arXiv:2005.13530v1 [math.AP] 27 May 2020

ON THE CONVERGENCE OF GRADIENT DESCENT TRAINING FOR TWO-LAYER RELU-NETWORKS IN THE MEAN FIELD REGIME

STEPHAN WOJTOWYTSCH

Abstract. We describe a necessary and sufficient condition for the convergence to minimum Bayes risk when training two-layer ReLU-networks by gradient descent in the mean field regime with omni-directional initial parameter distribution. This article extends recent results of Chizat and Bach to ReLU-activated networks and to the situation in which there are no parameters which exactly achieve MBR. The condition does not depend on the initalization of parameters and concerns only the weak convergence of the realization of the neural network, not its parameter distribution.

1. Introduction

Practitioners have found that artificial neural networks can be trained by gradient descent-based algorithms to fit many complicated target functions. While it is well-understood why the function class is sufficiently expressive for this purpose [Bar93, Cyb89, Hor91, LLPS93], the choice of optimal network parameters in applications is a highly non-convex problem. It is not fully understood why cleverly initialized gradient descent achieves good performance in practice.

In [CB18b], the authors prove that if the parameter distributions of neural network-like models converge to a limiting distribution, then the limit is in fact a global minimizer. The result is true for a class of 'spread out' initial conditions (implying very wide networks) and under some technical assumptions. One of the most prominent cases concerns function models with the same homogeneity as neural networks with a single hidden layer and ReLU (or leaky ReLU) activation. However, the result does not apply directly due to the lack of differentiability of these activation functions.

In this article, we extend the main result of [CB18b] for this ReLU-like setting and generalize previous results for some cases. The results proved here improve upon previous work in two ways:

(1) Our analysis applies to ReLU-networks with suitable initial conditions rather than toy models with similar properties.
(2) We only assume that a limiting object, whose existence is guaranteed by compactness, is unique.

In the analysis we exploit that ReLU activation

$$\phi(a, w, b; x) = a \max\left\{w^T x + b, 0\right\}$$

is not only positively two-homogeneous on parameter space $\mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$, but in fact the product of two one-homogeneous functions (one of which is the identity map on $\mathbb{R}$). This product structure allows us to avoid certain 'bad' points in the gradient flow. We state a special case of the main result informally.

Date: October 27, 2021.
2020 Mathematics Subject Classification. 35Q68, 68T07, 49Q22, 35F20.
Key words and phrases. Two-layer network, population risk, Wasserstein gradient flow, mean field model, ReLU activation, global minimizer, convergence, minimum Bayes risk.

Theorem. *Consider a two-layer mean field network model*

$$f_\pi(x) = \int_{\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}} a\,\sigma(w^T x + b)\,\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

with activation function $\sigma = ReLU$ (or leaky ReLU) and parameter distribution π on $\mathbb{R}\times\mathbb{R}^d\times\mathbb{R}$. Let π_t evolve by the 2-Wasserstein gradient flow

$$\frac{d}{dt}\pi_t = \operatorname{div}\left(\pi_t \nabla \frac{\delta \mathcal{R}}{\delta \pi}\right)$$

of a population risk functional

$$\mathcal{R}(\pi) = \int_{\mathbb{R}^d\times\mathbb{R}} \ell\big(f_\pi(x), y\big)\,\mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

with a convex and Lipschitz-continuous loss function ℓ. If the initial distribution π_0 is uniform on $[-1,1]\times S^d \subseteq \mathbb{R}\times\mathbb{R}^{d+1}$, then the velocity potential

$$\frac{\delta \mathcal{R}}{\delta \pi}(\pi_t; a, w, b) = \int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1\ell)\big(f_{\pi_t}(x), y\big)\,a\,\sigma(w^T x + b)\,\mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

lies in a compact subset of $C^0_{loc}(\mathbb{R}^{d+2})$ as $t\to\infty$. The following are equivalent.

(1) As time goes to infinity, $\mathcal{R}(\pi_t)$ converges to minimum Bayes risk.
(2) The velocity potential $\frac{\delta\mathcal{R}}{\delta\pi}$ converges to a unique limit g locally uniformly as $t\to\infty$.

If $\mathcal{R}(\pi_t)$ approaches MBR, we can additionally identify the limit

$$\lim_{t\to\infty} \frac{\delta\mathcal{R}}{\delta\pi} = 0.$$

If additionally π_t converges to a limit π_∞ in 2-Wasserstein distance, π_∞ minimizes $\mathbb{R}$.

The study of Wasserstein gradient flows in the context of (infinitely wide) shallow neural networks is motivated below. The same problem exactly captures the gradient descent training of finite neural networks for infinitesimal learning rate and asymptotically captures stochastic gradient descent when the learning rate is so small/batch size is so large that stochastic effects are negligible.

A further technical condition has to be imposed on $\mathbb{P}$, and the result only holds conditionally on a Morse-Sard type property for suitable two-layer neural networks. The Sard-condition is purely technical, but has neither been verified nor disproved in general. In the appendix, we show that the condition holds in dimension $d = 2$, but may fail in certain cases if $d \geq 8$. Without the Morse-Sard condition, we can show that the following are equivalent.

(1) As time goes to infinity, $\mathcal{R}(\pi_t)$ converges to minimum Bayes risk.
(2) The velocity potential $\frac{\delta\mathcal{R}}{\delta\pi}$ converges to 0 locally uniformly as $t\to\infty$.

The loss function is assumed to be C^1-smooth (with bounded derivative). We can consider more general loss functions (e.g. mean squared error), but only for uniformly bounded data. The initial condition can be generalized, but we require that at time $t = 0$ we have $-a^2 + |w|^2 + b^2 \geq 0$ π-almost surely. The Minkowski inner product is preserved along the gradient evolution and the condition $(w, b) \neq 0$ unless also $a = 0$ allows us to avoid issues of non-differentiability. To establish existence of a minimizer, it suffices that a certain projection of π_t converges.

The significance of our result is as follows.

- We do not need to assume the existence of a limit (guaranteed by compactness), but only its uniqueness.

- One of the main complications of the question of convergence to minimal energy for this gradient flow is the dependence on the initial condition. The question whether the limit of the velocity potential is generically unique can be asked independently of the initial condition and seems more approachable by standard means of analysis. We therefore believe that this perspective might be a first step towards the convergence theory for Wasserstein gradient flows for shallow neural networks.

In the proof, we show directly that convergence to minimal energy implies convergence of the velocity potentials $\frac{\delta \mathcal{R}}{\delta \pi}$ to zero. This resembles the first order optimality condition in classic calculus and does not require deep geometric insight. In the converse direction, the geometry of the energy landscape is used via the homogeneity of the activation function. While we formulate all results for ReLU-activation, they hold equivalently for leaky ReLU networks. We prove that the second moments of the evolving parameter distribution π_t grow at most sublinearly under general conditions and show that if the unique limit g does not vanish identically, they grow exponentially. Thus only $g \equiv 0$ is admissible as a unique limit. In this situation the second moments – which are bounded below by zero – decrease linearly at a positive rate unless $\mathcal{R}(\pi_t)$ decays to minimum Bayes risk.

The article is structured as follows. In the remainder of this section, we briefly review previous work and collect some notation. In Section 2, we describe the setting, state our main assumptions, and review some additional background information on Wasserstein gradient flows and continuity equations. In Section 3, we state our main results. Some concrete examples of data distributions and loss functionals to which the results apply are listed in Section 4. We conclude with a discussion of our results in Section 5. Longer proofs and some technical details are collected in Appendix A. Appendix B is dedicated to the technical condition of Morse-Sard type which we assume.

1.1. **Related Work.** Introductions to machine learning in general and neural networks in specific can be found for example in [SSBD14, HH19, GBC16, MBW$^+$19].

The question why (stochastic) gradient descent starting at a suitable random initialization finds good parameters for artificial neural networks despite the fact that the energy landscape is highly non-convex has attracted much attention and several competing explanations have emerged.

One avenue of research aims to uncover a description of gradient descent in neural networks in the infinite neuron limit in the mean field scaling regime [CB18b, RVE18, SS20, MMN18]. Convergence criteria in the shallow network setting are developed for example in [AKSG19] and [CB18b], but have not been established generally.

Another line of articles [EMWW19, EMW19d, BM19, DZPS18, DLL$^+$18, JGH18, ADH$^+$19] considers a heavily overparametrized regime with large initialization. In this setting, the gradient flow for neural networks behaves is proved to behave like the gradient flow of a very wide random feature model [EMW19d, EMWW19] with high probability over the choice of (suitable) initial condition. In particular, the direction and bias of a neuron barely change from their (random) initialization in this regime. Chizat and Bach dub this the 'lazy training regime' since neurons hardly move. They show that the underlying analysis is due rather to a (usually implicit) scaling assumption on initialization rather than the specific structure of neural networks [CB18a].

The success of neural networks in practical applications has been explained by the observation that – unlike any linear theory – neural networks can beat the 'curse of dimensionality' [Bar93]. It thus seems unlikely that linearization is able to explain their recent success. Furthermore, in these studies parameters are usually initialized so large that the path norm a two-layer network with m hidden neurons scales like $\sqrt{m}$ at initialization. Natural generalization bounds as derived in [EMW19a, EMW18] therefore do not apply.

Practitioners tend to train neural networks using stochastic gradient descent rather than full gradient descent. For small learning rate (time step size), the evolution can be described by SDEs with Gaussian noise [HLLL19, LTE15]. The noise coefficient is given by the covariance of the gradient, which is usually neither isotropic nor homogeneous. However, in the mean field regime and assuming that the noise is standard Gaussian, one can prove that the parameter distribution converges to a good value (minimizer of a regularized risk functional), see [HRSS19]. The derivation is built on the link of the heat equation to both stochastic analysis and optimal transport theory [JKO98]. In this case, the parameter distribution approaches the stationary measure of a Markov process as time approaches infinity. If the noise is sufficiently large, the convergence is exponential, while for small noise, the stationary measure approaches a minimizer of the mean field risk functional.

Rigorous convergence results in realistic settings can be obtained under strong assumptions on the initial condition (or a state which arises along the gradient flow). One example is [BJ18], where the authors show that if a neural network has well chosen parameters at some time, the parameters improve along the gradient flow. In [AKSG19] the authors study a mean field gradient flow and show that under a regularity/closeness assumption, the gradient flow converges. Unfortunately, the condition cannot be verified in practice. Global convergence for a toy model with similar properties is established in [EMW19c, Section 7].

Some results in [CB18b] also apply to neural networks with smooth activation and more than one hidden layer. However, the imposition of a linear structure results in network-like models where each neuron in the outermost layer has its own trainable weights for the deeper layers. More recent works in the mean field setting consider deep neural networks whose parameters are initialized independently across the layers [AOY19, NP20, SS19]. The independence is preserved through time ('propagation of chaos') and a mean field description is available in different scaling limits. The theory is entirely different from that of shallow networks. While a shallow networks can be described by indexed particles (a_i, w_i, b_i), the paths in a deep network through multiple layers have a more complicated interacting multi-index structure (a_i, b_{ij}, c_j).

1.2. **Notations and Terminology.** We denote by $|\cdot|$ the Euclidean norm on any finite-dimensional vector space. For a map $\phi : A \to B$ and a measure μ on A, we denote the push-forward of μ along ϕ by $\phi_\sharp\mu$ (which is a measure on B), see e.g. [EG15]. By $\mathrm{spt}(\mu)$, we denote the support of the measure μ. For a measure μ and a locally μ-integrable function f we denote by $f \cdot \mu$ the measure

$$(f \cdot \mu)(B) = \int_B f \mathrm{d}\mu$$

which has density f with respect to μ (occasionally denoted by $\mu|_f$ in other texts). We denote probability measures on data space by $\mathbb{P}$, projection operators by P and the space of probability measures by $\mathcal{P}$. Parameter distributions are denoted by π and assumed to be elements of the space $\mathcal{P} = \mathcal{P}_2$ of probability measures with finite second moments (Wasserstein space).

2. Background and Assumptions

2.1. **Activation, Loss and Data.** In this section, we describe the objects under consideration in this article. Some assumptions will be relaxed in Section 3.4.

Activation function. We consider the *parameter domain*

$$\Theta = \big\{(a, w, b) \in \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \,:\, a^2 < |w|_{\ell^2}^2 + b^2\big\}$$

and the *activation function* $\phi : \overline{\Theta} \to \mathbb{R}$

$$\phi(\theta; x) = a\,\sigma(w^T x + b) \qquad \text{where} \quad \theta = (a, w, b), \quad \sigma(z) = z_+ = \max\{z, 0\}.$$

For this introduction, the choice of cone Θ is inessential and will be motivated later in Section 2.5. Note that the family

$$\{\phi(\theta; x) : \theta \in \Theta\}$$

has the universal approximation property, i.e. for any continuous function f on a compact set $K \subseteq \mathbb{R}^d$ and every $\varepsilon > 0$, there exists a finite collection of parameters $\{\theta_1, \dots, \theta_m\}$ such that

$$\left\| f - \frac{1}{m} \sum_{i=1}^{m} \phi(\theta_i, \cdot) \right\|_{C^0(K)} < \varepsilon. \tag{2.1}$$

The property is known to hold when the parameters θ_i are not constrained to a cone, the factor $\frac{1}{m}$ is not present and coefficients a_i are included before $\phi(\theta_i, \cdot)$ [Cyb89]. None of these differences are significant as

$$\phi\left(\lambda a; \frac{w}{\lambda}, \frac{b}{\lambda}; \cdot\right) \equiv \phi(a, w, b; \cdot) \quad \forall\ \lambda > 0 \qquad \text{and} \quad \lambda\, \phi(a, w, b; \cdot) \equiv \phi(\lambda a, w, b; \cdot) \quad \forall\ \lambda \in \mathbb{R}.$$

Parameter distribution. A *parameter distribution* π is a Borel probability measure on $\overline{\Theta}$ with finite second moments

$$N(\pi) := \int_{\overline{\Theta}} |\theta|^2 \, \pi(\mathrm{d}\theta) = \int_{\overline{\Theta}} |a|^2 + |w|_{\ell^2}^2 + |b|^2 \, \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b),$$

i.e. $\pi \in \mathcal{P} := \mathcal{P}_2(\overline{\Theta})$ is an element of Wasserstein space over the cone $\overline{\Theta}$. We denote the *realization* of π as

$$f_\pi(x) = \int_{\overline{\Theta}} \phi(\theta; x) \, \pi(\mathrm{d}\theta)$$

and note that f_π is $N(\pi)$-Lipschitz and defined on the whole space $\mathbb{R}^d$.

Loss function. We further assume that the *loss function* $\ell : \mathbb{R} \times \mathbb{R} \to [0, \infty)$ satisfies the following.

- (L1) ℓ is jointly continuous.
- (L2) ℓ is once continuously differentiable in the first argument and there exists $C_\ell > 0$ such that
$$|\partial_1 \ell|(y, y') \le C_\ell \qquad \forall\ y, y' \in \mathbb{R}.$$
- (L3) The function $y \mapsto \ell(y, \bar{y})$ is convex for all $\bar{y} \in \mathbb{R}$.

Except for the Lipschitz condition (L2), the assumptions are common and non-restrictive. The assumption that $|\partial_1 \ell|$ is uniformly bounded is required for technical purposes to compensate the singularity of the ReLU function. It will be relaxed in Section 3.4.

Data distribution. Finally, we consider *data-distributions* $\mathbb{P}$ on $\mathbb{R}^d \times \mathbb{R}$ such that

- (P1) $\mathbb{P}$ is a Borel measure, i.e. every continuous function on $\mathbb{R}^d \times \mathbb{R}$ is $\mathbb{P}$-measurable.
- (P2) The first moments
$$\int_{\mathbb{R}^d \times \mathbb{R}} |x| + |y| \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$
are finite.
- (P3) The space spanned by the collection of functions $f(\theta_i; \cdot)$ is dense in $L^1(\overline{\mathbb{P}})$ where $\overline{\mathbb{P}} := P^x_\sharp \mathbb{P}$ denotes the projection of $\mathbb{P}$ on the first component.
- (P4) We assume that the map
$$S^d \to L^1\big(\sqrt{1 + |x|^2} \cdot \overline{\mathbb{P}}\big), \qquad (w, b) \mapsto 1_{\{x | w^T x + b > 0\}}$$
is Lipschitz continuous.

(P1) is a non-restrictive technical condition. (P2) is required to show that the risk of linearly growing functions is always finite. (P3) is always met if $\overline{\mathbb{P}}$ is compactly supported due to the Universal Approximation Theorem [Cyb89] and the density of continuous functions in L^1 for any Radon measure [FL07, Theorem 2.11]. (P4) imposes a high degree of smoothness on the data distribution which is used to compensate for the lack of differentiability of the activation function ϕ at points (a, w, b) and x for which $w^T x + b = 0$. Intuitively, $\overline{\mathbb{P}}$ cannot concentrate on or close to lower-dimensional *linear* objects. A full discussion of admissible data distributions is beyond the scope of this article, but we give some examples.

Theorem 2.1. *(1) Assume that $\overline{\mathbb{P}}$ has a density ρ with respect to Lebesgue measure on $\mathbb{R}^d$ such that*

$$\rho(x) \leq C\big(1 + |x|^2\big)^{-\frac{d+2+\varepsilon}{2}}$$

for some $C, \varepsilon > 0$. Then $\overline{\mathbb{P}}$ satisfies (P4).

(2) The uniform distribution on the unit sphere satisfies (P4).

(3) The set of data-distributions $\overline{\mathbb{P}}$ satisfying (P4) is convex.

(4) For any $L > 0$ the set of data-distributions $\overline{\mathbb{P}}$ satisfying

$$\big\|1_{\{w^T x + b > 0\}} - 1_{\tilde{w}^T x + \tilde{b} > 0\}}\big\|_{L^1(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}})} \leq L\big[|w - \tilde{w}| + |b - \tilde{b}|\big]$$

is convex and closed under weak convergence of Radon measures on $\mathbb{R}^d$.

We note that as a finite Borel measure on the locally compact Polish space $\mathbb{R}^d$, $\mathbb{P}$ is in fact a Radon measure. Due to [ABM14, Theorem 4.2.4] $\mathbb{P}$ can be decomposed into conditional probabilities $\mathbb{P}^x$ and a distribution $\overline{\mathbb{P}} = P^x_\sharp \mathbb{P}$ on $\mathbb{R}^d$ like in (P3) such that

(1) for any $\mathbb{P}$-measurable function $g : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$, the map

$$x \mapsto \int_{\mathbb{R}} g(x, y)\, \mathbb{P}^x(\mathrm{d}y)$$

is $\overline{\mathbb{P}}$-measurable and

(2) the equality

$$\int_{\mathbb{R}^d \times \mathbb{R}} g(x, y)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) = \int_{\mathbb{R}^d} \left(\int_{\mathbb{R}} g(x, y)\, \mathbb{P}^x(\mathrm{d}y) \right) \overline{\mathbb{P}}(\mathrm{d}x)$$

holds for all $\mathbb{P}$-measurable functions g.

Consider the *augmented loss function*

$$L : \mathbb{R}^d \times \mathbb{R} \to [0, \infty), \qquad L_x(\alpha) := \int_{\mathbb{R}} \ell(\alpha, y)\, \mathbb{P}^x(\mathrm{d}y)$$

which encodes many important properties of the problem.

2.2. **The Risk Functional.** Combining all previous notions, we define the risk functional $\mathcal{R} : \mathcal{P}_2(\overline{\Theta}) \to [0, \infty]$,

$$\mathcal{R}(\pi) := \int_{\mathbb{R}^d \times \mathbb{R}} \ell\big(f_\pi(x), y\big)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) = \int_U L_x\big(f_\pi(x)\big)\, \overline{\mathbb{P}}(\mathrm{d}x). \tag{2.2}$$

We impose the following compatibility conditions between ℓ and $\mathbb{P}$ to control the behaviour of $\mathcal{R}$.

(LP1) *Growth of augmented loss*:

$$\lim_{|\alpha| \to \infty} L_x(\alpha) = \infty$$

for $\overline{\mathbb{P}}$-almost every $x \in \mathbb{R}^d$.

(LP2) *Finite risk*:

$$\int_{\mathbb{R}^d\times\mathbb{R}} \ell(0,y)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y) < \infty.$$

Remark 2.2. A few observations are in order. Recall that $\mathcal{P}$ or $\mathcal{P}_2$ denotes the space of Radon measures with finite second moments (on the natural space in a given context, which usually is $\overline{\Theta}$).

(1) Since ℓ and f_π are Lipschitz continuous, we can use (LP2) and (P2) to deduce that $\mathcal{R}(\pi) < \infty$ for all $\pi \in \mathcal{P}_2$.
(2) For any fixed x, the function L_x is convex and differentiable in α if it is finite (i.e. for almost every x). If ℓ is strictly convex, so is L_x.
(3) The function $(x,\alpha) \to L_x(\alpha)$ is measurable (see Lemma 2.3 below), but cannot generally be assumed to be continuous.

We can extend the risk functional to the larger space of $\overline{\mathbb{P}}$-measurable functions by setting

$$\widetilde{\mathcal{R}}(f) := \int_{\mathbb{R}^d\times\mathbb{R}} \ell\big(f(x),y\big)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y) = \int_U L_x\big(f(x)\big)\,\overline{\mathbb{P}}(\mathrm{d}x)$$

with possibly infinite values. L_x is convex and continuous in α for $\overline{\mathbb{P}}$-almost every x and satisfies the growth condition (LP1), so there exists a compact convex set $M_x \subset \mathbb{R}$ such that $L_x(\alpha) = \inf_{\alpha'\in\mathbb{R}} L_x(\alpha')$ for all $\alpha \in M_x$. If L_x ist strictly convex (for example because ℓ is strictly convex), then the minimum of L_x is unique.

We note that there exists a measurable selection of minimizer for L_x. This is not entirely immediate even when L_x is strictly convex.

Lemma 2.3. *(1) The function $(x,\alpha) \to L_x(\alpha)$ is measurable on $\mathbb{R}^d\times\mathbb{R}$ with respect to the product σ-algebra generated by $\overline{\mathbb{P}}$ and the Borel sigma algebra on $\mathbb{R}$.*
(2) There exists a $\overline{\mathbb{P}}$-measurable function $f^*:\mathbb{R}^d\to\mathbb{R}$ such that $f^*(x) \in M_x$ for $\overline{\mathbb{P}}$-almost every $x\in\mathbb{R}^d$. f^* satisfies

$$(\nabla L_x)\big(f^*(x)\big) = \int_{\mathbb{R}} \big(\partial_1\ell\big)\big(f^*(x),y\big)\,\mathbb{P}^x(\,\mathrm{d}y) = 0 \qquad \overline{\mathbb{P}}-a.e.$$

It is clear that

$$\mathcal{R}(\pi) = \widetilde{\mathcal{R}}(f_\pi) \geq \widetilde{\mathcal{R}}(f^*)$$

and that equality is achieved if and only if $f_\pi(x) \in M_x$ $\overline{\mathbb{P}}$-almost everywhere. In particular, if L_x is strictly convex, equality holds if and only if $f_\pi = f^*$ $\overline{\mathbb{P}}$-almost everywhere. We add the following compatibility condition between $\mathbb{P}$ and ℓ.

(LP3) *Minimum Bayes risk*: There exists a version of f^* in $L^1(\overline{\mathbb{P}})$.

Since ℓ is Lipschitz-continuous in the first argument, assumption (P3) implies that

$$\inf_{\pi\in\mathcal{P}_2(\overline{\Theta})} \mathcal{R}(\pi) = \widetilde{\mathcal{R}}(f^*)$$

as f^* can be approximated arbitrarily well in $L^1(\overline{\mathbb{P}})$ by functions of type f_π. Thus $\widetilde{\mathcal{R}}(f^*)$ is the minimum Bayes risk of the problem.

Corollary 2.4. *The functional $\mathcal{R}$ admits a minimizer if and only if there exists a measure $\pi\in\mathcal{P}$ such that*

$$f_\pi = f^* \quad \overline{\mathbb{P}}-\textit{almost everywhere.}$$

In particular, as f_π is $N(\pi)$-Lipschitz (where again $N(\pi)$ denotes the second moments of π), there must be a version of f^* which is Lipschitz-continuous. The Lipschitz condition is far from sufficient [EW20b]. Since

$$f_{\lambda\pi+(1-\lambda)\pi'} = \lambda f_\pi + (1-\lambda)\, f_{\pi'} \qquad \forall\, \pi, \pi' \in \mathcal{P},\ \lambda \in [0,1],$$

the set of minimizers is a convex subset of $\mathcal{P}$.

Remark 2.5. Write $T(a,w,b) = (-a,w,b)$. The fact that $\phi(T\theta,\cdot) \equiv -\phi(\theta,\cdot)$ implies that the measure π representing a function f_π cannot be unique: Any measure π such that $T_\sharp\pi = \pi$ (i.e. $\pi(T^{-1}(V)) = \pi(V)$ for all measurable $V \subseteq \Theta$) represents the function 0 since

$$f_\pi(x) = \int \phi(\theta,x)\,\pi(\mathrm{d}\theta) = \int \phi(\theta,x)\,T_\sharp\pi(\mathrm{d}\theta) = \int \phi(T\theta,x)\,\pi(\mathrm{d}\theta) = -\int \phi(\theta,x)\,\pi(\mathrm{d}\theta) = -f_\pi(x).$$

In ReLU networks, another source of non-uniqueness is the identity

$$0 = (x+1) - x - 1 = \sigma(x+1) - \sigma\big(-(x+1)\big) - \sigma(x) + \sigma(-x) - \sigma(1).$$

More generally, if π represents f_π and π' represents 0, then for any $\lambda \in (0,1)$ consider the probability measure

$$\pi_\lambda = \lambda\, D(\lambda)_\sharp\pi + (1-\lambda)\pi'$$

where $D(\lambda) = \mathrm{diag}(\lambda^{-1/2})$ is a dilation. Then π_λ is a probability measure with finite second moments and

$$f_{\pi_\lambda} = \lambda f_{D(\lambda)_\sharp\pi} + (1-\lambda)\, f_{\pi'} = \lambda \int_\Theta f(\lambda^{-1/2}\theta,\cdot)\,\pi(\mathrm{d}\theta) = \int_\Theta f(\theta,\cdot)\,\pi(\mathrm{d}\theta) = f_\pi.$$

In particular, while the minimizer f^* of $\widetilde{\mathcal{R}}$ is unique if L_x is strictly convex, a minimizer of $\mathcal{R}$ must be highly non-unique. This is a major obstacle in linearization-based approaches to convergence.

2.3. **Wasserstein Gradient Flows.** The study of Wasserstein gradient flows in machine learning is motivated by the following observation: *The parameters* $\{\theta_i\}_{i=1}^m \in \overline{\Theta}^m$ *of a parametrized function*

$$f_{\theta_1,\dots\theta_m}(x) = \frac{1}{m}\sum_{i=1}^m \phi(x,\theta_i)$$

evolve by the time-accelerated Euclidean gradient flow

$$\dot\theta_i = -m\,\nabla_{\theta_i}\mathcal{R}\big(\theta_1(t),\dots,\theta_m(t)\big)$$

of

$$\mathcal{R}(\theta_1,\dots,\theta_n) = \int_{\mathbb{R}^d\times\mathbb{R}} \ell\big(f_{\theta_1,\dots,\theta_m}(x),y\big)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)$$

if and only if their distribution $\pi_m = \frac{1}{m}\sum_{i=1}^m \delta_{\theta_i}$ *follows the* 2*-Wasserstein gradient flow of the extended risk functional*

$$\mathcal{R}(\pi) = \int_{\mathbb{R}^d\times\mathbb{R}} \ell\big(f_\pi(x),y\big)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)$$

where

$$f_\pi(x) = \int_{\overline{\Theta}} \phi(\theta;x)\,\pi(\mathrm{d}\theta).$$

A key observation is that the individual particles θ_i are irrelevant and only their distribution matters when computing $f_{\theta_1,\dots,\theta_m}$. We refer to two-layer network functions of the form $f(x) = \frac{1}{m}\sum_{i=1}^m a_i\,\sigma(w_i^T x + b_i)$ as *mean field networks* in contrast to *classical two-layer networks* $f(x) = \sum_{i=1}^m a_i\,\sigma(w_i^T x + b_i)$. Both classes are identical from the perspective of approximation theory,

but lead to different dynamic models in the infinite-width limit. Classical networks are described by the linearized dynamics of neural tangent kernels, while mean field networks evolve truly non-linearly by Wasserstein gradient flows.

Thus optimizing mean field network parameters by the gradient flow of a risk functional is equivalent to optimizing their distribution by a Wasserstein gradient flow, see e.g. [CB18b, Proposition B.1]. An expanded heuristic can also be found in Appendix A.

The gradient flow and the map $\pi \mapsto f_\pi$ (so also the risk functional $\mathcal{R}$) are naturally defined on the Wasserstein space $\mathcal{P}_2$ of probability measures with finite second moments. Background on optimal transport theory and Wasserstein gradient flows can be found e.g. in [AGS08, San15, Vil08].

In this article, we mostly consider Wasserstein gradient flows of continuous distributions. By continuity, the Wasserstein gradient flows $\pi_m(t)$ starting at π_m^0 converge to the solution of the Wasserstein gradient flow $\pi(t)$ starting at $\pi^0 = \lim_{m\to\infty} \pi_m^0$ for all $t > 0$. Even more, the limits

$$\lim_{m\to\infty} \lim_{t\to\infty} \pi_m(t) = \lim_{t\to\infty} \lim_{m\to\infty} \pi_m(t) = \lim_{t\to\infty} \pi(t)$$

commute (if they exist). A proof under an additional technical condition (which can be eliminated if one considers gradient flows of unregularized risk or the second moment regularizer) is given in Appendix B of [CB18b]. The complication arises from regularizing functionals which do not have the correct homogeneity, which leads the authors to consider compactly supported initial conditions.

Hence if the gradient flow starting at a measure π^0 converges to a minimizer of risk, then gradient flows starting at closeby empirical measures asymptotically achieve low risk. The theory, at this point, is purely qualitative, but nonetheless indicative for practical applications.

The Wasserstein gradient flow π_t is described by the continuity equation

$$\frac{d}{dt}\pi_t = \operatorname{div}_\theta\big(\pi_t \, \nabla_\theta(\delta_\pi \mathcal{R})(\pi_t, \cdot)\big) \tag{2.3}$$

in the distributional sense, i.e.

$$\frac{d}{dt}\int_\Theta g(\theta)\, \pi_t(\mathrm{d}\theta) = -\int_\Theta \big\langle \nabla g(\theta), \nabla_\theta(\delta_\pi \mathcal{R})\big\rangle \, \pi_t(\mathrm{d}\theta)$$

for $g \in C_c^\infty(\Theta)$ where

$$\delta_\pi \mathcal{R}(\pi, \theta) = \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)(f_\pi(x), y)\, \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

$$\nabla_\theta\, \delta_\pi \mathcal{R}(\pi, \theta) = \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)(f_\pi(x), y)\, \nabla_\theta \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

denote the variational derivative of $\mathcal{R}$ with respect to π and its spatial gradient respectively.

2.4. **Continuity Equations.** Consider a general continuity equation

$$\begin{cases} \frac{d\mu}{dt} &= \operatorname{div}(\mu V) \quad t > 0 \\ \mu &= \mu_0(\theta) \qquad t = 0 \end{cases}$$

in the space of finite Radon measures on the whole space $\mathbb{R}^d$ and the associated flow map

$$\begin{cases} \frac{dX}{dt}(t, \theta) &= V\big(t, X(t, \theta)\big) \quad t > 0 \\ X(0, \theta) &= \theta \end{cases}. \tag{2.4}$$

Assume for the moment that V lies in the Bochner space $L^1\big((0,T), W^{1,\infty}(\mathbb{R}^d; \mathbb{R}^d)\big)$, i.e. V is Lipschitz-continuous in space at almost every time $t > 0$ and

$$\int_0^T \|V(t,\cdot)\|_{L^\infty(\mathbb{R}^d)} + \int_0^T \|\nabla_\theta V(t,\cdot)\|_{L^\infty(\mathbb{R}^d)}\, \mathrm{d}t < \infty.$$

Then according to [Amb08, Proposition 4], we have

$$\mu_t = X(t,\cdot)_\sharp \mu_0, \quad \text{i.e.} \int_{\mathbb{R}^d} g(\theta)\,\mu_t(\mathrm{d}\theta) = \int_{\mathbb{R}^d} g\big(X(t,\theta)\big)\,\mu_0(\mathrm{d}\theta) \tag{2.5}$$

for all $g \in C_c^\infty(\mathbb{R}^d)$. Note that by Grönwall's Lemma we also have

$$\mathrm{Lip}(X(t,\cdot)) \leq \exp\left(\int_0^t \mathrm{Lip}(V(s,\cdot))\,\mathrm{d}s\right). \tag{2.6}$$

The Lemma does not apply directly to the situation which we will consider since the flow field V will be positively one-homogeneous and thus unbounded in most cases. However, the result also applies under the weaker assumption

$$\frac{V(t,\theta)}{1+|\theta|} \in L^1\big([0,T],\, L^\infty(\mathbb{R}^d;\mathbb{R}^d)\big),$$

see [Amb08, Remark 7]. The condition of at most linear growth is required to prevent particles from escaping to infinity in finite time. Thus (2.5) and (2.6) apply also in our situation.

2.5. **Initial Condition.** There are two considerations concerning the initial parameter distribution of the gradient flow. The first is specific to ReLU activation and of a technical nature while the second one is more geometric and concerns energy decay to minimum Bayes risk/convergence to minimizers.

2.5.1. *The cone of good parameters.* We can formally compute the parameter gradient of the activation function

$$\nabla_\theta \phi(\theta;x) = \begin{pmatrix} \sigma(w^T x + b) \\ a\,\sigma'(w^T x + b)\,x \\ a\,\sigma'(w^T x + b) \end{pmatrix}$$

which is defined at all $\theta = (a,w,b)$ and x such that $w^T x + b \neq 0$ with $\sigma'(z) = 1_{\{z>0\}}$. In particular, if for every $(w,b) \neq 0$ the hyperplane $\{x : w^T x + b = 0\}$ is a $\overline{\mathbb{P}}$-null set, then the parameter gradient of risk

$$\nabla(\delta_\pi \mathcal{R})(\pi;\theta) = \int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1 \ell)\big(f_\pi(x),y\big)\,(\nabla_\theta\phi)(\theta;x)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)$$

is well-defined at $(w,b) \neq 0$. This is a regularity condition on $\overline{\mathbb{P}}$ which is satisfied for example whenever $\overline{\mathbb{P}}$ has a density with respect to Lebesgue measure or the uniform distribution on a sphere. For technical reasons, we need the gradient to even be Lipschitz continuous, which is why we imposed assumption (P4) on the data distribution $\mathbb{P}$.

There is a more subtle problem of regularity. Assumption (P4) guarantees that

$$(w,b) \mapsto \int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1 \ell)\big(f_\pi(x),y\big)\,1_{\{w^T x+b\}}\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)$$

is a Lipschitz-function on the sphere $S^d \subseteq \mathbb{R}^{d+1}$. However, close to a point where $(w,b) = 0$ and $a \neq 0$, the half-spaces $\{x : w^T x + b > 0\}$ oscillate rapidly and the gradient

$$(a,w,b) \mapsto \int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1 \ell)\big(f_\pi(x),y\big)\,1_{\{w^T x+b\}}\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)$$

generally fails to even be continuous. This difficulty can only be circumvented if a is close to zero whenever (w,b) is close to zero. At this point, the geometry of the ReLU function as the product of two positively one-homogeneous functions comes into play. Namely, we can use the fact that

$$\sigma(z) = \sigma'(z)z$$

to show that along trajectories of the flow map X in (2.4), the Minkowski norm $-a^2+|w|_{\ell^2}^2+b^2$ remains constant:

$$\begin{aligned}\frac{d}{dt}\big[-a^2+|w|_{\ell^2}^2+b^2\big] &= 2\big[-a\dot a+\langle w,\dot w\rangle+b\dot b\big]\\ &= 2\big[a\,\partial_a(\delta_\pi\mathcal{R})-\langle w,\nabla_w(\delta_\pi\mathcal{R})\rangle-b\,\partial_b(\delta_\pi\mathcal{R})\big]\\ &= 2\int_{\mathbb{R}^d}(\partial_1\ell)(f_\pi(x),y)\big[a\,\sigma(w^Tx+b)-a\,\sigma'(w^Tx+b)\cdot(w^Tx+b)\big]\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y)\\ &= 0.\end{aligned}$$

In particular, the cone of space-like vectors

$$\Theta=\{(a,w,b):-a^2+|w|^2+b^2>0\}$$

is preserved along the gradient flow evolution. Inside $\overline{\Theta}$, we have $|a|\le\sqrt{|w|^2+b^2}$, which allows us to save the Lipschitz property. We therefore restate an assumption on the initial condition more explicitly.

(IC1) *Small linear variable*: $\pi_0\in\mathcal{P}_2(\overline{\Theta})$, i.e. π_0 is a probability measure on $\mathbb{R}^{d+2}$ with finite second moments such that $\mathrm{spt}(\pi_0)\subseteq\overline{\Theta}$.

Remark 2.6. It is possible to consider initial conditions supported on other super-level sets like $\Theta_\varepsilon=\{(a,w,b):-a^2+|w|^2+b^2>\varepsilon^2\}$ instead. A flow starting at π_0 supported on Θ_ε will never reach a point where $(w,b)=0$, but on the other hand does not allow us to exploit the positive two-homogeneity of ϕ in θ as easily since Θ_ε is not a cone. Positive two-homogeneity is essential to many arguments below, so we proceed with $\Theta=\Theta_0$ instead.

2.5.2. *Omni-directional initial conditions.* It is well-known that the functional $\mathcal{R}$ is not sufficiently convex in Wasserstein geometry to guarantee convergence to global minimizers from any initial condition. In particular, if a global minimizer π exists but cannot be written as an empirical measure with m atoms, any initial condition corresponding to an empirical measure with m atoms *cannot* converge to π since the continuity equation has no smoothing effect and preserves atomic measures.

The following class of initial conditions is successful in theory and applications.

Definition 2.7 (Omni-directional measure)**.** We call a probability measure π on Θ *omni-directional* if every open cone in Θ has positive measure.

(IC2) *Omnidirectional initial condition*: π_0 is omni-directional.

Clearly, an omni-directional initial condition is an abstraction available only in the infinite-width limit.

Remark 2.8. A measure π_0 can be omni-directional in $\Theta=\Theta_0$ and yet supported on a smaller set Θ_ε for $\varepsilon>0$ since for any $\theta\in\Theta$ and $\lambda\in\mathbb{R}$ the Minkowski norm of $\lambda\theta$ is

$$\lambda^2(-a^2+|w|^2+b^2)>\varepsilon^2\quad\forall\ \lambda>\frac{\varepsilon}{\sqrt{-a^2+|w|^2+b^2}}.$$

Thus for any $\theta\in\Theta$, there exists $\lambda>0$ such that $\lambda\theta\in\Theta_\varepsilon$. In particular, if $\mathrm{spt}(\pi_0)=\overline{\Theta_\varepsilon}$ for any $\varepsilon>0$, then π_0 is omni-directional.

2.6. **Morse-Sard Property.** Below, we prove existence and uniqueness for gradient flow training of ReLU-activated two-layer networks under the assumptions listed above. For technical reasons, we add an assumption which has only been established in full generality in dimension $d=2$. The assumption is only required when discussing the limiting behavior of the gradient flow.

(M-S) *Morse-Sard condition*: Assume that

$$(\delta_\pi \mathcal{R})(\pi_{t_n}, \cdot) \to g, \qquad \nabla(\delta_\pi \mathcal{R})(\pi_{t_n}, \cdot) \to V$$

locally uniformly on $\overline{\Theta}$. Then the restriction of g to the unit sphere $S^{d+1} = \{a^2 + |w|^2 + b^2 = 1\}$ has the property that $P_*^{S^d} V = \nabla^S g$ does not vanish anywhere on the level set $\{g = t\}$ for Lebesgue-almost every t, where ∇^S denotes the gradient of g tangent to the sphere and $P_*^{S^d}$ the tangent map to the projection onto the sphere.

We discuss the Morse-Sard condition below in Appendix B. Using smoothness, we establish the property in dimension two and provide evidence on the other hand that (M-S) may not be expected to hold in high dimension in full generality. We isolate the point in the proof where the condition is used and where an argument would need to be adapted in order to avoid it.

A weaker version of the main result holds without this condition.

3. Evolution of the Parameter Distribution

3.1. **Existence and Uniqueness.** We first establish that solutions to gradient flow training exist. Assume all conditions outlined above except for (IC2) and (M-S), which are only required for statements about limiting objects.

Lemma 3.1. *Let $\pi_0 \in \mathcal{P}_2(\overline{\Theta})$. Then there exists a unique solution to the Wasserstein gradient flow* (2.3).

The next statement will be useful to understand the behavior of limits as $t \to \infty$.

Lemma 3.2. *If $\pi_0 \in \mathcal{P}_2(\overline{\Theta})$ is omni-directional and π_t denotes the Wasserstein gradient flow starting at π_0, then π_t is omni-directional for all $t > 0$.*

This Lemma is a simpler version of [CB18b, Theorem 3.3]. The proof is based on the flow map representation. Since the activation function is positively two-homogeneous in the network parameters, the flow field $-\nabla(\delta_\pi \mathcal{R})$ is positively one-homogeneous, which means that half-rays move as half-rays and cones are preserved under the flow. Both results also apply directly to explicitly regularized risk functions with suitable homogeneity such as

$$F_\varepsilon(\pi) = \mathcal{R}(\pi) + \varepsilon \int_{\overline{\Theta}} |a|^2 + |w|^2 + b^2 \, \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b), \qquad \varepsilon \geq 0.$$

The proofs are given in the appendix.

3.2. **Growth of Second Moments.** Recall that we had denoted the second moment of π by

$$N(\pi) := \int_{\overline{\Theta}} |\theta|^2 \, \pi(\mathrm{d}\theta) = \int_{\overline{\Theta}} a^2 + |w|_{\ell^2}^2 + b^2 \, \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b).$$

Under gradient flow training, N can only grow sublinearly in time.

Lemma 3.3. *If π_t evolves by the Wasserstein-gradient flow of $\mathcal{R}$, then*

$$N(\pi_t) \leq 2[N(\pi_0) + \mathcal{R}(\pi_0)\, t] \qquad \text{and} \qquad \lim_{t \to \infty} \frac{N(\pi_t)}{t} = 0.$$

The proof of the Lemma is fairly general and does not require the specific structure of ϕ, ℓ and $\mathbb{P}$. The second moments of π_t control the path-norm (or *Barron norm*) $\|f_{\pi_t}\|_{\mathcal{B}}$, where

$$\|f\|_{\mathcal{B}} = \inf \left\{ \int_{\overline{\Theta}} |a| \left[|w| + |b|\right] \pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b) \,\middle|\, \pi \in \mathcal{P}_2 \text{ and } f_\pi = f \, \overline{\mathbb{P}} - \text{a.e.} \right\}.$$

Neural networks with low Barron norm are poor approximators for general Lipschitz functions in the $L^2(\mathbb{P})$-topology if the data distribution $\overline{\mathbb{P}}$ is truly high-dimensional [EW20b]. The slow

growth of the Barron norm implies that gradient flows cannot decrease L^2-population risk at rates faster than $t^{-4/(d-2)}$ for general target functions f^* [WE20].

On the other hand, the key insight of [CB18b] is that the homogeneity of ϕ implies that if the velocity potentials $\delta_\pi \mathcal{R}$ were to converge to a non-trivial limit as $t \to \infty$, it would lead to exponential growth of $N(\pi_t)$.

Lemma 3.4. *Assume that*

$$\delta_\pi \mathcal{R}(\pi_t, \cdot) = \int (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, \phi(\cdot, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

$$\nabla_\theta\, \delta_\pi \mathcal{R}(\pi_t, \cdot) = \int (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, \nabla_\theta \phi(\cdot, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

converge to a function g and a vector field $V = \nabla g$ respectively, locally uniformly on $\overline{\Theta}$. If $g \not\equiv 0$, then

$$N(\pi_t) \geq c_0\, e^{c_1 t}$$

for some $c_0, c_1 > 0$ and $t \gg 1$.

We give both proofs in the appendix for the reader's convenience. The proof of Lemma 3.4 is the only point in the document where the assumptions (IC2) and (M-S) are used.

3.3. **Main Results: Lipschitz Loss.** We can now characterize the convergence of Wasserstein gradient flows for the risk functional $\mathcal{R}$ with omni-directional initial conditions entirely. We consider the *simultaneous ω-limit set*

$$\Omega^{lim} := \{(g, V) \,\big|\, \exists\ t_n \to \infty \text{ s.t. } \delta_\pi \mathcal{R}(\pi_{t_n}, \cdot) \to g,\ \nabla(\delta_\pi \mathcal{R})(\pi_{t_n}, \cdot) \to V \text{ locally uniformly}\}.$$

Theorem 3.5. *Assume the above conditions (L1), (L2), (L3), (P1), (P2), (P3), (P4), (LP1), (LP2), (IC1), (IC2), (M-S). Then*

(1) Ω^{lim} is not empty and
(2) $\lim_{t\to\infty} \mathcal{R}(\pi_t) = \inf_\pi \mathcal{R}(\pi)$ if and only if Ω^{lim} consists of only one element.

If Ω^{lim} consists of only one element, then that element is $\lim_{t\to\infty}(\delta_\pi \mathcal{R})(\pi_t; \theta) \equiv 0$.

The first statement of the Theorem does not require (IC2) and (M-S). Verifying that there exists only a single element in Ω^{lim} is non-trivial. Generally uniqueness is proved by showing that the limit satisfies an equation which can be studied separately. Even assuming that there exists a limiting measure π_∞ such that $\pi_t \to \pi_\infty$, the natural 'zero dissipation in the limit' condition

$$0 = \int_{\overline{\Theta}} |V|^2\, \pi_\infty(\mathrm{d}\theta) = \int_{\overline{\Theta}} \big|\nabla(\delta_\pi \mathcal{R})\big|^2(\pi_\infty, \theta)\, \pi_\infty(\mathrm{d}\theta)$$

only shows that $g(\theta) = 0$ for π_∞-almost all θ. This is significantly weaker than *every* θ if π_∞ concentrates on a small set. Functions satisfying the zero-dissipation property are stationary points of the flow, and there are many of them which are not the global minimum. Not assuming the existence of a limit π_∞, linearization becomes highly involved since the map $\pi \mapsto f_\pi$ is far from injective. This makes the question of uniqueness non-trivial. Despite these obstacles, we believe the result to be a first step towards a general convergence theory. Without assumptions (IC2), (M-S), a weaker result holds.

Theorem 3.6. *Assume that (L1), (L2), (L3), (P1), (P2), (P3), (P4), (LP1), (LP2) and (IC1) hold. Then*

(1) Ω^{lim} is not empty and
(2) $\lim_{t\to\infty} \mathcal{R}(\pi_t) = \inf_\pi \mathcal{R}(\pi)$ if and only if $|\Omega^{lim}|$ contains only the function $g = 0$.

Having found a characterization for risk converging to zero for omni-directional initial conditions, we now turn our attention to the question of whether we can find a minimizer of risk. The corollary is almost immediate.

Corollary 3.7. *If in addition to the assumptions of Theorem 3.5 we assume that there exist a sequence of times $t_n \to \infty$ and a measure π_∞ on Θ such that $\pi_{t_n} \to \pi_\infty$ in 2-Wasserstein distance and $|\Omega^{lim}| = 1$, then $f_{\pi_\infty} = f^*$, i.e. π_∞ is a risk minimizing measure.*

The existence of a suitable subsequence is guaranteed if the p-th moments of π_t remain uniformly bounded for some $p > 2$. A version of Corollary 3.7 also holds under the same assumptions ad Theorem 3.6. Under a weaker condition, we obtain a weaker statement.

Corollary 3.8. *If in addition to the assumptions of Theorem 3.5 we assume that the second moments $N(\pi_t)$ remain uniformly bounded (at least along a subsequence of times $t_n \to \infty$) and $|\Omega^{lim}| = 1$, there exists a measure $\pi_\infty \in \mathcal{P}$ such that $f_{\pi_\infty} = f^*$ (i.e. a risk minimizing measure).*

Again, there is a version of Corollary 3.8 under the conditions of Theorem 3.6. The measures π_{t_n} may not have a limit in 2-Wasserstein distance, but we can explicitly construct π_∞ from π_{t_n} by a reparametrization argument. If the second moments of π_t remain uniformly bounded in time, there exists a subsequence π_{t_n} which converges to a limit in p-Wasserstein distance for all $p < 2$. This is not sufficient since the map $\pi \mapsto f_\pi$ is not continuous in this topology, as the following example shows:

$$\pi_n := \frac{1}{n}\,\delta_{\theta=\sqrt{n}e_1} + \left(1 - \frac{1}{n}\right)\,\delta_{\theta=0}, \quad \pi_\infty = \delta_{\theta=0}, \qquad f_{\pi_n} \equiv \phi(e_1, \cdot) \neq 0 = f_{\pi_\infty}.$$

3.4. **Main Results: Smooth Loss and Bounded Data.** The previous results can be used to prove a different version of the main theorems in which a regularity assumption is shifted from the loss function to the data distribution. We make the following observation: If $\overline{\mathbb{P}}$ is supported on $B_R(0)$ for some $R > 0$, then

$$\|f_{\pi_t}\|_{L^\infty(\overline{\mathbb{P}})} \leq (1+R)\,N(\pi_t) \leq 2\,(1+R)\big[N(\pi_0) + \mathcal{R}(\pi_0)\,t\big]$$

by Lemma 3.3. Assume further that $|y| \leq R$ $\mathbb{P}$-almost surely and that the loss function ℓ satisfies

$$\sup_{|y|\leq S,\,|y'|\leq R} \big[|\ell| + |\partial_1 \ell|\big](y, y') < \infty$$

for any $S < \infty$. We consider the modified loss function

$$\ell_S(y, y') = \begin{cases} \ell(y, y') & |y| \leq S \\ \ell(S, y') + (\partial_1\ell)(S, y') \cdot (y - S) & y > S \\ \ell(-S, y') + (\partial_1\ell)(-S, y') \cdot (y + S) & y < -S \end{cases}$$

for $S > 0$ and the modified risk functional

$$\mathcal{R}_S(\pi) = \int_{\mathbb{R}^d \times \mathbb{R}} \ell_S(f_\pi(x), y)\,\mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y).$$

Then

(1) ℓ_S satisfies all conditions previously imposed and
(2) $\mathcal{R}_S(\pi) = \mathcal{R}(\pi)$ for all π such that $N(\pi) \leq \frac{S}{1+R}$.

In particular, the gradient flow π_t^S of $\mathcal{R}_S$ starting at a fixed measure π_0 independent of S exists and $\pi_t^S = \pi_t^{S'}$ describes the gradient flow of $\mathcal{R}$ for $t < c\,\min\{S, S'\}/(1+R)$. Note that the growth rate of N does not depend on S. Thus we may take S to infinity to find that the gradient flow π_t of $\mathcal{R}$ exists on $[0, \infty)$. We have shown the following.

Lemma 3.9. *Replace assumption (P2) by the stronger assumption*

(P2') $\mathbb{P}$ *has compact support in* $\mathbb{R}^{d+1}$.

and condition (L2) by the weaker assumption that

(L2') ℓ *satisfies the growth condition*

$$\sup_{|y|\le S,\,(x,y')\in \operatorname{spt}(\mathbb{P})} \big[|\ell| + |\partial_1 \ell|\big](y,y') < \infty$$

for all $S > 0$.

Then the gradient flow π_t *of* $\mathcal{R}$ *starting at* $\pi_0 \in \mathcal{P}_2(\overline{\Theta})$ *exists. If* π_0 *is omni-directional, so is* π_t *for all* $t > 0$.

Also the results on the convergence of gradient flows generalize under slightly stronger assumptions on the loss function. Examining the proof of Theorem 3.5, we find that the relevant properties of the proof are the following, which we postulate as conditions for smooth loss functions.

(SL1) $\inf_\pi \mathcal{R}(\pi) = \widetilde{\mathcal{R}}(f^*)$.

(SL2) For any $\pi \in \mathcal{P}, \theta \in \Theta$, the functions $(\partial_1\ell)(f_\pi(x),y)\, f^*(x)$ and $(\partial_1\ell)(f_\pi(x),y)\,\phi(\theta;x)$ are $\mathbb{P}$-integrable.

(SL3) If

$$\int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1\ell)(f_\pi(x),y)\,\phi(\theta;x)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y) \to 0 \quad \forall \theta\in\Theta$$

then also

$$\int_{\mathbb{R}^d\times\mathbb{R}} (\partial_1\ell)(f_\pi(x),y)\, f^*(x)\,\mathbb{P}(\mathrm{d}x\otimes\mathrm{d}y) \to 0$$

The smooth loss conditions have interpretations as follows. (SL1) shows that there exist measures $\pi \in \mathbb{P}$ which have finite loss and that f^* can be suitably approximated by f_π in risk, i.e. there exists no 'Lavrentiev gap' for this risk functional. (SL2) is achieved in the previous case by the fact that $\partial_1\ell \in L^\infty$ and $f^*, \phi(\theta,\cdot)$ in L^1. In the examples below, this will be extended to more general $L^p/L^{p'}$ dual pairings. (SL3) is a density result for the functions $\phi(\theta,\cdot)$ similar to (SL1), but for dual pairings instead of energies.

Theorem 3.10. *Assume the conditions (L1), (L2'), (L3), (P1), (P2'), (P3), (P4), (LP1), (LP2), (IC1), (IC2), (M-S), (SL1), (SL2) and (SL3). Then*

(1) Ω^{lim} *is not empty and*

(2) $\lim_{t\to\infty} \mathcal{R}(\pi_t) = \inf_\pi \mathcal{R}(\pi)$ *if and only if* Ω^{lim} *consists of only one element.*

If Ω^{lim} *consists of only one element, then that element is* $\lim_{t\to\infty}(\delta_\pi\mathcal{R})(\pi_t;\theta) \equiv 0$.

Corollaries 3.7 and 3.7 remain true also in this case, and Theorem 3.6 generalizes in the same way.

4. Examples

We imposed a number of abstract conditions on the loss function, data distribution, and initial condition. In this section we consider concrete situations where the conditions are met.

Example 4.1 (Huber-type loss functions). Loss functions like Huber loss

$$\ell_H(y,y') = \begin{cases} \frac{1}{2}\,|y-y'|^2 & \text{if } |y-y'|<1 \\ |y-y'| - \frac{1}{2} & \text{else} \end{cases}$$

and pseudo-Huber loss

$$\ell_{psH}(y,y') = \sqrt{|y-y'|^2+1} - 1$$

are Lipschitz-continuous and have Lipschitz continuous first derivatives. They grow at infinity for any fixed $y' \in \mathbb{R}$, so for distributions $\mathbb{P}$ with finite first moments, they are well-defined and coercive. Pseudo-Huber loss is strictly convex, so f^* is uniquely defined. Huber loss is only strictly convex around its minimum and not strictly convex away from it, so if

$$\mathbb{P}^x = \delta_{y=-2} + \delta_{y=2},$$

then any $f^*(x) \in [-1, 1]$ is an admissible minimizer. More generally, we can observe that Huber loss has a characteristic smoothing length scale. If data is very spotty on a larger length scale, uniqueness of the minimum may not hold.

Huber-type loss functions interpolate between finding the mean (quadratic loss) and median (linear loss) of the conditional distribution $\mathbb{P}^x$. They are less sensitive to outliers than mean squared loss.

Example 4.2 (Binary classification). Consider the soft-plus loss function

$$\ell(y, y') = \log\big(1 + \exp\big(-yy'\big)\big).$$

Clearly ℓ is Lipschitz-continuous and grows linearly if $yy' \to -\infty$. The loss decays to zero exponentially fast on the other hand if $yy' \to \infty$. In the case of binary classification, the data distribution $\mathbb{P}$ is given by a measure of data samples $\bar{\mathbb{P}}$ and marginals $\mathbb{P}^x = \lambda(x)\,\delta_{\{y=1\}} + (1 - \lambda(x))\,\delta_{\{y=-1\}}$ for $\lambda(x) \in [0, 1]$, thus

$$f^*(x) = \operatorname*{argmin}_{\alpha \in [-\infty, \infty]} \lambda(x)\, \log(1 + e^{-\alpha}) + (1 - \lambda(x))\, \log(1 + e^{\alpha}).$$

If there exists a minimal uncertainty $0 < \lambda_{min} \leq \lambda(x)$ in the prediction of labels for $\overline{\mathbb{P}}$-almost every x, then there is a unique minimizing $\alpha \in \mathbb{R}$ which satisfies

$$\begin{aligned} 0 &= \frac{d}{d\alpha}\big[\lambda\, \log(1 + e^{-\alpha}) + (1 - \lambda)\, \log(1 + e^{\alpha})\big] \\ &= \frac{-\lambda\, e^{-\alpha}}{1 + e^{-\alpha}} + (1 - \lambda)\frac{e^{\alpha}}{1 + e^{\alpha}} \\ &= (1 - \lambda - \lambda\, e^{-\alpha})\, \frac{e^{\alpha}}{1 + e^{\alpha}}. \end{aligned}$$

Hence

$$e^{-\alpha} = \frac{1 - \lambda}{\lambda}, \qquad \text{i.e.} \quad f^*(x) = \alpha(\lambda(x)) = \log\left(\frac{\lambda(x)}{1 - \lambda(x)}\right)$$

is (essentially) uniformly bounded in the presence of minimal uncertainty and continuous in x if λ is, while for perfectly predictable labels (i.e. when $\lambda \in \{0, 1\}$ $\overline{\mathbb{P}}$-almost surely)

$$f^* = \infty \cdot \big(1_{\{\lambda=1\}} - 1_{\{\lambda=0\}}\big) = (2\lambda - 1) \cdot \infty$$

is almost surely not finite-valued. Our results apply directly in the first situation, but not the second. In this case, the minimizer is simple enough to be understood directly. Namely,

$$\begin{aligned} \mathcal{R}(\pi_t) \to 0 &\Leftrightarrow \log\big(1 + \exp\big(-y\, f_{\pi_t}(x)\big)\big) \to 0 && \mathbb{P} - \text{almost everywhere} \\ &\Leftrightarrow y\, f_{\pi_t}(x) \to \infty && \mathbb{P} - \text{almost everywhere} \\ &\Leftrightarrow \frac{\exp(-y\, f_{\pi_t}(x))}{1 + \exp(-y f_{\pi_t}(x))} \to 0 && \mathbb{P} - \text{almost everywhere} \\ &\Leftrightarrow \delta_\pi \mathcal{R}(\theta) \to 0 \qquad \forall\, \theta \in \Theta \end{aligned}$$

Remark 4.3. If $\mathbb{P} = \frac{1}{N}\sum_{j=1}^{N} \delta_{(x_j, y_j)}$ is an empirical measure with $x_i \neq x_j$ for $i \neq j$, we can achieve perfect classification. The decay of loss only characterizes the behavior $f_\pi \to f^*$ at data points. A more careful analysis shows that classifiers converge to maximum margin solutions [CB20],

which also characterizes the behavior in between data points to be, in some sense, optimal. This corresponds to the observation in the proof or Theorem 3.5 that if the gradient of risk is uniformly small and the excess risk is large, then the second moments N are decreasing.

The model considered in [CB20] is neural network-like with the same homogeneity as ReLU-activated two layer networks. Our results for ReLU activation cannot be applied directly since $\overline{\mathbb{P}}$ must not be an empirical measure. Even assuming existence, the proof of Lemma 3.4 hinges on the analysis of the flow map and requires smoothness. However, very localized measures of the form

$$\mathbb{P} = \sum_{j=1}^{N} \varepsilon^{-(d+1)}\, \eta\left(\frac{\cdot - (x_j, y_j)}{\varepsilon}\right) \cdot \mathcal{L}^{d+1}$$

would be expected to lead to similar behavior. Here η is any compactly supported probability density, $\varepsilon > 0$ is a small parameter and $\mathcal{L}$ denotes Lebesgue measure. Empirical measures are recovered in the singular limit $\varepsilon \to 0$.

Example 4.4. Huber-type and softplus loss allow for data distributions $\mathbb{P}$ with bounded first moments. Under the (much stronger) assumption that the support of the data distribution $\mathbb{P}$ is compact, algebraic loss functions

$$\ell_p(y, y') = \frac{|y - y'|^p}{p}, \qquad p \in (1, \infty)$$

are admissible. Conditions (SL1) and (SL3) are met since functions of the form f_π are uniformly dense in the space of continuous functions on compact sets (Universal Approximation Theorem) and continuous functions are dense in Lebesgue spaces $L^q(\overline{\mathbb{P}})$ for any Radon probability measure on $\mathbb{R}^d$ due [FL07, Theorem 2.11].

(SL2) is met since $f^* \in L^p(\overline{\mathbb{P}})$ due to the boundedness assumption on the support of $\mathbb{P}$ and the growth condition on ℓ_p. Observe furthermore that $f_\pi, \phi(\theta, \cdot) \in L^\infty(\mathbb{P})$ for fixed $\pi \in \mathcal{P}$.

Example 4.5 (Data distributions). Using Theorem 2.1, many data distributions of practical importance are admissible for Lipschitz loss, including all distributions which have a bounded density with respect to Lebesgue measure and decay suitably fast at infinity, e.g.

(1) distributions with continuous density on a bounded open set or
(2) Gaussian mixture models with uniformly bounded means and variances.

For general loss functions, the first class is admissible, while Gaussian mixture models are excluded for purely technical reasons.

The question which lower-dimensional objects have admissible geometry remains open at this point.

Example 4.6 (Initial condition). Initial conditions need to be omni-directional and satisfy the scaling property

$$|a|^2 \leq |w|_{\ell^2}^2 + |b^2|.$$

The easiest choice of admissible condition is to independently choose (w, b) uniformly distributed on S^d and a uniformly distributed on $[-1, 1]$. In applications, a popular choice is to initialize (w, b) according to a standard normal distribution with mean zero and unit variance. It is easily computed that

$$\mathbb{E}\big[|w|^2 + b^2\big] = d + 1$$

and due to [Ver18, Theorem 3.1.1], the norm of (w, b) *concentrates* close to $\sqrt{d}$ in the sense that

$$\mathbb{P}_{\mathbb{N}}\big(\big\{(w, b) : \big|\sqrt{|w|^2 + b^2} - \sqrt{d}\big| > \varepsilon\big\}\big) \leq 2\exp\left(-c\, d\, \varepsilon\right)$$

for a universal constant $c > 0$. Thus if a is distributed on a domain sufficiently smaller than $O(\sqrt{d})$ and d is reasonably large with respect to the network width, then with high probability π_0 satisfies (IC1).

5. Conclusion

We have shown that the convergence of gradient descent training with infinitesimal step size for two-layer networks with ReLU or leaky ReLU activation starting at omni-directional initial conditions is equivalent to the convergence of the velocity potential to a unique limit (under certain technical conditions). The result holds for a fairly general class of loss functions and data distributions. Convergence along subsequences is guaranteed by compactness.

A number of questions remain open.

- We have shown that convergence to minimal Bayes risk is equivalent to the convergence of the velocity potentials $\delta_\pi \mathcal{R}$ to a unique limit (for suitable initial conditions). Whether the limiting potential is generally unique remains one of the most relevant open questions in theoretical machine learning.
- To prove existence and make use of the flow map representation, we are restricted to population risk for suitable data distributions. Especially for the case of low-dimensional data in high-dimensional spaces, the regularity condition on the data manifold is hard to understand and check.

  Even if a data distribution $\mathbb{P}$ is admissible and the gradient flows of empirical risk functionals sampled from $\mathbb{P}$ exist, it is not obvious if they approach the gradient flow of population risk. Convergence of gradient flows is known to hold under certain (usually hard to check) conditions [Ser11, SS04], but there are non-convergence results in applications of practical importance in other fields [Mie12, DKW19].
- It remains open whether a Morse-Sard type property like (M-S) generally holds in higher dimension (possibly under additional assumptions), or whether it can be eliminated from the proof of Lemma 3.4. The property is guaranteed to hold by the Morse-Sard theorem for sub-analytic functions [BDL06] if the measure $\mathbb{P}$ is a finite sum of point-masses. Unfortunately, highly concentrated measures are at odds with the regularity assumption (P4) which is required in the ReLU setting.
- Even if a risk-minimizing measure π exists and risk decays to its minimum, it is unclear whether
  (1) $\mathcal{R}(\pi_t)$ decays at a rate,
  (2) the second moments $N(\pi_t)$ remain bounded, and
  (3) the measures π_t converge to a minimizer weakly or in 2-Wasserstein distance (assuming that $N(\pi_t)$ remains bounded).
- State-of-the-art neural network architectures can have hundreds or even thousands of layers, far from the two-layer situation considered here. In [CB18b], the authors consider also the case of smooth bounded activation functions which are linear in one direction (and sufficiently smooth). These results apply to network architectures in which every node in the outermost layer is given its own set of parameters for deeper layers. Other models for mean field training of deep networks [AOY19, Ngu19, NP20, SS19] are very different from models for shallow networks. No analogous result is available in this setting.

Appendix A. Proofs

In this section we collect the previously omitted proofs.

A.1. **Proofs from Section 2.** We begin with a proof of Theorem 2.1 on admissible data distributions for our analysis. Recall that we say that (P4) holds for $\overline{\mathbb{P}}$ if the map

$$S^d \to L^1\big(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}}\big), \qquad (w,b)\mapsto 1_{\{x|w^Tx+b>0\}}$$

is Lipschitz continuous.

Proof of Theorem 2.1. Consider an open half-space H in $\mathbb{R}^d$. Since ∂H is convex, there exists a closest point $x_H\in\partial H$ to the origin in $\mathbb{R}^d$. If $0\notin\partial H$, we can characterize

$$H=\begin{cases}\{x\in\mathbb{R}^d\mid\langle x,x_H\rangle>|x_H|^2\} & 0\notin H\\ \{x\in\mathbb{R}^d\mid\langle x,x_H\rangle<|x_H|^2\} & 0\in H\end{cases}.$$

Thus H is represented by as $H=\{x|w^Tx+b>0\}$ for

$$(w,b)\in\mathbb{R}^d, \qquad (w,b)=\lambda\,(\pm x_H,|x_H|^2), \qquad \lambda>0.$$

In particular, there is a canonical representative in the unit sphere

$$(w,b)=\left(\pm\frac{x_H}{\sqrt{|x_H|^2+|x_H|^4}},\ \frac{|x_H|^2}{\sqrt{|x_H|^2+|x_H|^4}}\right).$$

For the derivative estimates below, we can assume without loss of generality that $b\notin\{0,\pm1\}$, as the Lipschitz estimate extends to the boundary points by uniform continuity. Given a density ρ, denote

$$\tilde\rho(x)=\rho(x)\sqrt{1+|x|^2}, \qquad \tilde\rho^+(x)=\limsup_{y\to x}\tilde\rho(y)$$

and observe that $\tilde\rho^+$ satisfies the same decay estimate as $\sqrt{1+|x|^2}\cdot\rho$ and is integrable with respect to Lebesgue measure. $\tilde\rho^+$ is upper semi-continuous.

We prove the first claim. Let $(w,b)\in S^d$. Without loss of generality, $w=\sqrt{1-b^2}\,e_1$. Here, we can even take $\varepsilon=0$ in the decay condition and compute

$$\begin{aligned}
&\limsup_{h\to0^+}\left\|\frac{1_{\{x|w^Tx+b+h>0\}}-1_{\{x|w^Tx+b>0\}}}{h}\right\|_{L^1(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}})}\\
&\quad=\limsup_{h\to0^+}\frac1h\int_{\left\{x\,\middle|\,-(b+h)<w^Tx<-b\right\}}\tilde\rho(x)\,\mathrm{d}x\\
&\quad=\limsup_{h\to0^+}\frac1h\int_0^{\frac{h}{\sqrt{1-b^2}}}\int_{\{x_1=-b/\sqrt{1-b^2}\}}\tilde\rho(x-te_1)\,\mathcal{H}^{d-1}(\mathrm{d}x)\,\mathrm{d}t\\
&\quad=\frac{1}{\sqrt{1-b^2}}\limsup_{h\to0}\int_0^1\int_{\{x_1=-b/\sqrt{b}\}}\tilde\rho\left(x-s\,\frac{h}{\sqrt{1-b^2}}e_1\right)\mathcal{H}^{d-1}(\mathrm{d}x)\,\mathrm{d}s\\
&\quad\le\frac{1}{\sqrt{1-b^2}}\int_{\{x_1=-b/\sqrt{1-b^2}\}}\tilde\rho^+(x)\,\mathcal{H}^{d-1}(\mathrm{d}x)\\
&\quad\le\frac{1}{\sqrt{1-b^2}}\int_{\mathbb{R}^{d-1}}\frac{C}{\left(1+|x|^2+\frac{b^2}{1-b^2}\right)^{\frac{d+1}{2}}}\,\mathrm{d}x \qquad \text{(A.1)}\\
&\quad=\frac{\sqrt{1-b^2}}{b^2}\int_{\mathbb{R}^{d-1}}\frac{C}{\left(\frac{1-b^2}{b^2}+\left|\frac{\sqrt{1-b^2}\,x}{b}\right|^2+1\right)^{\frac{d+1}{2}}}\left(\frac{\sqrt{1-b^2}}{b}\right)^d\mathrm{d}x\\
&\quad\le\frac{\sqrt{1-b^2}}{b^2}\int_{\mathbb{R}^{d-1}}\frac{1}{\left(1+|y|^2\right)^{\frac{d+1}{2}}}\,\mathrm{d}y \qquad \text{(A.2)}
\end{aligned}$$

This is bounded from above uniformly when b is bounded away from ± 1 by (A.1) and close to $b = \pm 1$ by (A.2). The passage to the limit can be justified using Fatou's lemma. A similar estimate holds for the limit $h \to 0^-$, and the derivative in direction of increasing/decreasing w radially is the same as that in direction of changing b.

Now consider $v \in \mathbb{R}^d$ such that $\langle v, w\rangle = 0$. Taking the derivative in v, the half-spaces are no longer aligned and rather than a slab, their symmetric difference is the union of two wedges. Again without loss of generality, we assume that $w = \sqrt{1-b^2}\, e_1$ and $v = e_2$. Then

$$\begin{aligned}
&\left\{x : \left|1_{\{w^Tx+b>0\}} - 1_{\{(w+hv)^Tx+b>0\}}\right| = 1\right\}\\
&\quad = \{w^Tx + b > 0\} \,\Delta\, \{(w+hv)^Tx + b > 0\}\\
&\quad = \left\{x : \sqrt{1-b^2}\, x_1 + hx_2 < -b < \sqrt{1-b^2}\, x_1\right\} \cup \left\{x : \sqrt{1-b^2}\, x_1 < -b < \sqrt{1-b^2}\, x_1 + hx_2\right\}\\
&\quad = \left\{x : -\frac{b}{\sqrt{1-b^2}} < x_1 < -\frac{b+hx_2}{\sqrt{1-b^2}}\right\} \cup \left\{x : -\frac{b+hx^2}{\sqrt{1-b^2}} < x_1 < -\frac{b}{\sqrt{1-b^2}}\right\}\\
&\quad \subset \left\{x : \frac{-b-|hx_2|}{\sqrt{1-b^2}} < x_1 < \frac{-b+|hx_2|}{\sqrt{1-b^2}}\right\}
\end{aligned}$$

Then we can therefore bound the tangential derivative as follows.

$$\begin{aligned}
&\limsup_{h\to 0^+} \frac{\|1_{\{x\,|\,(w+hv)^Tx+b>0\}} - 1_{\{x|w^Tx+b>0\}}\|_{L^1(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}})}}{h}\\
&\qquad \le \limsup_{h\to 0} \frac{1}{|h|} \int_{\mathbb{R}^{d-1}} \int_{\frac{-b-|hx_2|}{\sqrt{1-b^2}}}^{\frac{-b+|hx_2|}{\sqrt{1-b^2}}} \tilde{\rho}(y_1, x_2, \dots, x_d)\, \mathrm{d}y_1\, \mathrm{d}x_2 \dots \mathrm{d}x_d\\
&\qquad = \limsup_{h\to 0} \int_{\mathbb{R}^{d-1}} \int_{-1}^{1} \frac{|x|_2}{\sqrt{1-b^2}}\, \tilde{\rho}\left(-\frac{b}{\sqrt{1-b^2}} + \frac{th|x_2|}{\sqrt{1-b^2}}, x_2, \dots, x_d\right) \mathrm{d}t\, \mathrm{d}x_2 \dots \mathrm{d}x_d\\
&\qquad \le \frac{1}{\sqrt{1-b^2}} \int_{\mathbb{R}^{d-1}} |x_2|\, \tilde{\rho}^+\left(-\frac{b}{\sqrt{1-b^2}}, x_2, \dots, x_d\right)\, \mathrm{d}x_2 \dots \mathrm{d}x_d.
\end{aligned}$$

The remainder of the proof proceeds as above, except for the weight of $|x_2|$ in front of the density. While previously a decay as $|x|^{-(d+1+\varepsilon)}$ for $\varepsilon > 0$ would have been sufficient, here we need decay as $|x|^{-(d+2+\varepsilon)}$ at infinity.

The proof of the second statement is similar to the first. The third statement is immediate from the definition. To consider the fourth statement, let $\mathcal{M}$ be the set of Radon probability measures

$$\mathcal{M} = \left\{\mathbb{P} \,:\, \left\|1_{\{w^T\cdot+b>0\}} - 1_{\{\tilde{w}^T\cdot+\tilde{b}>0\}}\right\|_{L^1(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}})} \le L\big[|w-\tilde{w}| + |b-\tilde{b}|\big]\right\}.$$

Assume that $\mathbb{P}_n$ is a sequence in $\mathcal{M}$ and $\mathbb{P}$ is a probability measure such that $\mathbb{P}_n \to \mathbb{P}$ in the sense of Radon measures. Then

$$\mathbb{P}(U) \le \liminf_{n\to\infty} \mathbb{P}_n(U)$$

for all open sets U. In particular, the symmetric differences of open sets are open. This establishes the Lipschitz condition by the previous analysis since

$$\|1_H - 1_{H'}\|_{L^1(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}}_n)} = \big(\sqrt{1+|x|^2}\cdot\overline{\mathbb{P}}_n\big)(H\Delta H') \le L\big[|w-w'| + |b-b'|\big]$$

where H, H' are the half-spaces

$$H = \{x : w^Tx + b > 0\}, \qquad H' = \{x : (w')^Tx + b' > 0\}.$$

□

In particular, sums of Gaussian distributions or compactly supported regular distributions as they occur in density estimation are admissible data distributions for our purposes. They can be computed from a given finite data sample and mollify the problem sufficiently for our convergence result.

Many full-dimensional data distributions satisfy (P4), but distributions with a bounded density on the hypersphere is admissible. If data is concentrated on a manifold of dimension k, the intuition is that k should not have any 'straight' segments which lie mostly in a lower-dimensional affine subspace.

Remark A.1. It is easy to see that no measure $\overline{\mathbb{P}}$ such that $\overline{\mathbb{P}}(\partial H) > 0$ for any half-space H satisfies (P4). In fact, the map $(w, b) \mapsto 1_{\{w^T x + b > 0\}}$ is not even continuous for such distributions. This in particular excludes empirical measures, but also lower-dimensional data manifolds which are entirely contained in an affine subspace.

A complete characterization of admissible measures is beyond the scope of this article. We move on to the proof of the measurability of f^*.

Proof of Lemma 2.3. **First claim.** Note that for every $\alpha \in \mathbb{R}$, the map $x \mapsto L_x(\alpha)$ is $\overline{\mathbb{P}}$-measurable by the construction of the marginal measures in [ABM14, Theorem 4.2.4]. Furthermore, $L_x(\alpha) < \infty$ for $\overline{\mathbb{P}}$-almost every x and all $\alpha \in \mathbb{R}$ since $L_x(0)$ is integrable and L is Lipschitz continuous in α (with the same Lipschitz constant as ℓ).

Hence the function $(x, \alpha) \mapsto L_x(\alpha)$ is a Caratheodory integrand (finite, continuous in α and measurable in x) where α belongs to a second countable complete space. It is well-known that such functions are jointly measurable, see [AB94, Theorem 14.75]. *Sketch of proof:* Define $U = \{x : L_x(0) < \infty\}$ and the sequence of functions

$$\Psi_k : \mathbb{R} \times U \to [0, \infty), \qquad \Psi_k(\alpha, x) = L_x\left(k \cdot \left\lfloor \frac{\alpha}{k} \right\rfloor\right)$$

which are product measurable. Since L_x is continuous in α, $L_x(\alpha) = \lim_{k\to\infty} \Psi_k(\alpha, x)$ and thus $(x, \alpha) \mapsto L_x(\alpha)$ is product measurable on $\mathbb{R} \times U$. Since $\overline{\mathbb{P}}(\mathbb{R}^d \setminus U) = 0$, we find that the map is product measurable on the entire space. Note that we establish product-measurability with respect to the Borel σ-algebra on $\mathbb{R}$ and the σ-algebra on $\mathbb{R}^d$ generated by $\overline{\mathbb{P}}$, which contains all $\overline{\mathbb{P}}$-null sets and is in general larger than the Borel σ-algebra.

Second claim. To find f^* we use the Kuratowski-Ryll-Nardzewski Selection Theorem [AB94, Theorem 14.86] which states that if the (possibly multi-valued) map

$$M : U \to 2^{\mathbb{R}}, \qquad M(x) = \operatorname*{argmin}_{\alpha} L_x(\alpha)$$

is a weakly measurable correspondence (see [AB94, Chapter 14]) with nonempty closed values in a Polish space, then it admits a measurable selector. The only non-trivial fact is the weak measurability of M, which means that we need to check that the set

$$\{x \in U \mid \phi(x) \cap V \neq \emptyset\} = \left\{ x \in U \;\middle|\; \exists\ \alpha \in V \text{ s.t. } L_x(\alpha) = \inf_{\alpha' \in \mathbb{R}} L_x(\alpha') \right\}$$

is measurable in U whenever V is open. Any open set V admits a countable dense subset $\widehat{V}$, which means that

$$m_V : U \to [0, \infty), \quad m(x) = \inf_{\alpha \in V} L_x(\alpha) = \inf_{\alpha \in \widehat{V}} L_x(\alpha)$$

is measurable due to the continuity of L_x in α. We conclude that

$$\{x \in U \mid M(x) \cap V \neq \emptyset\} = \{x \in U \mid m_V(x) = m_{\mathbb{R}}(x)\}$$

is measurable. □

For the reader's convenience, we link Wasserstein gradient flows to classical gradient flows.

Proof from Section 2.3. Let $\Theta = (\theta_1, \dots, \theta_m)$ and $f_\Theta(x) = \frac{1}{m}\sum_{i=1}^m \phi(\theta_i; x)$. Then

$$\begin{aligned}
\nabla_{\theta_i} \mathcal{R}(\Theta) &= \nabla_{\theta_i} \int_\Theta \ell\big(f_\Theta(x), y\big)\, \mathbb{P}(\mathrm{d}x) \\
&= \int_\Theta (\partial_1 \ell)\big(f_\Theta(x), y\big)\, \nabla_{\theta_i} \left(\frac{1}{m} \sum_{i=1}^m \phi(\theta_i; x) \right) \mathbb{P}(\mathrm{d}x) \\
&= \frac{1}{m} \int_\Theta (\partial_1 \ell)\big(f_\Theta(x), y\big)\, (\nabla_\theta \phi)(\theta_i; x)\, \mathbb{P}(\,\mathrm{d}x) \\
&= \frac{1}{m} \nabla(\delta_\pi \mathcal{R}) \left(\frac{1}{m} \sum_{j=1}^m \delta_{\theta_j}; \theta_i \right).
\end{aligned}$$

Thus if the parameters θ_i evolve by the law $\dot\theta_i = -m\, \nabla_{\theta_i} \mathcal{R}(\Theta)$ for all $i = 1, \dots, m$, then their distribution $\pi^m := \frac{1}{m}\sum_{i=1}^m \delta_{\theta_i}$ satisfies the transport equation

$$\dot\pi_t^m = \operatorname{div}\big(\pi_t^m\, \nabla(\delta_\pi \mathcal{R})(\pi_t^m; \cdot)\big)$$

by the flow map representation. This is precisely the PDE formulation of the 2-Wasserstein gradient flow. □

Now we show that gradient flow of $\mathcal{R}$ exists for any initial condition $\pi_0 \in \mathcal{P}_2(\overline{\Theta})$ and that the omni-directionality of the initial measure is preserved along the gradient flow evolution (for finite time). Except for technical issues stemming from the lack of regularity in ReLU activation, the analysis follows [CB18b, Appendix B].

For technical purposes, it is necessary to consider ReLU activation on the whole space and not just the (non-convex) cone Θ. The natural extension which is linear in a faces the aforementioned issues of non-differentiability. We therefore define ϕ on $\mathbb{R}^{d+2}$ as

$$\phi(a, w, b; x) = \eta\left(\frac{a^2 - |w|^2 - b^2}{a^2 + |w|^2 + b^2} \right) a\, \sigma\left(w^T x + b\right)$$

where

$$\eta(z) = \begin{cases} 1 & z \leq 0 \\ 0 & z \geq \frac{1}{2} \end{cases}, \qquad \eta' \leq 0$$

is a smooth cut-off function. In particular, note that

$$\phi(a, w, b; x) = a\, \sigma(w^T x + b)$$

for all $(a, w, b) \in \Theta$.

Lemma A.2. *Consider $\phi : \mathbb{R}^{d+2} \times \mathbb{R}^d \to \mathbb{R}$ as above. Then*

(1) For any $\rho \in L^\infty(\mathbb{P})$, the vector field

$$V_\rho : \mathbb{R}^{d+2} \to \mathbb{R}, \quad V_\rho(\theta) = \int_{\mathbb{R}^d} \rho(x, y)\, \nabla_\theta \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

is Lipschitz continuous with Lipschitz constant $\|\rho\|_{L^\infty}$.

(2) For any $\theta \in \partial\Theta$ and any $\rho \in L^\infty(\mathbb{P})$, V_ρ is tangent to $\partial\Theta$ at θ.

Proof. **Lipschitz-regularity.** We observe that

$$\nabla_\theta \phi(\theta; x) = \eta'\left(\frac{a^2 - |w|^2 - b^2}{a^2 + |w|^2 + b^2} \right) a\, \sigma(w^T x + b)$$

$$\left\{ \frac{2}{a^2+|w|^2+b^2} \begin{pmatrix} a \\ -w \\ -b \end{pmatrix} - 2\frac{a^2-|w|^2-b^2}{\left[a^2+|w|^2+b^2\right]^2} \begin{pmatrix} a \\ w \\ b \end{pmatrix} \right\}$$
$$+ \eta \left(\frac{a^2-|w|^2-b^2}{a^2+|w|^2+b^2} \right) \begin{pmatrix} \sigma(w^T x+b) \\ a\, 1_{\{w^T x+b>0\}} x \\ a\, 1_{\{w^t x+b>0\}} \end{pmatrix}$$
$$= \sigma(w^T x+b)\, V_1(a,w,b) + \eta \left(\frac{a^2-|w|^2-b^2}{a^2+|w|^2+b^2} \right) a\, 1_{\{w^T x+b>0\}} \begin{pmatrix} 0 \\ x \\ 1 \end{pmatrix}$$

where V_1 is a positively 0-homogeneous vector field which is smooth on the sphere $a^2+|w|^2+b^2$. Thus for any $x \in \mathbb{R}^d$, the product $\sigma(w^T x+b) \cdot V_1$ is Lipschitz-continuous on the sphere with Lipschitz-constant $\leq C\sqrt{1+|x|^2}$ and positively one-homogeneous. The triangle-inequality implies that the term is $C\sqrt{1+|x|^2}$-Lipschitz on the whole space. Since the first moments of $\mathbb{P}$ are finite, we deduce that

$$(a,w,b) \mapsto \int_{\mathbb{R}^d \times \mathbb{R}} \rho(x,y)\, \sigma(w^T x+b)\, V_1(a,w,b)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

is $C\,\|\rho\|_{L^\infty(\mathbb{P})}$-Lipschitz continuous. Estimating the second term requires higher regularity of $\mathbb{P}$ as expressed in condition (P4). We note that if $\eta > 0$, then

$$\frac{a^2-|w|^2-b^2}{a^2+|w|^2+b^2} \leq \frac{1}{2} \qquad \Rightarrow \quad a^2-|w|^2-b^2 \leq \frac{a^2+|w|^2+b^2}{2} \qquad \Rightarrow \quad a^2 \leq 3\left[|w|^2+b^2\right].$$

For any fixed (a,w,b), $(\tilde{a},\tilde{w},\tilde{b})$, denote by $\eta, \tilde{\eta}$ the cut-off function evaluated at the respective point. Without loss of generality, we assume that $|(\tilde{w},\tilde{b})| \leq |(w,b)|$ and compute that

$$\left| \int_{\mathbb{R}^d \times \mathbb{R}} \rho(x,y) \eta\, a\, 1_{\{w^T x+b>0\}} \begin{pmatrix} 0 \\ x \\ 1 \end{pmatrix} \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) - \int_{\mathbb{R}^d \times \mathbb{R}} \rho(x,y) \tilde{\eta}\, \tilde{a}\, 1_{\{\tilde{w}^T x+\tilde{b}>0\}} \begin{pmatrix} 0 \\ x \\ 1 \end{pmatrix} \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \right|$$
$$\leq \|\rho\|_{L^\infty(\mathbb{P})} \int_{\mathbb{R}^d \times \mathbb{R}} \left| \eta a\, 1_{\{w^T x+b>0\}} - \tilde{\eta}\tilde{a}\, 1_{\{\tilde{w}^T x+\tilde{b}>0\}} \right| \cdot \sqrt{1+|x|^2}\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$
$$\leq \|\rho\|_{L^\infty} \int_{\mathbb{R}^d} \left\{ |\eta a - \tilde{\eta}\tilde{a}|\, 1_{\{w^T x+b>0\}} + |\tilde{\eta}\tilde{a}| \left| 1_{\{w^T x+b>0\}} - 1_{\{w^T x+b>0\}} \right| \right\} \sqrt{1+|x|^2}\, \overline{\mathbb{P}}(\mathrm{d}x)$$
$$\leq \|\rho\|_{L^\infty} \left\{ |\eta a - \tilde{\eta}\tilde{a}| \int_{\mathbb{R}^d} \sqrt{1+|x|^2}\, \overline{\mathbb{P}}(\mathrm{d}x) + C|\tilde{\eta}\tilde{a}| \left| \frac{(w,b)}{\sqrt{|w|^2+b^2}} - \frac{(\tilde{w},\tilde{b})}{\sqrt{|\tilde{w}|^2+\tilde{b}^2}} \right| \right\}$$
$$\leq C\,\|\rho\|_{L^\infty} \left[|\eta a - \tilde{\eta}\tilde{a}| + \left| \frac{|(\tilde{w},\tilde{b})|}{|(w,b)|}(w,b) - (\tilde{w},\tilde{b}) \right| \right]$$

The first term on the right is bounded by

$$|\eta a - \tilde{\eta}\tilde{a}| \leq C\,\|(a,w,b) - (\tilde{a},\tilde{w},\tilde{b})\|$$

since $(a,w,b) \mapsto a\,\eta(a,w,b)$ is a positiively one-homogeneous function which is Lipschitz-continuous on the sphere and thus Lipschitz continuous. The second term on the right is

$$\left| \frac{|(\tilde{w},\tilde{b})|}{|(w,b)|}(w,b) - (\tilde{w},\tilde{b}) \right| = \left| P^{S_r^{d+1}}(w,b) - P^{S_r^{d+1}}(\tilde{w},\tilde{b}) \right|$$

where $P^{S_r^{d+1}}$ denotes the projection onto the sphere of radius $r = |(\tilde{w},\tilde{b})|$ around the origin. Since both (w,b) and $(\tilde{w},\tilde{b})$ lie on the exterior of the sphere, this map is one-Lipschitz.

Tangency condition. We can write $\Theta = \{(a, w, b) \mid -a^2 + |w|^2 + b^2 > 0\}$. Then the normal to $\partial\Theta$ is parallel to $\nabla(-a^2 + |w|^2 + b^2) = 2(-a, w, b)$. Fix (a, w, b) and consider any x such that $w^T x + b \neq 0$. Since $\eta \equiv 1$ close to $\partial\Theta$, we have

$$\begin{aligned} 2|\theta| \, \langle \nabla_\theta \phi(\theta; x), \nu_{\partial\Theta} \rangle &= \begin{pmatrix} \sigma(w^T x + b) \\ a \, 1_{\{w^T x + b > 0\}} x \\ a \, 1_{\{w^t x + b > 0\}} \end{pmatrix} \cdot \begin{pmatrix} -a \\ w \\ b \end{pmatrix} \\ &= -a \, \sigma(w^T x + b) + a \, \sigma'(w^T x + b) \, (w^T x + b) \\ &= 0 \end{aligned}$$

since σ is positively one-homogeneous. The equality also holds trivially at x for which $w^T x + b = 0$, so in particular after integration. □

In the calculus of variations (which encompasses the study of gradient flows), different notions of convexity play a key role. In vector spaces, convexity is a condition along straight lines, which (at least for Hilbert spaces) are the same as length-minimizing curves. The natural generalization to (geodesically complete) metric spaces is to consider the notion of convexity where a functional is 'convex' if it is convex along constant-speed length minimizing geodesics. The analogy is particularly strong in 2-Wasserstein space, which carries the formal structure of a Hilbert manifold, see [Ott01] or [Vil08, Chapter 15]. This concept of convexity is referred to as *displacement convexity* and has been recognized since [McC97] as a useful notion when considering gradient flows.

The functionals we consider are not convex in Wasserstein space – in fact, convergence to a global minimizer is not guaranteed. For the existence of gradient flows, a weaker concept suffices. Recall that the functional $\mathcal{R}$ is called λ-displacement convex, if the following holds: If π_t is a geodesic in Wasserstein space, then

$$\mathcal{R}(\pi_t) \leq t \, \mathcal{R}(\pi_0) + (1 - t) \, \mathcal{R}(\pi_1) + \frac{\lambda}{2} \, t(1 - t) \, W_2^2(\pi_0, \pi_1).$$

By analogy with the smooth Euclidean case, we can think of the condition as a lower bound on the Hessian $D^2\mathcal{R} \geq -\lambda I$. Convexity corresponds to 0-convexity and uniform convexity to λ-convexity for $\lambda < 0$.

Proof of Lemma 3.1. **Step 1.** We first show that $\mathcal{R}$ is λ-displacement convex for a suitable $\lambda \in \mathbb{R}$ depending on $\overline{\mathbb{P}}, \ell$ and η. Unlike Wasserstein space $\mathcal{P}_2(\mathbb{R}^{d+2})$, the space of measures on the cone $\mathbb{P}_2(\overline{\Theta})$ not geodesically convex due to the non-convexity of Θ. We therefore use Lemma A.2 to extend ϕ to the whole space $\mathbb{R}^{d+2}$. Without the extension, Wasserstein space is not geodesically convex and we cannot explicitly characterize geodesics. We will later show that if $\mathrm{spt}(\pi_0) \subseteq \overline{\Theta}$, then $\mathrm{spt}(\pi_t) \subseteq \overline{\Theta}$ for all $t \geq 0$.

Let $\pi_0, \pi_1 \in \mathcal{P}_2(\mathbb{R}^{d+2})$. Denote by π_t, $t \in [0, 1]$ any unit speed geodesic between π_0 and π_1, i.e.

$$\pi_t = \left[t \, P^\theta + (1 - t) \, P^{\theta'} \right]_\sharp \gamma$$

where $P^\theta, P^{\theta'}$ are the projection from $\mathbb{R}^{d+2} \times \mathbb{R}^{d+2}$ to the first and second components respectively, $F_\sharp \mu$ denotes the push-forward of the measure μ along the map F, and γ is an optimal transport plan between π_0 and π_1.

It suffices to show that $h(t) := \mathcal{R}(\pi_t)$ is $\lambda \, W_2^2(\pi_0, \pi_1)$-convex on $[0, 1]$, i.e. $h'' \geq -\lambda \, W_2^2(\pi_0, \pi_1)$. We prove a stronger statement, namely that h' is $\lambda \, W_2^2(\pi_0, \pi_1)$-Lipschitz, which can be thought of as a type of Hessian bound from both sides instead of just one side. Note that

$$\frac{d}{dt} f_{\pi_t}(x) = \frac{d}{dt} \int_{\mathbb{R}^{d+2} \times \mathbb{R}^{d+2}} \phi(t\theta + (1 - t)\theta'; x) \, \gamma(\mathrm{d}\theta \otimes \mathrm{d}\theta')$$

$$
\begin{aligned}
&= \int_{\mathbb{R}^{d+2}\times\mathbb{R}^{d+2}} \big\langle (\nabla_\theta \phi)(t\theta + (1-t)\theta'; x),\, \theta - \theta' \big\rangle \, \gamma(\mathrm{d}\theta \otimes \mathrm{d}\theta') \\
h'(t) &= \frac{d}{dt} \mathcal{R}(\pi_t) \\
&= \frac{d}{dt} \int_{\mathbb{R}^d \times \mathbb{R}} \ell\big(f_{\pi_t}(x), y\big) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\
&= \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, \frac{d}{dt} f_{\pi_t}(x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\
&= \int_{\left(\mathbb{R}^{d+2}\right)^2} \left\langle \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}, y\big)(\nabla_\theta \phi)(t\theta + (1-t)\theta'; x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y),\, \theta - \theta' \right\rangle \gamma(\mathrm{d}\theta \otimes \mathrm{d}\theta').
\end{aligned}
$$

Due to Lemma A.2, the map

$$
\theta \mapsto \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)(f_\pi(x), y) \, (\nabla_\theta \phi)(\theta, x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)
$$

is $L \, \|\partial_1 \ell\|_{L^\infty(\mathbb{P})}$-Lipschitz, and $\|\partial_1 \ell\|_{L^\infty}$ is bounded uniformly by the Lipschitz-condition on the loss function. Consequently

$$
\begin{aligned}
|h'(t_1) - h'(t_0)| &\le C \int \big| t_1\theta + (1-t_1)\theta' - t_0\theta - (1-t_0)\theta' \big| \, |\theta - \theta'| \, \gamma(\mathrm{d}\theta \otimes \mathrm{d}\theta') \\
&= C \, |t_1 - t_0| \, W_2^2(\pi, \pi').
\end{aligned}
$$

It follows that $\mathcal{R}$ is in fact λ-displacement convex for a universal constant λ which only depends on ϕ, ℓ.

Step 2. Existence of the gradient flow follows directly from [AGS08, Theorem 11.2.1].

Step 3. In this step, we show that the gradient flow of $\mathcal{R}$ preserves the cone Θ, i.e. if $\pi_0 \in \mathcal{P}_2(\overline{\Theta})$, then $\pi_t \in \mathcal{P}_2(\overline{\Theta})$ for all $t \ge 0$. Note that the flow field

$$
(t, \theta) \mapsto -\nabla(\delta_\pi \mathcal{R})\big(\pi_t; \theta\big)
$$

is Lipschitz-continuous in θ with a uniform Lipschitz constant for all times by Lemma A.2. Like in Section 2.4, we find that $\pi_t = X(t, \cdot)_\sharp \pi_0$ where X is the flow map defined by

$$
\begin{cases} \frac{d}{dt} X(t, \theta) &= -\nabla(\delta_\pi \mathcal{R})\big(\pi_t; X(t, \theta)\big) \quad t > 0 \\ X(0, \theta) &= \theta \end{cases}.
$$

It thus suffices to show that $X(t, \theta) \in \Theta$ for every $\theta \in \Theta$. This is immediate since

$$
\begin{aligned}
\dot{X}(t, \theta) &= -\nabla(\delta_\pi \mathcal{R})(\pi_t; X(t, \theta)) \\
&= -\int_\Theta (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, (\nabla_\theta \phi)(\theta, x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)
\end{aligned}
$$

is parallel to $\partial\Theta$ on the boundary since $\nabla_\theta \phi(\theta, x)$ is parallel to $\partial\Theta$ at θ whenever it is defined – see also Section 2.5. □

We now prove that the gradient flow preserves the omni-directionality of measures (for finite positive time).

Proof of Lemma 3.2. **Preliminary analysis.** Again, we use the flow map X. Since $[\nabla(\delta_\pi \mathcal{R})]_{\mathrm{Lip}} \le C$, we deduce from Lemma [Amb08, Lemma 4] that X satisfies

$$
[X(t, \cdot)]_{\mathrm{Lip}} \le \exp\left(\int_0^t C \, \mathrm{d}s \right) \le \exp\big(Ct\big).
$$

Since we can solve the ordinary differential equation backwards in time for any $\theta \in \Theta$, $X(t, \cdot) : \Theta \to \Theta$ is a bi-Lipschitz homeomorphism.

Finally, we note that $X(t,\lambda\theta) = \lambda X(t,\theta)$ for all $\lambda, t > 0$ since

$$\frac{d}{dt}\lambda\, X(t,\theta) = -\lambda\, \nabla(\delta_\pi \mathcal{R})(\pi_t;\theta) = -\nabla(\delta_\pi\mathcal{R})(\pi_t;\lambda\theta)$$

due to the homogeneity of ϕ. Thus, the flow preserves half-rays and cones.

Proof of omni-directionality. Consider an open cone $C \subseteq \Theta$. Then by [Amb08, Lemma 4], we have

$$\pi_t(C) = \int_\Theta 1_C(X(t,\theta))\,\pi_0(\mathrm{d}\theta) = \int_\Theta 1_{X(t,\cdot)^{-1}(C)}(\theta)\,\pi_0(\mathrm{d}\theta) = \pi_0\big(X(t,\cdot)^{-1}(C)\big) > 0$$

since also $X(t,\cdot)^{-1}(C)$ is an open cone in Θ. □

Remark A.3. The analysis of the flow map shows more. Since rays are preserved, the projected measures $P^{S^{d+1}}_\sharp(\pi_t)$ on the unit sphere $S^{d+1}\cap\Theta$ evolve by the continuity equation

$$\frac{d}{dt}P^{S^{d+1}}(\pi_t) = \operatorname{div}\Big(\Big(P^{S^d}_\sharp \pi_t\Big)\cdot \nabla^{S^d}(\delta_\pi\mathcal{R})(\pi_t,\cdot)\Big)$$

where $\nabla^{S^{d+1}} f(\theta) = (I-\theta\otimes\theta)\nabla f(\theta)$ is the tangential gradient to the unit sphere. In particular, if $P^{S^{d+1}}_\sharp\pi_0$ has a density ρ_0 with respect to the uniform measure on $S^{d+1}\cap\Theta$, then $P^{S^{d+1}}_\sharp\pi_t$ has a density ρ_t and

$$\exp\big(-Ct\big)\inf_{\theta\in\Theta\cap S^{d+1}}\rho_0(\theta) \le \inf_{\theta\in\Theta\cap S^{d+1}}\rho_t(\theta) \le \sup_{\theta\in\Theta\cap S^{d+1}}\rho_t(\theta) \le \exp\big(Ct\big)\big)\sup_{\theta\in\Theta\cap S^{d+1}}\rho_0(\theta).$$

Next we show that the second moment of π_t grows at most sublinearly in time.

Proof of Lemma 3.3. We compute

$$\begin{aligned}
\frac{d}{dt}N(\pi_t) &= \frac{d}{dt}\int_\Theta |\theta|^2\,\pi_t(\mathrm{d}\theta)\\
&= -\int_\Theta \langle\nabla|\theta|^2, \nabla_\theta(\delta_\pi\mathcal{R})\rangle\,\pi_t(\mathrm{d}\theta)\\
&= -2\int_\Theta \langle\theta, \nabla_\theta(\delta_\pi\mathcal{R})\rangle\,\pi_t(\mathrm{d}\theta)\\
&\le 2\left(\int_\Theta |\theta|^2\,\pi_t(\mathrm{d}\theta)\right)^{\frac12}\left(\int_\Theta |\nabla_\theta(\delta_\pi\mathcal{R})|^2\,\pi_t(\mathrm{d}\theta)\right)^{\frac12}\\
&= 2N(\pi_t)^{1/2}\left|\frac{d}{dt}\,\mathcal{R}(\pi_t)\right|^{1/2},
\end{aligned}$$

i.e.

$$\frac{d}{dt}\sqrt{N(\pi_t)} \le \left|\frac{d}{dt}\,\mathcal{R}(\pi_t)\right|^{1/2}.$$

Note that π_t has finite second moments, so the quadratic test function $|\cdot|^2$ in the variational formulation is admissible. If π_0 is compactly supported, then so is π_t for all $t>0$ by the flow map representation and the linear growth of the flow field at infinity. In this situation, the identity is obvious. Otherwise, the argument is easily justified by using approximating test functions $\eta(R-|\theta|)\,|\theta|^2$ where η is a smooth cutoff function satisfying $\eta'\ge 0$, $\eta(z)=0$ for $z\le 0$ and $z=1$ for $z\ge 1$. Thus for every $0<T<t$ we have

$$\sqrt{N(\pi_t)} = \sqrt{N(\pi_T)} + \int_T^t \frac{d}{ds}\sqrt{N(\pi_s)}\,\mathrm{d}s$$

$$\begin{aligned}
&\leq \sqrt{N(\pi_T)} + \int_T^t \left| \frac{d}{ds} \mathcal{R}(\pi_s) \right|^{1/2} \mathrm{d}s \\
&\leq \sqrt{N(\pi_T)} + \left(\int_T^t 1 \,\mathrm{d}s \right)^{\frac{1}{2}} \left(\int_T^t \left| \frac{d}{ds} \mathcal{R}(\pi_s) \right| \mathrm{d}s \right)^{\frac{1}{2}} \\
&= \sqrt{N(\pi_T)} + \sqrt{t-T} \, [\mathcal{R}(\pi_T) - \mathcal{R}(\pi_t)]^{1/2}
\end{aligned}$$

and therefore

$$N(\pi_t) \leq 2\big[N(\pi_T) + (t-T)\,[\mathcal{R}(\pi_T) - \mathcal{R}(\pi_t)]\,\big]. \tag{A.3}$$

Since $\mathcal{R}(\pi_t)$ is monotone decreasing and bounded from below (by zero), $\mathcal{R}(\pi_t)$ converges to a limit. Thus, for every $\varepsilon > 0$ we can choose $T > 0$ such that $0 < \mathcal{R}(\pi_T) - \mathcal{R}(\pi_t) < \varepsilon$ for every $t > T$ and hence

$$\limsup_{t\to\infty} \frac{N(\pi_t)}{t} \leq 2 \limsup_{t\to\infty} [\mathcal{R}(\pi_T) - \mathcal{R}(\pi_t)] \leq 2\varepsilon.$$

□

Note that the proof applies in great generality to models with a linear structure. The next proof concerns the exponential growth of second moments if the velocity potential converges to a non-trivial limit. It is adapted from [CB18b] and repeated in this context for the reader's convenience. The following proof is the only point at which the Morse-Sard property is used in this article. It is also the only argument which hinges on the omni-directionality of the initial parameter distribution.

Proof of Lemma 3.4. If $g \not\equiv 0$, there exists $(a, w, b) \in \mathbb{R}^{d+2}$ such that $g(a, w, b) \neq 0$. Since g is linear in a and positively two-homogeneous in (a, w, b), there exists $\theta \in \Theta$ such that $g(\theta) \neq 0$. Since $g(-a, w, b) = -g(a, w, b)$, there therefore exist $\varepsilon > 0, \theta \in \Theta$ such that $g(\theta) < -2\varepsilon\,|\theta|^2$. Thus there exists an open cone C in $\mathbb{R}^{d+2}$ such that

(1) $C \cap \Theta \neq \emptyset$,
(2) $g(\theta) < -\varepsilon\,|\theta|^2$ for all $\theta \in C$, and
(3) $\langle \nabla g, \nu_C \rangle > 0$ does not vanish on ∂C where ν is the inner normal vector to ∂C.

Using Assumption (M-S), C can for example be chosen as

$$C = \left\{ \theta \in \mathbb{R}^{d+2} : \frac{g(\theta)}{|\theta|^2} < -t \right\}$$

for some $t \in (\varepsilon, 2\varepsilon)$. We define the localized second moments

$$N_C(\pi) := \int_C |\theta|^2 \, \pi(\mathrm{d}\theta).$$

Since $\partial C \cap S^{d+1}$ is compact and $\nabla(\delta_\pi \mathcal{R}) \to \nabla g$ locally uniformly, we find that there exists $T_0 > 0$ such that $\langle \nabla(\delta_\pi \mathcal{R})(\pi_t; \cdot), \nu_C \rangle > 0$ on ∂C for all $t > T_0$. In particular, no mass flows out of C after time T_0: If $X_\theta(T_0) \in C$ then also $X_\theta(t) \in C$ for $t > T_0$. Thus

$$N_C(\pi_t) = \int_C |\theta|^2 \, X(t)_\sharp \pi_0 = \int_{X(t)^{-1}(C)} \big|X_\theta(t)\big|^2 \, \pi_0(\mathrm{d}\theta) \geq \int_{X(T_0)^{-1}(C)} \big|X_\theta(t)\big|^2 \, \pi_0(\mathrm{d}\theta).$$

Secondly since $(\delta_\pi \mathcal{R})(\pi_t, \cdot) \to g$ locally uniformly, there exists $T_1 > 0$ such that

$$(\delta_\pi \mathcal{R})(\pi_t; \theta) \leq -\frac{\varepsilon}{2} \, |\theta|^2$$

for all $t > T_1$. Without loss of generality, we assume that $T_1 = T_0$. In particular

$$\frac{d}{dt} \, |X_\theta(t)|^2 = \langle X_\theta(t), -\nabla(\delta_\pi \mathcal{R})\big(\pi_t; X_\theta(t)\big) = -2\,(\delta_\pi \mathcal{R})\big(\pi_t; X_\theta(t)\big) \geq \varepsilon \, \big|X_\theta(t)\big|^2$$

for $t > T_0$ and $\theta \in X(T_0)^{-1}(C)$, using the positive two-homogeneity of $(\delta_\pi \mathcal{R})$. Thus

$$|X_\theta(t)|^2 \geq |X_\theta(T)|^2 \ e^{\varepsilon(t-T_0)}$$

for $t > T_0$ and $\theta \in X(T_0)^{-1}(C)$, and consequently

$$\begin{aligned} N(\pi_t) &\geq N_C(\pi_t) \\ &\geq \int_{X(T_0)^{-1}(C)} \big|X_\theta(t)\big|^2 \, \pi_0(\mathrm{d}\theta) \\ &\geq \int_{X(T_0)^{-1}(C)} \big|X_\theta(T_0)\big|^2 \, e^{\varepsilon(t-T_0)} \, \pi_0(\mathrm{d}\theta) \\ &= e^{\varepsilon(t-T_0)} \, N_C(\pi_{T_0}). \end{aligned}$$

If π_0 is omni-directional, then so is π_{T_0} and $N_C(\pi_{T_0}) > 0$. □

Remark A.4. Morally, assumption (M-S) is used to control the sign of a boundary flux term. Without an assumption of this type, the term would at most be asymptotically non-negative. It would be necessary to control the size of a boundary term by a volume contribution and its asymptotic behavior.

A.2. **Proofs from Section 3.** We use these results to prove the main theorem.

Proof of Theorem 3.5. **Step 1.** Since $\delta_\pi \mathcal{R}$ and $\nabla \delta_\pi \mathcal{R}$ are positively two- and one-homogeneous respectively, we find that $(\delta_\pi \mathcal{R})(\pi, 0) = 0$ and $\nabla(\delta_\pi \mathcal{R})(\pi, 0) = 0$ for any $\pi \in \mathcal{P}_2$. According to Lemma A.2, we may assume a uniform Lipschitz bound on $\nabla(\delta_\pi \mathcal{R})(\pi_t, \cdot)$. This implies a uniform Lipschitz bound also on $\delta_\pi \mathcal{R}$ on bounded sets. We thus conclude that $(\delta_\pi \mathcal{R})(\pi_t, \cdot)$ and $\nabla(\delta_\pi \mathcal{R})(\pi_t, \cdot)$ have convergent subsequences, since Lipschitz-space embeds compactly into the space of continuous functions by the Arzelà-Ascoli theorem.

Step 2. First, assume that Ω^{lim} consists of only one element (g, V). Then either $g \not\equiv 0$ or $g \equiv 0$. In the first case, we have by Lemma 3.4 that $N(\pi_t)$ grows exponentially in time, contradicting Lemma 3.3. In the second case, we use homogeneity to show that

$$\begin{aligned} -\frac{d}{dt} N(\pi_t) &= \int_\Theta \langle \nabla |\theta|^2, \nabla(\delta_\pi \mathcal{R}) \rangle \, \pi_t(\mathrm{d}\theta) \\ &= 2 \int_\Theta \langle \theta, \nabla(\delta_\pi \mathcal{R}) \rangle \, \pi_t(\mathrm{d}\theta) \\ &= 4 \int_\Theta \delta_\pi \mathcal{R}(\theta) \, \pi_t(\mathrm{d}\theta) \\ &= 4 \int_\Theta \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, \phi(\theta, x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \, \pi_t(\mathrm{d}\theta) \\ &= 4 \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \int_\Theta \phi(\theta, x) \, \pi_t(\mathrm{d}\theta) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\ &= 4 \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, f_{\pi_t}(x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\ &= 4 \int_{\mathbb{R}^d \times \mathbb{R}} \ell\big(f^*(x), y\big) + (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, \big(f_{\pi_t} - f^*\big)(x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\ &\qquad - 4 \, \widetilde{\mathcal{R}}(f^*) + 4 \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big) \, f^*(x) \, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \\ &\geq 4 \int_{\mathbb{R}^d \times \mathbb{R}} \ell\big(f_{\pi_t}(x), y\big) \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) - 4 \, \widetilde{\mathcal{R}}(f^*) \end{aligned}$$

$$+ 4 \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, f^*(x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$
$$= 4 \left[\widetilde{\mathcal{R}}(f_{\pi_t}) - \widetilde{\mathcal{R}}(f^*)\right] + 4 \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, f^*(x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y).$$

due to the first-order convexity condition for ℓ in the first argument. The integrals converge since $\partial_1 \ell \in L^\infty(\overline{\mathbb{P}})$ and $f^* \in L^1(\overline{\mathbb{P}})$ by Assumption (LP3). We know that

$$\int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \to g(\theta) = 0 \qquad \forall\, \theta \in \Theta.$$

Using assumption (P3), we conclude that

$$\int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_t}(x), y\big)\, f^*(x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \to 0$$

as $t \to \infty$ since the span of $\{\phi(\theta, \cdot) : \theta \in \Theta\}$ is dense in $L^1(\mathbb{P})$. Thus

$$\liminf_{t \to \infty} \frac{d}{dt} N(\pi_t) \leq 4 \lim_{t \to \infty} \left[\widetilde{\mathcal{R}}(f_{\pi_t}) - \widetilde{\mathcal{R}}(f^*)\right] \leq 0$$

because $\widetilde{\mathcal{R}}(f^*) = \inf_\pi \mathcal{R}(\pi)$. Since N is bounded from below by zero, it cannot decrease linearly at a fixed non-zero rate for all large arguments. We conclude that

$$\lim_{t \to \infty} \mathcal{R}(\pi_t) = \widetilde{\mathcal{R}}(f^*) = \inf_\pi \mathcal{R}(\pi).$$

Step 3. Now, assume that $\lim_{t \to \infty} \mathcal{R}(\pi_t) = \widetilde{\mathcal{R}}(f^*)$, i.e.

$$\lim_{t \to \infty} \int_{\mathbb{R}^d} L_x(f_{\pi_t}(x)) - L_x(f^*(x))\, \overline{\mathbb{P}}(\mathrm{d}x) = 0$$

where L is the *augmented loss function* discussed in Section 2.2. Since the first integrand is always larger than the second one, their difference is positive and thus we conclude that

$$L_x(f_{\pi_t}(x)) - L_x(f^*(x)) \to 0 \qquad \text{in } L^1(\overline{\mathbb{P}}).$$

At least for a subsequence $t_n \to \infty$, this means that the convergence holds pointwise almost everywhere. Since $L_x(f_{\pi_{t_n}}(x)) \to L_x(f^*(x))$ $\overline{\mathbb{P}}$-almost everywhere and $L_x(\alpha) \to \infty$ if $|\alpha| \to \infty$, the minimizing sets M_x are compact and that $\operatorname{dist}(f_{\pi_{t_n}}(x), M_x) \to 0$. The pointwise limit of the functions is measurable, and thus without loss of generality $f^*(x) = \lim_{n \to \infty} f_{\pi_{t_n}}(x)$. By the continuity of $\partial_1 \ell$, we conclude that

$$(\partial_1 \ell)\big(f_{\pi_{t_n}}(x), y\big) \to (\partial_1 \ell)\big(f^*(x), y\big)$$

almost everywhere. We recall that

$$L_x'(f^*(x)) = \int_{\mathbb{R}} (\partial_1 \ell)\big(f^*(x), y\big)\, \mathbb{P}^x(\mathrm{d}y) = 0$$

for $\overline{\mathbb{P}}$-almost all x and thus in particular

$$\int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f^*(x), y\big)\, \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) = 0 \qquad \forall\, \theta \in \Theta.$$

Since $\phi(\theta, \cdot) \in L^1(\overline{\mathbb{P}})$ for any $\theta \in \overline{\Theta}$, we conclude from the dominated convergence theorem that

$$\int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f_{\pi_{t_n}}(x), y\big)\, \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) \to \int_{\mathbb{R}^d \times \mathbb{R}} (\partial_1 \ell)\big(f^*(x), y\big)\, \phi(\theta, x)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y) = 0$$

using the bound

$$(\partial_1 \ell)\big(f_{\pi_{t_n}}(x), y\big)\, \phi(\theta, \cdot) \leq \|\partial_1 \ell\|_{L^\infty}\, |\phi(\theta, \cdot)|.$$

Thus $\lim_{n\to\infty} \delta_\pi \mathcal{R}(\pi_{t_n}, \cdot) \equiv 0$. By compactness, we know that Ω^{lim} is non-empty, and we have showed that for any sequence, we can extract a subsequence for which $g = 0$. Since locally uniform convergence is generated by a topology, this means that g is the only possible limit point. The same argument can be used for the gradient. □

Remark A.5. Note that omni-directionality is only used to exclude the case that $g \not\equiv 0$, whereas $g \equiv 0$ is only an admissible limit if $\mathcal{R}(\pi_t)$ decays to MBR. This corresponds to the fact that $\delta_\pi \mathcal{R}(\pi, \cdot) \equiv 0$ if and only if π is a global minimizer of $\mathcal{R}$, see [CB18b, Proposition 3.1].

Remark A.6. Omni-directionality of the initial condition (IC2) and the Morse-Sard property are both involved only in excluding a unique limit $g \not\equiv 0$. They could therefore be replaced for example by the zero-limit assumption

(ZL) If $(\delta_\pi \mathcal{R})(\pi_t, \cdot) \to g$ locally uniformly on $\overline{\Theta}$ and $g \not\equiv 0$, then

$$\limsup_{t\to\infty} \frac{N(\pi_t)}{t} > 0.$$

The following two proofs establish the corollaries to the main theorem concerning minimizers.

Proof of Corollary 3.7. If we assume in addition that there exists a probability measure π_∞ and a sequence of times $t_n \to \infty$ such that $\lim_{n\to\infty} W_2(\pi_{t_n}, \pi_\infty) = 0$, then we find by [Vil08, Theorem 6.9] that

$$f_{\pi_\infty} = \lim_{n\to\infty} f_{\pi_{t_n}} = f^*.$$

Thus π_∞ minimizes $\mathcal{R}$. □

Proof of Corollary 3.8. Consider the measures μ_t on $S^{d+1} \cap \Theta$ defined by

$$\int_{S^{d+1}\cap\Theta} g(\theta)\, \mu_t(\mathrm{d}\theta) = \int_\Theta g\left(\frac{\theta}{|\theta|}\right) |\theta|^2\, \pi_t(\mathrm{d}\theta)$$

for $g \in C(\overline{S^{d+1} \cap \Theta})$, or equivalently

$$\mu_t = P^{S^{d+1}}_\sharp \left(|\theta|^2 \cdot \pi_t\right)$$

where $P^{S^{d+1}} : \Theta \to S^{d+1} \cap \Theta$ is the canonical projection. While $P^{S^{d+1}}$ is undefined at $\theta = 0$, the projection of the measure with weight $|\theta|^2$ is well-defined. Under the moment bound assumption, the measures μ_t are uniformly bounded and

$$\int_{S^{d+1}\cap\Theta} \phi(\theta, x)\, \mu_t(\mathrm{d}\theta) = \int_\Theta \phi(\theta, x)\, \pi_t(\mathrm{d}\theta) = f_{\pi_t}(x)$$

for all $x \in U$. By the compactness theorem of Radon measures, there exists a finite Radon measure μ_∞ on $\overline{\Theta} \cap S^{d+1}$ and a sequence of times $t_n \to \infty$ such that $\lim_{n\to\infty} \tilde{\pi}_{t_n} = \mu_\infty$. On the closed cone $\overline{\Theta}$, we can apply the weak convergence of Radon measures to conclude that

$$\int_{S^{d+1}\cap\overline{\Theta}} \phi(\theta, x)\, \mu_\infty(\mathrm{d}\theta) = \lim_{n\to\infty} \int_{S^{d+1}\cap\Theta} \phi(\theta, x)\, \mu_{t_n}(\mathrm{d}\theta) = f^*(x)$$

$\overline{\mathbb{P}}$-almost everywhere. The measure μ_∞ is non-negative since all μ_{t_n} are. We now distinguish two cases:

(1) $\mu_\infty = 0$. In this case $f^* \equiv 0$ and $\pi_\infty := \delta_{\theta=0}$ is a risk minimizer.
(2) $\mu_\infty \neq 0$. In this case, we define the dilation map $D = \sqrt{\mu_\infty(\Theta)}\, I$ and

$$\pi_\infty := \frac{1}{\mu_\infty(\Theta)}\, D_\sharp \mu_\infty.$$

Then by homogeneity

$$f_{\pi_\infty}(x) = \int_\Theta \phi(\theta, x) \mathrm{d}\pi_\infty = \frac{1}{\tilde{\pi}_\infty(\Theta)} \int_{S^{d+1}\cap\Theta} \phi\left(\sqrt{\tilde{\pi}_\infty(\Theta)}\,\theta, x\right) \mu_\infty(\mathrm{d}\theta) = f^*(x),$$

for all $x \in U$, so π_∞ is a risk minimizer.

□

Remark A.7. Considering the example

$$\pi_n := \frac{1}{n}\,\delta_{\theta=\sqrt{n}e_1} + \left(1 - \frac{1}{n}\right)\,\delta_{\theta=0}, \quad \pi_\infty = \delta_{\theta=0}, \qquad f_{\pi_n} \equiv \phi(e_1, \cdot) \neq 0 = f_{\pi_\infty}.$$

after Corollary 3.8, we observe that

$$\mu_n := P_\sharp^{S^{d+1}}\left(|\theta|^2 \cdot \pi_n\right) \equiv \delta_{\theta=e_1}.$$

Thus the projection to the unit sphere adds compactness beyond the moment bound.

Appendix B. The Morse-Sard Property

The classical Morse-Sard theorem states that the set of critical values of a function $f \in C^k(\mathbb{R}^n, \mathbb{R}^m)$ has zero Lebesgue measure if $k \geq \max\{n - m + 1, 1\}$, where $y \in \mathbb{R}^m$ is a critical value of f if there exists $x \in f^{-1}(y)$ such that $Df(x)$ does not have full rank. Combined with the regular value theorem, it is a powerful tool in smooth approximation.

The result is due to Morse for $m = 1$ [Mor39] and Sard for general $m \geq 1$ [Sar42]. It is easy to extend the result to sufficiently differentiable manifolds, and there are more precise statements available for the Hausdorff measure of $f(S_\nu)$ where

$$S_\nu = \{x \in \mathbb{R}^n \mid \operatorname{rank}(Df_x) \leq \nu\}$$

for $0 \leq \nu \leq n - 1$, see [Fed69, Theorem 3.4.3]. In the classical setting, the differentiability assumptions are almost sharp, and the condition can only be weakened to $C^{k-1,1}$ in place of C^k [Bat93]. A classical example of Whitney [Whi35] which has been generalized to higher dimension e.g. in [Haj03] shows that there exists $f \in C^1(\mathbb{R}^2, \mathbb{R})$ such that f is non-constant on a connected component of its set of critical points, meaning that there is an open interval of critical points.

Morse-Sard theorems are known to fail in infinite dimension even for infinitely smooth maps, unless additional assumptions are imposed. Under weak conditions, however, the set of functions for which the Theorem holds is dense in the C^0-topology [EM68].

Due to its fundamental importance, some effort has been made to establish Morse-Sard type properties in other function classes. Among these are

- Morse-Sard theorems in classes of weakly differentiable functions [Fig08, BKK13, BKK15, dP01, KK18]. Here the relation between differentiability and integrability may even be chosen low enough to ensure continuity, but not classical differentiability of the functions under consideration.
- Morse-Sard theorems for the distance from a submanifold [Rif04] or more generally Lipschitz functions which are given as suprema of smooth functions over suitable index sets [BDDR16].
- Morse-Sard theorems for subanalytic functions, see [BDL06].
- Morse-Sard theorems for dc functions in two dimensions [PZ06]. A function is dc if it can be written as differences of two convex functions. In particular, every C^2-function with bounded second derivatives is dc. Thus the result cannot be generalized to higher dimensions.

In non-smooth function classes, a notion of gradient almost everywhere (with respect to a suitable Hausdorff measure) or a sub-differential is used.

We show below that (M-S) holds unconditionally in dimension $d = 2$. In [CB20], the Morse-Sard theorem for subanalytic functions from [BDL06] has successfully been used to establish a condition of Morse-Sard property in a very similar application. The subanalytic function that the authors consider in [CB20] is a finite sum of ReLU-like terms and the subanalyticity stems from the finiteness of the sum. The number of summands corresponds to the number of data samples in an empirical measure. The approach is therefore incompatible with assumption (P4) for the ReLU case.

B.1. **Finitely many neurons.** While it does not apply in our situation, we briefly sketch the result and its application in a similar situation. Consider $f(x) = \sum_{i=1}^m a_i(w_i^T x + b_i)_+$. Then f is only Lipschitz-smooth and not C^1, but has the Morse-Sard property due to its sub-analyticity.

Definition B.1. (1) A set $A \subseteq \mathbb{R}^n$ is called semi-analytic if for every $x \in A$ there exists a neighbourhood U of x and a collection of real analytic functions f_{ij}, g_{ij} such that

$$A \cap U = \bigcup_{i=1}^{N} \bigcap_{j=1}^{M_i} \{f_{ij} > 0, g_{ij} = 0\}.$$

(2) A set $A \subseteq \mathbb{R}^n$ is called sub-analytic if for every $x \in A$ there exists a neighbourhood U of x, a semi-analytic set $\tilde{A}$ in $\mathbb{R}^m$ for $m \geq n$ and an affine map $L : \mathbb{R}^m \to \mathbb{R}^n$ such that $A = L(\tilde{A})$.

(3) A function f is called sub-analytic if its graph $G_f = \{(x, y) \in \mathbb{R}^{n+1} \mid y = f(x)\}$ is sub-analytic.

We recall the following Morse-Sard theorem.

Theorem B.2. [BDL06, Theorem 13] *Let $f : \mathbb{R}^n \to \mathbb{R}$ be a continuous sub-analytic function. Then f is constant on connected components of the set of critical points (defined by the sub-differential), and the set of critical values is countable.*

We show that this applies to finite ReLU networks.

Lemma B.3. *$f(x) = \sum_{i=1}^m a_i(w_i^T x + b_i)_+$ is continuous and sub-analytic.*

Proof. Continuity is clear. Let $(x, y) \in G_f$. Since f is affine linear on the set $\{x | w_i^T x + b_i \neq 0 \forall\, i\}$, the graph is a plane and thus sub-analytic here. Now assume that $w_1^T x + b_1 = 0$ and $w_i^T x + b_i \neq 0$ for all $i > 1$. Then, locally after a rotation and translation we have

$$f(x) = \alpha^T x + \max\{x_1, 0\}.$$

The graph of this function is sub-analytic since

$$G_f = \{x_1 > 0,\ \alpha^T x + x_1 - y = 0\} \cup \{x_1 < 0,\ \alpha^T x - y = 0\} \cup \big(\{x_1 = 0\} \cap \{y = 0\}\big).$$

The case when more terms $w_i^T x + b_i$ vanish can be treated similarly, but is somewhat tedious to write out. It is, however, crucial that the sum is finite to ensure that there are at most finitely many sets to be united and intersected. □

B.2. **ReLU Geometry on the Sphere.** We can exploit the special geometry of the problem to reduce the dimension slightly. We can write

$$g(a, w, b) = a\, h(w, b), \qquad h(w, b) = \int_{\mathbb{R}^{d+1}} \rho(x, y)\, \mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

where ρ is an L^p-weak limit of $(\partial_1 \ell)(f_\pi(x), y)$ for any $1 < p < \infty$. In particular, h is positively one-homogeneous and a ReLU Barron function [EMW19b, Inverse Approximation Theorem].

To simplify notation, from now on w has $d+1$ components and we identify w with (w,b) and x with $(x,1)$. Denote by $e_a = (1,0,\dots,0) \in \mathbb{R}^{d+2}$ the unit vector pointing in the a-direction. At $(a,w) \in S^{d+1}$, the tangential part of ∂_a is

$$e_a^{\parallel} := e_a - \langle e_a, (a,w)\rangle (a,w) = e_a - a \begin{pmatrix} a \\ w \end{pmatrix}.$$

We compute

$$\nabla_a^S g(a,w) = \partial_a g - a\,(a,w)\cdot \nabla g = \frac{g}{a} - 2a\,g = \left(\frac{1}{a} - 2a\right) g.$$

Thus on the set

$$\left\{(a,w,b) : \frac{1}{a} - 2a \neq 0\right\} = \left\{(a,w,b) : a^2 \neq 2\right\} = \left\{(a,w,b) \in \Theta : a^2 < |w|^2 + b^2\right\}$$

the only critical value is (possibly) zero. On the other hand, consider (a,w) in $\partial\Theta \cap S^{d+1}$. We note that

$$g(\lambda a, \mu w) = \lambda\mu\, g(a,w) \qquad \forall\ \lambda,\mu > 0$$

and

$$(\lambda a)^2 + |\mu w|^2 = \lambda^2 a^2 + \mu^2 |w|^2 = 1 \qquad \text{if} \quad \mu^2 = \frac{1-\lambda^2 a^2}{|w|^2} = 2 - \lambda^2.$$

We find that

$$g\big(\lambda a,\ \mu(\lambda)\, w\big) = \lambda\sqrt{2-\lambda^2}\, g(a,w)$$

is extremal at $\lambda = \pm 1$ since

$$0 = \frac{d}{d\lambda}\lambda\sqrt{2-\lambda^2} = 2\,\frac{1-\lambda^2}{\sqrt{2-\lambda^2}} \qquad \Leftrightarrow \quad \lambda = \pm 1.$$

Thus all level sets of g intersect $\partial\Theta \cap S^{d+1}$ and linearity in a is insufficient to establish the Morse-Sard property. The Morse-Sard property (M-S) holds if and only if the Morse-Sard property holds for the function $h : S^d \to \mathbb{R}$. Since $h \in C^{1,1}_{loc}(\overline{\Theta})$ by Lemma A.2, the Morse-Sard property holds if $d = 2$ by [Bat93, Theorem 1].

B.3. **A mild counterexample.** We show that functions with similar structural properties as g (or h) may not satisfy a Morse-Sard property in dimension $d \geq 8$.

Definition B.4. Let $\Omega \subseteq \mathbb{R}^d$ be open and bounded. $f : \Omega \to \mathbb{R}$ is called a Barron function if there exists $\pi \in \mathcal{P}_2(\mathbb{R}^{d+2})$ such that

$$f(x) = f_\pi(x) = \int_{\mathbb{R}^{d+2}} a\,\sigma(w^T x + b)\,\pi(\mathrm{d}a \otimes \mathrm{d}w \otimes \mathrm{d}b)$$

in Ω.

The space of Barron functions is discussed in detail in [EMW18, EMW19b, EW20a] and [EW20b, Appendix A]. The space is named after Andrew Barron who first established that a large class of functions could be represented in such a way. We cite a simplified version of Barron's main theorem.

Lemma B.5. [Bar93, Proposition 1] *Assume that $f : \mathbb{R}^d \to \mathbb{R}$ is such that its Fourier transform $\widehat{f}$ satisfies*

$$\int_{\mathbb{R}^d} |\xi|\,\big|\widehat{f}\big|(\xi)\,\mathrm{d}\xi < \infty.$$

Then f is Barron.

The original proof uses sigmoidal activation functions, but the result for ReLU activation follows immediately since $\mathrm{ReLU}(x) - \mathrm{ReLU}(x-1)$ is sigmoidal. The criterion is classical. A less well know consequence is the following, derived like [Bar93, Section IX, point 15].

Corollary B.6. *Assume that $f \in H^s(\mathbb{R}^d)$ for $s > \frac{d}{2} + 1$. Then f is a Barron function.*

Proof. Using the identity $\widehat{\partial_j f} = i\,\xi_j \widehat{f}$ and Parseval's identity, we compute

$$\begin{aligned}\int_{\mathbb{R}^d} |\xi|\,\big|\widehat{f}\big|(\xi)\,\mathrm{d}\xi &= \int_{\mathbb{R}^d} |\xi|\,\big|\widehat{f}\big|(\xi)\,\left(1+|\xi|^{2s}\right)^{\frac{1}{2}}\left(1+|\xi|^{2s}\right)^{-\frac{1}{2}}\,\mathrm{d}\xi \\ &\leq \left(\int_{\mathbb{R}^d} |\widehat{f}|^2\,\left(1+|\xi|^{2s}\right)\mathrm{d}\xi\right)^{\frac{1}{2}}\left(\int_{\mathbb{R}^d}\frac{|\xi|^2}{1+|\xi|^{2s}}\mathrm{d}\xi\right)^{\frac{1}{2}}\end{aligned}$$

By [Dob10, Satz 9.37], the first factor on the right is finite if and only if $f \in H^s(\mathbb{R}^d)$. The second factor is finite if $2s - 2 > d$. □

In particular, if f is C^k-smooth on a ball $B_r(0)$ for $k > \frac{d}{2} + 1$, then we can use standard extension theorems to show that f is Barron. In particular, if $d \geq 7$ then $\frac{d}{2} + 2 < d - 1$ and e.g. every C^5-function on $B_1(0) \subset \mathbb{R}^7$ is Barron. In particular, we have shown the following.

Corollary B.7. *In dimension $d \geq 7$ there exist Barron functions which do not have the Morse-Sard property.*

Denote by S^d the unit sphere in $\mathbb{R}^{d+1}$. The function

$$h : S^d \to \mathbb{R}, \qquad h(w,b) = \int \rho(x,y)\,\mathbb{P}(\mathrm{d}x \otimes \mathrm{d}y)$$

from the previous section shares many properties with general Barron functions. In a fixed application, the data distribution $\mathbb{P}$ is fixed, so a Sard property would need to be established in a random feature space (for bounded density ρ). Across different applications, the data distribution $\mathbb{P}$ may vary within the class satisfying assumptions (P1), (P2), (P3), (P4). While this class does not exhaust the entire Barron space, we note that (P4) only implies $C^{1,1}$-smoothness and no sufficient regularity to invoke a classical Sard theorem. We thus expect that there are cases of interest in which the Morse-Sard condition (M-S) does not hold.

References

[AB94] C. D. Aliprantis and K. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, 1994.

[ABM14] H. Attouch, G. Buttazzo, and G. Michaille. *Variational analysis in Sobolev and BV spaces: applications to PDEs and optimization*. SIAM, 2014.

[ADH+19] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pages 8139–8148, 2019.

[AGS08] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.

[AKSG19] M. Arbel, A. Korba, A. Salim, and A. Gretton. Maximum mean discrepancy gradient flow. In *Advances in Neural Information Processing Systems*, pages 6481–6491, 2019.

[Amb08] L. Ambrosio. Transport equation and Cauchy problem for non-smooth vector fields. In *Calculus of variations and nonlinear partial differential equations*, pages 1–41. Springer, 2008.

[AOY19] D. Araújo, R. I. Oliveira, and D. Yukimura. A mean-field limit for certain deep neural networks. *arXiv:1906.00193 [math.ST]*, 2019.

[Bar93] A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.

[Bat93] S. M. Bates. Toward a precise smoothness hypothesis in Sard's theorem. *Proceedings of the American Mathematical Society*, pages 279–283, 1993.

[BDDR16] L. Barbet, M. Dambrine, A. Daniilidis, and L. Rifford. Sard theorems for Lipschitz functions and applications in optimization. *Israel Journal of Mathematics*, 212(2):757–790, 2016.

[BDL06] J. Bolte, A. Daniilidis, and A. Lewis. A nonsmooth Morse-Sard theorem for subanalytic functions. *Journal of mathematical analysis and applications*, 321(2):729–740, 2006.

[BJ18] L. Berlyand and P.-E. Jabin. On the convergence of formally diverging neural net-based classifiers. *Comptes Rendus Mathematique*, 356(4):395–405, 2018.

[BKK13] J. Bourgain, M. V. Korobkov, and J. Kristensen. On the Morse-Sard property and level sets of Sobolev and BV functions. *Rev.Mat.Iberoam*, 29(1):1 – 23, 2013.

[BKK15] J. Bourgain, M. V. Korobkov, and J. Kristensen. On the Morse–Sard property and level sets of $W^{n,1}$-Sobolev functions on $\mathbb{R}^n$. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2015(700):93–112, 2015.

[BM19] A. Bietti and J. Mairal. On the inductive bias of neural tangent kernels. In *Advances in Neural Information Processing Systems*, pages 12873–12884, 2019.

[CB18a] L. Chizat and F. Bach. A note on lazy training in supervised differentiable programming. *arXiv:1812.07956 [math.OC]*, 2018.

[CB18b] L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.

[CB20] L. Chizat and F. Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arxiv:2002.04486 [math.OC]*, 2020.

[Cyb89] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[DKW19] P. W. Dondl, M. W. Kurzke, and S. Wojtowytsch. The effect of forest dislocations on the evolution of a phase-field model for plastic slip. *Archive for Rational Mechanics and Analysis*, 232(1):65–119, 2019.

[DLL+18] S. S. Du, J. D. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. *arXiv:1811.03804 [cs.LG]*, 2018.

[Dob10] M. Dobrowolski. *Angewandte Funktionalanalysis: Funktionalanalysis, Sobolev-Räume und elliptische Differentialgleichungen*. Springer-Verlag, 2010.

[dP01] L. de Pascale. The Morse-Sard theorem in Sobolev spaces. *Indiana University mathematics journal*, pages 1371–1386, 2001.

[DZPS18] S. S. Du, X. Zhai, B. Poczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv:1810.02054 [cs.LG]*, 2018.

[EG15] L. C. Evans and R. F. Gariepy. *Measure theory and fine properties of functions*. CRC press, 2015.

[EM68] J. Eells and J. McAlpin. An approximate Morse-Sard theorem. *Journal of Mathematics and Mechanics*, 17(11):1055–1064, 1968.

[EMW18] W. E, C. Ma, and L. Wu. A priori estimates of the population risk for two-layer neural networks. *Comm. Math. Sci.*, 17(5):1407 – 1425 (2019), arxiv:1810.06397 [cs.LG] (2018).

[EMW19a] W. E, C. Ma, and Q. Wang. A priori estimates of the population risk for residual networks. *arXiv:1903.02154 [cs.LG]*, 2019.

[EMW19b] W. E, C. Ma, and L. Wu. Barron spaces and the compositional function spaces for neural network models. *arXiv:1906.08039 [cs.LG]*, 2019.

[EMW19c] W. E, C. Ma, and L. Wu. Machine learning from a continuous viewpoint. *arxiv:1912.12777 [math.NA]*, 2019.

[EMW19d] W. E, C. Ma, and L. Wu. A comparative analysis of optimization and generalization properties of two-layer neural network and random feature models under gradient descent dynamics. *Sci. China Math.*, https://doi.org/10.1007/s11425-019-1628-5, arXiv:1904.04326 [cs.LG] (2019).

[EMWW19] W. E, C. Ma, Q. Wang, and L. Wu. Analysis of the gradient descent algorithm for a deep neural network model with skip-connections. *arXiv:1904.05263 [cs.LG]*, 2019.

[EW20a] W. E and S. Wojtowytsch. Barron functions and their representation. *In preparation*, 2020.

[EW20b] W. E and S. Wojtowytsch. Kolmogorov width decay and poor approximators in machine learning: Shallow neural networks, random feature models and neural tangent kernels. *arXiv:2005.10807 [math.FA]*, 2020.

[Fed69] H. Federer. *Geometric Measure Theory*. Springer, Berlin, 1969.

[Fig08] A. Figalli. A simple proof of the Morse-Sard theorem in Sobolev spaces. *Proceedings of the American Mathematical Society*, 136(10):3675–3681, 2008.

[FL07] I. Fonseca and G. Leoni. *Modern Methods in the Calculus of Variations: L^p Spaces*. Springer Science & Business Media, 2007.

[GBC16] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.

[Haj03] P. Hajłasz. Whitney's example by way of Assouad's embedding. *Proceedings of the American Mathematical Society*, 131(11):3463–3467, 2003.

[HH19] C. F. Higham and D. J. Higham. Deep learning: An introduction for applied mathematicians. *SIAM Review*, 61(4):860–891, 2019.

[HLLL19] W. Hu, C. J. Li, L. Li, and J.-G. Liu. On the diffusion approximation of nonconvex stochastic gradient descent. *Annals of Mathematical Sciences and Applications*, 4(1):3–32, 2019.

[Hor91] K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.

[HRSS19] K. Hu, Z. Ren, D. Siska, and L. Szpruch. Mean-field Langevin dynamics and energy landscape of neural networks. *arXiv:1905.07769 [math.PR]*, 2019.

[JGH18] A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.

[JKO98] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.

[KK18] M. V. Korobkov and J. Kristensen. The trace theorem, the Luzin N-and Morse-Sard properties for the sharp case of Sobolev-Lorentz mappings. *The Journal of Geometric Analysis*, 28(3):2834–2856, 2018.

[LLPS93] M. Leshno, V. Y. Lin, A. Pinkus, and S. Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.

[LTE15] Q. Li, C. Tai, and W. E. Dynamics of stochastic gradient algorithms. *arXiv:1511.06251 [cs.LG]*, 2015.

[MBW+19] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab. A high-bias, low-variance introduction to machine learning for physicists. *Physics reports*, 2019.

[McC97] R. J. McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179, 1997.

[Mie12] A. Mielke. Emergence of rate-independent dissipation from viscous systems with wiggly energies. *Continuum Mechanics and Thermodynamics*, 24(4-6):591–606, 2012.

[MMN18] S. Mei, A. Montanari, and P.-M. Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

[Mor39] A. P. Morse. The behavior of a function on its critical set. *Annals of Mathematics*, pages 62–70, 1939.

[Ngu19] P.-M. Nguyen. Mean field limit of the learning dynamics of multilayer neural networks. *arXiv:1902.02880 [cs.LG]*, 2019.

[NP20] P.-M. Nguyen and H. T. Pham. A rigorous framework for the mean field limit of multilayer neural networks. *arXiv:2001.11443 [cs.LG]*, 2020.

[Ott01] F. Otto. The geometry of dissipative evolution equations: the porous medium equation. *Comm. Partial Differential Equations*, 26, 2001.

[PZ06] D. Pavlica and L. Zajíček. Morse-sard theorem for dc functions and mappings on $\mathbb{R}^2$. *Indiana University mathematics journal*, pages 1195–1207, 2006.

[Rif04] L. Rifford. A Morse-Sard theorem for the distance function on Riemannian manifolds. *manuscripta mathematica*, 113(2):251–265, 2004.

[RVE18] G. M. Rotskoff and E. Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *arXiv:1805.00915 [stat.ML]*, 2018.

[San15] F. Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.

[Sar42] A. Sard. The measure of the critical values of differentiable maps. *Bulletin of the American Mathematical Society*, 48(12):883–890, 1942.

[Ser11] S. Serfaty. Gamma-convergence of gradient flows on Hilbert and metric spaces and applications. *Discrete Contin. Dyn. Syst*, 31(4):1427–1451, 2011.

[SS04] E. Sandier and S. Serfaty. Gamma-convergence of gradient flows with applications to Ginzburg-Landau. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 57(12):1627–1672, 2004.

[SS19] J. Sirignano and K. Spiliopoulos. Mean field analysis of deep neural networks. *arXiv:1903.04440 [math.PR]*, 2019.

[SS20] J. Sirignano and K. Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM J. Appl. Math*, 80(2):725–752, 2020.

[SSBD14] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[Ver18] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

[Vil08] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

[WE20] S. Wojtowytsch and W. E. Can shallow neural networks beat the curse of dimensionality? A mean field training perspective. *arXiv:2005.10815 [cs.LG]*, 2020.

[Whi35] H. Whitney. A function not constant on a connected set of critical points. *Duke Mathematical Journal*, 1(4):514–517, 1935.

Stephan Wojtowytsch, Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544

E-mail address: stephanw@princeton.edu