



HAL
open science

On the symmetries in the dynamics of wide two-layer neural networks

Karl Hajjar, Lenaic Chizat

► **To cite this version:**

Karl Hajjar, Lenaic Chizat. On the symmetries in the dynamics of wide two-layer neural networks. 2022. hal-03829400v2

HAL Id: hal-03829400

<https://hal.science/hal-03829400v2>

Preprint submitted on 24 Nov 2022 (v2), last revised 8 Feb 2023 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the symmetries in the dynamics of wide two-layer neural networks

Karl Hajjar^{*,1} and Lénaïc Chizat²

¹Laboratoire de Mathématiques d’Orsay, Université Paris-Saclay, 91405 Orsay, France

²Institut de Mathématiques, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

November 24, 2022

Abstract

We consider the idealized setting of gradient flow on the population risk for infinitely wide two-layer ReLU neural networks (without bias), and study the effect of symmetries on the learned parameters and predictors. We first describe a general class of symmetries which, when satisfied by the target function f^* and the input distribution, are preserved by the dynamics. We then study more specific cases. When f^* is odd, we show that the dynamics of the predictor reduces to that of a (non-linearly parameterized) linear predictor, and its exponential convergence can be guaranteed. When f^* has a low-dimensional structure, we prove that the gradient flow PDE reduces to a lower-dimensional PDE. Furthermore, we present informal and numerical arguments that suggest that the input neurons align with the lower-dimensional structure of the problem.

1 Introduction

The ability of neural networks to learn rich representations—or features—of their input data is commonly observed in state-of-the-art models (Zeiler and Fergus, 2014; Cammarata et al., 2020) and often thought to be the reason behind their good practical performance (Goodfellow et al., 2016, Chap. 1). Yet, our theoretical understanding of how feature learning arises from simple gradient-based training algorithms remains limited. Much progress (discussed in Section 1.3) has been made recently to understand the power and limitations of gradient-based learning with neural networks, showing in particular their superiority over fixed-feature methods on some difficult tasks. However, positive results are often obtained for algorithms that differ in substantial ways from plain (stochastic) gradient descent (e.g. the layers trained separately, or the algorithm makes just one truly non-linear step, etc).

In this work, we take the algorithm as a given and instead adopt a descriptive approach. Our goal is to improve our understanding of how neural networks behave in the presence of symmetries in the data with plain gradient descent (GD) on two-layer fully-connected ReLU neural networks. To this end, we investigate situations with strong symmetries on the data, the target function and on the initial parameters, and study the properties of the training dynamics and the learnt predictor in this context.

*Corresponding author: hajjarkarl@gmail.com

1.1 Problem setting

We denote by d the input dimension, ρ the input data distribution which we assume to be uniform over the unit sphere \mathbb{S}^{d-1} of \mathbb{R}^d , and by $\mathcal{P}_2(\Omega)$ the space of probability measures with finite second moments over a measurable space Ω . We call σ the activation function, which we take to be ReLU, that is $\sigma(z) = \max(0, z)$, $\ell(\cdot, \cdot)$ the loss function, which we assume to be continuously differentiable *w.r.t.* its second argument and we denote by $\partial_2 \ell$ this derivative.

Mean-field limit of two-layer networks. In this work, we consider the infinite-width limit in the mean-field regime of the training dynamics of two-layer networks without intercept with a ReLU activation function. Given a measure $\mu \in \mathcal{P}_2(\mathbb{R} \times \mathbb{R}^d)$, we consider the infinitely wide two-layer network parameterized by μ , defined, for any input $x \in \mathbb{R}^d$, by

$$f(\mu; x) = \int_{c \in \mathbb{R}^{1+d}} \phi(c; x) d\mu(c), \quad (1.1)$$

where, for any $c = (a, b) \in \mathbb{R} \times \mathbb{R}^d$, $\phi(c; x) = a\sigma(b^\top x)$. Note that width- m two-layer networks with input weights $(b_j)_{j \in [1, m]} \in (\mathbb{R}^d)^m$ and output weights $(a_j)_{j \in [1, m]} \in \mathbb{R}^m$ can be recovered by a measure $\mu_m = (1/m) \sum_{j=1}^m \delta_{(ma_j, b_j)}$ with m atoms.

Objective and Wasserstein gradient flow. We consider the problem of minimizing the *population loss* objective for a given target function $f^* : \mathbb{R}^d \rightarrow \mathbb{R}$, that is

$$\min_{\mu \in \mathcal{P}_2(\mathbb{R} \times \mathbb{R}^d)} \left(F(\mu) := \mathbb{E}_{x \sim \rho} [\ell(f^*(x), f(\mu; x))] \right). \quad (1.2)$$

The Fréchet derivative of the objective function F at μ is given by the function $F'_\mu(c) = \mathbb{E}_{x \sim \rho} [\partial_2 \ell(f^*(x), f(\mu; x)) \phi(c; x)]$ for any $c = (a, b) \in \mathbb{R} \times \mathbb{R}^d$ (for more details, see Appendix B.1). Starting from a given measure $\mu_0 \in \mathcal{P}_2(\mathbb{R} \times \mathbb{R}^d)$, we study the Wasserstein gradient flow (GF) of the objective (1.2) which is a path $(\mu_t)_{t \geq 0}$ in the space of probability measures satisfying, in the sense of distributions, the partial differential equation (PDE) known as the continuity equation:

$$\begin{aligned} \partial_t \mu_t &= -\operatorname{div}(v_t \mu_t) \\ v_t(c) &:= -\nabla F'_{\mu_t}(c) \end{aligned} \quad (1.3)$$

Initialization. We make the following assumption on the initial measure $\mu_0 \in \mathcal{P}_2(\mathbb{R} \times \mathbb{R}^d)$: μ_0 decomposes as $\mu_0 = \mu_0^1 \otimes \mu_0^2$ where $\mu_0^1, \mu_0^2 \in \mathcal{P}_2(\mathbb{R}) \times \mathcal{P}_2(\mathbb{R}^d)$. This follows the standard initialization procedure at finite width. Because no one direction should *a priori* be favored, we assume μ_0^2 to have spherical symmetry, *i.e.*, it is invariant under any orthogonal transformation, and we additionally assume that $|a| = \|b\|$ almost surely at initialization. It is shown in (Chizat and Bach, 2020, Lemma 26), and (Wojtowytsch, 2020, Section 2.5), that with this initialization, μ_t stays supported on the set $\{|a| = \|b\|\}$ for any $t \geq 0$. This last assumption is of a technical nature, and, along with the regularity conditions on the loss ℓ and the input data distribution ρ , ensures that the Wasserstein GF (1.3) is well-defined (Wojtowytsch, 2020, Lemma 3.1, Lemma 3.9).

Relationship with finite-width GD. If $\mu_0 = (1/m) \sum_{j=1}^m \delta_{(a_j(0), b_j(0))}$ is discrete, the Wasserstein GF (1.3) is exactly continuous-time GD on the parameters of a standard finite-width neural network, and discretization errors (*w.r.t.* the number of neurons) can be provided (Mei et al., 2018; Nguyen and Pham, 2020).

1.2 Summary of contributions

Our main object of study is the gradient flow of the *population* risk of *infinitely wide* two-layer ReLU neural networks without intercept. Our motivation to consider this idealistic setting— infinite data and infinite width—is that it allows, under suitable choices for ρ and μ_0 , the emergence of exact symmetries which are only approximate in the non-asymptotic setting¹.

Symmetries, structure, and convergence. In this work, we are interested in the structures learned by the predictor $f(\mu_t; \cdot)$ under GF as t grows large. Specifically, we make the following contributions:

- In Section 2, we prove that if f^* is invariant under some orthogonal linear map T , then $f(\mu_t; \cdot)$ inherits this invariance under GF (Theorem 2.1)
- In Section 3, we study the case when f^* is an *odd* function and show that the network converges to the best linear approximator of f^* at an exponential rate (Theorem 3.2). Linear predictors are optimal over the hypothesis class in that case, in particular because there is no intercept in our model.
- In Section 4, we consider the *multi-index model* where f^* depends *only* on the orthogonal projection of its input onto some sub-space H of dimension d_H . We prove that the dynamics can be reduced to a PDE in dimension d_H . If in addition, f^* is the Euclidean norm of the projection of the input, we show that the dynamics reduce to a one-dimensional PDE (Theorem 4.3). In the latter case, we were not able to prove theoretically the convergence of the neurons of the first layer towards H , and leave this as an open problem but we provide numerical evidence in favor of this result.

The code to reproduce the results of the numerical experiments can be found at:

<https://github.com/karl-hajjar/learning-structure>.

1.3 Related work

Infinite-width dynamics. It has been shown rigourously that for infinitely wide networks there is a clear distinction between a feature-learning regime and a kernel regime (Chizat et al., 2019; Yang and Hu, 2021). For shallow networks, this difference stems from a different scale (*w.r.t.* width) of the initialization where a large initialization leads to the Neural Tangent Kernel (NTK) (*a.k.a.* the “lazy regime”) which is equivalent to a kernel method with random features (Jacot et al., 2018) whereas a small initialization leads to the so-called *mean-field* (MF) limit where features are learned from the first layer (Chizat et al., 2019; Yang and Hu, 2021). However, it is unclear in this setting exactly what those features are and what underlying structures are learned by the network. The aim of the present work is to study this phenomenon from a theoretical perspective for infinitely wide networks and to understand the relationship between the ability of networks to learn specific structures and the symmetries of a given task.

A flurry of works study the dynamics of infinitely wide two-layer neural networks. (Chizat and Bach, 2018; Mei et al., 2018; Rotskoff and Vanden-Eijnden, 2018; Wojtowytsch, 2020; Sirignano and Spiliopoulos, 2020) study the gradient flow dynamics of the MF limit and show that they are well-defined in general settings and lead to convergence results (local or global depending on the assumptions). On the other hand, Jacot et al. (2018) study the dynamic of the NTK parameterization in the infinite-width limit and show that it amounts to learning a *linear* predictor on top of random features (fixed kernel), so that there is no feature learning.

¹In contrast, our focus on GF is only for theoretical convenience and most of our results could be adapted to the case of GD.

Convergence rates. In the MF limit, convergence rates are in general difficult to obtain in a standard setting. For instance, [Chizat and Bach \(2018\)](#); [Wojtowytsch \(2020\)](#) show the convergence of the GF to a global optimum in a general setting but this does not allow convergence rates to be provided. To illustrate the convergence of the parameterizing measure to a global optimum in the MF limit, [Ma et al. \(2020\)](#) prove *local* convergence (see Section 7) for one-dimensional inputs and a specific choice of target function in $O(t^{-1})$ where t is the time step. At finite-width, [Daneshmand and Bach \(2022\)](#) also prove convergence of the parameters to a global optimum in $O(t^{-1})$ using an algebraic idea which is specific to the ad-hoc structure they consider (inputs in two dimensions and target functions with finite number of atoms).

In Section 3, we show convergence of the MF limit at an exponential rate when the target function is odd. In the setting of this section, the training dynamics are degenerate and although input neurons move, the symmetries of the problem imply that the predictor is linear.

Low-dimensional structure. Studying how neural networks can adapt to hidden low-dimensional structures is a way of approaching theoretically the feature-learning abilities of neural networks. [Bach \(2017\)](#) studies the statistical properties of infinitely wide two-layer networks, and shows that when the target function only depends on the projection on a low-dimensional sub-space, these networks circumvent the curse of dimensionality with generalization bounds which only depend on the dimension of the sub-space. In a slightly different context, [Chizat and Bach \(2020\)](#) show that for a binary classification task, when there is a low-dimensional sub-space for which the projection of the data has sufficiently large inter-class distance, only the dimension of the sub-space (and not that of the ambient space) appears in the upper bound on probability of misclassification. Whether or not such a low-dimensional sub-space is actually learned by GD is not addressed in these works.

Similarly, [Cloninger and Klock \(2021\)](#); [Damian et al. \(2022\)](#) focus on learning functions which have a hidden low-dimensional structure with neural networks. They consider a single step of GD on the input layer weights and show that the approximation / generalization error adapts to the structure of the problem: they provide bounds on the number of data points / parameters needed to achieve negligible error, which depend on the reduced dimension and not the dimension of the ambient space.

In a setting close to ours but on a classification task with finite-data and at finite-width, [Paccolat et al. \(2021\)](#) compare the feature learning regime with the NTK regime in the presence of hidden low-dimensional structure and quantify for each regime the scaling law of the test error *w.r.t.* the number of training samples, mostly focusing on the case $d_H = 1$.

In a similar setting to that of ([Bach, 2017](#)), [Abbe et al. \(2022\)](#) study how GF for infinitely wide two-layer networks can learn specific classes of functions which have a hidden low-dimensional structure when the inputs are Rademacher variables. This strong symmetry assumption ensures that the learned predictor shares the same low-dimensional structure at any time step (from the $t = 0$) and this allows them to characterize precisely what classes of target functions can or cannot be learned by GF in this setting. In contrast, we are interested in how infinitely wide networks *learn* those low-dimensional structures during training, and in the role of symmetries in enabling such a behaviour after initialization.

Learning representations. An existing line of work ([Yehudai and Shamir, 2019](#); [Allen-Zhu et al., 2019](#); [Abbe et al., 2021](#); [Damian et al., 2022](#); [Ba et al., 2022](#)) studies in depth the representations learned by neural networks trained with (S)GD at finite-width from a different perspective focusing on the advantages of feature-learning in terms of performance comparatively to using random features. In contrast, our aim is to describe the representations themselves in relationship with the symmetries of the problem.

Symmetries. We stress that the line of work around symmetries of neural networks dealing with finding network architectures for which the output is invariant (*w.r.t.* to its input or parameters) by some *group* of transformations (see Bloem-Reddy and Teh, 2020; Ganev and Walters, 2021; Gluch and Urbanke, 2021, and references therein) is entirely different from what we are concerned with in the present work. In contrast, the setting of (Mei et al., 2018) is much closer to ours as they study how the invariances of the target function / input data can lead to simplifications in the dynamics of infinitely wide two-layer networks in the mean-field regime which allows them to prove global convergence results.

1.4 Notations

We denote by $\mathcal{M}_+(\Omega)$ the space of non-negative measures over a measurable space Ω . For any measure μ and measurable map T , $T_{\#}\mu$ denotes the pushforward measure of μ by T . We denote by $\mathcal{O}(p)$ and $\text{id}_{\mathbb{R}^p}$ respectively the orthogonal group and the identity map of \mathbb{R}^p for any $p \in \mathbb{N}$. Finally, $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and $\|\cdot\|$ the corresponding norm.

2 Invariance under orthogonal symmetries

In this section, we demonstrate that if the target function f^* is invariant under some orthogonal transformation T , since the input data distribution is also invariant under T , then $f(\mu_t; \cdot)$ is invariant under T as well for any $t \geq 0$. This stability property of the dynamics *w.r.t.* orthogonal symmetries is possible with an infinite number of neurons but is only approximate at finite-width. It is noteworthy that the results of this section hold for any activation function σ and input data distribution ρ which has the same symmetries as f^* , provided that the Wasserstein GF (1.3) is unique. We start with a couple of definitions:

Definition 2.1 (Function invariance). Let T be a map from \mathbb{R}^d to \mathbb{R}^d , and $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Then, f is said to be *invariant* (resp. *anti-invariant*) under T if for any $x \in \mathbb{R}^d$, $f(T(x)) = f(x)$ (resp. $f(T(x)) = -f(x)$).

Definition 2.2 (Measure invariance). Let $\Omega \subset \mathbb{R}^d$, T be a measurable map from Ω to Ω , and μ be a measure on Ω . Then, μ is said to be invariant under T if $T_{\#}\mu = \mu$, or equivalently, if for any measurable $\varphi : \Omega \rightarrow \mathbb{R}$, $\int \varphi(x)d\mu(x) = \int \varphi(T(x))d\mu(x)$.

We are now ready to state the two main results of this section.

Proposition 2.1 (Learning invariance). Let $T \in \mathcal{O}(d)$, and assume that f^* , ρ and μ_0^2 are invariant under T . Then, for any $t \geq 0$, the Wasserstein GF μ_t of Equation (1.3) is invariant under $\tilde{T} : (a, b) \in \mathbb{R} \times \mathbb{R}^d \mapsto (a, T(b))$, and the corresponding predictor $f(\mu_t; \cdot)$ is invariant under T .

Proposition 2.2 (Learning anti-invariance). Under the same assumptions as in Proposition 2.1 except now we assume f^* is anti-invariant under T , and assuming further that $\partial_2 \ell(-y, -\hat{y}) = -\partial_2 \ell(y, \hat{y})$ for any $y, \hat{y} \in \mathbb{R}$, and that μ_0^1 is symmetric around 0 (i.e., invariant under $a \in \mathbb{R} \mapsto -a$), we then have that for any $t \geq 0$, the Wasserstein GF μ_t in Equation (1.3) is invariant under $\tilde{T} : (a, b) \in \mathbb{R} \times \mathbb{R}^d \mapsto (-a, T(b))$, and the corresponding predictor $f(\mu_t; \cdot)$ is anti-invariant under T .

Remark. The results above also hold for networks with intercepts at both layers. In addition, the conditions of Proposition 2.2 are satisfied by both the squared loss and the logistic loss (*a.k.a.* the cross-entropy loss).

Essentially, those results show that training with GF preserves the orthogonal symmetries

of the problem: the invariance of the target function under an orthogonal transformation leads to the same invariance for μ_t and $f(\mu_t; \cdot)$. The proof, presented in Appendix C, relies crucially on the fact that T is an orthogonal map which combines well with the structure of $\phi(c; x)$ involving an inner product. The idea is essentially that the orthogonality of T allows us to relate the gradient of ϕ (and consequently of F'_{μ_t}) w.r.t. c at $(T(c); x)$ to the same gradient at $(c; T^{-1}(x))$ and then to use the invariance of f^* and ρ to conclude.

In the following sections we discuss the particular cases where functions are (anti-)invariant under $-\text{id}_{\mathbb{R}^d}$ (i.e., even or odd functions) or some sub-group of $\mathcal{O}(d)$.

3 Exponential convergence for odd target functions

We consider here an odd target, function, i.e., for any $x \in \mathbb{R}^d$, $f^*(-x) = -f^*(x)$, and we make the following assumptions:

Assumption 1 (Squared loss function). The loss function ℓ is the squared loss, i.e., $\ell(y, \hat{y}) = \frac{1}{2}(y - \hat{y})^2$, and thus satisfies the condition of Proposition 2.2.

Linearity of odd predictors. Proposition 2.2 ensures that the predictor $f(\mu_t; \cdot)$ associated with the Wasserstein GF of Equation (1.3) is also odd at any time $t \geq 0$, and we can thus write, for any x , $f(\mu_t; x) = \frac{1}{2}(f(\mu_t; x) - f(\mu_t; -x))$, which yields

$$f(\mu_t; x) = \frac{1}{2} \left(\int_{a,b} a \left[\sigma(b^\top x) - \sigma(-b^\top x) \right] d\mu_t(a, b) \right) = \frac{1}{2} \int_{a,b} a \left(b^\top x \right) d\mu_t(a, b),$$

where the last equality stems from the fact that for ReLU, $\sigma(x) - \sigma(-x) = x$. Put differently, the predictor is **linear**: it is the same as replacing σ by $\frac{1}{2}\text{id}_{\mathbb{R}^d}$, and $f(\mu_t; x) = w(t)^\top x$, where

$$w(t) := \frac{1}{2} \int_{a,b} a b d\mu_t(a, b) \in \mathbb{R}^d. \quad (3.1)$$

This degeneracy is not surprising as in fact, a linear predictor is the best one can hope for in this setting. Indeed, consider the following lemma:

Lemma 3.1 (Optimality of odd predictors). *Let f be a predictor in the hypothesis class $\mathcal{F} := \{x \mapsto \int a \sigma(b^\top x) d\mu(a, b); \mu \in \mathcal{P}_2(\mathbb{R} \times \mathbb{R}^d)\}$. Then, denoting $f_{\text{odd}}(x) := \frac{1}{2}(f(x) - f(-x))$ (resp. $f_{\text{even}} := \frac{1}{2}(f(x) + f(-x))$) the odd (resp. even) part of f , one has:*

- (i) $f_{\text{odd}} \in \mathcal{F}$,
- (ii) $L(f) := \mathbb{E}_{x \sim \rho} \left[(f^*(x) - f(x))^2 \right] \geq \mathbb{E}_{x \sim \rho} \left[(f^*(x) - f_{\text{odd}}(x))^2 \right] =: L(f_{\text{odd}})$,
- (iii) equality holds if and only if f is odd ρ -almost surely.

Proof. The result readily follows from the decomposition $f = f_{\text{odd}} + f_{\text{even}}$ which leads to

$$L(f) = L(f_{\text{odd}}) + \underbrace{\mathbb{E}_{x \sim \rho} \left[(f_{\text{even}}(x))^2 \right]}_{\geq 0} - \underbrace{2 \mathbb{E}_{x \sim \rho} \left[(f^*(x) - f_{\text{odd}}(x)) f_{\text{even}}(x) \right]}_{0 \text{ by symmetry}}.$$

We then get that $L(f) \geq L(f_{\text{odd}})$ with equality if and only if $\mathbb{E}_{x \sim \rho} \left[(f_{\text{even}}(x))^2 \right] = 0$, i.e., $f_{\text{even}}(x) = 0$ for ρ -almost every x . Finally, if $\mu \in \mathcal{P}_2(\mathbb{R}^{d+1})$, then $\nu := \frac{1}{2}(\mu + S\#\mu) \in \mathcal{P}_2(\mathbb{R}^{d+1})$, where $S : (a, b) \in \mathbb{R}^{d+1} \mapsto (-a, -b)$, and $f(\nu; \cdot) = f_{\text{odd}}(\mu; \cdot)$, which shows $f_{\text{odd}}(\mu; \cdot) \in \mathcal{F}$. \square

Since, as shown above, any odd predictor turns out to be linear because of the symmetries of ReLU, in this context, the best one can expect is thus to learn the best linear predictor.

Exponential convergence for linear networks. We are thus reduced to studying the dynamics of linear networks (which in our case are infinitely wide), which is an interesting object of study in its own right (Ji and Telgarsky, 2018 show a result similar to our result below in the finite-width case). In this case, the Wasserstein GF (1.3) (with ReLU replaced by $\frac{1}{2}\text{id}_{\mathbb{R}^d}$) is defined for more general input distributions $\mathbb{P} \in \mathcal{P}_2(\mathbb{R}^d)$ (e.g., empirical measures) and target functions f^* . The objective in this context is thus to learn:

$$w^* \in \underset{w \in \mathbb{R}^d}{\text{argmin}} \left(Q(w) := \frac{1}{2} \mathbb{E}_{x \sim \mathbb{P}} \left[(f^*(x) - \langle w, x \rangle)^2 \right] \right) \quad (3.2)$$

with the dynamics of linear infinitely wide two-layer networks described by the Wasserstein GF (1.3) where the activation function σ is replaced by $\frac{1}{2}\text{id}_{\mathbb{R}^d}$. Theorem 3.2 below shows exponential convergence to a global minimum of Q as soon as the problem is strongly convex. Note that although in this case both $\phi(\cdot; \cdot)$ and the predictor in the objective Q are linear w.r.t. the input, only the predictor in Q is linear in the parameters (ordinary least squares).

Theorem 3.2. *Assume that the smallest eigenvalue λ_{\min} of $\mathbb{E}_{x \sim \mathbb{P}}[xx^\top]$ is positive. Let $(\mu_t)_{t \geq 0}$ be the Wasserstein GF associated to (1.3) with activation function $\frac{1}{2}\text{id}_{\mathbb{R}^d}$ instead of $\sigma = \text{ReLU}$, and call $w(t) = \frac{1}{2} \int ab d\mu_t(a, b) \in \mathbb{R}^d$. Then, there exists $\eta > 0$ and $t_0 > 0$ such that, for any $t \geq t_0$,*

$$\left(Q(w(t)) - Q(w^*) \right) \leq e^{-2\eta\lambda_{\min}(t-t_0)} \left(Q(w(t_0)) - Q(w^*) \right).$$

Remark. Note that as soon as \mathbb{P} has spherical symmetry, the problem becomes strongly convex by Lemma A.3. Note that although $F(\mu_t) = Q(w(t))$, $(w(t))_{t \geq 0}$ is **not** a gradient flow for the (strongly) convex objective Q (which would immediately guarantee exponential convergence to the global minimum).

The proof, provided in Appendix D, proceeds in two steps: first it is shown that $w'(t) = -H(t)\nabla Q(w(t))$ for some positive definite matrix $H(t)$ whose smallest eigenvalue is always lower-bounded by a positive quantity, then we prove that this leads to exponential convergence. Figure 1 illustrates that the dynamics of GF on F remain non-linear in that they do not reduce to GF on Q (although the paths are close). The dashed curves are the trajectories of $w(t)$ (see Theorem 3.2) for 10 random initializations, which corresponds to GD on F , whereas the blue curve is the trajectory of GD on Q . $\mathbb{P} = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$ is the empirical distribution over $n = 100$ samples, with $x_i \sim \mathcal{U}([-1, 1]^d)$ and $y_i = f^*(x_i)$ drawn i.i.d. from $\mathcal{N}(y^*, 2)$ where $y^* \sim \mathcal{N}(0, 1)$; the learning rate is 10^{-2} .

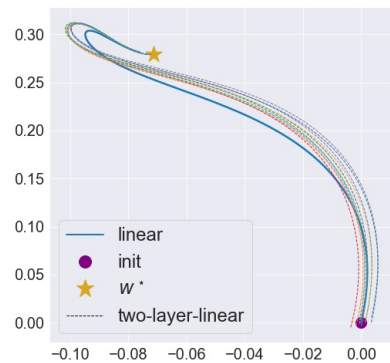


Figure 1: GD path for two coordinates: two-layer linear network vs pure linear model.

4 Learning the low-dimensional structure of the problem

Consider a linear sub-space H of dimension $d_H < d$ (potentially much smaller than the ambient dimension), and assume f^* has the following structure: $f^*(x) = f_H(p_H(x))$ where p_H is the orthogonal projection onto H (which we also write x^H for simplicity, and we reserve sub-scripts for denoting entries of vectors) and $f_H : H \rightarrow \mathbb{R}$ is a given function.

In this context it is natural to study whether the learned function shares the same structure as f^* . As observed in Figure 2 this is not the case in finite time, but it is reasonable

however to think that the learned predictor $f(\mu_t; \cdot)$ shares the same structure as f^* as $t \rightarrow \infty$, and we give numerical evidence in this direction. On the other hand, we prove rigorously that the structure of the problem allows to reduce the dynamics to a lower-dimensional PDE. In this section, we consider for simplicity that μ_0^1 is the uniform distribution over $\{-1, +1\}$ and that μ_0^2 is the uniform distribution over \mathbb{S}^{d-1} .

4.1 Symmetries and invariance

The structure of f^* implies that it is invariant by any $T \in \mathcal{O}(d)$ which preserves H , *i.e.*, such that the restrictions to H and H^\perp are $T|_H = \text{id}_H$ and $T|_{H^\perp} \in \mathcal{O}(d_\perp)$, where $\mathcal{O}(d_\perp)$ is the orthogonal group of H^\perp whose dimension is $d_\perp = d - d_H$. By Proposition 2.1, such transformations also leave the predictor $f(\mu_t; \cdot)$ invariant for any $t \geq 0$ since ρ is spherically symmetric. Lemma 4.1 below then ensures that $f(\mu_t; x)$ depends on the projection x^\perp onto H^\perp only through its norm, that is $f(\mu_t; x) = \tilde{f}_t(x^H, \|x^\perp\|)$, $\tilde{f}_t : H \times \mathbb{R}_+ \rightarrow \mathbb{R}$.

Lemma 4.1 (Invariance by a sub-group of $\mathcal{O}(d)$). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be invariant under any $T \in \mathcal{O}(d)$ such that $T|_H = \text{id}_H$ and $T|_{H^\perp} \in \mathcal{O}(d_\perp)$. Then, there exists some $\tilde{f} : H \times \mathbb{R}_+ \rightarrow \mathbb{R}$ such that for any $x \in \mathbb{R}^d$, $f(x) = \tilde{f}(x^H, \|x^\perp\|)$.*

Proof. Consider $\tilde{f} : (x^H, r) \in H \times \mathbb{R}_+ \mapsto f(x^H + re_1^\perp)$ where e_1^\perp is the first vector of an orthonormal basis of H^\perp , and let $x \in \mathbb{R}^d$. If $x^\perp = 0$, the result is obvious. Otherwise, consider an orthogonal linear map T_x such that $T_x|_H = \text{id}_H$ and T_x sends $x^\perp/\|x^\perp\|$ on e_1^\perp . The invariance of f under T_x implies $f(x) = f(T_x(x)) = f(x^H + \|x^\perp\|e_1^\perp) = \tilde{f}(x^H, \|x^\perp\|)$. \square

Figure 2 shows that the dependence in $\|x^\perp\|$ cannot be removed in finite time: $f(\mu_t; u_H + re_1^\perp)$ does depend on the distance $r \in \mathbb{R}_+$ to H , but this dependence tends to vanish as $t \rightarrow \infty$.




Figure 2: $f(\mu_t; u_H + re_1^\perp)$ vs r and t for a random $u_H \in \mathbb{S}^{d_H-1}$ with $d = 20$, $d_H = 5$.

Dynamics over the sphere \mathbb{S}^{d-1} . Using the positive 1-homogeneity of ReLU, and with the assumption on μ_0 , the dynamics on $\mu_t \in \mathcal{P}_2(\mathbb{R}^{d+1})$ can be reduced to dynamics on the space $\mathcal{M}_+(\mathbb{S}^{d-1})$ of non-negative measures over \mathbb{S}^{d-1} : only the direction of neurons matter and their norm only affects the total mass. From this point of view, neurons with positive and negative output weights behave differently and have separate dynamics. Indeed, consider the pair of measures $(\nu_t^+, \nu_t^-) \in \mathcal{M}_+(\mathbb{S}^{d-1})^2$ characterized by the property that for any continuous test function $\varphi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$,

$$\int_u \varphi(u) d\nu_t^\pm(u) = \int_{\pm a \geq 0, b} |a| \|b\| \varphi\left(\frac{b}{\|b\|}\right) d\mu_t(a, b), \quad (4.1)$$

where we have used the superscript \pm to denote either ν_t^+ or ν_t^- and the right-hand side is changed accordingly depending on the sign $+$ or $-$. Because ReLU is positively 1-homogeneous, we have $f(\mu_t; x) = \int \sigma(u^\top x) d(\nu_t^+ - \nu_t^-)(x)$. It is shown in Appendix E.1.1

that ν_t^\pm satisfies, in the sense of distributions, the equation

$$\partial_t \nu_t^\pm = -\operatorname{div}(\pm \tilde{v}_t \nu_t^\pm) \pm 2g_t \nu_t^\pm, \quad (4.2)$$

where, for any $u \in \mathbb{S}^{d-1}$,

$$\begin{aligned} g_t(u) &= - \int_y \partial_2 \ell(f^*(y), f(\mu_t; y)) \sigma(u^\top y) d\rho(y), \\ \tilde{v}_t(u) &= - \int_y \partial_2 \ell(f^*(y), f(\mu_t; y)) \sigma'(u^\top y) [y - (u^\top y)u] d\rho(y). \end{aligned} \quad (4.3)$$

This equation can be interpreted as a Wasserstein-Fisher-Rao GF (Gallouët et al., 2019) on the sphere since $\tilde{v}_t(u) = \operatorname{proj}_{\{u\}^\perp}(\nabla g_t(u))$.

Closed dynamics over $[0, \pi/2] \times \mathbb{S}^{d_H-1}$. The dynamics on the pair (ν_t^+, ν_t^-) can be further reduced to dynamics over $[0, \pi/2] \times \mathbb{S}^{d_H-1}$. Indeed, by positive 1-homogeneity of $f(\mu_t; \cdot)$ we may restrict ourselves to inputs $u \in \mathbb{S}^{d-1}$, and $f(\mu_t; u)$ depends only on u^H and $\|u^\perp\|$. However, because $\|u^H\|^2 + \|u^\perp\|^2 = 1$, this dependence translates into a dependence on the direction $u^H/\|u^H\|$ of the projection onto H and the norm $\|u^H\|$. The former is an element of \mathbb{S}^{d_H-1} while the latter is given by the angle θ between u and H , that is $\theta := \arccos(u^\top u^H/\|u^H\|) = \arccos(\|u^H\|)$. This simplification leads to the following lemma:

Lemma 4.2. *Define the measures τ_t^+, τ_t^- by $\tau_t^\pm = P_{\#} \nu_t^\pm \in \mathcal{M}_+([0, \pi/2] \times \mathbb{S}^{d_H-1})$ via $P : u \in \mathbb{S}^{d-1} \setminus H^\perp \mapsto (\arccos(\|u_H\|), u_H/\|u_H\|) \in [0, \pi/2] \times \mathbb{S}^{d_H-1}$. Then, the measures τ_t^+, τ_t^- satisfy the equation*

$$\partial \tau_t^\pm = -\operatorname{div}(\pm V_t \tau_t^\pm) \pm 2G_t \tau_t^\pm, \quad (4.4)$$

where $G_t : [0, \pi/2] \times \mathbb{S}^{d_H-1} \rightarrow \mathbb{R}$, and $V_t : [0, \pi/2] \times \mathbb{S}^{d_H-1} \rightarrow \mathbb{R}^{d_H+1}$ are functions depending **only** on (τ_t^+, τ_t^-) , and furthermore, $f(\mu_t; \cdot)$ can be expressed solely using τ_t^+, τ_t^- (exact formulas are provided in Appendix E.1.2).

Abbe et al. (2022) show a similar result with a lower-dimensional dynamics in the context of infinitely wide two-layer networks when the input data have i.i.d coordinates distributed uniformly over $\{-1, +1\}$ (i.e., Rademacher variables), except that they do not have the added dimension due to the angle θ as we do thanks to their choice of input data distribution.

Lemma 4.2 above illustrates how the GF dynamics of infinitely wide two-layer networks adapts to the lower-dimensional structure of the problem: the learned predictor and the dynamics can be described only in terms of the angle θ between the input neurons and H and that of their projection on the unit sphere of H .

4.2 One dimensional reduction

Since the predictors we consider are positively homogeneous, one cannot hope to do better than learn a positively homogeneous function. A natural choice of such a target function to learn is the Euclidean norm. With the additional structure that the target only depends on the projection onto H , this leads to considering $f^*(x) = \|x^H\|$ which has additional symmetries compared to the general case presented above: it is invariant by any linear map T such that $T|_H \in \mathcal{O}(d_H)$ and $T|_{H^\perp} \in \mathcal{O}(d_\perp)$. By Proposition 2.1 those symmetries are shared by μ_t and $f(\mu_t; \cdot)$, and we show that in this case the dynamic reduces to a one-dimensional dynamic over the angle θ between input neurons and H .

We prove a general disintegration result for the uniform measure on the sphere in the Appendix (see Lemma A.4) which allows, along with some spherical harmonics analysis, to describe the reduced dynamics and characterize the objective that they optimize. This leads to the following result:

Theorem 4.3 (1d dynamics over the angle θ). *Assume that $f^*(x) = \|x^H\|$, and define the measures $(\tau_t^+, \tau_t^-) \in \mathcal{M}_+([0, \pi/2])^2$ from (ν_t^+, ν_t^-) via $P : u \in \mathbb{S}^{d-1} \mapsto \arccos(\|u_H\|) \in [0, \pi/2]$: $\tau_t^\pm = P_{\#}\nu_t^\pm$. Then, the pair (τ_t^+, τ_t^-) follows the Wasserstein-Fisher-Rao GF for the objective $A(\tau^+, \tau^-) := (1/2)\mathbb{E}[(f^*(x) - f(\tau^+, \tau^-; x))^2]$ over the space $\mathcal{M}_+([0, \pi/2]) \times \mathcal{M}_+([0, \pi/2])$, where $f(\tau^+, \tau^-; x)$ is the expression (with a slight overloading of notations) of $f(\mu; x)$ in function of (τ^+, τ^-) (see Appendix E.2 for more details):*

$$\begin{aligned} d\tau_0^\pm(\theta) &= \frac{1}{B\left(\frac{d_H}{2}, \frac{d_\perp}{2}\right)} \cos(\theta)^{d_H-1} \sin(\theta)^{d_\perp-1} d\theta, \\ \partial_t \tau_t^\pm &= -\operatorname{div}(\pm V_t \tau_t^\pm) \pm 2G_t \tau_t^\pm, \end{aligned} \quad (4.5)$$

where B is the Beta function, and

$$\begin{aligned} G_t(\theta) &= -\int_y \partial_2 \ell\left(f^*(y), f(\mu_t; y)\right) \sigma\left(\cos(\theta)y_1^H + \sin(\theta)y_1^\perp\right) d\rho(y), \\ V_t(\theta) &= G_t'(\theta). \end{aligned}$$

Additionally, $f(\mu_t; \cdot)$, G_t , and V_t **only depend** on the pair (τ_t^+, τ_t^-) , and for any $t \geq 0$, it holds that $F(\mu_t) = A(\tau_t^+, \tau_t^-)$.

Remark. The result should still hold for general ρ which are spherically symmetric as long as the Wasserstein GF (1.3) is well-defined but the proof is more technical. In addition, this result shows that even with more structure than in Lemma 4.2, the dynamics of infinitely wide two-layer networks are still able to adapt to this setting: these dynamics, as well as the learned predictor, can be *fully characterized* solely by the one-dimensional dynamics over the angle θ between input neurons and H . This is noteworthy since this angle determines the alignment of the neurons with H , and thus measures how much the representations learned by the network have adapted to the structure of the problem. Furthermore, as discussed below, this reduction with exact formulas enables efficient numerical simulation in one dimension.

Daneshmand and Bach (2022) prove the global convergence of a reduced one-dimensional dynamics in a context similar to ours but their original problem is two-dimensional and with a choice of activation function that leads to specific algebraic properties.

Expression of $f(\mu_t; \cdot)$. Because of the symmetries of $f(\mu_t; \cdot)$, which result from that of f^* , $f(\mu_t; x)$ depends only on $\|x^H\|$ and $\|x^\perp\|$. What is more, since $f(\mu_t; \cdot)$ is positively 1-homogeneous (because ReLU is) it actually holds that $f(\mu_t; x) = \|x\| \tilde{f}_t(\varphi_x)$ where $\varphi_x = \arccos(\|x^H\|/\|x\|)$ is the angle between x and H , and $\tilde{f}_t(\varphi) := \int_\theta \tilde{\phi}(\theta; \varphi) d(\tau_t^+ - \tau_t^-)(\theta)$, $\tilde{\phi}$ depending only on σ and fixed probability measures (see Appendix E.2.3 for an exact formula).

Learning the low-dimensional structure as $t \rightarrow \infty$. Although, as shown in Figure 2, $f(\mu_t; \cdot)$ does not learn the low-dimensional structure in finite-time, it is reasonable to expect that as $t \rightarrow \infty$, the measures τ_t^\pm put mass only on $\theta = 0$, indicating that the only part of the space that the predictor is concerned with for large t is the sub-space H . Since we assume here that the target function f^* is non-negative, the most natural limits for τ_t^+ and τ_t^- are $\tau_t^+ \rightarrow \alpha \delta_0$ with $\alpha > 0$, and $\tau_t^- \rightarrow 0$ (in the sense that $\tau_t^-([0, \pi/2]) \rightarrow 0$) as $t \rightarrow \infty$, because then the “negative” output weights do not participate in the prediction in the large t limit.

The global convergence result of Chizat and Bach (2018); Wojtowytsch (2020) still holds but is not quantitative and moreover does not guarantee that the limit is the one described above. We leave the proof of this result as an open problem, but we provide numerical evidence supporting this conjecture. Indeed, we take advantage of the one-dimensional reduction from Theorem 4.3, and numerically simulate the resulting dynamics by parameterizing

τ_t^\pm via weight and position (Chizat, 2022) as $(1/m) \sum_{j=1}^m c_j^\pm(t) \delta_{\theta_j^\pm(t)}$, and simulating the corresponding dynamics for $c_j^\pm(t)$ and $\theta_j^\pm(t)$ (see more details in Appendix F). Figures 3a and 3b show that the mass of τ_t^+ tends to concentrate around 0 while that of τ_t^- tends to concentrate around $\pi/2$, indicating that τ_t^+ adapts to the part of the space relevant to learning f^* while τ_t^- puts mass close to the orthogonal to that space.

Total mass of particles at convergence. If $\tau_\infty^- = 0$ and $\tau_\infty^+ = \alpha \delta_0$ as described above, we have $f(\mu_\infty; x) = \alpha \|x\| \tilde{\phi}(0; \varphi_x) = \alpha \frac{\Gamma(d_H/2)}{2\sqrt{\pi}\Gamma((d_H+1)/2)} \|x\| \cos(\varphi_x) = \frac{\alpha \Gamma(d_H/2)}{2\sqrt{\pi}\Gamma((d_H+1)/2)} \|x^H\|$. To recover exactly f^* , it must hold that $\alpha = \tau_\infty^+([0, \pi/2]) = \frac{2\sqrt{\pi}\Gamma((d_H+1)/2)}{\Gamma(d_H/2)}$. Defining the normalized probability measure $\tilde{\tau}_t^\pm = \tau_t^\pm / \tau_t^\pm([0, \pi/2])$, we thus expect $\tilde{\tau}_t^+$ to grow close to δ_0 and $\tilde{\tau}_t^-$ to $\delta_{\pi/2}$. In terms of total mass, we expect that $\tau_t^+([0, \pi/2])$ gets closer to α while $\tau_t^-([0, \pi/2])$ gets closer to 0.

The numerical behaviour depicted in Figure 3c seems to follow our intuitive description, at least until a critical time t^* in the numerical simulation which corresponds to the first time t where $\tau_t^+([0, \pi/2]) > \alpha$. While the total mass of τ_t^\pm (dashed lines) seems to approach its limit rapidly before t^* it slowly moves further away from it for $t \geq t^*$. On the other hand, while the angles only slowly change before t^* , they start converging fast towards the corresponding Dirac measures after t^* . It is unclear whether this slight difference in behaviour (around the critical time t^*) between what we intuitively expected and the numerical simulation is an artefact of the finite width and finite step size or if it actually corresponds to some phenomenon present in the limiting model. For more details concerning the numerical experiments, see Appendix F.

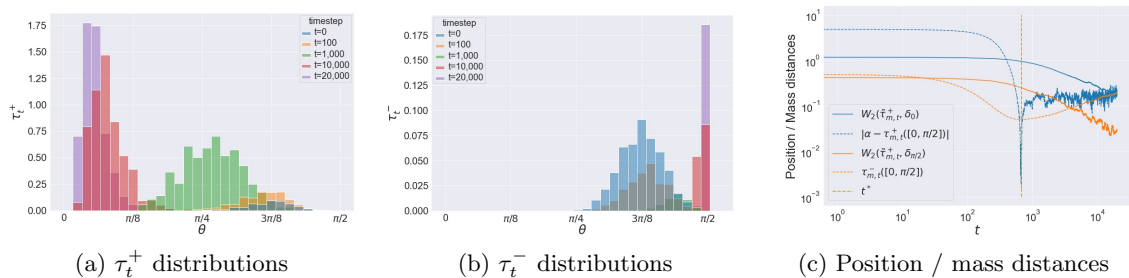


Figure 3: Angle distributions τ_t^+ / τ_t^- and position / mass distances with $m = 1024$, $d = 30$ and $d_H = 5$. (a) (*resp.* (b)) τ_t^+ (*resp.* τ_t^-) as a histogram for different t . (c) distances (in log-log scales) of the mass and positions of positive (blue) / negative (orange) particles to the intuitively expected limits: the distance in position the Wasserstein-2 distance of the normalized (probability) measures $\tilde{\tau}_t^\pm$ to the corresponding Dirac measures while the distance in mass is the absolute error to the expected mass as $t \rightarrow \infty$.

Note that there is *a priori* not a unique global optimum: τ_∞^+ and τ_∞^- (if they exist) can compensate on parts of the space $[0, \pi/2]$ and lead to the same optimal predictor for different choices of measures. Our numerical experiments suggest that the GF dynamics select a “simple” solution where τ_∞^+ is concentrated on $\{\theta = 0\}$ and τ_∞^- vanishes (puts 0 mass everywhere), which is a form of implicit bias.

5 Conclusion

We have explored the symmetries of infinitely wide two-layer ReLU networks and we have seen that: (i) they adapt to the orthogonal symmetries of the problem, (ii) they reduce to

the simple dynamics of a linear network in the case of an odd target function and (iii) when the target function depends only on the orthogonal projection onto a lower-dimensional subspace H , the dynamics can be reduced to a lower-dimensional PDE. In particular, when f^* is the Euclidean norm, this PDE is over a one dimensional space corresponding to the angle θ between the particles and H . We have presented numerical experiment indicating that the positive particles converge to the subspace H in this case and leave the proof of this result as an open problem. We also leave as an open question whether the results of Section 2 extend to deeper networks.

Acknowledgements

Karl Hajjar received full support from the Agence Nationale de la Recherche (ANR), reference ANR-19-CHIA-0021-01 “BiSCottE”. The authors thank Christophe Giraud for useful discussions and comments.

References

- Emmanuel Abbe, Enric Boix-Adsera, Matthew S. Brennan, Guy Bresler, and Dheeraj Nagaraj. The staircase property: How hierarchical structure can guide deep learning. *Advances in Neural Information Processing Systems*, 34:26989–27002, 2021.
- Emmanuel Abbe, Enric Boix-Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. *arXiv preprint arXiv:2202.08658*, 2022.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 6158–6169, 2019.
- Kendall Atkinson and Weimin Han. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, volume 2044. Springer, 01 2012. ISBN 978-3-642-25982-1. doi: 10.1007/978-3-642-25983-8.
- Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. *arXiv preprint arXiv:2205.01445*, 2022.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Benjamin Bloem-Reddy and Yee Whye Teh. Probabilistic symmetries and invariant neural networks. *J. Mach. Learn. Res.*, 21:90–1, 2020.
- Nick Cammarata, Shan Carter, Gabriel Goh, Chris Olah, Michael Petrov, Ludwig Schubert, Chelsea Voss, Ben Egan, and Swee Kiat Lim. Thread: Circuits. *Distill*, 2020. doi: 10.23915/distill.00024. <https://distill.pub/2020/circuits>.
- Lenaïc Chizat. Sparse optimization on measures with over-parameterized gradient descent. *Mathematical Programming*, 194(1):487–532, 2022.
- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3040–3050, 2018.

- Lenaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/ae614c557843b1df326cb29c57225459-Paper.pdf>.
- Alexander Cloninger and Timo Klock. A deep network construction that adapts to intrinsic dimensionality beyond the domain. *Neural Networks*, 141:404–419, 2021.
- Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.
- Hadi Daneshmand and Francis Bach. Polynomial-time sparse measure recovery. *arXiv preprint arXiv:2204.07879*, 2022.
- Thomas Gallouët, Maxime Laborde, and Leonard Monsaingeon. An unbalanced optimal transport splitting scheme for general advection-reaction-diffusion problems. *ESAIM: Control, Optimisation and Calculus of Variations*, 25:8, 2019.
- Iordan Ganev and Robin Walters. The qr decomposition for radial neural networks. *arXiv preprint arXiv:2107.02550*, 2021.
- Grzegorz Głuch and Rüdiger Urbanke. Noether: The more things change, the more stay the same. *arXiv preprint arXiv:2104.05508*, 2021.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *CoRR*, abs/1806.07572, 2018. URL <http://arxiv.org/abs/1806.07572>.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. *arXiv preprint arXiv:1810.02032*, 2018.
- Chao Ma, Lei Wu, et al. Machine learning from a continuous viewpoint, i. *Science China Mathematics*, 63(11):2233–2266, 2020.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multilayer neural networks. *CoRR*, abs/2001.11443, 2020. URL <https://arxiv.org/abs/2001.11443>.
- Jonas Paccolat, Leonardo Petrini, Mario Geiger, Kevin Tyloo, and Matthieu Wyart. Geometric compression of invariant manifolds in neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(4):044001, apr 2021. doi: 10.1088/1742-5468/abf1f3. URL <https://dx.doi.org/10.1088/1742-5468/abf1f3>.

- Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55 (58-63):94, 2015.
- Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020.
- Stephan Wojtowytsch. On the convergence of gradient descent training for two-layer relu-networks in the mean field regime. *arXiv preprint arXiv:2005.13530*, 2020.
- Greg Yang and Edward J. Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 11727–11737. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21c.html>.
- Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.

Appendix

A Additional notations and preliminary results

A.1 Notations for the appendix

We introduce in this section additional notation that we use throughout the Appendix.

Residual: we call $R_t(y) := -\partial_2 \ell(f^*(y), f(\mu_t; y))$, the “residual”, which is equal to the difference $f^*(y) - f(\mu_t; y)$ when ℓ is the squared loss.

Identity matrix: we denote by I_p the identity matrix in $\mathbb{R}^{p \times p}$ for any $p \in \mathbb{N}$.

Indicator functions: we denote by $\mathbf{1}_A$ the indicator of a set A , that is $\mathbf{1}_A(z) = 1 \iff z \in A$, and $\mathbf{1}_A(z) = 0$ otherwise.

Total variation: for any measure ν , we denote by $|\nu|$ its total variation, which should cause no confusion with the absolute value given the context.

Beta / Gamma function and distribution: for $\alpha, \beta > 0$, we denote by $B(\alpha, \beta)$ the Beta function equal to $\Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$ where Γ is the Gamma function, and by $\text{Beta}(\alpha, \beta)$ the beta law with density equal to $u^{\alpha-1}u^{\beta-1}/B(\alpha, \beta)$ on $[0, 1]$.

Gaussian / spherical measures: we call ρ_p the standard Gaussian measure in \mathbb{R}^p (corresponding to $\mathcal{N}(0, I_p)$) for any $p \in \mathbb{N}$.

Whenever $\tau \in \mathcal{M}_+(\Omega)$ has finite and non-zero total variation, we denote by $\tilde{\tau} \in \mathcal{P}_2(\Omega)$ its normalized counterpart (which is a probability measure), that is $\tilde{\tau} = \tau/\tau(\Omega) = \tau/|\tau|$.

For any $p \in \mathbb{N}$, we call ω_p the Lebesgue (spherical) measure over the unit sphere \mathbb{S}^{p-1} of \mathbb{R}^p , that is the measure such that $\tilde{\omega}_p$ is the *uniform* measure on \mathbb{S}^{p-1} . We then denote by $|\mathbb{S}^{p-1}|$ the surface area of \mathbb{S}^{p-1} , that is $|\mathbb{S}^{p-1}| := |\omega_p| = \omega_p(\mathbb{S}^{p-1}) = 2\pi^{p/2}/\Gamma(p/2)$.

Smooth functions: we denote by $\mathcal{C}(\Omega)$ (*resp.* $\mathcal{C}_c^1(\Omega)$) the set of continuous (*resp.* continuously differentiable and compactly supported) functions from a set Ω to \mathbb{R} .

A.2 General results on invariance for measures and functions

In this section, we list a number of lemmas related to symmetries of measures and functions which will prove helpful in the proofs presented in the Appendix.

Lemma A.1 (Invariance under invertible maps). *Let μ be a measure invariant under some measurable and invertible map T . Then, assuming T^{-1} is also measurable, one has that μ is also invariant under T^{-1} .*

Remark. A similar result holds for a function f invariant under an invertible map.

Proof. Because μ is invariant under T , we have for any measurable set A , $\mu(A) = \mu(T^{-1}(A))$. Since T^{-1} is assumed to be measurable, for any measurable set A , $T(A)$ is also measurable ($T(A) = (T^{-1})^{-1}(A)$) and thus $\mu(T(A)) = \mu(T^{-1}(T(A))) = \mu(A)$ which shows μ is invariant under T^{-1} . \square

Lemma A.2 (Invariance of the density). *Let ν be a measure with density p w.r.t. some measure μ , and assume both ν and μ are σ -finite and invariant under some measurable and invertible map T , whose inverse T^{-1} is also measurable. Then p is also invariant under T μ -almost everywhere, i.e., $p(T(x)) = p(x)$ for μ -almost every x .*

Proof. For any measurable φ (w.r.t. μ , and thus w.r.t. ν as well), $\varphi \circ T^{-1}$ is also measurable, and we have, on the one hand

$$\int \varphi \circ T^{-1} d\nu = \int (\varphi \circ T^{-1}) p d\mu = \int \varphi (p \circ T) d\mu,$$

and on the other hand

$$\int \varphi \circ T^{-1} d\nu = \int \varphi d\nu = \int \varphi p d\mu,$$

which shows that $\int \varphi (p \circ T) d\mu = \int \varphi p d\mu$, and thus that $p \circ T = p$ μ -almost everywhere. \square

Lemma A.3 (Projected variance with spherical symmetry). *Let ζ be a spherically symmetric measure on \mathbb{R}^p (i.e., such that for any orthogonal linear map $T \in \mathcal{O}(p)$, $T_{\#}\zeta = \zeta$), with finite second moment. Then we have the following matrix identity:*

$$\int_z z z^\top d\zeta(z) = v_\zeta I_p, \quad v_\zeta := \int_z (z_1)^2 d\zeta(z) = \frac{1}{p} \int_z \|z\|^2 d\zeta(z).$$

Proof. The (i, j) -th entry of the matrix on the left-hand-side is $\int_z z_i z_j d\zeta(z)$, and it is readily seen that the terms outside the diagonal are 0. Indeed, let $(i, j) \in [1, p]^2$ with $i \neq j$, and consider the orthogonal map $T_j : z \in \mathbb{R}^p \mapsto (z_1, \dots, z_{j-1}, -z_j, z_{j+1}, \dots, z_p)^\top$. The spherical symmetry of ρ implies that it is invariant under T_j , which yields $\int_z z_i z_j d\rho(z) = -\int_z z_i z_j d\rho(z)$, thereby showing that the latter is 0. To see that the diagonal terms are all equal, it suffices to consider the orthogonal map S_i which swaps the 1st and i -th coordinates of a vector z . The invariance of ρ under S_i yields $\int_z (z_1)^2 d\rho(z) = \int_z (z_i)^2 d\rho(z)$, which concludes the proof. \square

A.3 A disintegration result on the unit sphere \mathbb{S}^{d-1}

Consider a $u \in \mathbb{S}^{d-1}$. u is determined by: (i) its angle $\theta := \arccos(\|u^H\|) \in [0, \pi/2]$ with H (i.e., its angle with its projection u^H onto H), (ii) the direction $z^H = u^H/\|u^H\| \in \mathbb{S}^{d_H-1}$ of its projection u^H onto H , and finally (iii) the direction $z^\perp = u^\perp/\|u^\perp\| \in \mathbb{S}^{d_\perp-1}$ of its projection u^\perp onto H^\perp . Since $\|u^H\|^2 + \|u^\perp\|^2 = 1$, the angle θ gives both the norms of the projections onto H and H^\perp : $\|u^H\| = \cos(\theta)$ and $\|u^\perp\| = \sin(\theta)$.

When z ranges over the unit sphere \mathbb{S}^{d-1} , the angle θ and the directions z^H, z^\perp range over $[0, \pi/2]$, \mathbb{S}^{d_H-1} , and $\mathbb{S}^{d_\perp-1}$ respectively. We wish to understand what measures we obtain on these three sets when z is distributed on the sphere according to the Lebesgue measure ω_d . We show below that after the change of coordinates described above (from $u \in \mathbb{S}^{d-1}$ to $(\theta, z^H, z^\perp) \in [0, \pi/2] \times \mathbb{S}^{d_H-1} \times \mathbb{S}^{d_\perp-1}$) that the corresponding measures over \mathbb{S}^{d_H-1} and $\mathbb{S}^{d_\perp-1}$ are uniform and the measure over θ is given by a push-forward of a Beta distribution as defined below:

Definition A.1 (Distribution γ of the angle θ). We define the measure γ on $[0, \pi/2]$ with the following density w.r.t. the Lebesgue measure on $[0, \pi/2]$:

$$d\gamma(\theta) := \cos(\theta)^{d_H-1} \sin(\theta)^{d_\perp-1} d\theta.$$

Remark. γ is in fact simply given by $(\arccos \circ \sqrt{\cdot})_\# \text{Beta}(d_H/2, d_\perp/2)$. Note that the total variation of gamma is $|\gamma| = \gamma([0, \pi/2]) = \frac{1}{2} B(\frac{d_H}{2}, \frac{d_\perp}{2})$, and the corresponding normalized (probability) measure is $d\tilde{\gamma}(\theta) = d\gamma(\theta)/|\gamma| = \frac{2}{B(\frac{d_H}{2}, \frac{d_\perp}{2})} \cos(\theta)^{d_H-1} \sin(\theta)^{d_\perp-1} d\theta$.

We now state the disintegration theorem and give its proof:

Theorem A.4 (Disintegration of the Lebesgue measure on the sphere). *Let ω_d denote the Lebesgue measure on the sphere measure on the sphere of \mathbb{R}^d , and let γ be the measure of Definition A.1. Then, one has*

$$\omega_d = \Phi_\#(\omega_{d_H} \otimes \omega_{d_\perp} \otimes \gamma)$$

where

$$\begin{aligned} \Phi : [0, \pi/2] \times \mathbb{S}^{d_H-1} \times \mathbb{S}^{d_\perp-1} &\rightarrow \mathbb{S}^{d-1} \\ (\theta, z_H, z_\perp) &\mapsto \cos(\theta)z_H + \sin(\theta)z_\perp. \end{aligned}$$

Proof. Denoting $\tilde{\omega}_d$ the uniform measure on the sphere, $|\mathbb{S}^{d-1}| := \frac{2\pi^{d/2}}{\Gamma(d/2)}$ the surface area of the sphere in dimension d , and ρ_p the standard Gaussian distribution in \mathbb{R}^p for any p . Using the well-known fact that $\tilde{\omega}_d = \Pi_\# \rho_d$ with $\Pi : x \in \mathbb{R}^d \setminus \{0\} \mapsto x/\|x\| \in \mathbb{S}^{d-1}$, we have, for any measurable test function $\varphi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$,

$$\begin{aligned} \int \varphi d\omega_d &= |\mathbb{S}^{d-1}| \int \varphi d\tilde{\omega}_d \\ &= |\mathbb{S}^{d-1}| \int_x \varphi \left(\frac{x}{\|x\|} \right) d\rho_d(x) \\ &= |\mathbb{S}^{d-1}| \int_{x_H, x_\perp} \varphi \left(\frac{x_H + x_\perp}{\|x_H + x_\perp\|} \right) d\rho_{d_H}(x_H) d\rho_{d_\perp}(x_\perp) \\ &= C_d \int \varphi \left(\frac{r_H z_H + r_\perp z_\perp}{\|r_H z_H + r_\perp z_\perp\|} \right) r_H^{d_H-1} e^{-r_H^2/2} r_\perp^{d_\perp-1} e^{-r_\perp^2/2} dr_H dr_\perp d\omega_{d_H}(z_H) d\omega_{d_\perp}(z_\perp) \\ &= C_d \int_{z_H, z_\perp} \int_{r_H, r_\perp} \varphi \left(\frac{r_H z_H + r_\perp z_\perp}{\sqrt{r_H^2 + r_\perp^2}} \right) r_H^{d_H-1} r_\perp^{d_\perp-1} e^{-(r_H^2 + r_\perp^2)/2} dr_H dr_\perp d\omega_{d_H}(z_H) d\omega_{d_\perp}(z_\perp), \end{aligned}$$

with

$$C_d := \frac{|\mathbb{S}^{d-1}|}{(2\pi)^{d_H/2}(2\pi)^{d_\perp/2}} = \frac{|\mathbb{S}^{d-1}|}{(2\pi)^{d/2}} = \frac{2\pi^{d/2}}{2^{d/2}\pi^{d/2}\Gamma(d/2)} = \frac{1}{2^{(d-2)/2}\Gamma(d/2)}.$$

Doing the change of variables $(r_H, r_\perp) \in \mathbb{R}_+^2 \rightarrow (R, \theta) \in \mathbb{R}_+ \times [0, \pi/2]$, we get:

$$\int \varphi d\omega_d = C'_d \int_{z_H, z_\perp} \int_\theta \varphi(\cos(\theta)z_H + \sin(\theta)z_\perp) \cos(\theta)^{d_H-1} \sin(\theta)^{d_\perp-1} d\theta d\omega_{d_H}(z_H) d\omega_{d_\perp}(z_\perp)$$

where

$$\begin{aligned} C'_d &:= C_d \int_0^{+\infty} R^{d-2} e^{-R^2/2} R dR \\ &= C_d \int_0^{+\infty} R^{d-1} e^{-R^2/2} dR \\ &= C_d \times 2^{(d-2)/2} \Gamma(d/2) \\ &= 1. \end{aligned}$$

which concludes the proof. \square

Remark. A similar disintegration result holds for the uniform measure $\tilde{\omega}_d$ on the sphere. The corresponding measures which are then pushed-forward by the same Φ are the normalized counterparts of the measures in the theorem above: $\tilde{\omega}_d = \Phi_\#(\tilde{\omega}_{d_H} \otimes \tilde{\omega}_{d_\perp} \otimes \tilde{\gamma})$. This readily comes from noting that a simple calculation yields $|\omega_d| = |\omega_{d_H}| |\omega_{d_\perp}| |\gamma|$.

B Gradient flows on the space of probability measures

Say somewhere here what the distributional sense of the continuity equation is.

B.1 First variation of a functional over measures

Given a functional $F : \mathcal{P}_2(\mathbb{R}^p) \rightarrow \mathbb{R}$, its *first variation* or *Fréchet derivative* at $\mu \in \mathcal{P}_2(\mathbb{R}^p)$ is defined as a measurable function, denoted $\frac{\delta F}{\delta \mu}(\mu) : \mathbb{R}^p \rightarrow \mathbb{R}$, such that, for any $\nu \in \mathcal{P}_2(\mathbb{R}^p)$ for which $\mu + t\nu \in \mathcal{P}_2(\mathbb{R}^p)$ in a neighborhood (in t) of $t = 0$,

$$\left. \frac{d}{dt} F(\mu + t\nu) \right|_{t=0} = \int_z \frac{\delta F}{\delta \mu}(\mu)[z] d\nu(z).$$

See (Santambrogio, 2015, Definition 7.12), or (Santambrogio, 2017, p.29) for more details on the first variation.

In the case of the functional defined in Equation (1.2) corresponding to the population loss objective, using the differentiability of the loss ℓ *w.r.t.* its second argument, one readily has that

$$F'_\mu(c) := \frac{\delta F}{\delta \mu}(\mu)[c] = \int_x \partial_2 \ell(f^*(x), f(\mu; x)) \phi(c; x) d\rho(x)$$

since

$$\frac{d}{dt} \ell(f^*(x), f(\mu; x) + tf(\nu; x)) = \partial_2 \ell(f^*(x), f(\mu; x) + tf(\nu; x)) \int_c \phi(c; x) d\nu(c).$$

B.2 Wasserstein gradient flows in the space $\mathcal{P}_2(\mathbb{R}^{d+1})$

A Wasserstein gradient flow for the objective F defined in Equation (1.2) is a path $(\mu_t)_{t \geq 0}$ in the space of probability measures $\mathcal{P}_2(\mathbb{R}^{d+1})$ which satisfies the continuity equation with a vector field v_t which is equal to the opposite of the gradient of the first variation of the functional F . This means that we have, in the sense of distributions,

$$\partial_t \mu_t = -\operatorname{div} \left(-\nabla \left(\frac{\delta F}{\delta \mu} \right) \mu_t \right).$$

That a pair $((\mu_t)_{t \geq 0}, v_t)$ consisting of a path in $\mathcal{P}_2(\mathbb{R}^p)$ and a (time-dependent) vector field in \mathbb{R}^p satisfies the continuity equation $\partial_t \mu_t = -\operatorname{div}(v_t \mu_t)$ in the sense of the distributions simply means that for any test function $\varphi \in \mathcal{C}_c^1(\mathbb{R}^p)$,

$$\partial_t \int \varphi \, d\mu_t = \int v_t^\top \nabla \varphi \, d\mu_t,$$

where ∂_t stands for the time derivative $\frac{d}{dt}$. Similarly, when we say that the advection-reaction equation $\partial_t \mu_t = -\operatorname{div}(v_t \mu_t) + g_t \mu_t$ is satisfied for some function $g_t : \mathbb{R}^p \rightarrow \mathbb{R}$, we mean that it is in the sense of distributions: for any test function $\varphi \in \mathcal{C}_c^1(\mathbb{R}^p)$,

$$\partial_t \int \varphi \, d\mu_t = \int (v_t^\top \nabla \varphi + g_t) \, d\mu_t.$$

An alternative description of the Wasserstein gradient flow of the objective F is to consider a flow $X_\bullet(\cdot)$ in $\mathbb{R}_+ \times \mathbb{R}^{d+1}$ such that, for any $c \in \mathbb{R}^{d+1}$,

$$\begin{aligned} X_0(c) &= c \\ \frac{d}{dt} X_t(c) &= -\nabla \left(\frac{\delta F}{\delta \mu} \right) (X_t(c)) \end{aligned}$$

and to define $\mu_t = (X_t)_\# \mu_0$.

For more details on Wasserstein gradient flows in the space of probability measures see (Santambrogio, 2015, Section 5.3), and (Santambrogio, 2017, Section 4), and for more details on the equivalence between the continuity equation and the flow-based representation of the solution see (Santambrogio, 2015, Theorem 4.4).

C Proofs of the symmetry results of Section 2

There are two main ideas behind the proof. Call $\tilde{T} : (a, b) \in \mathbb{R} \times \mathbb{R}^d \mapsto (\pm a, T(b))$ (depending on whether f^* is invariant or anti-invariant under T) and consider the following two facts:

Structure of $\phi((a, b); x)$. Since T is orthogonal, so is \tilde{T} , and the structure of $\phi((a, b); x) = a\sigma(b^\top x)$ is such that $\phi(\tilde{T}(a, b); x) = \pm\phi((a, b); T^{-1}(x))$ because T is orthogonal (its adjoint is thus its inverse).

Conjugate gradients. Computing the gradient of a function whose input has been transformed by \tilde{T}^{-1} is the same as the conjugate action of \tilde{T} on the gradient: $\nabla(\varphi \circ \tilde{T}^{-1}) = \tilde{T} \circ (\nabla \varphi) \circ \tilde{T}^{-1}$ (this is due to the fact that the adjoint of \tilde{T}^{-1} is \tilde{T} because \tilde{T} is orthogonal). Note that we similarly get $\nabla(\varphi \circ \tilde{T}) = \tilde{T}^{-1} \circ (\nabla \varphi) \circ \tilde{T}$.

C.1 Preliminaries

We present here arguments that are present in both the proofs of Proposition 2.1 and 2.2. Let T be a linear orthogonal map such that $f^*(T(x)) = \pm f^*(x)$, where the \pm is because we deal with both cases at the same time since the logic is the same. Let $t \geq 0$, and define $\nu_t := \tilde{T}_\#^{-1} \mu_t$. We aim to show that $(\nu_t)_{t \geq 0}$ is also a Wasserstein gradient flow for the same objective as $(\mu_t)_{t \geq 0}$.

Prediction function. Let $x \in \mathbb{R}^d$. We have, using the fact that T is orthogonal (and thus that $\langle T(x), y \rangle = \langle x, T^{-1}(y) \rangle$),

$$\begin{aligned} f(\nu_t; x) &= \int_{a,b} a \sigma(b^\top x) d\nu_t(a, b) \\ &= \int_{a,b} \pm a \sigma(T^{-1}(b)^\top x) d\mu_t(a, b) \\ &= \pm \int_{a,b} a \sigma(b^\top T(x)) d\mu_t(a, b) \\ &= \pm f(\mu_t; T(x)). \end{aligned}$$

Time derivative. Let $\varphi \in C_c^1(\mathbb{R}^d)$. Because μ_t satisfies the continuity Equation (1.3) in the sense of distributions, and using the remark above on conjugate gradients as well as the orthogonality of \tilde{T} , we have:

$$\begin{aligned} \partial_t \int \varphi d\nu_t &= \partial_t \int \varphi \circ \tilde{T}^{-1} d\mu_t \\ &= \int \langle \nabla(\varphi \circ \tilde{T}^{-1}), v_t \rangle d\mu_t \\ &= \int \langle \tilde{T} \circ \nabla \varphi \circ \tilde{T}^{-1}, v_t \rangle d\mu_t \\ &= \int \langle \nabla \varphi \circ \tilde{T}^{-1}, \tilde{T}^{-1} \circ v_t \rangle d\mu_t \\ &= \int \langle \nabla \varphi, \tilde{T}^{-1} \circ v_t \circ \tilde{T} \rangle d\nu_t. \end{aligned}$$

Conjugate velocity field. The equality above actually shows that ν_t satisfies the continuity equation with the conjugate velocity field $\tilde{T}^{-1} \circ v_t \circ \tilde{T}$ instead of v_t . We show below that the former is closely related to the latter (and in fact equal to $-\nabla F'_{\nu_t}$ with sufficient assumptions on $\partial_2 \ell$, which is the step proven in Appendices C.2 and C.3). Indeed, because v_t is a gradient $v_t = -\nabla F'_{\mu_t}$, we have using again the remark above on conjugate gradients:

$$\tilde{T}^{-1} \circ v_t \circ \tilde{T} = -\nabla \left(F'_{\mu_t} \circ \tilde{T} \right).$$

Computing the function on the right-hand-side, for any $(a, b) \in \mathbb{R} \times \mathbb{R}^d$, we get, using the remark above on the structure of ϕ ,

$$\begin{aligned} F'_{\mu_t}(\tilde{T}(a, b)) &= \int_y \partial_2 \ell \left(f^*(y), f(\mu_t; y) \right) \phi \left(\tilde{T}(a, b); y \right) d\rho(y) \\ &= \pm \int_y \partial_2 \ell \left(f^*(y), f(\mu_t; y) \right) \phi \left((a, b); T^{-1}(y) \right) d\rho(y). \end{aligned}$$

ρ is invariant under T since it is spherically symmetric by assumption (and thus invariant under any orthogonal map) and we can therefore replace y by $T(y)$ in the integral above, which yields

$$\begin{aligned} F'_{\mu_t}(\tilde{T}(a, b)) &= \pm \int_y \partial_2 \ell(f^*(T(y)), f(\mu_t; T(y))) \phi((a, b); y) d\rho(y) \\ &= \pm \int_y \partial_2 \ell(\pm f^*(y), \pm f(\nu_t; y)) \phi((a, b); y) d\rho(y), \end{aligned}$$

and thus we get

$$\nabla (F'_{\mu_t} \circ \tilde{T})(a, b) = \pm \int_y \partial_2 \ell(\pm f^*(y), \pm f(\nu_t; y)) \nabla_{(a, b)} \phi((a, b); y) d\rho(y).$$

One can already notice that if f^* is invariant under T (as opposed to anti-invariant), that is if we keep the “+” in \pm , we get $\tilde{T}^{-1} \circ \nu_t \circ \tilde{T} = -\nabla F'_{\nu_t}$.

C.2 Proof of Proposition 2.1

Proof. We first prove $\nu_0 = \mu_0$ and then prove that both $(\mu_t)_{t \geq 0}$ and $(\nu_t)_{t \geq 0}$ are Wasserstein gradient flows of the objective F defined in Equation (1.2), starting from the initial condition μ_0 at $t = 0$. The unicity of such a gradient flow then guarantees that $\mu_t = \nu_t$ and thus $f(\mu_t; T(x)) = f(\mu_t; x)$ by the preliminaries above on the prediction function (see Appendix C.1).

Initialization: $\nu_0 = \mu_0$. By definition, $\tilde{T}(a, b) = (a, T(b))$. Since $\mu_0 = \mu_0^1 \otimes \mu_0^2$ by assumption, and μ_0^2 is invariant under T since it has spherical symmetry, it is clear that μ_0 is invariant under \tilde{T} , and thus under \tilde{T}^{-1} by Lemma A.1, which gives $\nu_0 = \mu_0$ because $\nu_t = \tilde{T}_{\#}^{-1} \mu_t$ by definition. \square

Time derivative. From the preliminary results above (see Appendix C.1) we have

$$\partial_t \nu_t = -\operatorname{div}(-\nabla F'_{\nu_t} \nu_t),$$

which shows that $(\nu_t)_{t \geq 0}$ is also a Wasserstein gradient flow of the objective F . By unicity of the latter (starting from the initial condition μ_0), it must hold that $\mu_t = \nu_t$ for any $t \geq 0$ which concludes the proof.

C.3 Proof of Proposition 2.2

The proof follows the exact same pattern as that of Proposition 2.1 (see Appendix C.2). We now have by definition, $\tilde{T}(a, b) = (-a, T(b))$ and the added symmetry assumption on μ_0^1 ensures that $\nu_0 = \mu_0$ still holds in this case. As for the time derivative, the preliminaries above (see Appendix C.1) ensure that

$$\begin{aligned} \nabla (F'_{\mu_t} \circ \tilde{T})(a, b) &= - \int_y \partial_2 \ell(-f^*(y), -f(\nu_t; y)) \nabla_{(a, b)} \phi((a, b); y) d\rho(y) \\ &= \int_y \partial_2 \ell(f^*(y), f(\nu_t; y)) \nabla_{(a, b)} \phi((a, b); y) d\rho(y), \end{aligned}$$

where we have used the extra assumption that $\partial_2 \ell(-y, -\hat{y}) = -\partial_2 \ell(y, \hat{y})$. This yields

$$\partial_t \nu_t = -\operatorname{div}(-\nabla F'_{\nu_t} \nu_t)$$

and the conclusion follows from the same logic as for Proposition 2.1.

D Proof of the exponential convergence for linear networks: Theorem 3.2

Proof. The proof is divided in three steps: (i) we derive the dynamics in time of the vector $w(t) = \frac{1}{2} \int ab d\mu_t(a, b)$, (ii) we show that the positive definite matrix $H(t)$ appearing in these dynamics has its smallest eigenvalue lower-bounded by some positive constant after some $t_0 > 0$, and (iii) we show that this implies the exponential convergence to the global minimum.

Generalities on the objective Q . Expanding the square in the definition of Q (3.2), we have

$$\begin{aligned} Q(w) &= \frac{1}{2} \left[\mathbb{E}_{x \sim \mathbb{P}} [f^*(x)^2] - 2\beta^\top w + w^\top C w \right], \\ C &:= \mathbb{E}_{x \sim \mathbb{P}} [xx^\top] \in \mathbb{R}^{d \times d}, \\ \beta &:= \mathbb{E}_{x \sim \mathbb{P}} [f^*(x)x] \in \mathbb{R}^d. \end{aligned}$$

Q is a convex function and is lower-bounded by 0 and thus admits at least one global minimum. This minimizer is unique as soon as Q is strongly convex, *i.e.*, C is definite positive, which holds in this case as we have assumed the smallest eigenvalue λ_{\min} of C to be > 0 . Note that $\nabla Q(w) = Cw - \beta = \int_x ((x^\top w) - f^*(x)) x d\mathbb{P}(x) \in \mathbb{R}^d$.

First step: dynamics of $w(t)$. Let $k \in \{1, \dots, d\}$, the k -th coordinate $w_k(t)$ of $w(t)$ is given by $w_k(t) = \int ab_k d\mu_t(a, b)$, and its time derivative is given by

$$w'_k(t) = \frac{1}{2} \int \left(\nabla_{(a,b)}(ab_k) \right)^\top v_t(a, b) d\mu_t(a, b)$$

where v_t is given by Equation (1.3) except we replace σ by $\frac{1}{2}\text{id}_{\mathbb{R}^d}$ and ρ by \mathbb{P} in F_{μ_t} , that is

$$v_t(a, b) = \frac{1}{2} \int_y R_t(y) \begin{pmatrix} b^\top y \\ ay \end{pmatrix} d\mathbb{P}(y) \in \mathbb{R}^{1+d}.$$

On the other hand, $\nabla_{(a,b)}(ab_k) = \begin{pmatrix} b_k \\ ae_k \end{pmatrix} \in \mathbb{R}^{1+d}$ where e_k is the k -th element of the canonical orthonormal basis of \mathbb{R}^d . Note that here, $R_t(y) = f^*(y) - \langle w(t), y \rangle$. We thus get

$$\begin{aligned} w'_k(t) &= \frac{1}{4} \left\langle \int_{a,b} b_k b d\mu_t(a, b), \int_y \left(f^*(y) - (w(t)^\top y) \right) y d\mathbb{P}(y) \right\rangle + \\ &\quad \frac{1}{4} \left\langle \int_{a,b} a^2 e_k d\mu_t(a, b), \int_y \left(f^*(y) - (w(t)^\top y) \right) y d\mathbb{P}(y) \right\rangle. \end{aligned}$$

Note that the term on the right in the inner products is in fact equal to $-\nabla Q(w(t))$, which yields the following dynamics for the vector $w(t)$:

$$\begin{aligned} w'(t) &= -H(t) \nabla Q(w(t)), \\ H(t) &:= \frac{1}{4} \left(\int bb^\top d\mu_t(a, b) + \int a^2 d\mu_t(a, b) I_d \right) \in \mathbb{R}^{d \times d}. \end{aligned}$$

Second step: lower bound on the smallest eigenvalue of $H(t)$. At initialization, by symmetry one has $w(0) = 0$, and using Lemma A.3, one has that $H(0) = \frac{1}{4}(\frac{1}{d} + 1)I_d$, so that

$$\begin{aligned}\frac{d}{dt}Q(w(t))\Big|_{t=0} &= \langle w'(0), \nabla Q(w(0)) \rangle \\ &= -\frac{d+1}{4d}\|\nabla Q(0)\|^2 \\ &= -\frac{d+1}{4d}\|\beta\|^2\end{aligned}$$

If $\beta = 0$, then $\nabla Q(0) = 0$ and since $w(0) = 0$, $w(t)$ starts at the global optimum and thus stays constant equal to 0. Otherwise, if $\|\beta\| > 0$, one has $\frac{d}{dt}Q(w(t))\Big|_{t=0} < 0$, which ensures that there is a $t_0 > 0$ such that $Q(w(t)) < Q(w(0)) = Q(0)$ for any $t \in (0, t_0]$. Call $\varepsilon := [Q(0) - Q(w(t_0))]/2 > 0$. The continuity of Q at 0 guarantees that there is a $\delta > 0$ such that for any $w \in \mathbb{R}^d$, if $\|w\| < \delta$, then $|Q(w) - Q(0)| \leq \varepsilon$.

Now assume that there exists $t_1 \geq t_0$ such that $\int a^2 d\mu_{t_1}(a, b) \leq \delta$. Then, one has

$$\begin{aligned}\|w(t_1)\| &= \left\| \frac{1}{2} \int ab d\mu_{t_1}(a, b) \right\| \\ &\leq \frac{1}{2} \int |a| |b| d\mu_{t_1}(a, b) \\ &\leq \frac{1}{2} \int a^2 d\mu_{t_1}(a, b) \\ &\leq \frac{\delta}{2} < \delta,\end{aligned}$$

where we have used in the penultimate inequality that μ_{t_1} is supported on the set $\{|a| = |b|\}$ because of the assumptions on the initialization μ_0 (see Section 1.1). This ensures that $|Q(w(t_1)) - Q(0)| \leq \varepsilon$. Since $t \mapsto Q(w(t))$ is decreasing ($Q(w(t)) = F(\mu_t)$ and it is classical that the objective is decreasing along the gradient flow path, see third step below) and $t_1 \geq t_0$, this means that

$$0 < Q(0) - Q(w(t_0)) \leq Q(0) - Q(w(t_1)) \leq \varepsilon = [Q(0) - Q(w(t_0))]/2$$

which is a contradiction. Therefore, for any $t \geq t_0$, $\int a^2 d\mu_t(a, b) \geq \delta$. Calling $\eta := \delta/4 > 0$, we thus have that for any $t \geq t_0$, the smallest eigenvalue of $H(t)$ is larger than η .

Third step: exponential convergence. We have:

$$\begin{aligned}\frac{d}{dt}Q(w(t)) &= \langle w'(t), \nabla Q(w(t)) \rangle \\ &= -\nabla Q(w(t))^\top H(t) \nabla Q(w(t)) \leq 0,\end{aligned}$$

which shows that because $H(t)$ is positive definite, the objective Q is decreasing along the path $(w(t))_{t \geq 0}$. Since after $t_0 > 0$, the smallest eigenvalue of $H(t)$ is lower bounded by a constant $\eta > 0$, we have that, for any $t \geq t_0$:

$$\frac{d}{dt}Q(w(t)) \leq -\eta \|\nabla Q(w(t))\|^2. \quad (\text{D.1})$$

Because Q is λ_{\min} -strongly convex (as the smallest eigenvalue of C is $\lambda_{\min} > 0$), one has the classical inequality

$$\frac{1}{2} \|\nabla Q(w)\|^2 \geq \lambda_{\min} (Q(w) - Q(w^*)).$$

Plugging this into Equation (D.1) gives

$$\frac{d}{dt} \left(Q(w(t)) - Q(w^*) \right) \leq -2\eta\lambda_{\min} \left(Q(w(t)) - Q(w^*) \right),$$

which by Gronwall's lemma in turn yields for any $t \geq t_0$

$$0 \leq Q(w(t)) - Q(w^*) \leq e^{-2\eta\lambda_{\min}(t-t_0)} \left(Q(w(t_0)) - Q(w^*) \right),$$

thereby proving exponential convergence.

Exponential convergence in distance. The convergence in distance to the optimum follows from the classical inequality for convex functions

$$\langle -\nabla Q(w(t)), w^* - w(t) \rangle \leq Q(w(t)) - Q(w^*).$$

Given that $\nabla Q(w) = Cw - \beta$ and that $Cw^* = \beta$ because w^* is an optimum, which implies $\nabla Q(w^*) = 0$, the inequality above yields

$$\langle C(w^* - w(t)), w^* - w(t) \rangle \leq Q(w(t)) - Q(w^*),$$

and given that the left-hand-side is larger than $\lambda_{\min} \|w(t) - w^*\|^2$, we get

$$\|w(t) - w^*\|^2 \leq \frac{1}{\lambda_{\min}} \left(Q(w(t)) - Q(w^*) \right),$$

and the exponential decrease of the right-hand-side allows to conclude. \square

E Proofs of Section 4: f^* depends only on the projection on a sub-space H

E.1 The general case

E.1.1 Closed dynamics on the sphere \mathbb{S}^{d-1}

We wish to show here that the pair of measures (ν_t^+, ν_t^-) defined through Equation (4.1) satisfy Equation (4.2) and that the corresponding dynamic is closed in the sense that it can be expressed solely using (ν_t^+, ν_t^-) (without requiring to express quantities in function of μ_t). Below, we use $\kappa(z) = \max(0, z)$. We do this by differentiating it from the activation function σ (which is also equal to ReLU) so as to avoid confusion because the κ which appears below has nothing to do with the activation function of the network and simply comes from the calculations which occur.

Equations of the dynamics on the sphere. Let $\varphi \in \mathcal{C}_c^1(\mathbb{S}^{d-1})$. One has

$$\begin{aligned} \partial_t \int \varphi d\nu_t^\pm &= \partial_t \int_{\pm a \geq 0, b} |a| |b| \varphi \left(\frac{b}{\|b\|} \right) d\mu_t(a, b) \\ &= \partial \int_{a, b} \kappa(\pm a) |b| \varphi \left(\frac{b}{\|b\|} \right) d\mu_t(a, b) \\ &= \int_{a, b} \nabla_{(a, b)} \left(\kappa(\pm a) |b| \varphi \left(\frac{b}{\|b\|} \right) \right)^\top v_t(a, b) d\mu_t(a, b) \end{aligned}$$

Let us compute the components of the gradient above. We have

$$\nabla_a \left(\kappa(\pm a) \|b\| \varphi \left(\frac{b}{\|b\|} \right) \right) = \pm \kappa'(\pm a) \|b\| \varphi \left(\frac{b}{\|b\|} \right) = \mathbf{1}_{\{\pm a \geq 0\}} \|b\| \varphi \left(\frac{b}{\|b\|} \right).$$

The Jacobian of the map $b \in \mathbb{R}^d \mapsto b/\|b\|$ is equal to $(I_d - bb^\top/\|b\|^2)/\|b\|$ which is a symmetric (or self-adjoint) matrix, so that the gradient *w.r.t.* b is

$$\nabla_b \left(\kappa(\pm a) \|b\| \varphi \left(\frac{b}{\|b\|} \right) \right) = \mathbf{1}_{\{\pm a \geq 0\}} |a| \left[\varphi \left(\frac{b}{\|b\|} \right) \frac{b}{\|b\|} + \left(I_d - \frac{b}{\|b\|} \left(\frac{b}{\|b\|} \right)^\top \right) \nabla \varphi \left(\frac{b}{\|b\|} \right) \right].$$

On the other hand, the first component of $v_t(a, b)$ (corresponding to the gradient *w.r.t.* a) is

$$v_t^1(a, b) = \int_y R_t(y) \kappa(b^\top y) d\rho(y) = \|b\| \int_y R_t(y) \kappa \left(\left(\frac{b}{\|b\|} \right)^\top y \right) d\rho(y),$$

and the last d components (corresponding to the gradient *w.r.t.* b) are

$$v_t^2(a, b) = \int_y R_t(y) a \kappa'(b^\top y) y d\rho(y) = a \int_y R_t(y) \kappa' \left(\left(\frac{b}{\|b\|} \right)^\top y \right) y d\rho(y).$$

When computing the inner product $\nabla_{(a,b)} \left(\kappa(\pm a) \|b\| \varphi \left(\frac{b}{\|b\|} \right) \right)^\top v_t(a, b)$, we can re-arrange the terms to keep one term where φ appears and the other where $\nabla \varphi$ appears. Using the facts that the Jacobian computed above is symmetric, that $\kappa(z) = \kappa'(z)z$ for any $z \in \mathbb{R}$, and that $\mathbf{1}_{\{\pm a \geq 0\}} a = \pm \mathbf{1}_{\{\pm a \geq 0\}} |a|$, we get,

$$\begin{aligned} \nabla_{(a,b)} \left(\kappa(\pm a) \|b\| \varphi \left(\frac{b}{\|b\|} \right) \right)^\top v_t(a, b) &= \pm \mathbf{1}_{\{\pm a \geq 0\}} \|b\| \|b\| \varphi \left(\frac{b}{\|b\|} \right) g_t \left(\frac{b}{\|b\|} \right) + \\ &\pm \mathbf{1}_{\{\pm a \geq 0\}} |a| |a| \varphi \left(\frac{b}{\|b\|} \right) g_t \left(\frac{b}{\|b\|} \right) + \\ &\pm \mathbf{1}_{\{\pm a \geq 0\}} |a| |a| \nabla \varphi \left(\frac{b}{\|b\|} \right)^\top \tilde{v}_t \left(\frac{b}{\|b\|} \right), \end{aligned}$$

where, for $u \in \mathbb{S}^{d-1}$

$$\begin{aligned} g_t(u) &:= \int_y R_t(y) \sigma(u^\top y) d\rho(y), \\ \tilde{v}_t(u) &:= \int_y R_t(y) \sigma'(u^\top y) [y - (u^\top y)u] d\rho(y). \end{aligned}$$

Finally, because μ_t stays on the cone $\{(a, b) \in \mathbb{R}^{d+1}; |a| = \|b\|\}$ for any t (see [Chizat and Bach, 2020](#), Lemma 26, [Wojtowytsch, 2020](#), Section 2.5), when integrating against μ_t , we can replace $\|b\|$ by $|a|$ and vice-versa. We thus get that the time derivative we initially computed is the sum of two terms:

$$\begin{aligned} \partial_t \int \varphi d\nu_t^\pm &= 2 \int_{\pm a \geq 0, b} |a| \|b\| \varphi \left(\frac{b}{\|b\|} \right) g_t \left(\frac{b}{\|b\|} \right) d\mu_t(a, b) + \\ &\int_{\pm a \geq 0, b} |a| \|b\| \nabla \varphi \left(\frac{b}{\|b\|} \right)^\top \tilde{v}_t \left(\frac{b}{\|b\|} \right) d\mu_t(a, b) \\ &= 2 \int_{u \in \mathbb{S}^{d-1}} \varphi(u) g_t(u) d\nu_t^\pm(u) + \\ &\int_{u \in \mathbb{S}^{d-1}} \nabla \varphi(u)^\top \tilde{v}_t(u) d\nu_t^\pm(u), \end{aligned}$$

which shows that ν_t^\pm satisfies Equation (4.2) in the sense of distributions.

Closed dynamics. We want to show that g_t and \tilde{v}_t can be expressed using only ν_t^+ and ν_t^- . Both these quantities depend on t only through the residual R_t , which itself only depends on t through $f(\mu_t; \cdot)$. We thus show that the latter can be expressed using only ν_t^+ and ν_t^- , which easily follows from writing, for any $y \in \mathbb{R}^d$,

$$\begin{aligned} f(\mu_t; y) &= \int a \sigma(b^\top y) \, d\mu_t(a, b) \\ &= \int a \|b\| \sigma\left(\left\langle \frac{b}{\|b\|}, y \right\rangle\right) \, d\mu_t(a, b) \\ &= \int_{a \geq 0, b} |a| \|b\| \sigma\left(\left\langle \frac{b}{\|b\|}, y \right\rangle\right) \, d\mu_t(a, b) - \int_{a \leq 0, b} |a| \|b\| \sigma\left(\left\langle \frac{b}{\|b\|}, y \right\rangle\right) \, d\mu_t(a, b) \\ &= \int_{u \in \mathbb{S}^{d-1}} \sigma(u^\top y) \, d\nu_t^+(u) - \int_{u \in \mathbb{S}^{d-1}} \sigma(u^\top y) \, d\nu_t^-(u) \end{aligned}$$

E.1.2 Closed dynamics in $d_H + 1$ dimensions

Proof. We first prove that the Equation (4.4) for τ_t^\pm holds in the sense of distributions, and then show that the corresponding dynamics are **closed** because the V_t and g_t appearing in Equation (4.4) can be expressed with (τ_t^+, τ_t^-) (and not only with (ν_t^+, ν_t^-) for instance). We show this by expressing $f(\mu_t; \cdot)$ only in function of the pair (τ_t^+, τ_t^-) .

The pair (τ_t^+, τ_t^-) satisfy Equation (4.4). First, we show that g_t and \tilde{v}_t defined in Equation (4.3) admit modified expressions that match the structure of the pushforward transforming ν_t^\pm into τ_t^\pm . Indeed, since ρ is assumed to be spherically symmetric, it is invariant by any orthogonal transformation. In particular, for a fixed $u \in \mathbb{S}^{d-1}$ such that $u^\perp \neq 0$, we consider the orthogonal map $T^u : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $T^u|_H = \text{id}_H$ and $T^u|_{H^\perp}$ sends the canonical orthonormal basis $(e_1^\perp, \dots, e_{d_\perp}^\perp)$ of H^\perp on $(u^\perp/\|u^\perp\|, u_2, \dots, u_{d_\perp})$ where $(u_2, \dots, u_{d_\perp}) \in (H^\perp)^{d_\perp-1}$ is an orthonormal family, orthogonal to u^\perp , so that for any $y^\perp \in H^\perp$ with coordinates $y_1^\perp, \dots, y_{d_\perp}^\perp$ in the basis $(e_1^\perp, \dots, e_{d_\perp}^\perp)$, $T^u|_{H^\perp}(y^\perp) = y_1^\perp u^\perp/\|u^\perp\| + h_u(y_\perp)$ with $h_u(y_\perp) \perp u^\perp$.

Note that since $f^*(y) = f_H(y^H)$ and $f(\mu_t; y) = f_t(y^H, \|y^\perp\|)$, the residual $R_t(y) = f^*(y) - f(\mu_t; y)$ is invariant by any orthogonal transformation which preserves H (and in particular by T^u). We thus have

$$\begin{aligned} g_t(u) &= \int_y R_t(y) \sigma\left(\langle u^H, y^H \rangle + y_1^\perp \|u^\perp\|\right) \, d\rho =: \tilde{g}_t(u^H, \|u^\perp\|), \\ \tilde{v}_t(u) &= \int_y R_t(y) \sigma'\left(\langle u^H, y^H \rangle + y_1^\perp \|u^\perp\|\right) \left[y^H + T^u|_{H^\perp}(y^\perp) - \left(\langle u^H, y^H \rangle + y_1^\perp \|u^\perp\|\right) u \right] \, d\rho. \end{aligned}$$

Now consider, for any $(\theta, z^H) \in [0, \pi/2] \times \mathbb{S}^{d_H-1}$,

$$\begin{aligned} G_t(\theta, z^H) &:= \tilde{g}_t(\cos(\theta)z^H, \sin(\theta)) \\ V_t(\theta, z^H) &:= \int_y R_t(y) \sigma'\left(\cos(\theta)\langle z^H, y^H \rangle + y_1^\perp \sin(\theta)\right) \left(\frac{y_1^\perp \cos(\theta) - \sin(\theta)\langle z^H, y^H \rangle}{\cos(\theta)} - \langle z^H, y^H \rangle \frac{z^H}{\cos(\theta)} \right) \, d\rho \end{aligned}$$

We show below that (τ_t^+, τ_t^-) satisfy Equation (4.4) with the G_t and V_t defined above. Let $\varphi \in \mathcal{C}_c^1([0, \pi/2] \times \mathbb{S}^{d_H-1})$. Since τ_t^\pm is defined as a push-forward measure obtained from ν_t^\pm

we have:

$$\begin{aligned} \partial_t \int \varphi(\theta, z^H) d\tau_t^\pm(\theta, z^H) &= \partial_t \int \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) d\nu_t^\pm(u) \\ &= \pm 2 \int \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \tilde{g}_t(u^H, \|u^\perp\|) d\nu_t^\pm(u) + \\ &\quad \pm \int \nabla_u \left(\varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \right)^\top \tilde{v}_t(u) d\nu_t^\pm(u). \end{aligned}$$

By definition of the pushforward, and since $u^H = \cos(\arccos(\|u^H\|))u^H/\|u^H\|$ and $\|u^\perp\| = \sin(\arccos(\|u^H\|))$ for $u \in \mathbb{S}^{d-1}$, the first integral is equal $\int \varphi(\theta, z^H) G_t(\theta, z^H) d\tau_t^\pm(u)$. For the second integral, let us first compute the gradient. One has

$$\begin{aligned} \nabla_u \left(\varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \right) &= \partial_\theta \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \frac{-1}{\sqrt{1-\|u^H\|^2}} \frac{u^H}{\|u^H\|} + \\ &\quad \frac{1}{\|u^H\|} \left[I_{d_H} - \frac{u^H(u^H)^\top}{\|u^H\|^2} \right] (\nabla_{z^H} \varphi) \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right). \end{aligned}$$

We observe that the gradient above belongs to H which implies that when computing its inner product with $\tilde{v}_t(u)$ we can consider only the component of the latter along H . Additionally, we note that $I_{d_H} - u^H(u^H)^\top/\|u^H\|^2$ is actually the orthogonal projection onto $\{u^H\}^\perp$, so that it yields 0 when applied to u . Using that $\|u^\perp\| = \sqrt{1-\|u^H\|^2}$ for $u \in \mathbb{S}^{d-1}$, we then get:

$$\nabla_u \left(\varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \right)^\top \tilde{v}_t(u) = \nabla \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right)^\top V_t \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right).$$

where $\nabla \varphi(\theta, z^H) = \begin{pmatrix} \partial_\theta \varphi(\theta, z^H) \\ \nabla_{z^H} \varphi(\theta, z^H) \end{pmatrix}$. This shows that

$$\begin{aligned} \partial_t \int \varphi(\theta, z^H) d\tau_t^\pm(\theta, z^H) &= \pm 2 \int \varphi(\theta, z^H) G_t(\theta, z^H) d\tau_t^\pm(\theta, z^H) + \\ &\quad \pm \int \nabla \varphi(\theta, z^H)^\top V_t(\theta, z^H) d\tau_t^\pm(\theta, z^H), \end{aligned}$$

which proves that τ_t^\pm indeed satisfies Equation (4.3) in the sense of distributions.

The dynamics are closed in the pair (τ_t^+, τ_t^-) . The only thing left to prove to show that the dynamics are closed for the pair (τ_t^+, τ_t^-) is that G_t and V_t can be expressed using only the pair (τ_t^+, τ_t^-) . The only dependence of these quantities on t is through the residual R_t which itself depends on t only through $f(\mu_t; \cdot)$. Let $y \in \mathbb{R}^d$. By definition of ν_t^+ and ν_t^- , we have

$$\begin{aligned} f(\mu_t; y) &= \int_{a,b} a \sigma(b^\top y) d\mu_t(a, b) \\ &= \int_{a \geq 0} |a| |b| \sigma(b^\top y) d\mu_t(a, b) - \\ &\quad \int_{a \leq 0} |a| |b| \sigma(b^\top y) d\mu_t(a, b) \\ &= \int_{u \in \mathbb{S}^{d-1}} \sigma(u^\top y) d(\nu_t^+ - \nu_t^-)(u). \end{aligned}$$

On the other hand, we show below that the integral of any measurable function $\varphi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ against ν_t^\pm can be expressed as an integral against τ_t^\pm in the case where ν_t^\pm admits a density *w.r.t.* the uniform measure on \mathbb{S}^{d-1} (which is the case for ν_0^\pm), the case of a general measure ν_t^\pm being a simple extension via a weak convergence argument. Thus call p_t^\pm the density of ν_t^\pm *w.r.t.* $\tilde{\omega}_d$. Since ν_t^\pm is invariant by any linear map T such that $T|_H = \text{id}_H$ $T|_{H^\perp} \in \mathcal{O}(d_\perp)$ (because of the symmetries on μ_t given by Proposition 2.1), and since this is also the case for $\tilde{\omega}_d$ because $\tilde{\omega}_d$ has spherical symmetry and T is orthogonal, we have by Lemma A.2 that p_t^\pm is invariant by any such T , which then leads to p_t having the form $p_t(u) = \tilde{p}_t^\pm(u^H, \|u^\perp\|)$ by Lemma 4.1.

First step. We show that τ_t^\pm has the density $q_t^\pm(\theta, z^H) = |\mathbb{S}^{d_\perp-1}| \tilde{p}_t^\pm(\cos(\theta)z^H, \sin(\theta))$ *w.r.t.* $\tilde{\gamma} \otimes \tilde{\omega}_{d_H}$ where the measure $\tilde{\gamma}$ is the normalized counterpart of the measure in Definition A.1. Indeed, let $\varphi : [0, \pi/2] \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ be any measurable function *w.r.t.* τ_t^\pm . Using the disintegration Lemma A.4 on $\tilde{\omega}_d$, one has that

$$\begin{aligned} \int \varphi(\theta, z^H) d\tau_t^\pm(\theta, z^H) &= \int \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) d\nu_t^\pm(u) \\ &= \int \varphi \left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|} \right) \tilde{p}_t^\pm(u^H, \|u^\perp\|) d\tilde{\omega}_d(u) \\ &= \int \varphi(\theta, z^H) \tilde{p}_t^\pm(\cos(\theta)z^H, \sin(\theta)) d\tilde{\gamma}(\theta) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) \\ &= \int \varphi(\theta, z^H) \tilde{p}_t^\pm(\cos(\theta)z^H, \sin(\theta)) d\tilde{\gamma}(\theta) d\tilde{\omega}_{d_H}(z^H), \end{aligned}$$

which proves the desired density for τ_t^\pm .

Second step. Consider a measurable $\varphi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ *w.r.t.* ν_t^\pm . One has with similar calculations as above

$$\begin{aligned} \int_u \varphi(u) d\nu_t^\pm(u) &= \int_u \varphi(u) \tilde{p}_t^\pm(u^H, \|u^\perp\|) d\tilde{\omega}_d(u) \\ &= \int \varphi(\cos(\theta)z^H + \sin(\theta)z^\perp) \tilde{p}_t^\pm(\cos(\theta)z^H, \sin(\theta)) d\tilde{\gamma}(\theta) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) \\ &= \int_{\theta, z^H} \left(\int_{z^\perp} \varphi(\cos(\theta)z^H + \sin(\theta)z^\perp) d\tilde{\omega}_{d_\perp}(z^\perp) \right) q_t^\pm(\theta, z^H) d\tilde{\gamma}(\theta) d\tilde{\omega}_{d_H}(z^H) \\ &= \int_{\theta, z^H} \left(\int_{z^\perp} \varphi(\cos(\theta)z^H + \sin(\theta)z^\perp) d\tilde{\omega}_{d_\perp}(z^\perp) \right) d\tau_t^\pm(\theta, z^H). \end{aligned}$$

Applying this to $f(\mu_t; y)$ shows that the latter quantity can be expressed solely using (τ_t^+, τ_t^-) , which proves that the dynamics is indeed closed and therefore concludes the proof when ν_t^\pm has a density.

Third step: extending to any measure. It is known that for any measure ν over \mathbb{S}^{d-1} , there exists a sequence of measure $(\nu_n)_{n \in \mathbb{N}}$ such that: (i) ν_n has a density p_n *w.r.t.* the uniform measure $\tilde{\omega}_d$ over \mathbb{S}^{d-1} , and (ii) the sequence $(\nu_n)_{n \in \mathbb{N}}$ converges weakly to ν , that is, for any continuous (and thus automatically bounded because the unit sphere is compact) φ , $\int \varphi d\nu_n \xrightarrow{n \rightarrow \infty} \int \varphi d\nu$. Let thus $\nu \in \mathcal{M}_+(\mathbb{S}^{d-1})$, and consider a sequence $(\nu_n)_{n \in \mathbb{N}}$ with density converging weakly towards ν . Let τ (*resp.* τ_n) be defined from ν (*resp.* ν_n) as τ_t^\pm is defined

from ν_t^\pm , that is for any measurable $\varphi : [0, \pi/2] \times \mathbb{S}^{d_H-1} \rightarrow \mathbb{R}$,

$$\begin{aligned}\int \varphi(\theta, z^H) d\tau(\theta, z^H) &= \int \varphi\left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|}\right) d\nu(u), \\ \int \varphi(\theta, z^H) d\tau_n(\theta, z^H) &= \int \varphi\left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|}\right) d\nu_n(u).\end{aligned}$$

Let thus φ be a continuous map from $\mathbb{S}^{d-1} \rightarrow \mathbb{R}$ (having in mind the example of $u \mapsto \sigma(u^\top y)$ for a fixed y). By the result of Step 2, since ν_n has a density for every n , we have that

$$\int \varphi(u) d\nu_n(u) = \int_{\theta, z^H} \left(\int_{z^\perp} \varphi(\cos(\theta)z^H + \sin(\theta)z^\perp) d\tilde{\omega}_{d_\perp}(z^\perp) \right) d\tau_n(\theta, z^H), \quad (\text{E.1})$$

and taking the limit $n \rightarrow \infty$, the left-hand-side converges of Equation (E.1) to $\int \varphi d\nu$ by assumption. Now let us look at the right-hand-side of (E.1). Calling $\psi(\theta, z^H) = \int_{z^\perp} \varphi(\cos(\theta)z^H + \sin(\theta)z^\perp) d\tilde{\omega}_{d_\perp}(z^\perp)$ and $\Phi(u) = \int_{z^\perp} \varphi(u^H + \|u^\perp\|z^\perp) d\tilde{\omega}_{d_\perp}(z^\perp)$, the right-hand-side is in fact $\int \psi d\tau_n$ and, for any $n \in \mathbb{N}$, is equal to:

$$\begin{aligned}\int \psi d\tau_n &= \int_u \psi\left(\arccos(\|u^H\|), \frac{u^H}{\|u^H\|}\right) d\nu_n(u) \\ &= \int_u \int_{z^\perp} \varphi\left(\|u^H\| \frac{u^H}{\|u^H\|} + \|u^\perp\|z^\perp\right) d\tilde{\omega}_{d_\perp}(z^\perp) d\nu_n(u) \\ &= \int_u \int_{z^\perp} \varphi\left(u^H + \|u^\perp\|z^\perp\right) d\tilde{\omega}_{d_\perp}(z^\perp) d\nu_n(u) \\ &= \int \Phi d\nu_n,\end{aligned}$$

and a similar result holds for τ and ν . Now, the continuity of Φ is readily obtained from that of φ , and thus the right-hand-side in the last equality above converges to $\int \Phi d\nu$ which is also equal to $\int \psi d\tau$ by the same calculations as above. The right-hand-side in (E.1) therefore converges to $\int \psi d\tau$, which is the claim of Step 2 for a general measure ν which does not necessarily admit a density, thereby concluding the proof. \square

E.2 Case when f^* is the euclidean norm: [Theorem 4.3](#)

Here, we give the proof of [Theorem 4.3](#) which shows that when $f^*(x) = \|x^H\|$ the dynamics can be reduced to a single variable: the angle $\theta \in [0, \pi/2]$ between particles and the subspace H .

We decompose the proof in three steps: first we show that the pair of measures $(\tau_t^+, \tau_t^-) \in \mathcal{M}_+([0, \pi/2])$ as defined in [Section 4.2](#) indeed follows [Equation \(4.5\)](#); then we show that the dynamics are indeed closed by proving that the terms V_t and G_t appearing in the GF depend only on (τ_t^+, τ_t^-) ; and finally, we show that [Equation \(4.5\)](#) indeed corresponds to a Wasserstein-Fisher-Rao GF on a given objective functional over $\mathcal{M}_+([0, \pi/2])^2$.

E.2.1 Proof of the GF equation

Proof. We first use the added symmetry to simplify the terms g_t and \tilde{v}_t which appear in the GF with (ν_t^+, ν_t^-) (see [Section 4.1](#)) and express them only with $\|u^H\|$ and $\|u^\perp\|$. Then we use the equations satisfied by (ν_t^+, ν_t^-) to obtain equations for (τ_t^+, τ_t^-) .

Equations for (τ_t^+, τ_t^-) . Let $\varphi \in \mathcal{C}_c^1([0, \pi/2])$. We have

$$\begin{aligned} \partial_t \int \varphi d\tau_t^\pm &= \partial_t \int \varphi (\arccos(\|u^H\|)) d\nu_t^\pm(u) \\ &= \pm \int \nabla_u (\varphi (\arccos(\|u^H\|)))^\top \tilde{v}_t(u) d\nu_t^\pm(u) \\ &\quad \pm 2 \int \varphi (\arccos(\|u^H\|)) g_t(u) d\nu_t^\pm(u). \end{aligned}$$

One has that

$$\nabla_u (\varphi (\arccos(\|u^H\|))) = \varphi' (\arccos(\|u^H\|)) \times \frac{-1}{\sqrt{1 - \|u^H\|^2}} \frac{u^H}{\|u^H\|},$$

which belongs to H . We recall here the expressions of \tilde{v}_t and g_t : for any $u \in \mathbb{S}^{d-1}$, we have

$$\begin{aligned} g_t(u) &= \int_y R_t(y) \sigma(u^\top y) d\rho(y), \\ \tilde{v}_t(u) &= \int_y R_t(y) \sigma'(u^\top y) [y - (u^\top y)u] d\rho(y). \end{aligned}$$

Since, $f^*(x) = \|x^H\|$, f^* is now invariant under any orthogonal map T preserving H and H^\perp , that is such that the restrictions $T|_H \in \mathcal{O}(d_H)$ and $T|_{H^\perp} \in \mathcal{O}(d_\perp)$. Proposition 2.1 then ensures that so is $f(\mu_t, \cdot)$, which in turn implies that the residual $R_t = f^* - f(\mu_t; \cdot)$ also shares that invariance property. Using a similar change of variable as in Appendix E.1.2, and because ρ is spherically symmetric, one gets that g_t can be re-written

$$g_t(u) = \int_y R_t(y) \sigma \left(y_1^H \|u^H\| + y_1^\perp \|u^\perp\| \right) d\rho(y).$$

Calling

$$G_t(\theta) := \int_y R_t(y) \sigma \left(y_1^H \cos(\theta) + y_1^\perp \sin(\theta) \right) d\rho(y),$$

one has $g_t(u) = G_t(\arccos(\|u^H\|))$ because $u \in \mathbb{S}^{d-1}$, so that $\|u^\perp\| = \sqrt{1 - \|u^H\|^2}$. Then, by definition of τ_t^\pm , the second integral in the time derivative above is equal to $\int \varphi(\theta) G_t(\theta) d\tau_t^\pm$. For the first integral appearing in that time derivative, we get

$$\nabla_u (\varphi (\arccos(\|u^H\|)))^\top \tilde{v}_t(u) = \frac{\varphi' (\arccos(\|u^H\|))}{\|u^\perp\| \|u^H\|} \int_y R_t(y) \sigma'(u^\top y) [(u^\top y)u - y]^\top u^H d\rho.$$

Expanding the inner product inside the integral, we have

$$\begin{aligned} [(u^\top y)u - y]^\top u^H &= (\langle u^H, y^H \rangle + \langle u^\perp, y^\perp \rangle) \|u^H\|^2 - \langle u^H, y^H \rangle \\ &= \|u^H\|^2 \langle u^\perp, y^\perp \rangle - (1 - \|u^H\|^2) \langle u^H, y^H \rangle \\ &= \|u^H\|^2 \langle u^\perp, y^\perp \rangle - \|u^\perp\|^2 \langle u^H, y^H \rangle. \end{aligned}$$

Calling

$$V_t(\theta) := \int_y R_t(y) \sigma' \left(y_1^H \cos(\theta) + y_1^\perp \sin(\theta) \right) [y_1^\perp \cos(\theta) - y_1^H \sin(\theta)] d\rho(y) = G'(\theta),$$

and using again the spherical symmetry of ρ , with the same change of variable in the integral as for g_t , we get that

$$\nabla_u (\varphi (\arccos(\|u^H\|)))^\top \tilde{v}_t(u) = \varphi' (\arccos(\|u^H\|)) V_t(\arccos(\|u^H\|)).$$

Finally, this combined with the previous result on the integral with g_t yields

$$\partial_t \int \varphi d\tau_t^\pm = \pm \int \varphi'(\theta) V_t(\theta) d\tau_t^\pm(\theta) \pm 2 \int \varphi(\theta) G_t(\theta) d\tau_t^\pm(\theta),$$

which leads to the desired equation

$$\partial \tau_t^\pm = -\operatorname{div} (\pm V_t \tau_t^\pm) \pm 2 G_t \tau_t^\pm.$$

□

E.2.2 Proof that the dynamics on the angle θ are closed

The proof follow closely that of Appendix E.1.2 (where we prove closed dynamics), except here we take advantage of the added symmetry of the dynamics. As in Appendix E.1.2, we have

$$f(\mu_t; y) = \int_{u \in \mathbb{S}^{d-1}} \sigma(u^\top y) d(\nu_t^+ - \nu_t^-)(u),$$

and the only thing to prove is that this quantity can be expressed using only (τ_t^+, τ_t^-) . As in Appendix E.1.2, we first prove this when ν_t^\pm has a density, which is the case for ν_0^\pm and should thus remain so during the dynamics.

Similarly to what occurs in Appendix E.1.2, ν_t^\pm is invariant by any orthogonal map T which preserves H and H^\perp because μ_t has those symmetries given by Proposition 2.1, and if ν_t^\pm has a density $p_{\nu_t^\pm}$ w.r.t. $\tilde{\omega}_d$, then p_t^\pm is also invariant by any such map T , and thus depends only on the norms $\|u^H\|$ and $\|u^\perp\|$ of its input $u \in \mathbb{S}^{d-1}$. But since its input is on the sphere, those norms are determined by the angle $\theta = \arccos(\|u^H\|)$ between the input u and H . Calling q_t^\pm such that $p_t^\pm(u) = q_t^\pm(\arccos(\|u^H\|))$, this will lead τ_t^\pm to have the density q_t^\pm w.r.t. $\tilde{\gamma}$. Then, we show below that similarly to Appendix E.1.2, the integral of any measurable $\varphi : \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ against ν_t^\pm can be expressed as an integral against τ_t^\pm . Indeed, using the disintegration Lemma A.4,

$$\begin{aligned} \int \varphi d\nu_t^\pm &= \int_{\theta \in [0, \pi/2]} \varphi(u) q_t^\pm(\arccos(\|u^H\|)) d\tilde{\omega}_d(u) \\ &= \int_u \varphi \left(\cos(\theta) z^H + \sin(\theta) z^\perp \right) q_t^\pm(\theta) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) d\tilde{\gamma}(\theta) \\ &= \int_{\theta \in [0, \pi/2]} \tilde{\varphi}(\theta) q_t^\pm(\theta) d\tilde{\gamma}(\theta) \\ &= \int_{\theta \in [0, \pi/2]} \tilde{\varphi}(\theta) d\tau_t^\pm(\theta) \end{aligned}$$

where

$$\tilde{\varphi}(\theta) := \int_{z^H, z^\perp} \varphi \left(\cos(\theta) z^H + \sin(\theta) z^\perp \right) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp),$$

which concludes the proof if ν_t^\pm has a density w.r.t. the uniform measure $\tilde{\omega}_d$ on the sphere \mathbb{S}^{d-1} . The general case is obtained by a weak convergence argument (of measures with density) as in the third step of Section E.1.2.

E.2.3 Proof of the Wasserstein-Fisher-Rao GF

Proof. Recall that γ is the measure in Definition A.1, and consider the following objective functional over $\mathcal{M}([0, \pi/2])^2$:

$$\begin{aligned} A(\tau^+, \tau^-) &:= \int_{\varphi \in [0, \pi/2]} \ell(\cos(\varphi), \tilde{f}(\tau^+, \tau^-; \varphi)) d\tilde{\gamma}(\varphi), \\ \tilde{f}(\tau^+, \tau^-; \varphi) &:= \int_{\theta \in [0, \pi/2]} \tilde{\phi}(\theta; \varphi) d(\tau^+ - \tau^-)(\theta), \\ \tilde{\phi}(\theta; \varphi) &:= \int_{r, s \in [-1, 1]} \sigma(r \cos(\varphi) \cos(\theta) + s \sin(\varphi) \sin(\theta)) d\tilde{\gamma}_{d_H}(r) d\tilde{\gamma}_{d_\perp}(s) \end{aligned}$$

where, for any $p \in \mathbb{N}$, $d\gamma_p(r) = (1-r^2)^{(p-3)/2} dr$, and $\tilde{\gamma}_p = \gamma_p/|\gamma_p|$ with the normalizing factor $|\gamma_p| = B(1/2, (p-1)/2) = \sqrt{\pi} \Gamma((p-1)/2) / \Gamma(p/2) = |\mathbb{S}^{p-1}| / |\mathbb{S}^{p-2}|$. Note that $\tilde{\gamma}_p$ can be simply expressed as the law of $\epsilon \times \sqrt{X}$ where $\epsilon \sim \mathcal{U}(\{-1, +1\})$ and $X \sim \text{Beta}(1/2, (p-1)/2)$.

Computing the first variation or Fréchet derivative of the functional A *w.r.t.* to its first and second argument yields, for any $\theta \in [0, \pi/2]$,

$$\frac{\delta A}{\delta \tau^\pm}(\tau^+, \tau^-)[\theta] = \pm \int_{\varphi} \partial_2 \ell(\cos(\varphi), \tilde{f}(\tau^+, \tau^-; \varphi)) \tilde{\phi}(\theta; \varphi) d\tilde{\gamma}(\varphi).$$

To conclude one needs only observe that the quantity above is simply equal to $G_t(\theta)$, up to a fixed multiplicative constant. Since we have assumed ρ to be the uniform measure over \mathbb{S}^{d-1} to ensure that the Wasserstein GF (1.3) is well-defined, the constant is one here but in the case of a general ρ with spherical symmetry, the result should also hold (as long as the Wasserstein GF (1.3) is well-defined) but the proof is more technical and different constants might appear.

Simplifying $f(\mu_t; \cdot)$. Using the results from Appendix E.2.2, we have for any $\varphi, z^H, z^\perp \in [0, \pi/2] \times \mathbb{S}^{d_H-1} \times \mathbb{S}^{d_\perp-1}$ (so that $u = \cos(\varphi)z^H + \sin(\varphi)z^\perp \in \mathbb{S}^{d-1}$)

$$\begin{aligned} f(\mu_t; \cos(\varphi)z^H + \sin(\varphi)z^\perp) &= \\ &\int_{\psi} \int_{\xi^H, \xi^\perp} \sigma(\cos(\psi) \cos(\varphi) \langle \xi^H, z^H \rangle + \sin(\psi) \sin(\varphi) \langle \xi^\perp, z^\perp \rangle) d\tilde{\omega}_{d_H}(\xi^H) d\tilde{\omega}_{d_\perp}(\xi^\perp) d(\tau_t^+ - \tau_t^-)(\psi) \end{aligned}$$

Now, because of the integration against uniform measures on the unit spheres, and the inner products involved, we can use some spherical harmonics theory to simplify those calculations. Using The Funk-Hecke formula (see Atkinson and Han, 2012, Theorem 2.22, $n = 0$, $d = d_H$ or $d = d_\perp$), we get

$$\begin{aligned} f(\mu_t; \cos(\varphi)z^H + \sin(\varphi)z^\perp) &= \\ &\frac{|\mathbb{S}^{d_H-2}| |\mathbb{S}^{d_\perp-2}|}{|\mathbb{S}^{d_H-1}| |\mathbb{S}^{d_\perp-1}|} \int_{\psi} \int_{r, s} \sigma(r \cos(\psi) \cos(\varphi) + s \sin(\psi) \sin(\varphi)) d\gamma_{d_H}(r) d\gamma_{d_\perp}(s) d(\tau_t^+ - \tau_t^-)(\psi) \\ &= \frac{1}{|\gamma_{d_H}| |\gamma_{d_\perp}|} \int_{\psi \in [0, \pi/2]} \tilde{\phi}(\psi; \varphi) d(\tau_t^+ - \tau_t^-)(\psi) \\ &= \tilde{f}(\tau_t^+, \tau_t^-; \varphi). \end{aligned}$$

Simplifying $f^*(\cos(\varphi)z^H + \sin(\varphi)z^\perp)$. Because $f^*(y) = \|y^H\|$, $f^*(\cos(\varphi)z^H + \sin(\varphi)z^\perp)$ is simply $\|\cos(\varphi)z^H\| = \cos(\varphi)$ because $z^H \in \mathbb{S}^{d_H-1}$.

With the previous expressions for $f(\mu_t; \cdot)$ and f^* we have that for any function $\Phi : \mathbb{R}^2 \rightarrow \mathbb{R}$,

$$\Phi\left(f^*(\cos(\varphi)z^H + \sin(\varphi)z^\perp), f(\mu_t; \cos(\varphi)z^H + \sin(\varphi)z^\perp)\right) = \Phi\left(\cos(\varphi), \tilde{f}(\tau_t^+, \tau_t^-; \varphi)\right).$$

Note that this applies both to $\Phi(y, \hat{y}) = \ell(y, \hat{y})$ and $\Phi(y, \hat{y}) = -\partial_2 \ell(y, \hat{y})$.

Proof that $F(\mu_t) = A(\tau_t^+, \tau_t^-)$. Using the disintegration Lemma A.4 for the uniform measure on the unit sphere \mathbb{S}^{d-1} , we have

$$\begin{aligned} F(\mu_t) &= \int_y \ell\left(f^*(y), f(\mu_t; y)\right) d\rho(y) \\ &= \int \ell \circ (f^*(\cdot), f(\mu_t; \cdot)) \left(\cos(\varphi)z^H + \sin(\varphi)z^\perp\right) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) d\tilde{\gamma}(\varphi) \\ &= \int \ell\left(\cos(\varphi), \tilde{f}(\tau_t^+, \tau_t^-; \varphi)\right) d\tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) d\tilde{\gamma}(\varphi) \\ &= \int \ell\left(\cos(\varphi), \tilde{f}(\tau_t^+, \tau_t^-; \varphi)\right) d\tilde{\gamma}(\varphi), \end{aligned}$$

where we have used in the last equality the fact that the integrand does not depend on z^H or z^\perp and that $\tilde{\omega}_{d_H}$ and $\tilde{\omega}_{d_\perp}$ are probability measures.

Simplifying G_t . Using the disintegration Lemma A.4, we have:

$$\begin{aligned} G_t(\theta) &= \int_y R_t(y) \sigma\left(y_1^H \cos(\theta) + y_1^\perp \sin(\theta)\right) d\rho(y) \\ &= \int R_t(\cos(\varphi)z^H + \sin(\varphi)z^\perp) \sigma\left(z_1^H \cos(\varphi) \cos(\theta) + z_1^\perp \sin(\varphi) \sin(\theta)\right) \tilde{\omega}_{d_H}(z^H) d\tilde{\omega}_{d_\perp}(z^\perp) d\tilde{\gamma}(\varphi). \end{aligned}$$

Similarly to what we did for simplifying $f(\mu_t; \cdot)$, we can simplify the integrals against $\tilde{\omega}_{d_H}$ and $\tilde{\omega}_{d_\perp}$ using spherical harmonics theory to get:

$$G_t(\theta) = - \int_{\varphi \in [0, \pi/2]} \partial_2 \ell\left(\cos(\varphi), \tilde{f}(\tau_t^+, \tau_t^-; \varphi)\right) \tilde{\phi}(\varphi; \theta) d\tilde{\gamma}(\varphi).$$

This shows that

$$\begin{aligned} -\frac{\delta A}{\delta \tau^+}(\tau_t^+, \tau_t^-)[\theta] &= G_t(\theta) \\ \frac{\delta A}{\delta \tau^-}(\tau_t^+, \tau_t^-)[\theta] &= G_t(\theta), \end{aligned}$$

which proves that Equation (4.5) indeed describes the evolution of the Wasserstein-Fisher-Rao for the objective functional A over $\mathcal{M}([0, \pi/2])^2$, given by the pair (τ_t^+, τ_t^-) . \square

F Numerical simulations in one dimension

Measure discretization. Discretizing μ_t via $\mu_{m,t} = \frac{1}{m} \sum_{j=1}^m \delta_{(a_j(t), b_j(t))}$, we get that $\tau_{m,t} := \tau_{m,t}^+ - \tau_{m,t}^- = \frac{1}{m} \sum_{j=1}^m c_j(t) \delta_{\theta_j(t)}$ where

$$\begin{aligned} c_j(t) &= \varepsilon_j |a_j(t)| \|b_j(t)\|, \\ \varepsilon_j &= \text{sign}(a_j(0)), \\ \theta_j(t) &= \arccos\left(\frac{b_j(t)}{\|b_j(t)\|}\right). \end{aligned}$$

Initializing through $a_j(0) \sim \mathcal{U}\{-1, +1\}$ and $b_j(0) \sim \tilde{\omega}_d = \mathcal{U}(\mathbb{S}^{d-1})$, yields $c_j(0) \sim \mathcal{U}\{-1, +1\}$ and $\theta_j(0) \sim \tilde{\gamma}$, i.i.d. over j . The gradient flows of Equation (4.5) translates into the following ODEs on $(c_j)_{j \in [1, m]}$ and $(\theta_j)_{j \in [1, m]}$:

$$\begin{aligned}\frac{d}{dt}c_j(t) &= 2\varepsilon_j G_t(\theta_j(t))c_j(t), \\ \frac{d}{dt}\theta_j(t) &= \varepsilon_j V_t(\theta_j(t)).\end{aligned}$$

where $\varepsilon_j = a_j(0) \in \{-1, +1\}$ denotes whether the corresponding quantity appears in $\tau_{t, m}^+$ ($\varepsilon = +1$) or $\tau_{t, m}^-$ ($\varepsilon = -1$).

Time discretization. Simulating these ODEs via the discrete Euler scheme with step $\eta > 0$, leads, for any iteration $k \in \mathbb{N}$, to:

$$\begin{aligned}c_j(k+1) &= \left(1 + 2\eta\varepsilon_j G_k(\theta_j(k))\right)c_j(k) \\ \theta_j(k+1) &= \theta_j(k) + \eta\varepsilon_j V_k(\theta_j(k)).\end{aligned}\tag{F.1}$$

Approximating integrals numerically. The only thing that needs to be dealt with numerically is estimating the values of G_t and V_t which are defined by integrals. With the discretization of the measures, we have:

$$\begin{aligned}G_k(\theta) &= \int_{\varphi} \left(\cos(\varphi) - \tilde{f}(\tau_k^+, \tau_k^-; \varphi)\right) \tilde{\phi}(\varphi; \theta) d\tilde{\gamma}(\varphi), \\ \tilde{f}(\tau_k^+, \tau_k^-; \varphi) &= \sum_{j=1}^m c_j(k) \tilde{\phi}(\theta_j(k); \varphi), \\ \tilde{\phi}(\theta; \varphi) &= \int_{r, s \in [-1, 1]} \sigma\left(r \cos(\varphi) \cos(\theta) + s \sin(\varphi) \sin(\theta)\right) d\tilde{\gamma}_{d_H}(r) d\tilde{\gamma}_{d_{\perp}}(s).\end{aligned}$$

We thus get:

$$G_k(\theta) = \int \frac{\psi(r, s; \theta, \varphi)}{m} \sum_{j=1}^m \left(\cos(\varphi) - c_j(k) \psi(r', s'; \theta_j(k), \varphi)\right) d\tilde{\gamma}(\varphi) (d\tilde{\gamma}_{d_H})^2(r, r') (d\tilde{\gamma}_{d_{\perp}})^2(s, s'),$$

with

$$\psi(r, s; \theta, \varphi) := \sigma\left(r \cos(\varphi) \cos(\theta) + s \sin(\varphi) \sin(\theta)\right).$$

Similarly, we have:

$$V_k(\theta) = \int \frac{\chi(r, s; \theta, \varphi)}{m} \sum_{j=1}^m \left(\cos(\varphi) - c_j(k) \psi(r', s'; \theta_j(k), \varphi)\right) d\tilde{\gamma}(\varphi) (d\tilde{\gamma}_{d_H})^2(r, r') (d\tilde{\gamma}_{d_{\perp}})^2(s, s'),$$

with

$$\begin{aligned}\chi(r, s; \theta, \varphi) &:= \frac{\partial}{\partial \theta} \psi(r, s; \theta, \varphi) \\ &= \sigma'\left(\cos(\theta) \cos(\varphi) r + \sin(\theta) \sin(\varphi) s\right) \left[-\sin(\theta) \cos(\varphi) r + \cos(\theta) \sin(\varphi) s\right].\end{aligned}$$

We use Monte-Carlo estimation through sampling to approximate the integrals against the five variables (φ, r, r', s, s') by drawing N samples from the corresponding distributions. We get:

$$\begin{aligned} G_k(\theta_j(k)) &\approx \frac{1}{mN} \sum_{i=1}^N \sum_{l=1}^m \Psi_{ji} \left(\cos(\Phi_i) - c_l(k) \tilde{\Psi}_{li} \right), \\ \Psi_{ji}(k) &= \psi(R_i, S_i; \theta_j(k), \Phi_i) \\ \tilde{\Psi}_{ji}(k) &= \psi(R'_i, S'_i; \theta_j(k), \Phi_i), \end{aligned}$$

and similarly

$$\begin{aligned} V_k(\theta_j(k)) &\approx \frac{1}{mN} \sum_{i=1}^N \sum_{j=1}^m \chi_{ji} \left(\cos(\Phi_i) - c_l(k) \tilde{\Psi}_{li} \right), \\ \chi_{ji}(k) &= \chi(R_i, S_i; \theta_j(k), \Phi_i), \end{aligned}$$

where we have drawn the samples i.i.d. over $i \in [1, N]$:

$$\begin{aligned} \Phi_i &\sim \tilde{\gamma}, \\ R_i, R'_i &\sim \tilde{\gamma}_{d_H}, \\ S_i, S'_i &\sim \tilde{\gamma}_{d_\perp}. \end{aligned}$$

Iterations in the numerical simulation. Defining the vectors $c(k) = (c_j(k))_{j \in [1, m]}$, $\theta(k) = (\theta_j)_{j \in [1, m]}$, and $\varepsilon = (\varepsilon_j)_{j \in [1, m]}$, the update Equations (F.1) can then be written in terms of update rules using the matrices $\Psi(k) = (\Psi_{ji}(k))_{j, i \in [1, m] \times [1, N]}$, $\tilde{\Psi}(k) = (\tilde{\Psi}_{ji}(k))_{j, i \in [1, m] \times [1, N]}$, and finally $\chi = (\chi_{ji}(k))_{j, i \in [1, m] \times [1, N]}$, and the vectors $(\Phi, R, R', S, S') = (\Phi_i, R_i, R'_i, S_i, S'_i)_{i \in [1, N]}$, which are re-sampled at each iteration $k \in [0, K]$, where $K \in \mathbb{N}$:

$$\begin{aligned} c(k+1) &= (1 + 2\eta\varepsilon \odot \hat{G}_k) \odot c(k), \\ \theta(k+1) &= \theta(k) + \eta\varepsilon \odot \hat{V}_k, \end{aligned}$$

where

$$\begin{aligned} \hat{G}_k &= \frac{d}{N_b} \Psi \left(\cos(\Phi) - \frac{1}{m} \tilde{\Psi}^\top c(k) \right), \\ \hat{V}_k &= \frac{d}{N_b} \chi \left(\cos(\Phi) - \frac{1}{m} \tilde{\Psi}^\top c(k) \right), \end{aligned}$$

and \odot denotes the Hadamard (element-wise) product of two vectors. One can compute the loss through sampling in a similar way.

Experimental value for α and parameters of the numerical simulation. For the numerical simulations, we fix the number of atoms of the measure (or equivalently the width of the network) to $m = 1,024$, the learning rate to $\eta = 5.10^{-3}$, the number of samples for the Monte-Carlo scheme to $N = 1,000$, and the total number of iterations to $K = 20,000$. The experimental value for α (see Section 4.2) is computed through $\alpha_{\text{exp}} = \tau_{m, K}^+([0, \pi/2])$, that is

$$\begin{aligned} \alpha_{\text{exp}} &= \frac{1}{m} \sum_{j \in J^+} c_j(K), \\ J^+ &:= \{j \in [1, m] ; \varepsilon_j = 1\}. \end{aligned}$$

As mentioned in the main text, the behaviour of the numerical simulation depends a lot on the step-size η . Some of the differences between our observations and our intuitive description of the limiting model (infinite-width and continuous time) can come from too big a step-size. We have thus run the numerical simulation with $\eta = 2 \cdot 10^{-5}$ as well, for $K = 230,000$ steps but the same differences still appear (*e.g.*, $\tau_{m,k}^+([0, \pi/2])$ still grows larger than the theoretically expected limit α after some time, albeit by a smaller margin) and after the critical t^* , some negative particles seem to go beyond $\pi/2$, even with a very small step-size, a fact which cannot happen for the limiting model. Consequently, in Figure 3, the first histogram bin right after $\pi/2$ has been merged with the one before.