



A Note on Lazy Training in Supervised Differentiable Programming

Lenaïc Chizat, Francis Bach

► To cite this version:

Lenaïc Chizat, Francis Bach. A Note on Lazy Training in Supervised Differentiable Programming. 2019. hal-01945578v3

HAL Id: hal-01945578

<https://inria.hal.science/hal-01945578v3>

Preprint submitted on 21 Feb 2019 (v3), last revised 7 Jan 2020 (v6)

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.

A Note on Lazy Training in Supervised Differentiable Programming

Lénaïc Chizat and Francis Bach
INRIA, ENS, PSL Research University
Paris, France

`lenaic.chizat@inria.fr`, `francis.bach@inria.fr`

February 21, 2019

Abstract

In a series of recent theoretical works, it has been shown that strongly over-parameterized neural networks trained with gradient-based methods could converge linearly to zero training loss, with their parameters hardly varying. In this note, our goal is to exhibit the simple structure that is behind these results. In a simplified setting, we prove that “lazy training” essentially solves a kernel regression. We also show that this behavior is not so much due to over-parameterization than to a choice of scaling, often implicit, that allows to linearize the model around its initialization. These theoretical results complemented with simple numerical experiments make it seem unlikely that “lazy training” is behind the many successes of neural networks in high dimensional tasks.

1 Introduction

Differentiable programming is becoming an important paradigm in certain areas of signal processing and machine learning that consists in building parameterized models, sometimes with a complex architecture and a large number of parameters, and adjusting these parameters so that the model fits training data using gradient-based optimization methods. The resulting problem is in general highly non-convex: it has been observed empirically that seemingly innocuous changes in the parameterization, optimization procedure, or in the initialization could lead to a selection of very different models, even though they sometimes all fit perfectly the training data [33]. Our goal is to showcase this effect by studying *lazy training*, which refers to training algorithms that select parameters that are close to the initialization.

This note is motivated by a series of recent articles [13, 21, 12, 2, 3, 35]¹ where it is shown that certain over-parameterized neural networks converge linearly to zero training loss with their parameters hardly varying. With a slightly different viewpoint, it was shown in [16] that the first phase of training behaves like a kernel regression in the over-parameterization limit, with a kernel built from the linearization of the neural network around its initialization. Blending these two points of views together, we remark that lazy training essentially amounts to kernel-based regression with a specific kernel. Importantly, in all these papers, a specific and somewhat implicit choice of scaling is made. We argue that lazy training is not so much due to over-parameterization than to this choice of scaling. By introducing a scale factor $\alpha > 0$, we see that essentially any

¹Other works in this line of research that have appeared after a first version of this note was communicated include [4, 26, 6].

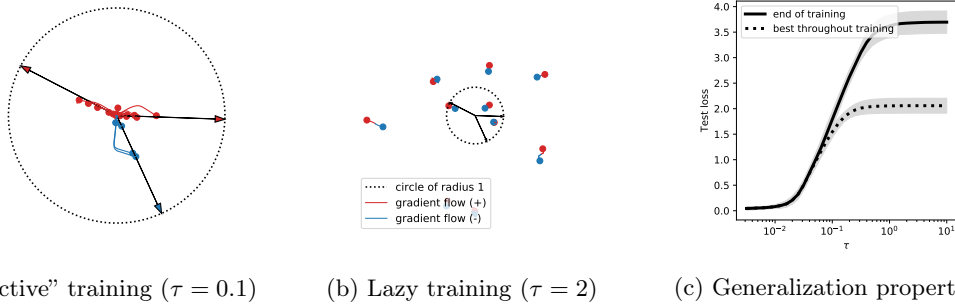


Figure 1: Training a 2-layers ReLU neural network initialized with normal random weights of variance τ^2 : lazy training occurs when τ is large. (a)-(b) Trajectory of weights during gradient descent in 2-d (color shows sign of output layer). (c) Generalization in 100-d: worsens as τ increases. Ground truth is generated with 3 neurons (arrows in (a)-(b)). Details in Section 5.

parametric model can be trained in this lazy regime if initialized close enough to 0 in the space of predictors. This remark allows to better understand lazy training: its generalization properties, when it occurs, and its downsides.

The takeaway is that guaranteed fast training is indeed possible, but at the cost of recovering a linear method². On the upside, as noticed in [16], this draws an interesting link between neural networks and kernel methods, which rely on an established theory with many performance guaranties. On the downside, we believe that the most competitive neural networks are not trained in this regime: for those, neurons weights seem to change significantly (at least in some layers [34]), and the first layer of convolutional neural networks tend to learn Gabor-like filters when randomly initialized [14, Chap. 9]. Instead, they seem to be able to perform high dimensional, non-linear, feature selection and this remains a fundamental theoretical challenge to understand how and when they do so.

The situation is illustrated in Figure 1 where lazy training for a two-layer neural network with rectified linear units (ReLU) is achieved by increasing the variance τ^2 at initialization (see Section 4). While in panel (a) the ground truth features are identified, this is not the case for lazy training on panel (b) that tries to interpolate the observations with the least displacement in parameter space (in both cases, near 0 training loss was achieved). As seen on panel (c), this behavior hinders good generalization. The plateau reached for large τ corresponds exactly to the performance of the corresponding kernel method.

Setting of supervised differentiable programming. In this note, a *model* is a black box function that maps an input variable $x \in \mathbb{R}^d$ to an output variable³ $y \in \mathbb{R}$ in a way that is consistent with observations. A common way to build models is to use a *parametric model*, that is a function $f : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}$ associated with a *training algorithm* designed to select a parameter $w \in \mathbb{R}^p$ given a sequence of n observations $(x_i, y_i) \in \mathbb{R}^d \times \mathbb{R}$, $i \in \{1, \dots, n\}$. When the function f is differentiable in the parameters (at least in a weak sense), most training algorithms are *gradient-based*. They consist in choosing a smooth and often strongly-convex *loss function* $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$, a (possibly random) initialization $w_0 \in \mathbb{R}^p$ and defining recursively a sequence of parameters w_0, w_1, \dots by accessing the model only through the gradient of $\ell(f(w_k, x_i), y_i)$ with

²Here we mean a prediction function linearly parameterized by a potentially infinite-dimensional vector.

³The case $y \in \mathbb{R}^k$ with $k > 1$ adds no difficulty but makes notations more complicated.

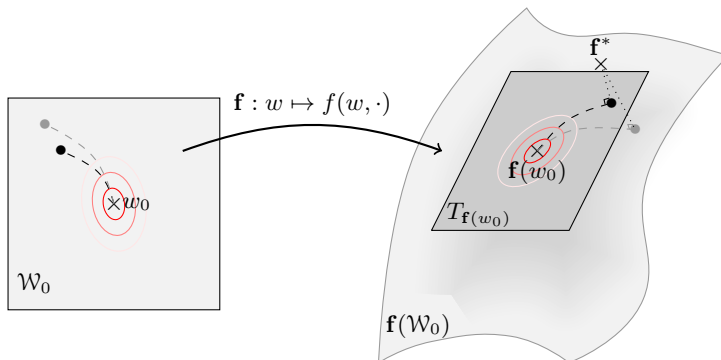


Figure 2: There is a small neighborhood $\mathcal{W}_0 \subset \mathbb{R}^p$ of the initialization w_0 , which image by $\mathbf{f} : w \mapsto f(w, \cdot)$ is a differentiable manifold in the space of predictors. Lazy training (dashed gray paths) and training a linear model in the tangent space $T_{\mathbf{f}(w_0)}$ (dashed black paths) are essentially equivalent if the later converges fast enough (Section 3). This figure illustrates the under-parameterized case as here the optimal model \mathbf{f}^* does not belong to $T_{\mathbf{f}(w_0)}$.

respect to w . For instance, stochastic gradient descent (SGD) updates the parameters as follows

$$w_{k+1} = w_k - \eta_k \ell'(f(w_k, x_k), y_k) \nabla f(w_k, x_k), \quad (1)$$

where ℓ' and ∇f are the gradients with respect to the first arguments and η_0, η_1, \dots , a specified sequence of step-sizes. When the model is linear in the parameters, or in other specific cases [11, 22], such methods are known to find models which are optimal in some statistical sense, but this property is not well understood in the general case.

1.1 Content of this note

Lazy training can be introduced as follows. We start by looking at the SGD update of Eq. (1) in the space of predictors: for $x \in \mathbb{R}^d$ and small η_k , one has the first order approximation

$$f(w_{k+1}, x) \approx f(w_k, x) - \eta_k \nabla f(w_k, x)^\top \nabla f(w_k, x_k) \ell'(f(w_k, x_k), y_k).$$

This is an SGD update for unregularized kernel-based regression [17, 10] with kernel $k(x, x') = \nabla f(w_k, x)^\top \nabla f(w_k, x')$. The key point is that if the iterates remain in a neighborhood of w_0 then this kernel is roughly constant throughout training. When $f(w_0, \cdot) \approx 0$, this behavior naturally arises when scaling the model as αf with a large scaling factor $\alpha > 0$. Indeed, this scaling does not change the tangent model and brings the iterates of SGD closer to w_0 by a factor $1/\alpha$. This scaling is not artificial: rather, it is often implicit in practice (for instance, hidden in the choice of initialization, see Section 4). Another depiction of lazy training with a geometrical point of view is given in Figure 2. There, α can be interpreted as a way to stretch the manifold $\mathbf{f}(\mathcal{W}_0)$ to bring it closer to $T_{\mathbf{f}(w_0)}$.

This note is organized as follows:

- in Section 2, we describe the tangent model introduced in [16] and detail the case of neural networks with a single hidden layer where the tangent model turns out to be equivalent to a random feature method [27];
- in Section 3, we give simple proofs that gradient flows in the lazy regime converge linearly, to a global minimizer for over-parameterized models, or to a local minimizer for under-

parameterized models. We also prove that they are identical to gradient flows associated to the tangent model, up to higher order terms in the scaling factor α .

- in Section 4, we emphasize that lazy training is just a specific regime: it occurs for a specific range of initialization or hyper-parameters. We give criteria to check whether a given parametric model is likely to exhibit this behavior;
- finally, Section 5 shows simple numerical experiments on synthetic cases to illustrate how lazy training differs from other regimes of training (see also Figure 1).

The main motivation for this note is to present in a simple setting the phenomenon underlying a series of recent results [16, 13, 21, 12, 2, 3, 35] and to emphasize that they (i) are not fundamentally related to over-parameterization nor to specific neural networks architectures, and (ii) correspond however to a very specific training regime which is likely not the most useful in practice. Our focus is on general principles and qualitative description so we make the simplifying assumption that f is differentiable⁴ in w and we sometimes provide statements without constants.

2 Training the tangent model

2.1 Tangent model

In first order approximation around the initial parameters $w_0 \in \mathbb{R}^p$, the parametric model $f(w, x)$ reduces to the following *tangent model* :

$$T_f(w, x) = f(w_0, x) + (w - w_0) \cdot \nabla_w f(w_0, x). \quad (2)$$

The corresponding hypothesis class is *affine* in the space of predictors. It should be stressed that when f is a neural network, T_f is generally not a linear neural network because it is not linear in $x \in \mathbb{R}^d$, but in the features $\nabla_w f(w_0, x) \in \mathbb{R}^p$ which generally depend non-linearly on x . For large neural networks, the dimension of the features might be much larger than d , which makes T_f similar to non-parametric methods. Finally, if f is already a linear model, then f and T_f are identical.

Kernel method with an offset. In the case where $\nabla \ell(y, y')$ only depends on $y - y'$, such as the quadratic loss, training the affine model (2) is equivalent to training a linear model in the variables

$$(\tilde{x}, \tilde{y}) := (\nabla_w f(w_0, x), y - f(w_0, x)).$$

This is equivalent to a kernel method with the *tangent kernel* (see [16]):

$$k(x, x') = \nabla_w f(w_0, x)^\top \nabla_w f(w_0, x'). \quad (3)$$

This kernel is different from the one generally associated to neural networks [28, 9] which involve the derivative with respect to the output layer only. Also, the output data is shifted by the initialization of the model $f(w_0, x)$. This term inherits from the randomness due to the initialization: it is for instance shown in [19, 24] that $f(w_0, \cdot)$ converges to a Gaussian process for certain over-parameterized neural networks initialized with random normal weights. For neural networks, we can make sure that $f(w_0, \cdot) = 0$ even with a random initialization by using a simple “doubling trick”: neurons in the last layer are duplicated, with the new neurons having the same input weights and opposite output weights.

⁴This assumption is relevant to some extent, as it is morally true for large networks or with a large amount of data (see, e.g., [7, App. D.4]).

Algorithm 1 SGD (lazy when $\alpha \gg 1$)

1. set $w_0 \in \mathbb{R}^p$
2. for $k \in \{1, \dots, n\}$ do
 $w_{k+1} = w_k - \frac{\eta_k}{\alpha} \ell'(\alpha f(w_k, x_k), y_k) \nabla f(w_k, x_k)$

Algorithm 2 SGD on the tangent model

1. set $w_0 \in \mathbb{R}^p$
2. for $k \in \{1, \dots, n\}$ do
 $w_{k+1} = w_k - \frac{\eta_k}{\alpha} \ell'(\alpha T_f(w_k, x_k), y_k) \nabla f(w_0, x_k)$

Figure 3: Stochastic Gradient Descent (SGD) for the parametric model f and the tangent model T_f of Eq. (2), with step-size sequence $(\eta_k)_{k \geq 0}$ and with a scaling $\alpha > 0$. In Section 3, we show that these algorithms are essentially equivalent when α is large and the step-sizes η_k are small.

Computational differences. We write on Figure 3 the SGD algorithm for the parametric model f and for the tangent model T_f in order to highlight the small differences. The (optional) scaling factor $\alpha > 0$ allows to recover the lazy regime when set to a large value (its role is detailed next section). For neural networks, the computational complexity per iteration is of the same order, the main difference being that for the tangent model, the forward and backward pass are done with the weights at initialization instead of the current weights. Another difference is that in the lazy regime, all the training information lies in the small bits of w which might make this regime unstable, e.g., for network compression [15].

2.2 Limit kernels and random feature

In this section, we show that the tangent kernel is a *random feature kernel* for neural networks with a single hidden layer. Consider a hidden layer of size $m \in \mathbb{N}$ and an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$,

$$f_m(w, x) = \frac{1}{\sqrt{m}} \sum_{j=1}^m b_j \cdot \sigma(a_j \cdot x),$$

with parameters⁵ $w = (a_1, \dots, a_m, b) \in (\mathbb{R}^d)^m \times \mathbb{R}^m$, so here $p = m(d+1)$. This scaling by $1/\sqrt{m}$ is the same as in [13] and leads to a non-degenerated limit of the kernel⁶ as $m \rightarrow \infty$. The associated tangent kernel in Eq. (3) is the sum of two kernels $k_m(x, x') = k_m^{(a)}(x, x') + k_m^{(b)}(x, x')$, one for each layer, where

$$k_m^{(a)}(x, x') = \frac{1}{m} \sum_{j=1}^m (x \cdot x') b_j^2 \sigma'(a_j \cdot x) \sigma'(a_j \cdot x') \quad \text{and} \quad k_m^{(b)}(x, x') = \frac{1}{m} \sum_{j=1}^m \sigma(a_j \cdot x) \sigma(a_j \cdot x').$$

If we assume that the initial weights a_j (resp. b_j) are independent samples of a distribution on \mathbb{R}^d (resp. a distribution on \mathbb{R}), these are random feature kernels [27] that converge as $m \rightarrow \infty$ to the kernels

$$k^{(a)}(x, x') = \mathbb{E}_{(a,b)} [(x \cdot x') b^2 \sigma'(a \cdot x) \sigma'(a \cdot x')] \quad \text{and} \quad k^{(b)}(x, x') = \mathbb{E}_a [\sigma(a \cdot x) \sigma(a \cdot x')].$$

⁵We have omitted the bias/intercept, which is recovered by fixing the last coordinate of x to 1.

⁶Since the definition of gradients depends on the choice of a metric, this scaling is not of intrinsic importance. Rather, it reflects that we work with the Euclidean metric on \mathbb{R}^p . Scaling this metric by $1/\sqrt{m}$, a natural choice suggested by the Wasserstein metric, would call for a scaling of the model by $1/m$ (see, e.g., [7, Sec. 2.2]). The choice of scaling will however become important when dealing with training (see also discussion in Section 4.2.2).

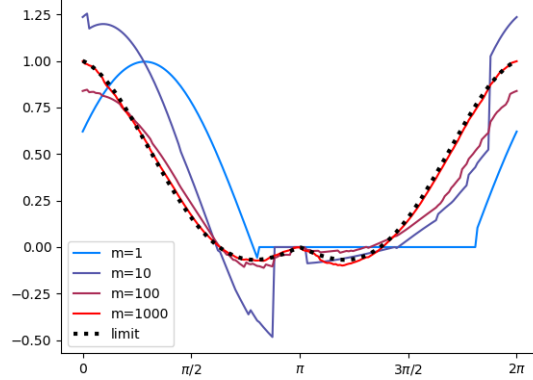


Figure 4: Random realizations of the kernels k_m and the limit kernel k of Eq. (4). We display the value of $k(x, x')$ as a function of $\varphi = \text{angle}(x, x')$ with x fixed, on a section of the sphere in \mathbb{R}^{10} . Parameters are normal random variables of variance 1, so $\mathbb{E}(b^2) = 1$ and $\mathbb{E}(\|a\|^2) = d$.

The second component $k^{(b)}$, corresponding to the differential with respect to the output layer, is the one traditionally used to make the link between these networks and random features [28]. When $\sigma(s) = \max\{s, 0\}$ is the rectified linear unit activation and the distribution of the weights a_j is rotation invariant in \mathbb{R}^d , one has the following explicit formulae [8]:

$$k^{(a)}(x, x') = \frac{(x \cdot x')\mathbb{E}(b^2)}{2\pi}(\pi - \varphi), \quad k^{(b)}(x, x') = \frac{\|x\|\|x'\|\mathbb{E}(\|a\|^2)}{2\pi d}((\pi - \varphi)\cos\varphi + \sin\varphi), \quad (4)$$

where $\varphi \in [0, \pi]$ is the angle between the two vectors x and x' . See Figure 4 for an illustration of this kernel. The link with random sampling is lost for deeper neural networks, because the non-linearities do not commute with expectations, but it is shown in [16] that tangent kernels still converge when the size of networks increase, for certain architectures.

3 Analysis of lazy training dynamics

3.1 Theoretical setting

This section is devoted to the theoretical analysis of lazy training dynamics under simple assumptions. Our goal is to show that they are essentially the same as the training dynamics for the tangent model of Eq. 2 when the scaling factor α is large. For theoretical purposes, the space of predictors \mathcal{F} is endowed with the structure of a separable Hilbert space and we consider an objective function $R : \mathcal{F} \rightarrow \mathbb{R}_+$ with a global minimizer \mathbf{f}^* . Our assumptions are the following:

Assumption 3.1. *The parametric model $\mathbf{f} : w \mapsto f(w, \cdot) \in \mathcal{F}$ is differentiable with a $L_{D\mathbf{f}}$ -Lipschitz derivative⁷ $D\mathbf{f}$. Moreover, R is μ -strongly convex and has a ν -Lipschitz derivative.*

This setting covers two cases of interest, where R is built from a loss function $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}_+$:

⁷Note that $D\mathbf{f}_w$ is a continuous linear map from \mathbb{R}^p to \mathcal{F} and the Lipschitz constant of $D\mathbf{f} : w \mapsto D\mathbf{f}_w$ is defined with respect to the operator norm. When $\mathcal{F} = \mathbb{R}^n$, $D\mathbf{f}_w$ can be identified with the Jacobian matrix $(\partial_{w_j} f(w, x_i))_{i,j}$ and we adopt this matrix notations throughout for simplicity.

- (interpolation) given a finite data set of n input/outputs $(x_i, y_i)_{i=1}^n$, we wish to find a model that fits these points. Here we define the objective function as $R(g) = \frac{1}{n} \sum_{i=1}^n \ell(g(x_i), y_i)$ and \mathcal{F} can be identified with \mathbb{R}^n with the Euclidean structure;
- (statistical learning) if instead, one desires good fitting on a (hypothetically) infinite data set, one may model the data as independent samples of a couple of random variables (X, Y) distributed according to a probability distribution $\rho \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R})$. The objective is the expected or population loss $R(g) = \mathbb{E}_{(X, Y)} \ell(g(X), Y)$ and \mathcal{F} is the space $L^2(\rho_X)$ of functions which are squared-integrable with respect to ρ_X , the marginal of ρ on the input space.

Scaled objective function. For a scale factor $\alpha > 0$, we introduce the scaled functional

$$F_\alpha(w) := R(\alpha \mathbf{f}(w))/\alpha^2. \quad (5)$$

This scaling factor was motivated in Section 1.1 as a mean to reach the lazy regime, when set to a large value. Here we have also multiplied the objective by $1/\alpha^2$, for the limit $\alpha \rightarrow \infty$ of training algorithms to be well-behaved. This scaling does not change the minimizers, but suggests in practice a step-size of the order $1/\alpha$, when α is large. For a quadratic loss $R(g) = \|g - g^*\|^2$ for some $g^* \in \mathcal{F}$, this objective reduces to $F_\alpha(w) = \|\alpha \mathbf{f}(w) - g^*\|^2/\alpha^2 = \|\mathbf{f}(w) - g^*/\alpha\|^2$ which amounts to learning a signal that is close to 0.

Lazy and tangent gradient flows. In the rest of this section, we study the *gradient flow* of the objective function F_α . This gradient flow is expected to reflect the behavior of first-order descent algorithms with small step-sizes, as the latter are known to approximate the former (see, e.g., [30] for gradient descent and [18, Thm. 2.1] for SGD). With an initialization $w_0 \in \mathbb{R}^p$, the gradient flow of F_α is the path $(w(t))_{t \geq 0}$ in the space of parameters \mathbb{R}^p that satisfies $w(0) = w_0$ and solves the ordinary differential equation:

$$w'(t) = -\nabla F_\alpha(w(t)) = -\frac{1}{\alpha} D\mathbf{f}_{w(t)}^\top \nabla R(\alpha \mathbf{f}(w(t))), \quad (6)$$

where $D\mathbf{f}^\top$ denotes the transposed/adjoint differential. We will study this dynamic for itself, but will also compare it to the gradient flow of the objective for the tangent model (Eq. 2). The objective function for the tangent model is defined as

$$\bar{F}_\alpha(w) := R(\alpha T_{\mathbf{f}}(w))/\alpha^2.$$

and the tangent gradient flow is the path $(\bar{w}(t))_{t \geq 0}$ that satisfies $\bar{w}(0) = w_0$ and

$$\bar{w}'(t) = -\nabla \bar{F}_\alpha(\bar{w}(t)) = -\frac{1}{\alpha} D\mathbf{f}_{w_0}^\top \nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t))).$$

As the gradient flow of a function \bar{F}_α that is strictly convex on the orthogonal complement of $\ker D\mathbf{f}_{w_0}$, $(\bar{w}(t))$ converges linearly to the unique global minimizer of \bar{F}_α . In particular, if $\mathbf{f}(w_0) = 0$ then this global minimizer does not depend on α and $\sup_{t \geq 0} \|\bar{w}(t) - w_0\| = O(1/\alpha)$.

3.2 Over-parameterized case

One generally says that a model is over-parameterized when the number of parameters p exceeds the number n of points to fit. The following proposition gives the main properties of lazy training under the slightly more stringent condition that $D\mathbf{f}_{w_0}$ is surjective (equivalently, has rank n). As $\text{rank } D\mathbf{f}_{w_0}$ gives the number of effective parameters or degrees of freedom of the model around w_0 , this over-parameterization assumption guarantees that any model around $\mathbf{f}(w_0)$ can be fitted.

Theorem 3.2 (Over-parameterized lazy training). *Let $\kappa := \nu/\mu$ denote the condition number of R and σ_{\min} the smallest singular value of $D\mathbf{f}_{w_0}$. Assume that $\sigma_{\min} > 0$ and that $\|\mathbf{f}(w_0)\| \leq M := \sigma_{\min}^3/(32\kappa^{3/2}\|D\mathbf{f}_{w_0}\|L_{D\mathbf{f}})$. If $\alpha > \|\mathbf{f}^*\|/M$, then for $t \geq 0$, it holds*

$$\|\alpha\mathbf{f}(w(t)) - \mathbf{f}^*\| \leq \sqrt{\kappa}\|\alpha\mathbf{f}(w_0) - \mathbf{f}^*\| \exp(-\mu\sigma_{\min}^2 t/4).$$

Moreover, as $\alpha \rightarrow \infty$, it holds $\sup_{t \geq 0} \|w(t) - w_0\| = O(1/\alpha)$,

$$\sup_{t \geq 0} \|\alpha\mathbf{f}(w(t)) - \alpha T_{\mathbf{f}}(\bar{w}(t))\| = O(1/\alpha) \quad \text{and} \quad \sup_{t \geq 0} \|w(t) - \bar{w}(t)\| = O(\log \alpha / \alpha^2).$$

The comparison with the tangent gradient flow in infinite time horizon is new and follows mostly from Lemma A.1 in appendix where constants are given. Otherwise, we do not claim an improvement over [13, 21, 12, 2, 3, 35]. The idea is rather to exhibit the key arguments behind lazy training in a simplified setting.

Proof. The trajectory in predictor space solves the differential equation

$$\frac{d}{dt} \alpha\mathbf{f}(w(t)) = -D\mathbf{f}_{w(t)} D\mathbf{f}_{w(t)}^T \nabla R(\alpha\mathbf{f}(w(t))),$$

that involves the covariance of the tangent kernel [16] evaluated at the current point $w(t)$, instead of w_0 . Consider the radius $r_0 := \sigma_{\min}/(2L_{D\mathbf{f}})$. By smoothness of \mathbf{f} , it holds $D\mathbf{f}_w D\mathbf{f}_w^T \succeq \sigma_{\min}^2 \text{Id}/4$ as long as $\|w - w_0\| < r_0$. Thus Lemma 3.3 below guarantees that $\alpha\mathbf{f}$ converges linearly, up to time $T := \inf\{t \geq 0; \|w(t) - w_0\| > r_0\}$. It only remains to find conditions on α so that $T = +\infty$. The variation of the parameters $w(t)$ can be bounded as follows for $0 \leq t \leq T$:

$$\|w'(t)\| \leq \frac{1}{\alpha} \|D\mathbf{f}_{w(t)}\| \|\nabla R(\alpha\mathbf{f}(w(t)))\| \leq \frac{2\nu}{\alpha} \|D\mathbf{f}_{w_0}\| \|\alpha\mathbf{f}(w(t)) - \mathbf{f}^*\|.$$

By Lemma 3.3, it follows that for $0 \leq t \leq T$,

$$\|w(t) - w(0)\| \leq \frac{2\nu^{3/2}}{\alpha\mu} \|D\mathbf{f}_{w_0}\| \|\alpha\mathbf{f}(w_0) - \mathbf{f}^*\| \int_0^t e^{-(\mu\sigma_{\min}^2/4)s} ds \leq \frac{8\kappa^{3/2}}{\alpha\sigma_{\min}^2} \|D\mathbf{f}_{w_0}\| \|\alpha\mathbf{f}(w_0) - \mathbf{f}^*\|.$$

This quantity is smaller than r_0 , and thus $T = \infty$, if $\|\alpha\mathbf{f}(w_0) - \mathbf{f}^*\| \leq 2\alpha M$. This is in particular guaranteed by the conditions on $\mathbf{f}(w_0)$ and α in the theorem. This also implies the “laziness” property $\sup_{t \geq 0} \|w(t) - w_0\| = O(1/\alpha)$.

For the comparison with the tangent gradient flow, the first bound is obtained by applying Lemma A.1 with $y(t) = \alpha\mathbf{f}(w(t))$ and $\bar{y}(t) = \alpha T_{\mathbf{f}}(\bar{w}(t))$, and noticing that the quantity denoted by K in that lemma is in $O(1/\alpha)$ thanks to the previous bound on $\|w(t) - w(0)\|$. For the last bound, we compute the integral over $[0, +\infty)$ of the bound

$$\begin{aligned} \alpha\|w'(t) - \bar{w}'(t)\| &= \|D\mathbf{f}_{w(t)} \nabla R(\alpha\mathbf{f}(w(t))) - D\mathbf{f}_{w_0} \nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t)))\| \\ &\leq \|D\mathbf{f}_{w(t)} - D\mathbf{f}_{w_0}\| \|\nabla R(\alpha\mathbf{f}(w(t)))\| + \|D\mathbf{f}_{w_0}\| \|\nabla R(\alpha\mathbf{f}(w(t))) - \nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t)))\|. \end{aligned}$$

It is easy to see from the derivations above that the integral of the first term is in $O(1/\alpha)$. For the second term, we define $t_0 := 4 \log \alpha / (\mu\sigma_{\min}^2)$ and on $[0, t_0]$ we use the smoothness bound

$$\|\nabla R(\alpha\mathbf{f}(w(t))) - \nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t)))\| \leq \nu \|\alpha\mathbf{f}(w(t)) - \alpha T_{\mathbf{f}}(\bar{w}(t))\|$$

which integral over $[0, t_0]$ is in $O(\log \alpha / \alpha)$, while on $[t_0, +\infty)$ we use the crude bound

$$\|\nabla R(\alpha\mathbf{f}(w(t))) - \nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t)))\| \leq \|\nabla R(\alpha\mathbf{f}(w(t)))\| + \|\nabla R(\alpha T_{\mathbf{f}}(\bar{w}(t)))\|$$

which integral over $[t_0, +\infty)$ is in $O(1/\alpha)$ thanks to the definition of t_0 and the exponential decrease of ∇R along both trajectories. \square

In geometrical term, the proof above can be summarized as follows. It is a general fact that the parametric model \mathbf{f} induces a pushforward-metric on \mathcal{F} [20]. With a certain choice of scaling, this metric hardly changes during training, equalling the inverse covariance of the tangent kernel. This makes the loss landscape in \mathcal{F} essentially convex and allows to call the following lemma, that shows linear convergence of strongly-convex gradient flows in a time-dependent metric.

Lemma 3.3 (Strongly-convex gradient flow in a time-dependent metric). *Let $F : \mathcal{F} \rightarrow \mathbb{R}$ be a μ -strongly-convex function with ν -Lipschitz continuous gradient and with global minimizer y^* and let $\Sigma(t) : \mathcal{F} \rightarrow \mathcal{F}$ be a time dependent continuous auto-adjoint linear operator with eigenvalues lower bounded by $\lambda > 0$ for $0 \leq t \leq T$. Then solutions on $[0, T]$ to the differential equation*

$$y'(t) = -\Sigma(t)\nabla F(y(t)),$$

satisfy, for $0 \leq t \leq T$,

$$\|y(t) - y^*\| \leq (\nu/\mu)^{1/2} \|y(0) - y^*\| \exp(-\mu\lambda t).$$

Proof. By strong convexity, it holds $\bar{F}(y) := F(y) - F(y^*) \leq \frac{1}{2\mu} \|\nabla F(y)\|^2$. It follows

$$\frac{d}{dt} \bar{F}(y(t)) = -\nabla F(y(t))^\top \Sigma(t) \nabla F(y(t)) \leq -\lambda \|\nabla F(y(t))\|^2 \leq -2\mu\lambda \bar{F}(y),$$

and thus $\bar{F}(y(t)) \leq \exp(-2\mu\lambda t) \bar{F}(y(0))$ by Grönwall's Lemma. We now use the strong convexity inequality $\|y - y^*\|^2 \leq \frac{2}{\mu} \bar{F}(y)$ in the left-hand side and the smoothness inequality $\bar{F}(y) \leq \frac{1}{2}\nu \|y - y^*\|^2$ in the right-hand side. This yields $\|y(t) - y^*\|^2 \leq \frac{\nu}{\mu} \exp(-2\mu\lambda t) \|y(0) - y^*\|^2$. \square

3.3 Under-parameterized case

In this section, we do not make the over-parameterization assumption and show that the lazy regime still occurs for large values of α . This covers for instance the case of the population loss, where $n = \infty$. For this setting, we content ourselves with a qualitative statement⁸ proved in Appendix B.

Theorem 3.4 (Under-parameterized lazy training). *Assume that $\mathbf{f}(w_0) = 0$ and that $\text{rank } D\mathbf{f}_{w_0}$ is constant on a neighborhood of w_0 . Then there exists $\alpha_0 > 0$ such that for all $\alpha \geq \alpha_0$ the gradient flow (6) converges at a geometric rate (asymptotically independent of α) to a local minimum of F_α .*

The assumption that the rank is locally constant holds generically due to lower-semicontinuity of the rank function. In particular, it holds with probability 1 if w_0 is randomly chosen according to an absolutely continuous probability measure. In this under-parameterized case, the limit is in general not a global minimum because the image of $D\mathbf{f}_{w_0}$ is a proper subspace of \mathcal{F} which may not contain the global minimizer of R , as pictured in Figure 2. Thus it cannot be excluded that there are models $\mathbf{f}(w)$ with w far from w_0 that have a smaller loss. This fact is clearly observed experimentally in Section 5 (Figure 6-(b)). Finally, a comparison with the tangent gradient flow as in Theorem 3.2 could be shown along the same lines, but would be technically slightly more involved because differential geometry comes into play.

⁸Quantitative statements would involve the smallest positive singular value of $D\mathbf{f}_{w_0}$, which is anyways hard to control.

Relationship to the global convergence result in [7]. A consequence of Theorem 3.4 is that when training a neural network with SGD to minimize a population loss (i.e., $n = \infty$) then lazy training gets stuck in a local minimum. In contrast, it is shown in [7] that gradient flows of neural networks with a single hidden layer converge to global optimality in the over-parameterization limit if initialized with enough diversity in the weights. This is not a contradiction since Theorem 3.4 assumes a finite number p of parameters. For lazy training, the population loss also converges to its minimum when p increases if the tangent kernel converges to a universal kernel as $p \rightarrow \infty$. However, this convergence might be unreasonably slow and does not compare to the actual practical performance of neural networks, as Figure 1-(c) suggests. As a side note, we stress that the global convergence result in [7] is not limited to lazy dynamics but also covers highly non-linear dynamics, including “active” learning behaviors such as seen on Figure 1.

4 Range of the lazy regime

In this section, we derive a simplified rule to determine if a certain choice of initialization leads to lazy training. Our aim is to emphasize the fact that this regime is due to the choice of scaling, which is often implicit. We proceed informally and do not claim mathematical rigorousness.

4.1 Informal view

Suppose that the training algorithm on the tangent model T_f from Eq. 2 converges to w^* . This tangent model is an accurate approximation of f throughout training when the second-order remainder remains negligible until convergence, i.e., $L_{Df}\|w^* - w_0\|^2 \ll \|Df_{w_0}(w^* - w_0)\|$ where L_{Df} is the Lipschitz constant of Df . With the rough simplification that $\|Df_{w_0}(w^* - w_0)\| \approx \|Df_{w_0}\| \cdot \|w^* - w_0\|$ this is equivalent to the property

$$\|w^* - w_0\| \ll \|Df_{w_0}\|/L_{Df}.$$

Making also the approximation $\|T_f(w^*) - T_f(w_0)\| \approx \|Df_{w_0}\| \cdot \|w^* - w_0\|$ leads to the following rule of thumb: lazy training occurs if

$$\|T_f(w^*) - f(w_0)\| \ll \|Df_{w_0}\|^2/L_{Df}. \quad (7)$$

This bound compares how close the model at initialization is to the best linear model, with the extent of validity of the linear approximation (it is a simplified version of the hypothesis of Theorem 3.2).

Scaling factor. The interest of Eq. (7) is that it allows for simplified informal considerations. For instance, consider the scaling factor $\alpha > 0$ and an initialization w_0 such that $\|f(w_0)\| = O(1/\alpha)$. Then for the scaled model αf , the left-hand side of (7) does not grow with α while the right-hand side is proportional to α . Thus, one is bound to reach the lazy regime when α is large.

Dealing with non-differentiability. In practice, one often uses models that are only differentiable in a weaker sense, as neural networks with ReLU activation or max-pooling. Then the quantity L_{Df} involved in Eq. (7) does not make sense immediately and should be translated into estimates on the stability of Df . It is one of the important technical contributions of [13, 21, 12, 2, 3, 35] to have studied rigorously such aspects in the non-differentiable setting.

4.2 Examples

4.2.1 Homogeneous models

A parametric model f is said q -homogeneous in the parameters w ($q \geq 2$), if for all $\lambda > 0$ it holds

$$f(\lambda w, x) = \lambda^q f(w, x).$$

For such models, changing the magnitude of the initialization by a factor λ is equivalent to changing the scale factor, with the relationship $\alpha = \lambda^q$. Indeed, the i -th derivative of a q -homogeneous function is $(q - i)$ -homogeneous. It follows that if the initialization w_0 is multiplied by λ , then the ratio of Eq. (7) is multiplied by $\lambda^{2q-2}/\lambda^{q-2} = \lambda^q$. This suggests an explanation of why neural networks initialized with large weights, but at the same time close to 0 in the space \mathcal{F} of predictors, display a lazy regime. In our experiments, large weights correspond to a high variance τ^2 at initialization (see Figures 1 and 6(b) for 2-homogeneous examples).

4.2.2 Large shallow neural networks

In order to understand how lazy training can be intertwined with over-parameterization, let us consider the most simple non-trivial example of a parametric model of the form

$$f(w, x) = \alpha(m) \sum_{i=1}^m \phi(\theta_i, x),$$

where the parameters are $w = (\theta_1, \dots, \theta_m)$ and $\Phi : \theta \mapsto \phi(\theta, \cdot)$ is assumed differentiable with a Lipschitz derivative. This includes neural networks with a single hidden layer, where $\phi(\theta = (a, b, c), x) = a \cdot \sigma(b + c \cdot x)$ (when σ is the ReLU activation function, differentiability requires $(b, c) \neq (0, 0)$, see, e.g., [7, App. D.4]). Our notation indicates that α should depend on m . Indeed, taking the limit $m \rightarrow \infty$ makes the quantities in Eq.(7) explode or vanish, depending on $\alpha(m)$, thus leading to lazy training or not.

Now consider an initialization $w_0 = (\theta_1, \dots, \theta_m)$ with independent and identically distributed variables θ_i such that $\mathbb{E}\Phi(\theta_i) = 0$. The initial model $\mathbf{f}(w_0) = \alpha(m) \sum_{i=1}^m \Phi(\theta_i)$ satisfies:

$$\mathbb{E}\|\mathbf{f}(w_0)\|^2 = m\alpha(m)^2\mathbb{E}\|\Phi(\theta)\|^2, \quad \mathbb{E}\|D\mathbf{f}_{w_0}\|^2 = m\alpha(m)^2\mathbb{E}\|D\Phi_\theta\|^2 \quad \text{and} \quad L_{D\mathbf{f}} = \alpha(m)L_{D\Phi}.$$

Thus the terms involved in the criterion of Eq.(7) satisfy (hiding constants depending on R or Φ):

$$\mathbb{E}\|T_{\mathbf{f}}(w^*) - \mathbf{f}(w_0)\| \lesssim \max\{1, \sqrt{m}\alpha(m)\}, \quad \mathbb{E}\|D\mathbf{f}_{w_0}\|^2/L_{D\mathbf{f}} \approx m\alpha(m).$$

It follows, looking at Eq.(7), that the criterion to obtain lazy training is satisfied when m grows if $m \cdot \alpha(m) \rightarrow \infty$. For instance, Du et al. [13] consider $\alpha(m) = 1/\sqrt{m}$ and Li et al. [21] consider an initialization with standard deviation $1/\sqrt{m}$ with a 1-homogeneous Φ , which amounts to the same scaling $\alpha(m) = 1/\sqrt{m}$ and leads to lazy training. It is more difficult to analyze this ratio for deeper neural networks as studied in [12, 2, 3, 35].

Mean-field limit. Note that this scaling contrasts with the scaling $\alpha(m) = 1/m$ chosen in a series of recent works that study the mean-field limit of neural networks with a single hidden layer [25, 7, 29, 31]. This scaling is precisely the one that maintains a ratio $\|D\mathbf{f}_{w_0}\|^2/L_{D\mathbf{f}}$ of order 1 and allows to converge as $m \rightarrow \infty$ to a non-degenerate dynamic, described by a partial differential equation. As a side note, this scaling leads to a vanishing derivative $D\mathbf{f}_{w_0}$ as $m \rightarrow \infty$: this might appear ill-posed, but this is actually not an intrinsic fact. Indeed, this is due to the choice of the Euclidean metric on \mathbb{R}^p , which could be itself scaled by $1/\sqrt{m}$ to give a non-degenerate limit [7]. In contrast, the ratio $\|D\mathbf{f}_{w_0}\|^2/L_{D\mathbf{f}}$ is independent of the choice of the scaling of the metric on the space of parameters.

5 Numerical experiments

In our numerical experiments, we consider the following neural network with a single-hidden layer

$$f(w, x) = \sum_{j=1}^m b_j \max(a_j \cdot x, 0),$$

with parameters $w = (a_1, \dots, a_m, b) \in (\mathbb{R}^d)^m \times \mathbb{R}^m$, so here $p = m(d+1)$. These parameters are initialized randomly and independently according to a normal distribution $\mathcal{N}(0, \tau^2)$, except when using the “doubling trick” mentioned in Section 2. In this case (assuming m even), the parameters with index $1, \dots, m/2$ are random as above and we set, for $j > m/2$, $a_j = a_{j-m/2}$ and $b_j = -b_{j-m/2}$. The input data are uniformly random on the unit sphere in \mathbb{R}^d and the output data is given by the output of a neural network with $m_0 = 3$ hidden units and with random parameters normalized so that $\|b_j a_j\| = 1$ for $j \in \{1, 2, 3\}$. We chose the quadratic loss and performed batch gradient descent with a small step-size on a finite data set of size n , except for Figure 6-(b) which has been obtained with mini-batch SGD with a small step-size and where each sample is only used once.

Cover illustration. Figure 1 in Section 1 was used to motivate this note and was discussed in Section 1. Here we give more details on the setting. On panels (a)-(b), we have taken $n = 15$ samples and $m = 20$ neurons with the doubling trick, in order to show a simple dynamic. To obtain a 2-d representation, we plot $|b_j(t)|a_j(t)$ throughout training (lines) and at convergence (dots) for $j \in \{1, \dots, m\}$. The blue or red colors stand for the sign of $b_j(t)$. The unit circle is displayed to help visualizing the change of scale. On panel (c), we have taken $n = 1000$, $m = 50$, used the “doubling trick” and averaged the result over 10 experiments. To make sure that the bad performance for large τ is not due to a lack of regularization, we have displayed the best test error throughout training (for kernel methods, early stopping is a form of regularization [32]).

Many neurons dynamics. Figure 5 is similar to panels (a)-(b) in Figure 1 except that there are $m = 200$ neurons and $n = 200$ training samples. We also show the trajectory of the parameters without the “doubling trick” where we see that the neurons need to move slightly more than in panel (c) in order to compensate for the non-zero initialization $\mathbf{f}(w_0)$. The good behavior for small τ (also observed on Figure 1) can be related to the results in [23] where it is shown that gradient flows of this form initialized close to zero quantize features.

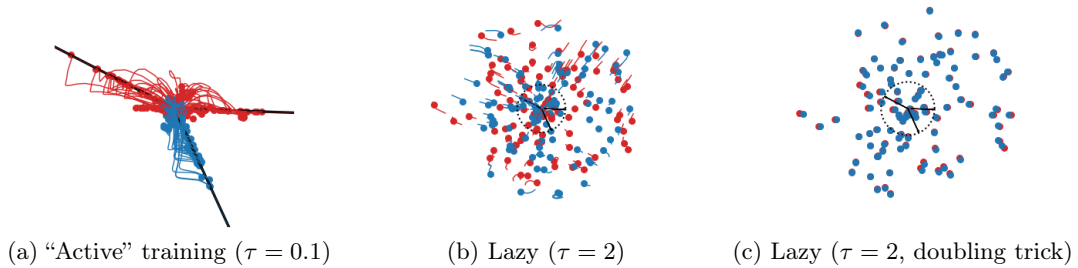


Figure 5: Training a 2-layers ReLU neural network initialized with normal random weights of variance τ^2 , as in Figure 1, but with more neurons. In this 2-homogeneous setting, changing τ^2 is equivalent to changing α by the same amount (Section 4.2.1) so lazy training occurs for large τ .

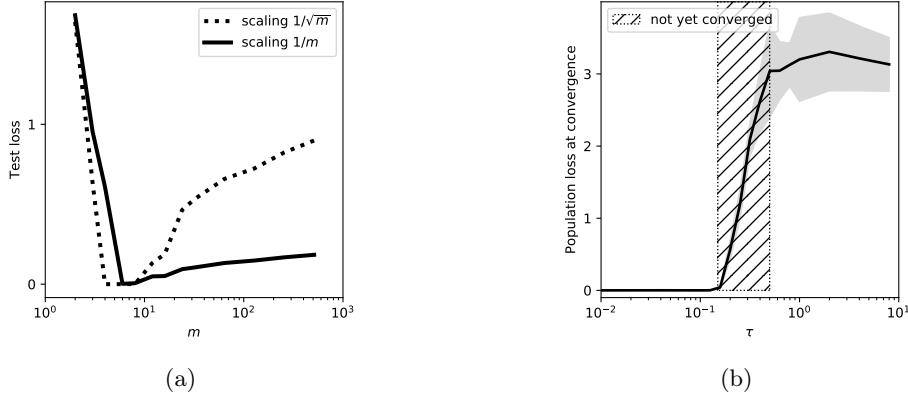


Figure 6: (a) Generalization when α depends on m as $\alpha = 1/m$ or $\alpha = 1/\sqrt{m}$, the latter leading to lazy training for large m . (b) Population loss at convergence versus τ for SGD with a random $\mathcal{N}(0, \tau^2)$ initialization. In the hatched area the loss was still slowly decreasing.

Increasing number of parameters. Figure 6-(a) shows the evolution of the test error when increasing m as discussed in Section 4.2.2, for two choices of scaling functions $\alpha(m)$, averaged over 5 experiments in dimension $d = 100$, with “doubling trick”. The scaling $1/\sqrt{m}$ leads to lazy training, with a poor generalization as m increases. This is in contrast to the scaling $1/m$ for which the test error remains relatively close to 0 for large m . The high test error for small values of m is due to the fact that the ground truth model has $m_0 = 3$ neurons: gradient descent seems to need a slight over-parameterization to perform well (more experiments with the scaling $1/m$ can be found in [7, 29, 25]).

Under-parameterized with SGD. Finally, Figure 6-(b) illustrates the under-parameterized case. We consider a random $\mathcal{N}(0, \tau^2)$ initialization with “doubling trick” in dimension $d = 100$ and $m = 50$ neurons. We used SGD with a batch-size 200, and displayed the final population loss (estimated with 2000 samples) averaged over 5 experiments. As shown in Theorem 3.4, SGD remains stuck in a local minimum in the lazy regime (i.e., here for large τ). As in Figure 1, it behaves intriguingly well when τ is small. There is also an intermediate regime (hatched area) where convergence is very slow and the loss was still decreasing when the algorithm was stopped.

6 Discussion

By connecting a series of recent works, we have exhibited the simple structure behind lazy training, which is a situation when a non-linear parametric model behaves like a linear one. We have studied under which conditions this regime occurs and shown that it is not limited to over-parameterized models. While the lazy training regime provides some of the first optimization-related theoretical insights for deeper models [12, 2, 3, 35, 16], we believe it does not explain yet the many successes of neural networks that have been observed in various challenging, high-dimensional tasks in machine learning. This is corroborated by numerical experiments where it is seen that networks trained in the lazy regime are those that perform worst. Instead, the intriguing phenomenon that still defies theoretical understanding is the one displayed on Figures 1(c) and 6(b) for small τ : neural networks trained with gradient-based methods (and neurons that move) have the ability to perform high-dimensional feature selection.

Acknowledgments

We acknowledge supports from grants from Région Ile-de-France and the European Research Council (grant SEQUOIA 724063).

References

- [1] Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. *Manifolds, tensor analysis, and applications*, volume 75. Springer Science & Business Media, 2012.
- [2] Zeyuan Allen-Zhu, Li Yuanzhi, and Liang Yingyu. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.04918*, 2018.
- [3] Zeyuan Allen-Zhu, Li Yuanzhi, and Liang Yingyu. Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*, 2018.
- [4] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- [5] Youness Boutaib. On Lipschitz maps and their flows. *arXiv preprint arXiv:1510.07614*, 2015.
- [6] Yuan Cao and Quanquan Gu. A generalization theory of gradient descent for learning over-parameterized deep relu networks. *arXiv preprint arXiv:1902.01384*, 2019.
- [7] Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, 2018.
- [8] Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. In *Advances in neural information processing systems*, pages 342–350, 2009.
- [9] Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In *Advances In Neural Information Processing Systems*, pages 2253–2261, 2016.
- [10] Aymeric Dieuleveut and Francis Bach. Nonparametric stochastic approximation with large step-sizes. *The Annals of Statistics*, 44(4):1363–1399, 2016.
- [11] Simon Du and Jason Lee. On the power of over-parametrization in neural networks with quadratic activation. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1329–1338, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- [12] Simon S. Du, Lee Jason D., Li Haochuan, Wang Liwei, and Zhai Xiyu. Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*.
- [13] Simon S. Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- [14] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.

- [15] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *International Conference on Learning Representations*, 2016.
- [16] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, 2018.
- [17] Jyrki Kivinen, Alexander J. Smola, and Robert C. Williamson. Online learning with kernels. *IEEE Transactions on Signal Processing*, 52(8):2165–2176, 2004.
- [18] Harold Kushner and G. George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.
- [19] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018.
- [20] John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–29. Springer, 2003.
- [21] Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8167–8176, 2018.
- [22] Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with ReLU activation. In *Advances in Neural Information Processing Systems*, pages 597–607, 2017.
- [23] Hartmut Maennel, Olivier Bousquet, and Sylvain Gelly. Gradient descent quantizes ReLU network features. *arXiv preprint arXiv:1803.08367*, 2018.
- [24] Alexander G. de G. Matthews, Mark Rowland, Jiri Hron, Richard E. Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018.
- [25] 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.
- [26] Samet Oymak and Mahdi Soltanolkotabi. Overparameterized nonlinear learning: Gradient descent takes the shortest path? *arXiv preprint arXiv:1812.10004*, 2018.
- [27] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [28] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in neural information processing systems*, pages 1313–1320, 2009.
- [29] Grant M. Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. In *Advances in neural information processing systems*, 2018.

- [30] Damien Scieur, Vincent Roulet, Francis Bach, and Alexandre d’Aspremont. Integration methods and optimization algorithms. In *Advances in Neural Information Processing Systems*, pages 1109–1118, 2017.
- [31] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *arXiv preprint arXiv:1808.09372*, 2018.
- [32] Yuan Yao, Lorenzo Rosasco, and Andrea Caponnetto. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289–315, 2007.
- [33] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- [34] Chiyuan Zhang, Samy Bengio, and Yoram Singer. Are all layers created equal? *arXiv preprint arXiv:1902.01996*, 2019.
- [35] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *arXiv preprint arXiv:1811.08888*, 2018.

A Stability lemma

The following stability lemma is at the basis of the equivalence between lazy training and linearized model training. We limit ourselves to a rough estimate sufficient for our purposes.

Lemma A.1. *Let $R : \mathcal{F} \rightarrow \mathbb{R}_+$ be a μ -strongly convex function and let $\Sigma(t)$ be a time dependent positive definite operator on \mathcal{F} such that $\Sigma(t) \succeq \lambda \text{Id}$ for $t \geq 0$. Consider the paths $y(t)$ and $\bar{y}(t)$ on \mathcal{F} that solve for $t \geq 0$,*

$$y'(t) = -\Sigma(t)\nabla R(y(t)) \quad \text{and} \quad \bar{y}'(t) = -\Sigma(0)\nabla R(\bar{y}(t)).$$

Defining $K := \sup_{t \geq 0} \|(\Sigma(t) - \Sigma(0))\nabla R(y(t))\|$, it holds for $t \geq 0$,

$$\|y(t) - \bar{y}(t)\| \leq \frac{K\|\Sigma(0)\|^{1/2}}{\lambda^{3/2}\mu}.$$

Proof. Let $\Sigma_0^{1/2}$ be the positive definite square root of $\Sigma(0)$, let $z(t) = \Sigma_0^{-1/2}y(t)$, $\bar{z}(t) = \Sigma_0^{-1/2}\bar{y}(t)$ and let $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be the function defined as $h(t) = \frac{1}{2}\|z(t) - \bar{z}(t)\|^2$. It holds

$$\begin{aligned} h'(t) &= \langle z'(t) - \bar{z}'(t), z(t) - \bar{z}(t) \rangle \\ &= -\langle \Sigma_0^{-1/2}\Sigma(t)\nabla R(\Sigma_0^{1/2}z(t)) - \Sigma_0^{1/2}\nabla R(\Sigma_0^{1/2}\bar{z}(t)), z(t) - \bar{z}(t) \rangle \\ &= -\langle \Sigma_0^{1/2}\nabla R(\Sigma_0^{1/2}z(t)) - \Sigma_0^{1/2}\nabla R(\Sigma_0^{1/2}\bar{z}(t)), z(t) - \bar{z}(t) \rangle \quad (A(t)) \\ &\quad - \langle \Sigma_0^{-1/2}(\Sigma(t) - \Sigma(0))\nabla R(\Sigma_0^{1/2}z(t)), z(t) - \bar{z}(t) \rangle. \quad (B(t)) \end{aligned}$$

Since the function $z \mapsto R(\Sigma_0^{1/2}z)$ is $\mu\sqrt{\lambda}$ -strongly convex, one has that $A(t) \leq -2\sqrt{\lambda}\mu h(t)$. Using the quantity K introduced in the statement, one has also $\|B(t)\| \leq K\|z(t) - \bar{z}(t)\|/\sqrt{\lambda} = K\sqrt{2h(t)}/\lambda$. Summing these two terms yields the bound

$$h'(t) \leq K\sqrt{2h(t)}/\lambda - 2\lambda\mu h(t).$$

The right-hand side is a concave function of $h(t)$ which is nonnegative for $h(t) \in [0, K^2/(2\lambda^3\mu^2)]$ and negative for higher values of $h(t)$. Since $h(0) = 0$ it follows that for all $t \geq 0$, one has $h(t) \leq K^2/(2\lambda^3\mu^2)$ and the result follows since $\|y(t) - \bar{y}(t)\| \leq \|\Sigma(0)\|^{1/2}\sqrt{2h(t)}$. \square

B Proof of Theorem 3.4 (under-parameterized case)

By the rank theorem (a result of differential geometry, see [20, Thm. 4.12] or [1] for an infinite dimensional statement), there exists open sets $\mathcal{W}_0, \bar{\mathcal{W}}_0 \subset \mathbb{R}^p$ and $\mathcal{F}_0, \bar{\mathcal{F}}_0 \subset \mathcal{F}$ and diffeomorphisms $\varphi : \mathcal{W}_0 \rightarrow \bar{\mathcal{W}}_0$ and $\psi : \mathcal{F}_0 \rightarrow \bar{\mathcal{F}}_0$ such that $\varphi(w_0) = 0$, $\psi(f(w_0)) = 0$ and $\psi \circ \mathbf{f} \circ \varphi^{-1} = \pi_r$, where π_r is the map that writes, in suitable bases, $(x_1, \dots, x_p) \mapsto (x_1, \dots, x_r, 0, \dots)$. Up to restricting these domains, we may assume that \mathcal{F}_0 is convex. We also denote by Π_r the r -dimensional hyperplan in \mathcal{F} that is spanned by the first r vectors of the basis. The situation is summarized in the following commutative diagram:

$$\begin{array}{ccc} \mathcal{W}_0 & \xrightarrow{\mathbf{f}} & \mathcal{F}_0 \\ \varphi \downarrow & & \downarrow \psi \\ \bar{\mathcal{W}}_0 & \xrightarrow{\pi_r} & \bar{\mathcal{F}}_0 \end{array}$$

In the rest of the proof, we denote by $C > 0$ any quantity that depends on μ, ν and the Lipschitz and Lipschitz smoothness constants of $\mathbf{f}, \psi, \varphi, \psi^{-1}, \varphi^{-1}$, but not on α . Although we do not do so, this could be translated into explicit constants that depends on $\mu, \nu, L_{\mathbf{f}}, L_{D\mathbf{f}}$ and on the smallest positive singular value of $D_f(w_0)$ using quantitative versions of the rank theorem [5, Thm. 2.7].

Step 1. Our proof is along the same lines as that of Theorem 3.2, but performed in Π_r which can be thought of as a straighten up version of $\mathbf{f}(\mathcal{W}_0)$. Consider the function G_α defined for $\mathbf{g} \in \bar{\mathcal{F}}_0$ as $G_\alpha(\mathbf{g}) = R(\alpha\psi^{-1}(\mathbf{g}))/\alpha^2$. The gradient and Hessian of G_α satisfy, for $v_1, v_2 \in \mathbb{R}^p$,

$$\nabla G_\alpha(\mathbf{g}) = \frac{1}{\alpha} (D\psi_{\mathbf{g}}^{-1})^\top \nabla R(\alpha\psi^{-1}(\mathbf{g})),$$

$$D^2 G_\alpha(\mathbf{g})(v_1, v_2) = v_1^\top (D\psi_{\mathbf{g}}^{-1})^\top \nabla^2 R(\alpha\psi^{-1}(\mathbf{g})) D\psi_{\mathbf{g}}^{-1} v_2 + \frac{1}{\alpha} D^2 \psi_{\mathbf{g}}^{-1}(v_1, v_2)^\top \nabla R(\alpha\psi^{-1}(\mathbf{g})).$$

The second order derivative of G_α is the sum of a first term with eigenvalues in an interval $[C^{-1}, C]$, and a second term that goes to 0 as α increases. It follows that G_α is smooth and strongly convex for α large enough. Note that if R or ψ^{-1} are not twice continuously differentiable, then the Hessian computations should be understood in the distributional sense (this is sufficient because the functions involved are Lipschitz smooth). Also, let \mathbf{g}^* be a minimizer of the lower-semicontinuous closure of G_α on the closure of $\bar{\mathcal{F}}_0$. By strong convexity of R and our assumptions, it holds

$$\|\mathbf{g}^*\|^2 \leq \frac{2}{\mu} (G_\alpha(0) - G_\alpha(\mathbf{g}^*)) \leq \frac{2R(0)}{\alpha^2 \mu},$$

so \mathbf{g}^* is in the interior of $\bar{\mathcal{F}}_0$ for α large enough and is then the unique minimizer of G_α .

Step 2. Now consider $T := \inf\{t \geq 0 ; w(t) \notin \mathcal{W}_0\}$. For $t \in [0, T)$, the trajectory $w(t)$ of the gradient flow (6) has “mirror” trajectories in the four spaces in the diagram above. Let us look more particularly at $\mathbf{g}(t) := \pi_r \circ \varphi(w(t)) = \psi \circ f(w(t))$ for $t < T$. In the following computation, we write $D\varphi$ for the value of the Jacobian at the corresponding point of the dynamic $D\varphi_{w(t)}$ (and similarly for other Jacobian matrices). By noticing that $D\mathbf{f} = D\psi^{-1} D\pi_r D\varphi$, we have

$$\begin{aligned} \mathbf{g}'(t) &= -\frac{1}{\alpha} D\psi D\mathbf{f} D\mathbf{f}^\top \nabla R(\alpha\psi^{-1}(\mathbf{g}(t))) \\ &= -\frac{1}{\alpha} D\pi_r D\varphi D\varphi^\top D\pi_r^\top (D\psi^{-1})^\top \nabla R(\alpha\psi^{-1}(\mathbf{g}(t))). \end{aligned}$$

so $\mathbf{g}(t)$ remains in Π_r . Also, the first $r \times r$ block of $D\pi_r D\varphi D\varphi^\top D\pi_r^\top$ is positive definite on Π_r , with a positive lower bound (up to taking \mathcal{W}_0 and \mathcal{F}_0 smaller if necessary). Thus by Lemma 3.3, there are constants $C_1, C_2 > 0$ independent of α such that, for $t \in [0, T)$, $\|\mathbf{g}(t) - \mathbf{g}^*\| \leq C_1 \|\mathbf{g}(0) - \mathbf{g}^*\| \exp(-C_2 t)$.

Step 3. Now we want to show that $T = +\infty$ for α large enough. It holds

$$w'(t) = -\frac{1}{\alpha} D\mathbf{f}^\top \nabla R(\alpha\mathbf{f}(w(t))) = D\varphi^\top D\pi_r^\top \nabla G_\alpha(\mathbf{g}(t))$$

and, by Lipschitz-smoothness of G_α (Step 1), $\|\nabla G_\alpha(\mathbf{g}(t))\| \leq \frac{C}{\alpha} \|\mathbf{g}(t) - \mathbf{g}^*\|$ hence

$$\|w(t) - w(0)\| \leq \frac{C}{\alpha} \int_0^t \exp(-C_2 s) ds \leq \frac{C}{\alpha C_2}.$$

Thus, by choosing α large enough, one has $w(t) \in \mathcal{W}_0$ for all $t \geq 0$, so $T = \infty$ and the theorem follows.