
Implicit Bias in Deep Linear Classification: Initialization Scale vs Training Accuracy

Edward Moroshko

edward.moroshko@gmail.com
Technion

Suriya Gunasekar

suriya@ttic.edu
Microsoft Research

Blake Woodworth

blake@ttic.edu
Toyota Technological Institute at Chicago

Jason D. Lee

jasonlee@princeton.edu
Princeton University

Nathan Srebro

nati@ttic.edu
Toyota Technological Institute at Chicago

Daniel Soudry

daniel.soudry@gmail.com
Technion

Abstract

We provide a detailed asymptotic study of gradient flow trajectories and their implicit optimization bias when minimizing the exponential loss over “diagonal linear networks”. This is the simplest model displaying a transition between “kernel” and non-kernel (“rich” or “active”) regimes. We show how the transition is controlled by the relationship between the initialization scale and how accurately we minimize the training loss. Our results indicate that some limit behaviors of gradient descent only kick in at ridiculous training accuracies (well beyond 10^{-100}). Moreover, the implicit bias at reasonable initialization scales and training accuracies is more complex and not captured by these limits.

1 Introduction

The optimization trajectory, and in particular the “implicit bias” determining which predictor the optimization algorithm leads to, plays a crucial role in learning with massively under-determined models, including deep networks, where many zero-error predictors are possible [e.g., Telgarsky, 2013, Neyshabur et al., 2015, Zhang et al., 2017, Neyshabur, 2017]. Indeed, in several models we now understand how rich and natural implicit bias, often inducing sparsity of some form, can arise when training a multi-layer network with gradient descent [Telgarsky, 2013, Gunasekar et al., 2017, Li et al., 2018, Gunasekar et al., 2018b, Ji and Telgarsky, 2019b, Arora et al., 2019a, Nacson et al., 2019a, Lyu and Li, 2020]. This includes low ℓ_1 norm [Woodworth et al., 2020], sparsity in the frequency domain [Gunasekar et al., 2018b], low nuclear norm [Gunasekar et al., 2017, Li et al., 2018], low rank [Arora et al., 2019a, Razin and Cohen, 2020], and low higher order total variations [Chizat and Bach, 2020]. A different line of works focuses on how, in a certain regime, the optimization trajectory of neural networks, and hence also the implicit bias, stays near the initialization and mimics that of a kernelized linear predictor (with the kernel given by the tangent kernel) [Li and Liang, 2018, Du et al., 2019b, Jacot et al., 2018, Chizat et al., 2019, Allen-Zhu et al., 2019b, Du et al., 2019a, Zou et al., 2019, Allen-Zhu et al., 2019a, Arora et al., 2019b, Cao and Gu, 2019]. In such a “kernel regime” the implicit bias corresponds to minimizing the norm in some Reproducing Kernel Hilbert

	$\alpha < \infty$	$\alpha \rightarrow \infty$
$\epsilon > 0$		Kernel [Chizat et al., 2019]
$\epsilon \rightarrow 0$	Rich [Lyu and Li, 2020]	This Work

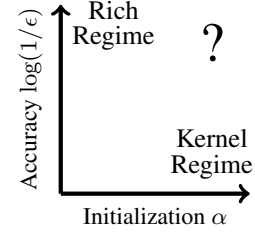


Figure 1: Kernel and rich limits.

Space (RKHS), and cannot yield the rich sparsity-inducing¹ inductive biases discussed above, and is perhaps not as “adaptive” as one might hope. It is therefore important to understand when and how learning is in the kernel regime, what hyper-parameters (*e.g.*, initialization, width, etc.) control the transition out and away from the kernel regime, and how the implicit bias (and thus the inductive bias driving learning) changes in different regimes and as a function of different hyper-parameters.

Initial work identified the **width** as a relevant hyper-parameter, where the kernel regime is reached when the width grows towards infinity [Li and Liang, 2018, Du et al., 2019b, Jacot et al., 2018, Allen-Zhu et al., 2019b, Du et al., 2019a, Zou et al., 2019]. But subsequent work by Chizat et al. [2019] pointed to the **initialization scale** as the relevant hyper-parameter, showing that models of any width enter the kernel regime as the scale of initialization tends towards infinity. Follow-up work by Woodworth et al. [2020] studied this transition in detail for regression with “diagonal linear networks” (see Section 3), showing how it is controlled by the interactions of **width**, **scale of initialization** and also **depth**. They obtained an exact expression for the implicit bias in terms of these hyper-parameters, showing how it transitions from ℓ_2 (kernel) implicit bias when the width or initialization go to infinity, to ℓ_1 (“rich” or “active” regime) for infinitesimal initialization, showing how a width of k has an equivalent effect to increasing initialization scale by a factor of \sqrt{k} . Studying the effect of initialization scale can therefore also be understood as studying the effect of width.

In this paper, we explore the transition and implicit bias for classification (as opposed to regression) problems, and highlight the effect of another hyper-parameter: **training accuracy**. A core distinction is that in classification with an exponentially decaying loss (such as the cross-entropy/logistic loss), the “rich regime” (*e.g.*, ℓ_1) implicit bias is attained with *any finite initialization*, and not only with infinitesimal initialization as in regression [Nacson et al., 2019a, Lyu and Li, 2020], and the scale of initialization does *not* effect the asymptotic implicit bias [Telgarsky, 2013, Soudry et al., 2018, Ji and Telgarsky, 2019a, Gunasekar et al., 2018a]. This is in seeming contradiction to the fact that even for classification problems, the kernel regime *is* attained when the scale of initialization grows [Chizat et al., 2019]. How can one reconcile the implicit bias not depending on the scale of initialization, with the kernel regime (and thus RKHS implicit bias) being reached as a function of initialization scale? The answer is discussed in detail in Section 2 and depicted in Figure 1: with the decaying losses used in classification, we can never get the loss to zero, only drive it arbitrarily close to zero. If we train infinitely long, then indeed we will eventually reach “rich” bias even for arbitrarily large initialization. On the other hand, if we only consider optimizing to within some arbitrarily high accuracy (*e.g.*, to loss 10^{-10}), we will find ourselves in the kernel regime as the initialization scale grows. But how do these hyper-parameters interact? Where does this transition happen? How accurately do we need to optimize to exit the “kernel” regime? To get the “rich” implicit bias? And what happens in between? What happens if both initialization scale and training accuracy go to infinity? And how is this affected by the depth?

To answer these questions quantitatively, we consider a concrete and simple model, and study classification using diagonal linear networks. This is arguably the simplest model that displays such

¹Sparsity in an implicit space can also be understood as feature search or “adaptation”. *E.g.* eigenvalue sparsity is equivalent to finding features which are linear combinations of input features, and sparsity in the space of ReLU units [*e.g.*, Savarese et al., 2019] corresponds to finding new non-linear features.

transitions², and as such has already been used to understand non-RKHS implicit bias and dynamics [Woodworth et al., 2020, Gissin et al., 2020].

We consider minimizing the exp-loss of a D -layer diagonal linear network on separable data, starting from initialization at scale α and optimizing until the training loss reaches a value of ϵ . In this case, the two extreme regimes depicted in Figure 1 correspond to implicit biases given by the ℓ_2 norm in the kernel regime and by the $\ell_{2/D}$ quasi-norm when $\epsilon \rightarrow 0$ for fixed α . We consider $\alpha \rightarrow \infty$ and optimizing to within $\epsilon(\alpha) \rightarrow 0$ and ask what happens for different joint behaviours of α and ϵ .

- We identify $\epsilon(\alpha) = \exp(-\Theta(\alpha^D))$ as the boundary of the kernel regime. When $\epsilon(\alpha) = \exp(-o(\alpha^D))$, the optimization trajectory follows that of a kernel machine, and the implicit bias is given by the ℓ_2 norm. But when $\epsilon(\alpha) = \exp(-\Omega(\alpha^D))$, the trajectory deviates from this kernel behaviour.
- Under additional condition (concerning stability of the support vectors), we characterize the behaviour at the transition regime, when $\epsilon(\alpha) = \exp(-\Theta(\alpha^D))$, and show that at this scaling the implicit bias is given by Woodworth et al.’s Q^D regularizer, which interpolates between the ℓ_1 -norm and the ℓ_2 -norm. This indicates roughly the following optimization trajectory for large enough initialization scale and under the support vector stability condition: we will first pass through the minimum $Q_\infty^D = \ell_2$ -norm predictor (or more accurately, max margin w.r.t. the ℓ_2 -norm), then traverse the path of minimizers of Q_μ^D , for $\mu \in (\infty, 0)$, until we reach the minimum $Q_0^D = \ell_1$ -norm predictor. We confirm such behaviour in simulations, as well as deviations from it when the condition does not hold.
- As suggested by our asymptotic theory, simulations in Section 5 show that even at moderate initialization scales, extreme training accuracy is needed in order to reach the asymptotic behaviour where the implicit bias is well understood. Rather, we see that even in tiny problems, at moderate and reasonable initialization scales and training accuracies, the implicit bias behaves rather different from both the “kernel” and “rich” extremes, suggesting that further work is necessary to understand the effective implicit bias in realistic settings.

Notation For vectors \mathbf{z}, \mathbf{v} , we denote by \mathbf{z}^k , $\exp(\mathbf{z})$, and $\log(\mathbf{z})$ the element-wise k th power, exponential, and natural logarithm, respectively; $\mathbf{z} \circ \mathbf{v}$ denotes element-wise multiplication; and $\mathbf{z} \propto \mathbf{v}$ implies that $\mathbf{z} = \gamma \mathbf{v}$ for a positive scalar γ . $\mathbf{1}$ denotes the vector of all ones.

2 Kernel and rich regimes in classification

We consider models as mappings $f : \mathbb{R}^p \times \mathcal{X} \rightarrow \mathbb{R}$ from trainable parameters $\mathbf{u} \in \mathbb{R}^p$ and input $\mathbf{x} \in \mathcal{X}$ to predictions. We denote by $F(\mathbf{u}) : \mathbf{x} \mapsto f(\mathbf{u}, \mathbf{x})$ the function implemented by the parameters \mathbf{u} . We will focus on models that are D -homogeneous in \mathbf{u} for some $D > 0$, i.e., such that $\forall c > 0, F(c\mathbf{u}) = c^D F(\mathbf{u})$. This includes depth- D linear and ReLU networks.

We consider minimizing $\mathcal{L}(\mathbf{u}) = \frac{1}{N} \sum_{n=1}^N \ell(f(\mathbf{u}, \mathbf{x}_n), y_n)$ for a given dataset $\{(\mathbf{x}_n, y_n) : n = 1, 2, \dots, N\}$ where $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a loss function. We will be mostly focus on binary classification problems, where $y_n \in \{-1, 1\}$, and with the exp-loss $\ell(\hat{y}, y) = \exp(-\hat{y}y)$, which has the same tail behaviour and thus similar asymptotic properties as the logistic or cross-entropy loss [e.g., Telgarsky, 2013, Soudry et al., 2018, Lyu and Li, 2020]. All our results and discussion refer to the exp-loss unless explicitly stated otherwise. We are concerned with understanding the trajectory of gradient descent, which we consider at the limit of infinitesimal stepsize, yielding the gradient flow dynamics,

$$\dot{\mathbf{u}}(t) = -\nabla \mathcal{L}(\mathbf{u}(t)), \quad (1)$$

where here and throughout $\dot{\mathbf{u}} = \frac{d\mathbf{u}}{dt}$.

Along the gradient flow path, $\mathcal{L}(\mathbf{u}(t))$ is monotonically decreasing and we consider cases where the loss is indeed minimized, i.e., converges to 0 for $t \rightarrow \infty$. However, if we stop the optimization trajectory at a large but finite t , which is what we do in practice, we optimize to some positive training loss $\epsilon(t) = \mathcal{L}(\mathbf{u}(t))$. We define $\tilde{\gamma}(t) = -\log(\epsilon(t))$ as the *training accuracy*. $\tilde{\gamma}(t)$ can also

²Our guiding principal is that when studying a new or not-yet-understood phenomena, we should first carefully and fully understand it in the simplest model that shows it, so as not to get distracted by possible confounders, and to enable a detailed analytic understanding.

be interpreted as the number of digits of the precision representing the training loss. This is related to the prediction margin $\gamma(t) = \min_n y_n \mathbf{x}_n^\top \mathbf{w}(t)$ as $\gamma(t) \leq \tilde{\gamma}(t) \leq \gamma(t) + \log(N)$ and was introduced as *smoothed margin* in Lyu and Li [2020].

For classification problems, we consider **separable data**, i.e., $\exists \mathbf{u}_* \in \mathbb{R}^p : \forall_n y_n f(\mathbf{u}_*, \mathbf{x}_n) > 0$, and so $\mathcal{L}(\gamma \mathbf{u}_*) \xrightarrow{\gamma \rightarrow \infty} 0$. But especially in high dimensions, there are many such separating predictors. If $\mathcal{L}(\mathbf{u}(t)) \rightarrow 0$, which of these does $\mathbf{u}(t)$ converge to? Of course $\mathbf{u}(t)$ does not converge, since to approach zero error it must *diverge*. Therefore, instead of the limit of the parameters, we study the limit of the *decision boundary* of the resulting classifier, which is given by $F\left(\frac{\mathbf{u}(t)}{\|\mathbf{u}(t)\|}\right)$.

Kernel Regime When the gradients $\nabla_{\mathbf{u}} f(\mathbf{u}, \mathbf{x}_n)$ do not change much during optimization, then $\mathbf{u}(t)$ behaves as if optimizing over a linearized model \mathbf{u} : $\tilde{f}(\mathbf{u}, \mathbf{x}) = f(\mathbf{u}(0), \mathbf{x}) + \langle \mathbf{u}, \phi(\mathbf{x}) \rangle$ where $\phi(\mathbf{x}) = \nabla_{\mathbf{u}} f(\mathbf{u}(0), \mathbf{x})$ is the feature map corresponding to the Tangent Kernel at initialization $K(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$ [Jacot et al., 2018, Allen-Zhu et al., 2019b, Du et al., 2019a]. Consider the trajectory $\bar{\mathbf{u}}(t)$ of gradient flow $\dot{\bar{\mathbf{u}}}(t) = -\nabla \tilde{\mathcal{L}}(\bar{\mathbf{u}}(t))$ on the loss of this linearized model $\tilde{\mathcal{L}}(\mathbf{u}) = \frac{1}{N} \sum_{n=1}^N \ell(\tilde{f}(\mathbf{u}, \mathbf{x}_n), y_n)$. Chizat et al. [2019] showed that any D -homogeneous model enters the kernel regime (i.e., behaves like a linearized model) when the scale of the initialization is large:

Theorem 1 (Adapted from Theorem 2.2 in Chizat et al. [2019]). *For any fixed time horizon $T > 0$, and any \mathbf{u}_0 such that $F(\mathbf{u}_0) = 0$, and for the exp-loss, consider the two gradient flow trajectories $\mathbf{u}(t)$ and $\bar{\mathbf{u}}(t)$, respectively, both initialized with $\mathbf{u}(0) = \bar{\mathbf{u}}(0) = \alpha \mathbf{u}_0$, for $\alpha > 0$. Then $\lim_{\alpha \rightarrow \infty} \sup_{t \in [0, T]} \|F(\mathbf{u}(t)) - \bar{F}(\bar{\mathbf{u}}(t))\| = 0$.*

For a linear model like \tilde{f} , the gradient flow $\bar{\mathbf{u}}$ converges in direction to the maximum margin solution in the corresponding RKHS norm [Soudry et al., 2018]. Combining this with Theorem 1, we have

$$\lim_{t \rightarrow \infty} \lim_{\alpha \rightarrow \infty} F\left(\frac{\mathbf{u}(t)}{\|\mathbf{u}(t)\|}\right) \propto \operatorname{argmin}_{f: \mathcal{X} \rightarrow \mathbb{R}} \|f\|_K \text{ s.t. } y_n f(\mathbf{x}_n) \geq 1, \quad (2)$$

where recall that K is the Tangent Kernel at the initialization and $\|f\|_K$ is the RKHS norm with respect to this kernel. It is important to highlight the crucial difference here compared to the corresponding statement for the squared loss [Chizat et al., 2019, Theorem 2.3]. For the squared loss we have that $\lim_{\alpha \rightarrow \infty} \sup_{t \in [0, \infty)} \|F(\mathbf{u}(t)) - \bar{F}(\bar{\mathbf{u}}(t))\| = 0$, i.e. the entire optimization trajectory converges uniformly to $\bar{\mathbf{u}}(t)$. But for the exp-loss, Theorem 1 only ensures convergence for prefixes of the path, up to finite time horizons T . The order of limits in (2) is thus crucial, while for the square loss the order of limits can be reversed.

Rich regime On the other hand, for or any finite initialization $\alpha \mathbf{u}_0$, the limit direction of gradient flow, when optimized indefinitely, gives rise to a different limit solution [Gunasekar et al., 2018a, Nacson et al., 2019a, Lyu and Li, 2020]:

Theorem 2 (Paraphrasing Theorem 4.4. in Lyu and Li [2020]). *Assume that the gradient flow trajectory in (1) minimizes the loss, i.e., $\mathcal{L}(\mathbf{u}(t)) \rightarrow 0$. Then, any limit point of $\left\{\frac{\mathbf{u}(t)}{\|\mathbf{u}(t)\|} : t > 0\right\}$ is along the direction of a KKT point of the following constrained optimization problem:*

$$\min_{\mathbf{u}} \|\mathbf{u}\|_2 \text{ s.t. } y_n f(\mathbf{u}, \mathbf{x}_n) \geq 1. \quad (3)$$

Compared to the kernel regime in (2), Theorem 2 suggests³ that

$$\lim_{\alpha \rightarrow \infty} \lim_{t \rightarrow \infty} F\left(\frac{\mathbf{u}(t)}{\|\mathbf{u}(t)\|}\right) \propto \text{stationary points of } \min_{f: \mathcal{X} \rightarrow \mathbb{R}} \mathcal{R}(f) \text{ s.t. } y_n f(\mathbf{x}_n) \geq 1, \quad (4)$$

where $\mathcal{R}(f) = \min_{\mathbf{u}} \|\mathbf{u}\|_2 \text{ s.t. } F(\mathbf{u}) = f$.

To understand this double limit, note that Theorem 2 ensures convergence for every α separately, and so also as we take $\alpha \rightarrow \infty$. For neural networks including linear networks, $\mathcal{R}(f)$ captures rich and

³Theorem 2 is suggestive of $\mathcal{R}(f)$ in (4) as the implicit induced bias in rich regime. However, although global minimizers of (3) and the RHS of (4) are equivalent, the same is not the case for stationary points. For the special cases of certain linear networks and the infinite width univariate ReLU network, stronger results for convergence in direction to the KKT points of (4) can be shown [Gunasekar et al., 2018b, Ji and Telgarsky, 2019b, Chizat and Bach, 2020].

often sparsity inducing inductive biases (e.g., nuclear norm, higher order total variations, ℓ_p bridge penalty for $p \leq 1$) that are not captured by RKHS norms [Gunasekar et al., 2017, Ji and Telgarsky, 2019a, Savarese et al., 2019, Ongie et al., 2019, Gunasekar et al., 2018b].

Contrasting (2) and (4) we see that if both the initialization scale α and the optimization time t go to infinity, the order in which we take the limits is crucial in determining the implicit bias, matching the depiction in Figure 1. Roughly speaking, if $\alpha \rightarrow \infty$ first (i.e., faster) and then $t \rightarrow \infty$, we end up in the kernel regime, but if $t \rightarrow \infty$ first (i.e., faster), we can end up with rich implicit bias corresponding to \mathcal{R} . The main question we ask is: *where the transition from kernel and rich regime happens, and what is the implicit bias when $\alpha \rightarrow \infty$ and $t \rightarrow \infty$ together?*

Since the time t for gradient flow does not directly correspond to actual “runtime”, and is perhaps less directly meaningful, we instead consider the optimization trajectory in terms of the training loss $\epsilon(t) = \mathcal{L}(\mathbf{u}(t))$. For some loss tolerance parameter ϵ , we follow the gradient flow trajectory until time t such that $\mathcal{L}(t) = \epsilon$ and ask *what is the implicit bias when $\alpha \rightarrow \infty$ and $\epsilon \rightarrow 0$ together?*

3 Diagonal linear network of depth D

In the remainder of the paper we focus on depth- D *diagonal linear networks*. This is a D -homogeneous model with parameters $\mathbf{u} = \begin{bmatrix} \mathbf{u}_+ \\ \mathbf{u}_- \end{bmatrix} \in \mathbb{R}^{2d}$ specified by:

$$f(\mathbf{u}, \mathbf{x}) = \langle \mathbf{u}_+^D - \mathbf{u}_-^D, \mathbf{x} \rangle \quad (5)$$

where recall that the exponentiation is element-wise.

As depicted in Figure 2, the model can be thought of as a depth- D network, with $D - 1$ hidden linear layers (i.e. the output is a weighted sum of the inputs), each consisting of $2d$ units, with the first hidden layer connected to the d inputs and their negations (depicted in the figure as another fixed layer), each unit in subsequent hidden layers connected to only a single unit in the preceding hidden layer, and the single output unit connected to all units in the final hidden layer. That is, the weight matrix at each layer $i = 1..D$ is a *diagonal* matrix $\text{diag}(\mathbf{u}^i)$. This presentation has $2dD$ parameters, as every layer has a different weight matrix. However, it is easy to verify that if we initialize all layers to the same weight matrix, i.e. with $\mathbf{u}^i = \mathbf{u}$, then the weight matrices will remain equal to each other throughout training, and so we can just use the $2d$ parameters in $\mathbf{u} \in \mathbb{R}^{2d}$ and take the weight matrix in every layer to be $\text{diag}(\mathbf{u})$, recovering the model (5).

Since all operations in a linear neural net are linear, the model just implements a linear mapping from the input \mathbf{x} to the output, and can therefore be viewed as an alternate parametrization of linear predictors. That is, the functions $F(\mathbf{u}) : \mathcal{X} \rightarrow \mathbb{R}$ implemented by the model is a linear predictor $F(\mathbf{u}) \in \mathcal{X}^*$ (since $\mathcal{X} = \mathbb{R}^d$, we also take $\mathcal{X}^* = \mathbb{R}^d$), and we can write $f(\mathbf{u}, \mathbf{x}) = \langle F(\mathbf{u}), \mathbf{x} \rangle$, where for diagonal linear nets $F(\mathbf{u}) = \mathbf{u}_+^D - \mathbf{u}_-^D$. In particular, for a trajectory $\mathbf{u}(t)$ in parameter space, we can also describe the corresponding trajectory $\mathbf{w}(t) = F(\mathbf{u}(t))$ of linear predictors.

The reason for using both the input features and their negation, and thus $2d$ units per layer, instead of a simpler model with $\mathbf{u} \in \mathbb{R}^d$ and $F(\mathbf{u}) = \mathbf{u}^D$ is two-fold: first, this allows the model to capture mixed sign linear predictors even with even depth. Second, this allows for scaling the parameters while still initializing at the zero predictor. In particular, we will consider initializing to $\mathbf{u}(0) = \alpha \mathbf{1}$, which for the diagonal neural net model (5) corresponds to $\mathbf{w}(0) = F(\mathbf{u}(0)) = 0$ regardless of the scale of α . Such *unbiased initialization* was suggested by Chizat et al. [2019] in order to avoid scaling a bias term when scaling the initialization.

Woodworth et al. [2020] provided a detailed study of diagonal linear net regression using the square loss $\ell(\hat{y}, y) = (\hat{y} - y)^2$. They showed that for an underdetermined problem (i.e. with multiple zero-error solutions), for any finite α the gradient flow trajectory with squared loss and initialization $\mathbf{u}(0) = \alpha \mathbf{1}$ converges to a zero-error (interpolating) solution minimizing the penalty $Q_{\alpha^D}^D$, where Q_{μ}^D is defined as:

$$Q_{\mu}^D(\mathbf{w}) = \sum_{i=1}^d q^D\left(\frac{\mathbf{w}_i}{\mu}\right), \text{ where } q^D(z) = \begin{cases} 2 - \sqrt{4 + z^2} + z \cdot \text{arcsinh}\left(\frac{z}{2}\right) & \text{for } D = 2 \\ \int_0^z h_D^{-1}(s) ds & \text{for } D > 2 \end{cases} \quad (6)$$

$$\text{and } h_D(s) = (1 - s)^{-\frac{D}{D-2}} - (1 + s)^{-\frac{D}{D-2}}.$$

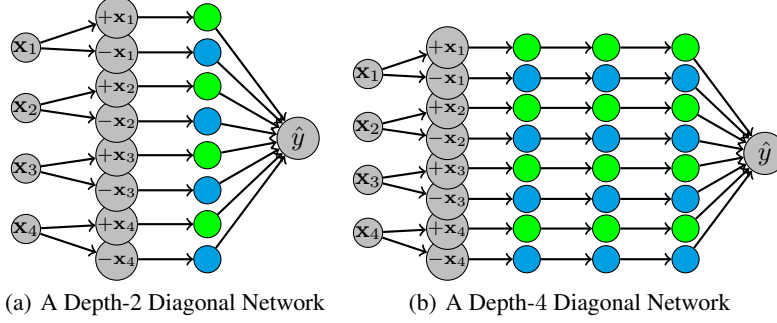


Figure 2: Diagonal linear networks.

For all $D \geq 2$, Q_μ^D with $\mu = \alpha^D$ interpolates between the ℓ_1 norm as $\alpha \rightarrow 0$, which corresponds to the rich limit previously shown by Gunasekar et al. [2017] and Arora et al. [2019a], and the ℓ_2 norm as $\alpha \rightarrow \infty$, which is the RKHS norm defined by the tangent kernel at initialization.

Can we identify a similar transition behavior between kernel and rich regimes with exponential loss?

4 Theoretical Analysis: Between the Kernel and Rich Regimes

We are now ready to state our main result, which describe the limit behaviour of gradient flow for classification with linear diagonal networks and the exp-loss, on separable data, and where the initialization scale α and the training accuracy $1/\epsilon$ go to infinity together.

First, we establish that if the data is separable, even though the objective $\mathcal{L}(\mathbf{u})$ is non-convex, gradient flow will minimize it, i.e. we will have $\mathcal{L}(\mathbf{u}(t)) \rightarrow 0$. In particular, gradient flow will lead to a separating predictor. We furthermore obtain a quantitative bound on how fast the training error decreases as a function of the initialization scale α , the depth D and the ℓ_2 separation margin of the data, $\gamma_2 = \max_{\|\mathbf{w}\|_2=1} \min_n y_n \mathbf{x}_n^\top \mathbf{w}$:

Lemma 3. For $D \geq 2$, any fixed α , and $\forall t$, $\mathcal{L}(\mathbf{u}(t)) \leq \frac{1}{1+2D^2\alpha^{2D-2}\gamma_2^2 t}$.

The proof appears in Appendix D. We now turn to ask which separating classifier we'd get to, if optimizing to within training accuracy ϵ and initialization of scale α , in terms of the relationship between these two quantities. To do so we consider different mapping $\epsilon(\alpha)$ between the initialization scale and training accuracy, such that $\epsilon : \mathbb{R}_{++} \rightarrow (0, 1]$ is strictly monotonic and $\lim_{\alpha \rightarrow \infty} \epsilon(\alpha) = 0$. We call $\epsilon(\alpha)$ the *stopping accuracy function*. For each α , we follow the gradient flow trajectory until time T_α such that $\mathcal{L}(\mathbf{u}(T_\alpha)) = \epsilon(\alpha)$ and denote $\hat{\mathbf{w}}_\alpha = \frac{\mathbf{w}(T_\alpha)}{\gamma(T_\alpha)}$. We are interested in characterizing the limit point $\hat{\mathbf{w}} = \lim_{\alpha \rightarrow \infty} \hat{\mathbf{w}}_\alpha$ for different stopping accuracy functions $\epsilon(\alpha)$, assuming this limit exists.

The Kernel Regime We start with showing that if $\epsilon(\alpha)$ goes to zero slowly enough, namely if $\log 1/\epsilon$ is sublinear in α^D , then with large initialization we obtain the ℓ_2 bias of the kernel regime:

Theorem 4. For $D \geq 2$, if $\epsilon(\alpha) = \exp(-o(\alpha^D))$ then $\hat{\mathbf{w}} = \arg\min_{\mathbf{w}} \|\mathbf{w}\|_2$ s.t. $\forall n, y_n \mathbf{w}^\top \mathbf{x}_n \geq 1$.

The proof for $D = 2$ appears in Appendix F.1 and the proof for $D > 2$ appears in Appendix G.1.

Escaping the kernel regime Theorem 4 shows that escaping the kernel regime requires optimizing to higher accuracy, such that $\log 1/\epsilon$ is at least linear in the initialization scale α^D . Complementing Theorem 4, we show a converse: that indeed the linear scaling $\log 1/\epsilon = \Theta(\alpha^D)$ is the transition point out of kernel regime, and once $\log 1/\epsilon = \Omega(\alpha^D)$ we no longer obtain the kernel ℓ_2 bias.

To show this, we first state a condition about the stability of support vectors⁴. Given a dataset $S = \{(\mathbf{x}_n, y_n) : n = 1, 2, \dots, N\}$ and a stopping accuracy function $\epsilon(\alpha)$:

Condition 5 (Stability condition). For all $k \in [N]$ such that $y_k \mathbf{x}_k^\top \hat{\mathbf{w}} > 1$, and large enough α , there exists $\epsilon^*(\alpha) = \exp(-o(\alpha^D))$ and $\rho_0 > 1$ such that $\forall t$ with $\mathcal{L}(t) \in [\epsilon(\alpha), \epsilon^*(\alpha)] : \frac{y_k \mathbf{x}_k^\top \mathbf{w}(t)}{\gamma(t)} \geq \rho_0$.

⁴Data point (\mathbf{x}_k, y_k) is a support vector at time t if $y_k \mathbf{x}_k^\top \mathbf{w}(t) = \min_n y_n \mathbf{x}_n^\top \mathbf{w}(t) = \gamma(t)$.

We say that Condition 5 holds *uniformly* if given a dataset S it holds for *all* stopping functions $\epsilon(\alpha) = \exp(-\Omega(\alpha^D))$.

For the linearized model Condition 5 holds almost surely (uniformly), see details in Appendix E. Therefore, if Condition 5 does not hold, it follows that we exit the kernel regime.

If Condition 5 does hold we show in Theorem 6 that when $\epsilon = \exp(-\Theta(\alpha^D))$ we will be in the intermediate regime, leading to max-margin solution with respect to Q^D function in (6), and again deviate from the kernel regime.

Theorem 6. *Under Condition 5, for $D \geq 2$, if $\lim_{\alpha \rightarrow \infty} \frac{\alpha^D}{\log(1/\epsilon(\alpha))} = \mu > 0$, then for Q_μ^D as defined in (6), $\hat{\mathbf{w}} = \operatorname{argmin}_{\mathbf{w}} Q_\mu^D(\mathbf{w})$ s.t. $\forall n, y_n \mathbf{x}_n^\top \mathbf{w} \geq 1$.*

To prove Theorem 6 we show that the following KKT conditions hold:

$$\exists \nu \in \mathbb{R}_{\geq 0}^N \text{ s.t. } \nabla Q_\mu^D(\hat{\mathbf{w}}) = \sum_{n=1}^N \nu_n y_n \mathbf{x}_n, \quad \forall n : y_n \mathbf{x}_n^\top \hat{\mathbf{w}} \geq 1, \quad \forall n : \nu_n (y_n \mathbf{x}_n^\top \hat{\mathbf{w}} - 1) = 0.$$

The result then follows from convexity of $Q_\mu^D(\mathbf{w})$.

The proof for $D = 2$ appears in Appendix F.2 and the proof for $D > 2$ appears in Appendix G.2.

The Rich Limit Above we saw that $\log 1/\epsilon$ being linear in α is enough to leave the kernel regime, i.e. with this training accuracy and beyond the trajectory no longer behaves as if we were training a kernel machine. We also saw that under Condition 5, when $\log 1/\epsilon$ is exactly linear in α , we are in a sense in an “transition” regime, with bias given by the Q_μ^D penalty which interpolates between ℓ_2 (kernel) and ℓ_1 (the rich limit for $D = 2$). Next we see that once the accuracy $\log 1/\epsilon$ is superlinear, and again under Condition 5, we are firmly at the rich limit:

Theorem 7. *Under Condition 5, for $D \geq 2$ if $\epsilon(\alpha) = \exp(-\omega(\alpha^D))$ then $\hat{\mathbf{w}} = \operatorname{argmin}_{\mathbf{w}} \|\mathbf{w}\|_1$ s.t. $\forall n : y_n \mathbf{x}_n^\top \mathbf{w} \geq 1$.*

For $D = 2$ the result holds also with a weaker condition, when $\epsilon^*(\alpha)$ in Condition 5 is replaced with $\epsilon^*(\alpha) = \exp\left[-o\left(\alpha^2 \log \frac{\log(1/\epsilon(\alpha))}{\alpha^2}\right)\right]$.

The proof for $D = 2$ appears in Appendix F.3 and the proof for $D > 2$ appears in Appendix G.3.

For $D > 2$ we know from Theorem 2 that the implicit bias in the rich limit is given by an $\ell_{2/D}$ quasi norm penalty, and not by the ℓ_1 penalty as in Theorem 7. It follows that when Condition 5 holds, the max ℓ_1 -margin predictor must also be a first order stationary point of the $\ell_{2/D}$ max margin problem. As we demonstrate in Section 5, this is certainly not always the case, and for many problems the max ℓ_1 -margin is not a stationary point for the $\ell_{2/D}$ max margin problem—in those cases Condition 5 does not hold. It might well be possible to show that a super-linear scaling is sufficient to reach the rich limit, be it the max ℓ_1 -margin for depth two, or the max $\ell_{2/D}$ -margin (or a stationary point for this non-convex criteria) for higher depth, and we hope future work will address this issue.

Role of depth From Theorem 6 we have that asymptotically $\epsilon(\alpha) = \exp(-\alpha^D/\mu)$. Woodworth et al. [2020] analyzed the Q^D function and concluded that in order to have δ approximation to ℓ_1 limit (achieved for $\mu \rightarrow 0$) we need to have $\mu = \exp(-1/\delta)$ for $D = 2$, and $1/\mu = \text{poly}(1/\delta)$ for $D > 2$. We conclude that in order to have δ approximation to ℓ_1 limit we need the training accuracy to be $\epsilon = \exp(-\alpha^D \exp(1/\delta))$ for $D = 2$, and $\epsilon = \exp(-\alpha^D \text{poly}(1/\delta))$ for $D > 2$. Thus, depth can mitigate the need to train to extreme accuracy. We confirm such behaviour in simulations.

5 Numerical Simulations and Discussion

We numerically study optimization trajectories to see whether we can observe the asymptotic phenomena studied at finite initialization and accuracy. We focus on low dimensional problems, where we can plot the trajectory in the space of predictors. In all our simulations we employ the Normalized GD algorithm, where the gradient is normalized by the loss itself, to accelerate convergence [Nacson et al., 2019b]. In all runs the learning rate was small enough to ensure gradient flow-like dynamics (always below 10^{-3}).

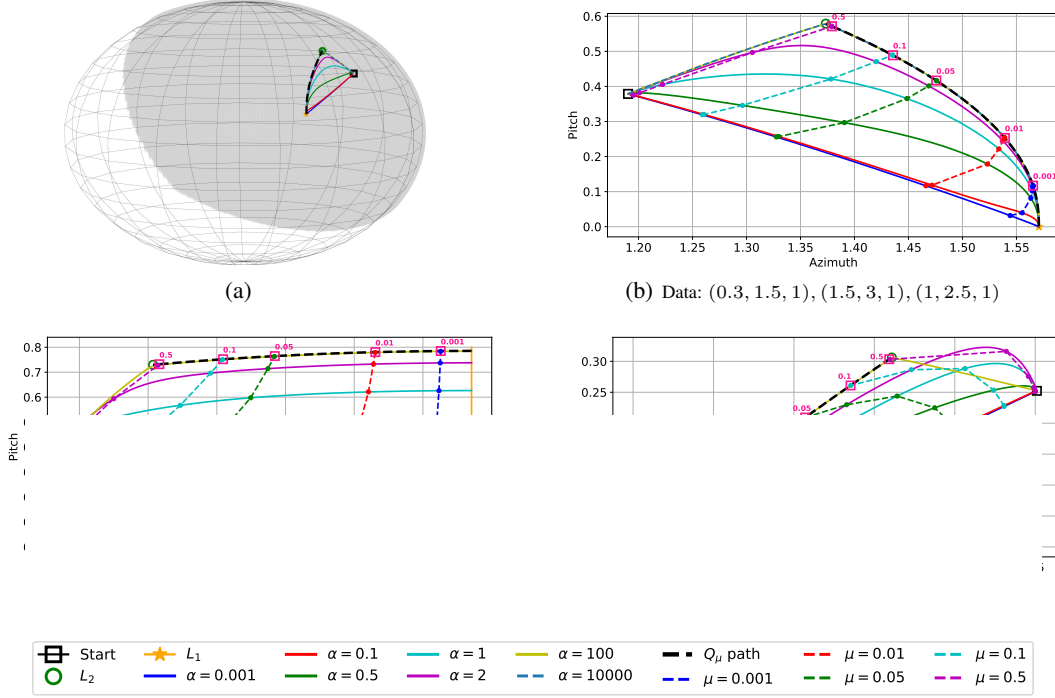


Figure 3: Optimization trajectories for 3 simple datasets in depth 2 linear diagonal network (b-d). Each point in Azimuth-Pitch plane represents a normalized classifier $\mathbf{w}/\|\mathbf{w}\|_2$. The curves corresponding to α in the legend are the entire gradient flow trajectories initialized with the respective α . The curves corresponding to μ are end points of gradient flow trajectories for different α with stopping criteria set as $\epsilon(\alpha) = \exp(-\alpha^2/\mu)$. The pink squares represent the directions along the Q_μ max-margin path for the appropriate μ marked near the square. The dynamics in (b) takes place on a small part of the sphere as shown in (a), where the grey area represents all separating directions.

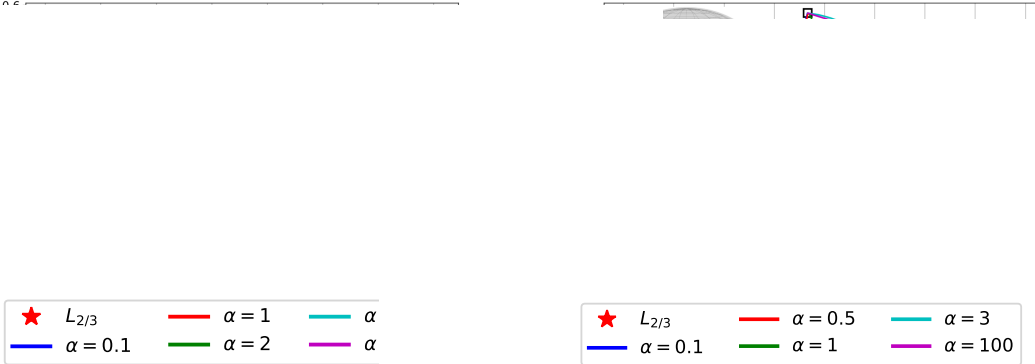


Figure 4: Optimization trajectories for $D = 3$. The legend for ℓ_1 , ℓ_2 , start and Q_μ path is the same as in Figure 3

Gradient flow trajectories In Figure 3 we plot trajectories for training depth $D = 2$ diagonal linear networks in dimension $d = 3$, on several constructed datasets, each consisting of three points. The trajectory $\mathbf{w}(t)$ in this case is in \mathbb{R}^3 , and so the corresponding binary predictor given by the normalization $\mathbf{w}(t)/\|\mathbf{w}(t)\|_2$ lies on the sphere (panel (a)). We zoom in on a small section of the sphere, and plot the trajectory of $\mathbf{w}(t)/\|\mathbf{w}(t)\|_2$ — the axes correspond to coordinates on the sphere (given as azimuth and pitch). The first step taken by gradient flow will be to the predictor proportional to the average of the data, $\frac{1}{N} \sum_n y_n \mathbf{x}_n$, and we denote this as the “start” point. The grey area in the sphere represents classifiers separating the data (with zero misclassification error), and thus directions where the loss can be driven to zero. The question of “implicit bias” is which of these classifiers the trajectory will converge to. With infinitesimal (small) stepsizes, the trajectories always remain inside this area (*i.e.*, just finding a separating direction is easy), and in a sense the entire optimization trajectory is driven by the implicit bias.

Panel 3(b) corresponds to a simple situation, with a unique ℓ_1 -max-margin solution, and where the support vectors for the ℓ_2 -max-margin and ℓ_1 -max-margin are the same (although the solutions are not the same!), and so the support vectors do not change throughout optimization and Condition 5 holds uniformly. For large initialization scales ($\alpha = 100$ and $\alpha = 10000$, which are indistinguishable here), the trajectory behaves as the asymptotic theory tells us: from the starting point (average of the data), we first go to the ℓ_2 -max-margin solution (green circle, and recall that this is also the Q_∞^2 -max-margin solution), and then follow the path of Q_μ^2 -max-margin predictors for μ going from ∞ to zero (this path is indicated by the dashed black lines in the plots), finally reaching the ℓ_1 -max-margin predictor (orange star, and this is also the Q_0^2 -max-margin solution). For smaller initialization scales, we still always reach the same endpoint as $\epsilon \rightarrow 0$ (as assured by the theory), but instead of first visiting the ℓ_2 -max-margin solution and traversing the Q_μ^2 path, we head more directly to the ℓ_1 -max-margin predictor. This can be thought of as the effect of initialization on the implicit bias and kernel regime transition: with small initialization we will never see the kernel regime, and go directly to the “rich” limit, but with large initialization we will initially remain in the kernel regime (heading to the ℓ_2 -max-margin), and then, only when the optimization becomes very accurate, escape it gradually.

To see the relative effect of scale and training accuracy, and following our theory, we plot for different values of μ , the different points along trajectories with initialization α such that we fix the stopping criteria as $\epsilon(\alpha) = \exp(-\alpha^2/\mu)$ (dashed cross-lines in the plots). Our theory indicates that for any value of μ , as $\alpha \rightarrow \infty$, the dashes line would converge to the Q_μ^2 -max-margin (a specific point on the Q_μ^2 path), and this is indeed confirmed in Panel 3(b), where the dashed lines converge to pink squares, which correspond to points along the Q_μ^2 path for the appropriate μ values. The clear correspondence between points with the same relationship μ between initialization and accuracy confirms that also at relatively small initialization scales, this parameter is the relevant relationship between them.

From the value μ we can also extract the actual training accuracy. E.g., we see that for a relatively large initialization scale $\alpha = 100$, escaping the kernel regime (getting to the first pink square just removed from the ℓ_2 -max-margin solution, with $\mu = 0.5$), requires optimizing to loss $\epsilon = \exp(-10^4/0.5) \approx 10^{-8700}$, while getting close to the asymptotic limit of the trajectory (the last pink square, with $\mu = 0.001$) requires $\epsilon \approx 10^{-4 \cdot 10^6}$ (that’s four millions digits of precision). Even with reasonable initialization at scale $\alpha = 1$, getting to this limit requires $\epsilon \approx 10^{-434}$.

Panel (c) displays the trajectories for another dataset where Condition 5 holds uniformly, but the ℓ_1 -max-margin solution is not unique. Although will always (eventually) converge to a ℓ_1 -max-margin solution, which one we converge to, and thus the implicit bias, *does* depends on the initialization. Panel (d) shows a situation where the support vectors for the ℓ_2 -max-margin and ℓ_1 -max-margin solutions are different, and so change during optimization, and Condition 5 does not hold uniformly. The condition does hold for $\mu > \mu_0$ where $\mu_0 \approx 0.04$, but does not hold otherwise. In this case, even for very large initialization scales α , the trajectory does not follow the Q_μ^2 path entirely. It is interesting to note that the Q_μ^2 is not smooth, which perhaps makes it difficult to follow. Around the kink in the path (at $\mu = \mu_0$), it does seem that larger initialization scales are able to follow it a bit longer, and so perhaps with huge initialization (beyond our simulation ability), the trajectory *does* follow the Q_μ^2 path. Whether the trajectory follows the path for sufficiently large initialization scales even when Condition 5 fails, or whether we can characterize its limit behaviour otherwise, remains an open question.

In Figure 4 we show optimization trajectories for depth 3 linear diagonal network where Condition 5 does not hold uniformly. As we can observe, in both cases the trajectory for large α will first go to ℓ_2 max-margin, then stay near the Q_μ path, until the trajectory starts to deviate in a direction of $\ell_{2/3}$ max-margin solution, whereas the Q_μ path continues to the ℓ_1 max-margin point. In Figure 4(b) we observe that there is a local minimum point at (0,0) and for small α the trajectory converges to it. Further discussion about convergence to local minima in high dimension appears in Appendix H.

Initialization Scale vs Training Accuracy Using the same dataset from Figure 3(b) we examine the question: given some initialization scale, how long we need to optimize to be in the rich regime? In Figure 5(a) we demonstrate how the initialization and depth affect the convergence rate to the rich regime. Specifically, we chose two reasonable initialization scales α^D , namely 0.1 and 1, and show their convergence to the ℓ_1 max-margin solution as a function of the training accuracy $\tilde{\gamma}$ (recall that $\tilde{\gamma} = \log(1/\epsilon)$), for depths $D = 2, 3, 10$. We chose this dataset so that the rich regime is the same

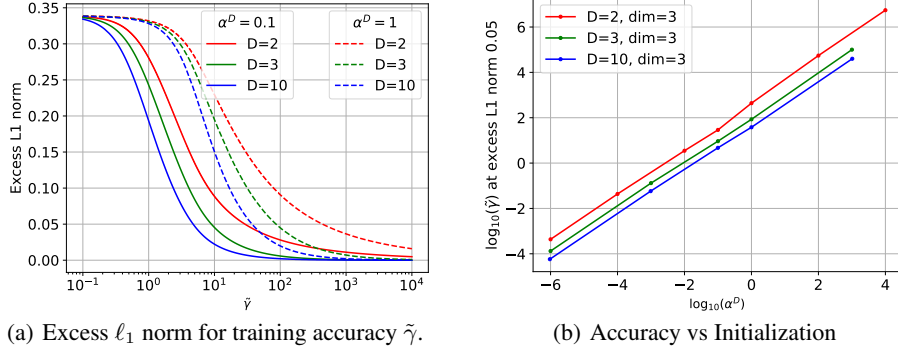


Figure 5: In (a) we plot the excess ℓ_1 norm, defined as $\|\mathbf{w}(t)\|_1 / \|\mathbf{w}_{\ell_1}\|_1 - 1$ where \mathbf{w}_{ℓ_1} is the ℓ_1 max-margin (minimum norm) solution, as a function of $\tilde{\gamma}$. For a fixed excess ℓ_1 norm of 0.05, in (b) we plot how long we need to optimize, given some initialization scale, to obtain the 5% closeness to ℓ_1 max-margin solution.

for all depths (i.e., the minimum ℓ_1 norm solution corresponds also to the minimum $\ell_{2/3}$ and $\ell_{2/10}$ quasi-norm solutions).

As previously discussed, even on such a small dataset (with three datapoints) we are required to optimize to an incredibly high precision — in order to converge near the rich regime. Notably, the situation is improved and we converge faster to the rich regime when the initialization scale is smaller and/or the depth is larger. Unfortunately, taking the initialization scale to 0 will increase the time needed to escape the vicinity of the saddle point $\mathbf{u} = 0$. For example, Shamir [2019] showed that exiting the initialization may take exponential-in-depth time.

In Figure 5(b) we examine the relative scale between (log) accuracy and (log) initialization needed to obtain 5% closeness to ℓ_1 max-margin. Based on the asymptotic result in Theorem 6, we expect that $\tilde{\gamma} \propto \alpha^D$ or $\log(\tilde{\gamma}) = a \log(\alpha^D) + b$, for some constants $a = 1$ and $b \in \mathbb{R}$. And indeed, in Figure 5(b) we obtain $a = 1$, as expected. Note that although our theoretical results are valid for $\alpha \rightarrow \infty$, we obtain the same accuracy vs initialization rate also for small α . Moreover, the intercept of the lines decreases when increasing the depth D , which matches the observed behavior on Figure 5(a) and to the discussion about the effect of depth in Section 4.

Broader Impact

The goal of this work is to shed light on the implicit bias hidden in the training process of deep networks. These results may enable a better understanding of how hyperparameters select the types of solutions that deep networks converge to, which in turn affect their final generalization performance and hidden biases. This could lead to better performance guarantees or to improved training algorithms which quickly converge to beneficial types of biases. Eventually, we believe progress on these fronts can transform deep learning from the current nascent “alchemy” age (where all the “knobs and levers” of the model and the training algorithm are tuned mostly heuristically during research and development), to a more mature field (like “chemistry”), which can be seamlessly integrated in many real world applications that require high performance, safety, and fair decisions.

Our guiding principal is that when studying a new or not-yet-understood phenomena, we should first study it in the simplest model that shows it, so as not to get distracted by possible confounders, and to enable a detailed analytic understanding, *e.g.*, when understanding or teaching many statistical issues, we would typically start with linear regression, understand the phenomena there, and *then* move on to more complex models. In the specific case here, one of the few models where we have an analytic handle on the implicit bias in the “rich” regime are linear diagonal networks, and it would be very optimistic to hope to get a detailed analytic description of the more complex phenomena we study in models where we can’t even understand the endpoint.

Acknowledgements The research of DS was supported by the Israel Science Foundation (grant No. 31/1031), and by the Taub Foundation. This work was partially done while SG, JDL, NS, and DS were visiting the Simons Institute for the Theory of Computing. BW is supported by a

Google Research PhD fellowship. JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0303, the Sloan Research Fellowship, and NSF CCF 2002272.

References

- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6155–6166, 2019a.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019b.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pages 7411–7422, 2019a.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, 2019b.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, pages 10835–10845, 2019.
- Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arXiv preprint arXiv:2002.04486*, 2020.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems 32*, pages 2937–2947, 2019.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685, 2019a.
- Simon S. Du, Xiyu Zhai, Barnabas Póczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b.
- Daniel Gissin, Shai Shalev-Shwartz, and Amit Daniely. The implicit bias of depth: How incremental learning drives generalization. In *International Conference on Learning Representations*, 2020.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems 30*, pages 6151–6159, 2017.
- Suriya Gunasekar, Jason D. Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, 2018a.
- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In *Advances in Neural Information Processing Systems 31*, pages 9461–9471, 2018b.
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems 31*, pages 8571–8580, 2018.
- Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*, 2019a.
- Ziwei Ji and Matus Jan Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations, ICLR 2019*, 2019b.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.

- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Proceedings of the 31st Conference On Learning Theory*, pages 2–47, 2018.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020.
- Mor Shpigel Nacson, Suriya Gunasekar, Jason Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In *International Conference on Machine Learning*, pages 4683–4692, 2019a.
- Mor Shpigel Nacson, Jason D. Lee, Suriya Gunasekar, Pedro Henrique Pamplona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, pages 3420–3428, 2019b. URL <http://proceedings.mlr.press/v89/nacson19b.html>.
- Behnam Neyshabur. Implicit regularization in deep learning. *arXiv preprint arXiv:1709.01953*, 2017.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations, workshop track*, 2015.
- Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A function space view of bounded norm infinite width relu nets: The multivariate case. *arXiv preprint arXiv:1910.01635*, 2019.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *arXiv preprint arXiv:2005.06398*, 2020.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In *Conference on Learning Theory*, pages 2667–2690, 2019.
- Ohad Shamir. Exponential convergence time of gradient descent for one-dimensional deep linear neural networks. In *Proceedings of the Thirty-Second Conference on Learning Theory*, pages 2691–2713, 2019.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 2018.
- Matus Telgarsky. Margins, shrinkage and boosting. In *International Conference on International Conference on Machine Learning*, pages II–307. JMLR. org, 2013.
- Blake Woodworth, Suriya Gunasekar, Pedro Savarese, Edward Moroshko, Itay Golan, Jason Lee, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, 2020.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Stochastic gradient descent optimizes over-parameterized deep relu networks. *Machine Learning Journal*, 2019.

Appendix

A Outline

This appendix is organized as follows: In Section B we provide preliminaries and notations used in the proofs. In Section C we prove auxiliary lemmas that characterize the dynamics of $\mathbf{w}(t)$ and bound the norm of $\mathbf{w}(t)$. In Section D we prove the loss bound in Lemma 3. In Section E we prove that Condition 5 holds for the linearized model. In the proofs we distinguish between the case $D = 2$ and $D > 2$ since the dynamics is different. In Section F we prove the results for $D = 2$ and in Section G we prove the results for $D > 2$. Finally, in Section H we provide additional simulation results and implementation details.

B Preliminaries and Notations for Proofs

To simplify notation in the proofs, without loss of generality we assume that $\forall n : y_n = 1$, as equivalently we can re-define $y_n \mathbf{x}_n$ as \mathbf{x}_n .

Path parametrization: In the proofs we parameterize the optimization path in terms of $\tilde{\gamma}$. Recall that $\tilde{\gamma} = -\log \epsilon$ and $\tilde{\gamma}(t)$ is monotonically increasing along the gradient flow path starting from $\tilde{\gamma}(0) = -\log \epsilon(0) = 0$. Accordingly, the stopping criteria is $\tilde{\gamma}(\alpha) = \tilde{\gamma}(T_\alpha) = -\log \epsilon(\alpha)$. We also overload notation and denote $\mathbf{w}(\tilde{\gamma}') = \mathbf{w}(t_{\tilde{\gamma}'})$ and $\gamma(\tilde{\gamma}') = \gamma(t_{\tilde{\gamma}'})$ where $t_{\tilde{\gamma}'}$ is the unique t such that $\tilde{\gamma}(t) = \tilde{\gamma}'$. Moreover, in this appendix we restate the conditions and theorems in terms of $\tilde{\gamma}$ rather than ϵ .

Notation: We use the following notations:

- $X = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{R}^{d \times N}$ denotes the data matrix.
- $\tilde{X} = [\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N] \in \mathbb{R}^{2d \times N}$ denotes the augmented data matrix where $\tilde{\mathbf{x}}_n = \begin{bmatrix} \mathbf{x}_n \\ -\mathbf{x}_n \end{bmatrix} \in \mathbb{R}^{2d}$.
- $\bar{x}_i = \sum_{n=1}^N |x_{n,i}|$ where $x_{n,i}$ is the coordinate i of \mathbf{x}_n . Also $\bar{x} = \max_i (\bar{x}_i)$.
- $x_{\max} = \max_n \|\mathbf{x}_n\|_2$.
- For some vector \mathbf{z} we denote by $\text{diag}(\mathbf{z})$ the diagonal matrix with diagonal \mathbf{z} , and $[\mathbf{z}]_i$ is the i coordinate.
- The ℓ_2 margin at time t is $\gamma_2(t) = \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(t))}{\|\mathbf{w}(t)\|_2}$. Recall that $\gamma_2 = \max_{\|\mathbf{w}\|_2=1} \min_n \mathbf{x}_n^\top \mathbf{w}$.
- ∂° denotes the local sub-differential (Clarke's sub-differential) operator defined as
$$\partial^\circ h(\mathbf{z}) = \text{conv} \{ \mathbf{v} : \exists \mathbf{z}_k \text{ s.t. } \mathbf{z}_k \rightarrow \mathbf{z} \text{ and } \nabla h(\mathbf{z}_k) \rightarrow \mathbf{v} \}.$$
Specifically, for $h(\mathbf{z}) = \|\mathbf{z}\|_1$:
$$\partial^\circ \|\mathbf{z}\|_1 = \{ \mathbf{v} \in \mathbb{R}^d : \forall i = 1, \dots, d : -1 \leq v_i \leq 1 \text{ and } z_i \neq 0 \Rightarrow v_i = \text{sign}(z_i) \}.$$
- We denote $\mathbf{r}(t) = \frac{1}{N} \exp(-X^\top \mathbf{w}(t))$. Note that $\|\mathbf{r}(t)\|_1 = \mathcal{L}(t) = \exp(-\tilde{\gamma}(t))$.
- We denote $A(t) = \text{diag} \left(4\sqrt{\mathbf{w}^2(t) + 4\alpha^4 \mathbf{1}} \right)$. This matrix is used in the proofs for $D = 2$.
- For $D > 2$ let:

$$h_D(z) = (1-z)^{-\frac{D}{D-2}} - (1+z)^{-\frac{D}{D-2}}, \quad z \in (-1, 1). \quad (7)$$

Note that $h_D(z)$ is monotonically increasing, where $h_D(z) \xrightarrow{z \rightarrow -1} -\infty$ and $h_D(z) \xrightarrow{z \rightarrow 1} \infty$, and thus the inverse h_D^{-1} is well defined, $h_D^{-1} : (-\infty, \infty) \rightarrow (-1, 1)$. In addition, it is easy to verify that for $z \in (-1, 1)$

$$h_D'(z) \doteq \frac{dh_D(z)}{dz} \geq \frac{2D}{D-2} \quad (8)$$

and

$$\lim_{z \rightarrow 0} \frac{h_D(z)}{z} = \frac{2D}{D-2}. \quad (9)$$

- We denote $A_D(t) = \text{diag} \left(\alpha^{2D-2} D(D-2) h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) \right)$. This matrix is used in the proofs for $D > 2$.

Useful inequalities: From the definitions of $\mathcal{L}(t)$ and $\tilde{\gamma}(t)$ we have that

$$\mathcal{L}(t) = \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \mathbf{w}(t)) = \exp(-\tilde{\gamma}(t))$$

and thus

$$\begin{aligned} \frac{1}{N} \exp(-\gamma(t)) &\leq \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \mathbf{w}(t)) \leq \frac{1}{N} N \exp(-\gamma(t)) \\ &\Rightarrow \frac{1}{N} \exp(-\gamma(t)) \leq \exp(-\tilde{\gamma}(t)) \leq \exp(-\gamma(t)) \\ &\Rightarrow \gamma(t) \leq \tilde{\gamma}(t) \leq \gamma(t) + \log(N). \end{aligned} \quad (10)$$

From eq. (10) we have that $\lim_{t \rightarrow \infty} \frac{\tilde{\gamma}(t)}{\gamma(t)} = 1$ and thus

$$\lim_{\alpha \rightarrow \infty} \frac{\tilde{\gamma}(T_\alpha)}{\gamma(T_\alpha)} = 1. \quad (11)$$

In addition, using $\mathbf{x}_n^\top \mathbf{w}(t) \leq x_{\max} \|\mathbf{w}(t)\|_2$ we derive a lower bound on $\|\mathbf{w}(t)\|_2$ as following:

$$\begin{aligned} \mathcal{L}(t) &= \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \mathbf{w}(t)) \\ &\geq \frac{1}{N} N \exp(-x_{\max} \|\mathbf{w}(t)\|_2) \\ &= \exp(-x_{\max} \|\mathbf{w}(t)\|_2) \end{aligned}$$

and thus

$$\begin{aligned} x_{\max} \|\mathbf{w}(t)\|_2 &\geq \log \frac{1}{\mathcal{L}(t)} = \tilde{\gamma}(t) \\ &\Rightarrow \|\mathbf{w}(t)\|_2 \geq \frac{\tilde{\gamma}(t)}{x_{\max}}. \end{aligned} \quad (12)$$

Conditions: We restate Condition 5 in terms of $\tilde{\gamma}$, i.e. we substitute $\epsilon = \exp(-\tilde{\gamma})$. We consider two cases:

Condition 8. For all $k \in [N]$ such that $\mathbf{x}_k^\top \hat{\mathbf{w}} > 1$, and large enough α , there exists $\tilde{\gamma}^*(\alpha) = o(\alpha^D)$ and $\rho_0 > 1$ such that $\forall \tilde{\gamma} \in [\tilde{\gamma}^*(\alpha), \tilde{\gamma}(\alpha)] : \frac{\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})}{\gamma(\tilde{\gamma})} \geq \rho_0$.

Condition 9. For all $k \in [N]$ such that $\mathbf{x}_k^\top \hat{\mathbf{w}} > 1$, and large enough α , there exists $\tilde{\gamma}^*(\alpha) = o\left(\alpha^2 \log \frac{\tilde{\gamma}(\alpha)}{\alpha^2}\right)$ and $\rho_0 > 1$ such that $\forall \tilde{\gamma} \in [\tilde{\gamma}^*(\alpha), \tilde{\gamma}(\alpha)] : \frac{\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})}{\gamma(\tilde{\gamma})} \geq \rho_0$.

We prove the intermediate regime for $D \geq 2$ and the rich regime for $D > 2$ under Condition 8. To prove the rich regime for $D = 2$ the weaker Condition 9 will suffice.

C Auxiliary lemmas

C.1 The case $D = 2$

Lemma 10. For $D = 2$ and all t ,

$$\mathbf{w}(t) = 2\alpha^2 \sinh \left(4X \int_0^t \mathbf{r}(s) ds \right) \quad (13)$$

and

$$\frac{d\mathbf{w}(t)}{dt} = \frac{4}{N} \sqrt{\mathbf{w}^2(t) + 4\alpha^4 \mathbf{1}} \circ X \exp(-X^\top \mathbf{w}(t)) = A(t) X \mathbf{r}(t) \quad (14)$$

where $A(t) = \text{diag}\left(4\sqrt{\mathbf{w}^2(t) + 4\alpha^4 \mathbf{1}}\right)$.

Proof. The gradient flow dynamics in the parameters space is given by

$$\dot{\mathbf{u}}(t) = -\nabla_{\mathbf{u}} \mathcal{L}(\mathbf{u}(t)) = \frac{2}{N} \mathbf{u}(t) \circ \tilde{X} \exp(-\tilde{X}^\top \mathbf{u}^2(t)). \quad (15)$$

It is easy to verify that the solution to eq. (15) can be written as

$$\mathbf{u}(t) = \mathbf{u}(0) \circ \exp\left(\frac{2}{N} \tilde{X} \int_0^t \exp(-\tilde{X}^\top \mathbf{u}^2(s)) ds\right) = \alpha \exp\left(\frac{2}{N} \tilde{X} \int_0^t \exp(-\tilde{X}^\top \mathbf{u}^2(s)) ds\right). \quad (16)$$

From (16) and $\mathbf{w} = \mathbf{u}_+^2 - \mathbf{u}_-^2$ we get eq. (13). Taking the derivative of eq. (13) we have

$$\dot{\mathbf{w}}(t) = \frac{8}{N} \alpha^2 \cosh\left(4X \int_0^t \mathbf{r}(s) ds\right) \circ X \exp(-X^\top \mathbf{w}(t)). \quad (17)$$

By combining eqs. (13) and (17) we get

$$\dot{\mathbf{w}}(t) = \frac{8}{N} \alpha^2 \cosh\left(\text{arcsinh}\left(\frac{\mathbf{w}(t)}{2\alpha^2}\right)\right) \circ X \exp(-X^\top \mathbf{w}(t)).$$

Since $\cosh(\text{arcsinh}(x)) = \sqrt{x^2 + 1}$ we get eq. (14). \square

Lemma 11. For $D = 2$ and all t ,

$$\|\mathbf{w}(t)\|_\infty \leq 2\alpha^2 \sinh\left(\frac{\bar{x}}{2\gamma_2^2 \alpha^2} \tilde{\gamma}(t)\right).$$

Proof. Note that

$$\begin{aligned} \frac{d\mathcal{L}(t)}{dt} &= (\nabla_{\mathbf{w}} \mathcal{L}(t))^\top \frac{d\mathbf{w}(t)}{dt} = -(X\mathbf{r}(t))^\top A(t) X\mathbf{r}(t) \\ &\Rightarrow \frac{d\tilde{\gamma}(t)}{dt} = -\frac{1}{\mathcal{L}(t)} \frac{d\mathcal{L}(t)}{dt} = \frac{(X\mathbf{r}(t))^\top A(t) X\mathbf{r}(t)}{\|\mathbf{r}(t)\|_1}. \end{aligned}$$

From $A_{i,i}(t) \geq 8\alpha^2$ we have

$$\frac{d\tilde{\gamma}(t)}{dt} \geq \frac{8\alpha^2 \|X\mathbf{r}(t)\|_2^2}{\|\mathbf{r}(t)\|_1}. \quad (18)$$

From Lemma 2 of Nacson et al. [2019b] we have that

$$\|X\mathbf{r}(t)\|_2 \geq \gamma_2 \|\mathbf{r}(t)\|_1. \quad (19)$$

Combining eqs. (18) and (19) we get

$$\frac{d\tilde{\gamma}(t)}{dt} \geq 8\alpha^2 \gamma_2^2 \|\mathbf{r}(t)\|_1 = 8\alpha^2 \gamma_2^2 \exp(-\tilde{\gamma}(t)). \quad (20)$$

We employ the dynamics equation $\dot{\mathbf{w}}(t) = \frac{4}{N} \sqrt{\mathbf{w}^2(t) + 4\alpha^4 \mathbf{1}} \circ X \exp(-X^\top \mathbf{w}(t))$ and change variables $t \rightarrow \tilde{\gamma}(t)$. Using eq. (20) we get that

$$\left| \frac{dw_i(\tilde{\gamma})}{d\tilde{\gamma}} \right| = \left| \frac{dw_i(t)}{dt} \frac{dt}{d\tilde{\gamma}} \right| \leq \left(\sqrt{w_i^2(\tilde{\gamma}) + 4\alpha^4} \right) |[X \exp(-X^\top \mathbf{w}(\tilde{\gamma}))]_i| \frac{1}{2N\alpha^2 \gamma_2^2 \exp(-\tilde{\gamma})}.$$

Using $\exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})) \leq N \exp(-\tilde{\gamma})$ which follows from eq. (10) we get

$$\left| \frac{dw_i(\tilde{\gamma})}{d\tilde{\gamma}} \right| \leq \frac{\bar{x}_i}{2\alpha^2 \gamma_2^2} \sqrt{w_i^2(\tilde{\gamma}) + 4\alpha^4}$$

and by the Grönwall's inequality we get the desired bound

$$|w_i(\tilde{\gamma})| \leq 2\alpha^2 \sinh\left(\frac{\bar{x}_i}{2\gamma_2^2 \alpha^2} \tilde{\gamma}\right) \leq 2\alpha^2 \sinh\left(\frac{\bar{x}}{2\gamma_2^2 \alpha^2} \tilde{\gamma}\right).$$

\square

C.2 The case $D > 2$

Lemma 12. For $D > 2$ and all t ,

$$\mathbf{w}(t) = \alpha^D h_D \left(\alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds \right)$$

and

$$\frac{d\mathbf{w}(t)}{dt} = A_D(t) X \mathbf{r}(t)$$

where $A_D(t) = \text{diag} \left(\alpha^{2D-2} D(D-2) h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) \right)$.

Proof. The gradient flow dynamics in the parameters space is given by

$$\dot{\mathbf{u}}(t) = -\nabla_{\mathbf{u}} \mathcal{L}(\mathbf{u}(t)) = \frac{D}{N} \mathbf{u}^{D-1}(t) \circ \tilde{X} \exp \left(-\tilde{X}^\top \mathbf{u}^D(t) \right). \quad (21)$$

It is easy to verify that the solution to eq. (21) is

$$\begin{aligned} \mathbf{u}(t) &= \left(\mathbf{u}^{2-D}(0) - \frac{D(D-2)}{N} \tilde{X} \int_0^t \exp \left(-\tilde{X}^\top \mathbf{u}^D(s) \right) ds \right)^{-\frac{1}{D-2}} \\ &= \left(\alpha^{2-D} \mathbf{1} - \frac{D(D-2)}{N} \tilde{X} \int_0^t \exp \left(-\tilde{X}^\top \mathbf{u}^D(s) \right) ds \right)^{-\frac{1}{D-2}} \\ &= \alpha \left(\mathbf{1} - \frac{\alpha^{D-2} D(D-2)}{N} \tilde{X} \int_0^t \exp \left(-\tilde{X}^\top \mathbf{u}^D(s) \right) ds \right)^{-\frac{1}{D-2}}. \end{aligned} \quad (22)$$

From eq. (22) and $\mathbf{w} = \mathbf{u}_+^D - \mathbf{u}_-^D$ we get

$$\begin{aligned} \mathbf{w}(t) &= \alpha^D \left[\left(\mathbf{1} - \alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds \right)^{-\frac{D}{D-2}} \right. \\ &\quad \left. - \left(\mathbf{1} + \alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds \right)^{-\frac{D}{D-2}} \right]. \end{aligned} \quad (23)$$

As $u_i(t) \geq 0$ for all i (because $u_i(0) = \alpha > 0$; the gradient flow dynamics are continuous; and $u_i(t) = 0 \Rightarrow \dot{u}_i(t) = 0$) we get from eq. (22) that

$$-1 \leq \frac{\alpha^{D-2} D(D-2)}{N} X \int_0^t \exp \left(-\tilde{X}^\top \mathbf{u}^D(s) \right) ds \leq 1. \quad (24)$$

Therefore we can write eq. (23) as

$$\begin{aligned} \mathbf{w}(t) &= \alpha^D h_D \left(\alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds \right) \\ &\Rightarrow \alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds = h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right). \end{aligned} \quad (25)$$

Taking the derivative of eq. (25) we get

$$\begin{aligned} \dot{\mathbf{w}}(t) &= \alpha^D h_D' \left(\alpha^{D-2} D(D-2) X \int_0^t \mathbf{r}(s) ds \right) \circ \left(\alpha^{D-2} D(D-2) X \mathbf{r}(t) \right) \\ &= \alpha^{2D-2} D(D-2) h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) \circ (X \mathbf{r}(t)). \end{aligned} \quad (26)$$

□

Lemma 13. For $D > 2$ and all t ,

$$\|\mathbf{w}(t)\|_\infty \leq \alpha^D h_D \left(\frac{(D-2) \bar{x}}{2D \gamma_2^2 \alpha^D} \tilde{\gamma}(t) \right).$$

Proof. Note that

$$\begin{aligned} \frac{d\mathcal{L}(t)}{dt} &= (\nabla_{\mathbf{w}} \mathcal{L}(t))^{\top} \frac{d\mathbf{w}(t)}{dt} = - (X\mathbf{r}(t))^{\top} A_D(t) X\mathbf{r}(t) \\ \Rightarrow \frac{d\tilde{\gamma}(t)}{dt} &= -\frac{1}{\mathcal{L}(t)} \frac{d\mathcal{L}(t)}{dt} = \frac{(X\mathbf{r}(t))^{\top} A_D(t) X\mathbf{r}(t)}{\|\mathbf{r}(t)\|_1}. \end{aligned} \quad (27)$$

From eq. (8) we get a lower bound on the entries of $A_D(t)$, $A_D(t) \geq 2D^2\alpha^{2D-2}$. Combining with eq. (27) we get

$$\frac{d\tilde{\gamma}(t)}{dt} \geq \frac{2D^2\alpha^{2D-2} \|X\mathbf{r}(t)\|_2^2}{\|\mathbf{r}(t)\|_1}. \quad (28)$$

From eqs. (28), (19) we get

$$\frac{d\tilde{\gamma}(t)}{dt} \geq 2D^2\alpha^{2D-2}\gamma_2^2 \|\mathbf{r}(t)\|_1 = 2D^2\alpha^{2D-2}\gamma_2^2 \exp(-\tilde{\gamma}(t)). \quad (29)$$

We employ the dynamics equation eq. (26) and change variables $t \rightarrow \tilde{\gamma}(t)$. Using eq. (29) we get that

$$\left| \frac{dw_i(\tilde{\gamma})}{d\tilde{\gamma}} \right| = \left| \frac{dw_i(t)}{dt} \frac{dt}{d\tilde{\gamma}} \right| \leq \alpha^{2D-2} D(D-2) h_D' \left(h_D^{-1} \left(\frac{w_i(\tilde{\gamma})}{\alpha^D} \right) \right) \frac{|[X \exp(-X^{\top} \mathbf{w}(\tilde{\gamma}))]_i|}{2ND^2\alpha^{2D-2}\gamma_2^2 \exp(-\tilde{\gamma})}.$$

Using $\exp(-\mathbf{x}_k^{\top} \mathbf{w}(\tilde{\gamma})) \leq N \exp(-\tilde{\gamma})$ which follows from eq. (10) we get

$$\left| \frac{dw_i(\tilde{\gamma})}{d\tilde{\gamma}} \right| \leq \frac{(D-2)\bar{x}_i}{2D\gamma_2^2} h_D' \left(h_D^{-1} \left(\frac{w_i(\tilde{\gamma})}{\alpha^D} \right) \right)$$

and by the Grönwall's inequality we get the desired bound

$$|w_i(\tilde{\gamma})| \leq \alpha^D h_D \left(\frac{(D-2)\bar{x}_i}{2D\gamma_2^2\alpha^D} \tilde{\gamma} \right) \leq \alpha^D h_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma} \right).$$

□

D Proof of Lemma 3

We prove the loss bound for $D \geq 2$, any fixed α , and $\forall t$:

$$\mathcal{L}(t) \leq \frac{1}{1 + 2D^2\alpha^{2D-2}\gamma_2^2 t}.$$

Proof. We employ the Grönwall's inequality. For $D = 2$ from eq. (20) we get

$$\tilde{\gamma}(t) \geq \log(1 + 8\alpha^2\gamma_2^2 t) \quad (30)$$

and thus

$$\mathcal{L}(t) \leq \frac{1}{1 + 8\alpha^2\gamma_2^2 t}. \quad (31)$$

For $D > 2$ from eq. (29) we get

$$\tilde{\gamma}(t) \geq \log(1 + 2D^2\alpha^{2D-2}\gamma_2^2 t) \quad (32)$$

and thus

$$\mathcal{L}(t) \leq \frac{1}{1 + 2D^2\alpha^{2D-2}\gamma_2^2 t}. \quad (33)$$

Note that by substituting $D = 2$ in eq. (33) we get eq. (31), so (33) is correct for $D \geq 2$. □

E Condition 5 holds for the linearized model

We show that Condition 8, which is equivalent to Condition 5, holds for the linearized model. The linearized model is

$$\bar{f}(\bar{\mathbf{u}}, \mathbf{x}) = f(\mathbf{u}(0), \mathbf{x}) + \nabla_{\mathbf{u}}^{\top} f(\mathbf{u}(0), \mathbf{x}) \bar{\mathbf{u}}.$$

For the diagonal linear network $f(\mathbf{u}, \mathbf{x}) = \mathbf{x}^{\top} (\mathbf{u}_+^D - \mathbf{u}_-^D)$, where $\mathbf{u} = \begin{bmatrix} \mathbf{u}_+ \\ \mathbf{u}_- \end{bmatrix} \in \mathbb{R}^{2d}$. We consider the initialization $\mathbf{u}(0) = \bar{\mathbf{u}}(0) = \alpha \mathbf{1}$, thus $f(\mathbf{u}(0), \mathbf{x}) = 0$, $\nabla_{\mathbf{u}} f(\mathbf{u}(0), \mathbf{x}) = D\alpha^{D-1} \begin{bmatrix} \mathbf{x} \\ -\mathbf{x} \end{bmatrix}$ and

$$\bar{f}(\bar{\mathbf{u}}, \mathbf{x}) = D\alpha^{D-1} \mathbf{x}^{\top} (\bar{\mathbf{u}}_+ - \bar{\mathbf{u}}_-).$$

Let $\bar{\mathbf{w}} = D\alpha^{D-1} (\bar{\mathbf{u}}_+ - \bar{\mathbf{u}}_-)$. Then $\bar{f}(\bar{\mathbf{w}}, \mathbf{x}) = \bar{\mathbf{w}}^{\top} \mathbf{x}$. We consider gradient flow $\frac{d\bar{\mathbf{u}}(t)}{dt} = -\nabla \bar{\mathcal{L}}(\bar{\mathbf{u}}(t))$ where

$$\bar{\mathcal{L}}(\bar{\mathbf{u}}(t)) = \frac{1}{N} \sum_{n=1}^N \exp(-\bar{f}(\bar{\mathbf{u}}(t), \mathbf{x}_n)).$$

Thus

$$\begin{aligned} \frac{d\bar{\mathbf{u}}_+(t)}{dt} &= \frac{1}{N} D\alpha^{D-1} \sum_{n=1}^N \exp(-D\alpha^{D-1} \mathbf{x}_n^{\top} (\bar{\mathbf{u}}_+ - \bar{\mathbf{u}}_-)) \mathbf{x}_n \\ \frac{d\bar{\mathbf{u}}_-(t)}{dt} &= -\frac{1}{N} D\alpha^{D-1} \sum_{n=1}^N \exp(-D\alpha^{D-1} \mathbf{x}_n^{\top} (\bar{\mathbf{u}}_+ - \bar{\mathbf{u}}_-)) \mathbf{x}_n \end{aligned}$$

and

$$\begin{aligned} \frac{d\bar{\mathbf{w}}(t)}{dt} &= D\alpha^{D-1} \left(\frac{d\bar{\mathbf{u}}_+(t)}{dt} - \frac{d\bar{\mathbf{u}}_-(t)}{dt} \right) \\ &= \frac{2}{N} D^2 \alpha^{2D-2} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n. \end{aligned} \quad (34)$$

It follows that

$$\begin{aligned} \frac{d\bar{\mathcal{L}}(t)}{dt} &= (\nabla_{\bar{\mathbf{w}}} \bar{\mathcal{L}}(t))^{\top} \frac{d\bar{\mathbf{w}}(t)}{dt} \\ &= \left(-\frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n \right)^{\top} \left(\frac{2}{N} D^2 \alpha^{2D-2} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n \right) \\ &= -2D^2 \alpha^{2D-2} \left\| \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n \right\|_2^2. \end{aligned}$$

Let $\bar{\gamma}(t) = \log \frac{1}{\bar{\mathcal{L}}(t)}$. Then

$$\frac{d\bar{\gamma}(t)}{dt} = -\frac{1}{\bar{\mathcal{L}}(t)} \frac{d\bar{\mathcal{L}}(t)}{dt} = \frac{1}{\bar{\mathcal{L}}(t)} 2D^2 \alpha^{2D-2} \left\| \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n \right\|_2^2. \quad (35)$$

From Lemma 2 of Nacson et al. [2019b] we know that

$$\left\| \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \mathbf{x}_n \right\|_2^2 \geq \gamma_2^2 \left(\frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^{\top} \bar{\mathbf{w}}(t)) \right)^2 = \gamma_2^2 \bar{\mathcal{L}}^2(t). \quad (36)$$

Combining eqs. (35) and (36) we get

$$\frac{d\bar{\gamma}(t)}{dt} \geq 2D^2 \alpha^{2D-2} \gamma_2^2 \bar{\mathcal{L}}(t) = 2D^2 \alpha^{2D-2} \gamma_2^2 \exp(-\bar{\gamma}(t)). \quad (37)$$

In addition,

$$\left\| \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) \mathbf{x}_n \right\|_2 \leq \frac{1}{N} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) \|\mathbf{x}_n\|_2 \leq x_{\max} \bar{\mathcal{L}}(t). \quad (38)$$

Combining eqs. (35) and (38) we get

$$\frac{d\bar{\gamma}(t)}{dt} \leq 2D^2 \alpha^{2D-2} x_{\max}^2 \bar{\mathcal{L}}(t) = 2D^2 \alpha^{2D-2} x_{\max}^2 \exp(-\bar{\gamma}(t))$$

and by the Grönwall's inequality we get

$$\begin{aligned} \bar{\gamma}(t) &\leq \log(1 + 2D^2 \alpha^{2D-2} x_{\max}^2 t) \\ \Rightarrow t &\geq \frac{\exp(\bar{\gamma}) - 1}{2D^2 \alpha^{2D-2} x_{\max}^2}. \end{aligned} \quad (39)$$

The ℓ_2 max-margin solution is $\mathbf{w}_{\ell_2} = \sum_{n \in S_2} \nu_n \mathbf{x}_n$ where S_2 denotes the set of support vectors of \mathbf{w}_{ℓ_2} . Let $\tilde{\mathbf{w}}$ be a vector that satisfies $\exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}}) = \nu_n$ for $n \in S_2$. Such $\tilde{\mathbf{w}}$ exists for almost all datasets, where the support vectors of \mathbf{w}_{ℓ_2} are associated with positive dual variables ν_n [Soudry et al., 2018]. Let

$$\kappa(t) = \bar{\mathbf{w}}(t) - \log\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right) \mathbf{w}_{\ell_2} - \tilde{\mathbf{w}}. \quad (40)$$

Then

$$\frac{d\kappa(t)}{dt} = \frac{d\bar{\mathbf{w}}(t)}{dt} - \frac{1}{t} \mathbf{w}_{\ell_2}$$

and thus

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|\kappa(t)\|_2^2 &= \left(\frac{d\kappa(t)}{dt} \right)^\top \kappa(t) \\ &= \left(\frac{d\bar{\mathbf{w}}(t)}{dt} - \frac{1}{t} \mathbf{w}_{\ell_2} \right)^\top \kappa(t) \\ &\stackrel{(34)}{=} \frac{2}{N} D^2 \alpha^{2D-2} \sum_{n=1}^N \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) \mathbf{x}_n^\top \kappa(t) - \frac{1}{t} \mathbf{w}_{\ell_2}^\top \kappa(t) \\ &= \left[\frac{2}{N} D^2 \alpha^{2D-2} \sum_{n \in S_2} \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) \mathbf{x}_n^\top \kappa(t) - \frac{1}{t} \mathbf{w}_{\ell_2}^\top \kappa(t) \right] \\ &\quad + \left[\frac{2}{N} D^2 \alpha^{2D-2} \sum_{n \notin S_2} \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) \mathbf{x}_n^\top \kappa(t) \right] \end{aligned} \quad (41)$$

For $n \in S_2$ we have that $\mathbf{x}_n^\top \mathbf{w}_{\ell_2} = 1$, thus

$$\begin{aligned} \exp(-\mathbf{x}_n^\top \bar{\mathbf{w}}(t)) &= \exp\left(-\mathbf{x}_n^\top \left(\log\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right) \mathbf{w}_{\ell_2} + \tilde{\mathbf{w}} + \kappa(t)\right)\right) \\ &= \frac{1}{\frac{2}{N} D^2 \alpha^{2D-2} t} \exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}}) \exp(-\mathbf{x}_n^\top \kappa(t)) \\ &= \frac{1}{\frac{2}{N} D^2 \alpha^{2D-2} t} \nu_n \exp(-\mathbf{x}_n^\top \kappa(t)) \end{aligned}$$

and the first bracketed term in eq. (41) can be written as

$$\begin{aligned} &\frac{2}{N} D^2 \alpha^{2D-2} \sum_{n \in S_2} \frac{1}{\frac{2}{N} D^2 \alpha^{2D-2} t} \nu_n \exp(-\mathbf{x}_n^\top \kappa(t)) \mathbf{x}_n^\top \kappa(t) - \frac{1}{t} \sum_{n \in S_2} \nu_n \mathbf{x}_n^\top \kappa(t) \\ &= \frac{1}{t} \sum_{n \in S_2} [\nu_n (\exp(-\mathbf{x}_n^\top \kappa(t)) - 1) \mathbf{x}_n^\top \kappa(t)] \\ &\leq 0 \end{aligned} \quad (42)$$

since $(e^{-z} - 1)z \leq 0$ for all z .

Let $\theta = \min_{n \notin S_2} (\mathbf{x}_n^\top \mathbf{w}_{\ell_2}) > 1$ and $c_1 = \max_{n \notin S_2} \exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}})$. For $n \notin S_2$ we have that

$$\begin{aligned} \exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}}(t)) &= \exp\left(-\mathbf{x}_n^\top \left(\log\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right) \mathbf{w}_{\ell_2} + \tilde{\mathbf{w}} + \kappa(t)\right)\right) \\ &\leq \frac{1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right)^\theta} \exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}}) \exp(-\mathbf{x}_n^\top \kappa(t)) \\ &\leq \frac{c_1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right)^\theta} \exp(-\mathbf{x}_n^\top \kappa(t)) \end{aligned}$$

and thus the second bracketed term in eq. (41) can be bounded as following

$$\begin{aligned} &\frac{2}{N} D^2 \alpha^{2D-2} \sum_{n \notin S_2} \exp(-\mathbf{x}_n^\top \tilde{\mathbf{w}}(t)) \mathbf{x}_n^\top \kappa(t) \\ &\leq \frac{\frac{2}{N} D^2 \alpha^{2D-2} c_1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right)^\theta} \sum_{n \notin S_2} \exp(-\mathbf{x}_n^\top \kappa(t)) \mathbf{x}_n^\top \kappa(t) \\ &\leq \frac{2 D^2 \alpha^{2D-2} c_1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right)^\theta} \end{aligned} \tag{43}$$

since $e^{-z} z \leq 1$ for all z . Substituting eqs. (42) and (43) in eq. (41) we get

$$\frac{1}{2} \frac{d}{dt} \|\kappa(t)\|_2^2 \leq \frac{2 D^2 \alpha^{2D-2} c_1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} t\right)^\theta}.$$

Using (39) we get

$$\frac{1}{2} \frac{d}{dt} \|\kappa(t)\|_2^2 \leq \frac{2 D^2 \alpha^{2D-2} c_1}{\left(\frac{2}{N} D^2 \alpha^{2D-2} \frac{\exp(\bar{\gamma}(t)) - 1}{2 D^2 \alpha^{2D-2} x_{\max}^2}\right)^\theta} = \frac{2 D^2 \alpha^{2D-2} c_1}{\left(\frac{1}{N} \frac{\exp(\bar{\gamma}(t)) - 1}{x_{\max}^2}\right)^\theta}.$$

We change variables $t \rightarrow \bar{\gamma}$ and get

$$\begin{aligned} \frac{1}{2} \frac{d}{d\bar{\gamma}} \|\kappa(\bar{\gamma})\|_2^2 &= \frac{1}{2} \frac{d}{dt} \|\kappa(\bar{\gamma}(t))\|_2^2 \frac{dt}{d\bar{\gamma}} \\ &\stackrel{(37)}{\leq} \frac{2 D^2 \alpha^{2D-2} c_1}{\left(\frac{1}{N} \frac{\exp(\bar{\gamma}) - 1}{x_{\max}^2}\right)^\theta} \frac{1}{2 D^2 \alpha^{2D-2} \gamma_2^2 \exp(-\bar{\gamma})} \\ &= \frac{c_1 \exp(-(\theta - 1) \bar{\gamma})}{\gamma_2^2 \left(\frac{1}{N} \frac{1 - \exp(-\bar{\gamma})}{x_{\max}^2}\right)^\theta} \\ &\leq C \frac{\exp(-(\theta - 1) \bar{\gamma})}{(1 - \exp(-\bar{\gamma}))^\theta} \end{aligned}$$

where C is a constant. Integrating we have that for all $\bar{\gamma}_0 > 0, \bar{\gamma} > \bar{\gamma}_0$

$$\begin{aligned} \|\kappa(\bar{\gamma})\|_2^2 - \|\kappa(\bar{\gamma}_0)\|_2^2 &\leq C \int_{\bar{\gamma}_0}^{\bar{\gamma}} \frac{\exp(-(\theta - 1) \bar{\gamma}_1)}{(1 - \exp(-\bar{\gamma}_1))^\theta} d\bar{\gamma}_1 \\ &\leq C \int_{\bar{\gamma}_0}^{\infty} \frac{\exp(-(\theta - 1) \bar{\gamma}_1)}{(1 - \exp(-\bar{\gamma}_1))^\theta} d\bar{\gamma}_1 \\ &= C \frac{1}{(\theta - 1) (\exp(\bar{\gamma}_0) - 1)^{\theta-1}} \end{aligned}$$

and thus

$$\|\kappa(\bar{\gamma})\|_2 \leq C' \tag{44}$$

where C' is a constant. Finally, for $k \notin S_2$ we have that $\mathbf{x}_k^\top \mathbf{w}_{\ell_2} \geq \theta > 1$ and thus

$$\begin{aligned}
\frac{\mathbf{x}_k^\top \bar{\mathbf{w}}(\tilde{\gamma})}{\tilde{\gamma}} &\stackrel{(40)}{=} \frac{\mathbf{x}_k^\top \left(\log \left(\frac{2}{N} D^2 \alpha^{2D-2} t \right) \mathbf{w}_{\ell_2} + \tilde{\mathbf{w}} + \kappa(\tilde{\gamma}) \right)}{\tilde{\gamma}} \\
&\stackrel{(39)}{\geq} \frac{\log \left(\frac{2}{N} D^2 \alpha^{2D-2} \frac{\exp(\tilde{\gamma})-1}{2D^2 \alpha^{2D-2} x_{\max}^2} \right) \mathbf{x}_k^\top \mathbf{w}_{\ell_2} + \mathbf{x}_k^\top \tilde{\mathbf{w}} + \mathbf{x}_k^\top \kappa(\tilde{\gamma})}{\tilde{\gamma}} \\
&\stackrel{(44)}{\geq} \frac{\log \left(\frac{\exp(\tilde{\gamma})-1}{N x_{\max}^2} \right) \theta - \log c_1 - x_{\max} C'}{\tilde{\gamma}} \\
&= \frac{\log \left(\frac{\exp(\tilde{\gamma})-1}{N x_{\max}^2} \right)}{\tilde{\gamma}} \theta - \frac{\log c_1 + x_{\max} C'}{\tilde{\gamma}}.
\end{aligned}$$

Note that $\frac{\log \left(\frac{\exp(\tilde{\gamma})-1}{N x_{\max}^2} \right)}{\tilde{\gamma}}$ is monotonically increasing and

$$\frac{\log \left(\frac{\exp(\tilde{\gamma})-1}{N x_{\max}^2} \right)}{\tilde{\gamma}} \xrightarrow{\tilde{\gamma} \rightarrow \infty} 1.$$

Therefore there exists $\tilde{\gamma}^*$ (independent of α !) such that for $\tilde{\gamma} \geq \tilde{\gamma}^*$

$$\frac{\log \left(\frac{\exp(\tilde{\gamma})-1}{N x_{\max}^2} \right)}{\tilde{\gamma}} \geq \frac{3\theta + 1}{4\theta}$$

and

$$\frac{\log c_1 + x_{\max} C'}{\tilde{\gamma}} \leq \frac{\theta - 1}{4}.$$

It follows that for $\tilde{\gamma} \geq \tilde{\gamma}^*$

$$\frac{\mathbf{x}_k^\top \bar{\mathbf{w}}(\tilde{\gamma})}{\tilde{\gamma}} \geq \frac{3\theta + 1}{4\theta} \theta - \frac{\theta - 1}{4} = \frac{\theta + 1}{2} \doteq \rho_0 > 1.$$

From $\tilde{\gamma} \geq \bar{\gamma}$ (where $\bar{\gamma}(t) = \min_n (\mathbf{x}_n^\top \bar{\mathbf{w}}(t))$) we get

$$\frac{\mathbf{x}_k^\top \bar{\mathbf{w}}(\tilde{\gamma})}{\tilde{\gamma}} \geq \frac{\mathbf{x}_k^\top \bar{\mathbf{w}}(\bar{\gamma})}{\bar{\gamma}} \geq \rho_0 > 1$$

for $\tilde{\gamma} \geq \tilde{\gamma}^* = o(\alpha^D)$ since $\tilde{\gamma}^*$ is independent of α .

F Proofs for $D = 2$

F.1 Kernel Regime Proof

Theorem 14 (Theorem 4 for $D = 2$). *For $D = 2$, if $\tilde{\gamma}(\alpha) = o(\alpha^2)$ then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{w}\|_2 \quad \text{s.t. } \forall n : \mathbf{x}_n^\top \mathbf{w} \geq 1.$$

Proof. We show convergence of the ℓ_2 margin $\gamma_2(T_\alpha) = \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(T_\alpha))}{\|\mathbf{w}(T_\alpha)\|_2}$ to the max margin γ_2 when $\alpha \rightarrow \infty$. Note that by definition $\gamma_2(T_\alpha) \leq \gamma_2$. Next we show that $\gamma_2(T_\alpha) \geq \gamma_2$ when $\alpha \rightarrow \infty$. From eq. (10) we have that

$$\gamma_2(t) = \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(t))}{\|\mathbf{w}(t)\|_2} \geq \frac{\tilde{\gamma}(t) - \log(N)}{\|\mathbf{w}(t)\|_2}. \quad (45)$$

In order to lower bound $\gamma_2(t)$ we derive a lower bound on $\tilde{\gamma}(t)$ and an upper bound on $\|\mathbf{w}(t)\|_2$.

Lower bound on $\tilde{\gamma}(t)$: Combining (18) and (19) we get

$$\begin{aligned} \frac{d\tilde{\gamma}(t)}{dt} &\geq 8\alpha^2\gamma_2 \|X\mathbf{r}(t)\|_2 \\ \Rightarrow \tilde{\gamma}(t) &\geq 8\alpha^2\gamma_2 \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \end{aligned} \quad (46)$$

Upper bound on $\|\mathbf{w}(t)\|_2$: We decompose $\dot{\mathbf{w}}(t)$ to two terms:

$$\begin{aligned} \dot{\mathbf{w}}(t) &= 4\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} \circ X\mathbf{r}(t) \\ &= 4\left(\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right) \circ X\mathbf{r}(t) + 8\alpha^2 X\mathbf{r}(t) \\ \Rightarrow \|\dot{\mathbf{w}}(t)\|_2 &\leq 4\left\|\left(\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right) \circ X\mathbf{r}(t)\right\|_2 + 8\alpha^2 \|X\mathbf{r}(t)\|_2 \\ \Rightarrow \|\mathbf{w}(t)\|_2 &\leq 4 \int_0^t \left\|\left(\sqrt{\mathbf{w}^2(\tau) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right) \circ X\mathbf{r}(\tau)\right\|_2 d\tau + 8\alpha^2 \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \end{aligned} \quad (47)$$

Let $v(t) = \left\|\left(\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right) \circ X\mathbf{r}(t)\right\|_2$. Then

$$\begin{aligned} v(t) &\leq \left\|\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right\|_\infty x_{\max} \|\mathbf{r}(t)\|_1 \\ &= \left\|\sqrt{\mathbf{w}^2(t) + 4\alpha^4\mathbf{1}} - 2\alpha^2\mathbf{1}\right\|_\infty x_{\max} \exp(-\tilde{\gamma}(t)). \end{aligned}$$

Using Lemma 11 we get

$$\begin{aligned} v(t) &\leq \left(2\alpha^2 \sqrt{\sinh^2\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}(t)\right) + 1} - 2\alpha^2\right) x_{\max} \exp(-\tilde{\gamma}(t)) \\ &= 2\alpha^2 \left[\cosh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}(t)\right) - 1\right] x_{\max} \exp(-\tilde{\gamma}(t)). \end{aligned}$$

We are interested in bounding $\int_0^t v(\tau) d\tau$. We change variables $t \rightarrow \tilde{\gamma}(t)$ and proceed using (20),

$$\begin{aligned} \int_0^t v(\tau) d\tau &\leq \int_0^{\tilde{\gamma}(t)} 2\alpha^2 \left[\cosh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}\right) - 1\right] x_{\max} \exp(-\tilde{\gamma}) \frac{1}{8\alpha^2\gamma_2^2 \exp(-\tilde{\gamma})} d\tilde{\gamma} \\ &= \frac{x_{\max}}{4\gamma_2^2} \int_0^{\tilde{\gamma}(t)} \left[\cosh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}\right) - 1\right] d\tilde{\gamma} \\ &= \frac{x_{\max}}{4\gamma_2^2} \left[\frac{2\alpha^2\gamma_2^2}{\bar{x}} \sinh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}(t)\right) - \tilde{\gamma}(t)\right]. \end{aligned} \quad (48)$$

Plugging eqs. (48) in (47) we get

$$\|\mathbf{w}(t)\|_2 \leq \frac{x_{\max}}{\gamma_2^2} \left[\frac{2\alpha^2\gamma_2^2}{\bar{x}} \sinh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}(t)\right) - \tilde{\gamma}(t)\right] + 8\alpha^2 \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \quad (49)$$

Putting things together: From eqs. (45) and (12) we have

$$\gamma_2(t) \geq \frac{\tilde{\gamma}(t) - \log(N)}{\|\mathbf{w}(t)\|_2} \geq \frac{\tilde{\gamma}(t)}{\|\mathbf{w}(t)\|_2} - \frac{\log(N)x_{\max}}{\tilde{\gamma}(t)}. \quad (50)$$

Next we set $t = T_\alpha$ and take the limit $\alpha \rightarrow \infty$. Note that $\tilde{\gamma}(T_\alpha) \xrightarrow{\alpha \rightarrow \infty} \infty$ since $\epsilon(T_\alpha) \xrightarrow{\alpha \rightarrow \infty} 0$, and thus the right term in eq. (50) is vanishing. Using eq. (49) we get

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} \frac{1}{\gamma_2(T_\alpha)} &\leq \lim_{\alpha \rightarrow \infty} \frac{\|\mathbf{w}(T_\alpha)\|_2}{\tilde{\gamma}(T_\alpha)} \\ &\leq \lim_{\alpha \rightarrow \infty} \left[\frac{x_{\max}}{\gamma_2^2} \left[\frac{2\alpha^2\gamma_2^2}{\bar{x}} \sinh\left(\frac{\bar{x}}{2\alpha^2\gamma_2^2} \tilde{\gamma}(T_\alpha)\right) - 1\right] + \frac{8\alpha^2}{\tilde{\gamma}(T_\alpha)} \int_0^{T_\alpha} \|X\mathbf{r}(\tau)\|_2 d\tau \right]. \end{aligned}$$

We use $\frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} \xrightarrow{\alpha \rightarrow \infty} 0$, $\lim_{z \rightarrow 0} \frac{\sinh z}{z} = 1$ and eq. (46) to get

$$\lim_{\alpha \rightarrow \infty} \frac{1}{\gamma_2(T_\alpha)} \leq \frac{1}{\gamma_2}.$$

It follows that $\lim_{\alpha \rightarrow \infty} \gamma_2(T_\alpha) = \gamma_2$. \square

F.2 Intermediate Regime Proof

Theorem 15 (Theorem 6 for $D = 2$). *Under Condition 8, for $D = 2$ if $\lim_{\alpha \rightarrow \infty} \frac{\alpha^2}{\tilde{\gamma}(\alpha)} = \mu > 0$, then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} Q_\mu^2(\mathbf{w}) \text{ s.t. } \forall n, \mathbf{x}_n^\top \mathbf{w} \geq 1$$

where $Q_\mu^2(\mathbf{w}) = \sum_{i=1}^d q_2\left(\frac{w_i}{\mu}\right)$ and $q_2(s) = 2 - \sqrt{4 + s^2} + s \cdot \operatorname{arcsinh}\left(\frac{s}{2}\right)$.

Proof. We show that the KKT conditions hold in the limit $\alpha \rightarrow \infty$. The KKT conditions are that there exists $\boldsymbol{\nu} \in \mathbb{R}_{\geq 0}^N$ such that

$$\nabla Q_\mu^2(\hat{\mathbf{w}}) = X\boldsymbol{\nu} \quad (51)$$

$$\forall n : \mathbf{x}_n^\top \hat{\mathbf{w}} \geq 1 \quad (52)$$

$$\forall n : \nu_n (\mathbf{x}_n^\top \hat{\mathbf{w}} - 1) = 0. \quad (53)$$

Primal feasibility (52): The condition (52) follows by definition of $\hat{\mathbf{w}}$,

$$\forall n : \mathbf{x}_n^\top \hat{\mathbf{w}} = \lim_{\alpha \rightarrow \infty} \frac{\mathbf{x}_n^\top \mathbf{w}(T_\alpha)}{\gamma(T_\alpha)} \geq \lim_{\alpha \rightarrow \infty} \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(T_\alpha))}{\gamma(T_\alpha)} = 1. \quad (54)$$

Stationarity condition (51): To show the condition (51) let

$$\boldsymbol{\nu} = \frac{4}{\mu} \limsup_{\alpha \rightarrow \infty} \int_0^{T_\alpha} \mathbf{r}(s) ds \in \mathbb{R}_{\geq 0}^N. \quad (55)$$

We need to show that

$$\nabla Q_\mu(\hat{\mathbf{w}}) = \frac{1}{\mu} \operatorname{arcsinh}\left(\frac{\hat{\mathbf{w}}}{2\mu}\right) = X\boldsymbol{\nu}.$$

Indeed from eqs. (13) and (11) we have

$$\begin{aligned} \hat{\mathbf{w}} &= \lim_{\alpha \rightarrow \infty} \frac{2\alpha^2 \sinh\left(4X \int_0^{T_\alpha} \mathbf{r}(s) ds\right)}{\gamma(T_\alpha)} \\ &= 2 \lim_{\alpha \rightarrow \infty} \frac{\tilde{\gamma}(T_\alpha)}{\gamma(T_\alpha)} \lim_{\alpha \rightarrow \infty} \frac{\alpha^2}{\tilde{\gamma}(T_\alpha)} \limsup_{\alpha \rightarrow \infty} \sinh\left(4X \int_0^{T_\alpha} \mathbf{r}(s) ds\right) \\ &= 2\mu \sinh\left[\mu X \left(\frac{4}{\mu} \limsup_{\alpha \rightarrow \infty} \int_0^{T_\alpha} \mathbf{r}(s) ds\right)\right] \\ &= 2\mu \sinh(\mu X \boldsymbol{\nu}) \end{aligned}$$

and thus $\frac{1}{\mu} \operatorname{arcsinh}\left(\frac{\hat{\mathbf{w}}}{2\mu}\right) = X\boldsymbol{\nu}$, as desired.

Complementary slackness (53): To show the condition (53) let $k \in [N]$ such that

$$\mathbf{x}_k^\top \hat{\mathbf{w}} > 1. \quad (56)$$

We need to show that $\nu_k = 0$. We change variables $t \rightarrow \tilde{\gamma}(t)$ and using eq. (20) we get

$$\int_0^{T_\alpha} \exp(-\mathbf{x}_k^\top \mathbf{w}(s)) ds \leq \frac{1}{8\alpha^2 \gamma_2^2} \int_0^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma}. \quad (57)$$

From Condition 8 we know that there exists $\tilde{\gamma}^*(\alpha) = o(\alpha^2)$ and $\rho_0 > 1$ such that for large enough α and $\tilde{\gamma} \in [\tilde{\gamma}^*(\alpha), \tilde{\gamma}(\alpha)]$, $\frac{\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})}{\gamma(\tilde{\gamma})} \geq \rho_0$. Let

$$\tilde{\gamma}_1^*(\alpha) = \max\left(\frac{2\rho_0 \log N}{\rho_0 - 1}, \tilde{\gamma}^*(\alpha)\right) = o(\alpha^2)$$

and $\rho_1 = \frac{\rho_0 + 1}{2} > 1$. Then for large enough α and $\tilde{\gamma} \in [\tilde{\gamma}_1^*(\alpha), \tilde{\gamma}(\alpha)]$, using $\tilde{\gamma} \leq \gamma + \log N$ we get

$$\begin{aligned} \frac{\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})}{\tilde{\gamma}} &= \frac{\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma})}{\gamma} \frac{\gamma}{\tilde{\gamma}} \\ &\geq \rho_0 \frac{\tilde{\gamma} - \log N}{\tilde{\gamma}} \\ &= \rho_0 - \rho_0 \frac{\log N}{\tilde{\gamma}} \\ &\geq \rho_0 - \rho_0 \frac{\log N}{\frac{2\rho_0 \log N}{\rho_0 - 1}} \\ &= \rho_1. \end{aligned}$$

Next we decompose the RHS of eq. (57) as following

$$\begin{aligned} \frac{1}{8\alpha^2\gamma_2^2} \int_0^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} &= \frac{1}{8\alpha^2\gamma_2^2} \int_0^{\tilde{\gamma}_1^*(\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &\quad + \frac{1}{8\alpha^2\gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &= (I) + (II) \end{aligned} \quad (58)$$

From eq. (10) we have that $\exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) \leq N$ and thus

$$(I) \leq \frac{1}{8\alpha^2\gamma_2^2} \int_0^{\tilde{\gamma}_1^*(\alpha)} N d\tilde{\gamma} = \frac{N\tilde{\gamma}_1^*(\alpha)}{8\alpha^2\gamma_2^2} \xrightarrow{\alpha \rightarrow \infty} 0 \quad (59)$$

since $\tilde{\gamma}_1^*(\alpha) = o(\alpha^2)$. For the second term in eq. (58) we have for large enough α ,

$$\begin{aligned} (II) &= \frac{1}{8\alpha^2\gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp\left[-\left(\mathbf{x}_k^\top \frac{\mathbf{w}(\tilde{\gamma})}{\tilde{\gamma}} - 1\right) \tilde{\gamma}\right] d\tilde{\gamma} \\ &\leq \frac{1}{8\alpha^2\gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp[-(\rho_1 - 1) \tilde{\gamma}] d\tilde{\gamma} \\ &\leq \frac{1}{8\alpha^2\gamma_2^2} \int_0^\infty \exp[-(\rho_1 - 1) \tilde{\gamma}] d\tilde{\gamma} \\ &= \frac{1}{8\alpha^2\gamma_2^2(\rho_1 - 1)} \xrightarrow{\alpha \rightarrow \infty} 0. \end{aligned} \quad (60)$$

By substituting eqs. (59) and (60) in eq. (58) we get that $\nu_k = 0$. \square

F.3 Rich Regime Proof

Theorem 16 (Theorem 7 for $D = 2$). *Under Condition 9, for $D = 2$ if $\tilde{\gamma}(\alpha) = \omega(\alpha^2)$ then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{w}\|_1 \quad \text{s.t. } \forall n : \mathbf{x}_n^\top \mathbf{w} \geq 1. \quad (61)$$

Proof. We show that the KKT conditions for the ℓ_1 max-margin problem (61) hold in the limit $\alpha \rightarrow \infty$. The KKT conditions are that there exists $\boldsymbol{\nu}^{(\ell_1)} \in \mathbb{R}_{\geq 0}^N$ such that

$$X\boldsymbol{\nu}^{(\ell_1)} \in \partial^\circ \|\hat{\mathbf{w}}\|_1 \quad (62)$$

$$\forall n : \mathbf{x}_n^\top \hat{\mathbf{w}} \geq 1 \quad (63)$$

$$\forall n : \nu_n^{(\ell_1)} (\mathbf{x}_n^\top \hat{\mathbf{w}} - 1) = 0. \quad (64)$$

To this end let

$$\boldsymbol{\nu}^{(\ell_1)} = \limsup_{\alpha \rightarrow \infty} \frac{4}{\log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{T_\alpha} \mathbf{r}(s) ds \in \mathbb{R}_{\geq 0}^N. \quad (65)$$

The proof for the primal feasibility condition (63) appears in eq. (54).

Stationarity condition (62):

$$\begin{aligned} \hat{\mathbf{w}} &= \lim_{\alpha \rightarrow \infty} \frac{2\alpha^2 \sinh \left(4X \int_0^{T_\alpha} \mathbf{r}(s) ds \right)}{\gamma(T_\alpha)} \\ &\stackrel{(11)}{=} \lim_{\alpha \rightarrow \infty} \frac{2\alpha^2 \sinh \left(\log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} \frac{4X}{\log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{T_\alpha} \mathbf{r}(s) ds \right)}{\tilde{\gamma}(T_\alpha)} \\ &= \lim_{\alpha \rightarrow \infty} \frac{2 \sinh \left(\log \left(\frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} \right)^{\frac{4X}{\log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}}} \int_0^{T_\alpha} \mathbf{r}(s) ds \right)}{\frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \\ &= \lim_{\alpha \rightarrow \infty} \frac{2 \sinh \left(\log(g(\alpha))^{z(\alpha)} \right)}{g(\alpha)}, \end{aligned} \quad (66)$$

where we defined

$$\begin{aligned} g(\alpha) &= \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} \in \mathbb{R} \\ z(\alpha) &= \frac{4X}{\log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{T_\alpha} \mathbf{r}(s) ds \in \mathbb{R}^d. \end{aligned}$$

Note that from $\lim_{\alpha \rightarrow \infty} \frac{\alpha^2}{\tilde{\gamma}(T_\alpha)} = 0$ we have $\lim_{\alpha \rightarrow \infty} g(\alpha) = \infty$ and from (65) we get $\limsup_{\alpha \rightarrow \infty} z(\alpha) = X\boldsymbol{\nu}^{(\ell_1)}$. In addition, for some $f > 0$ and $a \in \mathbb{R}$:

$$\frac{2 \sinh(\log f^a)}{f} = \frac{f^a - \frac{1}{f^a}}{f} = f^{a-1} - \frac{1}{f^{a+1}}.$$

Therefore in (66) we have,

$$\hat{\mathbf{w}} = \lim_{\alpha \rightarrow \infty} \left(g(\alpha)^{z(\alpha)-1} - \frac{1}{g(\alpha)^{z(\alpha)+1}} \right).$$

Next, it is easy to verify that for all $i = 1, \dots, d$:

$$\begin{aligned} \hat{w}_i > 0 &\Rightarrow \limsup_{\alpha \rightarrow \infty} z_i(\alpha) = 1 \\ \hat{w}_i < 0 &\Rightarrow \limsup_{\alpha \rightarrow \infty} z_i(\alpha) = -1 \\ \hat{w}_i = 0 &\Rightarrow -1 \leq \limsup_{\alpha \rightarrow \infty} z_i(\alpha) \leq 1 \end{aligned}$$

and so $X\boldsymbol{\nu}^{(\ell_1)} \in \partial^\circ \|\hat{\mathbf{w}}\|_1$.

Complementary slackness (64): We perform similar steps to the proof of the intermediate regime in Appendix F.2. We change variables $t \rightarrow \tilde{\gamma}(t)$ and use the weaker Condition 9, where we replace $\tilde{\gamma}^*(\alpha)$ with $\tilde{\gamma}_1^*(\alpha)$ and ρ_0 with ρ_1 (see the proof of the intermediate regime in Appendix F.2). We get that

$$\begin{aligned} \frac{4}{N \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{T_\alpha} \exp(-\mathbf{x}_k^\top \mathbf{w}(s)) ds &\leq \frac{1}{2N\alpha^2\gamma_2^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &\leq \frac{1}{2N\alpha^2\gamma_2^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_0^{\tilde{\gamma}_1^*(\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &\quad + \frac{1}{2N\alpha^2\gamma_2^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &= (I) + (II) \end{aligned} \quad (67)$$

Using $\tilde{\gamma}_1^*(\alpha) = o\left(\alpha^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}\right)$ we bound the first term similarly to eq. (59):

$$(I) \leq \frac{\tilde{\gamma}_1^*(\alpha)}{2N\alpha^2\gamma_2^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2}} \xrightarrow{\alpha \rightarrow \infty} 0. \quad (68)$$

The second term is bounded similarly to eq. (60):

$$(II) \leq \frac{1}{2N\alpha^2\gamma_2^2 \log \frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} (\rho_1 - 1)} \xrightarrow{\alpha \rightarrow \infty} 0. \quad (69)$$

By substituting eqs. (68) and (69) in eq. (67) we get that $\nu_k^{(\ell_1)} = 0$. \square

G Proofs for $D > 2$

G.1 Kernel Regime Proof

Theorem 17 (Theorem 4 for $D > 2$). *For $D > 2$, if $\tilde{\gamma}(\alpha) = o(\alpha^D)$ then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{w}\|_2 \quad \text{s.t. } \forall n : \mathbf{x}_n^\top \mathbf{w} \geq 1.$$

The proof is similar in spirit to the proof for the case $D = 2$ (see Appendix F.1).

Proof. We show convergence of the ℓ_2 margin $\gamma_2(T_\alpha) = \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(T_\alpha))}{\|\mathbf{w}(T_\alpha)\|_2}$ to the max margin γ_2 as $\alpha \rightarrow \infty$. We have that

$$\gamma_2(t) = \frac{\min_n (\mathbf{x}_n^\top \mathbf{w}(t))}{\|\mathbf{w}(t)\|_2} \geq \frac{\tilde{\gamma}(t) - \log(N)}{\|\mathbf{w}(t)\|_2}. \quad (70)$$

Lower bound on $\tilde{\gamma}(t)$: Combining eqs. (28) and (19) we get

$$\begin{aligned} \frac{d\tilde{\gamma}(t)}{dt} &\geq 2D^2\alpha^{2D-2}\gamma_2\|X\mathbf{r}(t)\|_2 \\ \Rightarrow \tilde{\gamma}(t) &\geq 2D^2\alpha^{2D-2}\gamma_2 \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \end{aligned} \quad (71)$$

Upper bound on $\|\mathbf{w}(t)\|_2$: We decompose $\dot{\mathbf{w}}(t)$ to two terms:

$$\begin{aligned} \dot{\mathbf{w}}(t) &= \alpha^{2D-2}D(D-2)h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) \circ X\mathbf{r}(t) \\ &= \alpha^{2D-2}D^2 \left[\frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right] \circ X\mathbf{r}(t) + 2D^2\alpha^{2D-2}X\mathbf{r}(t) \\ \Rightarrow \|\dot{\mathbf{w}}(t)\|_2 &\leq \alpha^{2D-2}D^2 \left\| \left[\frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right] \circ X\mathbf{r}(t) \right\|_2 + 2D^2\alpha^{2D-2}\|X\mathbf{r}(t)\|_2 \\ \Rightarrow \|\mathbf{w}(t)\|_2 &\leq \alpha^{2D-2}D^2 \int_0^t \left\| \left[\frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(\tau)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right] \circ X\mathbf{r}(\tau) \right\|_2 d\tau \\ &\quad + 2D^2\alpha^{2D-2} \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \end{aligned} \quad (72)$$

Let $v(t) = \left\| \left[\frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right] \circ X\mathbf{r}(t) \right\|_2$. Then

$$\begin{aligned} v(t) &\leq \left\| \frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right\|_\infty x_{\max} \|\mathbf{r}(t)\|_1 \\ &= \left\| \frac{D-2}{D}h_D' \left(h_D^{-1} \left(\frac{\mathbf{w}(t)}{\alpha^D} \right) \right) - 2 \cdot \mathbf{1} \right\|_\infty x_{\max} \exp(-\tilde{\gamma}(t)). \end{aligned}$$

Using Lemma 13 we get

$$v(t) \leq \left[\frac{D-2}{D} h'_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma}(t) \right) - 2 \right] x_{\max} \exp(-\tilde{\gamma}(t)).$$

We are interested in bounding $\int_0^t v(\tau) d\tau$. We change variables $t \rightarrow \tilde{\gamma}(t)$ and proceed using eq. (29),

$$\begin{aligned} \int_0^t v(\tau) d\tau &\leq \int_0^{\tilde{\gamma}(t)} \left[\frac{D-2}{D} h'_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma} \right) - 2 \right] \frac{x_{\max} \exp(-\tilde{\gamma})}{2\alpha^{2D-2} D^2 \gamma_2^2 \exp(-\tilde{\gamma})} d\tilde{\gamma} \\ &= \frac{x_{\max}}{2\alpha^{2D-2} D^2 \gamma_2^2} \int_0^{\tilde{\gamma}(t)} \left[\frac{D-2}{D} h'_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma} \right) - 2 \right] d\tilde{\gamma} \\ &= \frac{x_{\max}}{2\alpha^{2D-2} D^2 \gamma_2^2} \left[\frac{2\gamma_2^2\alpha^D}{\bar{x}} h_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma}(t) \right) - 2\tilde{\gamma}(t) \right]. \end{aligned} \quad (73)$$

Plugging eq. (73) in eq. (72) we get

$$\|\mathbf{w}(t)\|_2 \leq \frac{x_{\max}}{2\gamma_2^2} \left[\frac{2\gamma_2^2\alpha^D}{\bar{x}} h_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma}(t) \right) - 2\tilde{\gamma}(t) \right] + 2D^2\alpha^{2D-2} \int_0^t \|X\mathbf{r}(\tau)\|_2 d\tau. \quad (74)$$

Putting things together: From eqs. (70) and (12) we have

$$\gamma_2(t) \geq \frac{\tilde{\gamma}(t) - \log(N)}{\|\mathbf{w}(t)\|_2} \geq \frac{\tilde{\gamma}(t)}{\|\mathbf{w}(t)\|_2} - \frac{\log(N)x_{\max}}{\tilde{\gamma}(t)}. \quad (75)$$

Next we set $t = T_\alpha$ and take the limit $\alpha \rightarrow \infty$. Note that $\tilde{\gamma}(T_\alpha) \xrightarrow{\alpha \rightarrow \infty} \infty$ since $\epsilon(T_\alpha) \xrightarrow{\alpha \rightarrow \infty} 0$, and thus the right term in eq. (75) is vanishing. Using eq. (74) we get

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} \frac{1}{\gamma_2(T_\alpha)} &\leq \lim_{\alpha \rightarrow \infty} \frac{\|\mathbf{w}(T_\alpha)\|_2}{\tilde{\gamma}(T_\alpha)} \\ &\leq \lim_{\alpha \rightarrow \infty} \left[\frac{x_{\max}}{2\gamma_2^2} \left[\frac{2\gamma_2^2\alpha^D}{\bar{x}\tilde{\gamma}(T_\alpha)} h_D \left(\frac{(D-2)\bar{x}}{2D\gamma_2^2\alpha^D} \tilde{\gamma}(T_\alpha) \right) - 2 \right] + \frac{2D^2\alpha^{2D-2}}{\tilde{\gamma}(T_\alpha)} \int_0^{T_\alpha} \|X\mathbf{r}(\tau)\|_2 d\tau \right]. \end{aligned}$$

We use $\frac{\tilde{\gamma}(T_\alpha)}{\alpha^2} \xrightarrow{\alpha \rightarrow \infty} 0$, eq. (9) and eq. (71) to get

$$\lim_{\alpha \rightarrow \infty} \frac{1}{\gamma_2(T_\alpha)} \leq \frac{1}{\gamma_2}.$$

It follows that $\lim_{\alpha \rightarrow \infty} \gamma_2(T_\alpha) = \gamma_2$. \square

G.2 Intermediate Regime Proof

Theorem 18 (Theorem 6 for $D > 2$). *Under Condition 8, for $D > 2$ if $\lim_{\alpha \rightarrow \infty} \frac{\alpha^D}{\tilde{\gamma}(\alpha)} = \mu > 0$, then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} Q_\mu^D(\mathbf{w}) \text{ s.t. } \forall n, \mathbf{x}_n^\top \mathbf{w} \geq 1$$

where $Q_\mu^D(\mathbf{w}) = \sum_{i=1}^d q_D\left(\frac{w_i}{\mu}\right)$ and $q_D(s) = \int_0^s h_D^{-1}(z) dz$ for $h_D(z) = (1-z)^{-\frac{D}{D-2}} - (1+z)^{-\frac{D}{D-2}}$.

Proof. The proof is similar in spirit to the proof for the case $D = 2$ (see Appendix F.2). We show that the KKT conditions hold in the limit $\alpha \rightarrow \infty$. The KKT conditions are that there exists $\boldsymbol{\nu} \in \mathbb{R}_{\geq 0}^N$ such that

$$\nabla Q_\mu^D(\hat{\mathbf{w}}) = X\boldsymbol{\nu} \quad (76)$$

$$\forall n : \mathbf{x}_n^\top \hat{\mathbf{w}} \geq 1 \quad (77)$$

$$\forall n : \nu_n (\mathbf{x}_n^\top \hat{\mathbf{w}} - 1) = 0. \quad (78)$$

The proof of primal feasibility (77) for $D = 2$ applies also here.

Stationarity condition (76): Let

$$\boldsymbol{\nu} = \frac{D(D-2)}{\mu} \limsup_{\alpha \rightarrow \infty} \left(\alpha^{D-2} \int_0^{T_\alpha} \mathbf{r}(s) ds \right) \in \mathbb{R}_{\geq 0}^N. \quad (79)$$

We need to show that

$$\nabla Q_\mu^D(\mathbf{w}) = \frac{1}{\mu} h_D^{-1} \left(\frac{\mathbf{w}}{\mu} \right) = X \boldsymbol{\nu}.$$

Indeed using Lemma 12 and eq. (11) we have

$$\begin{aligned} \hat{\mathbf{w}} &= \lim_{\alpha \rightarrow \infty} \frac{\alpha^D h_D \left(\alpha^{D-2} D(D-2) X \int_0^{T_\alpha} \mathbf{r}(s) ds \right)}{\gamma(T_\alpha)} \\ &= \lim_{\alpha \rightarrow \infty} \frac{\tilde{\gamma}(T_\alpha)}{\gamma(T_\alpha)} \lim_{\alpha \rightarrow \infty} \frac{\alpha^D}{\tilde{\gamma}(T_\alpha)} \limsup_{\alpha \rightarrow \infty} h_D \left(\alpha^{D-2} D(D-2) X \int_0^{T_\alpha} \mathbf{r}(s) ds \right) \\ &= \mu h_D \left[\mu X \left(\frac{D(D-2)}{\mu} \limsup_{\alpha \rightarrow \infty} \left(\alpha^{D-2} \int_0^{T_\alpha} \mathbf{r}(s) ds \right) \right) \right] \\ &= \mu h_D (\mu X \boldsymbol{\nu}) \end{aligned}$$

and thus $\frac{1}{\mu} h_D^{-1} \left(\frac{\hat{\mathbf{w}}}{\mu} \right) = X \boldsymbol{\nu}$, as desired.

Complementary slackness (78): Let $k \in [N]$ such that

$$\mathbf{x}_k^\top \hat{\mathbf{w}} > 1. \quad (80)$$

We have to show that $\nu_k = 0$. We change variables $t \rightarrow \tilde{\gamma}(t)$ and using eq. (29) we get

$$\alpha^{D-2} \int_0^{T_\alpha} \exp(-\mathbf{x}_k^\top \mathbf{w}(s)) ds \leq \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_0^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma}. \quad (81)$$

Next we decompose the RHS of eq. (81) and employ Condition 8, where similarly to the case $D = 2$ we replace $\tilde{\gamma}^*(\alpha)$ with $\tilde{\gamma}_1^*(\alpha)$ and ρ_0 with ρ_1 (see the proof of the intermediate regime for $D = 2$ in Appendix F.2). We get that

$$\begin{aligned} \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_0^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} &= \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_0^{\tilde{\gamma}_1^*(\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &\quad + \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) d\tilde{\gamma} \\ &= (I) + (II) \end{aligned} \quad (82)$$

From eq. (10) we have that $\exp(-\mathbf{x}_k^\top \mathbf{w}(\tilde{\gamma}) + \tilde{\gamma}) \leq N$ and thus

$$(I) \leq \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_0^{\tilde{\gamma}_1^*(\alpha)} N d\tilde{\gamma} = \frac{N \tilde{\gamma}_1^*(\alpha)}{2D^2 \alpha^D \gamma_2^2} \xrightarrow{\alpha \rightarrow \infty} 0 \quad (83)$$

since $\tilde{\gamma}_1^*(\alpha) = o(\alpha^D)$. For the second term in eq. (82) we get for large enough α ,

$$\begin{aligned} (II) &= \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp \left[- \left(\mathbf{x}_k^\top \frac{\mathbf{w}(\tilde{\gamma})}{\tilde{\gamma}} - 1 \right) \tilde{\gamma} \right] d\tilde{\gamma} \\ &\leq \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_{\tilde{\gamma}_1^*(\alpha)}^{\tilde{\gamma}(T_\alpha)} \exp [- (\rho_1 - 1) \tilde{\gamma}] d\tilde{\gamma} \\ &\leq \frac{1}{2D^2 \alpha^D \gamma_2^2} \int_0^\infty \exp [- (\rho_1 - 1) \tilde{\gamma}] d\tilde{\gamma} \\ &= \frac{1}{2D^2 \alpha^D \gamma_2^2 (\rho_1 - 1)} \xrightarrow{\alpha \rightarrow \infty} 0. \end{aligned} \quad (84)$$

By substituting eqs. (83) and (84) in eq. (82) and back in eq. (81) we get that $\nu_k = 0$. \square

G.3 Rich Regime Proof

Theorem 19 (Theorem 7 for $D > 2$). *Under Condition 8, for $D > 2$ if $\tilde{\gamma}(\alpha) = \omega(\alpha^D)$ then*

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{w}\|_1 \quad \text{s.t. } \forall n : \mathbf{x}_n^\top \mathbf{w} \geq 1. \quad (85)$$

Proof. We show that the KKT conditions (62), (63), (64) for the ℓ_1 max-margin problem (85) hold in the limit $\alpha \rightarrow \infty$. To this end let

$$\boldsymbol{\nu}^{(\ell_1)} = D(D-2) \limsup_{\alpha \rightarrow \infty} \left(\alpha^{D-2} \int_0^{T_\alpha} \mathbf{r}(s) ds \right) \in \mathbb{R}_{\geq 0}^N. \quad (86)$$

Note that this definition is similar to eq. (79), and if $\nu_k = 0$ for some k then also $\nu_k^{(\ell_1)} = 0$. Therefore it is left to show the stationarity condition $X\boldsymbol{\nu}^{(\ell_1)} \in \partial^\circ \|\hat{\mathbf{w}}\|_1$. Indeed, from eq. (24) we know that $-1 \leq [X\boldsymbol{\nu}^{(\ell_1)}]_i \leq 1$ for all i . In addition

$$\begin{aligned} \hat{\mathbf{w}} &= \lim_{\alpha \rightarrow \infty} \frac{\alpha^D h_D \left(\alpha^{D-2} D(D-2) X \int_0^{T_\alpha} \mathbf{r}(s) ds \right)}{\gamma(T_\alpha)} \\ &\stackrel{(11)}{=} \lim_{\alpha \rightarrow \infty} \left(\frac{\alpha^D}{\tilde{\gamma}(T_\alpha)} h_D \left(\alpha^{D-2} D(D-2) X \int_0^{T_\alpha} \mathbf{r}(s) ds \right) \right). \end{aligned}$$

Assume that $\hat{w}_i > 0$. As $\lim_{\alpha \rightarrow \infty} \left(\frac{\alpha^D}{\tilde{\gamma}(T_\alpha)} \right) = 0$ we must have that

$$\left[\alpha^{D-2} D(D-2) X \int_0^{T_\alpha} \mathbf{r}(s) ds \right]_i \xrightarrow{\alpha \rightarrow \infty} 1$$

and thus $[X\boldsymbol{\nu}^{(\ell_1)}]_i \xrightarrow{\alpha \rightarrow \infty} 1$. Similarly, if $\hat{w}_i < 0$ we get $[X\boldsymbol{\nu}^{(\ell_1)}]_i \xrightarrow{\alpha \rightarrow \infty} -1$. It follows that $X\boldsymbol{\nu}^{(\ell_1)} \in \partial^\circ \|\hat{\mathbf{w}}\|_1$. \square

H Additional Simulation Results and Details

H.1 Optimization trajectories with $\tilde{\gamma}$ indicators

In Figure 6 we repeat the optimization trajectories from Figure 3, but we add indicators that indicate the value of $\tilde{\gamma}$ along the path. Recall that $\tilde{\gamma}(t) = -\log \epsilon(t)$. For example, a number 10 near some point on the path means that the loss at this point is $\exp(-10)$.

In all three examples we observe that for $\alpha = 100$, where the trajectory first visits the ℓ_2 predictor, around the ℓ_2 predictor we have $\tilde{\gamma} = 10^4 = \alpha^2$, as suggested by our theoretical results. In the top figure we also plot the path for $\alpha = 10000$, and again around the ℓ_2 predictor it holds that $\tilde{\gamma} = 10^8 = \alpha^2$.

In addition we can see that in order to be rather close to ℓ_1 with large initialization, we need very large $\tilde{\gamma}$, corresponding to extremely small loss ϵ . For example, consider the center plot. For $\alpha = 100$ to be close to ℓ_1 direction we need $\tilde{\gamma} = 10^8$, or $\epsilon = \exp(-10^8)$! However, with small initialization, e.g. $\alpha = 0.001$, $\tilde{\gamma}$ can be as small as 0.1, or $\epsilon = \exp(-0.1) \approx 0.9$, and we are close to ℓ_1 .

H.2 Understanding the non-unique ℓ_1 case

In Figure 3(c) we showed an example of optimization trajectories for data with non-unique ℓ_1 predictor. We can observe that for different initializations the selected ℓ_1 direction, and thus the implicit bias, is different.

It is interesting to understand what are the properties of different ℓ_1 solutions. To this end, in Figure 7 we plot the optimization trajectory in a different way. Instead of looking at the direction of the predictor (as in Figure 3(c)), we consider the excess ℓ_1 and ℓ_2 norms along the path, defined as

$\|\mathbf{w}(t)\|_1/\|\mathbf{w}_{\ell_1}\|_1 - 1$ and $\|\mathbf{w}(t)\|_2/\|\mathbf{w}_{\ell_2}\|_2 - 1$ where \mathbf{w}_{ℓ_1} and \mathbf{w}_{ℓ_2} are the ℓ_1 and ℓ_2 max-margin (minimum norm) solutions accordingly.

We can observe that for large initialization, where we follow the Q_μ path, the selected ℓ_1 predictor has the smallest ℓ_2 norm. Moreover, for small initialization the selected ℓ_1 predictor has the largest ℓ_2 norm. Thus, in this case we see an example where the asymptotic (at long time/small loss) implicit bias is affected by the initialization. This in contrast to previous results for exp-tailed losses (e.g. Soudry et al. [2018], Ji and Telgarsky [2019b], Gunasekar et al. [2018b], Nacson et al. [2019a], Lyu and Li [2020]), where the asymptotic bias was independent of the initialization.

H.3 Local minima in high dimension

In Figure 8 we consider optimization trajectories for data in dimension 10. In this case we cannot show the direction of the predictor $\frac{\mathbf{w}(t)}{\|\mathbf{w}(t)\|_2}$ on a sphere, as we did for data in dimension 3. Instead, we take the approach similar to Figure 7, where we show the excess margins.

We consider two datasets in dimension 10. The first is a random, yet separable, data composed of 10 points where the coordinates are drawn from $\sim \mathcal{U}(0, 1)$. The second dataset is a sparse dataset of 4 points, where the first coordinate is 1 and the other 9 coordinates are random noise $\sim \mathcal{U}(0, 0.5)$. This dataset allows a large separation between the ℓ_2 max-margin and ℓ_1 max-margin solutions.

We train depth-3 linear diagonal network and plot the optimization trajectories in $\ell_{2/3}$ - ℓ_2 plane. We observe that for the random data (Figure 8(a)) there are many local minima of the max $\ell_{2/3}$ margin, and depending on initialization we are biased towards different local-minima points. However, with a large initialization we converge to a local point, quite close to the Q_μ path and ℓ_1 .

For the sparse data (Figure 8(b), and a zoom-in shown in Figure 8(c)) the local minima are quite far away from the paths. Also in this case the ℓ_1 and $\ell_{2/3}$ max-margin solutions are the same, and with a large initialization we converge to them, along the Q_μ path. This seem to suggest that for certain structure data, like sparse data, we tend to converge to the *global* max $\ell_{2/3}$ -margin predictor.

H.4 Tangent kernel during training

In Figure 5(a) we showed how the excess ℓ_1 -norm depends on α and depth D , and measured closeness to the rich limit by excess ℓ_1 -norm. An alternative and complementary approach is to look at the tangent kernel $K_t(\mathbf{x}, \mathbf{x}') = \langle \nabla_{\mathbf{u}} f(\mathbf{u}(t), \mathbf{x}), \nabla_{\mathbf{u}} f(\mathbf{u}(t), \mathbf{x}') \rangle$, which is directly related to closeness to the kernel regime. As discussed in Section 2, the tangent kernel is almost fixed in the kernel regime, yet can change significantly when we exit the kernel regime.

In Figure 9 we show the kernel distance during optimization for the same data and network (depth 2) as in Figure 5(a). The kernel distance is defined as $1 - \text{CosineSimilarity}(K(t), K(0))$ where $K(0) \in \mathbb{R}^{N \times N}$ is the tangent kernel at initialization, $K(t) \in \mathbb{R}^{N \times N}$ is the tangent kernel at time t , and

$$\text{CosineSimilarity}(K(t), K(0)) = \frac{\langle K(t), K(0) \rangle}{\|K(t)\|_2 \|K(0)\|_2} = \frac{\text{Tr}(K(t)K^\top(0))}{\|K(t)\|_2 \|K(0)\|_2}.$$

Here we focus on exiting the kernel regime, rather than closeness to the rich regime. We observe that increasing depth will help to exit the kernel regime (where *distance* ≈ 0) earlier, at a larger loss value ϵ . Decreasing the initialization has a similar effect, and this is consistent with Figure 5(a).

H.5 Addressing Numerical Issues

In our simulations we employ the normalized gradient descent update rule, given by

$$\mathbf{u}(t+1) = \mathbf{u}(t) - \eta \frac{\nabla \mathcal{L}(\mathbf{u}(t))}{\mathcal{L}(\mathbf{u}(t))}$$

where $\mathbf{u} \in \mathbb{R}^{2d}$ is the vector of parameters and

$$\mathcal{L}(\mathbf{u}(t)) = \frac{1}{N} \sum_{n=1}^N \exp(-\tilde{\mathbf{x}}_n^\top \mathbf{u}^D(t)).$$

This algorithm effectively enlarges the learning rate according to the current loss, and for single layer linear models Nacson et al. [2019b] showed that the loss decreases exponentially faster.

Let $G(\mathbf{u}(t)) = \frac{\nabla \mathcal{L}(\mathbf{u}(t))}{\mathcal{L}(\mathbf{u}(t))}$. During training the loss can become extremely small, e.g. well beyond 10^{-1000} , and in this case also the gradient is very small. This can cause numerical issues in calculating G . In order to have a numerically stable evaluation of G , and avoid cases like $0/0$, we take an approach similar to Lyu and Li [2020]. Specifically, let

$$\bar{\gamma}_n(t) = \tilde{\mathbf{x}}_n^\top \mathbf{u}^D(t) \quad , \quad \gamma(t) = \min_n \bar{\gamma}_n(t)$$

Then we have that

$$\nabla \mathcal{L}(\mathbf{u}(t)) = -\frac{D}{N} \mathbf{u}^{D-1}(t) \circ \sum_{n=1}^N \exp(-\bar{\gamma}_n(t)) \tilde{\mathbf{x}}_n$$

and

$$\begin{aligned} G(\mathbf{u}(t)) &= -D \mathbf{u}^{D-1}(t) \circ \frac{\sum_{n=1}^N \exp(-\bar{\gamma}_n(t)) \tilde{\mathbf{x}}_n}{\sum_{n=1}^N \exp(-\bar{\gamma}_n(t))} \\ &= -D \mathbf{u}^{D-1}(t) \circ \frac{\sum_{n=1}^N \exp(-(\bar{\gamma}_n(t) - \gamma(t))) \tilde{\mathbf{x}}_n}{\sum_{n=1}^N \exp(-(\bar{\gamma}_n(t) - \gamma(t)))} . \end{aligned} \quad (87)$$

We calculate G according to (87). Note that $\max_n \exp(-(\bar{\gamma}_n(t) - \gamma(t))) = 1$ so the denominator is at least 1 and the sum in the numerator will contain at least one support vector $\tilde{\mathbf{x}}_{n:\bar{\gamma}_n(t)=\gamma(t)}$.

It is important to note that we never represent the loss values, but only the parameters \mathbf{u} . Thus, as long as \mathbf{u} can be represented by `float64` precision the simulation can continue and we get extremely large parameters corresponding to extremely small loss.

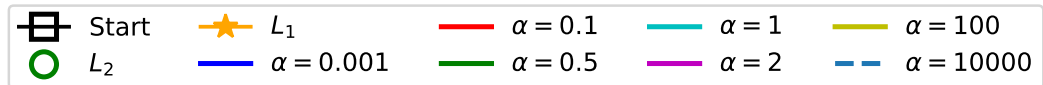
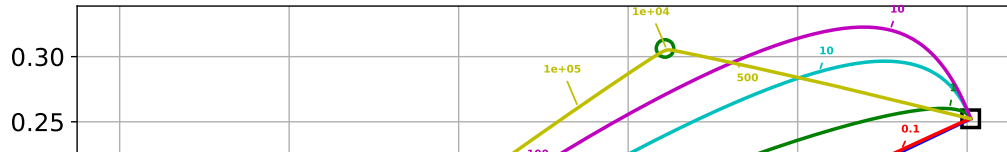
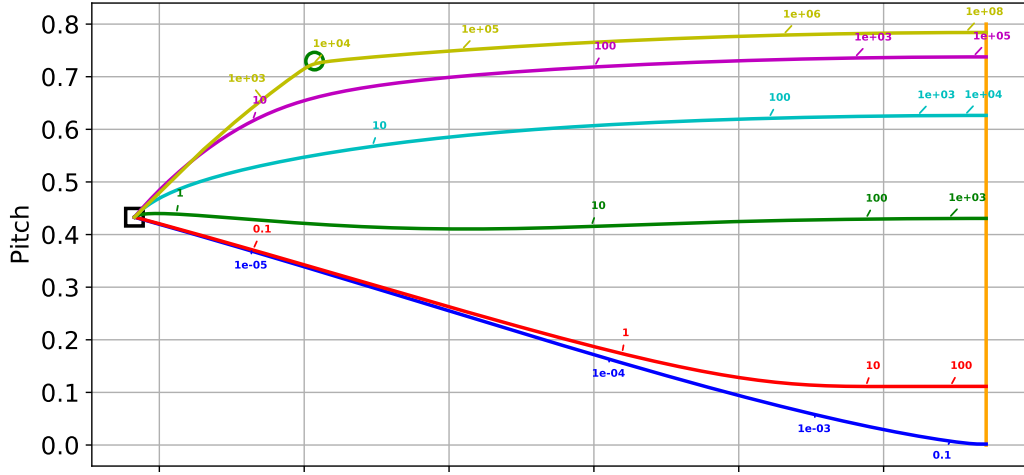
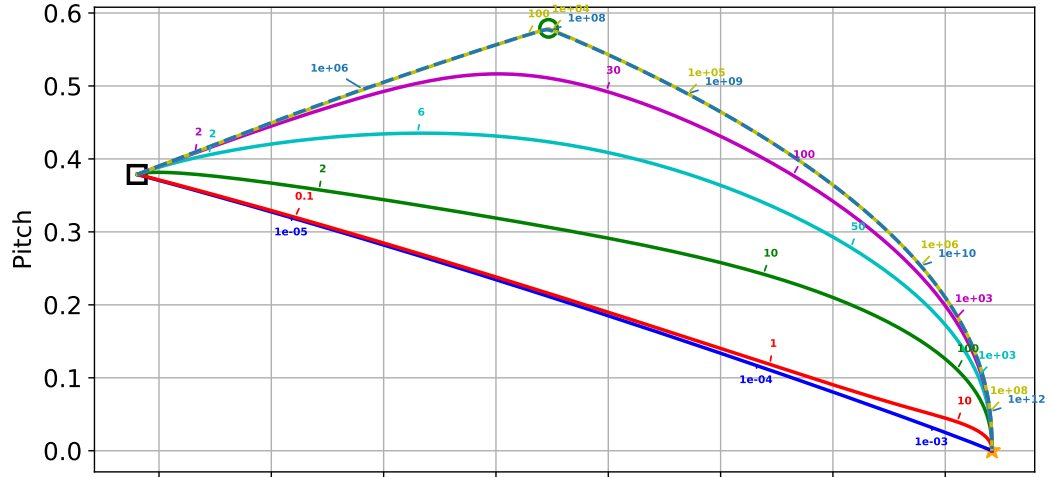


Figure 6: The same optimization trajectories from Figure 3 with $\tilde{\gamma}$ values indications.

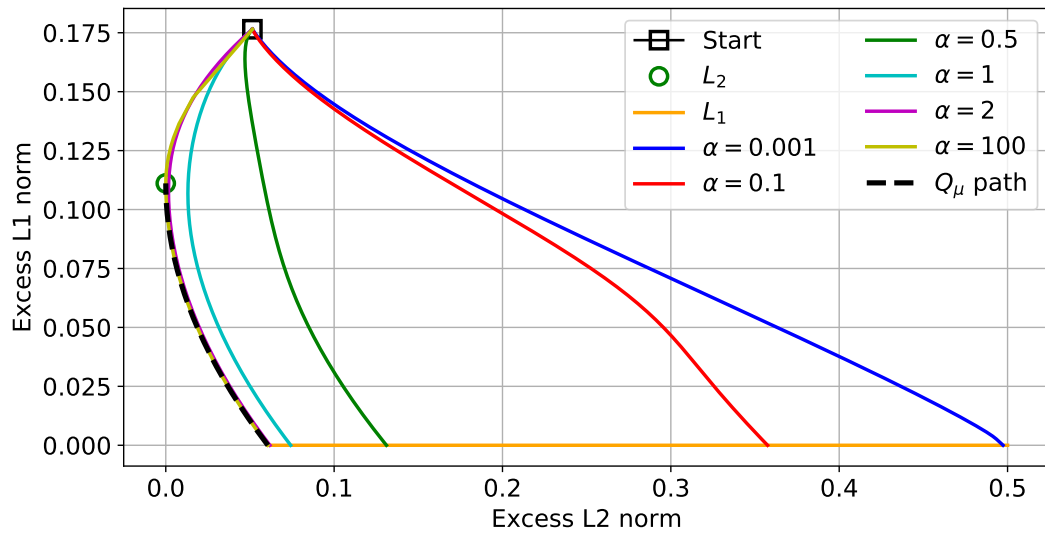
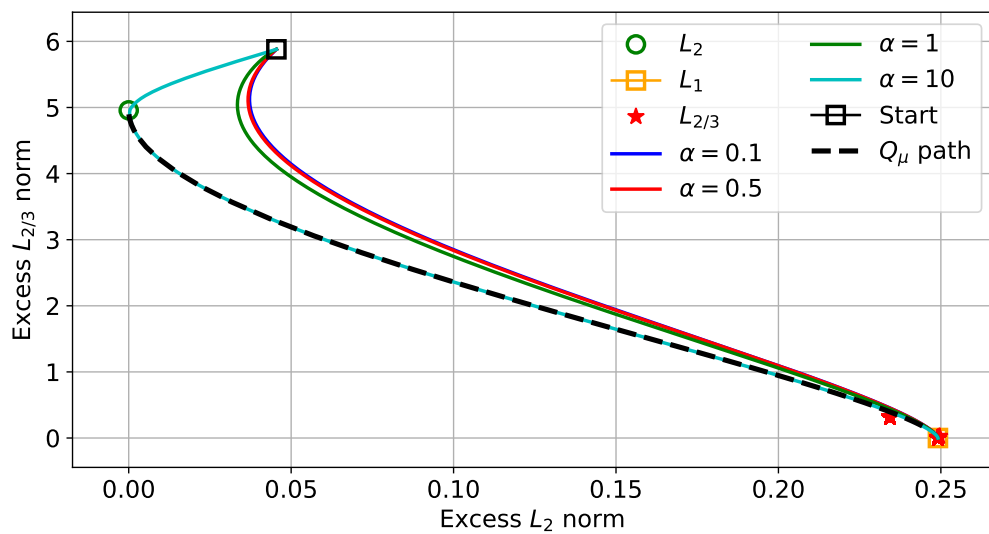
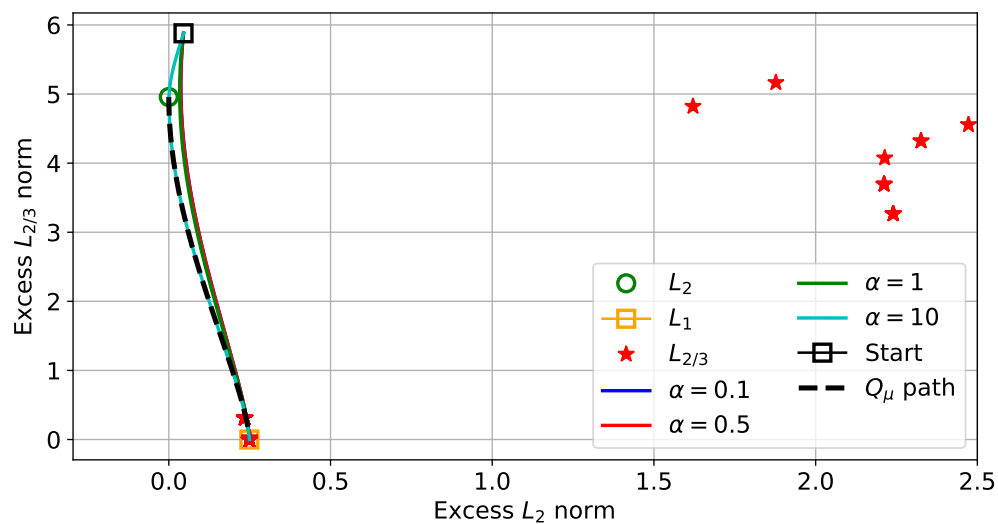
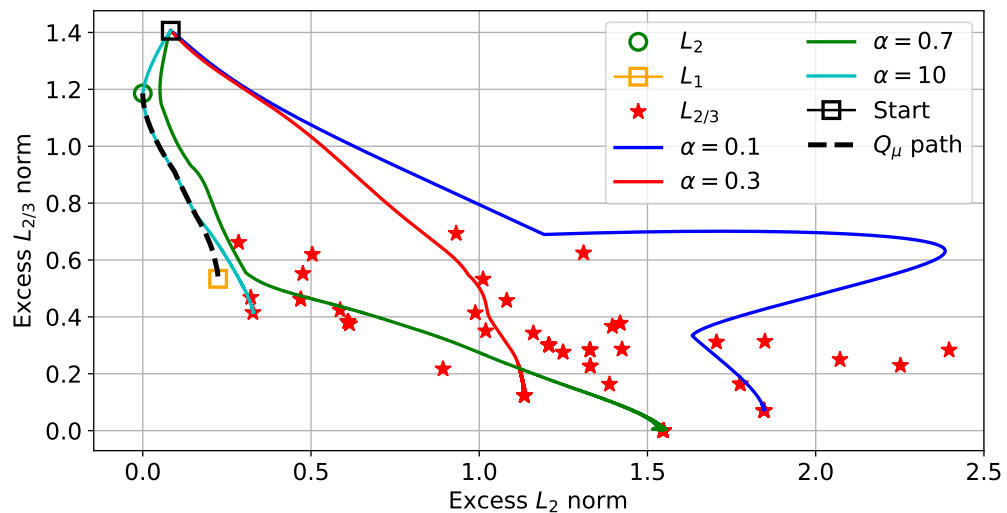


Figure 7: Optimization trajectories for data in Figure 3(c) in excess ℓ_1 -norm - excess ℓ_2 -norm plane.



(c) Zoom-in of Figure 8(b)

Figure 8: Optimization trajectories for data in dimension 10.

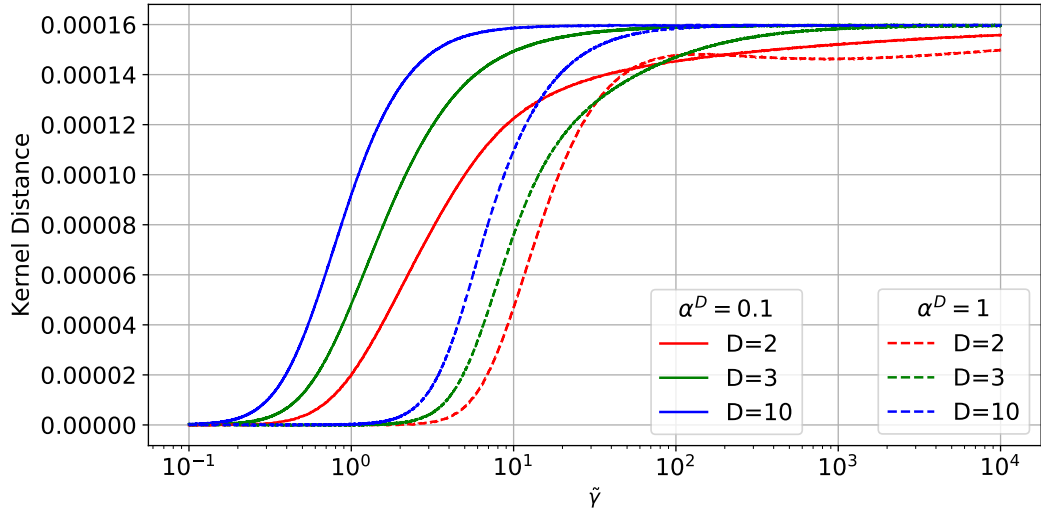


Figure 9: The Kernel distance is defined as $1 - \text{CosineSimilarity}(K(t), K(0))$ where $K(0)$ is the tangent kernel at initialization and $K(t)$ is the tangent kernel at time t .