
Dynamics of learning in deep linear neural networks

Andrew M. Saxe (asaxe@stanford.edu)

Department of Electrical Engineering

James L. McClelland (mcclelland@stanford.edu)

Department of Psychology

Surya Ganguli (sganguli@stanford.edu)

Department of Applied Physics

Stanford University, Stanford, CA 94305 USA

Abstract

Despite the widespread practical success of deep learning methods, our theoretical understanding of the dynamics of learning in deep neural networks remains quite sparse. We attempt to bridge the gap between the theory and practice of deep learning by systematically analyzing learning dynamics for the restricted case of deep linear neural networks. Despite the linearity of their input-output map, such networks have nonlinear gradient descent dynamics that change with the addition of each new hidden layer. We show that deep linear networks exhibit nonlinear learning phenomena similar to those seen in simulations of nonlinear networks, including long plateaus followed by rapid transitions to lower error solutions, and faster convergence from greedy unsupervised pretraining initial conditions than from random initial conditions. We provide an analytical description of these phenomena by finding new exact solutions to the nonlinear dynamics of deep learning. Our theoretical analysis also reveals the surprising finding that infinitely deep networks can be learned in finite time: for a special class of initial conditions on the weights, very deep networks incur only a finite delay in learning speed relative to shallow networks. We further show that, under certain conditions on the training data, unsupervised pre-training can find this special class of initial conditions, thereby providing analytical insight into the success of unsupervised pre-training in deep supervised learning tasks.

Deep learning methods have realized impressive performance in a range of applications, from visual object classification [1, 2] to speech recognition [3] and natural language processing [4, 5]. These successes have been achieved despite the noted difficulty of training such deep architectures [6, 7, 8, 9]. Indeed, many explanations for the difficulty of deep learning have been advanced in the literature, including the presence of many local minima, low curvature regions due to saturating nonlinearities, and exponential growth or decay of back-propagated gradients [10, 11, 12, 13]. Furthermore, many neural network simulations have observed strikingly nonlinear learning dynamics, including long plateaus of little apparent improvement followed by almost stage-like transitions to better performance. However, a quantitative, analytical understanding of the rich dynamics of deep learning remains elusive. For example, what determines the time scales over which deep learning unfolds? How does training speed retard with depth? Under what conditions will greedy unsupervised pretraining speed up learning? And how do the final learned internal representations depend on the statistical regularities inherent in the training data?

Here we provide an exact analytical theory of learning in deep linear neural networks that quantitatively answers these questions for this restricted setting. Because of its linearity, the input-output map of a deep linear network can always be rewritten as a shallow network. In this sense, a linear network does not gain expressive power from depth, and hence will underfit and perform poorly on complex real world problems. But while it lacks this important aspect of practical deep learning systems, a deep linear network can nonetheless exhibit highly nonlinear learning dynamics, and these dynamics change with increasing depth. Indeed, the training error, as a function of the network weights, is non-convex, and gradient descent dynamics on this non-convex error surface exhibits a subtle interplay between different weights across multiple layers of the network. Hence deep linear networks provide a useful starting point for understanding deep learning dynamics.

To answer these questions, we derive and analyze a set of nonlinear coupled differential equations describing learning dynamics on weight space as a function of the statistical structure of the inputs and outputs. We find exact time-dependent solutions to these nonlinear equations, as well as find conserved quantities in the weight dynamics arising from symmetries in the error function. These solutions provide intuition into how a deep network successively builds up information about the statistical structure of the training data and embeds this information into its weights and internal representations. Moreover, we compare our analytical solutions of learning dynamics in deep linear networks to numerical simulations of learning dynamics in deep non-linear networks, and find that our analytical solutions provide a reasonable approximation. Our solutions also reflect nonlinear phenomena seen in simulations, including alternating plateaus and sharp periods of rapid improvement. Indeed, it has been shown previously [14] that this nonlinear learning dynamics in deep linear networks is sufficient to qualitatively capture aspects of the progressive, hierarchical differentiation of conceptual structure seen in infant development. Finally, we apply these solutions to investigate the commonly used greedy layer-wise pretraining strategy for training deep networks [15, 16], and recover conditions under which such pretraining speeds learning. We show that these conditions are approximately satisfied for the MNIST dataset, and that unsupervised pretraining therefore confers an optimization advantage for deep linear networks applied to MNIST.

1 General learning dynamics of gradient descent

We begin by analyzing learning in a three layer network (input, hidden, and output) with linear activation functions (Fig 1), for which some theory has been developed previously [17, 18]. We let N_i be the number of neurons in layer i . The input-output map of the network is $y = W^{32}W^{21}x$. We wish to train the network to learn a particular input-output map from a set of P training examples $\{x^\mu, y^\mu\}$, $\mu = 1, \dots, P$. Training is accomplished via gradient descent on the squared error $\sum_{\mu=1}^P \|y^\mu - W^{32}W^{21}x^\mu\|^2$ between the desired feature output, and the network's feature output. This gradient descent procedure yields the batch learning rule

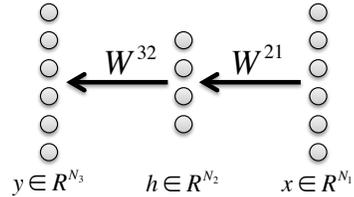


Figure 1: The three layer network analyzed in this section.

$$\Delta W^{21} = \lambda \sum_{\mu=1}^P W^{32T} (y^\mu x^{\mu T} - W^{32}W^{21}x^\mu x^{\mu T}), \quad \Delta W^{32} = \lambda \sum_{\mu=1}^P (y^\mu x^{\mu T} - W^{32}W^{21}x^\mu x^{\mu T}) W^{21T}, \quad (1)$$

where λ is a small learning rate. As long as λ is sufficiently small, we can take a continuous time limit to obtain the dynamics,

$$\tau \frac{d}{dt} W^{21} = W^{32T} (\Sigma^{31} - W^{32}W^{21}\Sigma^{11}), \quad \tau \frac{d}{dt} W^{32} = (\Sigma^{31} - W^{32}W^{21}\Sigma^{11}) W^{21T}, \quad (2)$$

where $\Sigma^{11} \equiv \sum_{\mu=1}^P x^\mu x^{\mu T}$ is an $N_1 \times N_1$ input correlation matrix, $\Sigma^{31} \equiv \sum_{\mu=1}^P y^\mu x^{\mu T}$ is an $N_3 \times N_1$ input-output correlation matrix, and $\tau \equiv \frac{P}{\lambda}$. Here t measures time in units of learning epochs; as t varies from 0 to 1, the network has seen P examples corresponding to one learning epoch. Despite the linearity of the network's input-output map, the gradient descent learning dynamics given in Eqn (2) constitutes a complex set of coupled nonlinear differential equations with up to cubic interactions in the weights.

1.1 Learning dynamics with orthogonal inputs

Our fundamental goal is to understand the dynamics of learning in (2) as a function of the input statistics Σ^{11} and input-output statistics Σ^{31} . In general, the outcome of learning will reflect an interplay between input correlations, described by Σ^{11} , and the input-output correlations described by Σ^{31} . To begin, though, we further simplify the analysis by focusing on the case of orthogonal input representations where $\Sigma^{11} = I$. This assumption will hold exactly for whitened input data, a widely used preprocessing step. We further focus on the case in which $N_3 > N_1$, so that there are more output units than input units.

Because we have assumed orthogonal input representations ($\Sigma^{11} = I$), the input-output correlation matrix contains all of the information about the dataset used in learning, and it plays a pivotal role in the learning dynamics. We consider its singular value decomposition (SVD)

$$\Sigma^{31} = U^{33} S^{31} V^{11 T} = \sum_{\alpha=1}^{N_1} s_\alpha u_\alpha v_\alpha^T, \quad (3)$$

which will be central in our analysis. Here V^{11} is an $N_1 \times N_1$ orthogonal matrix whose columns contain *input-analyzing* singular vectors v_α that reflect independent modes of variation in the input, U^{33} is an $N_3 \times N_3$ orthogonal matrix whose columns contain *output-analyzing* singular vectors u_α that reflect independent modes of variation in the output, and S^{31} is an $N_3 \times N_1$ matrix whose only nonzero elements are on the diagonal; these elements are the singular values $s_\alpha, \alpha = 1, \dots, N_1$ ordered so that $s_1 \geq s_2 \geq \dots \geq s_{N_1}$.

Now, performing the change of variables on synaptic weight space, $W^{21} = \overline{W}^{21} V^{11 T}$, $W^{32} = U^{33} \overline{W}^{32}$, the dynamics in (2) simplify to

$$\tau \frac{d}{dt} \overline{W}^{21} = \overline{W}^{32 T} (S^{31} - \overline{W}^{32} \overline{W}^{21}), \quad \tau \frac{d}{dt} \overline{W}^{32} = (S^{31} - \overline{W}^{32} \overline{W}^{21}) \overline{W}^{21 T}. \quad (4)$$

To gain intuition for these equations, note that while the matrix elements of W^{21} and W^{32} connected neurons in one layer to neurons in the next layer, we can think of the matrix element $\overline{W}^{21}_{i\alpha}$ as connecting input mode v_α to hidden neuron i , and the matrix element $\overline{W}^{32}_{\alpha i}$ as connecting hidden neuron i to output mode u_α . Let a^α be the α^{th} column of W^{21} , and let $b^{\alpha T}$ be the α^{th} row of \overline{W}^{32} . Intuitively, a^α is a column vector of N_2 synaptic weights presynaptic to the hidden layer coming from input mode α , and b^α is a column vector of N_2 synaptic weights postsynaptic to the hidden layer going to output mode α . In terms of these variables, or connectivity modes, the learning dynamics in (4) become

$$\tau \frac{d}{dt} a^\alpha = (s_\alpha - a^\alpha \cdot b^\alpha) b^\alpha - \sum_{\gamma \neq \alpha} b^\gamma (a^\alpha \cdot b^\gamma), \quad \tau \frac{d}{dt} b^\alpha = (s_\alpha - a^\alpha \cdot b^\alpha) a^\alpha - \sum_{\gamma \neq \alpha} a^\gamma (b^\alpha \cdot a^\gamma). \quad (5)$$

Note that $s_\alpha = 0$ for $\alpha > N_1$. These dynamics arise from gradient descent on the energy function

$$E = \frac{1}{2\tau} \sum_{\alpha} (s_\alpha - a^\alpha \cdot b^\alpha)^2 + \frac{1}{2\tau} \sum_{\alpha \neq \beta} (a^\alpha \cdot b^\beta)^2, \quad (6)$$

and display an interesting combination of cooperative and competitive interactions. Consider the first terms in each equation. In these terms, the connectivity modes from the two layers, a^α and b^α associated with the same input-output mode of strength s^α , cooperate with each other to drive each other to larger magnitudes

as well as point in similar directions in the space of hidden units; in this fashion these terms drive the product of connectivity modes $a^\alpha \cdot b^\alpha$ to reflect the input-output mode strength s^α . The second terms describe competition between the connectivity modes in the first (a^α) and second (b^β) layers associated with different input modes α and β . This yields a symmetric, pairwise repulsive force between all distinct pairs of first and second layer connectivity modes, driving the network to a decoupled regime in which the different connectivity modes become orthogonal.

1.2 The final outcome of learning

The fixed point structure of gradient descent learning was worked out in [17]. In the language of the connectivity modes, a necessary condition for a fixed point is $a^\alpha \cdot b^\beta = s_\alpha \delta_{\alpha\beta}$, while a^α and b^α are zero whenever $s_\alpha = 0$. To satisfy these relations, a^α and b^α can be nonzero for at most N_2 values of α . Since there are N_1 nonzero values of s_α , there are $\binom{N_1}{N_2}$ families of fixed points. However, all of these fixed points are unstable, except for the one in which only the first N_2 strongest modes, i.e. a^α and b^α for $\alpha = 1, \dots, N_2$ are active. Thus remarkably, the dynamics in (5) has only saddle points and no non-global local minima. In terms of the original synaptic variables W^{21} and W^{32} , all globally stable fixed points satisfy

$$W^{32}W^{21} = \sum_{\alpha=1}^{N_2} s_\alpha u_\alpha v_\alpha^T. \quad (7)$$

Hence when learning has converged, the network will represent the closest rank N_2 approximation to the true input-output correlation matrix. In this work, we are interested in understanding the dynamical weight trajectories and learning time scales that lead to this final fixed point.

1.3 The time course of learning

It is difficult though to exactly solve (5) starting from arbitrary initial conditions because of the competitive interactions between different input-output modes. Therefore, to gain intuition for the general dynamics, we restrict our attention to a special class of initial conditions of the form a^α and $b^\alpha \propto r^\alpha$ for $\alpha = 1, \dots, N_2$, where $r^\alpha \cdot r^\beta = \delta_{\alpha\beta}$, with all other connectivity modes a^α and b^α set to zero (see [18] for solutions to a partially overlapping but distinct set of initial conditions). Here r^α is a fixed collection of N_2 vectors that form an orthonormal basis for synaptic connections from an input or output mode onto the set of hidden units. Thus for this set of initial conditions, a^α and b^α point in the same direction for each alpha and differ only in their scalar magnitudes, and are orthogonal to all other connectivity modes. It is straightforward to verify that starting from these initial conditions, a^α and b^α will remain parallel to r^α for all future time. Furthermore, because the different active modes are orthogonal to each other, they do not compete, or even interact with each other (all dot products in the second terms of (5)-(6) are 0). Thus this class of conditions defines an invariant manifold in weight space where the modes evolve independently of each other.

If we let $a = a^\alpha \cdot r^\alpha$, $b = b^\alpha \cdot r^\alpha$, and $s = s^\alpha$, then the dynamics of the scalar projections (a, b) obeys,

$$\tau \frac{d}{dt} a = b(s - ab), \quad \tau \frac{d}{dt} b = a(s - ab). \quad (8)$$

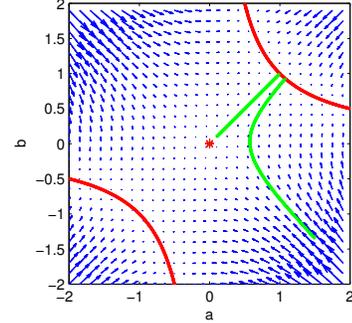


Figure 2: Vector field (blue), stable manifold (red) and two solution trajectories (green) for the two dimensional dynamics of a and b in (8), with $\tau = 1, s = 1$.

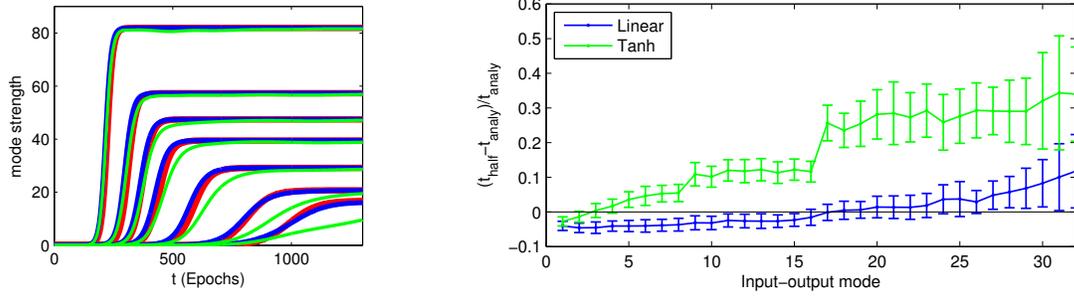


Figure 3: **Left:** Dynamics of learning in a three layer neural network. Curves show the strength of the network’s representation of seven modes of the input-output correlation matrix over the course of learning. Red traces show analytical curves from Eqn. 12. Blue traces show simulation of full dynamics of a linear network (Eqn. (2)) from small random initial conditions. Green traces show simulation of a nonlinear three layer network with tanh activation functions. To generate mode strengths for the nonlinear network, we computed the nonlinear network’s evolving input-output correlation matrix, and plotted the diagonal elements of $U^{33T} \Sigma_{\tanh}^{31} V^{11}$ over time. The training set consists of 32 orthogonal input patterns, each associated with a 1000-dimensional feature vector generated by a hierarchical diffusion process described in [14] with a five level binary tree and flip probability of 0.1. Modes 1, 2, 3, 5, 12, 18, and 31 are plotted with the rest excluded for clarity. Network training parameters were $\lambda = 0.5e^{-3}$, $N_2 = 32$, $u_0 = 1e^{-6}$. **Right:** Delay in learning due to competitive dynamics and sigmoidal nonlinearities. Vertical axis shows the difference between simulated time of half learning and the analytical time of half learning, as a fraction of the analytical time of half learning. Error bars show standard deviation from 100 simulations with random initializations.

Thus our ability to decouple the connectivity modes yields a dramatically simplified two dimensional non-linear system. These equations can be solved by noting that they arise from gradient descent on the error,

$$E(a, b) = \frac{1}{2\tau}(s - ab)^2. \quad (9)$$

This implies that the product ab monotonically approaches the fixed point s from its initial value. Moreover, $E(a, b)$ satisfies a symmetry under the one parameter family of scaling transformations $a \rightarrow \lambda a$, $b \rightarrow \frac{b}{\lambda}$. This symmetry implies, through Noether’s theorem, the existence of a conserved quantity, namely $a^2 - b^2$, which is a constant of motion. Thus the dynamics simply follows hyperbolas of constant $a^2 - b^2$ in the (a, b) plane until it approaches the hyperbolic manifold of fixed points, $ab = s$. The origin $a = 0, b = 0$ is also a fixed point, but is unstable. Fig. 2 shows a typical phase portrait for these dynamics.

As a measure of the timescale of learning, we are interested in how long it takes for ab to approach s from any given initial condition. The case of unequal a and b is treated in the Supplementary Materials due to space constraints. Here we pursue an explicit solution with the assumption that $a = b$, a reasonable limit when starting with small random initial conditions. We can then track the dynamics of $u \equiv ab$, which from (8) obeys

$$\tau \frac{d}{dt} u = 2u(s - u). \quad (10)$$

This equation is separable and can be integrated to yield

$$t = \tau \int_{u_0}^{u_f} \frac{du}{2u(s - u)} = \frac{\tau}{2s} \ln \frac{u_f(s - u_0)}{u_0(s - u_f)}. \quad (11)$$

Here t is the time it takes for u to travel from u_0 to u_f . If we assume a small initial condition $u_0 = \epsilon$, and ask when u_f is within ϵ of the fixed point s , i.e. $u_f = s - \epsilon$, then the learning timescale in the limit $\epsilon \rightarrow 0$ is $t = \tau/s \ln(s/\epsilon) = O(\tau/s)$ (with a weak logarithmic dependence on the cutoff). This yields a key result: the timescale of learning of each input-output mode α of the correlation matrix Σ^{31} is inversely proportional

to the correlation strength s_α of the mode. Thus the stronger an input-output relationship, the quicker it is learned.

We can also find the entire time course of learning by inverting (11) to obtain

$$u_f(t) = \frac{se^{2st/\tau}}{e^{2st/\tau} - 1 + s/u_0}. \quad (12)$$

This time course describes the temporal evolution of the product of the magnitudes of all weights from an input mode (with correlation strength s) into the hidden layers, and from the hidden layers to the same output mode. If this product starts at a small value $u_0 < s$, then it displays a sigmoidal rise which asymptotes to s as $t \rightarrow \infty$. This sigmoid can exhibit sharp transitions from a state of no learning to full learning. This analytical sigmoid learning curve is shown in Fig. 3 to yield a reasonable approximation to learning curves in linear networks that start from random initial conditions that are not on the orthogonal, decoupled invariant manifold—and that therefore exhibit competitive dynamics between connectivity modes—as well as in nonlinear networks solving the same task. We note that though the nonlinear networks behaved similarly to the linear case for this particular task, this is likely to be problem dependent.

2 Deeper multilayer dynamics

The network analyzed in Section 1 is the minimal example of a multilayer net, with just a single layer of hidden units. How does gradient descent act in much deeper networks? We make an initial attempt in this direction based on initial conditions that yield particularly simple gradient descent dynamics.

In a linear neural network with N_l layers and hence $N_l - 1$ weight matrices indexed by $W^l, l = 1, \dots, N_l - 1$, the gradient descent dynamics can be written as

$$\tau \frac{d}{dt} W^l = \left(\prod_{i=l+1}^{N_l-1} W^i \right)^T \left[\Sigma^{31} - \left(\prod_{i=1}^{N_l-1} W^i \right) \Sigma^{11} \right] \left(\prod_{i=1}^{l-1} W^i \right)^T, \quad (13)$$

where $\prod_{i=a}^b W^i = W^b W^{(b-1)} \dots W^{(a-1)} W^a$ with the caveat that $\prod_{i=a}^b W^i = I$, the identity, if $a > b$.

To describe the initial conditions, we suppose that there are N_l orthogonal matrices R_l that diagonalize the starting weight matrices, that is, $R_{l+1}^T W_l(0) R_l = D_l$ for all l , with the caveat that $R_1 = V^{11}$ and $R_{N_l} = U^{33}$. This requirement essentially demands that the output singular vectors of layer l be the input singular vectors of the next layer $l + 1$, so that a change in mode strength at any layer propagates to the output without mixing into other modes. Making the change of variables $W_l = R_{l+1} \bar{W}_l R_l^T$ along with the assumption that $\Sigma^{11} = I$ leads to a set of decoupled connectivity modes that evolve independently of each other. In analogy to the simplification occurring in the three layer network from (2) to (8), each connectivity mode in the N_l layered network can be described by $N_l - 1$ scalars a^1, \dots, a^{N_l-1} , whose dynamics obeys gradient descent on the energy function (the analog of (9)),

$$E(a_1, \dots, a_{N_l-1}) = \frac{1}{2\tau} \left(s - \prod_{i=1}^{N_l-1} a_i \right)^2. \quad (14)$$

This dynamics also has a set of conserved quantities $a_i^2 - a_j^2$ arising from the energetic symmetry w.r.t. the transformation $a_i \rightarrow \lambda a_i, a_j \rightarrow \frac{a_j}{\lambda}$, and hence can be solved exactly. We focus on the invariant submanifold in which $a_i(t = 0) = a_0$ for all i , and track the dynamics of $u = \prod_{i=1}^{N_l-1} a_i$, the overall strength of this mode, which obeys (i.e. the generalization of (10)),

$$\tau \frac{d}{dt} u = (N_l - 1) u^{2-2/(N_l-1)} (s - u). \quad (15)$$

This can be integrated for any positive integer N_l , though the expression is complicated. Once the overall strength increases sufficiently, learning explodes rapidly.

Eqn. (15) lets us study the dynamics of learning in an infinitely deep network. In particular, as $N_l \rightarrow \infty$ we have the dynamics

$$\tau \frac{d}{dt} u = N_l u^2 (s - u) \quad (16)$$

which can be integrated to obtain

$$t = \frac{\tau}{N_l} \left[\frac{1}{s^2} \log \left(\frac{u_f(u_0 - s)}{u_0(u_f - s)} \right) + \frac{1}{su_0} - \frac{1}{su_f} \right]. \quad (17)$$

Remarkably this implies that, for a fixed learning rate, the learning time tends to zero as N_l goes to infinity. This result depends on the continuous time formulation, however. Any implementation will operate in discrete time and must choose a finite learning rate that yields stable dynamics. An estimate of the optimal learning rate can be derived from the maximum eigenvalue of the Hessian over the region of interest. For linear networks with $a_i = a_j = a$, this optimal learning rate α_{opt} decays with depth as $O\left(\frac{1}{N_l s^2}\right)$ for large N_l (see supplementary material). Incorporating this dependence of the learning rate on depth, the learning time in an infinitely deep network still surprisingly remains finite: with the optimal learning rate, the difference between learning times for an $N_l = 3$ network and an $N_l = \infty$ network is $t_\infty - t_3 \sim O(s/\epsilon)$ for small ϵ (see supplementary material).

To verify these predictions, we trained deep linear networks on the MNIST dataset with depths ranging from $N_l = 3$ to $N_l = 100$. We used hidden layers of size 1000, and calculated the epoch at which training error fell below a fixed threshold of 1.3×10^4 corresponding to nearly complete learning. We optimized the learning rate separately for each

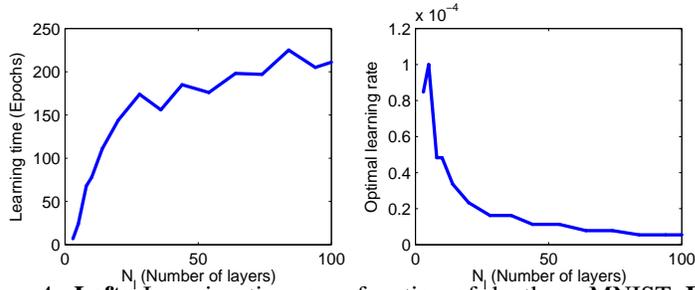


Figure 4: **Left:** Learning time as a function of depth on MNIST. **Right:** Empirically optimal learning rates as a function of depth.

depth by training each network with twenty rates logarithmically spaced between 10^{-4} and 10^{-7} and picking the fastest. Networks were initialized with decoupled initial conditions and starting initial mode strength $u_0 = 0.001$. Fig. 4 shows the resulting learning times, which saturate, and the empirically optimal learning rates, which scale like $O(1/N_l)$ as predicted.

Thus learning times in deep linear networks that start with decoupled initial conditions are only a finite amount slower than a shallow network regardless of depth. Moreover, the delay incurred by depth scales inversely with the size of the initial strength of the association. Hence finding a way to initialize the mode strengths to large values is crucial for fast deep learning.

3 Efficacy of greedy unsupervised pretraining

The breakthrough in training deep neural networks started with the discovery that greedy layer-wise unsupervised pre-training could substantially speed up and improve the generalization performance of standard gradient descent [15, 16]. Unsupervised pre-training has been shown to help with the optimization of deep networks, and also to act as a special regularizer towards solutions with better generalization performance

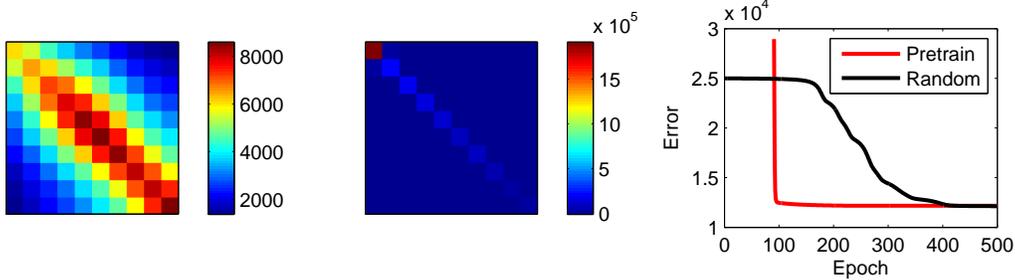


Figure 5: MNIST satisfies the consistency condition for greedy pretraining. **Left:** Submatrix from the raw MNIST input correlation matrix Σ^{11} . **Center:** Submatrix of $V^{11}\Sigma^{11}V^{11T}$ which is approximately diagonal as required. **Right:** Learning curves on MNIST for a five layer linear network starting from random (black) and pretrained (red) initial conditions. Pretrained curve starts with a delay due to pretraining time.

[16, 10, 11, 12]. Here we show analytically how unsupervised pre-training achieves an optimization advantage, at least, in deep linear networks by finding the special class of orthogonalized initial conditions in the previous section that allow for rapid supervised deep learning, for input-output tasks with a certain precise structure.

In particular, we consider the effect of using autoencoders as the unsupervised pretraining module [16, 10], for which the input-output correlation matrix Σ^{31} is simply the input correlation matrix Σ^{11} . Hence the SVD of Σ^{31} is PCA on the input correlation matrix, since $\Sigma^{31} = \Sigma^{11} = Q\Lambda Q^T$, where Q are eigenvectors of Σ^{11} and Λ is a diagonal matrix of variances. After pretraining, the weights thus converge to $W^{32}W^{21} = Q\Lambda Q^T$. Recall that a^α denotes the strength of mode α in W^{21} , and b^α denotes the strength in W^{32} . If we further take the assumption that $a^\alpha \approx b^\alpha$, as is typical when starting from small random weights, then the input-to-hidden mapping will be $W^{21} = R_2\sqrt{\Lambda}Q^T$ where R_2 is an arbitrary orthogonal matrix. Now consider fine-tuning on a task with input-output correlations $\Sigma^{31} = U^{33}S^{31}V^{11}$. The pretrained initial condition $W^{21} = R_2\sqrt{\Lambda}Q^T$ will be a decoupled initial condition for the task, $W^{21} = R_2D_1V^{11T}$, provided

$$Q = V^{11}. \quad (18)$$

Hence we can state the underlying condition required for successful greedy pretraining in deep linear networks: the right singular vectors of the ultimate input-output task of interest V^{11} must be similar to the principal components of the input data Q . This gives a simple empirical criterion that can be evaluated on any new dataset: given Σ^{31} and Σ^{11} , compute the right singular vectors V^{11} and check that $V^{11}\Sigma^{11}V^{11T}$ is approximately diagonal. If the condition in Eqn. (18) holds, autoencoder pretraining will have properly set up decoupled initial conditions for W^{21} , with an appreciable initial association strength of $\sqrt{\Lambda}$. This argument also goes through straightforwardly for layer-wise pretraining of deeper networks. Fig. 5 shows that this consistency condition empirically holds on MNIST, and that a pretrained deep linear neural network learns faster than one started from random initial conditions, even accounting for pretraining time.

4 Discussion

Despite the simplicity of their input-output map, the dynamics of learning in deep linear networks reveals a surprising amount of rich mathematical structure, including nonlinear hyperbolic dynamics, plateaus and sudden performance transitions, and most importantly, a sensitive but computable dependence of learning time scales on input statistics, initial conditions, and network depth. With the right initial conditions, deep linear networks can be only a finite amount slower than shallow networks, and unsupervised pretraining can find these initial conditions for tasks with the right structure. At the cost of expressivity, deep linear networks gain tractability and may prove fertile for addressing other phenomena in deep learning, such as the impact of carefully scaled initializations [11, 19], momentum [19], dropout regularization [1], and sparsity constraints [2]. Although extensions to deep nonlinear networks remain daunting, our work fulfills an important pre-requisite for progress towards a more general quantitative theory of deep learning.

References

- [1] A. Krizhevsky and G.E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, 2012.
- [2] Q.V. Le, M.A. Ranzato, R. Monga, M. Devin, K. Chen, G.S. Corrado, J. Dean, and A.Y. Ng. Building high-level features using large scale unsupervised learning. In *29th International Conference on Machine Learning*, 2012.
- [3] A. Mohamed, G.E. Dahl, and G. Hinton. Acoustic Modeling Using Deep Belief Networks. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(1):14–22, January 2012.
- [4] R. Collobert and J. Weston. A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning. In *Proceedings of the 25th International Conference on Machine Learning*, 2008.
- [5] R. Socher, J. Bauer, C.D. Manning, and A.Y. Ng. Parsing with Compositional Vector Grammars. In *Association for Computational Linguistics Conference*, 2013.
- [6] Y. Bengio, P. Simard, and P. Frasconi. Learning Long-Term Dependencies with Gradient Descent is Difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
- [7] Y. LeCun, L. Bottou, G.B. Orr, and K.R. Müller. Efficient BackProp. *Neural networks: Tricks of the trade*, 1998.
- [8] Y. Bengio and Y. LeCun. Scaling learning algorithms towards AI. In L. Bottou, O. Chapelle, D. De-Coste, and J. Weston, editors, *Large-Scale Kernel Machines*, number 1, pages 1–41. MIT Press, 2007.
- [9] D. Erhan, P.A. Manzagol, Y. Bengio, S. Bengio, and P. Vincent. The Difficulty of Training Deep Architectures and the Effect of Unsupervised Pre-Training. In *12th International Conference on Artificial Intelligence and Statistics*, volume 5, 2009.
- [10] Y. Bengio. Learning Deep Architectures for AI. 2009.
- [11] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. *13th International Conference on Artificial Intelligence and Statistics*, 2010.
- [12] D. Erhan, Y. Bengio, A. Courville, P.A. Manzagol, and P. Vincent. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11:625–660, 2010.
- [13] Y.N. Dauphin and Y. Bengio. Big Neural Networks Waste Capacity. In *International Conference on Learning Representations*, 2013.
- [14] A.M. Saxe, J.L. McClelland, and S. Ganguli. Learning hierarchical category structure in deep neural networks. In *Proceedings of the 35th Annual Conference of the Cognitive Science Society*, 2013.
- [15] G.E. Hinton and R.R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–7, July 2006.
- [16] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy Layer-Wise Training of Deep Networks. *Advances in Neural Information Processing Systems 20*, 2007.
- [17] P. Baldi and K. Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural Networks*, 2(1):53–58, January 1989.
- [18] 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.
- [19] I. Sutskever, J. Martens, G. Dahl, and G.E. Hinton. On the importance of initialization and momentum in deep learning. In *30th International Conference on Machine Learning*, 2013.