

Convergence of Anisotropic Consensus-Based Optimization in Mean-Field Law

Massimo Fornasier^{1,2*}, Timo Klock^{3†}, and Konstantin Riedl^{1‡}

¹ Technical University of Munich, Department of Mathematics, Munich, Germany

² Munich Data Science Institute, Munich, Germany

³ Simula Research Laboratory, Department of Numerical Analysis and Scientific Computing, Oslo, Norway

Abstract. In this paper we study anisotropic consensus-based optimization (CBO), a population-based metaheuristic derivative-free optimization method capable of globally minimizing nonconvex and nonsmooth functions in high dimensions. CBO is based on stochastic swarm intelligence, and inspired by consensus dynamics and opinion formation. Compared to other metaheuristic algorithms like Particle Swarm Optimization, CBO is of a simpler nature and therefore more amenable to theoretical analysis. By adapting a recently established proof technique, we show that anisotropic CBO converges globally with a dimension-independent rate for a rich class of objective functions under minimal assumptions on the initialization of the method. Moreover, the proof technique reveals that CBO performs a convexification of the optimization problem as the number of particles goes to infinity, thus providing an insight into the internal CBO mechanisms responsible for the success of the method. To motivate anisotropic CBO from a practical perspective, we further test the method on a complicated high-dimensional benchmark problem, which is well understood in the machine learning literature.

Keywords: High-Dimensional Global Optimization · Metaheuristics · Consensus-Based Optimization · Mean-Field Limit · Anisotropy

1 Introduction

Several problems arising throughout all quantitative disciplines are concerned with the global unconstrained optimization of a problem-dependent objective function $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}$ and the search for the associated minimizing argument

$$v^* = \arg \min_{v \in \mathbb{R}^d} \mathcal{E}(v),$$

which is assumed to exist and be unique in what follows. Because of nowadays data deluge such optimization problems are usually high-dimensional. In machine learning, for instance, one is interested in finding the optimal parameters

* Email: massimo.fornasier@ma.tum.de

† Email: timo@simula.no

‡ Email: konstantin.riedl@ma.tum.de

of a neural network (NN) to accomplish various tasks, such as clustering, classification, and regression. The availability of huge amounts of training data for various real-world applications allows practitioners to work with models involving a large number of trainable parameters aiming for a high expressivity and accuracy of the trained model. This makes the resulting optimization process a high-dimensional problem. Since typical model architectures consist of many layers with a large amount of neurons, and include nonlinear and potentially non-smooth activation functions, the training process is in general a high-dimensional nonconvex optimization problem and therefore a particularly hard task.

Metaheuristics have a long history as state-of-the-art methods when it comes to tackling hard optimization problems. Inspired by self-organization and collective behavior in nature or human society, such as the swarming of flocks of birds or schools of fish [3], or opinion formation [20], they orchestrate an interplay between locally confined procedures and global strategies, randomness and deterministic decisions, to ensure a robust search for the global minimizer. Some prominent examples are Random Search [19], Evolutionary Programming [7], Genetic Algorithms [11], Ant Colony Optimization [6], Particle Swarm Optimization [14] and Simulated Annealing [1].

CBO follows those guiding principles, but is of much simpler nature and more amenable to theoretical analysis. The method uses N particles V^1, \dots, V^N , which are initialized independently according to some law $\rho_0 \in \mathcal{P}(\mathbb{R}^d)$, to explore the domain and to form a global consensus about the minimizer v^* as time passes. For parameters $\alpha, \lambda, \sigma > 0$ the dynamics of each particle is given by

$$dV_t^i = -\lambda (V_t^i - v_\alpha(\widehat{\rho}_t^N)) dt + \sigma D(V_t^i - v_\alpha(\widehat{\rho}_t^N)) dB_t^i, \quad (1)$$

where $\widehat{\rho}_t^N$ denotes the empirical measure of the particles. The first term in (1) is a drift term dragging the respective particle towards the momentaneous consensus point, a weighted average of the particles' positions, computed as

$$v_\alpha(\widehat{\rho}_t^N) := \int v \frac{\omega_\alpha(v)}{\|\omega_\alpha\|_{L_1(\widehat{\rho}_t^N)}} d\widehat{\rho}_t^N(v), \quad \text{with} \quad \omega_\alpha(v) := \exp(-\alpha \mathcal{E}(v))$$

and motivated by the fact that $v_\alpha(\widehat{\rho}_t^N) \approx \arg \min_{i=1, \dots, N} \mathcal{E}(V_t^i)$ for $\alpha \gg 1$ if the arg min is unique. To feature the exploration of the energy landscape of \mathcal{E} , the second term in (1) is a diffusion injecting randomness into the dynamics through independent standard Brownian motions $((B_t^i)_{t \geq 0})_{i=1, \dots, N}$. The two commonly studied diffusion types are isotropic [2, 9, 18] and anisotropic [4] diffusion with

$$D(V_t^i - v_\alpha(\widehat{\rho}_t^N)) = \begin{cases} \|V_t^i - v_\alpha(\widehat{\rho}_t^N)\|_2 \text{Id}, & \text{for isotropic diffusion,} \\ \text{diag}(V_t^i - v_\alpha(\widehat{\rho}_t^N)), & \text{for anisotropic diffusion,} \end{cases}$$

where $\text{Id} \in \mathbb{R}^{d \times d}$ is the identity matrix and $\text{diag} : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ the operator mapping a vector onto a diagonal matrix with the vector as its diagonal. The term's scaling encourages in particular particles far from $v_\alpha(\widehat{\rho}_t^N)$ to explore larger regions. The coordinate-dependent scaling of anisotropic diffusion has proven to

be particularly beneficial for high-dimensional optimization problems by yielding dimension-independent convergence rates (see Figure 1) and therefore improving both computational complexity and success probability of the algorithm [4, 8].

A theoretical convergence analysis of the CBO dynamics is possible either on the microscopic level (1) or by analyzing the macroscopic behavior of the particle density through a mean-field limit. In the large particle limit a particle is not influenced by individual particles but only by the average behavior of all particles. As shown in [12], the empirical random particle measure $\hat{\rho}^N$ converges in law to the deterministic particle density $\rho \in \mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))$, which weakly (see Definition 1) satisfies the non-linear Fokker-Planck equation

$$\partial_t \rho_t = \lambda \operatorname{div}((v - v_\alpha(\rho_t)) \rho_t) + \frac{\sigma^2}{2} \sum_{k=1}^d \partial_{kk} \left(D(v - v_\alpha(\rho_t))_{kk}^2 \rho_t \right). \quad (2)$$

A quantitative analysis of the convergence rate remains, on non-compact domains, an open problem, see, e.g., [9, Remark 2]. Analyzing a mean-field limit such as (2) allows for establishing strong qualitative theoretical guarantees about CBO methods, paving the way to understand the internal mechanisms at play.

Prior Arts. The original CBO work [18] proposes the dynamics (1) with isotropic diffusion, which is analyzed in the mean-field sense in [2]. Under a stringent well-preparedness condition about ρ_0 and \mathcal{C}^2 regularity of \mathcal{E} the authors show consensus formation of the particles at some \tilde{v} close to v^* by first establishing exponential decay of the variance $\operatorname{Var}(\rho_t)$ and consecutively showing $\tilde{v} \approx v^*$ as a consequence of the Laplace principle [17]. This analysis is extended to the anisotropic case in [4].

Motivated by the surprising phenomenon that, on average, individual particles of the CBO dynamics follow the gradient flow of $v \mapsto \|v - v^*\|_2^2$, see [9, Figure 1b], the authors of [9] develop a novel proof technique for showing global convergence of isotropic CBO in mean-field law to v^* under minimal assumptions. Following [9, Definition 1], we speak of convergence in mean-field law to v^* for the interacting particle dynamics (1), if the solution ρ_t to the associated mean-field limit dynamics (2) converges narrowly to the Dirac delta δ_{v^*} at v^* for $t \rightarrow \infty$. The proof is based on showing an exponential decay of the energy functional $\mathcal{V} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}_{\geq 0}$, given by

$$\mathcal{V}(\rho_t) := \frac{1}{2} \int \|v - v^*\|_2^2 d\rho_t(v). \quad (3)$$

This simultaneously ensures consensus formation and convergence of ρ_t to δ_{v^*} .

Contribution. In view of the effectiveness and efficiency of CBO methods with anisotropic diffusion for high-dimensional optimization problems, a thorough understanding is of considerable interest. As we illustrate in Figure 1, anisotropic CBO [4] converges with a dimension-independent rate as opposed to




Fig. 1: A demonstration of the benefit of using anisotropic diffusion in CBO. For the Rastrigin function $\mathcal{E}(v) = \sum_{k=1}^d v_k^2 + \frac{5}{2}(1 - \cos(2\pi v_k))$ with $v^* = 0$ and spurious local minima (see (a)), we evolve the discretized system of isotropic and anisotropic CBO using $N = 320000$ particles, discrete time step size $\Delta t = 0.01$ and $\alpha = 10^{15}$, $\lambda = 1$, and $\sigma = 0.32$ for different dimensions $d \in \{4, 8, 12, 16\}$. We observe in (b) that the convergence rate of the energy functional $\mathcal{V}(\hat{\rho}_t^N)$ for isotropic CBO (dashed lines) is affected by the ambient dimension d , whereas anisotropic CBO (solid lines) converges independently from d with rate $(2\lambda - \sigma^2)$.

isotropic CBO [2, 9, 18], making it a particularly interesting choice for problems in high-dimensional spaces, e.g., from signal processing and machine learning applications. In this work we extend the analysis of [9] from isotropic CBO to CBO with anisotropic diffusion. More precisely, we show global convergence of the anisotropic CBO dynamics in mean-field law to the global minimizer v^* under minimal assumptions about the initial measure ρ_0 and for a rich class of objectives \mathcal{E} . Furthermore, utilizing some tweaks in the implementation of anisotropic CBO, such as a random mini-batch idea and a cooling strategy of the parameters as proposed in [4, 10], we show that CBO performs well, in fact, almost on par with state-of-the-art gradient-based methods, on a long-studied machine learning benchmark in 2000 dimensions, despite using just 100 particles and no gradient information. This encourages the use and further investigation of CBO as a training algorithm for challenging machine learning tasks.

Organization. In Section 2 we first recall details about the well-posedness of the mean-field dynamics (2) in the case of anisotropic diffusion before we present the main theoretical result about the convergence of anisotropic CBO in mean-field law. The proof follows in Section 3. Section 4 illustrates the practicability of the method on a benchmark problem and Section 5 concludes the paper.

For the sake of reproducible research, in the GitHub repository <https://github.com/KonstantinRiedl/CBOGlobalConvergenceAnalysis> we provide the Matlab code implementing the CBO algorithm used in this work.

Notation. $B_r^\infty(u)$ is a closed ℓ^∞ ball in \mathbb{R}^d with center u and radius r . For the space of continuous functions $f : X \rightarrow Y$ we write $\mathcal{C}(X, Y)$, with $X \subset \mathbb{R}^n$, $n \in \mathbb{N}$, and a suitable topological space Y . For $X \subset \mathbb{R}^n$ open and for $Y = \mathbb{R}^m$, $m \in \mathbb{N}$, the function space $\mathcal{C}_c^k(X, Y)$ contains functions $f \in \mathcal{C}(X, Y)$ that are k -times continuously differentiable and compactly supported. Y is omitted if $Y = \mathbb{R}$.

The objects of study are laws of stochastic processes, $\rho \in \mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))$, where $\mathcal{P}(\mathbb{R}^d)$ contains all Borel probability measures over \mathbb{R}^d . $\rho_t \in \mathcal{P}(\mathbb{R}^d)$ is a snapshot of such law at time t and ϱ some fixed distribution. Measures $\varrho \in \mathcal{P}(\mathbb{R}^d)$ with finite p -th moment are collected in $\mathcal{P}_p(\mathbb{R}^d)$. For any $1 \leq p < \infty$, W_p denotes the Wasserstein- p distance. $\mathbb{E}(\varrho)$ is the expectation of a probability measure ϱ .

2 Global Convergence in Mean-Field Law

In this section we first recite a well-posedness result about the Fokker-Planck equation (2) and then present the main result about global convergence.

2.1 Definition of Weak Solutions and Well-Posedness

We begin by defining weak solutions of the Fokker-Planck equation (2).

Definition 1. Let $\rho_0 \in \mathcal{P}(\mathbb{R}^d)$, $T > 0$. We say $\rho \in \mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))$ satisfies the Fokker-Planck equation (2) with initial condition ρ_0 in the weak sense in the time interval $[0, T]$, if we have for all $\phi \in \mathcal{C}_c^\infty(\mathbb{R}^d)$ and all $t \in (0, T)$

$$\begin{aligned} \frac{d}{dt} \int \phi(v) d\rho_t(v) &= -\lambda \int \sum_{k=1}^d (v - v_\alpha(\rho_t))_k \partial_k \phi(v) d\rho_t(v) \\ &\quad + \frac{\sigma^2}{2} \int \sum_{k=1}^d D(v - v_\alpha(\rho_t))_{kk}^2 \partial_{kk}^2 \phi(v) d\rho_t(v) \end{aligned} \quad (4)$$

and $\lim_{t \rightarrow 0} \rho_t = \rho_0$ pointwise.

In what follows the case of CBO with anisotropic diffusion is considered, i.e., $D(v - v_\alpha(\cdot)) = \text{diag}(v - v_\alpha(\cdot))$ in Equations (1), (2) and (4).

Analogously to the well-posedness results [2, Theorems 3.1, 3.2] for CBO with isotropic diffusion, we can obtain well-posedness of (2) for anisotropic CBO.

Theorem 1. Let $T > 0$, $\rho_0 \in \mathcal{P}_4(\mathbb{R}^d)$ and consider $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\underline{\mathcal{E}} := \mathcal{E}(v^*) > -\infty$, which, for some constants $C_1, C_2 > 0$, satisfies

$$\begin{aligned} |\mathcal{E}(v) - \mathcal{E}(w)| &\leq C_1(\|v\|_2 + \|w\|_2) \|v - w\|_2, \quad \text{for all } v, w \in \mathbb{R}^d, \\ \mathcal{E}(v) - \underline{\mathcal{E}} &\leq C_2(1 + \|v\|_2^2), \quad \text{for all } v \in \mathbb{R}^d. \end{aligned}$$

If in addition, either $\sup_{v \in \mathbb{R}^d} \mathcal{E}(v) < \infty$, or, for some $C_3, C_4 > 0$, \mathcal{E} satisfies

$$\mathcal{E}(v) - \underline{\mathcal{E}} \geq C_3 \|v\|_2^2, \quad \text{for all } \|v\|_2 \geq C_4,$$

then there exists a law $\rho \in \mathcal{C}([0, T], \mathcal{P}_4(\mathbb{R}^d))$ weakly satisfying Equation (2).

Proof. The proof is based on the Leray-Schauder fixed point theorem and uses the same arguments as the ones provided for [2, Theorems 3.1, 3.2].

Remark 1. As discussed in [9, Remark 7], the proof of Theorem 1 justifies an extension of the test function space $\mathcal{C}_c^\infty(\mathbb{R}^d)$ in Definition 1 to

$$\mathcal{C}_*^2(\mathbb{R}^d) := \left\{ \phi \in \mathcal{C}^2(\mathbb{R}^d) : |\partial_k \phi(v)| \leq c(1 + |v_k|) \text{ and } \|\partial_{kk}^2 \phi\|_\infty < \infty \right. \\ \left. \text{for all } k \in \{1, \dots, d\} \text{ and some constant } c > 0 \right\}.$$

2.2 Main Results

We now present the main result about global convergence in mean-field law for objective functions that satisfy the following conditions.

Definition 2 (Assumptions). *We consider functions $\mathcal{E} \in \mathcal{C}(\mathbb{R}^d)$, for which*

- A1 *there exists $v^* \in \mathbb{R}^d$ such that $\mathcal{E}(v^*) = \inf_{v \in \mathbb{R}^d} \mathcal{E}(v) =: \underline{\mathcal{E}}$, and*
- A2 *there exist $\mathcal{E}_\infty, R_0, \eta > 0$, and $\nu \in (0, \infty)$ such that*

$$\|v - v^*\|_\infty \leq \frac{1}{\eta} (\mathcal{E}(v) - \underline{\mathcal{E}})^\nu \quad \text{for all } v \in B_{R_0}^\infty(v^*), \quad (5)$$

$$\mathcal{E}_\infty < \mathcal{E}(v) - \underline{\mathcal{E}} \quad \text{for all } v \in (B_{R_0}^\infty(v^*))^c. \quad (6)$$

Assumption A2 can be regarded as a tractability condition of the energy landscape around the minimizer and in the farfield. Equation (5) requires the local coercivity of \mathcal{E} , whereas (6) prevents that $\mathcal{E}(v) \approx \underline{\mathcal{E}}$ far away from v^* .

Definition 2 covers a wide range of function classes, including for instance the Rastrigin function, see Figure 1a, and objectives related to various machine learning tasks, see, e.g., [10].

Theorem 2. *Let \mathcal{E} be as in Definition 2. Moreover, let $\rho_0 \in \mathcal{P}_4(\mathbb{R}^d)$ be such that*

$$\rho_0(B_r^\infty(v^*)) > 0 \quad \text{for all } r > 0. \quad (7)$$

Define $\mathcal{V}(\rho_t) := 1/2 \int \|v - v^\|_2^2 d\rho_t(v)$. Fix any $\varepsilon \in (0, \mathcal{V}(\rho_0))$ and $\tau \in (0, 1)$, parameters $\lambda, \sigma > 0$ with $2\lambda > \sigma^2$, and the time horizon*

$$T^* := \frac{1}{(1 - \tau)(2\lambda - \sigma^2)} \log \left(\frac{\mathcal{V}(\rho_0)}{\varepsilon} \right). \quad (8)$$

Then there exists $\alpha_0 > 0$, which depends (among problem dependent quantities) on ε and τ , such that for all $\alpha > \alpha_0$, if $\rho \in \mathcal{C}([0, T^], \mathcal{P}_4(\mathbb{R}^d))$ is a weak solution to the Fokker-Planck equation (2) on the time interval $[0, T^*]$ with initial condition ρ_0 , we have $\min_{t \in [0, T^*]} \mathcal{V}(\rho_t) \leq \varepsilon$. Furthermore, until $\mathcal{V}(\rho_t)$ reaches the prescribed accuracy ε , we have the exponential decay*

$$\mathcal{V}(\rho_t) \leq \mathcal{V}(\rho_0) \exp(-(1 - \tau)(2\lambda - \sigma^2)t)$$

and, up to a constant, the same behavior for $W_2^2(\rho_t, \delta_{v^})$.*

The rate of convergence $(2\lambda - \sigma^2)$ obtained in Theorem 2 is confirmed numerically by the experiments depicted in Figure 1. We emphasize the dimension-independent convergence of CBO with anisotropic diffusion, contrasting the dimension-dependent rate $(2\lambda - d\sigma^2)$ of isotropic CBO, cf. [9, Theorem 12].

Refining the argument of the proof of Theorem 2 allows to show that the time T , where $\mathcal{V}(\rho_T) = \varepsilon$ is achieved, satisfies $T \in [\frac{1-\tau}{(1+\tau/2)} T^*, T^*]$. For the technical details we refer to the proof of [9, Theorem 12], where an analogous statement is shown to hold in the setting of isotropic noise. A proof for the herein investigated anisotropic setting is provided in full detail in the dissertation of the third author.

3 Proof of Theorem 2

This section provides the proof details for Theorem 2, starting with a sketch in Section 3.1. Sections 3.2–3.4 present statements, which are needed in the proof and may be of independent interest. Section 3.5 completes the proof.

Remark 2. Without loss of generality we assume $\underline{\mathcal{E}} = 0$ throughout the proof.

3.1 Proof Sketch

The main idea is to show that $\mathcal{V}(\rho_t)$ satisfies the differential inequality

$$\frac{d}{dt}\mathcal{V}(\rho_t) \leq -(1 - \tau)(2\lambda - \sigma^2)\mathcal{V}(\rho_t) \quad (9)$$

until $\mathcal{V}(\rho_T) \leq \varepsilon$. The first step towards (9) is to derive a differential inequality for $\mathcal{V}(\rho_t)$ using the dynamics of ρ , which is done in Lemma 1. In order to control the appearing quantity $\|v_\alpha(\rho_t) - v^*\|_2$, we establish a quantitative Laplace principle. Namely, under the inverse continuity property A2, Proposition 1 shows

$$\|v_\alpha(\rho_t) - v^*\|_2 \lesssim \ell(r) + \frac{\sqrt{d}\exp(-\alpha r)}{\rho_t(B_r^\infty(v^*))}, \quad \text{for sufficiently small } r > 0,$$

where ℓ is decreasing with $\ell(r) \rightarrow 0^+$ as $r \rightarrow 0$. Thus, $\|v_\alpha(\rho_t) - v^*\|_2$ can be made arbitrarily small by suitable choices of $r \ll 1$ and $\alpha \gg 1$, as long as we can guarantee $\rho_t(B_r^\infty(v^*)) > 0$ for all $r > 0$ and at all times $t \in [0, T]$. The latter requires non-zero initial mass $\rho_0(B_r^\infty(v^*))$ as well as an active Brownian motion, as made rigorous in Proposition 2.

3.2 Evolution of the Mean-Field Limit

We now derive the evolution inequality of the energy functional $\mathcal{V}(\rho_t)$.

Lemma 1. *Let $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}$, and fix $\alpha, \lambda, \sigma > 0$. Moreover, let $T > 0$ and let $\rho \in \mathcal{C}([0, T], \mathcal{P}_4(\mathbb{R}^d))$ be a weak solution to Equation (2). Then $\mathcal{V}(\rho_t)$ satisfies*

$$\begin{aligned} \frac{d}{dt} \mathcal{V}(\rho_t) &\leq - (2\lambda - \sigma^2) \mathcal{V}(\rho_t) + \sqrt{2} (\lambda + \sigma^2) \sqrt{\mathcal{V}(\rho_t)} \|v_\alpha(\rho_t) - v^*\|_2 \\ &\quad + \frac{\sigma^2}{2} \|v_\alpha(\rho_t) - v^*\|_2^2. \end{aligned}$$

Proof. Noting that $\phi(v) = 1/2 \|v - v^*\|_2^2$ is in $\mathcal{C}_*^2(\mathbb{R}^d)$ and recalling that ρ satisfies the identity (4) for all test functions in $\mathcal{C}_*^2(\mathbb{R}^d)$, see Remark 1, we obtain

$$\frac{d}{dt} \mathcal{V}(\rho_t) = -\lambda \int \langle v - v^*, v - v_\alpha(\rho_t) \rangle d\rho_t(v) + \frac{\sigma^2}{2} \int \|v - v_\alpha(\rho_t)\|_2^2 d\rho_t(v),$$

where we used $\partial_k \phi(v) = (v - v^*)_k$ and $\partial_{kk}^2 \phi(v) = 1$ for all $k \in \{1, \dots, d\}$. Following the steps taken in [9, Lemma 17] yields the statement. \square

3.3 Quantitative Laplace Principle

The Laplace principle asserts that $-\log(\|\omega_\alpha\|_{L^1(\varrho)})/\alpha \rightarrow \mathcal{E}$ as $\alpha \rightarrow \infty$ as long as the global minimizer v^* is in the support of ϱ . Under the assumption of the inverse continuity property this can be used to qualitatively characterize the proximity of $v_\alpha(\varrho)$ to the global minimizer v^* . However, as it neither allows to quantify this proximity nor gives a suggestion on how to choose α to reach a certain approximation quality, we introduced a quantitative version in [9, Proposition 21], which we now adapt suitably to satisfy the anisotropic setting.

Proposition 1. *Let $\mathcal{E} = 0$, $\varrho \in \mathcal{P}(\mathbb{R}^d)$ and fix $\alpha > 0$. For any $r > 0$ we define $\mathcal{E}_r := \sup_{v \in B_r^\infty(v^*)} \mathcal{E}(v)$. Then, under the inverse continuity property A2, for any $r \in (0, R_0]$ and $q > 0$ such that $q + \mathcal{E}_r \leq \mathcal{E}_\infty$, we have*

$$\|v_\alpha(\varrho) - v^*\|_2 \leq \frac{\sqrt{d}(q + \mathcal{E}_r)^\nu}{\eta} + \frac{\sqrt{d} \exp(-\alpha q)}{\varrho(B_r^\infty(v^*))} \int \|v - v^*\|_2 d\varrho(v).$$

Proof. Following the lines of the proof of [9, Proposition 22] but replacing all ℓ^2 balls and norms by ℓ^∞ balls and norms, respectively, we obtain

$$\|v_\alpha(\varrho) - v^*\|_\infty \leq \frac{(q + \mathcal{E}_r)^\nu}{\eta} + \frac{\exp(-\alpha q)}{\varrho(B_r^\infty(v^*))} \int \|v - v^*\|_\infty d\varrho(v).$$

The statement now follows noting that $\|\cdot\|_\infty \leq \|\cdot\|_2 \leq \sqrt{d} \|\cdot\|_\infty$. \square

3.4 A Lower Bound for the Probability Mass around v^*

In this section we provide a lower bound for the probability mass of $\rho_t(B_r^\infty(v^*))$, where $r > 0$ is a small radius. This is achieved by defining a mollifier ϕ_r so that $\rho_t(B_r^\infty(v^*)) \geq \int \phi_r(v) d\rho_t(v)$ and studying the evolution of the right-hand side.

Lemma 2. For $r > 0$ we define the mollifier $\phi_r : \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$\phi_r(v) := \begin{cases} \prod_{k=1}^d \exp\left(1 - \frac{r^2}{r^2 - (v - v^*)_k^2}\right), & \text{if } \|v - v^*\|_\infty < r, \\ 0, & \text{else.} \end{cases} \quad (10)$$

We have $\text{Im}(\phi_r) = [0, 1]$, $\text{supp}(\phi_r) = B_r^\infty(v^*)$, $\phi_r \in \mathcal{C}_c^\infty(\mathbb{R}^d)$ and

$$\begin{aligned} \partial_k \phi_r(v) &= -2r^2 \frac{(v - v^*)_k}{\left(r^2 - (v - v^*)_k^2\right)^2} \phi_r(v), \\ \partial_{kk}^2 \phi_r(v) &= 2r^2 \left(\frac{2 \left(2(v - v^*)_k^2 - r^2\right) (v - v^*)_k^2 - \left(r^2 - (v - v^*)_k^2\right)^2}{\left(r^2 - (v - v^*)_k^2\right)^4} \right) \phi_r(v). \end{aligned}$$

Proof. ϕ_r is a tensor product of classical well-studied mollifiers. \square

Proposition 2. Let $T > 0$, $r > 0$, and fix parameters $\alpha, \lambda, \sigma > 0$. Assume $\rho \in \mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))$ weakly solves the Fokker-Planck equation (2) in the sense of Definition 1 with initial condition $\rho_0 \in \mathcal{P}(\mathbb{R}^d)$ and for $t \in [0, T]$. Then, for all $t \in [0, T]$ we have

$$\rho_t(B_r^\infty(v^*)) \geq \left(\int \phi_r(v) d\rho_0(v) \right) \exp(-pt) \quad (11)$$

with

$$p := 2d \max \left\{ \frac{\lambda(cr + B\sqrt{c})}{(1-c)^2 r} + \frac{\sigma^2(cr^2 + B^2)(2c+1)}{(1-c)^4 r^2}, \frac{2\lambda^2}{(2c-1)\sigma^2} \right\}, \quad (12)$$

for any $B < \infty$ with $\sup_{t \in [0, T]} \|v_\alpha(\rho_t) - v^*\|_\infty \leq B$ and for any $c \in (1/2, 1)$ satisfying $(1-c)^2 \leq (2c-1)c$.

Remark 3. In order to ensure a finite decay rate $p < \infty$ in Proposition 2 it is crucial to have a non-vanishing diffusion $\sigma > 0$.

Proof (Proposition 2). By the properties of the mollifier in Lemma 2 we have

$$\rho_t(B_r^\infty(v^*)) \geq \int \phi_r(v) d\rho_t(v).$$

Our strategy is to derive a lower bound for the right-hand side of this inequality. Using the weak solution property of ρ and the fact that $\phi_r \in \mathcal{C}_c^\infty(\mathbb{R}^d)$, we obtain

$$\frac{d}{dt} \int \phi_r(v) d\rho_t(v) = \sum_{k=1}^d \int (T_{1k}(v) + T_{2k}(v)) d\rho_t(v) \quad (13)$$

$$\text{with } T_{1k}(v) := -\lambda (v - v_\alpha(\rho_t))_k \partial_k \phi_r(v)$$

$$\text{and } T_{2k}(v) := \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 \partial_{kk}^2 \phi_r(v)$$

for $k \in \{1, \dots, d\}$. We now aim for showing $T_{1k}(v) + T_{2k}(v) \geq -p\phi_r(v)$ uniformly on \mathbb{R}^d individually for each k and for $p > 0$ as in the statement. Since the mollifier ϕ_r and its derivatives vanish outside of $\Omega_r := \{v \in \mathbb{R}^d : \|v - v^*\|_\infty < r\}$ we restrict our attention to the open ℓ^∞ -ball Ω_r . To achieve the lower bound over Ω_r , we introduce for each $k \in \{1, \dots, d\}$ the subsets

$$K_{1k} := \{v \in \mathbb{R}^d : |(v - v^*)_k| > \sqrt{c}r\} \quad (14)$$

and

$$K_{2k} := \left\{ v \in \mathbb{R}^d : -\lambda(v - v_\alpha(\rho_t))_k (v - v^*)_k \left(r^2 - (v - v^*)_k^2 \right)^2 \right. \\ \left. > \tilde{c}r^2 \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (v - v^*)_k^2 \right\}, \quad (15)$$

where $\tilde{c} := 2c - 1 \in (0, 1)$. For fixed k we now decompose Ω_r according to

$$\Omega_r = (K_{1k}^c \cap \Omega_r) \cup (K_{1k} \cap K_{2k}^c \cap \Omega_r) \cup (K_{1k} \cap K_{2k} \cap \Omega_r), \quad (16)$$

which is illustrated in Figure 2 for different positions of $v_\alpha(\rho_t)$ and values of σ .

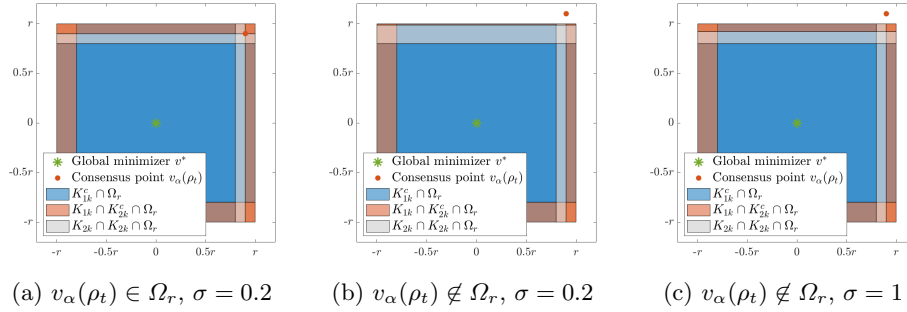


Fig. 2: Visualization of the decomposition of Ω_r as in (16) for different positions of $v_\alpha(\rho_t)$ and values of σ .

In the following we treat each of these three subsets separately.

Subset $K_{1k}^c \cap \Omega_r$: We have $|(v - v^*)_k| \leq \sqrt{c}r$ for each $v \in K_{1k}^c$, which can be used to independently derive lower bounds for both terms T_{1k} and T_{2k} . For T_{1k} , we insert the expression for $\partial_k \phi_r$ from Lemma 2 to get

$$T_{1k}(v) = 2r^2 \lambda (v - v_\alpha(\rho_t))_k \frac{(v - v^*)_k}{\left(r^2 - (v - v^*)_k^2 \right)^2} \phi_r(v) \\ \geq -2r^2 \lambda \frac{|(v - v_\alpha(\rho_t))_k| |(v - v^*)_k|}{\left(r^2 - (v - v^*)_k^2 \right)^2} \phi_r(v) \geq -\frac{2\lambda(\sqrt{c}r + B)\sqrt{c}}{(1-c)^2 r} \phi_r(v) \\ =: -p_1 \phi_r(v),$$

where $|(v - v_\alpha(\rho_t))_k| \leq |(v - v^*)_k| + |(v^* - v_\alpha(\rho_t))_k| \leq \sqrt{cr} + B$ is used in the last inequality. For T_2 we insert the expression for $\partial_{kk}^2 \phi_r$ from Lemma 2 to obtain

$$\begin{aligned} T_{2k}(v) &= \sigma^2 r^2 (v - v_\alpha(\rho_t))_k^2 \frac{2 \left(2(v - v^*)_k^2 - r^2 \right) (v - v^*)_k^2 - \left(r^2 - (v - v^*)_k^2 \right)^2}{\left(r^2 - (v - v^*)_k^2 \right)^4} \phi_r(v) \\ &\geq -\frac{2\sigma^2(cr^2 + B^2)(2c+1)}{(1-c)^4 r^2} \phi_r(v) =: -p_2 \phi_r(v), \end{aligned}$$

where the last inequality uses $(v - v_\alpha(\rho_t))_k^2 \leq (\sqrt{cr} + B)^2 \leq 2(cr^2 + B^2)$.

Subset $K_{1k} \cap K_{2k}^c \cap \Omega_r$: As $v \in K_{1k}$ we have $|(v - v^*)_k|_2 > \sqrt{cr}$. We observe that $T_{1k}(v) + T_{2k}(v) \geq 0$ for all v in this subset whenever

$$\begin{aligned} &\left(-\lambda (v - v_\alpha(\rho_t))_k (v - v^*)_k + \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 \right) \left(r^2 - (v - v^*)_k^2 \right)^2 \\ &\leq \sigma^2 (v - v_\alpha(\rho_t))_k^2 \left(2(v - v^*)_k^2 - r^2 \right) (v - v^*)_k^2. \end{aligned} \quad (17)$$

The first term on the left-hand side in (17) can be bounded from above exploiting that $v \in K_{2k}^c$ and by using the relation $\tilde{c} = 2c - 1$. More precisely, we have

$$\begin{aligned} &-\lambda (v - v_\alpha(\rho_t))_k (v - v^*)_k \left(r^2 - (v - v^*)_k^2 \right)^2 \leq \tilde{c} r^2 \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (v - v^*)_k^2 \\ &= (2c-1) r^2 \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (v - v^*)_k^2 \leq \left(2(v - v^*)_k^2 - r^2 \right) \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (v - v^*)_k^2, \end{aligned}$$

where the last inequality follows since $v \in K_{1k}$. For the second term on the left-hand side in (17) we can use $(1-c)^2 \leq (2c-1)c$ as per assumption, to get

$$\begin{aligned} &\frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 \left(r^2 - (v - v^*)_k^2 \right)^2 \leq \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (1-c)^2 r^4 \\ &\leq \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 (2c-1) r^2 c r^2 \leq \frac{\sigma^2}{2} (v - v_\alpha(\rho_t))_k^2 \left(2(v - v^*)_k^2 - r^2 \right) (v - v^*)_k^2. \end{aligned}$$

Hence, (17) holds and we have $T_{1k}(v) + T_{2k}(v) \geq 0$ uniformly on this subset.

Subset $K_{1k} \cap K_{2k} \cap \Omega_r$: As $v \in K_{1k}$ we have $|(v - v^*)_k|_2 > \sqrt{cr}$. We first note that $T_{1k}(v) = 0$ whenever $\sigma^2 (v - v_\alpha(\rho_t))_k^2 = 0$, provided $\sigma > 0$, so nothing needs to be done if $v_k = (v_\alpha(\rho_t))_k$. Otherwise, if $\sigma^2 (v - v_\alpha(\rho_t))_k^2 > 0$, we exploit $v \in K_{2k}$ to get

$$\begin{aligned} &\frac{(v - v_\alpha(\rho_t))_k (v - v^*)_k}{\left(r^2 - (v - v^*)_k^2 \right)^2} \geq \frac{-|(v - v_\alpha(\rho_t))_k| |(v - v^*)_k|}{\left(r^2 - (v - v^*)_k^2 \right)^2} \\ &> \frac{2\lambda (v - v_\alpha(\rho_t))_k (v - v^*)_k}{\tilde{c} r^2 \sigma^2 |(v - v_\alpha(\rho_t))_k| |(v - v^*)_k|} \geq -\frac{2\lambda}{\tilde{c} r^2 \sigma^2}. \end{aligned}$$

Using this, T_{1k} can be bounded from below by

$$T_{1k}(v) = 2r^2 \lambda (v - v_\alpha(\rho_t))_k \frac{(v - v^*)_k}{\left(r^2 - (v - v^*)_k\right)^2} \phi_r(v) \geq -\frac{4\lambda^2}{\tilde{c}\sigma^2} \phi_r(v) =: -p_3 \phi_r(v).$$

For T_{2k} , the nonnegativity of $\sigma^2 (v - v_\alpha(\rho_t))_k^2$ implies $T_{2k}(v) \geq 0$, whenever

$$2 \left(2(v - v^*)_k^2 - r^2\right) (v - v^*)_k \geq \left(r^2 - (v - v^*)_k\right)^2.$$

This holds for $v \in K_{1k}$, if $2(2c - 1)c \geq (1 - c)^2$ as implied by the assumption.

Concluding the proof: Using the evolution of ϕ_r as in (13) and the individual decompositions of Ω_r for the terms $T_{1k} + T_{2k}$, we now get

$$\begin{aligned} \frac{d}{dt} \int \phi_r(v) d\rho_t(v) &= \sum_{k=1}^d \left(\int_{K_{1k} \cap K_{2k}^c \cap \Omega_r} \underbrace{(T_{1k}(v) + T_{2k}(v))}_{\geq 0} d\rho_t(v) \right. \\ &\quad \left. + \int_{K_{1k} \cap K_{2k} \cap \Omega_r} \underbrace{(T_{1k}(v) + T_{2k}(v))}_{\geq -p_3 \phi_r(v)} d\rho_t(v) + \int_{K_{1k}^c \cap \Omega_r} \underbrace{(T_{1k}(v) + T_{2k}(v))}_{\geq -(p_1 + p_2) \phi_r(v)} d\rho_t(v) \right) \\ &\geq -d \max\{p_1 + p_2, p_3\} \int \phi_r(v) d\rho_t(v) = -p \int \phi_r(v) d\rho_t(v). \end{aligned}$$

An application of Grönwall's inequality concludes the proof. \square

3.5 Proof of Theorem 2

We now have all necessary tools to conclude the global convergence proof.

Proof (Theorem 2). Let us first choose the parameter α such that

$$\begin{aligned} \alpha > \alpha_0 := \frac{1}{q_\varepsilon} \left(\log \left(\frac{2^{d+1} \sqrt{2d\mathcal{V}(\rho_0)}}{c(\tau, \lambda, \sigma) \sqrt{\varepsilon}} \right) + \frac{p}{(1 - \tau)(2\lambda - \sigma^2)} \log \left(\frac{\mathcal{V}(\rho_0)}{\varepsilon} \right) \right. \\ \left. - \log \rho_0(B_{r_\varepsilon/2}^\infty(v^*)) \right), \end{aligned} \quad (18)$$

where we introduce the definitions

$$c(\tau, \lambda, \sigma) := \min \left\{ \frac{\tau}{2} \frac{(2\lambda - \sigma^2)}{\sqrt{2}(\lambda + \sigma^2)}, \sqrt{\tau \frac{(2\lambda - \sigma^2)}{\sigma^2}} \right\} \quad (19)$$

as well as

$$q_\varepsilon := \frac{1}{2} \min \left\{ \left(\eta \frac{c(\tau, \lambda, \sigma) \sqrt{\varepsilon}}{2\sqrt{d}} \right)^{1/\nu}, \mathcal{E}_\infty \right\} \text{ and } r_\varepsilon := \max_{s \in [0, R_0]} \left\{ \max_{v \in B_s^\infty(v^*)} \mathcal{E}(v) \leq q_\varepsilon \right\}.$$

Moreover, p is as defined in (12) in Proposition 2 with $B = c(\tau, \lambda, \sigma)\sqrt{\mathcal{V}(\rho_0)}$ and with $r = r_\varepsilon$. By construction, $q_\varepsilon > 0$ and $r_\varepsilon \leq R_0$. Moreover, recalling the notation $\mathcal{E}_r = \sup_{v \in B_{r_\varepsilon}^\infty(v^*)} \mathcal{E}(v)$ from Proposition 1, we have $q_\varepsilon + \mathcal{E}_{r_\varepsilon} \leq 2q_\varepsilon \leq \mathcal{E}_\infty$ by definition of r_ε . Furthermore, since $q_\varepsilon > 0$, the continuity of \mathcal{E} ensures that there exists $s_{q_\varepsilon} > 0$ such that $\mathcal{E}(v) \leq q_\varepsilon$ for all $v \in B_{s_{q_\varepsilon}}^\infty(v^*)$, yielding also $r_\varepsilon > 0$.

Let us now define the time horizon $T_\alpha \geq 0$ by

$$T_\alpha := \sup \{t \geq 0 : \mathcal{V}(\rho_{t'}) > \varepsilon \text{ and } \|v_\alpha(\rho_{t'}) - v^*\|_2 < C(t') \text{ for all } t' \in [0, t]\} \quad (20)$$

with $C(t) := c(\tau, \lambda, \sigma)\sqrt{\mathcal{V}(\rho_t)}$. Notice for later use that $C(0) = B$. Our aim is to show that $\min_{t \in [0, T_\alpha]} \mathcal{V}(\rho_t) \leq \varepsilon$ with $T_\alpha \leq T^*$ and that we have at least exponential decay until $\mathcal{V}(\rho_t)$ reaches the prescribed accuracy ε .

First, however, let us ensure that $T_\alpha > 0$, which follows from the continuity of the mappings $t \mapsto \mathcal{V}(\rho_t)$ and $t \mapsto \|v_\alpha(\rho_t) - v^*\|_2$ since $\mathcal{V}(\rho_0) > \varepsilon$ and $\|v_\alpha(\rho_0) - v^*\|_2 < C(0)$. While the former holds by assumption, the latter follows by applying Proposition 1 with q_ε and r_ε as defined before, which gives

$$\begin{aligned} \|v_\alpha(\rho_0) - v^*\|_2 &\leq \frac{\sqrt{d}(q_\varepsilon + \mathcal{E}_{r_\varepsilon})^\nu}{\eta} + \frac{\sqrt{d} \exp(-\alpha q_\varepsilon)}{\rho_0(B_{r_\varepsilon}^\infty(v^*))} \int \|v - v^*\|_2 d\rho_0(v) \\ &\leq \frac{c(\tau, \lambda, \sigma) \sqrt{\varepsilon}}{2} + \frac{\sqrt{d} \exp(-\alpha q_\varepsilon)}{\rho_0(B_{r_\varepsilon}^\infty(v^*))} \sqrt{2\mathcal{V}(\rho_0)} \\ &\leq c(\tau, \lambda, \sigma) \sqrt{\varepsilon} < c(\tau, \lambda, \sigma) \sqrt{\mathcal{V}(\rho_0)} = C(0), \end{aligned}$$

where the first inequality in the last line holds by the choice of α in (18).

Let us now show that $\mathcal{V}(\rho_t)$ decays at least exponentially fast up to time T_α . Lemma 1 provides an upper bound for the time derivative of $\mathcal{V}(\rho_t)$, given by

$$\begin{aligned} \frac{d}{dt} \mathcal{V}(\rho_t) &\leq -(2\lambda - \sigma^2) \mathcal{V}(\rho_t) + \sqrt{2} (\lambda + \sigma^2) \sqrt{\mathcal{V}(\rho_t)} \|v_\alpha(\rho_t) - v^*\|_2 \\ &\quad + \frac{\sigma^2}{2} \|v_\alpha(\rho_t) - v^*\|_2^2. \end{aligned} \quad (21)$$

Combining this with the definition of T_α in (20) we have by construction

$$\frac{d}{dt} \mathcal{V}(\rho_t) \leq -(1 - \tau) (2\lambda - \sigma^2) \mathcal{V}(\rho_t) \quad \text{for all } t \in (0, T_\alpha).$$

Grönwall's inequality implies the upper bound

$$\mathcal{V}(\rho_t) \leq \mathcal{V}(\rho_0) \exp(-(1 - \tau) (2\lambda - \sigma^2) t) \quad \text{for all } t \in [0, T_\alpha]. \quad (22)$$

Accordingly, we note that $\mathcal{V}(\rho_t)$ is decreasing in t , which implies the decay of the function $C(t)$ as well. Hence, recalling the definition of T_α , we may bound

$$\max_{t \in [0, T_\alpha]} \|v_\alpha(\rho_t) - v^*\|_2 \leq \max_{t \in [0, T_\alpha]} C(t) \leq C(0). \quad (23)$$

We now conclude by showing $\min_{t \in [0, T_\alpha]} \mathcal{V}(\rho_t) \leq \varepsilon$ with $T_\alpha \leq T^*$. For this we distinguish the following three cases.

Case $T_\alpha \geq T^*$: If $T_\alpha \geq T^*$, we can use the definition of T^* in (8) and the time-evolution bound of $\mathcal{V}(\rho_t)$ in (22) to conclude that $\mathcal{V}(\rho_{T^*}) \leq \varepsilon$. Hence, by definition of T_α in (20), we find $\mathcal{V}(\rho_{T_\alpha}) = \varepsilon$ and $T_\alpha = T^*$.

Case $T_\alpha < T^*$ and $\mathcal{V}(\rho_{T_\alpha}) \leq \varepsilon$: Nothing needs to be discussed in this case.

Case $T_\alpha < T^*$ and $\mathcal{V}(\rho_{T_\alpha}) > \varepsilon$: We shall show that this case can never occur by verifying that $\|v_\alpha(\rho_{T_\alpha}) - v^*\|_2 < C(T_\alpha)$ due to the choice of α in (18). Namely, by applying again Proposition 1 with q_ε and r_ε as defined before, we get

$$\begin{aligned} \|v_\alpha(\rho_{T_\alpha}) - v^*\|_2 &\leq \frac{\sqrt{d}(q_\varepsilon + \mathcal{E}_{r_\varepsilon})^\nu}{\eta} + \frac{\sqrt{d}\exp(-\alpha q_\varepsilon)}{\rho_{T_\alpha}(B_{r_\varepsilon}^\infty(v^*))} \int \|v - v^*\|_2 d\rho_{T_\alpha}(v) \\ &\leq \frac{c(\tau, \lambda, \sigma)\sqrt{\varepsilon}}{2} + \frac{\sqrt{d}\exp(-\alpha q_\varepsilon)}{\rho_{T_\alpha}(B_{r_\varepsilon}^\infty(v^*))} \sqrt{2\mathcal{V}(\rho_{T_\alpha})} \\ &< \frac{c(\tau, \lambda, \sigma)\sqrt{\mathcal{V}(\rho_{T_\alpha})}}{2} + \frac{\sqrt{d}\exp(-\alpha q_\varepsilon)}{\rho_{T_\alpha}(B_{r_\varepsilon}^\infty(v^*))} \sqrt{2\mathcal{V}(\rho_{T_\alpha})}. \end{aligned} \quad (24)$$

Since, thanks to (23), we have the bound $\max_{t \in [0, T_\alpha]} \|v_\alpha(\rho_t) - v^*\|_2 \leq B$ for $B = C(0)$, which is in particular independent of α , Proposition 2 guarantees that there exists a $p > 0$ independent of α (but dependent on B and r_ε) with

$$\begin{aligned} \rho_{T_\alpha}(B_{r_\varepsilon}^\infty(v^*)) &\geq \left(\int \phi_{r_\varepsilon}(v) d\rho_0(v) \right) \exp(-pT_\alpha) \\ &\geq \frac{1}{2^d} \rho_0(B_{r_\varepsilon/2}^\infty(v^*)) \exp(-pT^*) > 0, \end{aligned}$$

where we used (7) for bounding the initial mass ρ_0 and the fact that ϕ_r (as defined in Equation (10)) is bounded from below on $B_{r/2}^\infty(v^*)$ by $1/2^d$. With this we can continue the chain of inequalities in (24) to obtain

$$\begin{aligned} \|v_\alpha(\rho_{T_\alpha}) - v^*\|_2 &< \frac{c(\tau, \lambda, \sigma)\sqrt{\mathcal{V}(\rho_{T_\alpha})}}{2} + \frac{2^d \sqrt{d}\exp(-\alpha q_\varepsilon)}{\rho_0(B_{r_\varepsilon/2}^\infty(v^*)) \exp(-pT^*)} \sqrt{2\mathcal{V}(\rho_{T_\alpha})} \\ &\leq c(\tau, \lambda, \sigma)\sqrt{\mathcal{V}(\rho_{T_\alpha})} = C(T_\alpha), \end{aligned}$$

where the first inequality in the last line holds by the choice of α in (18). This establishes the desired contradiction, again as consequence of the continuity of the mappings $t \mapsto \mathcal{V}(\rho_t)$ and $t \mapsto \|v_\alpha(\rho_t) - v^*\|_2$. \square

4 A Machine Learning Example

In this section, we showcase the practicability of the implementation of anisotropic CBO as described in [4, Algorithm 2.1] for problems appearing in machine learning by training a shallow and a convolutional NN (CNN) classifier for the MNIST dataset of handwritten digits [16]. Let us emphasize that it is not our aim to challenge the state of the art for this task by employing the most sophisticated model or intricate data preprocessing. We merely believe that this

is a well-understood, complex, high-dimensional benchmark to demonstrate that CBO achieves good results already with limited computational capacities.

Let us now describe the NN architectures used in our numerical experiment, see also Figure 3. Each input image is represented by a matrix of dimension

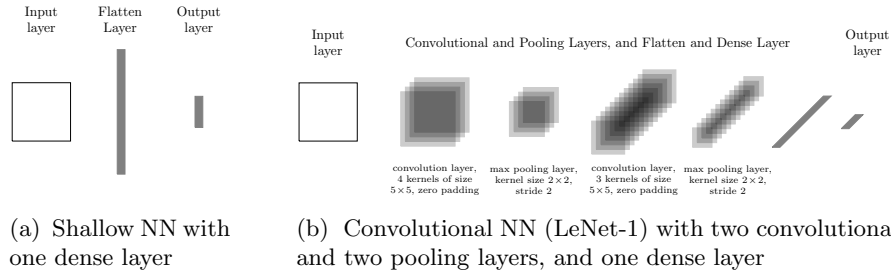


Fig. 3: Architectures of the NNs used in the experiments of Section 4.

28×28 with entries valued between 0 and 1 depending on the grayscale of the respective pixel. For the shallow neural net (see Figure 3a) the image is first reshaped to a vector $x \in \mathbb{R}^{728}$ before being passed through a dense layer of the form $\text{ReLU}(Wx + b)$ with trainable weight matrix $W \in \mathbb{R}^{10 \times 728}$ and bias vector $b \in \mathbb{R}^{10}$. The CNN (see Figure 3b) has learnable kernels and its architecture is similar to the one of the LeNet-1, cf. [15, Section III.C.7]. In both networks a batch normalization step is included after each ReLU activation, which entails a considerably faster training process. Moreover, in the final layers a softmax activation function is applied so that the output can be interpreted as a probability distribution over the digits. In total, the number of unknowns to be trained in case of the shallow NN is 7850, which compares to 2112 free parameters for the CNN. We denote the parameters of the NN by θ and its forward pass by f_θ .

As a loss function during training we use the categorical crossentropy loss $\ell(\hat{y}, y) = -\sum_{k=0}^9 y_k \log(\hat{y}_k)$ with $\hat{y} = f_\theta(x)$ denoting the output of the NN for a training sample (x, y) consisting of image and label. This gives rise to the objective function $\mathcal{E}(\theta) = \frac{1}{M} \sum_{m=1}^M \ell(f_\theta(x^m), y^m)$, where $(x^m, y^m)_{m=1}^M$ denote the M training samples. When evaluating the performance of the NN we determine the accuracy on a test set by counting the number of successful predictions.

The used implementation of anisotropic CBO combines ideas presented in [10, Section 2.2] with the algorithm proposed in [4]. More precisely, it employs random mini-batch ideas when evaluating the objective function \mathcal{E} and when computing the consensus point v_α , meaning that \mathcal{E} is only evaluated on a random subset of size $n_\mathcal{E}$ of the training dataset and v_α is only computed from a random subset of size n_N of all particles. While this reduces the computational complexity, it simultaneously increases the stochasticity, which enhances the ability to escape from local optima. Furthermore, inspired by Simulated Annealing, a cooling strategy for the parameters α and σ is used as well as a variance-based

particle reduction technique similar to ideas from Genetic Algorithms. More specifically, α is multiplied by 2 after each epoch, while the diffusion parameter σ follows the schedule $\sigma_{epoch} = \sigma_0 / \log_2(epoch + 2)$. For our experiments we choose the parameters $\lambda = 1$, $\sigma_0 = \sqrt{0.4}$ and $\alpha_{initial} = 50$, and discrete time step size $\Delta t = 0.1$ for training both the shallow and the convolutional NN. We use $N = 100$ particles, which are initialized according to $\mathcal{N}((0, \dots, 0)^T, \text{Id})$. The mini-batch sizes are $n_{\mathcal{E}} = 60$ and $n_N = 10$ and despite v_α being computed only on a basis of n_N particles, all N particles are updated in every time step, referred to as the full update in [4]. We emphasize that hyperparameters have not been tuned extensively.

In Figure 4 we report the results of our experiment. While achieving a test




Fig. 4: Comparison of the performances of a shallow (dashed lines) and convolutional (solid lines) NN with architectures as described in Figures 3a and 3b, when trained with a discretized version of the anisotropic CBO dynamics (1). Depicted are the accuracies on a test dataset (orange lines) and the values of the objective function \mathcal{E} (blue lines), which was chosen to be the categorical crossentropy loss on a random sample of the training set of size 10000.

accuracy of almost 90% for the shallow NN, we obtain around 97% accuracy with the CNN. For comparison, when trained with backpropagation with finely tuned parameters, a comparable CNN achieves 98.3% accuracy, cf. [15, Figure 9]. In view of these results, CBO can be regarded as a successful optimizer for machine learning tasks, which performs comparably to the state of the art. At the same time it is worth highlighting that CBO is extremely versatile and customizable, does not require gradient information or substantial hyperparameter tuning and has the potential to be parallelized.

5 Conclusion

In this paper we establish the global convergence of anisotropic consensus-based optimization (CBO) to the global minimizer in mean-field law with dimension-independent convergence rate by adapting the proof technique developed in [9].

It is based on the insight that the dynamics of individual particles follow, on average, the gradient flow dynamics of the map $v \mapsto \|v - v^*\|_2^2$. Furthermore, by utilizing the implementation of anisotropic CBO suggested in [4], we demonstrate the practicability of the method by training the well-known LeNet-1 on the MNIST data set, achieving around 97% accuracy after few epochs with just 100 particles.

In subsequent work we plan to extend our theoretical understanding of CBO to the finite particle regime, and aim to provide extensive numerical studies. We also intend to use this approach to explain the mean-field law convergence behavior of other metaheuristics such as Particle Swarm Optimization, see, e.g., [5, 13].

Acknowledgements. MF acknowledges the support of the DFG Project “Identification of Energies from Observations of Evolutions” and the DFG SPP 1962 “Non-smooth and Complementarity-Based Distributed Parameter Systems: Simulation and Hierarchical Optimization”. KR acknowledges the financial support from the Technical University of Munich – Institute for Ethics in Artificial Intelligence (IEAI).

References

1. Aarts, E., Korst, J.: Simulated annealing and Boltzmann machines. A stochastic approach to combinatorial optimization and neural computing. Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Ltd. (1989)
2. Carrillo, J.A., Choi, Y.P., Totzeck, C., Tse, O.: An analytical framework for consensus-based global optimization method. *Math. Models Methods Appl. Sci.* **28**(6), 1037–1066 (2018)
3. Carrillo, J.A., Fornasier, M., Toscani, G., Vecil, F.: Particle, kinetic, and hydrodynamic models of swarming. In: *Mathematical modeling of collective behavior in socio-economic and life sciences*, pp. 297–336. *Model. Simul. Sci. Eng. Technol.*, Birkhäuser Boston, Boston, MA (2010)
4. Carrillo, J.A., Jin, S., Li, L., Zhu, Y.: A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM Control Optim. Calc. Var.* **27**(suppl.), Paper No. S5, 22 (2021)
5. Cipriani, C., Huang, H., Qiu, J.: Zero-inertia limit: from particle swarm optimization to consensus based optimization. [arXiv:2104.06939](https://arxiv.org/abs/2104.06939) (2021)
6. Dorigo, M., Blum, C.: Ant colony optimization theory: a survey. *Theoret. Comput. Sci.* **344**(2-3), 243–278 (2005)
7. Fogel, D.B.: *Evolutionary computation. Toward a new philosophy of machine intelligence*. IEEE Press, Piscataway, NJ, second edn. (2000)
8. Fornasier, M., Huang, H., Pareschi, L., Sünnen, P.: Anisotropic diffusion in consensus-based optimization on the sphere. [arXiv:2104.00420](https://arxiv.org/abs/2104.00420) (2021)
9. Fornasier, M., Klock, T., Riedl, K.: Consensus-based optimization methods converge globally in mean-field law. [arXiv:2103.15130](https://arxiv.org/abs/2103.15130) (2021)
10. Fornasier, M., Pareschi, L., Huang, H., Sünnen, P.: Consensus-based optimization on the sphere: Convergence to global minimizers and machine learning. *Journal of Machine Learning Research* **22**(237), 1–55 (2021)

11. Holland, J.H.: Adaptation in natural and artificial systems. An introductory analysis with applications to biology, control, and artificial intelligence. University of Michigan Press, Ann Arbor, Mich. (1975)
12. Huang, H., Qiu, J.: On the mean-field limit for the consensus-based optimization. arXiv:2105.12919 (2021)
13. Huang, H., Qiu, J., Riedl, K.: On the global convergence of particle swarm optimization methods. arXiv:2201.12460 (2022)
14. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of ICNN'95 - International Conference on Neural Networks. vol. 4, pp. 1942–1948. IEEE (1995)
15. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE **86**(11), 2278–2324 (1998)
16. LeCun, Y., Cortes, C., Burges, C.: MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist/> (2010)
17. Miller, P.D.: Applied asymptotic analysis, Graduate Studies in Mathematics, vol. 75. American Mathematical Society, Providence, RI (2006)
18. Pinnau, R., Totzeck, C., Tse, O., Martin, S.: A consensus-based model for global optimization and its mean-field limit. Math. Models Methods Appl. Sci. **27**(1), 183–204 (2017)
19. Rastrigin, L.: The convergence of the random search method in the extremal control of a many parameter system. Automaton & Remote Control **24**, 1337–1342 (1963)
20. Toscani, G.: Kinetic models of opinion formation. Commun. Math. Sci. **4**(3), 481–496 (2006)