arXiv:2106.15933v2 [stat.ML] 31 Jan 2022

Saddle-to-Saddle Dynamics in Deep Linear Networks: Small Initialization Training, Symmetry, and Sparsity

Arthur Jacot [1] **François Ged** [1] **Berfin Şimşek** [1] **Clément Hongler** [1] **Franck Gabriel** [1]

Abstract

The dynamics of Deep Linear Networks (DLNs) is dramatically affected by the variance σ^2 of the parameters at initialization θ_0. For DLNs of width w, we show a phase transition w.r.t. the scaling γ of the variance $\sigma^2 = w^{-\gamma}$ as $w \to \infty$: for large variance ($\gamma < 1$), θ_0 is very close to a global minimum but far from any saddle point, and for small variance ($\gamma > 1$), θ_0 is close to a saddle point and far from any global minimum. While the first case corresponds to the well-studied NTK regime, the second case is less understood. This motivates the study of the case $\gamma \to +\infty$, where we conjecture a Saddle-to-Saddle dynamics: throughout training, gradient descent visits the neighborhoods of a sequence of saddles, each corresponding to linear maps of increasing rank, until reaching a sparse global minimum. We support this conjecture with a theorem for the dynamics between the first two saddles, as well as some numerical experiments.

1. Introduction

In spite of their widespread usage, the theoretical understanding of Deep Neural Networks (DNNs) remains limited. In contrast to more common statistical methods which are built (and proven) to recover the specific structure of the data, the development of DNNs techniques has been mostly driven by empirical results. This has led to a great variety of models which perform consistently well, but without a theory explaining why. In this paper, we provide a theoretical analysis of Deep Linear (Neural) Networks (DLNs), whose simplicity makes them particularly attractive as a first step towards the development of such a theory.

DLNs have a non-convex loss landscape and the behavior of training dynamics can be subtle. For shallow networks, the convergence of gradient descent is guaranteed by the fact that the saddles are strict and that all minima are global (Baldi & Hornik, 1989; Kawaguchi, 2016; Lee *et al.*, 2016; 2019b). In contrast, the deep case features non-strict saddles (Kawaguchi, 2016) and no general proof of convergence exists at the moment, though convergence to a global minimum can be guaranteed in some cases (Arora *et al.*, 2019a; Eftekhari, 2020).

A recent line of work focuses on the implicit bias of DLNs, and consistently reveals some form of incremental learning and implicit sparsity as in (Gissin *et al.*, 2020). Diagonal networks are known to learn minimal L_1 solutions (Moroshko *et al.*, 2020; Woodworth *et al.*, 2020). With a specific initialization and the MSE loss, DLNs learn the singular components of the signal one by one (Saxe *et al.*, 2014; Advani & Saxe, 2017; Saxe *et al.*, 2019; Gidel *et al.*, 2019; Arora *et al.*, 2019b). Recently, it has been shown that with losses such as the cross-entropy and the exponential loss, the parameters diverge towards infinity, but end up following the direction of the max-margin classifier w.r.t. the L_p-Schatten (quasi)norm (Gunasekar *et al.*, 2018a;b; Soudry *et al.*, 2018; Ji & Telgarsky, 2018; 2020; Chizat & Bach, 2020; Lyu & Li, 2020; Moroshko *et al.*, 2020; Yun *et al.*, 2021).

In parallel, recent works have shown the existence of two regimes in large-width DNNs: a kernel regime (also called NTK or lazy regime) where learning is described by the so-called Neural Tangent Kernel (NTK) guaranteeing linear convergence (Jacot *et al.*, 2018; Du *et al.*, 2019; Chizat & Bach, 2018a; Arora *et al.*, 2019c; Lee *et al.*, 2019a; Huang & Yau, 2019), and an active regime where the dynamics is nonlinear (Chizat & Bach, 2018b; Rotskoff & Vanden-Eijnden, 2018; Mei *et al.*, 2018; 2019; Chizat & Bach, 2020). For DLNs, both regimes can be observed as well, with evidence that while the linear regime exhibits no sparsity, the active regime favors solutions with some kind of sparsity (Woodworth *et al.*, 2020; Moroshko *et al.*, 2020).

1.1. Contributions

We study deep linear networks $x \mapsto A_\theta x$ of depth $L \geq 1$ and widths $n_0, \cdots, n_L$, that is $A_\theta = W_L \cdots W_1$ where $W_1, \ldots, W_L$ are matrices such that $W_i \in \mathbb{R}^{n_i \times n_{i-1}}$ and θ is a vector that consists of all the (learnable) parameters of the DLN, i.e. the components of the matrices $W_1, \cdots, W_L$. For

[*]Equal contribution [1]EPFL,Switzerland. Correspondence to: Arthur Jacot <arthur.jacot@netopera.net>.

any general convex cost $C : \mathbb{R}^{n_L \times n_0} \to \mathbb{R}$ on matrices such that the zero matrix is not a global minimum, we investigate the gradient flow minimizing the loss $\mathcal{L}(\theta) = C(A_\theta)$. To ease the notation, suppose that the hidden layers have the same size, that is $w = n_1 = \cdots = n_{L-1}$ for some $w \in \mathbb{N}$.

The variance of the parameters at initialization has a profound effect on the training dynamics. If the parameters are initialized with variance $\sigma^2 = w^{-\gamma}$, where w is the size of the hidden layers, we observe a phase transition in the infinite width limit as $w \to \infty$ and show in Theorem 1 that:

- when $\gamma < 1$, the random initialization θ_0 is (with high probability) very close to a global minimum and very far from any saddle,
- when $\gamma > 1$, the initialization is very close to a saddle and far from any global minimum.

The case $\gamma < 1$ corresponds to the NTK regime (or kernel/lazy regime, described in Section 4) and the case $\gamma = 1$ corresponds to the Mean-Field limit (or the Maximal Update parametrization of (Yang, 2019)). It appears that the case $\gamma > 1$ has been much less studied in previous works.

To understand this regime, we investigate in Section 5 the case $\gamma \to +\infty$. More precisely, we fix the width of the network and let the variance at initialization go to zero. We show in Theorem 2 that the gradient flow trajectory asymptotically goes from the saddle at the origin $\vartheta^0 = 0$ to a rank-one saddle ϑ^1, i.e. a saddle where the matrices $W_1, \dots, W_L$ are of rank 1. The proof is based on a new description (Theorem 5), in the spirit of the Hartman-Grobman theorem, of the so-called fast escape paths at the origin. This theorem may be of independent interest.

We propose the Conjecture 3, backed by numerical experiments, describing the full gradient flow when the variance at initialization is very small, suggesting that it goes from saddle to saddle, visiting the neighborhoods of a sequence of critical points $\vartheta^0, \dots, \vartheta^K$ (the first K ones being saddle points, the last one being either a global minimum or a point at infinity) corresponding to matrices of increasing ranks. This is consistent with (Gissin *et al.*, 2020) which shows that incremental learning occurs in a toy model of DLNs and that gradient-based optimization hence has an implicit bias towards simple (sparse) solutions.

In Section 5.3, we show how this Saddle-to-Saddle dynamics can be described using a greedy low-rank algorithm which bears similarities with that of (Li *et al.*, 2020) and leads to a low-rank bias of the final learned function. This is in stark contrast to the NTK regime which features no low-rank bias.

1.2. Related Works

The existence of distinct regimes in the training dynamics of DNNs has been explored in previous works, both theoretically (Chizat & Bach, 2018a; Yang, 2019) and empirically (Geiger *et al.*, 2019). The theoretical works (Chizat & Bach, 2018a; Yang, 2019) have mostly focused on the transition from the NTK regime ($\gamma < 1$) to the Mean-Field regime ($\gamma = 1$). This paper is focused on the regime beyond the critical one ($\gamma > 1$).

Our study of the Saddle-to-Saddle dynamics can also be understood as a generalization of the works (Saxe *et al.*, 2014; Advani & Saxe, 2017; Saxe *et al.*, 2019; Gidel *et al.*, 2019; Arora *et al.*, 2019b) which describe a similar plateau effect in a very specific setting and with a very carefully chosen initialization.

Shortly after the initial publication of this article, we came aware of the paper (Li *et al.*, 2020) which provides a similar description to our Saddle-to-Saddle dynamics. For shallow networks, the results are almost equivalent, although the techniques are very different, especially when dealing with the fact that the escape directions (and escape paths) are unique only up to rotations. The paper (Li *et al.*, 2020) uses a clever trick that allows them to both study the dynamics of the output matrix $A_{\theta(t)}$, without the need to keep track of the parameters, and obtain a unicity property for the asymptotic dynamics. Instead, we focus on the dynamics of the parameters, give an identification of all optimal escape paths, and show that the path followed by the parameters' dynamics is unique up to symmetries of the network. Note also that, as in our paper, (Li *et al.*, 2020) only proves the first step of the Saddle-to-Saddle regime: for the subsequent steps, it is assumed that the next saddle is not approached along a bad' direction (as we discuss in Section 5.5). For deep networks, our results are more general as they hold for more general initializations than in (Li *et al.*, 2020). Indeed, in order to avoid the non-uniqueness problem of the escape paths in the space of parameters, their analysis relies heavily on the assumption that the weights of the network are balanced at initialization, and thus during training. Because we do not rely on this trick, our analysis does not require a balanced initialization.

2. Deep Linear Networks

2.1. Setup

A DLN of depth L and widths $n_0, \dots, n_L$ is the composition of L matrices

$$A_\theta = W_L \cdots W_1$$

where $W_\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$. The number of parameters is $P = \sum_{\ell=1}^{L} n_{\ell-1} n_\ell$ and we denote by $\theta = (W_L, \dots, W_1) \in \mathbb{R}^P$ the vector of parameters. The input dimension, resp.

the output dimension is n_0, resp. n_L. All parameters are initialized as i.i.d. $\mathcal{N}(0, \sigma^2)$ Gaussian random variables.

We will focus on the so-called rectangular networks, in which the number of neurons in all hidden layers is the same, i.e. $n_1 = \cdots = n_{L-1} = w$. Such rectangular network is called a (L, w)-DLN, and its number of parameters is denoted by $P = P_{(L,w)} = n_0 w + (L-2)w^2 + w n_L$. The proofs given in this article can be extended to the non-rectangular case, but this leads to more complex notations.

We study the dynamics of gradient descent on the loss $\mathcal{L}(\theta) = C(A_\theta)$ for a general differentiable and convex cost C on $n_L \times n_0$ matrices. To ensure a non-trivial minimisation problem, we assume that the null matrix is not a global minimum of C : in this case, the origin in the parameter space is a saddle of $\mathcal{L}$. Given a starting point $\theta_0 \in \mathbb{R}^P$, we denote by $t \mapsto \Gamma(t, \theta_0)$ the gradient flow path on the cost $\mathcal{L}(\theta)$ starting from θ_0, i.e. $\Gamma(0, \theta_0) = \theta_0$ and $\partial_t \Gamma(t, \theta_0) = -\nabla \mathcal{L}(\Gamma(t, \theta_0))$.

While our analysis applies to general twice differentiable costs C, the typical costs used in practice are:

The *Mean-Squared Error* (MSE) loss $C(A) = \frac{1}{N} \|AX - Y\|_F^2$ for some inputs $X \in \mathbb{R}^{n_0 \times N}$ and labels $Y \in \mathbb{R}^{n_L \times N}$, where $|| \cdot ||_F$ is the Frobenius norm.

The *Matrix Completion* (MC) loss $C(A) = \frac{1}{N} \sum_{i=1}^N (A_{k_i,m_i} - A^*_{k_i,m_i})^2$ for some true matrix A^* of which we observe only the N entries $A^*_{k_1,m_1}, \ldots, A^*_{k_N,m_N}$.

2.2. Symmetries and Invariance

A key tool in this paper is the use of two important symmetries of the parametrization map $\theta \mapsto A_\theta$ in DLNs: rotations of hidden layers and inclusions in wider DLNs.

Rotations: A $L-1$ tuple $R = (O_1, \ldots, O_{L-1})$ of orthogonal $w \times w$ matrices is called a w-width network rotation, or in short a rotation. A rotation R acts on a parameter vector $\theta = (W_L, \ldots, W_1)$ as $R\theta = (W_L O_{L-1}^T, O_{L-1} W_{L-1} O_{L-2}^T, \ldots, O_1 W_1)$. The space of rotations is an important symmetry of DLN: indeed, for any parameter θ, and any cost C, the two following important properties hold:

$$A_{R\theta} = A_\theta, \quad \nabla_\theta C(A_{R\theta}) = R \nabla_\theta C(A_\theta),$$

where we considered $\nabla_\theta C(A_\theta) \in \mathbb{R}^{P_{L,w}}$ as another vector of parameters. These properties imply that if $\theta(t) = \Gamma(t, \theta_0)$ is a gradient flow path, then so is $R\theta(t) = \Gamma(t, R\theta_0)$.

Inclusion: The inclusion $I^{(w \to w')}$ of a network of width w into a network of width $w' > w$ (by adding zero weights on the new neurons) is defined as $I^{(w \to w')}(\theta) = (V_L, \ldots, V_1)$ with

$$V_1 = \begin{pmatrix} W_1 \\ 0 \end{pmatrix}, V_\ell = \begin{pmatrix} W_\ell & 0 \\ 0 & 0 \end{pmatrix}, V_L = \begin{pmatrix} W_L & 0 \end{pmatrix}.$$

For any parameters θ and any cost C, we have $A_{I^{(w \to w')}\theta} = A_\theta$ and $\nabla C(A_{I^{(w \to w')}\theta}) = I^{(w \to w')} \nabla C(A_\theta)$: the image of the inclusion map $I^{(w \to w')}$ (as well as any rotation $R \mathrm{Im} I^{(w \to w')}$ thereof) is invariant under gradient flow.

3. Proximity of Critical Points at Initialization

It has already been observed that in the infinite width limit, when the width w of the network grows to infinity, the scale at which the variance σ^2 of the parameters at initialization scales with the width can lead to very different behaviors (Chizat & Bach, 2018a; Geiger *et al.*, 2019; Yang, 2019). Let us consider scaling of the variance $\sigma^2 = w^{-\gamma}$ for $\gamma \geq 1 - \frac{1}{L}$. The reason we lower bound γ is that any smaller γ would lead to an explosion of the variance of the matrix A_θ at initialization as the width w grows.

Let d_{m} and d_{s} be the Euclidean distances between the initialization θ and, respectively, the set of global minima and the set of all saddles. For random variables $f(w), g(w)$ which depend on w, we write $f \asymp g$ if both $f(w)/g(w)$ and $g(w)/f(w)$ are stochastically bounded as $w \to \infty$. The following theorem studies how d_{m} and d_{s} scale as $w \to \infty$:

Theorem 1. *Suppose that the set of matrices that minimize C is non-empty, has Lebesgue measure zero, and does not contain the zero matrix. Let θ be i.i.d. centered Gaussian r.v. of variance $\sigma^2 = w^{-\gamma}$ where $1 - \frac{1}{L} \leq \gamma < \infty$. Then:*

1. *if $1 - \frac{1}{L} \leq \gamma < 1$, we have $d_{\mathrm{m}} \asymp w^{-\frac{(1-\gamma)(L-1)}{2}}$ and $d_{\mathrm{s}} \asymp w^{\frac{1-\gamma}{2}}$,*
2. *if $\gamma = 1$, we have $d_{\mathrm{m}}, d_{\mathrm{s}} \asymp 1$,*
3. *if $\gamma > 1$ we have $d_{\mathrm{m}} \asymp 1$ and $d_{\mathrm{s}} \asymp w^{-\frac{\gamma-1}{2}}$.*

This theorem shows an important change of behavior between the case $\gamma < 1$ and $\gamma > 1$. When $\gamma < 1$, the network is initialized very close to a global minimum and far from any saddle. When $\gamma > 1$, the parameters are initialized very close to a saddle but far away from any global minimum. The critical case $\gamma = 1$ is the unique limit where both types of critical points are at the same distance from the initialization.

Hence, the landscape of the loss near the initialization displays distinct features in the three regimes highlighted in the previous theorem. In fact, the dynamics of the gradient descent also exhibits very distinctive characteristics in the different regimes. In Appendix B.1.1, we show that the

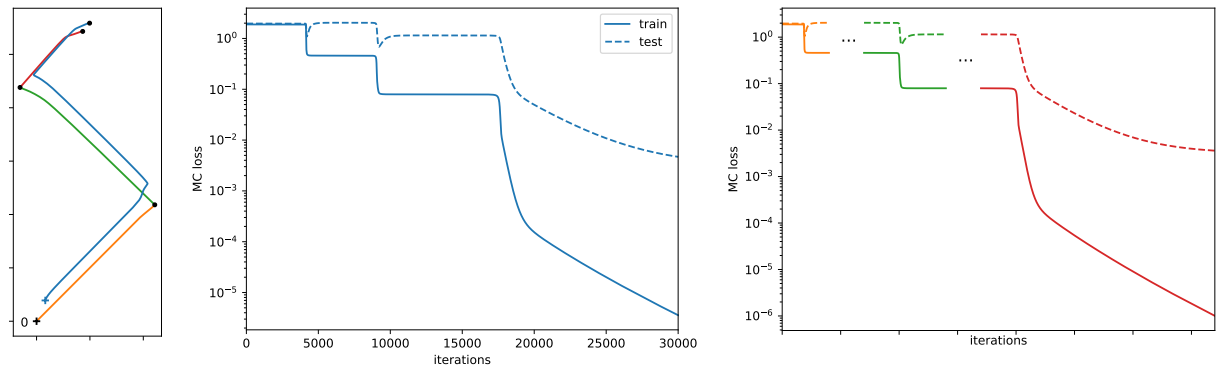

Figure 1. Saddle-to-Saddle dynamics: A DLN ($L = 4, w = 100$) with a small initialization ($\gamma = 2$) trained on a MC loss fitting a 10×10 matrix of rank 3. **Left:** Projection onto a plane of the gradient flow path θ_α in parameter space (in blue) and of the sequence of 3 paths $\theta^1, \theta^2, \theta^3$ (in orange, green and red), described by Algorithm $\mathcal{A}_{\epsilon,T,\eta}$, starting from the origin (+) and passing through 2 saddles ($\cdot$) before converging. **Middle:** Train (solid) and test (dashed) MC costs through training. We observe three plateaus, corresponding to the three saddles visited. **Right:** The train (solid) and test (dashed) losses of the three paths plotted sequentially, in the saddle-to-saddle limit; the dots represent an infinite amount of steps separating these paths.

largest initialization, corresponding to the choice $\gamma = 1 - \frac{1}{L}$, is equivalent to the so-called NTK parametrization of (Jacot *et al.*, 2018), up to a rescaling of the learning rate. In the range $1 - \frac{1}{L} < \gamma < 1$, (Yang & Hu, 2020) obtain a similar, yet slightly different, kernel regime. The initialization $\gamma = 1$ corresponds to the Mean-Field limit for shallow networks (Chizat & Bach, 2018b; Rotskoff & Vanden-Eijnden, 2018) or, more generally, to the Maximal Update parametrization (Yang & Hu, 2020) (see Appendix B.1.2). The case $\gamma > 1$ is however much less studied and is difficult to study since the initialization approaches a saddle as $w \to \infty$. Thus, in this regime, the wider the network, the longer it takes to escape this nearby saddle and, in the limit as $w \to \infty$, nothing happens over a finite number of gradient steps. With the right time parametrization, we will observe interesting Saddle-to-Saddle dynamics in this regime, leading to some low-rank bias.

4. NTK regime: $\gamma < 1$

The NTK for linear networks can be expressed easily using the tensor

$$\Theta^{(L)} = \sum_\theta \partial_\theta A \otimes \partial_\theta A,$$

which entries are given by $[\Theta^{(L)}]^{j,l}_{i,k} = (\nabla_\theta (A_\theta)_{i,j})^T (\nabla_\theta (A_\theta)_{k,l})$, for $i, k = 1, \dots, n_L$ and $j, l = 1, \dots, n_0$. For any x, y in $\mathbb{R}^{n_0}$, the value of the NTK at x and y is $\Theta^{(L)}(x \otimes y) = \sum_{j,l} [\Theta^{(L)}]^{j,l}_{\cdot,\cdot} x_j y_l$.

When the parameters evolve according to the gradient flow on $\mathcal{L}(\theta) = C(A_\theta)$, the dynamics of $A_{\theta(t)}$ is:

$$\begin{aligned} \partial_t A_{\theta(t)} &= -\Theta^{(L)} \cdot \nabla_A C(A_{\theta(t)}) \\ &= -\sum_{k,l} [\Theta^{(L)}]^{\cdot,l}_{\cdot,k} \frac{d}{dA_{k,l}} C(A_{\theta(t)}), \end{aligned}$$

where $\cdot$ denotes a contraction of the k, l indices of $\Theta^{(L)}$ with the two indices of $\nabla_A C(A_{\theta(t)})$.

At initialization, $\Theta^{(L)}$ concentrates around its expectation $\mathbb{E}\left[\Theta^{(L)}\right] = L w^{(1-\gamma)(L-1)} \delta_{i,k} \delta_{j,l}$ as the width grows. It was first proven in (Jacot *et al.*, 2018) that for an initialization equivalent to the case $\gamma = 1 - \frac{1}{L}$ (see Appendix B.1.1 for more details), as $w \to \infty$ the NTK remains constant during training. Recent results (Yang & Hu, 2020) have shown that the NTK is asymptotically fixed for all $\gamma \in (1 - \frac{1}{L}, 1)$. In this case, given the asymptotic behavior of the NTK, the evolution of $A_{\theta(t)}$ is the same (up to a change of learning rate) as the one obtained by performing directly a gradient flow on the cost C.

As a result, in the regime $\gamma < 1$, if the cost C is strictly convex (or satisfies the Polyak-Lojasievicz inequality (Liu *et al.*, 2020)), the loss decays exponentially fast. Besides, the depth of the network has no effect in the infinite width limit (except for a change of learning rate) and the DLN structure adds no specific bias to the global minimum learned with gradient descent. In particular, this regime leads to no low-rank bias.

5. Saddle-to-Saddle Dynamics: $\gamma \gg 1$

We now study the dynamics of DLN during training as the variance at initialization goes to zero. Specifically, we

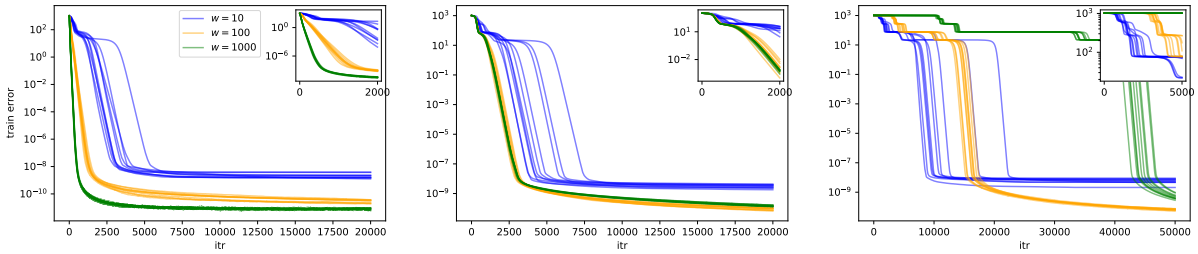

(a) $\gamma = 0.75$ (NTK) **(b)** $\gamma = 1$ (MF) **(c)** $\gamma = 1.5$ (S-S)

*Figure 2. Training in **(a)** the NTK regime, **(b)** mean-field, **(c)** saddle-to-saddle regimes in deep linear networks for three widths $w = 10, 100, 1000$, $L = 4$, and 10 seeds.* Parameters are initialized with variance $\sigma^2 = w^{-\gamma}$. We observe that **(a)** in the NTK regime, the training loss shows typical linear convergence behavior for $w = 1000$ and $w = 100$; **(b)** in the mean-field regime, we observe that even the large width networks approach to a saddle at the beginning of the training and that the length of the plateaus remains constant between widths $w = 1000$ and $w = 100$; **(c)** in the saddle-to-saddle regime, the plateaus become longer as the width grows. In all cases, we see a reduction in the variation between the different seeds as $w \to \infty$.

sample some random parameters θ_0 with i.i.d. $\mathcal{N}(0,1)$ entries, consider the gradient flow $\theta_\alpha(t) = \Gamma(t, \alpha\theta_0)$, and let $\alpha \searrow 0$. Since the origin is a saddle, for all fixed times t, $\lim_{\alpha\searrow 0} \theta_\alpha(t) = 0$. We will show however that there is an escape time t_α, which grows to infinity as $\alpha \searrow 0$, such that the limit $\lim_{\alpha\searrow 0} \theta_\alpha(t_\alpha + t)$ is non-trivial for all $t \in \mathbb{R}$.

The study of shallow networks ($L = 2$) is facilitated by the fact that the saddle at the origin is strict: its Hessian has negative eigenvalues. For deeper networks ($L > 2$), the saddle is highly degenerate: the $L - 1$ first order derivatives vanish. In Section 5.4, we develop new theoretical tools to analyze the two types of saddles and their escape paths.

5.1. First Path

It turns out that gradient flow paths naturally escape the saddle at the origin along so-called *optimal escape paths*. We say that a gradient flow path $\theta(\cdot) : \mathbb{R} \to \mathbb{R}^P$ is an *escape path* of a critical point θ^* if $\lim_{t\to-\infty} \theta(t) = \theta^*$. Informally, the optimal escape paths, whose precise definition is given in Section 5.4, are the escape paths that allow the fastest exit from a saddle. In DLNs, these optimal escape paths are of the form $RI^{(1\to w)}\underline{\theta}^1(t)$ where $\underline{\theta}^1(t)$ is a path of a width 1 DLN which escapes from the origin:

Theorem 2. *Assume that the largest singular value s_1 of the gradient of C at the origin $\nabla C(0) \in \mathbb{R}^{n_L \times n_0}$ has multiplicity 1. There is a deterministic gradient flow path $\underline{\theta}^1$ in the space of width-1 DLNs such that, with probability 1 if $L \leq 3$, and probability at least $1/2$ if $L > 3$, there exists an escape time t^1_α and a rotation R such that*

$$\lim_{\alpha\to 0} \theta_\alpha(t^1_\alpha + t) = RI^{(1\to w)}\underline{\theta}^1(t).$$

The unicity of the largest singular value of the gradient at the origin guarantees the unicity (up to rotation) of the optimal escape paths. For example, with the MSE loss, the gradient at the origin is $2YX^T$: for generic Y and X, the largest singular value of the gradient has a multiplicity of 1.

The reason why, for DLN with $L > 3$, we can only guarantee a probability of $\frac{1}{2}$ in the previous theorem, is that we need to ensure that gradient descent does not get stuck at the saddle at the origin or at other saddles connected to it. For $L = 2$, this follows from the fact that the saddle is strict. When $L > 2$, the saddle is not strict and we were only able to prove it in the case where $L = 3$. We conjecture that the behavior described in Theorem 2 happens with probability 1 for all $L \geq 2$.

As shown in the Appendix C.5, the escape time t_α is of order $-\log \alpha$ for shallow networks and of order $\alpha^{-(L-2)}$ for networks of depth $L > 2$. Hence, the deeper the network, the slower the gradient flow escapes the saddle.

Besides, as also discussed in the Appendix C.5, the norm $\left\|\theta^1(t)\right\|$ of the limiting escape path $\theta^1(t) = RI^{(1\to w)}\underline{\theta}^1(t)$ grows at an optimal speed: as $e^{s^*(t+T)}$ for some T when $L = 2$ and as $(s^*(L-2)(T-t))^{-\frac{1}{L-2}}$ for some T when $L > 2$, where s^* is the optimal escape speed $s^* = L^{-\frac{L-2}{2}} s_1$. These are optimal in the sense that given an other gradient flow path $\theta(t)$ which exits from the origin, there exists a ball B centered at the origin such that, for any small ϵ, if t_1 and t_2 are the times such that $\left\|\theta^1(t_1)\right\| = \epsilon = \|\theta(t_2)\|$, then $\left\|\theta^1(t + t_1)\right\| \geq \|\theta(t + t_2)\|$ for any positive t, until one of the paths exits the ball B.

5.2. Subsequent Paths

What happens after this first path? The width-1 gradient flow path $\underline{\theta}^1(t)$ converges to a width-1 critical point $\underline{\vartheta}^1$ as $t \to \infty$. While $\underline{\vartheta}^1$ may be a local minimum amongst

width-1 DLNs, its inclusion $\vartheta^1 = RI^{(1\to w)}(\underline{\vartheta}^1)$ will be a saddle assuming it is not a global minimum already and that the network is wide enough, since if $w \geq \min\{n_0, n_L\}$ all critical points are either global minima or saddles (Nouiehed & Razaviyayn, 2021).

Theorem 2 guarantees that, as $\alpha \searrow 0$, the gradient flow path $\theta_\alpha(t)$ will approach the saddle ϑ^1. It is then natural to assume that $\theta_\alpha(t)$ will escape this saddle along an optimal escape path (which is the inclusion of a width-2 path). Repeating this process, we expect gradient flow to converge as $\alpha \searrow 0$ to the concatenation of paths going from saddle to saddle of increasing width:

Conjecture 3. *With probability* 1*, there exist* $K+1$ *critical points* $\vartheta^0, \dots, \vartheta^K \in \mathbb{R}^{P_{L,w}}$ *(with* $\vartheta^0 = 0$*) and* K *gradient flow paths* $\theta^1, \dots, \theta^K : \mathbb{R} \to \mathbb{R}^{P_{L,w}}$ *connecting the critical points (i.e.* $\lim_{t\to-\infty} \theta^k(t) = \vartheta^{k-1}$ *and* $\lim_{t\to+\infty} \theta^k(t) = \vartheta^k$*) such that the path* $\theta_\alpha(t)$ *converges as* $\alpha \to 0$ *to the concatenation of* $\theta^1(t), \dots, \theta^K(t)$ *in the following sense: for all* $k < K$*, there exist times* t^k_α *(which depend on* θ_0*) such that*

$$\lim_{\alpha\to 0} \theta_\alpha(t^k_\alpha + t) = \theta^k(t).$$

Furthermore, for all $k < K$*, there is a deterministic path* $\underline{\theta}^k(t)$ *and a local minimum* $\underline{\vartheta}^k$ *of a width-*k *network such that for some rotation* R *(which depends on* θ_0*),* $\theta^k(t) = RI^{(k\to w)}(\underline{\theta}^k(t))$ *and* $\vartheta^k = RI^{(k\to w)}(\underline{\vartheta}^k)$ *for all* k *and* t*.*

This Saddle-to-Saddle behavior explains why for small initialization scale, the train error gets stuck at plateaus during training (Figures 1 and 2). Conjecture 3 suggests that these plateaus correspond to the saddle visited.

Note that for losses such as the cross-entropy, the gradient descent may diverge towards infinity, as studied in (Soudry *et al.*, 2018; Gunasekar *et al.*, 2018a). From now on, we focus on the case where ϑ^K is a finite global minimum. By the invariance under gradient flow of $\mathrm{Im}[I^{(k\to w)}]$ (the image of the inclusion map), the inclusion of a width-k local minimum $\underline{\vartheta}^k$ into a larger network is a saddle ϑ^k (if A_{ϑ^k} is not a global minimum of C). These types of saddles are closely related to the symmetry-induced saddles studied in (Simsek *et al.*, 2021) in non-linear networks.

Remark 4. Note that each of the limiting paths θ^k and critical points ϑ^k will be *balanced* (i.e. their weight matrices satisfy $W_\ell W_\ell^T = W_{\ell+1}^T W_{\ell+1}$ for all $\ell = 1, \dots, L-1$). The origin is obviously balanced and since balancedness is an invariant of gradient flow and all other paths and saddles are connected to the origin by a sequence of gradient flow paths, they must be balanced too. Note however that for all $\alpha > 0$, the path $\theta_\alpha(t)$ is almost surely not balanced.

Algorithm $\mathcal{A}_{\epsilon,T,\eta}$

Compute the first singular vectors of $\nabla C(0)$:
$u, s, v \leftarrow \mathrm{SVD}_1(\nabla C(0))$
$\theta \leftarrow (-\epsilon v^T, \epsilon, \dots, \epsilon u)$
$w \leftarrow 1$
while $C(A_\theta) < C_{min} + \epsilon$ **do**
  # T steps of GD on the loss of width-w DLN with lr η
  $\theta \leftarrow \mathrm{SGD}_{w,T,\eta}(\theta)$
  $u, s, v \leftarrow \mathrm{SVD}_1(\nabla C(A_\theta))$
  $\theta \leftarrow \left(\begin{pmatrix} W_1 \\ -\epsilon v^T \end{pmatrix}, \begin{pmatrix} W_2 & 0 \\ 0 & \epsilon \end{pmatrix}, \dots, \begin{pmatrix} W_L & \epsilon u \end{pmatrix} \right)$
  $w \leftarrow w + 1$
end while

5.3. Greedy Low-Rank Algorithm

Conjecture 3 suggests that the gradient flow with vanishing initialization implements a greedy low-rank algorithm which performs a greedy search for a lowest-rank solution: it first tries to fit a width 1 network, then a width 2 network and so on until reaching a solution. Thus, we expect that as $\alpha \searrow 0$, the dynamics of gradient flow corresponds, up to inclusion and rotation, to the limit of the algorithm $\mathcal{A}_{\epsilon,T,\eta}$ as sequentially $T \to \infty$, $\eta \to 0$ and $\epsilon \to 0$. In particular, we used the Algorithm $\mathcal{A}_{\epsilon,T,\eta}$, with large T and small η and ϵ to approximate the paths $\underline{\theta}^k$ and points $\underline{\vartheta}^k$ in Figure 1. Note how this limiting algorithm is deterministic. This implies that even for finite widths the dynamics of gradient flow converge to a deterministic limit (up to random rotations R) as the variance at initialization goes to zero.

A similar algorithm has already been described in (Li *et al.*, 2020), however thanks to our different proof techniques, we are able to give a more precise description of the evolution of the parameters.

5.4. Description of the paths that escape a saddle

Our proof relies on a theorem which relates the escape paths of the saddle at the origin of the cost $\mathcal{L}$ and the escape paths of the L-th order Taylor approximation H of $\mathcal{L}$. This correspondence only applies to paths which escape the saddle sufficiently fast.

We define the set of fast escaping paths $\mathcal{F}_{\mathcal{L}}(s)$ of the cost $\mathcal{L}$ with speed at least s as follows:

- for shallow networks ($L = 2$), it is the set of gradient flow paths that satisfy $\|\theta(t)\| = O(e^{st})$ as $t \to -\infty$,
- for deep networks ($L > 2$), it is the set of gradient flow paths that satisfy $\|\theta(t)\| \leq (s(L-2)(T-t))^{-\frac{1}{L-2}}$ for some T and any small enough t.

The optimal escape speed is $s^* = L^{-\frac{L-2}{2}} s_1$ where s_1 is the largest singular value of $\nabla C(0)$. It is the optimal escape speed in the sense that there are no faster escape paths: $\mathcal{F}_{\mathcal{L}}(s) = \emptyset$ if $s > s^*$. Escape paths which exit the saddle at the optimal escape speed are called optimal escape paths.

There is a bijection between fast escaping paths of the loss $\mathcal{L}$ and those of its Lth order Taylor approximation H:

Theorem 5. ***Shallow networks:*** *for all s s.t. $s > \frac{1}{3}s^*$ there is a unique bijection $\Psi : \mathcal{F}_{\mathcal{L}}(s) \to \mathcal{F}_H(s)$ such that for all paths $\theta \in \mathcal{F}_{\mathcal{L}}(s)$, $\|\theta(t) - \Psi(\theta)(t)\| = O(e^{3st})$ as $t \to -\infty$.*

Deep networks: *for all $s > \frac{L-1}{L+1}s^*$, there is a unique bijection $\Psi : \mathcal{F}_{\mathcal{L}}(s) \to \mathcal{F}_H(s)$ such that for all paths $\theta \in \mathcal{F}_{\mathcal{L}}(s)$, $\|\theta(t) - \Psi(\theta)(t)\| = O((-t)^{-\frac{L+1}{L-2}})$ as $t \to -\infty$.*

We believe that this theorem is of independent interest, and it is stated in a more general setting in the Appendix. Theorem 5 is similar to the Hartman-Grobman Theorem, which shows a bijection, in the vicinity of a critical point, between the gradient flow paths of $\mathcal{F}$ and of its linearization. The bijection in Theorem 5 holds only between fast escaping paths, but it gives stronger guarantees regarding how close the paths $\theta(\cdot)$ and $\Psi(\theta)(\cdot)$ are. In particular, Theorem 5 guarantees that a fast escaping path $\theta(\cdot)$ and its image $\Psi(\theta)(\cdot)$ have the same 'escape speed', whereas the correspondence between paths of in the Hartman-Grobman theorem does not in general conserve speed. This is due to the fact that the homeomorphism which allows to construct the bijection in the Hartman-Grobman theorem is only Hölder continuous. This suggests that fast escaping paths can be guaranteed to conserve their speed after the Taylor approximation while slower paths can change speed. Finally, our result has the significant advantage that it may be applied to higher order Taylor approximations, whereas the Hartman-Grobman Theorem only applies to the linearization of the flow (i.e. it could only be useful in the shallow case $L = 2$).

5.5. Sketch of Proof

In this section, we provide a sketch of proof for Theorem 2.

We fix some small $r > 0$ independent of α. The *escape time* t_α is the earliest time such that $\|\theta_\alpha(t_\alpha)\| = r$. We show that the limiting escape path $(\theta^1(t))_{t\in\mathbb{R}}$ as $\theta^1(t) = \lim_{\alpha\searrow 0} \theta_\alpha(t_\alpha + t)$ is well defined and non-trivial since $\theta^1(0) \neq 0$. The next step of the proof is to show that θ^1 escapes the saddle at an almost optimal speed: for any $\epsilon > 0$, for some T and any small enough t, for shallow network $\left\|\theta^1(t)\right\| = O(e^{(s^*-\epsilon)t})$, and for deeper networks $\left\|\theta^1(t)\right\| \leq [(L-2)(s^*-\epsilon)(T-t)]^{-\frac{1}{L-2}}$. We may therefore apply Theorem 5: there exists a unique optimal escape path for the L-th order Taylor approximation H around the origin which is 'close', in the sense given in Theorem 5, to θ^1.

For the Taylor approximation H, we have a precise description of the optimal escaping paths for the saddle at the origin. Assuming that the largest singular value s_1 of the gradient matrix $\nabla C(0)$ has multiplicity one, all optimal escape paths of H (i.e. the set of paths that escape with the largest speed) are of the form $\theta_H(t) = d(t)RI^{(1\to w)}(\rho)$ where R is some rotation, the scalar function $d(t)$ is equal to $e^{s^*(t+T)}$ for shallow networks and $(s^*(L-2)(T-t))^{-\frac{1}{L-2}}$ for deep networks, and the vector of parameters ρ is given by:

$$\rho = \left(v_1^T, 1, \ldots, 1, u_1\right)$$

with u_1, v_1 the left and right singular vectors of the largest singular value s_1 of the gradient matrix $\nabla C(0)$.

Let us consider the unique optimal escape path $\theta_H(t) = d(t)RI^{(1\to w)}(\rho)$ for H which is 'close' to θ^1. The path $\underline{\theta}_H(t) = d(t)\rho$ is also an optimal escape path for H: from Theorem 5, there exists a unique optimal escape path $\underline{\theta}(t)$ which is 'close' to $\underline{\theta}_H$. The former escape path corresponds to a 1-width DLN and it is easy to show that $RI^{(1\to w)}(\underline{\theta})$ is an optimal escape path for $\mathcal{L}$ which is 'close' to $RI^{(1\to w)}(\underline{\theta}_H) = \theta_H$.

In particular, we obtain that both θ^1 and $RI^{(1\to w)}(\underline{\theta})$ are optimal escape path for $\mathcal{L}$ which are 'close' to θ_H. By the unicity property in Theorem 5, we obtain that $\theta^1 = RI^{(1\to w)}(\underline{\theta})$ which allows us to conclude.

Remark 6. To prove Conjecture 3, one needs to apply a similar argument to understand how gradient flow escapes the subsequent saddles $\vartheta^1, \ldots, \vartheta^K$. There are two issues:

First, even though Theorem 2 guarantees that gradient descent will come arbitrarily close to the next saddle ϑ^1, it may not approach it along a generic direction: it could approach along a "bad" direction. For the first path, we relied on the fact that θ_0 is Gaussian to guarantee that these bad directions are avoided with probability 1 (or $1/2$). Note that this problem c

ould be addressed using the so-called perturbed stochastic gradient descent described in (Jin *et al.*, 2017; Du *et al.*, 2017) since, in this learning algorithm, once in the vicinity of the saddle, a small Gaussian noise is added to the parameters: as a consequence, they end up being in a generic position in the neighborhood of the saddle.

Second, for deep networks ($L > 2$), the saddle ϑ^1 has a different local structure to ϑ^0. Indeed, at the origin, the $L-1$ first derivatives vanish, leading to an (approximately) L-homogeneous saddle at the origin. On the contrary, at the rank 1 saddle $\vartheta^1 = RI^{(1\to w)}(\underline{\vartheta}^1)$, if $\underline{\vartheta}^1$ is a local minimum of the width 1 network, the Hessian is positive along the inclusion $\mathrm{Im}\left[RI^{(1\to w)}\right]$. This implies that the dynamics can only escape the saddle through the Hessian null-space, along which the first $L-1$ derivatives vanish. Although

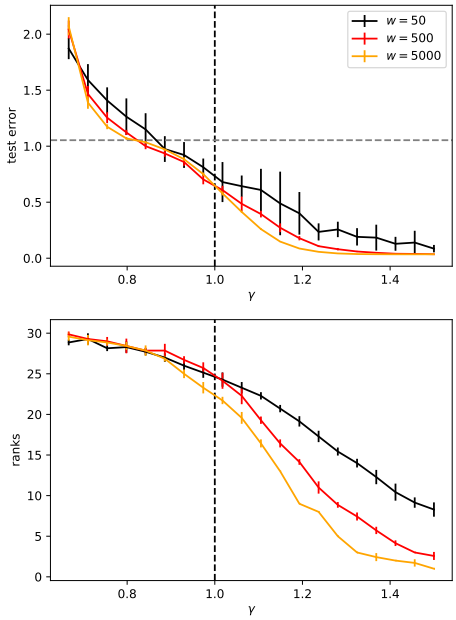

Figure 3. Test errors and ranks at convergence as a function of initialization scale γ, matrix completion task. The task is finding a matrix of size 30×30 and rank 1 from 20% of its entries. The test error and ranks are averaged over 7 seeds (± 1 standard deviations are reported in the error bar). In the NTK regime, the solutions at convergence are almost full-rank and the test error is roughly the same or worse than that of the zero predictor. On the other hand in the Saddle-to-Saddle regime the test error approaches zero. As the width grows the transition between regimes becomes sharper and the test error becomes more consistent within each regimes.

the loss restricted to this null-space around ϑ_1 has a similar structure to the loss around the origin, the fact that the Hessian at ϑ_1 is not null complexifies the analysis.

6. Characterization of the Regimes of Training

In light of the results presented in this paper, we discuss the three regimes that can be obtained by varying the initialization scale γ: the kernel regime ($\gamma < 1$), the Mean-Field regime ($\gamma = 1$) and the Saddle-to-Saddle regime ($\gamma > 1$).

The NTK limit ($\gamma = 1 - \frac{1}{L}$) (Jacot *et al.*, 2018; Lee *et al.*, 2019a) is representative of the other scalings $1 - \frac{1}{L} \leq \gamma < 1$ (Yang & Hu, 2020). The critical regime $\gamma = 1$ corresponds to the Mean-Field limit for shallow networks (Chizat & Bach, 2018b; Rotskoff & Vanden-Eijnden, 2018) or the Maximal Update parametrization for deep networks (Yang & Hu, 2020). Finally, we conjecture that the last regime where $\gamma > 1$, displays features very akin to the $\gamma = +\infty$ case studied in this article. Under this assumption, we obtain the following list of properties that characterize each of these regimes:

In the **NTK regime** ($1 - \frac{1}{L} \leq \gamma < 1$):

1. During training, the parameters converge to a nearby global minimum, and do not approach any saddle (Figure 2a shows how the plateaus disappear as w grows).
2. If the cost on matrices C is strictly convex, one can guarantee exponential decrease of the loss (i.e. linear convergence).
3. The NTK is asymptotically fixed during training.
4. No low-rank bias in the learned matrix - as a result the test error for matrix completion is the same (or even larger) than the zero predictor in the NTK regime, as shown in Figure 3.

The **Saddle-to-Saddle regime** ($\gamma > 1$):

1. The parameters start in the vicinity of a saddle and visit a sequence of saddles during training. They come closer to each of these saddles as the width grows.
2. As the width grows, it takes longer to escape each saddle, leading to long plateaus for the training error. The training time is therefore asymptotically infinite (see Figure 2c).
3. The rate of change $\|\Theta(\theta_T) - \Theta(\theta_0)\|$ (where $T \in \mathbb{R}$ is the stopping time) of the NTK is infinitely larger than the NTK at initialization $\|\Theta(\theta_0)\|$. This follows from the fact that the NTK at initialization goes to zero, while it has finite size at the end of training.
4. The learned matrix is the result of a greedy algorithm that finds the lowest rank solution.

The **Mean-Field regime** $\gamma = 1$ lies at the transition between the two previous regimes and is more difficult to characterize:

1. In this critical regime, the constant factor c in the variance at initialization $\sigma^2 = cw^{-\gamma}$ can have a strong effect on the dynamics.
2. Plateaus can still be observed (see Figure 2b), however in contrast to the Saddle-to-Saddle regime, the length of the plateaus does not increase as the width grows, but remains roughly constant.
3. The NTK and its rate of change are of same order.

In general, we observe some tradeoff: the NTK regime leads to fast convergence without low-rank bias, while the Saddle-to-Saddle regime leads to some low-rank bias, but at the cost of an asymptotically infinite training time.

7. Conclusion

We propose a simple criterion to identify three regimes in the training of large DLNs: the distances from the initialization to the nearest global minimum and to the nearest saddle. The NTK regime ($1 - \frac{1}{L} \leq \gamma < 1$) is characterized by an initialization which is close to a global minimum and far from any saddle, the Saddle-to-Saddle regime ($\gamma > 1$) is characterized by an initialization which is close to a saddle and (comparatively) far from any global minimum and, finally, in the critical Mean-Field regime ($\gamma = 1$), these two distances are of the same order as the width grows.

While the NTK and Mean-Field limits are well-studied, the Saddle-to-Saddle regime is less understood. We therefore investigate the case $\gamma = +\infty$ (i.e. we fix the width and let the variance at initialization go to zero). In this limit, the initialization converges towards the saddle at the origin $\vartheta^0 = 0$. We show that gradient flow naturally escapes this saddle along an 'optimal escape path' along which the network behaves as a width-1 network. This leads the gradient flow to subsequently visit a second saddle ϑ^1 which has the property that the matrix A_{ϑ^1} has rank 1. We conjecture that the gradient flow next visits a sequence of critical points $\vartheta^2, \dots, \vartheta^K$ of increasing rank, implementing some form of greedy low-rank algorithm. These saddles explain the plateaus in the loss curve which are characteristic of the Saddle-to-Saddle regime.

Similar plateaus can be observed in non-linear networks: this suggests that the regimes and dynamics described in this paper could be generalized to non-linear networks.

References

Advani, Madhu S., & Saxe, Andrew M. 2017. *High-dimensional dynamics of generalization error in neural networks*.

Arora, Sanjeev, Cohen, Nadav, Golowich, Noah, & Hu, Wei. 2019a. A Convergence Analysis of Gradient Descent for Deep Linear Neural Networks. *In: International Conference on Learning Representations*.

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

Arora, Sanjeev, Du, Simon S, Hu, Wei, Li, Zhiyuan, Salakhutdinov, Ruslan, & Wang, Ruosong. 2019c. On Exact Computation with an Infinitely Wide Neural Net. *arXiv preprint arXiv:1904.11955*.

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

Chizat, Lenaic, & Bach, Francis. 2018a. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*.

Chizat, Lénaïc, & Bach, Francis. 2018b. On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport. *Pages 3040–3050 of: Advances in Neural Information Processing Systems 31*. Curran Associates, Inc.

Chizat, Lénaïc, & Bach, Francis. 2020. Implicit Bias of Gradient Descent for Wide Two-layer Neural Networks Trained with the Logistic Loss. *Pages 1305–1338 of:* Abernethy, Jacob, & Agarwal, Shivani (eds), *Proceedings of Thirty Third Conference on Learning Theory*. Proceedings of Machine Learning Research, vol. 125. PMLR.

de G. Matthews, Alexander G., Hron, Jiri, Rowland, Mark, Turner, Richard E., & Ghahramani, Zoubin. 2018. Gaussian Process Behaviour in Wide Deep Neural Networks. *In: International Conference on Learning Representations*.

Du, Simon S., Jin, Chi, Lee, Jason D., Jordan, Michael I., Póczos, Barnabás, & Singh, Aarti. 2017. Gradient descent can take exponential time to escape saddle points. *Pages 1067–1077 of: Proceedings of the 31st International Conference on Neural Information Processing SystemsDecember 2017, NIPS'17*. Curran Associates, Inc.

Du, Simon S., Zhai, Xiyu, Poczos, Barnabas, & Singh, Aarti. 2019. Gradient Descent Provably Optimizes Over-parameterized Neural Networks. *In: International Conference on Learning Representations*.

Eftekhari, Armin. 2020. Training Linear Neural Networks: Non-Local Convergence and Complexity Results. *Pages 2836–2847 of:* III, Hal Daumé, & Singh, Aarti (eds), *Proceedings of the 37th International Conference on Machine Learning*. Proceedings of Machine Learning Research, vol. 119. PMLR.

Geiger, Mario, Spigler, Stefano, Jacot, Arthur, & Wyart, Matthieu. 2019. Disentangling feature and lazy learning in deep neural networks: an empirical study. *CoRR*, **abs/1906.08034**.

Gidel, Gauthier, Bach, Francis, & Lacoste-Julien, Simon. 2019. Implicit Regularization of Discrete Gradient Dynamics in Linear Neural Networks. *In:* Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., & Garnett, R. (eds), *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc.

Gissin, Daniel, Shalev-Shwartz, Shai, & Daniely, Amit. 2020. The Implicit Bias of Depth: How Incremental Learning Drives Generalization. *In: International Conference on Learning Representations*.

Gunasekar, Suriya, Lee, Jason, Soudry, Daniel, & Srebro, Nathan. 2018a. Characterizing Implicit Bias in Terms of Optimization Geometry. *Pages 1832–1841 of:* Dy, Jennifer, & Krause, Andreas (eds), *Proceedings of the 35th International Conference on Machine Learning*. Proceedings of Machine Learning Research, vol. 80. PMLR.

Gunasekar, Suriya, Lee, Jason D, Soudry, Daniel, & Srebro, Nati. 2018b. Implicit Bias of Gradient Descent on Linear Convolutional Networks. *In:* Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., & Garnett, R. (eds), *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc.

Huang, Jiaoyang, & Yau, Horng-Tzer. 2019. Dynamics of deep neural networks and neural tangent hierarchy. *arXiv preprint arXiv:1909.08156*.

Jacot, Arthur, Gabriel, Franck, & Hongler, Clément. 2018. Neural Tangent Kernel: Convergence and Generalization in Neural Networks. *Pages 8580–8589 of: Advances in Neural Information Processing Systems 31*. Curran Associates, Inc.

Ji, Ziwei, & Telgarsky, Matus. 2018. Gradient descent aligns the layers of deep linear networks. *CoRR*, **abs/1810.02032**.

Ji, Ziwei, & Telgarsky, Matus. 2020. Directional convergence and alignment in deep learning. *Pages 17176–17186 of:* Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., & Lin, H. (eds), *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc.

Jin, Chi, Ge, Rong, Netrapalli, Praneeth, Kakade, Sham M., & Jordan, Michael I. 2017. How to Escape Saddle Points Efficiently. *Page 1724–1732 of: Proceedings of the 34th International Conference on Machine Learning - Volume 70*. ICML'17. JMLR.org.

Kawaguchi, Kenji. 2016. Deep Learning without Poor Local Minima. *In:* Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., & Garnett, R. (eds), *Advances in Neural Information Processing Systems*, vol. 29. Curran Associates, Inc.

Lee, Jaehoon, Bahri, Yasaman, Novak, Roman, Schoenholz, Samuel S, Pennington, Jeffrey, & Sohl-Dickstein, Jascha. 2017. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*.

Lee, Jaehoon, Xiao, Lechao, Schoenholz, Samuel, Bahri, Yasaman, Novak, Roman, Sohl-Dickstein, Jascha, & Pennington, Jeffrey. 2019a. Wide neural networks of any depth evolve as linear models under gradient descent. *Pages 8572–8583 of: Advances in neural information processing systems*.

Lee, Jason D., Simchowitz, Max, Jordan, Michael I., & Recht, Benjamin. 2016. Gradient Descent Only Converges to Minimizers. *Pages 1246–1257 of:* Feldman, Vitaly, Rakhlin, Alexander, & Shamir, Ohad (eds), *29th Annual Conference on Learning Theory*. Proceedings of Machine Learning Research, vol. 49. Columbia University, New York, New York, USA: PMLR.

Lee, Jason D, Panageas, Ioannis, Piliouras, Georgios, Simchowitz, Max, Jordan, Michael I, & Recht, Benjamin. 2019b. First-order methods almost always avoid strict saddle points. *Mathematical programming*, **176**(1), 311–337.

Li, Zhiyuan, Luo, Yuping, & Lyu, Kaifeng. 2020. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. *arXiv preprint arXiv:2012.09839*.

Liu, Chaoyue, Zhu, Libin, & Belkin, Mikhail. 2020. Toward a theory of optimization for over-parameterized systems of non-linear equations: the lessons of deep learning. *arXiv preprint arXiv:2003.00307*.

Lyu, Kaifeng, & Li, Jian. 2020. Gradient Descent Maximizes the Margin of Homogeneous Neural Networks. *In: International Conference on Learning Representations*.

Mei, Song, Montanari, Andrea, & Nguyen, Phan-Minh. 2018. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, **115**(33), E7665–E7671.

Mei, Song, Misiakiewicz, Theodor, & Montanari, Andrea. 2019. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. *arXiv preprint arXiv:1902.06015*.

Moroshko, Edward, Woodworth, Blake E, Gunasekar, Suriya, Lee, Jason D, Srebro, Nati, & Soudry, Daniel. 2020. Implicit Bias in Deep Linear Classification: Initialization Scale vs Training Accuracy. *Pages 22182–22193 of:* Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., & Lin, H. (eds), *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc.

Nouiehed, Maher, & Razaviyayn, Meisam. 2021. Learning deep models: Critical points and local openness. *INFORMS Journal on Optimization*.

Rotskoff, Grant, & Vanden-Eijnden, Eric. 2018. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. *Pages 7146–7155 of: Advances in Neural Information Processing Systems 31*. Curran Associates, Inc.

Saxe, Andrew M., McClelland, James L., & Ganguli, Surya. 2014. *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*.

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

Simsek, Berfin, Ged, François, Jacot, Arthur, Spadaro, Francesco, Hongler, Clement, Gerstner, Wulfram, & Brea, Johanni. 2021. Geometry of the Loss Landscape in Overparameterized Neural Networks: Symmetries and Invariances. *Pages 9722–9732 of:* Meila, Marina, & Zhang, Tong (eds), *Proceedings of the 38th International Conference on Machine Learning*. Proceedings of Machine Learning Research, vol. 139. PMLR.

Soudry, Daniel, Hoffer, Elad, Nacson, Mor Shpigel, Gunasekar, Suriya, & Srebro, Nathan. 2018. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, **19**(1), 2822–2878.

Vershynin, Roman. 2010. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*.

Woodworth, Blake, Gunasekar, Suriya, Savarese, Pedro, Moroshko, Edward, Golan, Itay, Lee, Jason, Soudry, Daniel, & Srebro, Nathan. 2020. *Kernel and Rich Regimes in Overparametrized Models*.

Yang, Greg. 2019. Scaling Limits of Wide Neural Networks with Weight Sharing: Gaussian Process Behavior, Gradient Independence, and Neural Tangent Kernel Derivation. *arXiv e-prints*, Feb, arXiv:1902.04760.

Yang, Greg, & Hu, Edward J. 2020. *Feature Learning in Infinite-Width Neural Networks*.

Yun, Chulhee, Krishnan, Shankar, & Mobahi, Hossein. 2021. A unifying view on implicit bias in training linear neural networks. *In: International Conference on Learning Representations*.

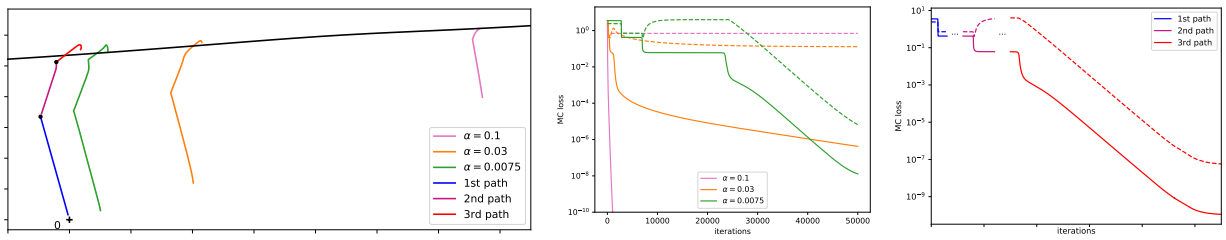

Figure 4. Matrix Completion in linear/lazy vs. saddle-to-saddle regimes. 3 DLNs ($L = 4, w = 100$) trained on a MC loss fitting a 10×10 matrix of rank 3 with initialization $\alpha\theta_0$ for a fixed random θ_0 and three values of α. **Left:** Train (solid) and test (dashed) MC cost for the three networks, for large α the network is in the linear/lazy regime and does not learn the low-rank structure. For smaller α plateaus appear and the network generalizes. **Middle:** Visualization of the gradient paths in parameter space. The black line represents the manifold of solutions to which all example paths converge. As $\alpha \to 0$ the training trajectory converges to a sequence of 3 paths (in blue, purple and red) starting from the origin (+) and passing through 2 saddles ($\cdot$) before converging. **Right:** The train (solid) and test (dashed) loss of the three paths plotted sequentially, in the saddle-to-saddle limit; $\cdots$ represent an infinite amount of steps separating these paths.

We organize the Appendix as follows:

- In Section A, we present the details for the numerical results presented in the main text together with some discussions.
- In Section B, we present the proofs for the result on the proximity of critical points, i.e. Theorem 1.
- In Section C, we present the proofs for the Saddle-to-Saddle regime, in particular Theorems 2 and 5.
- In Section D, we state and prove a few technical results.

A. Further Experimental Details

Experimental details of Fig. 2: A teacher network matrix of size 5×5 is generated as $10\mathrm{diag}([1, 2, 3, 4, 5])$. The input data is i.i.d. 5-dim. standard Gaussian samples, and the number of training samples is 100. The labels are generated by the teacher, no noise is added. Training is performed with gradient descent for 50000 epochs and a learning rate of $1e - 4$ is used.

Experimental details of Fig. 3: A random matrix A^* of size 30×30 is generated by multiplying two i.i.d. matrices of size 30×1 with i.i.d. standard Gaussian entries. 0.2 of the entries of this matrix was accecible in training, and the training objecive is the squared difference between the (observed) entries of the linear network matrix and those of the matrix A^*. The training is performed for 20000 gradient descent iterations with a learning rate of $\eta_0 = 0.05$ if $\gamma > 1$, and $\eta = \eta_0 w^{(L-1)(\gamma-1)}$ for $\gamma \leq 1$. The tolerance for computing the rank is set to 0.1.

Experimental details of Fig. 6 in the Appendix: We created a rank 3 teacher weight matrix $W_T = W_0 W_0^T$ of size 10×10 where W_0 is a 10×3 matrix with all entries independent Gaussian with zero mean and where all entries in i-th column has variance i for all $i \in \{1, 2, 3\}$. We corrupted the teacher weight matrix by an addition of a 10×10 matrix where each entry is i.i.d. centered Gaussian with standard deviation 0.2. Input points are isotropic Gaussians. The training outputs are generated by the noisy teacher, and the test outputs are generated by the noiseless teacher. We generated 100 training and 1000 test data points. Different runs of the same experiment yielded effectively the same figure. The learning rate is 0.001 both for the shallow and the deep case. Tolerance for the rank is set to 10^{-4} (i.e. eigenvalues smaller than 10^{-4} are set to 0 for the rank calculation).

B. Regimes of Training

In this section we describe the regimes of training depending on the scaling γ of the variance at initialization $\sigma^2 = w^{-\gamma}$.

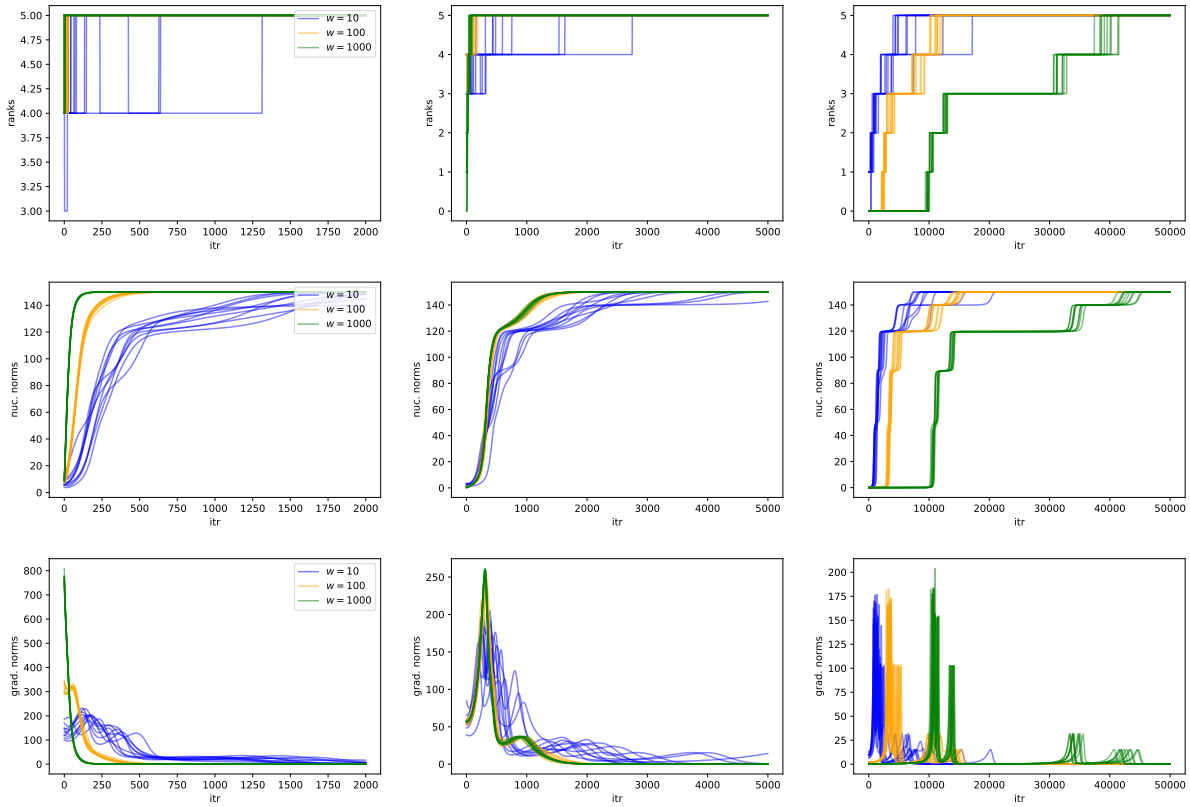

(a) $\gamma = 0.75$ (NTK) **(b)** $\gamma = 1$ (MF) **(c)** $\gamma = 1.5$ (S-S)

*Figure 5. Training in **(a)** the NTK regime, **(b)** mean-field, **(c)** saddle-to-saddle regimes in deep linear networks for three widths $w = 10, 100, 1000$, $L = 4$, and 10 seeds; extension of Fig. 2 in the main.* **Top:** The evolution of the rank of the network matrices during training. Tolerance of the matrix is set at $1e-1$. **Middle:** The evolution of the nuclear norm during training, we can see that the smooth jumps are aligned with the rank transitions. **Bottom:** The evolution of the gradient norm of the parameters. Decrease of the gradient norm down to zero indicates approaching to a saddle, and the following increase indicates escaping it.

B.1. Equivalence of Parametrization/Initializations

B.1.1. NTK Parametrization

Let us show that the NTK parametrization corresponds to a scaling of $\gamma = 1 - \frac{1}{L}$.

The NTK parametrization (Jacot *et al.*, 2018) for linear networks is

$$A_\theta^{NTK} = \frac{W_L}{\sqrt{n_{L-1}}} \cdots \frac{W_1}{\sqrt{n_0}} = \frac{1}{\sqrt{n_0 \cdots n_{L-1}}} W_L \cdots W_1$$

with all parameters initialized with a variance of 1. One can show that gradient flow $\theta^{NTK}(t)$ with the NTK parametrization, initialized at some parameters θ_0^{NTK} is equivalent (up to a rescaling of the learning rate) to gradient flow $\theta(t)$ with the classical parametrization with an initialization of $\theta_0 = (n_0 \cdots n_{L-1})^{-\frac{1}{2L}} \theta_0^{NTK}$:

Proposition 7. *Let $\theta^{NTK}(t)$ be gradient flow on the loss $\mathcal{L}^{NTK}(\theta) = C(A_\theta^{NTK})$ initialized at some parameters θ_0^{NTK} and $\theta(t)$ be gradient flow on the cost $\mathcal{L}(\theta) = C(A_\theta)$ initialized at $\theta_0 = (n_0 \cdots n_{L-1})^{-\frac{1}{2L}} \theta_0^{NTK}$. We have*

$$A_{\theta(t)} = A^{NTK}_{\theta^{NTK}(\sqrt{n_0 \cdots n_{L-1}} t)}.$$

Proof. We will show that $\theta(t) = (n_0 \cdots n_{L-1})^{-\frac{1}{2L}} \theta^{NTK}(\sqrt{n_0 \cdots n_{L-1}} t)$ which implies that $A_{\theta(t)} = A^{NTK}_{\theta^{NTK}(t)}$. This

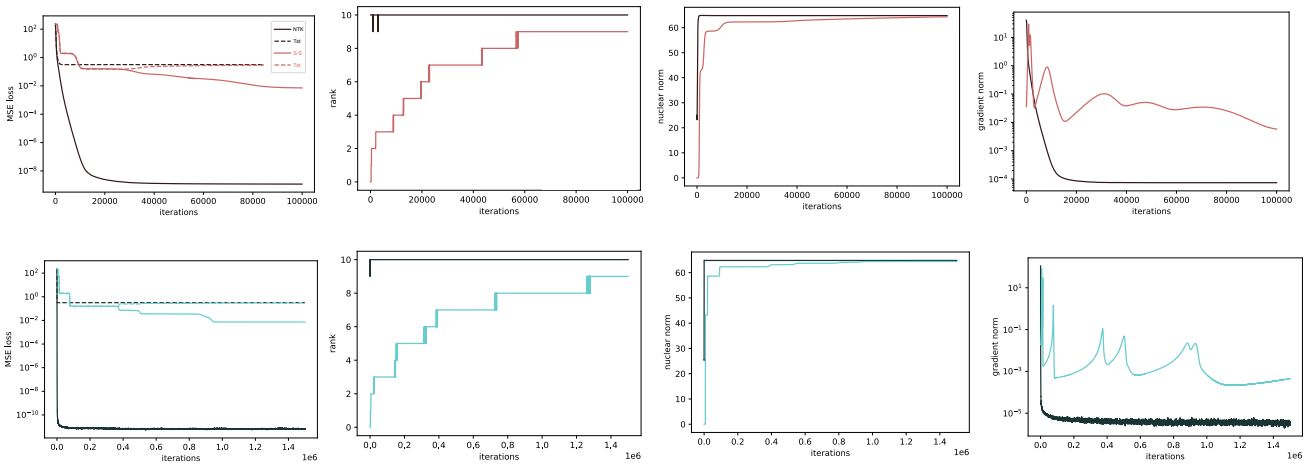

Figure 6. Training in the NTK vs. saddle-to-saddle regimes in shallow (top) and deep (bottom) networks when learning a low rank matrix corrupted with noise. Black lines (the NTK regime): the parameters are initialized with the standard deviation $\tilde{\sigma} = w^{-L-1/2L}$. The rank of the network matrix increases incrementally as the gradient trajectory follows the paths between the saddles. **Top/Shallow case:** $L = 2$ and $w = 50$; in the saddle-to-saddle regime (shown in red), the initialization scale is $\tilde{\sigma} = w^{-2}$. Bigger initialization scales result in shorter plateaus in the loss curve if the same learning rate is used. **Bottom/Deep case:** $L = 4$ and $w = 100$; in the saddle-to-saddle regime (shown in blue), the initialization scale is $\tilde{\sigma} = w^{-1}$. We observe that the transitions from saddles to saddles are sharper. We observe that the gradient norm of the parameters is highly non-monotonic; a decrease down to 0 indicates approaching to a saddle, and a following increase indicates escaping it. We note that the peaks of the gradient norm are sharper in the deep case, suggesting a different rate of escape. In the NTK regime, the gradient norm decreases down to 0 monotonically. In the deep case the GD training is implemented for 1500000 iterations whereas in the shallow case it is only 100000 iterations.The input data is standard Gaussian, the outputs are generated by a rank 3 teacher of size 10×10 corrupted with noise, and the loss is MSE.

is obviously true at $t = 0$. Now assuming it is true at a time t, we show that the time derivatives of $\theta(t)$ and $(n_0 \cdots n_{L-1})^{-\frac{1}{2L}} \theta^{NTK}(\sqrt{n_0 \cdots n_{L-1}} t)$ match:

$$\partial_t \theta^{NTK}(\sqrt{n_0 \cdots n_{L-1}} t) = \frac{\sqrt{n_0 \cdots n_{L-1}}}{\sqrt{n_0 \cdots n_{L-1}}} \partial_t \theta(t) = \partial_t \theta(t).$$

□

This implies that the NTK parametrization with $\mathcal{N}(0, 1)$ initialization is equivalent to the classical parametrization with $\mathcal{N}(0, (n_0 \cdots n_{L-1})^{-\frac{1}{L}})$ initialization, which for rectangular networks corresponds to a $\mathcal{N}(0, n_0^{-\frac{1}{L}} w^{-\frac{L-1}{L}})$ initialization with scaling $\gamma = \frac{L-1}{L} = 1 - \frac{1}{L}$.

B.1.2. Maximal Update Parametrization

The Maximal Update parametrization (or μ-parametrization) (Yang & Hu, 2020) is equivalent to $\gamma = 1$. The μ-parametrization for linear rectangular networks is the same the classical one, since

$$A_\theta^\mu = \frac{W_L}{\sqrt{w}} W_{L-1} \cdots W_2 \left(\sqrt{w} W_1\right) = W_L \cdots W_1$$

and the parameters are initialized with variance w^{-1}, i.e. $\gamma = 1$.

B.2. Distance to Different Critical Points

Let d_{m} and d_{s} be the Euclidean distances between the initialization θ and, respectively, the set of global minima and the set of all saddles. For random variables $f(w), g(w)$ which depend on w, we write $f \asymp g$ if both $f(w)/g(w)$ and $g(w)/f(w)$ are stochastically bounded as $w \to \infty$. The following theorem studies how d_{m} and d_{s} scale as $w \to \infty$:

Theorem 8 (Theorem 1 in the main)**.** *Suppose that the set of matrices that minimize C is non-empty, has Lebesgue measure zero, and does not contain the zero matrix. Let θ be i.i.d. centered Gaussian r.v. of variance $\sigma^2 = w^{-\gamma}$ where $1 - \frac{1}{L} \leq \gamma < \infty$. Then:*

1. *if $1 - \frac{1}{L} \leq \gamma < 1$, we have $d_{\mathrm{m}} \asymp w^{-\frac{(1-\gamma)(L-1)}{2}}$ and $d_{\mathrm{s}} \asymp w^{\frac{1-\gamma}{2}}$,*

2. *if $\gamma = 1$, we have $d_{\mathrm{m}}, d_{\mathrm{s}} \asymp 1$,*

3. *if $\gamma > 1$ we have $d_{\mathrm{m}} \asymp 1$ and $d_{\mathrm{s}} \asymp w^{-\frac{\gamma-1}{2}}$.*

To prove this result, we require a few Lemmas:

Lemma 9. *Let θ be the vector of parameters of a DLN with i.i.d. $\mathcal{N}(0, w^{-\gamma})$ Gaussian entries, and let $\mathcal{A}_{\min} = \{A \in \mathbb{R}^{n_L \times n_0} : C(A) = 0\}$ be the set of global minimizers of C. Under the same assumptions on the cost C as Proposition 8, we have $d(A_\theta, \mathcal{A}_{\min}) \asymp 1$ as $w \to \infty$.*

Proof. If $\gamma > 1 - \frac{1}{L}$ then A_θ converges in distribution to the zero matrix as $w \to \infty$, the distance $d(A_\theta, \mathcal{A}_{\min})$ therefore converges to the distance finite value $d(0, \mathcal{A}_{\min}) \neq 0$.

If $\gamma = 1 - \frac{1}{L}$, then A_θ converges in distribution to random Gaussian matrix with iid $\mathcal{N}(0, 1)$ entries (this can seen as a consequence of the more general results for non-linear networks (Lee *et al.*, 2017; de G. Matthews *et al.*, 2018)). As a result the distribution of $d(A_\theta, \mathcal{A}_{\min})$ converges to the distribution of $d(B, \mathcal{A}_{\min})$ for a matrix B with iid Gaussian $\mathcal{N}(0, 1)$ entries. Since $\mathbb{P}\left[d(B, \mathcal{A}_{\min}) = 0\right] = 0$ and $\mathbb{P}\left[d(B, \mathcal{A}_{\min}) > b\right] \to 0$ as $b \to \infty$ we have that $d(A_\theta, \mathcal{A}_{\min}) \asymp 1$ as needed. □

Lemma 10. *Let θ be the vector of parameters of a DLN with iid $\mathcal{N}(0, w^{-\gamma})$ Gaussian entries. For all ϵ, there is a constant $C_{\epsilon,L}$ that does not depend on w s.t. with prob. $1 - \epsilon$, we have for all $\theta' \in \mathbb{R}^P$ that*

$$\|A_{\theta'} - A_\theta\|_F^2 \leq C_{\epsilon,L} \sum_{k=1}^{L} \|\theta - \theta'\|^{2k} w^{(1-\gamma)(L-k)}.$$

Proof. By Corollary 5.35 in (Vershynin, 2010), reformulated as Theorem 11 below, we know that for all ϵ, there is a constant $c_\epsilon$that does not depend on w s.t. with prob. $1 - \epsilon$, we have for all ℓ

$$\|W_\ell\|_{op}^2 \leq c_\epsilon w^{1-\gamma}.$$

We now write $d\theta = \theta' - \theta$ (and the corresponding matrices $dW_\ell = W'_\ell - W_\ell$) so that we may write the difference $A_{\theta + d\theta} - A_\theta$ as the following sum

$$\sum_{\substack{a_1, \ldots, a_L \in \{0,1\} \\ \exists \ell, a_\ell \neq 0}} \left(\begin{cases} W_L & \text{if } a_L = 0 \\ dW_L & \text{if } a_L = 1 \end{cases} \right) \cdots \left(\begin{cases} W_1 & \text{if } a_1 = 0 \\ dW_1 & \text{if } a_1 = 1 \end{cases} \right)$$

where the indicator a_ℓ determines whether we take W_ℓ or dW_ℓ in the product. We can therefore bound

$$\begin{aligned}
\|A_{\theta+d\theta} - A_\theta\|_F^2 &\leq \left(\sum_{\substack{a_1,\dots,a_L \in \{0,1\} \\ \exists \ell, a_\ell \neq 0}} \left\| \left(\begin{cases} W_L & \text{if } a_L = 0 \\ dW_L & \text{if } a_L = 1 \end{cases} \right) \cdots \left(\begin{cases} W_1 & \text{if } a_1 = 0 \\ dW_1 & \text{if } a_1 = 1 \end{cases} \right) \right\|_F \right)^2 \\
&\leq (2^L - 1) \sum_{\substack{a_1,\dots,a_L \in \{0,1\} \\ \exists \ell, a_\ell \neq 0}} \left\| \left(\begin{cases} W_L & \text{if } a_L = 0 \\ dW_L & \text{if } a_L = 1 \end{cases} \right) \cdots \left(\begin{cases} W_1 & \text{if } a_1 = 0 \\ dW_1 & \text{if } a_1 = 1 \end{cases} \right) \right\|_F^2 \\
&\leq (2^L - 1) \sum_{\substack{a_1,\dots,a_L \in \{0,1\} \\ \exists \ell, a_\ell \neq 0}} \left(\begin{cases} \|W_L\|_{op}^2 & \text{if } a_L = 0 \\ \|dW_L\|_F^2 & \text{if } a_L = 1 \end{cases} \right) \cdots \left(\begin{cases} \|W_1\|_{op}^2 & \text{if } a_1 = 0 \\ \|dW_1\|_F^2 & \text{if } a_1 = 1 \end{cases} \right)
\end{aligned}$$

We now bound $\|W_L\|_{op}^2$ by $c_\epsilon w^{1-\gamma}$ and $\|dW_L\|_F^2$ by $\|d\theta\|^2$ so that we obtain the bound

$$\|A_{\theta+d\theta} - A_\theta\|_F^2 \leq (2^L - 1) \sum_{k=1}^{L} \begin{pmatrix} L \\ k \end{pmatrix} \|d\theta\|^{2k} c_\epsilon^{L-k} w^{(1-\gamma)(L-k)} \leq C_{\epsilon,L} \sum_{k=1}^{L} \|d\theta\|^{2k} w^{(1-\gamma)(L-k)}$$

for $C_{\epsilon,L} = (2^L - 1) \max_{k=1,\dots,L} \begin{pmatrix} L \\ k \end{pmatrix} c_\epsilon^{L-k}$. □

Let us now prove Theorem 8:

Proof. **(1) Distance to minimum:** Let us first give an lower bound on the distance from initialization to a global minimum. Let θ be the intialization and $\theta + d\theta$ be the closest minimum. By Lemma 10, we obtain

$$\|A_{\theta+d\theta} - A_\theta\|_F^2 \leq C'_L \sum_{k=1}^{L} \|d\theta\|^{2k} w^{(1-\gamma)(L-k)}.$$

If $\gamma > 1$, the term with $k = L$ dominates, in which case $\|A_{\theta+d\theta} - A_\theta\|_F^2 \leq \|d\theta\|^{2L}$ which implies that $\|d\theta\| \geq \|A_{\theta+d\theta} - A_\theta\|_F^{\frac{1}{L}} \geq d(A_\theta, \mathcal{A}_{\min})^{\frac{1}{L}} \asymp 1$ by Lemma 9.

If $\gamma < 1$, the term $k = 1$ dominates, which implies $\|A_{\theta+d\theta} - A_\theta\|_F^2 \leq \|d\theta\|^2 w^{(1-\gamma)(L-1)}$ which implies that $\|d\theta\| \geq \|A_{\theta+d\theta} - A_\theta\|_F w^{-\frac{(1-\gamma)(L-1)}{2}} = O(w^{-\frac{(1-\gamma)(L-1)}{2}})$, which decreases with width.

Let us now show upper bounds on $\|d\theta\|$. When $\gamma > 1$, we will construct a closeby minimum. Let us first define the parameters $\bar{\theta} = (\bar{W}_1, \dots, \bar{W}_L)$ where $\bar{W}_1 = 0$ and $\bar{W}_L = 0$ and $\bar{W}_{\ell,ij} = \begin{cases} W_{\ell,ij} & \text{if } i,j > \min\{n_0, n_L\} \\ 0 & \text{otherwise} \end{cases}$. Since we have set only $O(w)$ parameters to zero, we have $\left\|\theta - \bar{\theta}\right\|^2 = O(\sigma^2 w) = O(w^{1-\gamma})$. Now let the matrix A be a global minimum of the cost C with SVD $A = USV^T$ (with inner dimension equal to the rank k of A), we then set $\theta^* = \bar{\theta} + I^{(k\to w)}(S^{\frac{1}{L}}V^T, S^{\frac{1}{L}}, \dots, S^{\frac{1}{L}}, US^{\frac{1}{L}})$. The parameters θ^* are a global minimum since $A_{\theta^*} = A$ and $\|\theta^* - \theta\| \leq \left\|\theta^* - \bar{\theta}\right\| + \left\|\bar{\theta} - \theta\right\| = O(1) + O(w^{\frac{1-\gamma}{2}}) = O(1)$.

When $\gamma < 1$, with prob. $1 - \epsilon$, we have $s_{min}(W_{L-1} \cdots W_1) > \frac{1}{2}\sigma^{(L-1)} w^{\frac{L-1}{2}} = w^{\frac{(1-\gamma)(L-1)}{2}}$, we can reach a global minimum by only changing W_L, we need $dW_L W_{L-1} \cdots W_1 = A^* - A_\theta$ hence we take $dW_L = (A^* - A_\theta)(W_L \cdots W_1)^+$ with norm $\|d\theta\| = \|dW_L\|_F \leq \frac{\|A^* - A_\theta\|}{s_{min}(W_{L-1} \cdots W_1)} = O(w^{-\frac{(1-\gamma)(L-1)}{2}})$.

(2) Distance to saddles: Given parameters $\theta = (W_1, \dots, W_L)$, we can obtain a saddle θ^* by setting all entries of W_1 and W_L to zero. We have

$$\mathbb{E}\left[\|\theta - \theta^*\|^2\right] = \mathbb{E}\left[\|W_1\|_F^2\right] + \mathbb{E}\left[\|W_L\|_F^2\right] = \sigma^2(n_0 + n_L)w = O(w^{1-\gamma}).$$

This gives an upper bound of order $w^{1-\gamma}$ on the distance between θ and the set of saddles θ^* with $A_{\theta^*} = 0$.

Now let $\theta^* = \theta + d\theta$ be the saddle closest to θ, we know that

$$0 = \partial_{W_L}\mathcal{L}(\theta^*) = \nabla C(A_{\theta^*})\,(W_1^*)^T \cdots \left(W_{L-1}^*\right)^T.$$

Since A_{θ^*} is not a global minimum, $\nabla C(A_{\theta^*}) \neq 0$, for the above to be zero, we therefore need $(W_1^*)^T \cdots \left(W_{L-1}^*\right)^T$ to not have full column rank, i.e. $\operatorname{Rank}(W_1^*)^T \cdots \left(W_{L-1}^*\right)^T = n_0$.

We will show that at initialization $(W_1)^T \cdots (W_{L-1})^T$ has rank n_0 and its smallest non-zero singular value s_{min} is of order $w^{\frac{(1-\gamma)(L-1)}{2}}$. We will use the fact that $\left\|(W_1)^T \cdots (W_{L-1})^T - (W_1^*)^T \cdots \left(W_{L-1}^*\right)^T\right\|_F \geq s_{min}$ to lower bound the distance $\|\theta - \theta^*\|$ using Lemma 10.

The singular values of $W_1^T \cdots W_{L-1}^T$ are the squared root of the eigenvalues of the $n_0 \times n_0$ matrix $W_1^T \cdots W_{L-1}^T W_{L-1} \cdots W_1$. One can show that as $w \to \infty$ this matrix concentrates in its expectation

$$\mathbb{E}\left[W_1^T \cdots W_{L-1}^T W_{L-1} \cdots W_1\right] = \sigma^{2(L-1)} w^{L-1} = w^{(1-\gamma)(L-1)}.$$

which implies that $s_{\min}$ concentrates in $w^{\frac{(1-\gamma)(L-1)}{2}}$ and therefore $s_{\min} \asymp w^{\frac{(1-\gamma)(L-1)}{2}}$.

Now by Lemma 10 (applied to the depth $L-1$ this time), we have with prob. $1-\epsilon$

$$\begin{aligned} s_{\min}^2 &\leq \left\|(W_1)^T \cdots (W_{L-1})^T - (W_1^*)^T \cdots \left(W_{L-1}^*\right)^T\right\|_F^2 \\ &\leq C_{\epsilon,L-1} \sum_{k=1}^{L-1} \|\theta - \theta'\|^{2k} w^{(1-\gamma)(L-1-k)} \end{aligned}$$

and $\|\theta - \theta'\|$ needs to be at least of order $w^{\frac{(1-\gamma)}{2}}$ for any of the terms in the sum to be at least of order $w^{(1-\gamma)(L-1)}$ (actually all these become of the right order at the same time). □

B.3. Spectrum bounds

An important tool in our analysis is the following Theorem (which is a reformulation of Corollary 5.35 in (Vershynin, 2010))

Theorem 11. *Let A be a $m \times n$ matrix with i.i.d. $\mathcal{N}(0, \sigma^2)$ entries. For all $t \geq 0$, with probability at least $1 - 2e^{-\frac{t^2}{2}}$, it holds that*

$$\sigma(-\sqrt{m} - \sqrt{n} - t) \leq s_{min}(A) \leq s_{max}(A) \leq \sigma\left(\sqrt{m} + \sqrt{n} + t\right).$$

Corollary 12. *If the parameters θ are independent centered Gaussian with variance σ^2, for all $t \geq 0$, with probability at least $1 - 2Le^{-\frac{t^2}{2}}$, it holds that*

$$\|A_\theta\|_{op} \leq (1+t)^L \sigma^L \left(\sqrt{n_0} + \sqrt{w}\right)(4w)^{\frac{L-2}{2}}\left(\sqrt{w} + \sqrt{n_L}\right).$$

Proof. By Theorem 11, with probability greater than $1 - 2Le^{-\frac{t^2}{2}}$, for all $\ell = 1, \dots, L$, $\|W_\ell\|_{op} \leq \sigma\left(\sqrt{n_{\ell-1}} + \sqrt{n_\ell} + t\right)$, where $n_\ell = w$ for $\ell \in \{1, \cdots, L-1\}$. Hence

$$\|A_\theta\|_{op} \leq \|W_L\|_{op} \cdots \|W_1\|_{op} \leq \sigma^L \prod_{\ell=1}^{L}\left(\sqrt{n_{\ell-1}} + \sqrt{n_\ell} + t\right) \leq (1+t)^L \sigma^L \prod_{\ell=1}^{L}\left(\sqrt{n_{\ell-1}} + \sqrt{n_\ell}\right).$$

□

C. Proofs for the Saddle-to-Saddle regime

In this section, we prove Theorem 4 of the main. Given a saddle $\vartheta^* = RI^{(k\to w)}(\vartheta)$ where ϑ is a local minimum in a width k network, we want to describe the dynamics of gradient descent $\theta_\alpha(t) = \gamma(t, \vartheta^* + \alpha\theta_0)$, initialized close to ϑ^*. We shall consider $\vartheta^* = 0$ for convenience, though the same arguments could be applied for $\vartheta^* \neq 0$. We will start by studying the case of homogeneous costs, which will allow us to describe costs that locally look homogeneous around 0. Later on, after having defined the notion of *escape paths,* we will show that as $\alpha \to 0$, the path $(\theta_\alpha(t))_{t\in\mathbb{R}_+}$ converges to an escape path with specific direction and speed. We will then show that the escape paths which escape at this speed are unique in some aspects.

C.1. Homogeneous Costs

As in the main text, we use θ to denote an element in the parameter space $\mathbb{R}^P$. Let $k \geq 2$ be an integer. We say that a cost H is k-homogeneous if $H(\alpha\theta) = \alpha^k H(\theta)$ for all $\theta \in \mathbb{R}^P$ and all scalar $\alpha > 0$. Later in this paper, we will be particularly interested in the case where $H(\theta) = \mathrm{Tr}\,[GA_\theta]$ for a linear network A_θ of depth L and some $n_L \times n_0$ matrix G. Thus defined, H is a L-homogeneous polynomial.

Throughout, when studying a k-homogeneous cost H, we will always assume that it is twice differentiable.

A useful property of gradient descent on a homogeneous cost is that:

Lemma 13. *Gradient flow on a twice-differentiable k-homogeneous cost H satisfies*

$$\gamma_H(t, \lambda\theta_0) = \lambda\gamma_H(\lambda^{k-2}t, \theta_0)$$

for all $\theta_0 \in \mathbb{R}^P$,all $\lambda > 0$ and all $t \geq 0$.

Proof. We simply need show that for all $\theta_0 \in \mathbb{R}^P, \lambda > 0, t \geq 0$, we have $\frac{1}{\lambda}\gamma_H(\lambda^{2-k}t, \lambda\theta_0) = \gamma_H(t, \theta_0)$, i.e. that the path $t \mapsto \frac{1}{\lambda}\gamma_H(\lambda^{2-k}t, \lambda\theta_0)$ is the solution of gradient descent starting at θ_0. Clearly, the path starts at θ_0 and satisfies

$$\partial_t \frac{1}{\lambda}\gamma_H(\lambda^{2-k}t, \lambda\theta_0) = -\lambda^{1-k}\nabla H\left(\gamma_H(\lambda^{2-k}t, \lambda\theta_0)\right) = -\nabla H\left(\frac{1}{\lambda}\gamma_H(\lambda^{2-k}t, \lambda\theta_0)\right)$$

since, using the fact that H is k-homogeneous, for all scalar $\alpha > 0$, and any $\theta \in \mathbb{R}^P$, $\alpha\nabla H(\alpha\theta) = \alpha^k\nabla H(\theta)$. One concludes using Picard-Lindelöf Theorem, using that ∇H is locally Lipschitz around 0 since H is twice differentiable. □

An **Escape Direction** at 0 of H is a vector on the sphere $\rho \in \mathbb{S}^{P-1}$ such that $\nabla H(\rho) = -s\rho$ for some $s \in \mathbb{R}_+$ which we call the *escape speed* associated with ρ. A path $(\theta(t))_{t<0}$ indexed by negative times and following gradient flow on H such that $\theta(t)$ is on one escape direction for some $t < 0$ will remain along this direction (these paths are equal for $t < 0$ to $\theta(t) = \rho e^{st}$ when $k = 2$ and $\theta(t) = \rho\left(-(k-2)st\right)^{-\frac{1}{k-2}}$ when $k > 2$). Note that this entails that $\theta(t) \to 0$ as $t \to -\infty$. When $H(\theta) = \theta^T A\theta$ for some symmetric matrix A, the escape directions are simply the eigenvectors of the Hessian A and the escape speeds are twice the eigenvalues of A.

An **Optimal Escape Direction** $\rho^* \in \mathbb{S}^{P-1}$ is an escape direction with the largest speed $s^* > 0$. It is a minimizer of H restricted to $\mathbb{S}^{P-1}$:

$$\rho^* \in \arg\min_{\rho\in\mathbb{S}^{P-1}} H(\rho). \tag{1}$$

Indeed, critical points of $H(\rho)$ restricted to the sphere are the escape directions, and by Euler's condition (i.e. $\nabla H(\rho)^T\rho = kH(\rho)$ if H is k-homogeneous), if ρ is an escape direction with speed s, then $H(\rho) = -\frac{s}{k}$: optimal escape directions are thus global minimizers of H restricted on the unit sphere.

Under some conditions on the Hessian along the escape directions, one can guarantee that gradient descent will escape along an optimal escape path:

Proposition 14. *Assume that the optimal escape speed s^* is positive and that for all escape directions ρ which are not optimal, there is a vector $v \perp \rho$ such that $v^T\mathcal{H}H(\rho)v < -s^*v^Tv$. Let Ω be the set of θ_0 such that the direction $\frac{\gamma(t,\theta_0)}{\|\gamma(t,\theta_0)\|}$ of the gradient descent flow converges towards an optimal escape direction as $t \to T$, where T is the explosion time of the path (which can be infinite). The set $\mathbb{S}^{P-1} \setminus \Omega$ has spherical measure zero.*

Proof. Let $\Omega' \subset \mathbb{R}^P$ be the set of points $\theta \neq 0$ such that gradient flow on the 0-homogeneous cost $\bar{H}(\theta) = H\left(\frac{\theta}{\|\theta\|}\right)$ converges to a global minimum. Our proof is divided in two steps: (1) we show that $\Omega' \subset \Omega$, (2) we show that $\Omega' \cap \mathbb{S}^{P-1}$ has spherical measure 1.

(1) Note that both sets Ω and Ω' are cones: for any $\alpha > 0$, $\Omega = \alpha\Omega$ and $\Omega' = \alpha\Omega'$. Therefore, we only need to show that $\Omega' \cap \mathbb{S}^{P-1} \subset \Omega \cap \mathbb{S}^{P-1}$. Besides, note that since $\bar{H}$ is 0-homogenous, $\nabla\bar{H}(\theta)^T\theta = 0$ for all $\theta \in \mathbb{R}^P$ and thus, the norm is an invariant of the descent gradient flow for $\bar{H}$: for any $\theta_0 \in \mathbb{R}^P$, $t \mapsto \|\gamma_{\bar{H}}(t, \theta_0)\|$ is constant.

In particular, if $\theta_0 \in \Omega' \cap \mathbb{S}^{P-1}$, then $\bar{\theta}(t) = \gamma_{\bar{H}}(t, \theta_0)$ converges to an optimal escape direction ρ^* as $t \to \infty$, by Equation (1). The gradient flow path $\theta(t)$ can be obtained from the gradient flow path $\bar{\theta}(s)$ directly. First we define the function $\alpha(s) = e^{-k\int_0^s H(\bar{\theta}(r))dr}$ such that

$$\begin{aligned}\partial_s\left[\bar{\theta}(s)\alpha(s)\right] &= -(I - \bar{\theta}(s)\bar{\theta}(s)^T)\nabla H(\bar{\theta}(s))\alpha(s) - k\bar{\theta}(s)H(\bar{\theta}(s))\alpha(s)\\ &= -\nabla H(\bar{\theta}(s))\alpha(s)\end{aligned}$$

where we used the fact that $\theta^T\nabla H(\theta) = kH(\theta)$. Let us now define $\tau(t) = \int_0^t \alpha(s)^{k-2}ds$, we have

$$\bar{\theta}(\tau(t))\alpha(\tau(t)) = -\nabla H(\bar{\theta}(\tau(t)))\alpha(\tau(t))^{k-1} = -\nabla H\left(\bar{\theta}(\tau(t))\alpha(\tau(t))\right)$$

which implies that $\theta(t) = \bar{\theta}(\tau(t))\alpha(\tau(t))$. As $r \to \infty$, we have $H(\bar{\theta}(r)) \to -s^* < 0$, which implies that $\alpha(s) \to \infty$ as $s \to \infty$. This in turn implies that $\tau(t) \to \infty$ as $t \to \infty$. As a result, we obtain that

$$\lim_{t\to\infty}\frac{\theta(t)}{\|\theta(t)\|} = \lim_{t\to\infty}\frac{\bar{\theta}(\tau(t))\alpha(\tau(t))}{\left\|\bar{\theta}(\tau(t))\alpha(\tau(t))\right\|} = \lim_{t\to\infty}\bar{\theta}(\tau(t)) = \rho^*$$

and hence $\theta_0 \in \Omega$ as needed.

(2) We now show that $\Omega' \cap \mathbb{S}^{P-1}$ has spherical measure 1: this is a consequence of the fact that the critical points of $\bar{H}$ are global minima or strict saddle points. By taking the gradient of $\bar{H}$, one sees that the critical points of $\bar{H}$ on the sphere $\mathbb{S}^{P-1}$ are the points $\theta \in \mathbb{S}^{P-1}$ such that

$$\nabla H(\theta) = \theta\theta^T\nabla H(\theta).$$

Since $\theta\theta^T$ is the orthogonal projection on the line $\mathbb{R}.\theta$, the critical points of $\bar{H}$ on $\mathbb{S}^{P-1}$ are the escape directions. As explained before, global minima of $\overline{H}$ are optimal escape directions. The other escape directions are strict saddle points: consider such ρ and let $v \perp \rho$ be such that $v^T\mathcal{H}H(\rho)v < -s^*v^Tv$. Differentiating $\bar{H}$ twice and using that $v \perp \rho$, one can show that

$$v^T\mathcal{H}\bar{H}(\rho)v = v^T\mathcal{H}H(\rho)v - \nabla H(\rho)^T\rho v^Tv.$$

Since ρ is an escape direction, $\nabla H(\rho) = -s\rho$ with $s < s^*$ and $\|\rho\| = 1$: this implies $v^T\mathcal{H}\bar{H}(\rho)v < (s - s^*)\,v^Tv < 0$. In particular, the points such that the gradient descent on $\bar{H}$ converge to a saddle have spherical measure 0 on $\mathbb{S}^{P-1}$. This shows that $\Omega' \cap \mathbb{S}^{P-1}$ has spherical measure 1, and allows us to conclude. □

C.1.1. Deep Linear Networks

For a depth L DLN and the homogeneous cost $H(\theta) = \mathrm{Tr}\left[G^TA_\theta\right]$ with SVD decomposition $G = USV^T$, the escape directions ρ are of the form

$$\frac{1}{\sqrt{L}}(\pm u_iw_{L-1}^T, w_{L-1}w_{L-2}^T, \ldots, w_1v_i^T)$$

with speed $s = \mp\frac{s_i}{L^{\frac{L-2}{2}}}$, where u_i, v_i are the i-th columns of U, V respectively. The optimal speed is $\frac{s_1}{L^{\frac{L-2}{2}}}$, where s_1 is the largest singular value of G.

Furthermore this loss satisfies the property required to ensure convergence along the fastest escape path:

Lemma 15. *For a network of depth L and width $w \geq 1$, for any escape direction of the form $\rho = \frac{1}{\sqrt{L}}(\pm u_iw_{L-1}^T, w_{L-1}w_{L-2}^T, \ldots, w_1v_i^T)$ with speed $\mp\frac{s_i}{L^{\frac{L-2}{2}}} < \frac{s_1}{L^{\frac{L-2}{2}}}$ the vector $v = (-u_1w_{L-1}^T, 0, \ldots, 0, w_1v_1^T)$ satisfies*

$$v^T\mathcal{H}H(\rho)v < \mp\frac{s_i}{L^{\frac{L-2}{2}}}v^Tv.$$

Proof. We have $v^T \mathcal{H}H(\rho)v = -\frac{2s_1}{L^{\frac{L-2}{2}}}$ and $\pm\frac{s_i}{L^{\frac{L-2}{2}}}v^T v = \pm\frac{2s_i}{L^{\frac{L-2}{2}}}$ as needed. □

This guarantees that gradient flow will not escape along a non-optimal direction, but it does not rule out the possibility that it converges to a saddle of the loss $H(\theta) = \mathrm{Tr}\left[G^T A_\theta\right]$. Each non-zero saddle $\theta^* = (W_1, \dots, W_L)$ is technically proportional to an escape direction ρ with escape speed 0, since $\nabla H(\theta^*) = 0$. For shallow networks these saddles are strict (Kawaguchi, 2016) and so they are almost surely avoided, guaranteeing convergence in direction. For depth $L = 3$ we can apply Proposition 14 since we have:

Lemma 16. *Consider the cost $H(\theta) = \mathrm{Tr}\left[GA_\theta\right]$ for a rank $\min\{n_0, n_L\}$ matrix G and a network of depth $L = 3$ and width $w \geq 1$. For any escape direction ρ with speed 0 there is a vector v such that $v^T \mathcal{H}H(\rho)v < 0$.*

Proof. Since $\rho \neq 0$ there must a non-zero W_1,W_2 or W_3. We separate the case $W_2 \neq 0$ from W_1 or W_3 is non-zero.

Case $W_2 \neq 0$: let u_1, v_1 be the largest singular vectors of G and $\tilde{u}_1, \tilde{v}_1$ the largest singular vectors of W_2, then $v = (-\tilde{u}_1 v_1^T, 0, u_1 \tilde{v}_1^T)$ satisfies

$$v^T \mathcal{H}H(\rho)v = -\mathrm{Tr}\left[G u_1 \tilde{v}_1^T W_2 \tilde{u}_1 v_1^T\right] = -s_1 \tilde{s}_1 < 0.$$

Case $W_1 \neq 0$ (the case $W_3 \neq 0$ is similar): Let u_1, v_1 be the largest singular vectors of $W_1 G$ and b be any unitary w-dim vector, then the parameters $v = (0, b v_1^T, u_1 b^T)$ satisfy

$$v^T \mathcal{H}H(\rho)v = -\mathrm{Tr}\left[G u_1 b^T b v_1^T W_1\right] = -s_1 < 0.$$

□

For $L > 3$ we were not able to prove that the saddles can be avoided with prob. 1, we therefore introduce the assumption:

Assumption 17. *Let I be the set of initializations which converge to a saddle of the cost $H(\theta) = \mathrm{Tr}\left[GA_\theta\right]$. We shall work on the event $E = \theta_0 \notin I$.*

It can easily be proven for a Gaussian initialization that $\mathbb{P}(E) \geq 1/2$, i.e. that saddles can be avoided with probability at least $1/2$, since $P\left(H(\theta_0) < 0\right) = \frac{1}{2}$ at initialization (this follows from the fact that $H((W_1, \dots, W_L)) = -H((-W_1, \dots, W_L))$).

Another motivation for this assumption is the fact that if the network is initialized with balanced weights (Arora *et al.*, 2019a;b), i.e. if $W_\ell^T W_\ell = W_{\ell-1} W_{\ell-1}^T$ for $1 < \ell < L$, then necessarily $\theta_0 \notin I$. This is because the balancedness is conserved during training: if gradient flow converges to a saddle, this saddle must be balanced. However the only balanced saddle of H is the origin, which can only be approached along an escape direction ρ with positive speed $s > 0$, which are avoided with prob. 1 by Proposition 14.

C.2. Approximately Homogeneous Costs

In the previous section, we studied the escape paths for homogeneous costs H. We extend these results to more general cost functions, which are only locally homogeneous around a saddle ϑ^*, i.e. we consider costs of the form

$$C(\theta) = H(\theta - \vartheta^*) + e(\theta - \vartheta^*), \tag{2}$$

where H is a k-homogeneous cost H, and where e is infinitely differentiable such that its $m - 1$ first derivatives vanish at 0 for a given $m > k$. We call such costs (k, m)-*approximately homogeneous*. In the setting of a cost $C(A_\theta)$ for a neural network of depth L, the saddle at the origin $\theta = 0$ is $(L, 2L)$-approximately homogeneous, since the only non-vanishing derivatives are the kL-th derivatives for $k = 0, 1, 2, \dots$.

Since we are only interested in the local behaviour around the saddle ϑ^*, we localize the cost: let $h : \mathbb{R}_+ \to \mathbb{R}_+$ be a smooth cut-off function such that $h(x) = 1$ if $0 \leq x \leq 1$, $0 \leq h(x) \leq 1$ if $1 < x < 2$ and $h(x) = 0$ when $x \geq 2$. For $r > 0$, we define the localization C_r of the cost C as

$$C_r(\theta) = H(\theta - \vartheta^*) + e(\theta - \vartheta^*) h\left(\frac{\|\theta - \vartheta^*\|}{r}\right). \tag{3}$$

As usual, we assume for simplicity that $\vartheta^* = 0$. We note for later use that by assumption on e, for all compact set K containing 0, there exists a finite constant $c > 0$ such that

$$||\nabla e(\theta)|| \leq c||\theta||^k, \forall \theta \in K. \tag{4}$$

Lemma 18. *Let $h : \mathbb{R}_+ \to \mathbb{R}$ and $e : \mathbb{R}^P \to \mathbb{R}$ be as above. The correction $e_r(\theta) = e(\theta)h\left(\frac{1}{r}\|\theta\|\right)$ satisfies*

$$\left\|\partial_\theta^k [e_r]\right\|_\infty = O\left(r^{m-k}\right), \quad \text{as } r \to 0.$$

Proof. We have

$$\partial_\theta^k \left[e(\theta)h\left(\frac{1}{r}\|\theta\|\right)\right] = \sum_{k_1+k_2=k} \partial_\theta^{k_1} e(\theta) \partial_\theta^{k_2} h\left(\frac{1}{r}\|\theta\|\right)$$

Since $\partial_\theta^{k_2} h\left(\frac{1}{r}\|\theta\|\right) = 0$ whenever $\|\theta\| > 2r$ and $\left\|\partial_\theta^{k_2} h\left(\frac{1}{r}\|\theta\|\right)\right\|_\infty = O\left(r^{-k_2}\right)$, while $\left\|\partial_\theta^{k_1} e(\theta)\right\| = O(\|\theta\|^{m-k_1})$ we see from the above equation that

$$\left\|\partial_\theta^k \left[e(\theta)h\left(\frac{1}{r}\|\theta\|\right)\right]\right\|_\infty = O(r^{m-k_1} r^{-k_2}) = O(r^{m-k}),$$

as claimed. □

C.3. Escape Cones

We will approximate approximately homogeneous costs by homogeneous ones using the following approximation:

Lemma 19. *Suppose that $C(\theta) = H(\theta) + e(\theta)$ is (k, m)-approximately homogeneous around 0 as defined in 2. Let $\theta_0 \in \mathbb{R}^P$. It holds that for all $t \geq 0$ and all $\alpha > 0$, there is a finite constant $c_1(t)$ that does not depend on α such that*

$$\left\|\gamma_C(\alpha^{2-k}t, \alpha\theta_0) - \gamma_H(\alpha^{2-k}t, \alpha\theta_0)\right\| \leq c_1(t)\alpha^2.$$

Proof. Fix $\theta_0 \in \mathbb{R}^P$ and $t \geq 0$ and let $d_t = d(t, \theta_0) := \sup_{s \leq t} \gamma_C(s, \theta_0), \gamma_H(s, \theta_0)$. We can bound how fast the distance between the two paths γ_C and γ_H increases as follows:

$$\begin{aligned}
\partial_t \|\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0)\| &= -\frac{(\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0))^T}{\|\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0)\|} \left(\nabla H\left(\gamma_C(t, \theta_0)\right) + \nabla e\left(\gamma_C(t, \theta_0)\right) - \nabla H\left(\gamma_H(t, \theta_0)\right)\right) \\
&\leq \left(\sup_{\|\theta\| \leq d_t} \|\mathcal{H}H(\theta)\|_{op}\right) \|\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0)\| + \|\nabla e\left(\gamma_C(t, \theta_0)\right)\| \\
&\leq c' d_t^{k-2} \|\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0)\| + c d_t^k
\end{aligned}$$

where c comes from 4 and $c' = \sup_{\|x\| \leq 1} \|\mathcal{H}H(x)\|_{op}$. Applying Grönwall's inequality on $[0, t]$ to $A(s) = \|\gamma_C(s, \theta_0) - \gamma_H(s, \theta_0)\| + \frac{c}{c'} d_t^2$ (such that $\partial_s A(s) \leq c' d_t^{k-2} A(s)$), we obtain

$$\|\gamma_C(s, \theta_0) - \gamma_H(s, \theta_0)\| + \frac{c}{c'} d_t^2 = A(s) \leq A(0) e^{c' d_t^{k-2} s} = \frac{c}{c'} d_t^2 e^{c' d_t^{k-2} s}.$$

Hence $\|\gamma_C(t, \theta_0) - \gamma_H(t, \theta_0)\| \leq \frac{c}{c'} d_t^2 e^{c' d_t^{k-2} t}$ for all times $t \geq 0$. To finish the proof, one uses that for a fixed $t \geq 0$, $d(t, \alpha\theta_0) = O(\alpha)$ as $\alpha \to 0$, which is true because 0 is a saddle of C and H so their gradient tends to 0 with α. □

We define the ϵ-**Escape Cone** as the set $\mathcal{C}_\epsilon = \left\{\theta \in \mathbb{R}^P : \frac{H(\theta)}{\|\theta\|^k} < \frac{-s^* + \epsilon}{k}\right\}$ where we recall that s^* denotes the optimal escape speed.

Proposition 20. *For all $\epsilon > 0$ small enough there is a $r > 0$ such that*

1. *for any $\theta \in \partial\mathcal{C}_\epsilon$ with $\|\theta\| < r$, the negative of the gradient of C at θ points inside the cone, i.e. denoting by n the normal of $\mathcal{C}_\epsilon$ at θ pointing inside of the cone, we have $-\nabla C(\theta)^T n \geq 0$.*

2. *for any point θ inside the cone with $\|\theta\| < r$, we have* $\|\theta\|^{k-1}(-s^* - \epsilon) \leq \nabla C(\theta)^T \frac{\theta}{\|\theta\|} \leq \|\theta\|^{k-1}(-s^* + 2\epsilon)$.
3. *Let $\theta_0 \in \mathcal{C}_\epsilon$ and $\|\theta_0\| < r$, let T be the time when $\|\gamma_C(t, \theta_0)\| = r$. When $k = 2$ we have for all time $0 \leq t < T$*
$$\|\theta_0\| e^{-(-s^*+2\epsilon)t} \leq \|\gamma(t, \theta_0)\| \leq \|\theta_0\| e^{-(-s^*-\epsilon)t}$$
and when $k \neq 2$ we have for all time $0 \leq t < T$
$$\left[\|\theta_0\|^{-(k-2)} + (k-2)(-s^* + 2\epsilon)t\right]^{-\frac{1}{k-2}} \leq \|\gamma(t, \theta_0)\| \leq \left[\|\theta_0\|^{-(k-2)} + (k-2)(-s^* - \epsilon)t\right]^{-\frac{1}{k-2}}.$$

Proof. For all non-zero $\theta \in \mathbb{R}^P$, define $P_\theta = \left[I_d - \frac{\theta\theta^T}{\|\theta\|^2}\right]$, which is the orthogonal projection to the tangent space of $\mathbb{S}^{P-1}$ at $\frac{\theta}{\|\theta\|}$. Denote by $\partial\mathcal{C}_\epsilon$ the boundary of the cone and note that for any $\theta \in \partial\mathcal{C}_\epsilon$, it holds that $H(\theta) = (-s^* + \epsilon)/k$. Choose

$$r = r(\epsilon) = \min\left\{\frac{\inf_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon}\left\{\nabla H(\rho)^T P_\rho \nabla H(\rho)\right\}}{c \sup_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon}\left\{\sqrt{\nabla H(\rho)^T P_\rho \nabla H(\rho)}\right\}}, \sqrt{\frac{\epsilon}{c}}\right\},$$

where the constant c comes from 4.

(1) Let $\theta \in \partial\mathcal{C}_\epsilon$, so that $\frac{H(\theta)}{\|\theta\|^k} = H\left(\frac{\theta}{\|\theta\|}\right) = (-s^* + \epsilon)/k$. The normal pointing inside the cone is equal (up to a positive scaling) to

$$-\nabla_\theta\left(H\left(\frac{\theta}{\|\theta\|}\right)\right) = -\nabla H\left(\frac{\theta}{\|\theta\|}\right)\frac{P_\theta}{\|\theta\|}.$$

We then have that

$$\begin{aligned}
\left(-\nabla_\theta\left(\frac{H(\theta)}{\|\theta\|^k}\right)\right)^T(-\nabla C(\theta)) &= -\nabla H\left(\frac{\theta}{\|\theta\|}\right)^T \frac{P_\theta}{\|\theta\|}(-\nabla H(\theta) - \nabla e(\theta)) \\
&= \|\theta\|^{k-1} \nabla H\left(\frac{\theta}{\|\theta\|}\right)^T P_\theta \nabla H\left(\frac{\theta}{\|\theta\|}\right) + \nabla H\left(\frac{\theta}{\|\theta\|}\right)^T \frac{P_\theta}{\|\theta\|}\nabla e(\theta) \\
&\geq \|\theta\|^{k-1} \inf_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon}\left\{\nabla H(\rho)^T P_\rho \nabla H(\rho)\right\} \\
&\quad - \frac{1}{\|\theta\|} \sup_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon} \sqrt{\nabla H(\rho)^T P_\rho \nabla H(\rho)}\, \|\nabla e(\theta)\| \\
&\geq \|\theta\|^{k-1} \inf_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon}\left\{\nabla H(\rho)^T P_\rho \nabla H(\rho)\right\} \\
&\quad - c\|\theta\|^k \sup_{\rho \in \mathbb{S}^{P-1} \cap \partial\mathcal{C}_\epsilon}\left\{\sqrt{\nabla H(\rho)^T P_\rho \nabla H(\rho)}\right\},
\end{aligned}$$

where we used 4 for the last inequality. The right-hand side above is positive since $\|\theta\| < r(\epsilon) \leq \frac{\inf_{H(\rho)=s^*+\epsilon}\left\{\nabla H(\rho)^T P_\rho \nabla H(\rho)\right\}}{c \sup_{H(\rho)=s^*+\epsilon}\left\{\sqrt{\nabla H(\rho)^T P_\rho \nabla H(\rho)}\right\}}$.

(2) Let $\theta \in \mathcal{C}_\epsilon$. By 4, we have that

$$\begin{aligned}
\nabla C(\theta)^T\theta &= \partial_\lambda H(\lambda\theta)\Big|_{\lambda=1} + (\nabla e(\theta))^T\theta \\
&\leq kH(\theta) + c\|\theta\|^{k+2} \\
&= k\|\theta\|^k H\left(\frac{\theta}{\|\theta\|}\right) + c\|\theta\|^{k+2} \\
&\leq \|\theta\|^k\left(-s^* + \epsilon + c\|\theta\|^2\right) \\
&\leq \|\theta\|^k(-s^* + 2\epsilon)
\end{aligned}$$

where we used that $\|\theta\|^2 < r^2 \leq \frac{\epsilon}{c}$. In the other direction we obtain

$$\begin{aligned}\nabla C(\theta)^T\theta &= \partial_\lambda H(\lambda\theta)\Big|_{\lambda=1} + (\nabla e(\theta))^T\theta \\ &\geq kH(\theta) - c\|\theta\|^{k+2} \\ &= k\|\theta\|^k H\left(\frac{\theta}{\|\theta\|}\right) - c\|\theta\|^{k+2} \\ &\geq \|\theta\|^k\left(-s^* - c\|\theta\|^2\right) \\ &\geq \|\theta\|^k(-s^* - \epsilon).\end{aligned}$$

(3) Applying Grönwall's inequality generalized to polynomial bounds (Lemma 28), we have

$$\partial_t\|\theta\|^2 = -2\nabla C(\theta)^T\theta \leq c_1\left(\|\theta\|^2\right)^{\frac{k}{2}}.$$

□

Putting it all together, this guarantees that with probability 1 over the initialization, gradient flow escapes the saddle at a specific speed along a path θ^1:

Proposition 21. *Let $\theta_\alpha(t) = \gamma_C(t, \alpha\theta_0)$ for all $t \geq 0$. With prob. 1 over initialization (and under Assumption 17 when $L > 3$) there is a time horizon t^1_α that tends to ∞ as $\alpha \to 0$ and a path $(\theta^1(t))_{t\in\mathbb{R}}$ such that for all $t \in \mathbb{R}$, $\lim_{\alpha\to 0}\theta_\alpha(t^1_\alpha + t) = \theta^1(t)$. Furthermore, for all $\epsilon > 0$ s.t. $\epsilon < s^*/2$, there exists $T \in \mathbb{R}_+$ such that:*

(1) Shallow networks: $e^{(s^*-2\epsilon)(T+t)} \leq \left\|\theta^1(t)\right\| \leq e^{(s^*+\epsilon)(T+t)}$ for all $t \in \mathbb{R}$.

(2) Deep networks: $[(L-2)(s^*-2\epsilon)(T-t)]^{-\frac{1}{L-2}} \leq \left\|\theta^1(t)\right\| \leq [(L-2)(s^*+\epsilon)(T-t)]^{-\frac{1}{L-2}}$ for all $t < T$ (the path θ^1 is defined up to time T in this case).

Proof. We consider the gradient flow path $\tilde{\theta}_\alpha(t) = \gamma_H(t, \alpha\theta_0)$ on the k-homogeneous cost H. With prob. 1 (and under Assumption 17 when $L > 3$), we have $\frac{H(\tilde{\theta}_\alpha(t))}{\|\tilde{\theta}_\alpha(t)\|^L} \to -\frac{s^*}{k}$ as $t \to \infty$. In particular, for all $\epsilon > 0$, there exists a finite $t \in \mathbb{R}$ such that $\tilde{\theta}_{\alpha=1} \in \mathcal{C}_\epsilon$ and more generally, by Lemma 13, we have $\tilde{\theta}_\alpha(\alpha^{-(L-2)}t) = \alpha\tilde{\theta}_1(t) \in \mathcal{C}_\epsilon$. Lemma 19 then shows that there exists $\alpha_0 > 0$ such that for all $\alpha < \alpha_0$, it holds that $\theta_\alpha(\alpha^{-(L-2)}t) \in \mathcal{C}_\epsilon$. Define $t_0 := \inf\left\{t \in \mathbb{R} : \tilde{\theta}_{\alpha=1}(t_0) \in \mathcal{C}_\epsilon\right\} < +\infty$.

By Proposition 20, once the gradient flow path is inside $\mathcal{C}_\epsilon$, it cannot leave the escape cone until the norm $\|\theta_\alpha(t)\|$ is larger than some radius r. We define the time horizon $t^1_\alpha = \inf\left\{t \in \mathbb{R} : \|\theta_\alpha(t)\| = \frac{r}{2}\right\}$ and the escape path θ^1 as the limit $\theta^1(t) = \lim_{\alpha\to 0}\theta_\alpha(t^1_\alpha + t)$ for $t \in \mathbb{R}$ (the limit is well defined by continuity of $\theta \mapsto \gamma_C(t, \theta)$ and is an escape path by continuity of $\theta \mapsto \nabla\gamma_C(t, \theta)$). One can see that for any $t < 0$, there exists $\alpha > 0$ small enough such that $t_\alpha + t > \alpha^{-(L-2)}t_0$, thus it holds that $\theta^1(t) \in \mathcal{C}_\epsilon$ since for a small enough α, we have $\theta(t^1_\alpha + t) \in \mathcal{C}_\epsilon$. Proposition 20 then implies the escape rates for deep and shallow networks. □

C.4. Optimal Escape Paths

In this section, we define the notions of escape paths, optimal escape paths and we give a description of the optimal escape paths at the origin.

Proposition 21 shows that as $\alpha \searrow 0$ one has convergence to an escape path which escapes with an almost optimal speed $s^* - 2\epsilon$ for a small $\epsilon > 0$. We will show that the only such escape paths are the optimal escape paths, i.e. those that escape exactly at a speed of s^*, furthermore these escape paths are unique up to rotations of the network.

We understand well the escape paths of the homogeneous loss H, and want to use this knowledge to describe the escape paths of the locally homogogeneous loss C. We will show a bijection between the escape paths of H and those of C such that their speed is preserved, but only between the set of escape paths which escape faster than a certain speed. It seems that in general there is no speed-preserving bijection between escape paths, indeed while for shallow networks (when the saddle

is strict) one may apply the Hartman-Grobman Theorem to obtain a bijection, it does not preserve speed (since the bijection is in general not differentiable, only Hölder continuous).

This bijection is described by the following theorem (which is a more general version of Theorem 5 [] from the main - one simply needs to set $k = L$ and $m = 2L$ and $s^* = L \|H\|_\infty$ to recover theorem 5, i.e. the DLN case):

Theorem 22 (Theorem 5 of the main text)**.** *Let $C = H + e$ be a (k, m)-approximately homogeneous loss, where H is a polynomial.*

***When** $k = 2$**:** for all s_0 s.t. $s_0 > \frac{k\|H\|_\infty}{m-1}$ there is a unique bijection $\Psi : \mathcal{F}_H(s_0) \to \mathcal{F}_C(s_0)$ such that for all paths $x \in \mathcal{F}_C(s_0)$, we have $\|x(t) - \Psi(x)(t)\| = O(e^{(m-1)s_0 t})$ as $t \to -\infty$.*

***When** $k > 2$**:** for all $s_0 > \frac{k-1}{m-k+1} k \|H\|_\infty$ there is a unique bijection $\Psi : \mathcal{F}_H(s_0) \to \mathcal{F}_C(s_0)$ such that for all paths $x \in \mathcal{F}_C(s_0)$, we have $\|x(t) - \Psi(x)(t)\| = O((-t)^{-\frac{m-k+1}{k-2}})$ as $t \to -\infty$.*

Note that in the case $s_0 > k \|H\|_\infty$, the set $\mathcal{F}_H(s_0)$ is empty (and therefore so is $\mathcal{F}_C(s_0)$).

Proof. For $r > 0$, recall that C_r denotes the localization of the cost C as introduced in Section C.2. It is readily seen that for all $r > 0$, there is a bijection Ψ_r between $\mathcal{F}_C[s_0]$ and $\mathcal{F}_{C_r}[s_0]$ such that for all $x_0 \in \mathcal{F}_C[s_0]$, $\|x_0(t) - \Psi_r(x_0)(t)\| = O(0)$ (i.e. the difference is zero for small enough $t < 0$). We therefore only need to show a bijection between $\mathcal{F}_{C_r}[s_0]$ and $\mathcal{F}_H[s_0]$.

Consider a fast escaping path $x_0 \in \mathcal{F}_H(s_0)$ of the homogeneous approximation of the loss. The escape paths of the origin w.r.t. to gradient flow on the cost C_r are fixed points of the following map:

$$\Phi_{C_r} : x_0 \mapsto \left(t \mapsto \int_{-\infty}^{t} -\nabla C_r(x_0(u)) du \right).$$

Our strategy is simply to iterate this map starting from the path x_0 to find such a fixed point (note that any fixed point of Φ_{C_r} is differentiable by the fundamental theorem of calculus). We will show that this iteration converges to a gradient flow path x_0' of the cost C_r which is, as $t \to -\infty$, $O(e^{(m-1)s_0 t})$-close to x_0 when $k = 2$ and $O((-t)^{\frac{k-m-1}{k-2}})$-close to x_0 when $k > 2$.

For $c > 0$, let B_c be the *set of corrections*, that is the set of all paths $b : \mathbb{R}_- \to \mathbb{R}^P$ (which are Lebesgue measurable functions) such that when $k = 2$, $\|b(t)\| \leq c e^{(m-1)s_0 t}$ for all $t \leq 0$ and when $k > 2$, $\|b(t)\| \leq c(-t)^{\frac{k-m-1}{k-2}}$ for all $t \leq 0$.

The convergence of the iteration process follows from the fact that Φ is a contraction w.r.t. to some norm on the set of paths $x_0 + B_c$ (the set of possible corrections around x_0). Indeed, Lemma 24 (case $k = 2$, stated and proven in Section C.4.1) and Lemma 26 (case $k > 2$, stated and proven in Section C.4.2) show that for all $x_0 \in \mathcal{F}_H[s_0]$, there exist $r > 0$ small enough and $c > 0$ large enough, such that Φ_{C_r} is a contraction on $x_0 + B_c$ for some well-suited norm (defined below), hence guaranteeing the existence and uniqueness of a fixpoint x_0' of Φ_{C_r} (which is obtained by interating Φ_{C_r} infinitely many times). We thus define the map $\Psi : x_0 \mapsto \Psi(x_0) = x_0'$.

We need to show that Ψ has an inverse that maps a fast escaping path $y_0 \in \mathcal{F}_C(s_0)$ back to a path $\Psi^{-1}(y_0) \in \mathcal{F}_H(s_0)$. We iterate the map

$$\Phi_H : y_0 \mapsto \left(t \mapsto \int_{-\infty}^{t} -\nabla H(y_0(u)) du \right)$$

whose fixed points are the escape paths w.r.t. to gradient flow on the cost H. By a similar argument we can show that this map is a contraction on $x_0 + B_c$. Choosing $y_0 = x_0'$, this again implies the existence of a unique path $\Psi^{-1}(x_0')$ which is $O(e^{(m-1)s_0 t})$-close to x_0' when $k = 2$ and $O((-t)^{-\frac{m-k+1}{k-2}})$-close to x_0' when $k > 2$. Because $x_0' \in x_0 + B_c$ The uniqueness implies that since $x_0' = \Psi(x_0)$, the path $\Psi^{-1}(x_0')$ must be x_0. This shows that Ψ is a bijection and it is the only bijection between fast escaping paths with the property of mapping a path to a closeby path as in the statement of the theorem. □

C.4.1. SHALLOW NETWORKS

For the case $k = 2$, we consider the following norm for $\beta < (m-1)s_0$

$$\|b\|_\beta^2 = \int_{-\infty}^0 e^{-2\beta t} \|b(t)\|^2 \, dt$$

defined on the corrections $b \in B_c$, where the set of corrections B_c is the set of all paths $b : \mathbb{R}_- \to \mathbb{R}^P$ such that $\|b(t)\| \leq ce^{(m-1)s_0 t}$ for all $t \leq 0$. The condition $\beta < (m-1)s_0$ ensures that

$$\|b\|_\beta^2 = \int_{-\infty}^0 e^{-2\beta t} \|b(t)\|^2 \, dt \leq c^2 \int_{-\infty}^0 e^{2((m-1)s_0 - \beta)t} dt < \infty.$$

As a result the set $x_0 + B_c$ equipped with the distance induced by the norm $\|\cdot\|_\beta$ is a complete metric space.

We define the scalar product for two corrections $x, y \in B$

$$\langle x, y \rangle_\beta = \int_{-\infty}^\infty e^{-2\beta t} \langle x(t), y(t) \rangle \, dt.$$

We first state a few useful properties of $\langle \cdot, \cdot \rangle_\beta$. In the following, $\dot{x}$ is the path obtained by considering the derivative of x.

Lemma 23. *For any two corrections $x, y \in B$, we have*

1. $\langle x, \dot{y} \rangle_\beta = 2\beta \langle x, y \rangle_\beta - \langle \dot{x}, y \rangle_\beta$.
2. $\langle x, \dot{x} \rangle_\beta = \beta \|x\|_\beta^2$.
3. $\|x\|_\beta \leq \frac{1}{\beta} \|\dot{x}\|_\beta$.

Proof. The first point is obtained by integration by part:

$$\begin{aligned}\langle x, \dot{y} \rangle_\beta = \int_{-\infty}^\infty e^{-2\beta t} (x(t))^T \dot{y}(t) dt &= \int_{-\infty}^\infty 2\beta e^{-2\beta t} (x(t))^T y(t) dt - \int_{-\infty}^\infty e^{-2\beta t} \dot{x}(t) y(t) dt \\ &= 2\beta \langle x, y \rangle_\beta - \langle \dot{x}, y \rangle_\beta .\end{aligned}$$

The second point is a consequence of the first one, by taking $x = y$. Finally, the last point follows from the second one since $\|x\|_\beta^2 = \frac{1}{\beta} \langle x, \dot{x} \rangle_\beta \leq \frac{1}{\beta} \|x\|_\beta \|\dot{x}\|_\beta$, by Cauchy-Schwarz Inequality. □

We may now describe how for large enough β, one can guarantee that the map Φ_{C_r} is a contraction on the set $x_0 + B_c$:

Lemma 24. *Let $C_r = H + e_r$ be a localized $(2, m)$-approximately homogeneous loss as in 3, where H is a polynomial. Choose a $s_0 > \frac{2\|H\|_\infty}{m-1}$. There is a r small enough such that for any $x_0 \in \mathcal{F}_H[s_0]$ there is a constant c such that the map Φ_{C_r} is contraction on the set $x_0 + B_c = \left\{x_0 + b : \|b(t)\| \leq ce^{s_0(m-1)t}, \forall t < 0\right\}$ w.r.t. the norm on paths $\|\cdot\|_\beta$ for some β.*

Proof. We first show that for $r > 0$ small enough and $c > 0$ large enough, the image of $x_0 + B_c$ under Φ_{C_r} is contained in itself and then show that Φ_{C_r} is a contraction w.r.t. the norm $\|\cdot\|_\beta$ for an adequate β.

(1) **Self-map:** Let $x \in x_0 + B_c$, i.e. $x = x_0 + b$ for some $b \in B_c$, then using the linearity of ∇H and the fact x_0 is a gradient flow path of H, we obtain

$$\begin{aligned}\Phi_{C_r}(x)(t) &= -\int_{-\infty}^t \nabla H(x(u)) + \nabla e_r(x(u)) du \\ &= -\int_{-\infty}^t \nabla H(x_0(u)) + \nabla H(b(u)) + \nabla e_r(x(u)) du \\ &= x_0(t) - \int_{-\infty}^t \nabla H(b(u)) + \nabla e_r(x(u)) du.\end{aligned}$$

Writing $b'(t) = \int_{-\infty}^{t} [\nabla H(b(u)) + \nabla e_r(x(u))]\, du$ we need to show that $b' \in B_c$. We can bound $\|b'(t)\|$ by

$$\|b'(t)\| \leq \int_{-\infty}^{t} (\|\nabla H(b(u))\| + \|\nabla e_r(x_0(u))\| + \|\nabla e_r(x(u)) - \nabla e_r(x_0(u))\|)\, du.$$

Using the fact that for a map g with uniformly bounded Hessian $\|\mathcal{H}g\|_\infty < \infty$, we have $\|\nabla g(x) - \nabla g(y)\| \leq \|\mathcal{H}g\|_\infty \|x - y\|$, it follows that $\|\nabla H(b(u))\| \leq \sup_{z\in\mathbb{R}^P} \|\mathcal{H}H(z)\|_{\text{op}} \|b(u)\|$ and $\|\nabla e_r(x(u)) - \nabla e_r(x_0(u))\| \leq \sup_{z\in\mathbb{R}^P} \|\mathcal{H}e_r(z)\|_{\text{op}} \|b(u)\|$. The last term $\|\nabla e_r(x_0(u))\|$ can be bounded by $\frac{\sup_z \|\partial_z^m e_r(z)\|_{\text{op}} \|x_0(u)\|^{m-1}}{(m-1)!}$ since the first $m-1$ derivatives of e_r vanish at 0 (see point 1 of Lemma 29).

We therefore get

$$\|b'(t)\| \leq \int_{-\infty}^{t} \left(\left(\sup_{z\in\mathbb{R}^P} \|\mathcal{H}H(z)\|_{\text{op}} + \sup_{z\in\mathbb{R}^P} \|\mathcal{H}e_r(z)\|_{\text{op}} \right) \|b(u)\| + \sup_{z\in\mathbb{R}^P} \|\partial_z^m e_r(z)\|_{\text{op}} \|x_0(u)\|^{m-1} \right) du.$$

Since $\sup_{z\in\mathbb{R}^P} \|\mathcal{H}H(z)\|_{\text{op}} = 2\|H\|_{\text{op}}$ (where $\|\mathcal{H}H(x)\|_{\text{op}}$ is the operator norm of the Hessian $\mathcal{H}H$, while $\|H\|_{\text{op}} = \max_{x\in\mathbb{S}^{P-1}} |H(x)|$) and by Lemma 18 $\sup_{z\in\mathbb{R}^P} \|\mathcal{H}e_r(z)\|_\infty \leq \kappa_0 r^{m-2}$ and $\sup_{z\in\mathbb{R}^P} \|\partial_x^m e_r(z)\|_\infty \leq \kappa_1$

$$\begin{aligned}
&\leq \frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{s_0(m-1)} c e^{s_0(m-1)t} + \kappa_1 c_0 e^{s_0(m-1)t} \\
&\leq \left(\frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{s_0(m-1)} c + \kappa_1 c_0 \right) e^{s_0(m-1)t}
\end{aligned}$$

Since by assumption $s_0 > \frac{2\|H\|_{\text{op}}}{m-1}$, we can choose r small enough such that $\frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{s_0(m-1)} < 1$. We can then choose c large enough so that $\frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{s_0(m-1)} c + \kappa_1 c_0 \leq c$. With these choices of r and c, we obtain that $\|b'(t)\| \leq c e^{s_0(m-1)t}$ and therefore $b' \in B_c$ as needed.

(2) **Contraction:** We need to bound for any $x, y \in x_0 + B_c$

$$\|\Phi_{C_r}(x) - \Phi_{C_r}(y)\|_\beta^2 = \left\| t \mapsto \int_{-\infty}^{t} [\nabla C_r(x(u)) - \nabla C_r(y(u))]\, du \right\|_\beta^2,$$

for any $\beta < (m-1)s_0$.

From point (1), we know that $\Phi_{C_r}(x), \Phi_{C_r}(y) \in x_0 + B_c$ and hence $\|\Phi_{C_r}(x) - \Phi_{C_r}(y)\|_\beta^2 \leq \infty$ (since $\beta < (m-1)s_0$).

From point (3) of Lemma 23 we have:

$$\begin{aligned}
\left\| t \mapsto \int_{-\infty}^{t} [\nabla C_r(x(u)) - \nabla C_r(y(u))]\, du \right\|_\beta &\leq \frac{1}{\beta} \|t \mapsto \nabla C_r(x(t)) - \nabla C_r(y(t))\|_\beta \\
&\leq \frac{\sup_z \|\mathcal{H}C_r(z)\|_{\text{op}}}{\beta} \|x - y\|_\beta .
\end{aligned}$$

By the localization, we have $\sup_z \|\mathcal{H}C_r(z)\|_{\text{op}} \leq \sup_z \|\mathcal{H}H(z)\|_{\text{op}} + \sup_z \|\mathcal{H}e_r(z)\|_{\text{op}} \leq 2\|H\|_{\text{op}} + \kappa_0 r^{m-2}$.

Therefore to guarantee a contraction, we choose $\beta > 2\|H\|_{\text{op}} + \kappa_0 r^{m-2}$, so that $\frac{\sup_z \|\mathcal{H}C_r(z)\|_{\text{op}}}{\beta} < 1$. Therefore β lies in an open interval

$$\left(\frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{(m-1)s_0} \right) (m-1)s_0 < \beta < (m-1)s_0$$

which is non-empty since we have chosen r small enough in point (1) such that $\frac{2\|H\|_{\text{op}} + \kappa_0 r^{m-2}}{s_0(m-1)} < 1$. □

C.4.2. DEEP CASE

For the case $k > 2$, we consider the following norm

$$\|b\|_\alpha^2 = \int_{-\infty}^0 (-t)^{2\alpha-1} \|b(t)\|^2 \, dt.$$

If $\alpha < \frac{m-k+1}{k-2}$ then this norm is finite on any corrections $b \in B_c$ (i.e. if $\|b(t)\| \leq c(-t)^{-\frac{m-k+1}{k-2}}$), since

$$\|b\|_\alpha^2 \leq c^2 \int_{-\infty}^0 (-t)^{2(\alpha - \frac{m-k+1}{k-2})-1} dt < \infty.$$

The set $x_0 + B_c$ equipped with the distance $\|\cdot\|_\alpha$ therefore defines a complete metric space.

Again, for paths x, y such that $\|x\|_w, \|y\|_w < \infty$, we define the scalar product

$$\langle x, y \rangle_w = \int_{-\infty}^0 (-t)^{2\alpha-1} \langle x(t), y(t) \rangle \, dt.$$

Lemma 23 is now replaced by the following:

Lemma 25. *For any differentiable paths x, y with $\|x\|_w, \|y\|_w < \infty$, we have*

1. $\langle x, -t\dot{y} \rangle_w = 2\alpha \langle x, y \rangle_w - \langle -t\dot{x}, y \rangle_w$.
2. $\frac{1}{\alpha} \langle x, -t\dot{x} \rangle_w = \|x\|_w^2$.
3. $\|x\|_w \leq \frac{1}{\alpha} \|-t\dot{x}\|_w$.

Proof. The first point is obtained by integration by part:

$$\begin{aligned}
\langle x, -t\dot{y} \rangle_w &= \int_{-\infty}^0 (-t)^{2\alpha} x(t)\dot{y}(t) dt \\
&= \int_{-\infty}^0 2\alpha(-t)^{2\alpha-1} x(t)y(t) dt - \int_{-\infty}^0 (-t)^{2\alpha} \dot{x}(t)y(t) dt \\
&= 2\alpha \langle x, y \rangle_w - \langle -t\dot{x}, y \rangle_w .
\end{aligned}$$

Taking $x = y$, we obtain the second point. Finally, the last point follows from the second one since:

$$\|x\|_w^2 = \frac{1}{\alpha} \langle x, -t\dot{x} \rangle_w \leq \frac{1}{\alpha} \|x\|_w \|-t\dot{x}\|_w .$$

Under certain conditions, we can ensure that there is an α such that Φ is a contraction on $x_0 + B_c$ w.r.t. the norm $\|\cdot\|_\alpha$: □

Lemma 26. *Let $C_r = H + e_r$ be a localized (k, m)-approximately homogeneous loss as in 3, where H is a polynomial, with $k > 2$. Choose a $s_0 > \frac{k-1}{m-k+1} k \|H\|_\infty$. Let $x_0 \in \mathcal{F}_H[s_0]$, there exist $r > 0$ small enough, $c > 0$ large enough and $T < 0$ small enough, such that the map Φ is a contraction on the set $x_0 + B_{c,T} = \left\{x_0 + b : \|b(t)\| \leq ce^{s(m-1)t}, \forall t < T\right\}$ w.r.t. to the norm $\|\cdot\|_\alpha$ for some well-suited α.*

Proof. (1) **Self-map:** Let $x_0 + b \in x_0 + B_{c,T}$, we first show that $\Phi(x_0 + b) \in x_0 + B_{c,T}$. Let us rewrite

$$\begin{aligned}
\Phi(x_0 + b) &= -\int_{-\infty} \nabla C_r(x_0(u) + b(u)) du \\
&= -\int_{-\infty} \nabla H(x_0(u)) du + \int_{-\infty} \nabla H(x_0(u)) - \nabla C_r(x_0(u) + b(u)) du \\
&= x_0 + b'
\end{aligned}$$

where

$$\begin{aligned}
b'(t) &= \int_{-\infty}^{t} \nabla H(x_0(u)) - \nabla C_r(x_0(u) + b(u))du \\
&= \int_{-\infty}^{t} \nabla H(x_0(u)) - \nabla C_r(x_0(u))du \\
&+ \int_{-\infty}^{t} \nabla C_r(x_0(u)) - \nabla C_r(x_0(u) + b(u))du \\
&= -\int_{-\infty}^{t} \nabla e_r(x_0(u))du \\
&+ \int_{-\infty}^{t} \nabla C_r(x_0(u)) - \nabla C_r(x_0(u) + b(u))du.
\end{aligned}$$

Our goal is to show that $b'(t) \in B_c$, i.e. that $\|b'(t)\| \leq c(-t)^{\frac{k-m-1}{k-2}}$. We bound the two terms separately:

$$\|b'(t)\| \leq \left\| \int_{-\infty}^{t} \nabla e_r(x_0(u))du \right\| + \left\| \int_{-\infty}^{t} \left[\nabla C_r(x_0(u)) - \nabla C_r(x_0(u) + b(u)) \right] du \right\|.$$

The first term $\left\| \int_{-\infty}^{t} \nabla e_r(x_0(u))du \right\|$ is bounded by

$$\begin{aligned}
&\int_{-\infty}^{t} \|\nabla e_r(x_0(u))\| \, du \\
&\leq m \, \|\partial_x^m e_r\| \int_{-\infty}^{t} \|x_0(u)\|^{m-1} \, du \\
&\leq m\kappa s_0^{-\frac{m-1}{k-2}} (k-2)^{-\frac{m-1}{k-2}} \int_{-\infty}^{t} (-u)^{-\frac{m-1}{k-2}} \, du \\
&= m\kappa s_0^{-\frac{m-1}{k-2}} (k-2)^{-\frac{m-1}{k-2}} \frac{k-2}{m-k+1} (-t)^{\frac{k-m-1}{k-2}} \\
&= m\kappa s_0^{-\frac{m-1}{k-2}} \frac{(k-2)^{\frac{k-m-1}{k-2}}}{(m-k+1)} (-t)^{\frac{k-m-1}{k-2}}.
\end{aligned}$$

The second term $\left\| \int_{-\infty}^{t} \nabla C_r(x_0(u)) - \nabla C_r(x_0(u) + b(u))du \right\|$ is bounded by

$$\begin{aligned}
&\int_{-\infty}^{t} \|\nabla C_r(x_0(u)) - \nabla C_r(x_0(u) + b(u))\| \, du \\
&\leq \frac{\sup_z \left\| \partial_z^k C_r(z) \right\|_{\text{op}}}{(k-2)!} \int_{-\infty}^{t} \|b(u)\| \max\left\{ \|x_0(u)\|, \|x_0(u) + b(u)\| \right\}^{k-2} du
\end{aligned}$$

by Lemma 29. Let us first bound $\max\left\{ \|x_0(u)\|, \|x_0(u) + b(u)\| \right\}^{k-2}$ by

$$\begin{aligned}
\left(\|x_0(u)\| + \|b(u)\| \right)^{k-2} &\leq \left((s_0(k-2)(-u))^{-\frac{1}{k-2}} + c(-u)^{-\frac{m-k+1}{k-2}} \right)^{k-2} \\
&= (s_0(k-2)(-u))^{-1} + \sum_{i=1}^{k-2} \binom{k-2}{i} (s_0(k-2)(-u))^{-1+\frac{i}{k-2}} \, c(-u)^{-\frac{m-k+1}{k-2}i} \\
&= (s_0(k-2)(-u))^{-1} + (s_0(k-2)(-u))^{-1} \sum_{i=1}^{k-2} \binom{k-2}{i} (s_0(k-2))^{\frac{i}{k-2}} \, c^i (-u)^{-\frac{m-k}{k-2}i} \\
&\leq (s_0(k-2)(-u))^{-1} \left[1 + \sum_{i=1}^{k-2} \binom{k-2}{i} (s_0(k-2))^{\frac{i}{k-2}} \, c^i (-T)^{-\frac{m-k}{k-2}i} \right],
\end{aligned}$$

for any ϵ, we can choose $T < 0$ small enough so that $\max\{\|x_0(u)\|, \|x_0(u)+b(u)\|\}^{k-2}$ is bounded by $(s_0(k-2)(-u))^{-1}[1+\epsilon]$.

Using also the bounds $\frac{\sup_z \|\partial_z^k C_r(z)\|_{\text{op}}}{(k-2)!} \leq k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}$ and $\|b(u)\| \leq c(-u)^{-\frac{m-k+1}{k-2}}$, the second term $\left\|\int_{-\infty}^t \nabla C_r(x_0(u)) - \nabla C_r(x_0(u)+b(u))du\right\|$ can be bounded by

$$\begin{aligned}
&\frac{k(k-1)\|H\|_{\text{op}} + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(k-2)}(1+\epsilon)\int_{-\infty}^t c(-u)^{-\frac{m-k+1}{k-2}-1}du \\
&= \frac{k(k-1)\|H\|_{\text{op}} + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(m-k+1)}(1+\epsilon)c(-t)^{-\frac{m-k+1}{k-2}}.
\end{aligned}$$

We choose $r > 0$ small enough, $c > 0$ large enough, and $T < 0$ small enough so that $\frac{k(k-1)\|H\|_{\text{op}} + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(m-k+1)}(1+\epsilon) < 1$ and $\left[m\kappa s_0^{-\frac{m-1}{k-2}} \frac{(k-2)^{-\frac{m-k+1}{k-2}}}{m-k+1} + \frac{k(k-1)\|H\|_{\text{op}} + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(m-k+1)}(1+\epsilon)c\right] \leq c$ so that

$$\begin{aligned}
\|b'(t)\| &\leq \left[m\kappa s_0^{-\frac{m-1}{k-2}} \frac{(k-2)^{-\frac{m-k+1}{k-2}}}{m-k+1} + \frac{k(k-1)\|H\|_{\text{op}} + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(m-k+1)}(1+\epsilon)c\right](-t)^{-\frac{m-k+1}{k-2}} \\
&\leq c(-t)^{-\frac{m-k+1}{k-2}}
\end{aligned}$$

and therefore $b' \in B_c$.

(2) **Contraction:** We have, for any $x, y \in x_0 + B_c$

$$\begin{aligned}
\|\Phi(x) - \Phi(y)\|_w &= \left\|\int_{-\infty}^t \nabla C_r(x(s)) - \nabla C_r(y(s))ds\right\|_\alpha \\
&\leq \frac{1}{\alpha}\|-t(\nabla C_r(x) - \nabla C_r(y))\|_\alpha \\
&\leq \frac{\|\partial_k C_r(x_0)\|_\infty}{\alpha(k-2)!}\left\|-t\|x-y\|(\max\{\|x\|, \|y\|\})^{k-2}\right\|_\alpha
\end{aligned}$$

by Lemma 29. Using the same argument as in point (1) to bound $(\max\{\|x\|, \|y\|\})^{k-2}$, we obtain

$$\begin{aligned}
&\frac{k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}}{\alpha}\left\|-t\|x-y\|(s_0(k-2)(-t))^{-1}\right\|_\alpha \\
&\leq \frac{k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}}{\alpha s_0(k-2)}(1+\epsilon)\|x-y\|_\alpha
\end{aligned}$$

To obtain a contraction, we need to choose $\alpha > \frac{k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}}{s_0(k-2)}(1+\epsilon)$. To summarize α must lie within the two bounds:

$$\frac{k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}}{(m-k+1)s}(1-\epsilon)\frac{m-k+1}{k-2} < \alpha < \frac{m-k+1}{k-2}$$

which is possible since we have chosen r, c and T such that $\frac{k(k-1)\|H\|_\infty + \frac{\kappa}{(k-2)!} r^{m-k}}{(m-k+1)s}(1-\epsilon) < 1$. □

C.5. Proof of Theorem 2

We have now all the tools to prove the Theorem 4 of the main:

Theorem 27 (Theorem 2 of the main text)**.** *Assume that the largest singular value s_1 of the gradient of C at the origin $\nabla C(0) \in \mathbb{R}^{n_L \times n_0}$ has multiplicity 1. There is a deterministic gradient flow path $\underline{\theta}^1$ in the space of width-1 DLNs such that, with probability 1 if $L \leq 3$, and probability at least $1/2$ if $L > 3$, there exists an escape time t^1_α and a rotation R such that*

$$\lim_{\alpha \to 0} \theta_\alpha(t^1_\alpha + t) = RI^{(1\to w)}\underline{\theta}^1(t).$$

Proof. From Proposition 21 we know that with prob. 1 there is a time horizon t^1_α and an escape path such that $\lim_{\alpha\to 0} \theta_\alpha(t^1_\alpha + t) = \theta^1(t)$ which for any $\epsilon > 0$ escapes the origin at a rate of at least $e^{(s^*+\epsilon)t}$ for shallow networks and $[(k-2)(s^*+\epsilon)t]^{\frac{1}{2-k}}$ for deep networks, where $s^* = -L^{-\frac{L=2}{2}} s_1$.

Since the loss C^{NN} is $(L, 2L)$-approximately homogeneous, we can apply Theorem 22 to obtain that θ^1 must be in bijection with an escape path of the homogeneous loss H of the same speed. For small enough ϵ the only escape path of H of at least this speed are of the form $\rho^* e^{s^*(t+T)}$ for shallow networks and $\rho^* \left((k-1)s^*(-t-T)\right)^{-\frac{1}{k-1}}$ for some constant T and an optimal escape direction ρ^*. We therefore call θ^1 an optimal escape path since it belongs to the unique set of paths which escape at an optimal speed and are in bijection to the optimal escape directions.

Assuming that the largest singular value of s_1 of $\nabla C(0)$ has multiplicity 1, with singular vectors u_1, v_1, the optimal escape directions are of the form

$$\rho^* = RI^{(1\to w)}(\underline{\rho}^*) = \frac{1}{\sqrt{L}} RI^{(1\to w)}\left(-v_1^T, 1, \ldots, 1, u_1\right)$$

for any rotation R. In the width 1 network, there is an optimal escape path $\underline{\theta}^1$ corresponding to the escape direction $\underline{\rho}^*$, then by the unicity of the bijection of Theorem 22, the escape path $RI^{(1\to w)}(\underline{\theta}^1)$ is the unique optimal escape path escaping along $RI^{(1\to w)}(\underline{\rho}^*)$, as a result, we know that $\theta^1 = RI^{(1\to w)}(\underline{\theta}^1)$ for some rotation R. □

D. Technical Results

In this section, we state and prove a few technical lemmas used throughout the appendix.

Let us first prove a generalization of Grönwall's inequality for polynomial bounds:

Lemma 28. *Let $x : \mathbb{R}^+ \to \mathbb{R}$ which satisfy*

$$\partial_t x(t) \leq c x(t)^\alpha,$$

for some $c > 0$ and $\alpha > 1$. Then, for all $t < \frac{x(0)^{1-\alpha}}{c(\alpha-1)}$,

$$x(t) \leq \left[x(0)^{1-\alpha} - c(\alpha-1)t\right]^{-\frac{1}{\alpha-1}}.$$

Proof. Note that the function $y(t) = \left[x(0)^{1-\alpha} + c(1-\alpha)t\right]^{-\frac{1}{\alpha-1}}$ satisfies $y(0) = x(0)$ and for all $t < \frac{x(0)^{1-\alpha}}{c(\alpha-1)}$:

$$\partial_t y(t) = c y(t)^\alpha.$$

We conclude by showing that if $x(t) \leq y(t)$ then $x(s) \leq y(s)$ for all $t \leq s \leq \frac{x(0)^{1-\alpha}}{c(\alpha-1)}$: this follows from the fact that on the diagonal, i.e. when $x(t) = y(t)$, we have

$$\partial_t x(t) - \partial_t y(t) \leq c x(t)^\alpha - c y(t)^a = 0$$

which implies that the flow points towards the inside of $\{(x, y) : x \leq y\}$. □

Let us now state a lemma to bound the gradient of a cost C in terms of its high order derivatives:

Lemma 29. *Let C be a cost and k the largest integer such that $\partial_x^n C(0) = 0$ for all $n < k$ and $\left\|\partial_x^k C\right\|_\infty < \infty$, then*

1. For all x, $\|\nabla C(x)\| \leq \frac{\left\|\partial_x^k C\right\|_\infty \|x\|^{k-1}}{(k-1)!}$.
2. For all x, y, $\|\nabla C(x) - \nabla C(y)\| \leq \frac{1}{(k-2)!} \left\|\partial_x^k C\right\|_\infty \|x - y\| \left(\max\{\|x\|, \|y\|\}\right)^{k-2}$.

Proof. (1)

$$\begin{aligned}
\|\nabla C(x)\| &= \left\| \int_0^1 \mathcal{H}C(\lambda x)[x] d\lambda \right\| \\
&= \left\| \int_0^1 \int_0^{\lambda_1} \cdots \int_0^{\lambda_{k-2}} \partial_x^k C(\lambda_1 \cdots \lambda_{k-1} z_t)[x, \ldots, x] dt_1 \cdots dt_1 \right\| \\
&\leq \int_0^1 \int_0^{\lambda_1} \cdots \int_0^{\lambda_{m-2}} \left\| \partial_x^k C \right\|_\infty \|x\|^{k-1} dt_1 \cdots dt_1 \\
&\leq \frac{\left\| \partial_x^k C \right\|_\infty \|x\|^{k-1}}{(k-1)!}
\end{aligned}$$

(2) First note that $\nabla C(x) - \nabla C(y)$ is equal to

$$\int_0^1 \mathcal{H}C(z_t)[x-y] dt$$

where $z_t = tx + (1-t)y$. This can further be rewritten as

$$\int_0^1 \int_0^1 \partial_x^3 C(t_1 z_{t,t_1})[x-y, z_t] dt_1 \, dt.$$

Iterating this procedure, we obtain that $\nabla C(x) - \nabla C(y)$ equals

$$\int_0^1 \int_0^1 \int_0^{t_1} \cdots \int_0^{t_{k-3}} \partial_x^k C(t_1 \cdots t_{k-2} z_{t,t_1})[x-y, z_t, \ldots, z_t] dt_{k-2} \cdots dt_2 \, dt_1 \, dt.$$

Since the volume of the set $\{(t_1, \ldots, t_{k-2}) : 0 \geq t_1 \geq \cdots \geq t_{k-2} \geq 0\}$ is $\frac{1}{(k-2)!}$ we have

$$\begin{aligned}
\|\nabla C(x) - \nabla C(y)\| &\leq \int_0^1 \int_0^1 \int_0^{t_1} \cdots \int_0^{t_{k-3}} \left\| \partial_x^k C \right\| \|x-y\| \|z_t\|^{k-2} dt_{k-2} \cdots dt_2 \, dt_1 \, dt \\
&\leq \frac{1}{(k-2)!} \left\| \partial_x^k C \right\|_\infty \|x-y\| \left(\max\{\|x\|, \|y\|\} \right)^{k-2}.
\end{aligned}$$

□