peak magnitude as $t_{peak} = \frac{\tau}{4s} \ln \frac{s(s-a_1-a_2)}{a_1 a_2 - b^2}$ (Fig 5F green dotted line), which coincides approximately with the time at which the on-diagonal element is half learned. From small random weights, the off-diagonal remains near-zero throughout learning, reminiscent of the silent alignment effect [26]. For large initialisations, no peak is observed and the dynamics are exponential. At intermediate initialisations, the maximum of the off-diagonal is reached before the singular mode is fully learned (Appendix D). Intuitively, a particular input singular vector can initially project appreciably onto the wrong output singular vector, corresponding to initial misalignment, but this is only revealed when this link is amplified, at which point corrective dynamics remove the counter-productive coupling, as schematised in Fig. 5B. However, the alignment dynamics have a minimal effect on the on-diagonal trajectories. We report further measurements of alignment in Appendix D.

Effect of initialisation variance. Next, we revisit the impact of initialisation scale for the on-diagonal dynamics. As shown in figure 5D-F, as the initialisation variance grows the learning dynamics change from sigmoidal to exponential, possibly displaying more complex behavior at intermediate variance D. In this simple setting we can analyse this transition in detail. Taking $s_1 = s_2 = s$ as in 5F and $|a_1(0)|, |a_2(0)|, |b(0)| \ll 1$, we recover a sigmoidal trajectory,

$$a_1(t) = \frac{s a_1(0)}{e^{\frac{-2st}{\tau}} \left[s - a_1(0) - a_2(0)\right] + a_1(0) + a_2(0)}, \tag{11}$$

while for $|a_1(0)|, |a_2(0)|, |b(0)| \gg 0$ the dynamics of the on-diagonal element a_1 is close to exponential (Fig. 5D-F left and right columns). We examine larger networks in Appendix D.

6 Applications

The solutions derived in Sections 3 and 5 provide tools to examine the impact of prior knowledge on dynamics in deep linear networks. So far we have traced general features of the behaviour of these solutions. In this section, we use this toolkit to develop accounts of several specific phenomena.

Continual Learning Continual learning (see [12] for a review) and the pathology of catastrophic interference have long been a challenge for neural network models [22, 53, 54]. A variety of theoretical work has investigated aspects of continual learning [55, 56, 57, 58, 59]. In this setting, starting from an initial set of weights, a network is trained on a sequence of tasks with respective input-output correlations $\mathcal{T}_1 = \tilde{\mathbf{\Sigma}}_1^{yx}, \mathcal{T}_2 = \tilde{\mathbf{\Sigma}}_2^{yx}, \mathcal{T}_3 = \tilde{\mathbf{\Sigma}}_3^{yx}, \ldots$. As shown in Fig. 6A, our dynamics

immediately enable exact solutions for the full continual learning process, whereby the final state after training on one task becomes the initialisation for the next task. These solutions thus reveal the exact time course of forgetting for arbitrary sequences of tasks.

Training on later tasks can overwrite previously learned knowledge, a phenomenon known as catastrophic interference [22, 53, 54]. From Theorem 3.2 it follows that from any arbitrary zero-balanced initialisation 2.3, the network converges to the global optimum such that the initialisation is completely overwritten and forgetting is truly catastrophic. In particular, the loss of any other task j after training to convergence on task i is $\mathcal{L}(\mathcal{T}_j) = 1/2||\tilde{\Sigma}_j^{yx} - \tilde{\Sigma}_i^{yx}||_F^2$. Hence the amount of forgetting after convergence is fully determined by the similarity structure of the two tasks (Fig. 6B), and we can fully predict catastrophic forgetting for a sequence of tasks of arbitrary length before training onset (Fig. 6D). As expected, our results depend on our linear setting and $\tanh$ or ReLU nonlinearities can show different behaviour, typically increasing the amount of forgetting (Fig. 6B,D). Further, in nonlinear networks, weights become rapidly unbalanced and forgetting values that are calculated before the onset of training do not predict the actual outcome (Fig. 6C). In summary, our results link exact learning dynamics with catastrophic forgetting and thus provide an analytical tool to study the mechanisms and potential counter measures underlying catastrophic forgetting.

Revising structured knowledge Finally, we study knowledge revision, that is, how pre-existing knowledge can be altered. We consider the setting in which a network first learns about a hierarchical dataset (e.g., the living kingdom, Fig. 4A), and subsequently must alter this knowledge. First we examine reversal learning, which requires exchanging knowledge about two items, for instance as might arise when items are mislabeled and later corrected. We show analytically and numerically in Appendix E that reversal learning in fact does not succeed in deep linear networks; the first task lies exactly on the separatrix of a saddle point, and the dynamics converge to zero (Fig. 7A), though any noise will allow learning to proceed. However, the dynamics still slow in the vicinity of the saddle point, providing a theoretical explanation for catastrophic slowing in deep linear networks [60]. By contrast, shallow networks do not exhibit catastrophic slowing during reversal learning (Fig. 7B), as we show in Appendix E.

Next, we consider learning a different hierarchy for the same set of items, for instance learning a color hierarchy (Appendix E) after the living hierarchy. These tasks share structure at a highly abstract level, and differ purely in their singular vectors, not their singular values. We numerically and analytically show that deep linear networks are unable to exploit this similarity, and the network undergoes similar learning dynamics to those for random weights of similar scale (Appendix E). In other words, deep linear networks in our setting do not have an advantage in learning subsequent abstract structures of a similar form.

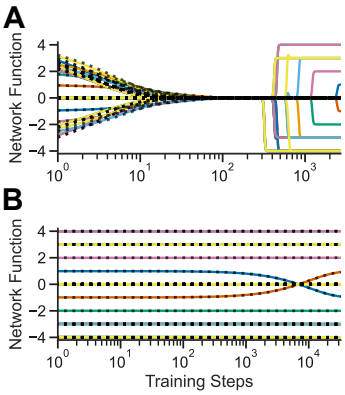

Figure 7: Catastrophic slowing **A** Numerical learning dynamics (coloured lines) and analytical solution (black dotted line) of a reversal learning task. The analytical solution gets stuck on a saddle point while the numerical solution reaches the target. **B** Learning dynamics (coloured lines) and analytical solution of a shallow network of the same task. The learning converges.

7 Discussion

We derive exact solutions to the dynamics of learning with rich prior knowledge in a tractable model class: deep linear networks. While our results broaden the class of two-layer linear network problems that can be described analytically, they remain limited and rely on a set of assumptions (2.1-3.2). In particular, weakening the requirement that the input covariance be white and the weights be zero-balanced would enable analysis of the impact of initialisation on internal representations. Nevertheless, these solutions reveal several insights into network behaviour. We show that there exists a large set of initial values, namely zero-balanced weights 2.3, which lead to task-specific representations; and that large initialisations lead to exponential rather than sigmoidal learning curves. We hope our results provide a mathematical toolkit that illuminates the complex impact of prior knowledge on deep learning dynamics.

References

[1] K. Fukumizu. Effect of Batch Learning In Multilayer Neural Networks. In *Proceedings of the 5th International Conference on Neural Information Processing*, pages 67–70, 1998.

[2] S.E. Carey. *Conceptual Change In Childhood*. MIT Press, Cambridge, MA, 1985.

[3] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419, 1995.

[4] Gregory Murphy. *The big book of concepts*. MIT press, 2004.

[5] James L McClelland. Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. *Journal of Experimental Psychology: General*, 142(4):1190, 2013.

[6] Timo Flesch, Jan Balaguer, Ronald Dekker, Hamed Nili, and Christopher Summerfield. Comparing continual task learning in minds and machines. *Proceedings of the National Academy of Sciences*, 115(44):E10313–E10322, 2018.

[7] Burak Erdeniz and Nart Bedin Atalay. Simulating probability learning and probabilistic reversal learning using the attention-gated reinforcement learning (agrel) model. In *The 2010 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE, 2010.

[8] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.

[9] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.

[10] Andrew K. Lampinen and Surya Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL `https://openreview.net/forum?id=ryfMLoCqtQ`.

[11] Federica Gerace, Luca Saglietti, Stefano Sarao Mannelli, Andrew Saxe, and Lenka Zdeborová. Probing transfer learning with a model of synthetic correlated datasets. *Machine Learning: Science and Technology*, 2022.

[12] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.

[13] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[14] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017.

[15] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.

[16] Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, pages 1820–1830, 2019.

[17] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.

[18] Jeffrey Pennington, Samuel Schoenholz, and Surya Ganguli. Resurrecting the sigmoid in deep learning through dynamical isometry: theory and practice. *Advances in neural information processing systems*, 30, 2017.

[19] Yasaman Bahri, Jonathan Kadmon, Jeffrey Pennington, Sam S Schoenholz, Jascha Sohl-Dickstein, and Surya Ganguli. Statistical mechanics of deep learning. *Annual Review of Condensed Matter Physics*, 11:501–528, 2020.

[20] Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32, 2019.

[21] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. *arXiv preprint arXiv:1902.07208*, 2019.

[22] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.

[23] Pierre Baldi and Kurt Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.

[24] Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, pages 244–253. PMLR, 2018.

[25] Andrew M Saxe, James L McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116 (23):11537–11546, 2019.

[26] Alexander Atanasov, Blake Bordelon, and Cengiz Pehlevan. Neural networks as kernel learners: The silent alignment effect. *arXiv preprint arXiv:2111.00034*, 2021.

[27] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*, 2018.

[28] Simon S. Du and Wei Hu. Width provably matters in optimization for deep linear neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 1655–1664. PMLR, 2019. URL `http://proceedings.mlr.press/v97/du19a.html`.

[29] Tomaso Poggio, Qianli Liao, Brando Miranda, Andrzej Burbanski, and Jack Hidary. Theory iiib: Generalization in deep networks. Technical report, Center for Brains, Minds and Machines (CBMM), arXiv. org, 2018.

[30] Dongsung Huh. Curvature-corrected learning dynamics in deep neural networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 4552–4560. PMLR, 2020. URL `http://proceedings.mlr.press/v119/huh20a.html`.

[31] Thomas Laurent and James von Brecht. Deep linear networks with arbitrary loss: All local minima are global. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 2908–2913. PMLR, 2018. URL `http://proceedings.mlr.press/v80/laurent18a.html`.

[32] Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL `https://openreview.net/forum?id=HJflg30qKX`.

[33] Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 9482–9491, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/0e98aeeb54acf612b9eb4e48a269814c-Abstract.html.

[34] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. *Advances in Neural Information Processing Systems*, 32, 2019.

[35] Michael Biehl and Holm Schwarze. Learning by on-line gradient descent. *Journal of Physics A: Mathematical and general*, 28(3):643, 1995.

[36] David Saad and Sara A Solla. Exact solution for on-line learning in multilayer neural networks. *Physical Review Letters*, 74(21):4337, 1995.

[37] Sebastian Goldt, Madhu Advani, Andrew M Saxe, Florent Krzakala, and Lenka Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. *Advances in neural information processing systems*, 32, 2019.

[38] Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 8580–8589, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/5a4be1fa34e62bb8a6ec6b91d2462f5a-Abstract.html.

[39] Jaehoon Lee, Lechao Xiao, Samuel S. Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8570–8581, 2019. URL https://proceedings.neurips.cc/paper/2019/hash/0d1a9651497a38d8b1c3871c84528bd4-Abstract.html.

[40] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *Advances in Neural Information Processing Systems*, 32, 2019.

[41] Song Mei, Andrea Montanari, and Phan-minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings Of The National Academy Of Sciences*, 115(33):E7665–e7671, 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1806579115. URL https://www.pnas.org/content/115/33/E7665. Publisher: National Academy Of Sciences Section: Pnas Plus.

[42] Grant M. Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 7146–7155, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/196f5641aa9dc87067da4ff90fd81e7b-Abstract.html.

[43] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.

[44] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.

[45] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[46] Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, pages 5393–5402. PMLR, 2018.

[47] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.

[48] Raman Arora, SANJEEV Arora, Joan Bruna, NADAV Cohen, SIMON DU, RONG GE, SURIYA GUNASEKAR, C Jin, JASON LEE, TENGYU MA, et al. Theory of deep learning, 2020.

[49] Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4, 2008.

[50] Timo Flesch, Keno Juechems, Tsvetomira Dumbalska, Andrew Saxe, and Christopher Summerfield. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron*, 0(0), 2022. ISSN 0896-6273. doi: 10.1016/j.neuron.2022.01.005. URL `https://www.cell.com/neuron/abstract/S0896-6273(22)00005-8`. Publisher: Elsevier.

[51] Andrew M. Saxe, James L. Mcclelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings Of The National Academy Of Sciences*, 116 (23):11537–11546, 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1820226116.

[52] Dmytro Mishkin and Jiri Matas. All you need is a good init. *arXiv preprint arXiv:1511.06422*, 2015.

[53] Roger Ratcliff. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285, 1990.

[54] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.

[55] Haruka Asanuma, Shiro Takagi, Yoshihiro Nagano, Yuki Yoshida, Yasuhiko Igarashi, and Masato Okada. Statistical mechanical analysis of catastrophic forgetting in continual learning with teacher and student networks. *Journal of the Physical Society of Japan*, 90(10):104001, 2021.

[56] Thang Doan, Mehdi Abbana Bennani, Bogdan Mazoure, Guillaume Rabusseau, and Pierre Alquier. A theoretical analysis of catastrophic forgetting through the ntk overlap matrix. In *International Conference on Artificial Intelligence and Statistics*, pages 1072–1080. PMLR, 2021.

[57] Sebastian Lee, Sebastian Goldt, and Andrew Saxe. Continual learning in the teacher-student setup: Impact of task similarity. In *International Conference on Machine Learning*, pages 6109–6119. PMLR, 2021.

[58] Gal Shachaf, Alon Brutzkus, and Amir Globerson. A theoretical analysis of fine-tuning with linear teachers. *Advances in Neural Information Processing Systems*, 34, 2021.

[59] Nilesh Tripuraneni, Michael I Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. *arXiv preprint arXiv:2006.11650*, 2020.

[60] Sebastian Lee, Stefano Sarao Mannelli, Claudia Clopath, Sebastian Goldt, and Andrew Saxe. Maslow's hammer for catastrophic forgetting: Node re-use vs node activation, 2022. URL `https://arxiv.org/abs/2205.09029`.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See section 7
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See section G
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See section2
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]