

Consensus-Based Optimization on the Sphere: Convergence to Global Minimizers and Machine Learning

Massimo Fornasier

*Department of Mathematics
Technical University of Munich
Boltzmannstraße 3, 85748 Garching (Munich), Germany*

MASSIMO.FORNASIER@MA.TUM.DE

Hui Huang

*Department of Mathematics
Technical University of Munich
Boltzmannstraße 3, 85748 Garching (Munich), Germany*

HUI.HUANG@TUM.DE

Lorenzo Pareschi

*Department of Mathematics & Computer Science
University of Ferrara
Via Machiavelli 30, Ferrara, 44121, Italy*

LORENZO.PARESCHI@UNIFE.IT

Philippe Sünnen

*Department of Mathematics
Technical University of Munich
Boltzmannstraße 3, 85748 Garching (Munich), Germany*

PHILIPPE.SUENNEN@MA.TUM.DE

Editor: Marc Schoenauer

Abstract

We investigate the implementation of a new stochastic Kuramoto-Vicsek-type model for global optimization of nonconvex functions on the sphere. This model belongs to the class of Consensus-Based Optimization. In fact, particles move on the sphere driven by a drift towards an instantaneous consensus point, which is computed as a convex combination of particle locations, weighted by the cost function according to Laplace's principle, and it represents an approximation to a global minimizer. The dynamics is further perturbed by a random vector field to favor exploration, whose variance is a function of the distance of the particles to the consensus point. In particular, as soon as the consensus is reached the stochastic component vanishes. The main results of this paper are about the proof of convergence of the numerical scheme to global minimizers provided conditions of well-preparation of the initial datum. The proof combines previous results of mean-field limit with a novel asymptotic analysis, and classical convergence results of numerical methods for SDE. We present several numerical experiments, which show that the algorithm proposed in the present paper scales well with the dimension and is extremely versatile. To quantify the performances of the new approach, we show that the algorithm is able to perform essentially as good as *ad hoc* state of the art methods in challenging problems in signal processing and machine learning, namely the phase retrieval problem and the robust subspace detection.

Keywords: global optimization, consensus-based optimization, asymptotic convergence analysis, stochastic Kuramoto-Vicsek model, mean-field limit, numerical methods for SDE.

1. Introduction

1.1 Derivative-free optimization and metaheuristics

Machine learning is about parametric nonlinear algorithms, whose parameters are optimized towards several tasks such as feature selection, dimensionality reduction, clustering, classification, regression, and generation. In view of the nonlinearity of the algorithms and the use of often nonconvex data misfits or penalizations/regularizations, the training phase is most commonly a nonconvex optimization. Moreover, the efficacy of such methods is often determined by considering a large amount of parameters, which makes the optimization problem high dimensional and therefore quite hard. Often first order methods, such as gradient descent methods, are preferred both because of speed and scalability and because they are considered generically able to escape the trap of saddle points (Lee et al., 2019), and in some cases they are able even to compute global minimizers (Chen et al., 2019; Liu et al., 2020; Bah et al., 2019). Nevertheless, for some models, such as training of certain feed-forward deep neural networks, the gradient tends to explode or vanish, (Bengio et al., 1994). For many other problems the derivative of the objective function can be extremely computational expensive to compute or the objective function may not be even differentiable at all. Finally, gradient descent methods do not offer in general guarantees of global convergence and, in view of high dimensionality and nonconvexity, a large amount of local minimizers are expected to possibly trap the dynamics (see Section 2.4.2 for concrete examples).

Long before the current uses in machine learning, nonconvex optimizations have been considered in optimal design of any sort of processes and several solutions have been proposed to tackle these problems. In this paper we are concerned with those which fall into the class of *metaheuristics* (Aarts and Korst, 1989; Back et al., 1997; Blum and Roli, 2003; Glover and Kochenberger, 2006), which provide empirically robust solutions to hard optimization problems with fast algorithms. Metaheuristics are methods that orchestrate an interaction between local improvement procedures and global/high level strategies, and combine random and deterministic decisions, to create a process capable of escaping from local optima and performing a robust search of a solution space. Starting with the groundbreaking work of Rastrigin on Random Search in 1963 (Rastrigin, 1963), numerous mechanisms for multi-agent global optimization have been considered, among the most prominent instances we recall the Simplex Heuristics (Nelder and Mead, 1965), Evolutionary Programming (Fogel, 2006), the Metropolis-Hastings sampling algorithm (Hastings, 1970), Genetic Algorithms (Holland, 1992), Particle Swarm Optimization (PSO) (Kennedy, 2010; Poli et al., 2007), Ant Colony Optimization (ACO) (Dorigo and Blum, 2005), Simulated Annealing (SA), (Holley and Stroock, 1988; Kirkpatrick et al., 1983). Despite the tremendous empirical success of these techniques, it is still quite difficult to provide guarantees of robust convergence to global minimizers, because of the random component of metaheuristics, which would require to discern the stochastic dependencies. Such analysis is often a very hard task, especially for those methods that combine instantaneous decisions with memory mechanisms.

1.2 Consensus-based optimization

Recent work by Pinnau, Carrillo et al. (Pinnau et al., 2017; Carrillo et al., 2018) on Consensus-based Optimization (CBO) focuses on instantaneous stochastic and deterministic decisions in order to establish a consensus among particles on the location of the global minimizers within a domain. In view of the instantaneous nature of the dynamics, the evolution can be interpreted as a system of first order stochastic differential equations (SDEs), whose large particle limit is approximated by a deterministic partial differential equation of mean-field type. The large time behavior of such a deterministic PDE can be analyzed by classical techniques of large deviation bounds and the global convergence of the mean-field model can be mathematically proven in a rigorous way for a large class of optimization problems, see (Fornasier et al., 2021). Certainly CBO is a significantly simpler mechanism with respect to more sophisticated metaheuristics, which can include different features including memory of past exploration. Nevertheless, it is general enough to explain other metaheuristics methods such as particle swarm optimization (Cipriani et al., 2021; Grassi and Pareschi, 2021) and powerful and robust enough to tackle many interesting nonconvex optimizations of practical relevance in machine learning. In particular CBO and variants have been recently tested as optimization methods for the training of artificial neural networks, showing competitive results over stochastic gradient descent, also in terms of generalization error, see (Benfenati et al., 2021, Section 5.5 and Figure 6 and Figure 7) and (Carrillo et al., 2021, Section 4.3 and Figure 6). From the theoretical side, one may refer, for instance, to the recent paper (Cao and Gu, 2020) for theoretical estimates of the generalization error in training deep neural networks. If one really inspects carefully the results and the proofs, one realizes that (Cao and Gu, 2020, Theorem 3.2) is all about the *global* optimization by gradient descent of the empirical risk. Since CBO methods are precisely designed to achieve global optimization, they also allow for same guarantees of generalization errors as gradient descent, as one could simply substitute (Cao and Gu, 2020, Theorem 3.2) with any global convergence result of CBO and obtain the same bounds. We mention also that CBO with adaptive momentum has been recently proposed in (Chen et al., 2020), providing better generalization results than the state-of-the-art method Adam in solving a deep learning task for partial differential equations with low-regularity solutions. Some theoretical gaps remain open in the analysis of CBO though, in particular the lack of a rigorous derivation of the mean-field limit as in (Carrillo et al., 2018), which has been very recently established just in a non-quantitative form in (Huang and Qiu, 2021).

1.3 Consensus-based optimization on the sphere

Motivated by lack of a quantitative mean-field limit and by the several potential applications in machine learning, in the companion paper (Fornasier et al., 2020) we introduced for the first time in the literature a new CBO approach to solve the following constrained optimization problem

$$v^* \in \operatorname{argmin}_{v \in \Gamma} \mathcal{E}(v), \tag{1.1}$$

where $\mathcal{E} : \mathbb{R}^d \rightarrow \mathbb{R}$ is a given continuous cost function, which we wish to minimize over a compact hypersurface Γ . In this paper we consider the particular case of $\Gamma = \mathbb{S}^{d-1}$ being the hypersphere, for which we formulate a system of N interacting particles $((V_t^i)_{t \geq 0})_{i=1, \dots, N}$

satisfying the following stochastic Kuramoto-Vicsek-type dynamics expressed in Itô's form

$$dV_t^i = \lambda P(V_t^i) v_{\alpha, \mathcal{E}}(\rho_t^N) dt + \sigma |V_t^i - v_{\alpha, \mathcal{E}}(\rho_t^N)| P(V_t^i) dB_t^i - \frac{\sigma^2}{2} (V_t^i - v_{\alpha, \mathcal{E}}(\rho_t^N))^2 \frac{(d-1)V_t^i}{|V_t^i|^2} dt, \quad (1.2)$$

where $\lambda > 0$ is a suitable drift parameter, $\sigma > 0$ is a diffusion parameter,

$$\rho_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{V_t^i} \quad (1.3)$$

is the empirical measure of the particles (δ_v is the Dirac measure at $v \in \mathbb{R}^d$), and

$$v_{\alpha, \mathcal{E}}(\rho_t^N) = \sum_{j=1}^N \frac{V_t^j \omega_\alpha^\mathcal{E}(V_t^j)}{\sum_{i=1}^N \omega_\alpha^\mathcal{E}(V_t^i)} = \frac{\int_{\mathbb{R}^d} v \omega_\alpha^\mathcal{E}(v) d\rho_t^N}{\int_{\mathbb{R}^d} \omega_\alpha^\mathcal{E}(v) d\rho_t^N}, \quad \omega_\alpha^\mathcal{E}(v) := e^{-\alpha \mathcal{E}(v)}. \quad (1.4)$$

Here and below we denote with $\int f(v) d\mu(v)$ or equivalently $\int f(v) \mu(dv)$ the integration of an arbitrary function with respect to a measure μ . This stochastic system is considered complemented with independent and identically distributed (i.i.d.) initial data $V_0^i \in \mathbb{S}^{d-1}$ with $i = 1, \dots, N$, and the common law is denoted by $\rho_0 \in \mathcal{P}(\mathbb{S}^{d-1})$. The trajectories $((B_t^i)_{t \geq 0})_{i=1, \dots, N}$ denote N independent standard Brownian motions in \mathbb{R}^d . In (1.2) the projection operator $P(\cdot)$ is defined by

$$P(v) = I - \frac{vv^T}{|v|^2}. \quad (1.5)$$

It is easy to check that

$$P(v)v = 0 \text{ and } v \cdot P(v)y = 0 \text{ for all } y \in \mathbb{R}^d. \quad (1.6)$$

The choice of the weight function $\omega_\alpha^\mathcal{E}$ in (1.4) comes from the well-known Laplace principle (Miller, 2006; Dembo and Zeitouni, 2010; Pinnau et al., 2017), a classical asymptotic method for integrals, which states that for any probability measure $\rho \in \mathcal{P}_{\text{ac}}(\mathbb{R}^d)$ (absolutely continuous), it holds

$$\lim_{\alpha \rightarrow \infty} \left(-\frac{1}{\alpha} \log \left(\int_{\mathbb{R}^d} e^{-\alpha \mathcal{E}(v)} d\rho(v) \right) \right) = \inf_{v \in \text{supp } \rho} \mathcal{E}(v). \quad (1.7)$$

Let us discuss the mechanism of the dynamics. The right-hand-side of the equation (1.2) is made of three terms. The first deterministic term $\lambda P(V_t^i) v_{\alpha, \mathcal{E}}(\rho_t^N) dt = -\lambda P(V_t^i) (V_t^i - v_{\alpha, \mathcal{E}}(\rho_t^N)) dt$, because of (1.6) $P(V_t^i) V_t^i = 0$, imposes a drift to the dynamics towards $v_{\alpha, \mathcal{E}}$, which is the current consensus point at time t as an approximation to the global minimizer, and the term disappears when $V_t^i = v_{\alpha, \mathcal{E}}$. In fact, the consensus point $v_{\alpha, \mathcal{E}}$ is explicitly computed as in (1.4) and it may lay in general outside \mathbb{S}^{d-1} . This choice of an embedded weighted barycenter is very simple, compatible with fast computations, and, for a compact manifold as \mathbb{S}^{d-1} , it is a good proxy for a minimizer v^* . One could alternatively consider the computation of a weighted barycenter on the manifold

$$v_{\alpha, \mathcal{E}}^{\mathbb{S}^{d-1}}(\rho_t^N) = \arg \min_{v \in \mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} d_{\mathbb{S}^{d-1}}(v, w)^p e^{-\alpha \mathcal{E}(w)} d\rho_t^N(w),$$

where $d_{\mathbb{S}^{d-1}}$ is the (Riemannian) distance on \mathbb{S}^{d-1} , $p > 0$, and ρ_t^N is again the particle distribution. However, the computation of $v_{\alpha, \mathcal{E}}^{\mathbb{S}^{d-1}}$ is in general not explicit and one may have to solve at each time t a nontrivial optimization problem over the sphere in order to compute $v_{\alpha, \mathcal{E}}^{\mathbb{S}^{d-1}}$, the so-called Weber problem. These are all good reasons for choosing the simpler embedded alternative (1.4).

The second stochastic term $\sigma|V_t^i - v_{\alpha, \mathcal{E}}(\rho_t^N)|P(V_t^i)dB_t^i$ introduces a random decision to favor the exploration, whose variance is a function of the distance of particles to the consensus points. In particular, as soon as the consensus is reached, then the stochastic component vanishes. The last term $-\frac{\sigma^2}{2}(V_t^i - v_{\alpha, \mathcal{E}}(\rho_t^N))^2 \frac{(d-1)V_t^i}{|V_t^i|^2} dt$, combined with $P(\cdot)$, it is needed to ensure that the dynamics stays on the sphere despite the Brownian motion component. Namely, this third term stems from Itô's formula to ensure $d|V_t^i|^2 = 0$, see (Fornasier et al., 2020, Theorem 2.1). We further notice that the dynamics does not make use of any derivative of \mathcal{E} , but only of its pointwise evaluations, which appear integrated in (1.4). Hence, the equation can be in principle numerically implemented at discrete times also for cost functions \mathcal{E} which are just continuous and with no further smoothness and the resulting numerical scheme is fully derivative-free. We require more regularity of \mathcal{E} exclusively to ensure formal well-posedness of the evolution and for the analysis of its large time behavior, but it is not necessary for its numerical implementation. A possible discrete-time approximation and resulting numerical scheme, which we consider in this paper is given by the projected Euler-Maruyama method as follows: generate i.i.d. V_0^i , $i = 1, \dots, N$ sample vectors according to $\rho_0 \in \mathcal{P}(\mathbb{S}^{d-1})$ and iterate for $n = 0, 1, \dots$

$$\begin{aligned} \tilde{V}_{n+1}^i &\leftarrow V_n^i + \Delta t \lambda P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i \\ &\quad - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1) V_n^i, \quad (1.8) \\ V_{n+1}^i &\leftarrow \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i|, \quad i = 1, \dots, N, \end{aligned}$$

for

$$V_n^{\alpha, \mathcal{E}} = \frac{1}{N_\alpha} \sum_{j=1}^N w_\alpha^\mathcal{E}(V_n^j) V_n^j, \quad N_\alpha = \sum_{j=1}^N w_\alpha^\mathcal{E}(V_n^j), \quad (1.9)$$

where $w_\alpha^\mathcal{E}(V_n^j) = \exp(-\alpha \mathcal{E}(V_n^j))$, and ΔB_n^i are independent normal random vectors $N(0, \Delta t)$. Let us stress however that this is by no means the only possible discretization and we refer to, e.g., (Platen, 1999), for picking a favorite alternative scheme.

1.4 Main result and sketch of its proof

The main result of the present paper establishes the convergence of the discrete- and continuous-time dynamics to global minimizers of \mathcal{E} under mild smoothness conditions and local coercivity of the function around global minimizers. The analysis goes in two steps: First of all, one needs to establish the large particle limit of the stochastic dynamics (1.2). This first step was already obtained in (Fornasier et al., 2020), whose main results are about the well-posedness of (1.2) and its rigorous mean-field limit - which is an open issue for unconstrained CBO (Carrillo et al., 2018) - to the following nonlocal, nonlinear Fokker-Planck

equation

$$\partial_t \rho_t = \lambda \nabla_{\mathbb{S}^{d-1}} \cdot (\langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - v_{\alpha, \mathcal{E}}(\rho_t)) \rho_t + \frac{\sigma^2}{2} \Delta_{\mathbb{S}^{d-1}} (|v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 \rho_t), \quad t > 0, v \in \mathbb{S}^{d-1}, \quad (1.10)$$

with the initial datum $\rho_0 \in \mathcal{P}(\mathbb{S}^{d-1})$. Here $\rho_t = \rho(t, v) \in \mathcal{P}(\mathbb{S}^{d-1})$ is a Borel probability measure on \mathbb{S}^{d-1} and

$$v_{\alpha, \mathcal{E}}(\rho_t) = \frac{\int_{\mathbb{S}^{d-1}} v \omega_{\alpha}^{\mathcal{E}}(v) d\rho_t}{\int_{\mathbb{S}^{d-1}} \omega_{\alpha}^{\mathcal{E}}(v) d\rho_t}.$$

The operators $\nabla_{\mathbb{S}^{d-1}} \cdot$ and $\Delta_{\mathbb{S}^{d-1}}$ denote the divergence and Laplace-Beltrami operator on the sphere \mathbb{S}^{d-1} respectively. The mean-field limit is achieved through the coupling method (Sznitman, 1991; Fetecau et al., 2019; Huang and Liu, 2017) by introducing the mean-filed dynamics satisfying

$$d\bar{V}_t^i = \lambda P(\bar{V}_t^i) v_{\alpha, \mathcal{E}}(\rho_t) dt + \sigma |\bar{V}_t^i - v_{\alpha, \mathcal{E}}(\rho_t)| P(\bar{V}_t^i) dB_t^i - \frac{\sigma^2}{2} (\bar{V}_t^i - v_{\alpha, \mathcal{E}}(\rho_t))^2 \frac{(d-1) \bar{V}_t^i}{|\bar{V}_t^i|^2} dt, \quad (1.11)$$

where $(\bar{V}_t^i)_{t \geq 0, i=1, \dots, N}$ are i.i.d. with common law $(\rho_t)_{t \geq 0}$ satisfying (1.10). It yields the following quantitative form of mean-field limit

$$\sup_{t \in [0, T]} \sup_{i=1, \dots, N} \mathbb{E} \left[|V_t^i - \bar{V}_t^i|^2 \right] \leq CN^{-1}, \quad N \rightarrow \infty, \quad (1.12)$$

for any $T > 0$ time horizon, see (Fornasier et al., 2020, Theorem 3.1 and Remark 3.1). The rate of convergence (1.12) is not affected by the curse of dimension and, for $\Gamma = \mathbb{S}^{d-1}$ the constant C depends at most linearly in d and, as a worst case analysis, exponentially in α and in T , see (Fornasier et al., 2020, Remark 3.2 and Lemma 3.1) respectively. Besides the well-posedness of (1.10) in the space of probability measures established in (Fornasier et al., 2020, Section 2.3), for more regular datum ρ_0 , we prove additionally in the present paper existence and uniqueness of distributional solutions $\rho \in L^2([0, T], H^1(\mathbb{S}^{d-1}))$ at any finite time $T > 0$, see Theorem 13. This auxiliary regularity results is needed in our convergence analysis.

The second step to establish global convergence, which is also carried out in the present paper, is about proving the large time asymptotics of the PDE solution $\rho_t(v) = \rho(t, v)$. In Theorem 6 we show that, for any $\epsilon > 0$ there exists suitable parameters α, λ, σ and well-prepared initial densities ρ_0 such that for $T^* > 0$ large enough the expected value of the distribution $E(\rho_{T^*}) = \int v d\rho_{T^*}(v)$ is near a global minimizers v^* of \mathcal{E} , i.e.,

$$|E(\rho_{T^*}) - v^*| \leq C\epsilon. \quad (1.13)$$

The convergence to $E(\rho_{T^*})$ is exponential in time and the rate depends on the parameters $\epsilon, \alpha, \lambda, \sigma$. We summarize the main result as follows.

Theorem 1 *Assume $\mathcal{E} \in C^2(\mathbb{S}^{d-1})$ and that for any $v \in \mathbb{S}^{d-1}$ there exists a minimizer $v^* \in \mathbb{S}^{d-1}$ of \mathcal{E} (which may depend on v) such that it holds*

$$|v - v^*| \leq C_0 |\mathcal{E}(v) - \underline{\mathcal{E}}|^\beta, \quad (1.14)$$

where β, C_0 are some positive constants and $\underline{\mathcal{E}} := \inf_{v \in \mathbb{S}^{d-1}} \mathcal{E}(v)$. We also denote $\bar{\mathcal{E}} := \sup_{v \in \mathbb{S}^{d-1}} \mathcal{E}(v)$, $C_{\alpha, \varepsilon} = e^{\alpha(\bar{\mathcal{E}} - \underline{\mathcal{E}})}$, and $C_{\sigma, d} = \frac{(d-1)\sigma^2}{2}$. Additionally for any $\varepsilon > 0$ assume that the initial datum ρ_0 and parameters λ, σ are well-prepared in the sense of Definition 5 for a time horizon $T^* > 0$ and a parameter $\alpha^* > 0$ large enough, depending on C_0 and β . Then the iterations $\{V_n^i := V_{\Delta t, n}^i : n = 0, \dots, n_{T^*}; i = 1 \dots N\}$ generated by a discrete-time approximation of (1.2) fulfill the following error estimate

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N V_{n_{T^*}}^i - v^* \right|^2 \right] \leq \underbrace{C_1(\Delta t)^{2m}}_{\text{Discr. err.}} + \underbrace{C_2 N^{-1}}_{\text{Mean-field lim.}} + \underbrace{C_3 \varepsilon^2}_{\text{Laplace princ.}}, \quad (1.15)$$

where m is the order of approximation of the numerical scheme. (For the projected Euler-Maruyama scheme the order is $m = 1/2$.) The constant C_1 depends linearly on the dimension d and the number of particles N , and possibly exponentially on T^* and the parameters λ and σ ; the constant C_2 depends linearly on the dimension d , polynomially in $C_{\alpha^*, \varepsilon}$, and exponentially in T^* ; the constant C_3 depends on C_0 and β . The convergence is exponential with rate

$$\lambda\theta - 4C_{\alpha^*, \varepsilon}C_{\sigma, d} > 0, \quad (1.16)$$

for a suitable $0 < \theta < 1$.

The detailed proof of this result is reported in Section 3.4. We provide here a sketch of it.

Proof (Sketch). We recall the definitions of expectation and variance of ρ_t as

$$E(\rho_t) := \int_{\mathbb{S}^{d-1}} v d\rho_t(v) \quad V(\rho_t) := \frac{1}{2} \int_{\mathbb{S}^{d-1}} |v - E(\rho_t)|^2 d\rho_t(v).$$

By combining the coercivity condition (1.14) and the Laplace principle (1.7) we show that

$$\left| \frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right| \leq C(C_0, \|\nabla \mathcal{E}\|_\infty, \beta) \left((C_{\alpha, \varepsilon})^\beta V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^\beta \right).$$

Hence, in order to prove the large time convergence to a global minimizer (1.13), we may want to show that the variance is monotonically decreasing to zero with an exponential rate. An explicit computation leveraging the PDE (1.10) yields

$$\begin{aligned} \frac{d}{dt} V(\rho_t) &= -\lambda V(\rho_t) \langle E(\rho_t), v_{\alpha, \varepsilon} \rangle - \frac{\lambda v_{\alpha, \varepsilon}^2 + 1}{2} 2V(\rho_t) + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \varepsilon})^2 d\rho_t \\ &\quad + C_{\sigma, d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha, \varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t \\ &\leq -\lambda V(\rho_t) \left(\langle E(\rho_t), v_{\alpha, \varepsilon} \rangle + \frac{v_{\alpha, \varepsilon}^2 + 1}{2} \right) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \varepsilon})^2 d\rho_t + 4C_{\alpha, \varepsilon} C_{\sigma, d} V(\rho_t). \end{aligned}$$

The idea is to balance all the terms on the right-hand side by using the parameters λ, σ in such a way of obtaining a negative sign. Under assumptions of well-preparation, $V(\rho_t)$ is

actually small enough for ensuring $|E(\rho_t)| \approx \langle E, v_{\alpha, \varepsilon} \rangle \approx |v_{\alpha, \varepsilon}| \approx 1$, and, thanks to Theorem 13, for any $\delta > 0$ arbitrarily small

$$\frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \varepsilon})^2 d\rho_t \leq \gamma \lambda V(\rho_t) + \delta,$$

for a suitable $0 < \gamma < 2$, γ depending on δ . Hence,

$$V(\rho_t) \leq V(\rho_0) e^{-(\lambda\theta - 4C_{\alpha, \varepsilon} C_{\sigma, d})t} + \delta$$

for $\theta \approx 2 - \gamma$, and one concludes the convergence in finite time T^* as in (1.13), under the given error threshold $\epsilon > 0$. By combining now classical results of convergence of numerical approximations¹ $(V_{\Delta t, n}^i)_{i=1, \dots, N}$ (Platen, 1999) with the the mean-field approximation (1.12) (Fornasier et al., 2020) and the large time behavior (1.13), which is proven in detail in Theorem 6 below, we obtain that the expected large time outcome of the numerical approximation to (1.2) is about a global minimizer of \mathcal{E}

$$\begin{aligned} & \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N V_{\Delta t, n_{T^*}}^i - v^* \right|^2 \right] \\ & \lesssim \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (V_{\Delta t, n_{T^*}}^i - V_{T^*}^i) \right|^2 \right] + \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (V_{T^*}^i - \bar{V}_{T^*}^i) \right|^2 \right] \\ & \quad + \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \bar{V}_{T^*}^i - E(\rho_{T^*}) \right|^2 \right] + |E(\rho_{T^*}) - v^*|^2 \\ & \lesssim (\Delta t)^{2m} + N^{-1} + \epsilon^2, \end{aligned} \tag{1.17}$$

where m is the order of strong convergence of the numerical method. ■

1.5 Discussion

Some comments about the result and its proof are in order. First of all, we stress that Theorem 1 is the first and so far the unique complete result of convergence of consensus-based optimizations in the literature. In fact, the results in (Carrillo et al., 2018, 2021) are exclusively addressing the large time behavior of the mean-field PDE, because, for consensus-based optimization in the Euclidean space, a mean-field approximation of the type (1.12) has not been established yet for unbounded \mathcal{E} . Similarly, the convergence proof of the purely numerical scheme in (Ha et al., 2020) establishes convergence in ess inf_ω over all possible realizations ω (significantly weaker than (1.15)). In particular it does not provide a rate of convergence in terms of number N of particles. Moreover the result is established under the simplified assumption that the noise is equal for all particles.

Our proof strategy described above made of a numerical approximation, mean-field limit,

1. In this paper we consider numerical approximations by Euler-Maruyama scheme, which converges strongly with order $m = 1/2$ (Platen, 1999), see Algorithm 1 in Section 2.2.

and asymptotic analysis parallels a similar approach by Montanari et al. (Mei et al., 2018; Javanmard et al., 2020) for proving the convergence of stochastic gradient descent to global minimizers in the training of two-layer neural networks.

The initial datum ρ_0 has to be interpreted as the uncertainty on the location of a global minimizer. The condition of Definition 5 of well-preparation of ρ_0 may have a locality flavour, i.e., they essentially require that ρ_0 has small variance and simultaneously it not centered too far from a global minimizers v^* of \mathcal{E} . However, in the case the function \mathcal{E} is symmetric, i.e., $\mathcal{E}(v) = \mathcal{E}(-v)$ (as it happens in numerous applications, in particular the ones we present in this paper) and $C_0 > 0$ is relatively large for $\beta \geq 1$, then the condition is generically/practically satisfied at least for one of the two global minimizers $\pm v^*$. The convergence result is based on proving the monotone decay of the variance $V(\rho_t) = \int |v - E(\rho_t)|^2 d\rho_t(v)$, see Proposition 12, and this cannot be achieved unless the initial condition is well-prepared. In fact, for a non-symmetric function \mathcal{E} , a given unique global minimizer v^* , and for a datum ρ_0 fully concentrated around the opposite vector $-v^*$, i.e., on the other side of the sphere, the variance may start small, but it must grow well before getting small again. Hence, it is not possible for arbitrary \mathcal{E} and initial datum to have monotone decay of the variance, and we conjecture that the result can be further improved to obtain even more generic initial conditions, but one needs to use a different proving technique.

Let us now discuss the interplay between the different approximations and the constants appearing in (1.15). While the constants C_1, C_2 in (1.15) depend explicitly only linearly on the dimension d , the constants C_2 may depend polynomially on $C_{\alpha^*, \mathcal{E}} = e^{\alpha^*(\bar{\mathcal{E}} - \underline{\mathcal{E}})}$, hence, exponentially in α^* . This exponential dependence stems from the worst case analysis due to (Fornasier et al., 2020, Lemma 3.1). Moreover, at this level of generality it is difficult to establish how α^* depends on d as such dependence is strongly affected by the particular objective function \mathcal{E} and ρ_0 : to clarify the predicament, in the extreme case where $\rho_0 \approx \delta_{v^*}$ it is

$$\left(-\frac{1}{\alpha} \log \left(\int_{\mathbb{R}^d} e^{-\alpha \mathcal{E}(v)} d\rho_0(v) \right) \right) \approx \underline{\mathcal{E}},$$

independently of α (also for α very small!). Also one does not expect a strong dependence of α^* on d for the case where $C_0 > 0$ is large for $\beta \geq 1$ with a symmetric behavior of \mathcal{E} around v^* . Instead, in the worst case scenario we may need $\alpha^* \geq d$ and our estimates may simply reflect the fact that the optimization problem at hand is NP-hard or intrinsically affected by the curse of dimensionality. However, this is by no means the typical situation, as in our numerical experiments in Section 2.4.2 we show that the method scales well with the dimension also for problems with $d \approx 3000$. Due to the worst case analysis and the use of Gronwall's inequalities in the literature, constants C_1, C_2 may depend also exponentially on T^* ; however T^* is fixed at the beginning and the initial datum and parameters are assumed to be well-prepared so that the algorithm reaches precisely at the time T^* the expected accuracy. Hence, T^* does not need to be very large if we assume that our initial datum ρ_0 offers already a reasonable confidence on the location of a global minimizer. In particular the rate of convergence $\lambda\theta - 4C_{\alpha, \mathcal{E}}C_{\sigma, d}$ is completely determined by the choices of λ and σ . The choice of $\lambda > 0$ large necessarily implies the discretization parameter Δt small, as C_1 depends by the worst case analysis and the use of Gronwall's inequalities exponentially in λ . Recent work (Carrillo et al., 2018) on unconstrained consensus-based optimization with




Figure 1: The Ackley function for $d = 2$ on $[-3, 3]^2$ and its representation for $d = 3$ in the constrained case over the half sphere \mathbb{S}^2 (right). The global minimum corresponds to the direction $v^* = (0, 0, 1)^T$.

anisotropic noise suggests the possibility of having parameters completely independent of the dimension.

1.6 Organization of the paper

The rest of the paper is organized as follows: in Section 2 we present and explain right away the numerical implementation, Algorithm 1, of the stochastic Kuramoto-Vicsek (sKV) system (1.2). We further propose a few relevant speed-ups, which will be implemented in Algorithm 2. As a warm up, we illustrate the behavior of the algorithms on the synthetic example of the Ackley function over the sphere (see Figure 1) in dimension $d = 3$. In the second part of this section, we present applications in signal processing and machine learning, namely the phase retrieval problem and the robust subspace detection and we provide comparisons with state of the art methods. For the robust subspace detection we test the algorithm also in dimension $d \approx 3000$ on the *Adult Faces Database* (Bainbridge et al., 2013) for the computation of eigenfaces. These experiments show that the algorithm scales well with the dimension and is extremely versatile (one just needs to modify the definition of the function \mathcal{E} and the rest goes with the same code!). The algorithm is able to perform essentially as good as *ad hoc* state of the art methods and in some instances it obtains quantitatively better results. For the sake of reproducible research, in the repository <https://github.com/PhilippeSu/KV-CBO> we provide the Matlab code, which implements the algorithms on the test cases of this paper. In Section 3 we provide the analysis of global optimization guarantees, which yield the main error estimate (1.15). In Section 4 we collect proofs of a few auxiliary results.

2. Numerical Implementation and Tests

In this section we report several tests and examples of application of the consensus based optimization (CBO) method based on the stochastic Kuramoto-Vicsek (sKV) system. First, we discuss fast first order discretization methods for the stochastic system, which preserve the dynamics on the multi-dimensional sphere. Implementation aspects and speed-ups are also analyzed. In particular, we derive fast algorithms, which permit to obtain an exponentially diminishing computational cost in time. Next, we test the method and its sensitivity to the choice of the computational parameters with respect to some well-known prototype test functions in high dimensions. Real-life applications are also provided to sustain the versatility and scalability of the method.

2.1 Discretization of the sKV system

We discuss the discretization of the sKV system in Itô's form

$$dV_t^i = \lambda P(V_t^i) V_t^{\alpha, \mathcal{E}} dt + \sigma |V_t^i - V_t^{\alpha, \mathcal{E}}| P(V_t^i) dB_t^i - \frac{\sigma^2}{2} (V_t^i - V_t^{\alpha, \mathcal{E}})^2 \frac{(d-1)V_t^i}{|V_t^i|^2} dt, \quad (2.1)$$

with $V_t^i \in \mathbb{S}^{d-1}$, $i = 1, \dots, N$, and

$$V_t^{\alpha, \mathcal{E}} = \sum_{j=1}^N \frac{V_t^j \omega_\alpha^\mathcal{E}(V_t^j)}{\sum_{i=1}^N \omega_\alpha^\mathcal{E}(V_t^i)} = v_{\alpha, \mathcal{E}}(\rho_t^N).$$

First let us remark that for $d = 2$ the problem is considerably simpler since the passage to spherical coordinates permits an easy integration of the system by preserving its geometrical nature of motion on \mathbb{S}^1 . However, for arbitrary dimensions this is more complicated and we must integrate the stochastic system in the vector form (2.1). We refer to (Platen, 1999) for an introduction to numerical methods for SDEs and to (Hairer et al., 2006) for deterministic time discretizations, which preserve some geometrical properties of the solution.

Let us denote $|V| = \|V\|_2 = \langle V, V \rangle^{\frac{1}{2}}$ the Euclidean norm. A simple geometrical argument allows to prove the following observation:

Lemma 2 *Let us consider a one step time discretization of (2.1) in the general form*

$$V_{n+1}^i = V_n^i + \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i) \quad (2.2)$$

where the function $\Phi(\Delta t, \cdot, \cdot, \xi_n^i) : \mathbb{R}^{2d} \rightarrow \mathbb{R}^d$ defines the method, $\Delta t > 0$ is the time step, $V_n^i \approx V_t^i|_{t=t^n}$, $t^n = n\Delta t$ and ξ_n^i are independent random variables.

Then

$$|V_{n+1}^i|^2 = |V_n^i|^2 \quad (2.3)$$

if and only if

$$\langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), V_{n+1}^i + V_n^i \rangle = 0. \quad (2.4)$$

This shows that $\Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i)$ must be orthogonal to $V_{n+1}^i + V_n^i$ in order to preserve the norm and, consequently, to obtain one step methods satisfying (2.3) we have to resort to implicit methods.

For example, it is immediate to verify that the Euler-Maruyama method

$$V_{n+1}^i = V_n^i + \Delta t \lambda P(V_n^i) V_n^{\alpha, \mathcal{E}} dt + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 \frac{(d-1)V_n^i}{|V_n^i|^2}, \quad (2.5)$$

where $\Delta B_n^i = B_{t_{n+1}}^i - B_{t_n}^i$ are independent normal random variables $N(0, \Delta t)$ with mean zero and variance Δt , is not invariant with respect to the norm of V_n^i .

A method that preserves the norm is obtained by modifying the Euler-Maryuama method as follows

$$V_{n+1}^i = V_n^i + \Delta t \lambda P(V_{n+\frac{1}{2}}^i) (V_n^{\alpha, \mathcal{E}} - V_n^i) dt + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_{n+\frac{1}{2}}^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 \frac{(d-1)V_{n+\frac{1}{2}}^i}{|V_{n+\frac{1}{2}}^i|^2},$$

where $V_{n+\frac{1}{2}}^i = V_{n+1}^i + V_n^i$ and, for consistency, we have the term $-V_n^i$ in the alignment process since now $P(V_{n+\frac{1}{2}}^i)V_n^i \neq 0$. By similar arguments, we can construct implicit methods of weak order higher than one which preserve the norm of the solution. Implicit methods, however, due to the nonlinearity of the projection operator $P(\cdot)$ require the inversion of a large nonlinear system. This represents a serious drawback for our purposes, where efficiency of the numerical solver is fundamental.

In order to promote efficiency, we consider instead explicit one-step methods that preserve the geometric properties by adopting a projection method at each time step for the iterations to stay on the sphere (Hairer et al., 2006). This corresponds to solve the stochastic differential problem under the algebraic constraint to preserve the norm. Since we are on the unit hypersphere, we simply divide the numerical approximation by its Euclidean norm to get a vector of length one. This class of schemes has the general form

$$\begin{cases} \tilde{V}_{n+1}^i = V_n^i + \Phi(\Delta t, V_n^i, \tilde{V}_{n+1}^i, \xi_n^i), \\ V_{n+1}^i = \frac{\tilde{V}_{n+1}^i}{|\tilde{V}_{n+1}^i|}. \end{cases} \quad (2.6)$$

We keep the dependence from \tilde{V}^{n+1} on the right hand side to include semi-implicit methods with better stability properties than the Euler-Maruyama scheme. One example is obtained by the following integration scheme

$$\tilde{V}_{n+1}^i = V_n^i + \Delta t \lambda P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1) \tilde{V}_{n+1}^i$$

which can be written explicitly as

$$\tilde{V}_{n+1}^i = \frac{1}{1 + \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1)} (V_n^i + \Delta t \lambda P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i). \quad (2.7)$$

In our experiments, since efficiency of the numerical solver is of paramount importance, we rely on projection methods of the type (2.6) based on the simple Euler-Maruyama scheme (2.5) or the semi-implicit scheme (2.7).

Remark 3 Another popular approach is based on simulating the two fundamental processes characterizing the dynamics by a splitting method on the time interval $[n\Delta t, (n+1)\Delta t]$

$$\begin{cases} d\tilde{V}_t^i = \lambda P(\tilde{V}_t^i) \tilde{V}_t^{\alpha, \mathcal{E}} dt, & \tilde{V}_0^i = V_t^i|_{t=n\Delta t}, \\ dV_t^i = \sigma |V_t^i - V_t^{\alpha, \mathcal{E}}| P(V_t^i) dB_t^i - \frac{\sigma^2}{2} (V_t^i - V_t^{\alpha, \mathcal{E}})^2 \frac{(d-1)V_t^i}{|V_t^i|^2} dt, & V_0^i = \tilde{V}_t^i|_{t=(n+1)\Delta t}, \end{cases} \quad (2.8)$$

where the first step is a standard alignment dynamics over the hypersphere and the second step corresponds to solve a Brownian motion with variance $\sigma^2(V_t^i - V_t^{\alpha, \mathcal{E}})^2$ on the unit hypersphere. Typically, the approximated value of $V_t^{\alpha, \mathcal{E}}$ is kept constant in a splitting time step to avoid computing it twice and increasing the computational cost. This approach would allow to solve the first step using standard structure preserving ODEs approaches (Hairer et al., 2006) and to use specific simulation methods for the Brownian motion over the hypersphere in the second step (Großmann et al., 2015; c et al., 2020). We will leave to further study the possibility to apply methods in the splitting form (2.8).

2.2 Implementation aspects and generalizations

First let us point out that the set of three computational parameters, Δt , σ and λ , defining the discretization scheme can be reduced since we can rescale the time by setting

$$\tau = \lambda \Delta t, \quad \nu^2 = \frac{\sigma^2}{\lambda},$$

to obtain a scheme which depends only on two parameters τ and ν . In practice, we can simply assume $\lambda = 1$ and keep the original notations. Starting from a set of computational parameters and a given objective function $\mathcal{E}(\cdot)$ defined on \mathbb{S}^{d-1} , the simplest KV-CBO method is described in Algorithm 1.

Algorithm 1: KV-CBO

Input: $\Delta t, \sigma, \alpha, d, N, n_T$ and the function $\mathcal{E}(\cdot)$

- 1 Generate $V_0^i, i = 1, \dots, N$ sample vectors uniformly on \mathbb{S}^{d-1} ;
- 2 **for** $n = 0$ **to** n_T **do**
- 3 Generate ΔB_n^i independent normal random vectors $N(0, \Delta t)$;
- 4 Compute $V_n^{\alpha, \mathcal{E}}$;
- 5 $\tilde{V}_{n+1}^i \leftarrow V_n^i + \Delta t P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1) V_n^i,$
 $V_{n+1}^i \leftarrow \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i|, i = 1, \dots, N;$
- 6 **end**

The approximation order of the projected Euler-Maruyama method as in KV-CBO is $m = 1/2$. In fact, the only difference with respect to the classical Euler-Maruyama method is the post-projection onto the sphere $V_{n+1}^i \leftarrow \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i|$. We show in the proof of Theorem 1 in Section 3.4 that this may introduce an error of at most order Δt , preserving the order of convergence.

Let us now discuss briefly about the complexity of the scheme as optimization method in order to place the discussion in the correct frame. It is of utmost importance to stress once again that the method is of 0-order, i.e., it is derivative-free. Hence, for those problems for which computing derivatives of the objective function is an unfeasible task, either because of complexity or because of non-differentiability, the KV-CBO is necessarily superior in terms of complexity than first order methods such as gradient descent or second order methods such as Newton method. Note, in particular, that the computational cost for a single time step of KV-CBO is $\mathcal{O}(N)$, the minimum cost to evolve a system of N particles since $V_n^{\alpha, \mathcal{E}}$ is the same for all agents. Let us however mention that the KV-CBO is highly and very easily parallelizable and therefore, on a multi-processor parallel machine, the method can be easily reduced to complexity $\mathcal{O}(N/\mathcal{P})$, where \mathcal{P} is the number of processors. The algorithm may be complemented with a suitable stopping criterion, for example checking consensus using the quantity

$$\frac{1}{N} \sum_{i=1}^N |V_n^i - V_n^{\alpha, \mathcal{E}}| \leq \varepsilon, \quad (2.9)$$

or checking, as in (Carrillo et al., 2021), for $p \geq 0$ that

$$|V_{n+1}^{\alpha, \mathcal{E}} - V_{n-p}^{\alpha, \mathcal{E}}| \leq \varepsilon, \quad (2.10)$$

for a given tolerance ε . In point 5 of Algorithm 1 we used the Euler-Maruyama discretization (2.5), similarly one could use the semi-implicit method (2.7). The computational parameters Δt , σ and α can in practice be adaptively modified from step to step to improve the performance of the method. In the sequel we analyze in more detail some computational aspects and speed ups related to Algorithm 1.

SAMPLING OVER \mathbb{S}^{d-1}

First let us discuss point 1 of algorithm 1, namely how to generate points uniformly over the d -dimensional sphere. Despite the fact that our theoretical results would suggest to use a more concentrated measure ρ_0 to generate the initial points, see Definition 5, the uniform distribution is likely the simplest to be realized and it does certainly not induce initial bias towards any direction. Even though many methods have been designed for low dimension $d \leq 3$, very few of them can be extended to large dimensions. Therefore, the one that is often used for a d -dimensional sphere is the method of normalized Gaussians first proposed by Muller and later by Marsaglia (Muller, 1959; Marsaglia, 1972). The method is extremely simple, and exploits the non-obvious relationship between a uniform distribution on the sphere and the normal distribution. More precisely, to pick a random point on a d -dimensional sphere one first generates d standard normal random variables $\xi_1, \xi_2, \dots, \xi_d \sim N(0, 1)$, then the distribution of the vectors of components

$$v_k = \frac{\xi_k}{\sqrt{\xi_1^2 + \dots + \xi_d^2}}, \quad k = 1, \dots, d \quad (2.11)$$

coincides with the uniform one over the hypersphere \mathbb{S}^{d-1} .

EVALUATION OF $V_n^{\alpha, \mathcal{E}}$

Let us observe that the computation of $V_n^{\alpha, \mathcal{E}}$, points 2 and 6 of Algorithm 1, is crucial and that a straightforward evaluation using

$$V_n^{\alpha, \mathcal{E}} = \frac{1}{N_\alpha} \sum_{j=1}^N w_\alpha^\mathcal{E}(V_n^j) V_n^j, \quad N_\alpha = \sum_{j=1}^N w_\alpha^\mathcal{E}(V_n^j), \quad (2.12)$$

where $w_\alpha^\mathcal{E}(V_n^j) = \exp(-\alpha \mathcal{E}(V_n^j))$, is generally numerically unstable since for large values of $\alpha \gg 1$ the value of N_α is close to zero. On the other hand, the use of large values of α is essential for the performance of the method. A practical way to overcome this issue is based on the following numerical trick

$$\begin{aligned} \frac{w_\alpha^\mathcal{E}(V_n^j)}{N_\alpha} &= \frac{\exp(-\alpha \mathcal{E}(V_n^j))}{\sum_{j=1}^N \exp(-\alpha \mathcal{E}(V_n^j))} \cdot \frac{\exp(\alpha \mathcal{E}(V_n^*))}{\exp(\alpha \mathcal{E}(V_n^*))} \\ &= \frac{\exp(-\alpha(\mathcal{E}(V_n^j) - \mathcal{E}(V_n^*)))}{\sum_{j=1}^N \exp(-\alpha(\mathcal{E}(V_n^j) - \mathcal{E}(V_n^*)))} \end{aligned}$$

where

$$V_n^* := \arg \min_{V \in \{V_n^i\}_{i=1}^N} \mathcal{E}(V) \quad (2.13)$$

is the location of the particle with the minimal function value in the current population. This ensures that for at least one particle $V_n^j = V_n^*$, we have $\mathcal{E}(V_n^j) - \mathcal{E}(V_n^*) = 0$ and therefore, $\exp(-\alpha(\mathcal{E}(V_n^j) - \mathcal{E}(V_n^*))) = 1$. For the sum this leads to $\sum_{j=1}^N \exp(-\alpha(\mathcal{E}(V_n^j) - \mathcal{E}(V_n^*))) \geq 1$, so that the division does not induce a numerical problem. In the numerical simulations we will always compute the weights by the above strategy. Note that, the evaluation of (2.13) has linear cost, and does not affect the overall cost. The computation of $V_n^{\alpha, \mathcal{E}}$ may be accelerated by using the random approach presented in (Albi and Pareschi, 2013) (see Algorithm 4.7). Namely, by considering a random subset J_M of size $M < N$ of the indexes $\{1, \dots, N\}$ and computing

$$V_n^{\alpha, \mathcal{E}, J_M} = \frac{1}{N_\alpha^{J_M}} \sum_{j \in J_M} w_\alpha^\mathcal{E}(V_n^j) V_n^j, \quad N_\alpha^{J_M} = \sum_{j \in J_M} w_\alpha^\mathcal{E}(V_n^j). \quad (2.14)$$

Similarly, we will stabilize the above computation by centering it to

$$V_n^{J_M, *} := \arg \min_{V \in \{V_n^j\}_{j \in J_M}} \mathcal{E}(V). \quad (2.15)$$

The random subset is typically chosen at each time step in the simulation.

Remark 4 *As a further randomization variant, at each time step, we may partition particles into disjoint subsets J_M^k , $k = 1, \dots, S$ of size M such that $SM = N$ and compute the evolution of each batch separately (see (Jin et al., 2020; Carrillo et al., 2021) for more details). Since the computational cost of the CBO method is linear, unlike (Albi and Pareschi, 2013; Jin et al., 2020; Lorenzo Pareschi and Visconti., 2020) these randomization techniques can accelerate the simulation process (and eventually improve the particles exploration dynamic thanks to additional stochasticity), but do not reduce the overall asymptotic cost $\mathcal{O}(N)$.*

FAST METHOD

Using a constant number of particles is not the most efficient way to simulate the trend towards equilibrium of a system, typically because we can use some (deterministic) information on the steady state to speed up the method. In the case of CBO methods, asymptotically the variance of the system tends to vanish because of the consensus dynamics, see Proposition 12. So, we may accelerate the simulation by discarding particles in time accordingly to the variance of the system (Albi and Pareschi, 2013). This also influences the computation of $V_n^{\alpha, \mathcal{E}}$ by increasing the randomness and reducing the possibilities to get trapped in a local minimum. For a set of N_n particles we define the empirical variance at time $t^n = n\Delta t$ as

$$\Sigma_n = \frac{1}{N_n} \sum_{j=1}^{N_n} (V_n^j - \bar{V}_n)^2, \quad \bar{V}_n = \frac{1}{N_n} \sum_{j=1}^{N_n} V_n^j.$$

When the trend to consensus is monotone, that is $\Sigma_{n+1} \leq \Sigma_n$, we can discard particles uniformly in the next time step $t^{n+1} = (n+1)\Delta t$ accordingly to the ratio $\Sigma_{n+1}/\Sigma_n \leq 1$, without affecting their theoretical distribution. One way to realize this is to define the new number of particles as

$$N_{n+1} = \left\lfloor N_n \left(1 + \mu \left(\frac{\hat{\Sigma}_{n+1} - \Sigma_n}{\Sigma_n} \right) \right) \right\rfloor \tag{2.16}$$

where $\lfloor \cdot \rfloor$ denotes the integer part, $\mu \in [0, 1]$ and

$$\hat{\Sigma}_{n+1} = \frac{1}{N_n} \sum_{j=1}^{N_n} (V_{n+1}^j - \hat{V}_{n+1})^2, \quad \hat{V}_{n+1} = \frac{1}{N_n} \sum_{j=1}^{N_n} V_{n+1}^j.$$

For $\mu = 0$ we have the standard algorithm where no particles are discarded whereas for $\mu = 1$ we achieve the maximum speed up. We implement the details of the method, which includes the speed-up techniques just discussed, in Algorithm 2. As before we fix $\lambda = 1$.

Typically, a minimum bound N_{min} of the number of particles is adopted to guarantee that $N_n \geq N_{min}$ during the simulation and the variance reduction test is performed every fixed amount of iterations to avoid fluctuations effects.

ADAPTIVE PARAMETERS

Our main theoretical result Theorem 6 and condition (3.6) establish that, once N is large, for σ small enough and α large enough, Algorithm 1 will converge near to a global minimizer. One important aspect, as in many metaheuristic algorithms, concerns the choice of the parameters in the method. The adaptation of hyperparameters in multi-particle optimization is a well-known problem, which deserves a proper discussion, see, e.g., (Escalante et al., 2009). In our case, we observed that decreasing σ and increasing α during the iterative process leads to improved results in term of convergence and accuracy. One strategy, therefore, would be to start with a large σ and to reduce it progressively over time as a function of a suitable indicator of convergence, for example the average variance of the solution or the relative variation of V_α over time. This can be realized starting from σ_0 and

Algorithm 2: Fast KV-CBO

- Input:** $\Delta t, \sigma, \alpha, d, N, n_T, \mu, M$ and the function $\mathcal{E}(\cdot)$
- 1 Generate $V_0^i, i = 1, \dots, N_0$ sample vectors uniformly on \mathbb{S}^{d-1} ;
 - 2 Compute the variance Σ_0 of V_0^i and set $N_0 = N$;
 - 3 **for** $n = 0$ **to** n_T **do**
 - 4 Generate ΔB_n^i independent normal random vectors $N(0, \Delta t)$;
 - 5 Compute $V_n^{\alpha, \mathcal{E}}$ from (2.14) if $M \leq N_n$ otherwise use (2.12);
 - 6 $\tilde{V}_{n+1}^i \leftarrow V_n^i + \Delta t P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1) V_n^i$
 $V_{n+1}^i \leftarrow \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i|, i = 1, \dots, N_n$;
 - 7 Compute the quantity $\hat{\Sigma}_{n+1}$ from V_{n+1}^i ;
 - 8 Set $N_{n+1} \leftarrow \llbracket N_n \left(1 + \mu \left(\frac{\hat{\Sigma}_{n+1} - \Sigma_n}{\Sigma_n} \right) \right) \rrbracket$ and discard uniformly $N_n - N_{n+1}$ samples;
 - 9 Compute the variance Σ_{n+1} of V_{n+1}^i ;
 - 10 **end**
-

by decreasing it as

$$\sigma_{n+1} = \frac{\sigma_n}{\tau}, \quad (2.17)$$

where $\tau > 1$ is a constant. Other techniques, of course, can be used to decrease σ , for example following a cooling strategy as in the Simulated Annealing approach (Holley and Stroock, 1988). In (Carrillo et al., 2021) it has been proposed to reduce σ independently of the solution behavior, as a function of the initial value σ_0 and the number of iterations. This corresponds to take $\sigma_{n+1} = \sigma_n / (\sigma_0 \log(n+1))$ in (2.17). As a result of these strategies, the noise level in the system will decrease in time. Note that, since we need $\lambda \gg \sigma^2(d-1)$ (see formula (3.6) below) to achieve consensus in the system, this approach allows to start initially with a larger σ which permits to explore the surrounding area well before entering the consensus regime.

Similarly, it might not be beneficial to start with a large α from the beginning. In fact, in this case the V_α would right away equal the particle with the lowest energy and all the other particles will be forced to move towards this particle, with a lower impact on the initial exploration mechanism. Therefore, we can start with an initial value α_0 and gradually increase it to a maximum value α_{\max} accordingly to an appropriate convergence indicator, or independently as a function of the number of iterations. In particular, large values of α at the end of the simulation process are essential to achieve high accuracy in the computation of the minimum.

2.3 Numerical experiments for the Ackley function

MINIMIZING THE ACKLEY FUNCTION IN DIMENSION $d = 3$

First we consider the behavior of the model and its mean field limit in the case $d = 3$ for computing the minimum of the Ackley function² constrained over the sphere

$$\mathcal{E}(V) = -A \exp \left(-a \sqrt{\frac{b^2}{d} \sum_{k=1}^d (V_k - v_k^*)^2} \right) - \exp \left(\frac{1}{d} \sum_{k=1}^d \cos(2\pi b(V_k - v_k^*)) \right) + e + B, \quad (2.18)$$

with $A = 20$, $a = 0.2$, $b = 3$, $B = 20$ and $V = (V_1, \dots, V_d)^T$ with $|V| = 1$.

The global minimum is attained at $V = v^*$. In Figure 1 we report the Ackley function for $d = 3$ over the half sphere $V_3 \geq 0$. Note that, this problem differs from the standard minimization of the Ackley function over the whole space \mathbb{R}^d since KV-CBO operates through unitary vectors over the hypersphere.

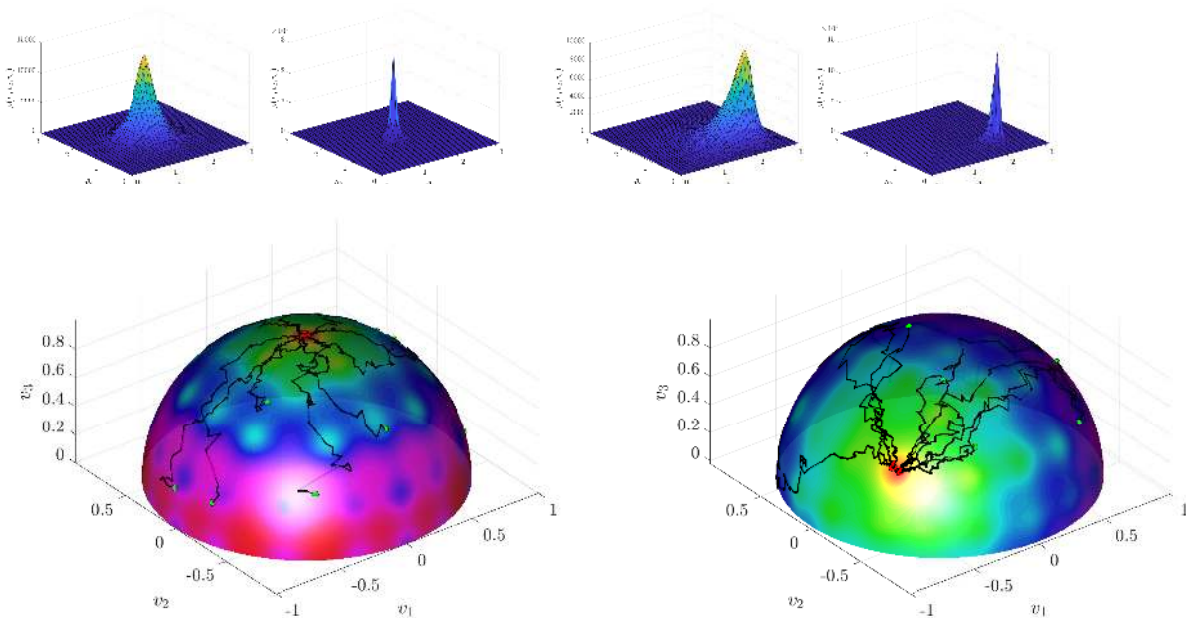


Figure 2: Particles trajectories along the simulation for the Ackley function in the case $d = 3$, $N = 20$ with minimum at $v^* = (0, 0, 1)^T$ (left) and $v^* = (-1/\sqrt{2}, -1/2, 1/2)^T$ (right). On the top corresponding time evolution of the particle distribution $\rho(v, t)$ in angular coordinates at $t = 1$ and $t = 2.5$ for $N = 10^6$. The simulation parameters are $\Delta t = 0.05$, $\sigma = 0.25$ and $\alpha = 50$.

In all our simulations we initialize the particles with a uniform distribution over the half sphere characterized by $V_3 \geq 0$ and employ the simple Euler-Maruyama scheme with

2. https://en.wikipedia.org/wiki/Test_functions_for_optimization

projection. We report in Figure 2 the particle trajectories for $t \in [0, 5]$ in the case of $N = 20$, $\Delta t = 0.05$, $\sigma = 0.25$ and $\alpha = 50$. On the left we consider the case with minimum at $v^* = (0, 0, 1)^T$, on the right the case with minimum at $v^* = (1/\sqrt{2}, -1/2, 1/2)^T$. The time evolution of the particle distribution $\rho(v, t)$ in the numerical mean field limit for $N = 10^6$ is also reported in the upper part of the same figure.

Next in Figure 3, we consider the convergence to consensus measured using various indicators for $N = 50$, $\Delta t = 0.1$, $v^* = (0, 0, 1)^T$ and various values of σ and α . The results have been averaged 1000 times with a success rate of 100% in all test cases considered. Following (Carrillo et al., 2021; Pinnau et al., 2017), we consider a run successful if $V_n^{\alpha, \mathcal{E}}$ at the final time is such that

$$\|V_{n_T}^{\alpha, \mathcal{E}} - v^*\|_\infty := \sup_{k=1, \dots, d} |(V_{n_T}^{\alpha, \mathcal{E}})_k - (v^*)_k| \leq 1/4.$$

We also compute the expected error in the computation of the minimum by considering time averages of $\|V^{\alpha, \mathcal{E}} - v^*\|_\infty$ and we report the quantity $|V^{\alpha, \mathcal{E}} - v^*|^2/d$ used in (Carrillo et al., 2021; Pinnau et al., 2017). As can be seen from Figure 3 (top) where $\sigma = 0.7$ the influence of large values α in the accuracy of the computation of the minimum is clear when passing from $\alpha = 5$ to $\alpha = 500$.

In Figure 3 (bottom) we show the same computations for a larger value $\sigma = 2$ of the diffusion coefficients, which violate the consensus bound $\sigma^2(d-1) \ll \lambda$, see (3.6). We compare our results with the ones computed using the CBO method in (Pinnau et al., 2017). Even if both methods yield a success rate of 100%, the methods clearly do not reach consensus, in the sense that the consensus error (2.9) is not diminishing in time. This behavior is common also to the CBO solvers in (Carrillo et al., 2021) where the above quantity may even diverge since it is not bounded by the geometry of the sphere.

MINIMIZING THE ACKLEY FUNCTION IN DIMENSION $d = 20$

Next we consider the more difficult case of the Ackley function in dimension $d = 20$.

In Table 1 we report the results for $\sigma = 0.3$, $\Delta t = 0.05$, $\alpha = 5 \times 10^4$, $T = 100$ and various values of N and M . The rate of success and the expectation of the error $|V^{\alpha, \mathcal{E}} - v^*|^2$ have been measured over 100 runs and the minimum has been considered in two different positions

$$v^* = (0, \dots, 0, 1)^T, \quad v^* = (d^{-1/2}, \dots, d^{-1/2})^T.$$

In the first case the minimum is at the center of our initial distribution (so $V^{\alpha, \mathcal{E}}$ initially is not too far from v^*) whereas the second choice is more difficult for the CBO solver since the minimum is shifted with respect to the center of the initial particle distribution, uniformly in all coordinates.

In all test cases considered the success rate is close to 100%. In particular, let us observe (see Table 2) that the fast method for $\mu = 0.3$ and $\mu = 0.2$ with $N_{min} = 10$ permits to achieve better performances for a given computational cost. We have selected a final computation time lower than the optimal computation time that would have allowed us to achieve maximum precision in the computation of the minimum, this to avoid unnecessary iterations with a small number of particles that would have created a bias in the final average particle number N_{avg} .




Figure 3: Behavior of various convergence indicators in time for the Ackley function in the case $d = 3$ for $N = 50$, $\Delta t = 0.1$. The two graphs on top show the accuracy of KV-CBO for $\sigma = 0.7$, which satisfies $\sigma^2(d-1) \ll \lambda = 1$ as dictated by (3.6). We chose $\alpha = 5$ (top, left) and $\alpha = 500$ (top, right); the yellow line on the top right plot is superimposed by the green line. It is seen that the accuracy is much better for the choice $\alpha = 500$. For the two graphs on the bottom we chose $\sigma = 2$, which violates $\sigma^2(d-1) \ll \lambda$, and used KV-CBO with $\alpha = 30$ (bottom, left) and the CBO method from (Pinnau et al., 2017) with $\alpha = 30$ (bottom, right). Again, the green line is superimposing the yellow line. The results have been averaged 1000 times with a success rate of 100% in all cases.

Table 1: Ackley function in $d = 20$: $\mu = 0$, $\sigma = 0.3$, $\Delta t = 0.05$ and $T = 100$

v^*		$N = 50$	$N = 100$	$N = 200$
		$M = 40$	$M = 70$	$M = 100$
$(0, \dots, 0, 1)^T$	Rate	100%	100%	100%
	Error	$2.24118e - 08$	$1.3364e - 09$	$3.51083e - 09$
$(d^{-1/2}, \dots, d^{-1/2})^T$	Rate	98%	99%	100%
	Error	$1.15704e - 06$	$1.476e - 09$	$5.09216e - 09$

Table 2: Ackley function in $d = 20$: $\sigma = 0.3$, $\Delta t = 0.05$ and $T = 100$

v^*		$N_0 = 100$	$N_0 = 200$	$N_0 = 400$
		$M = 70$	$M = 100$	$M = 150$
$(0, \dots, 0, 1)^T$ $\mu = 0.3$	Rate	100%	100%	100%
	Error	$1.20639e - 07$	$3.73419e - 08$	$2.24362e - 08$
	N_{avg}	21.6	38.7	71.4
$(d^{-1/2}, \dots, d^{-1/2})^T$ $\mu = 0.2$	Rate	100%	100%	100%
	Error	$1.34745e - 06$	$2.02787e - 08$	$8.06536e - 09$
	N_{avg}	27.3	53.1	103.0

2.4 Challenging applications in signal processing and machine learning

In this section we consider two applications of KV-CBO, namely, the phase retrieval problem and the robust subspace detection problem. For the former we consider only synthetic data, for the latter we consider synthetic as well as real-life data in dimension up to $d = 2880$. The solution to these problems can be reformulated in terms of a high dimensional nonconvex optimization over the sphere with unique symmetric solutions. Both these problems have by now *ad hoc* state of the art methods for their solution. The aim of this section is to demonstrate that Algorithms 1 or 2 can be used in a versatile and scalable way to solve several and diverse problems and achieve state of the art performances by comparison with the more specific methods.

2.4.1 PHASE RETRIEVAL

Recently there has been growing interest in recovering an input vector $z^* \in \mathbb{R}^d$ from quadratic measurements

$$y_i = |\langle z^*, a_i \rangle|^2 + w_i, \quad i = 1, \dots, M \quad (2.19)$$

where w_i is adversarial noise, and a_i are a set of known vectors. Since only the magnitude of $\langle z^*, a_i \rangle$ is measured, and not the phase (or the sign, in the case of real valued vectors), this problem is referred to as phase retrieval. Phase retrieval problems arise in many areas of optics, where the detector can only measure the magnitude of the received optical

wave. Important applications of phase retrieval include X-ray crystallography, transmission electron microscopy and coherent diffractive imaging (Quiney, 2010; Hurt; Harrison, 1993; Walther, 1963). Several algorithms have been devised for robustly computing z^* from measured information $y = (y_i)_{i=1,\dots,M}$ based on different principles, such as alternating projections, lifting and convex relaxation, and simple gradient descent for empirical risk minimization (Gerchberg, 1972; Fienup, 1982; Yang et al., 1994; Candés et al., 2013; Candés et al., 2015; Chen et al., 2019). Despite the wide range of solutions, most of these algorithms fail to tackle robustly the crystallographic problem which is both the leading application and one of the hardest forms of phase retrieval (Elser et al., 2018). One of the reasons is that the phase retrieval problem is intrinsically ill-posed for M small. Recent work (Mondelli and Montanari, 2019) explains even by information theoretical arguments that no estimator can do better than a random estimator for $M \leq d - o(d)$. Uniqueness results of the solution z^* of the real-valued phase retrieval problem in the case of no noise has been established in (Balan et al., 2006) for sets of measurement vector $\{a_i : i = 1, \dots, M\}$ forming a frame for \mathbb{R}^d , i.e., there are constants $0 < A \leq B < \infty$ such that

$$A|z|^2 \leq \sum_{i=1}^M |\langle z, a_i \rangle|^2 \leq B|z|^2 \tag{2.20}$$

holds for any $z \in \mathbb{R}^d$. Specifically, (Balan et al., 2006, Theorem 2.2) ensures that for generic frames unique identifiability occurs for $M \geq 2d - 1$, as the map $\mathbb{R}^d \setminus \{\pm 1\} \ni z \rightarrow y(z) := (|\langle z, a_i \rangle|^2)_{i=1,\dots,M}$ is in fact injective. In order to tackle the robust identifiability, empirical risk minimization has been considered in (Eldar and Mendelson, 2014), i.e., the minimization of the discrepancy

$$\mathcal{E}(z) = \sum_{i=1}^M \left| |\langle z, a_i \rangle|^2 - y_i \right|^2. \tag{2.21}$$

Guarantees of stable reconstruction via empirical risk minimization are obtained under the assumption that the measurements vectors $\{a_i : i = 1, \dots, M\}$ fulfill the stability property

$$\sum_{i=1}^M \left| |\langle z, a_i \rangle|^2 - |\langle \hat{z}, a_i \rangle|^2 \right| \geq \kappa \|z - \hat{z}\| \|z + \hat{z}\|, \tag{2.22}$$

for all $z, \hat{z} \in \mathbb{R}^d$ and some fixed $\kappa > 0$. In particular, (Eldar and Mendelson, 2014, Theorem 2.4) ensures that for measurement vectors $\{a_i : i = 1, \dots, M\}$ generated at random, e.g., as i.i.d. Gaussian vectors, for $M \geq \gamma d$, the stability estimate (2.22) holds for a suitable $\kappa > 0$ with high probability depending on the constant $\gamma > 0$. As a broad disquisition about the phase retrieval problem is not the focus of this paper, we omit here details about stability under adversarial noise and we refer to (Bandeira et al., 2014; Eldar and Mendelson, 2014) for further insights. However, we should notice at this point that the empirical risk \mathcal{E} in (2.21) fulfills then all the requests of Assumptions 3.1 below, in particular the stability estimate (2.22) naturally induces the *inverse continuity property* 4. of Assumptions 3.1. Hence, the minimization of (2.21) is a challenging nonconvex optimization problem, which falls precisely in the realm of problems for which Algorithm 1 or Algorithm 2 are expected

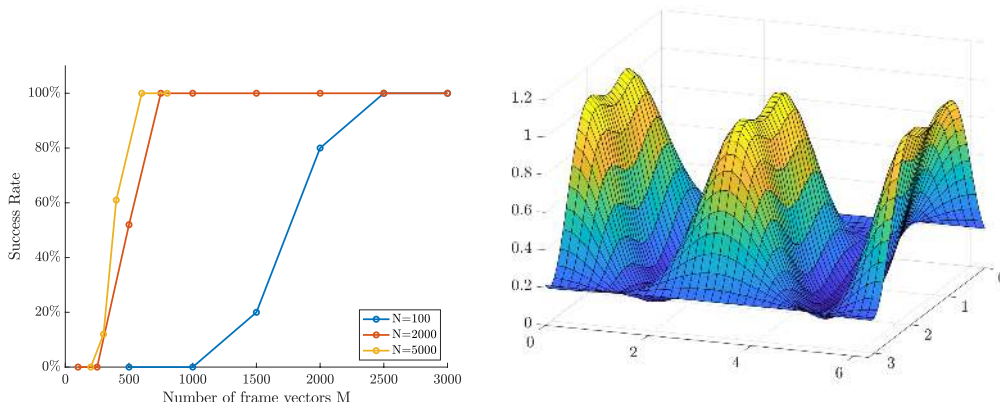


Figure 4: Left: Success rate for different numbers of frame vectors M in dimension $d = 100$. We have used the following parameters: $\alpha = 2 \cdot 10^{15}$, $\lambda = 1$, $\Delta t = 0.1$ and $\sigma = 0.11$ for $N = 100$ and $\sigma = 0.08$ for $N = 2000$. Right: Plot of the energy defined in (2.21) in $d = 2$. It is evident that the energy may exhibit saddle points, but no spurious minimizers appear. This is the reason for a vanilla gradient descent method to work so well for such a problem (Candes et al., 2015; Chen et al., 2019; Lee et al., 2019).

to work at best. Before presenting numerical experiments of the use of Algorithm 1 or Algorithm 2 and comparisons with state of the art methods, we should perhaps clarify that the empirical risk minimization can without loss of generality be restricted to vectors on the sphere as soon as the lower frame constant A is known: for the sake of simplicity, let us assume again that the noise $w \equiv 0$ and we observe that

$$\sum_{i=1}^M y_i = \sum_{i=1}^M |\langle z^*, a_i \rangle|^2 \geq A |z^*|^2 \quad \text{and} \quad |z^*| \leq \sqrt{\frac{1}{A} \sum_{i=1}^M y_i} =: R \quad (2.23)$$

where we take A to be the optimal lower frame bound. We define the vectors \tilde{a}_i by one zero padding, i.e.,

$$\tilde{a}_i = [a_i, 0] \in \mathbb{R}^{d+1}, \quad (2.24)$$

and we further denote

$$\tilde{z}^* = [z^*, \sqrt{R^2 - |z^*|^2}] \in R \mathbb{S}^d, \quad v^* = \frac{\tilde{z}^*}{R} \in \mathbb{S}^d, \quad \text{and} \quad \tilde{y}_i = \frac{y_i}{R^2}. \quad (2.25)$$

With these notations, (2.19) can be equivalently recast in the form

$$\tilde{y}_i = |\langle v^*, \tilde{a}_i \rangle|^2, \quad v^* \in \mathbb{S}^d.$$

Hence, the unconstrained minimization of \mathcal{E} can be equivalently solved by the constrained minimization of

$$\tilde{\mathcal{E}}(v) := \sum_{i=1}^M \left| |\langle v, \tilde{a}_i \rangle|^2 - \tilde{y}_i \right|^2, \quad (2.26)$$

over the sphere \mathbb{S}^d . In fact, the first d components of the minimizing vector v^* must coincide with z^*/R . So from now on we implicitly assume that the problem is transformed into one of the type (1.1).

We tested KV-CBO for dimension $d = 100$ for the function defined in (2.26), where the vectors a_i are sampled from a uniform distribution over the sphere. We computed the success rate for reconstructing the vector z^* in terms of the number M of vectors a_i . We count a run as successful if the computed \bar{z} by Algorithm 1 or Algorithm 2 fulfills

$$\min\{|z^* - \bar{z}|, |z^* + \bar{z}|\} < 0.05. \tag{2.27}$$

The phase transitions of success recovery are shown in on the left-hand-side of Figure 4. We can observe that the success rate improves with the number N of particles used by Algorithm 1 or Algorithm 2 and best success is obtained by $M \geq \gamma d$ as predicted by theory. We notice that the optimization via KV-CBO is evidently not affected by the curse of dimension. On the right-hand-side we depict the typical cost function landscape with saddle-points and symmetric global minimizers.

In the following, we compare Algorithm 2 with three relevant state of the art methods for phase retrieval:

- Wirtinger Flow (fast gradient descent method) (Candes et al., 2015; Chen et al., 2019);
- Hybrid Input Output/Gerchberg-Saxton’s Alternating Projections (alternating projection methods) (Gerchberg, 1972; Fienup, 1982; Yang et al., 1994);
- PhaseMax/PhaseLamp (convex relaxation and its multiple iteration version) (Candés et al., 2013).

For the comparison we used the Matlab toolbox PhasePack³ (Chandra et al., 2017) and our own code⁴.

In Figure 5 we demonstrate on the left that KV-CBO is exactly as robust as Wirtinger Flow with respect to adversarial noise and on the right we compare phase transition diagrams of success rate, which show that KV-CBO has a slight delay in perfect recovery with respect to Wirtinger Flow and PhaseMax/PhaseLamp, but it is comparable with Hybrid Input Output/Gerchberg-Saxton’s Alternating Projections. The delayed perfect recovery indirectly confirms that the *inverse continuity property* 4. of Assumptions 3.1 needs to be fulfilled for the method to work optimally. (We reiterate that if M is large enough, then the stability property (2.22) holds with high probability and as a consequence also the inverse continuity property.)

2.4.2 ROBUST SUBSPACE DETECTION

Let us consider a cloud of points $\mathcal{Q} = \{x^{(i)} \in \mathbb{R}^d : i = 1, \dots, M\}$ in an Euclidean space with $d \gg 1$. We assume without loss of generality that the point cloud is centered, that is, the mean of the point cloud is zero. Subspace detection is about finding a lower dimensional linear subspace $V \subset \mathbb{R}^d$ that fits the data at best, in the sense that the sum of

3. <https://www.cs.umd.edu/~tomg/projects/phasepack/>

4. <https://github.com/PhilippeSu/KV-CBO>

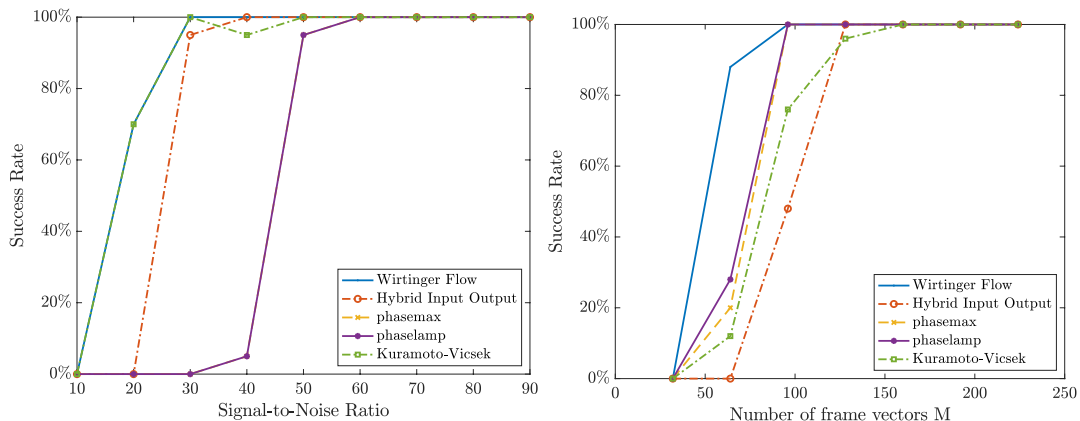


Figure 5: Left: Success rate in terms of the Signal-to-Noise Ratio in dimension $d = 32$ for $M = 4d$ Gaussian vectors. The green dashed curve representing KV-CBO is exactly superimposed with the light blue curve of the Wirtinger Flow. Right: Phase transitions for different numbers of Gaussian vectors M in dimension $d = 32$ (the yellow curve is superimposed by the purple curve). We used $\sigma = 0.2, \Delta t = 0.1, N = 10^4$ and chose the parameter α adaptively, with initial $\alpha_0 = 2000$ and final $\alpha_{max} = 1e15$. The results are averaged 25 times.

the squared norms of the orthogonal projection of the points $x^{(i)}$ to V^\perp is minimal. In the simplest case of a one-dimensional subspace, the cost function to be minimized is given by $\mathcal{E}_2(v) := \sum_{i=1}^M |x^{(i)} - \langle x^{(i)}, v \rangle v|^2$ where each summand is the squared norm of the orthogonal projection of one point $x^{(i)}$ to the space $\text{span}(v)^\perp$. It is well-known that the minimizer $v^* = \arg \min_{v \in \mathbb{S}^{d-1}} \mathcal{E}_2(v) = \arg \max_{v \in \mathbb{S}^{d-1}} |Xv|^2$ represents the direction of maximal variance of the point cloud, see, e.g., Figure 6 (left), and coincides with the right singular vector associated to the operator norm of the matrix $X = (x^{(i)T})_{i=1, \dots, M}$ whose rows are the vectors $x^{(i)}$'s. Despite the nonconvexity of the cost, the computation of the best fitting subspace can be conveniently done by singular value decomposition (SVD) also for subspaces of higher dimension. In this case the cost would simply read $\mathcal{E}_2(V) := \sum_{i=1}^M |P_{V^\perp} x^{(i)}|^2$. The drawback of the energy $\mathcal{E}_2(v)$ is the fact that it is quadratic, thus the summand $|x^{(i)} - \langle x^{(i)}, v \rangle v|^2$ will be particularly large if $x^{(i)}$ is an outlier, far from the subspace where most of the other points may cluster. The aim of *robust* subspace detection (Lerman et al., 2015; Lerman and Maunu, 2018; Maunu et al., 2019) is finding the principal direction of a point cloud without assigning too much weight to outliers. We therefore introduce the more general energy

$$\mathcal{E}_p(V) := \sum_{i=1}^M |P_{V^\perp} x^{(i)}|^p, \quad V \subset \mathbb{R}^d, \quad \dim(V) = k \ll d, \quad (2.28)$$

where $0 < p \leq 2$. Even in the simplest one dimensional case, the minimization of the energy

$$\mathcal{E}_p(v) := \sum_{i=1}^M |x^{(i)} - \langle x^{(i)}, v \rangle v|^p = \sum_{i=1}^M (|x^{(i)}|^2 - |\langle x^{(i)}, v \rangle|^2)^{p/2}, \quad v \in \mathbb{S}^{d-1},$$

turns out for $0 < p \ll 2$ to be a rather nontrivial nonconvex optimization problem. On the right of Figure 7, Figure 8, and Figure 9 we illustrate some cost function landscapes in dimension $d = 2$. One can immediately notice how \mathcal{E}_p becomes in fact rougher and exhibits all of the sudden several spurious local minimizers (compare with the case of $p = 2$ in Figure 6). Hence, the success of a simple gradient descent method is far less obvious than for the phase retrieval problem, where the energy may have saddle-points, but it has generically no local minimizers, see Figure 4 and refer to (Candes et al., 2015; Chen et al., 2019; Lee et al., 2019) for details.

In the following we test KV-CBO for clouds of synthetic data points and a cloud of real-life photos from the *10K US Adult Faces Database* (Bainbridge et al., 2013). We discuss the performance of the method both for $p = 2$ and $0 < p < 2$. In the former case, we can compute the exact minimizer of the energy $\mathcal{E}_2(v)$ by SVD. For $0 < p < 2$ we compare the result with the state of the art algorithm Fast Median Subspace (FMS) (Lerman and Maunu, 2018) as benchmark. We mention that FMS is proven in general to converge to stationary points of the cost function only, which are in special data models very close to global minimizers with high probability. The synthetic point cloud models we use for comparison below are in part fitting the existing guarantees of global optimization for FMS. In these cases, we analyze different sets of parameters and dimensionality of the problem and we discuss the success rate for different parameters such as numbers of particles and $\sigma > 0$. In fact, the choice of the parameter $\sigma > 0$ is perhaps a bit tricky. From our theoretical findings, it would be sufficient that $\lambda/(d-1) \gg \sigma^2$, see (3.6), thus $\sigma > 0$ needs simply to decrease with growing dimension d . However, in the pure particle simulation σ cannot be taken too small otherwise randomness won't be enough to explore the space in a reasonable computational time. In Figure 10 we report the success rate in terms of σ for different dimensions. We further chose $\alpha = 2 \cdot 10^{15}$ and $\Delta t = 0.25$.

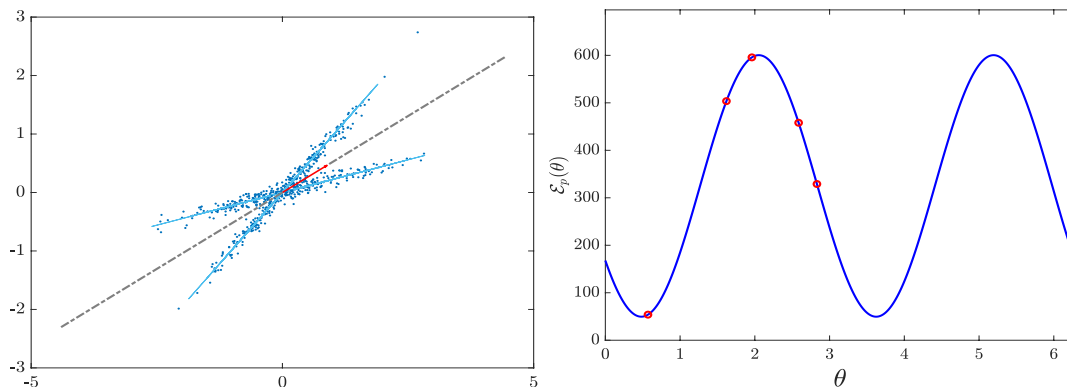


Figure 6: Left: Point cloud with $N_{sp} = 2$ one-dimensional subspaces in dimension $d = 2$ with Gaussian noise of 0.01. The red vector shows the principal direction. Right: Energy $\mathcal{E}_2(v(\theta))$ for the point cloud on the left for $\theta \in [0, 2\pi)$, where $v(\theta)$ is expressed in polar coordinates. The particles are shown in red.




Figure 7: Left: Point cloud in dimension $d = 2$ with $N_{sp} = 2$ subspaces with 100 points each and Gaussian noise of 0.25. We further have added 500 outliers. The orange/ red vector shows the principal component computed by SVD of the point cloud with/ without the outliers. The green vector is the principal component computed by KV-CBO. Right: Energy $\mathcal{E}_p(\theta)$ for $\theta \in [0, 2\pi)$ for the point cloud on the left and $p = 1$. The particles are shown in red, $V_0^{\alpha, \mathcal{E}}$ is shown in green.




Figure 8: Left: Point cloud in dimension $d = 2$ with Gaussian noise of 0.25 on a one-dimensional subspace with 100 points and 250 outliers. The orange/ red vector shows the principal component computed by SVD of the point cloud with/ without the 250 outliers. The green vector is the principal component computed by KV-CBO. Right: Energy $\mathcal{E}_p(\theta)$ for $\theta \in [0, 2\pi)$ for the point cloud on the left and $p = 0.5$. The particles are shown in red, $V_0^{\alpha, \mathcal{E}}$ is shown in green.




Figure 9: Left: Point cloud with $N_{sp} = 2$ one-dimensional subspaces with 100 points on the first subspace and 40 points on the second with Gaussian noise of 0.01. Further, we have added 100 outliers. The red vector shows the principal direction for the point cloud without the outliers. The green vectors shows the direction computed by KV-CBO. It matches the one-dimensional cluster with 100 points. Right: Energy $\mathcal{E}_p(\theta)$ for $\theta \in [0, 2\pi)$ and $p = 0.2$ for the point cloud on the left. The particles are shown in red. The initial $V_0^{\alpha, \mathcal{E}}$ is shown in green, superimposing the particle with the smallest energy.

SYNTHETIC DATA

In this section we discuss numerical tests for synthetic point clouds in dimensions up to $d = 200$ for $p = 2$ and $0 < p < 2$. In Figures 6 to 9 we report plots of energies in $d = 2$ for different values of p .

We test the method for point clouds laying on $N_{sp} = 25$ nearly parallel one dimensional subspaces and point clouds laying $N_{sp} = 25$ randomly chosen subspaces, each with Gaussian noise of 0.01. The latter point clouds do not have an obvious principal direction, as opposed to the case of nearly parallel subspaces (see Figure 10 on the right). In this case a larger number of initial particles is needed to find the minimizer.

Case $p = 2$

For the case $p = 2$ we compare the minimizer $V_{n_T}^{\alpha, \mathcal{E}}$ computed by KV-CBO with the minimizer v^* computed by SVD. In Figure 11 we plot the average error $|V_n^{\alpha, \mathcal{E}} - v^*|$ for $n = 0, \dots, n_T$ for 25 runs. In the plot on the right we show the success rate for different numbers of particles for different dimensions. We count a run as successful if

$$\min\{|V_{n_T}^{\alpha, \mathcal{E}} - v^*|, |V_{n_T}^{\alpha, \mathcal{E}} + v^*|\} \leq 0.01 \quad (2.29)$$

where n_T is the final time step. We observe that for point clouds with nearly parallel one-dimensional subspaces, a very small number of particles already yields good results. For the point clouds with randomly chosen one-dimensional subspaces, corresponding to a flatter spectrum, the number of particles N has to be chosen larger in order to obtain good results.

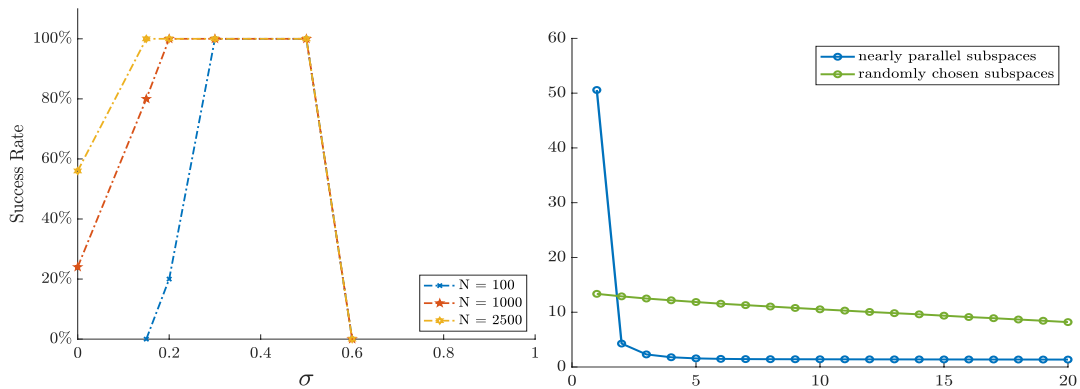


Figure 10: Left: Success rate for different values of σ and N in dimensions $d = 10$. Choosing a larger number of particles clearly widens the window from which σ can be chosen. Note that for $\sigma = 0$ (deterministic Kuramoto-Vicsek model) we have a success rate of nearly 60% for $N = 2500$ particles. Right: Singular value decay of the point cloud with $N_{sp} = 25$ nearly parallel (blue) and randomly chosen one-dimensional subspaces (green).

Still, KV-CBO can certainly be considered an interesting, robust, and efficient alternative method for computing SVD's.

Dimension	$d = 10$ $d = 100$ $d = 200$			
		$N_0 = 1000$	$N_0 = 2500$	$N_0 = 5000$
Relative Error $\leq 10^{-2}$	Rate	100%	100%	100%
$\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) \leq \mathcal{E}_{p,0}(V_{FMS})$	Rate	63%	13%	0%
	Absolute Error	$2.8413e - 12$	$2.0026e - 12$	—
	Relative Error	$7.6774e - 15$	$5.8669e - 15$	—
$\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) > \mathcal{E}_{p,0}(V_{FMS})$	Rate	37%	87%	100%
	Absolute Error	$2.9066e - 12$	$71218e - 12$	$1.3387e - 11$
	Relative Error	$7.9272e - 15$	$2.0745e - 14$	$3.9490e - 14$

Table 3: Numerical comparison of KV-CBO and the FMS method for a point cloud with $N_{sp} = 25$ nearly parallel one-dimensional subspaces with Gaussian noise of 0.01. The relative error is defined in (2.31). The results are averaged over 100 runs.

Dimension	$d = 10$			$d = 100$			$d = 200$		
	$N_0 = 1000$			$N_0 = 2500$			$N_0 = 5000$		
Relative Error $\leq 10^{-2}$	Rate	100%	100%	100%	100%	100%	100%	100%	100%
$\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) \leq \mathcal{E}_{p,0}(V_{FMS})$	Rate	79%	14%	14%	15%	15%	15%	15%	15%
	Absolute Error	0.0814	4.6205	4.6205	5.9582	5.9582	5.9582	5.9582	5.9582
	Relative Error	$4.4532e - 5$	0.0024	0.0024	0.0031	0.0031	0.0031	0.0031	0.0031
$\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) > \mathcal{E}_{p,0}(V_{FMS})$	Rate	21%	86%	86%	85%	85%	85%	85%	85%
	Absolute Error	0.3312	0.5628	0.5628	1.4965	1.4965	1.4965	1.4965	1.4965
	Relative Error	$1.8466e - 4$	$2.8669e - 4$	$2.8669e - 4$	$7.6525e - 4$	$7.6525e - 4$	$7.6525e - 4$	$7.6525e - 4$	$7.6525e - 4$

Table 4: Numerical comparison of KV-CBO and the FMS method for a point cloud with $N_{sp} = 25$ randomly chosen one-dimensional subspaces with Gaussian noise of 0.01. The relative error is defined in (2.31). The results are averaged over 100 runs.

Case $p = 1$

For $0 < p < 2$ the energy $\mathcal{E}_p(v)$ is not smooth enough to fulfill the regularity conditions of Assumptions 3.1 below. In order to fit the experiment to our theoretical findings, we may consider the smoothed energy

$$\mathcal{E}_{p,\delta}(v) = \sum_{i=1}^M (|x_i - \langle x_i, v \rangle v|^2 + \delta^2)^{p/2} \quad (2.30)$$

where we chose $\delta = 10^{-7}$ (as $\delta > 0$ is chosen so small, it is actually irrelevant from a numerical precision point view). We again test KV-CBO on synthetic point clouds with $N_{sp} = 25$ one-dimensional subspaces with 100 points each, thus $M = 2500$. We run the experiment 100 times in dimension $d \in \{10, 100, 200\}$ and count one run as successful if the relative error of the function values is less than 1%, that is,

$$\frac{|\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) - \mathcal{E}_{p,0}(V_{FMS})|}{\min\{\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}), \mathcal{E}_{p,0}(V_{FMS})\}} \leq 10^{-2}, \quad (2.31)$$

where V_{FMS} denotes the minimum of $\mathcal{E}_{p,0}(v)$ computed by the FMS method. We note that $V_{n_T}^{\alpha,\mathcal{E}}$ is the minimizer of the function $\mathcal{E}_{p,\delta}(v)$ for $\delta \neq 0$ computed by KV-CBO. We further report the average absolute and relative errors of the function values for the runs for which $\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) \leq \mathcal{E}_{p,0}(V_{FMS})$ as well as $\mathcal{E}_{p,0}(V_{n_T}^{\alpha,\mathcal{E}}) > \mathcal{E}_{p,0}(V_{FMS})$. In the stopping criterion for KV-CBO (2.9) we chose $\varepsilon = 10^{-10}$, as maximal amount of iterations $n_T = 10^4$, and use Algorithm 2 to speed up the method. For the FMS method we chose $\varepsilon = 10^{-10}$ and $n_T = 10^2$, as FMS method converges to a good minimizer after fewer iterations than KV-CBO. In Tables 3 and 4 we show that (2.31) is fulfilled in 100% of the cases. In other words: KV-CBO and state of the art FMS perform equally good on point clouds with nearly parallel

one-dimensional subspaces as well as randomly chosen one-dimensional subspaces. For the former the maximal relative error is in the order of 10^{-14} in dimension $d = 200$.

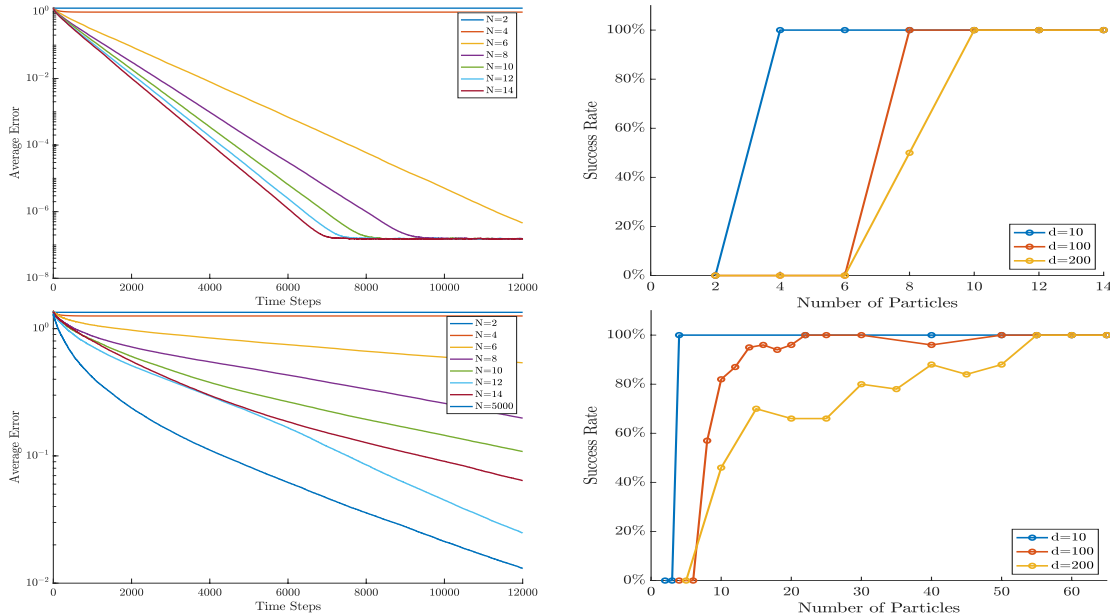


Figure 11: Average error for a point cloud with $N_{sp} = 25$ nearly parallel (top left) respectively randomly chosen (bottom left) one-dimensional subspaces in dimension $d = 200$ with Gaussian noise of 0.1 for different numbers of particles. We chose $\alpha = 2 \cdot 10^{15}$, $\sigma = 0.08$, $\Delta t = 0.25$. The curve for $N = 5000$ particles has been calculated with Algorithm 2. Right: success rate for the same point clouds in dimension $d \in \{10, 100, 200\}$. The results have been averaged 25 times and we count one run as successful if $|V_{n_T}^{\alpha, \mathcal{E}} - v^*| \leq 0.01$.

ROBUST COMPUTATION OF EIGENFACES

In this section we discuss the numerical results of KV-CBO on real-life data. We chose a subset of $M = 421$ similar looking pictures of the *10K US Adult Faces Database* (Bainbridge et al., 2013). We converted this subset to gray scale images and reduced the size of each picture by factor 4. We finally extract at a subset of $M = 421$ pictures of size 64×45 , which yields a point cloud $X \in \mathbb{R}^{2880 \times 421}$.



Figure 12: Samples from the *10K US Adult Faces Database* (Bainbridge et al., 2013) and one instance of outlier.

The eigenfaces computed by SVD and KV-CBO are shown in Figure 13(a) and Figure 13(b). The computed eigenfaces are visually indistinguishable and the final error is in the order of 10^{-3} . We then added 6 outliers (pictures of different plants and animals on a white background) to the point cloud and again computed the eigenface by SVD (see Figure 14(c)) and by KV-CBO with $p = 1$ and $N = 5000$ particles (see Figure 13(d)). The eigenface computed by SVD still retain some features, but the difference to the original eigenface (without outliers) is clearly perceivable. Instead, the eigenface computed by the KV-CBO still looks very similar to the eigenface of the point cloud without outliers. We quantify the accuracy of the results by Peak Signal-to-Noise Ratio (see caption of Figure 13). We then added further 6 outliers (amounting to a total of 12 outliers) to the point cloud and again computed the eigenface by SVD (see Figure 13(e)) and KV-CBO with $p = 0.5$ and $N = 50000$ particles (see Figure 13(f)). The difference of both eigenfaces to the original eigenface (without outliers) is clearly visible. The eigenface computed by SVD lost most of the original features. On the other hand, the eigenface computed by KV-CBO still retains the main features. We reiterate that the energy landscape $\mathcal{E}_{p,\delta}(v)$ is much more complex for $0 < p < 1$ than for $p \in [1, 2]$ (see Figures 6 to 9). An increase of the number of particles N did not yield better results.




Figure 13: Eigenface for the point cloud of faces without outliers computed by SVD (a), and KV-CBO (b). Eigenface for point cloud with 6 outliers computed by SVD (c), and KV-CBO with $p = 1$ (d). Eigenface for point cloud with 12 outliers computed by SVD (e), and KV-CBO with $p = 0.5$ (f). We used the following parameters: $\alpha = 2 \cdot 10^{15}$, $\lambda = 1$, $\sigma = 0.019$, $\Delta t = 0.25$, $T = 25000$, $N = 5000$ and $N_{min} = 150$ (see algorithm 2) for (b) and (d). For (f) we used $p = 0.5$, $N = 50000$ and $N_{min} = 5000$. For $p < 2$ we used $\delta = 10^{-7}$. For the case (b) the error to the SVD eigenface was 0.00071. The Peak Signal-to-Noise Ratio is: 61.4214 for (a) and (b), 15.9764 for (a) and (c), 20.7344 for (a) and (d), 12.3109 for (a) and (e) and 14.2892 for (a) and (f).




Figure 14: Average error (blue) and variance (red) for the computation of the eigenface (b) in the figure above.

3. Global optimization guarantees

3.1 Main result

In this section, we address the convergence of the stochastic Kuramoto-Vicsek particle system (1.2) to global minimizers of some cost function \mathcal{E} . In view of the already established mean-field limit result (1.12), it is actually sufficient to analyze the large time behavior of the solution $\rho(t, v)$ to the corresponding mean-field PDE (1.10). Let us rewrite (1.10) as

$$\partial_t \rho_t = \Delta_{\mathbb{S}^{d-1}}(\kappa_t \rho_t) + \nabla_{\mathbb{S}^{d-1}} \cdot (\eta_t \rho_t), \quad (3.1)$$

where $\kappa_t := \frac{\sigma^2}{2} |v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 \in \mathbb{R}$ and $\eta_t := \lambda \langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - \lambda v_{\alpha, \mathcal{E}}(\rho_t) \in \mathbb{R}^d$. We also introduce the auxiliary self-consistent nonlinear SDE

$$d\bar{V}_t = \lambda P(\bar{V}_t) v_{\alpha, \mathcal{E}}(\rho_t) dt + \sigma |\bar{V}_t - v_{\alpha, \mathcal{E}}(\rho_t)| P(\bar{V}_t) dB_t - \frac{(d-1)\sigma^2}{2} (\bar{V}_t - v_{\alpha, \mathcal{E}}(\rho_t))^2 \frac{\bar{V}_t}{|\bar{V}_t|^2} dt, \quad (3.2)$$

with the initial data \bar{V}_0 distributed according to $\rho_0 \in \mathcal{P}(\mathbb{S}^{d-1})$. Here $\rho_t = \text{law}(\bar{V}_t)$ is also the unique solution of the PDE (3.1), see (Fornasier et al., 2020, Section 2.3). The well-posedness of (3.2) is shown in (Fornasier et al., 2020, Theorem 2.2). We now define the expectation and variance of ρ_t as

$$E(\rho_t) := \int_{\mathbb{S}^{d-1}} v d\rho_t(v) \quad V(\rho_t) := \frac{1}{2} \int_{\mathbb{S}^{d-1}} |v - E(\rho_t)|^2 d\rho_t(v). \quad (3.3)$$

In the following, we show that, under suitable smoothness requirements, see Assumptions 3.1 below, for any $\epsilon > 0$ there exists suitable parameters α, λ, σ and well-prepared initial

distributions ρ_0 such that for $T^* > 0$ large enough the expected value of the distribution $E(\rho_{T^*}) = \int v d\rho_{T^*}(v)$ is in an ϵ -neighborhood of a global minimizers v^* of \mathcal{E} . The convergence rate is exponential in time and the rate depends on the parameters $\epsilon, \alpha, \lambda, \sigma$ (see Proposition 12). As mentioned in the introduction, this approximation together with (1.12) and classical results of the convergence of numerical methods for SDE (Platen, 1999) yield the convergence of Algorithm 1. In particular, we shall address the proof of the main result Theorem 1 and of the quantitative estimate (1.15) at the end of this section.

In order to formalize the result we state our fundamental assumptions: Throughout this section, the objective function $\mathcal{E} \in \mathcal{C}^2(\mathbb{S}^{d-1})$ satisfies the following properties

Assumption 3.1

1. \mathcal{E} is bounded and $0 \leq \underline{\mathcal{E}} := \inf \mathcal{E} \leq \mathcal{E} \leq \sup \mathcal{E} =: \bar{\mathcal{E}} < \infty$;
2. $\|\nabla \mathcal{E}\|_\infty \leq c_1$;
3. $\max \{ \|\nabla^2 \mathcal{E}\|_\infty, \|\Delta \mathcal{E}\|_\infty \} \leq c_2$;
4. For any $v \in \mathbb{S}^{d-1}$ there exists a minimizer $v^* \in \mathbb{S}^{d-1}$ of \mathcal{E} (which may depend on v) such that it holds

$$|v - v^*| \leq C_0 |\mathcal{E}(v) - \underline{\mathcal{E}}|^\beta,$$

where β, C_0 are some positive constants.

While the assumptions 1.-3. are all automatically fulfilled as soon as smoothness is provided, requirement 4. - which we call *inverse continuity assumption* - is a bit more technical and needs to be verified, depending on the specific application. In Section 2.4.1 we provided the concrete example of the phase retrieval problem for which all the conditions are in fact verifiable. The request of smoothness is exclusively functional to the proof of well-posedness and mean-field limit (Fornasier et al., 2020) and the proof of asymptotic convergence. As a matter of fact Algorithm 1 and Algorithm 2 are implementable even if \mathcal{E} admits just pointwise evaluations, e.g., \mathcal{E} is just a continuous function with no further regularity. Below we denote $C_{\alpha, \mathcal{E}} = e^{\alpha(\bar{\mathcal{E}} - \underline{\mathcal{E}})}$ and $C_{\sigma, d} = \frac{(d-1)\sigma^2}{2}$.

Definition 5 For any given $T > 0$ and $\alpha > 0$, we say that the initial datum and the parameters are well-prepared if $\rho_0 \in \mathcal{P}_{ac}(\mathbb{S}^{d-1}) \cap L^2(\mathbb{S}^{d-1})$, and parameters $\lambda, d, \beta > 0$, $0 < \epsilon \ll 1$, $0 < \delta \ll 1$, $0 < \theta < \delta$ satisfy

$$C_{\alpha, \mathcal{E}}^{2 \max\{1, \beta\}} \left(V(\rho_0) + \frac{\lambda C_T}{\lambda \theta - 4C_{\alpha, \mathcal{E}} C_{\sigma, d}} \delta^{\frac{d-2}{4}} \right)^{\frac{1}{2} \min\{1, \beta\}} + \epsilon^\beta < \frac{\delta - \theta}{C^*}, \quad (3.4)$$

$$V(\rho_0) + \frac{\lambda C_T}{\lambda \theta - 4C_{\alpha, \mathcal{E}} C_{\sigma, d}} \delta^{\frac{d-2}{4}} \leq \min \left\{ T^{-1} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2, T^{-1} \lambda^{-2} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^4, \frac{3}{8} \right\} \quad (3.5)$$

and

$$\lambda \theta - 4C_{\alpha, \mathcal{E}} C_{\sigma, d} > 0, \quad (3.6)$$

where C_T is a constant depending only on λ, σ, T and $\|\rho_0\|_2$, and $C^* > 0$ is a constant depending only on c_1, β, C_0 (c_1, β, C_0 are used in Assumption 3.1). Both C_T and C^* need to be subsumed from the proof of Proposition 12 and they are both dimension independent.

We shall prove first the following result.

Theorem 6 *Let us fix $\varepsilon_1 > 0$ small and assume that the initial datum ρ_0 and parameters $\{d, \beta, \varepsilon, \delta, \theta, \lambda, \sigma\}$ are well-prepared for a time horizon $T^* > 0$ and parameter $\alpha^* > 0$ large enough. Then $E(\rho_{T^*})$ well approximates a minimizer v^* of \mathcal{E} , and the following quantitative estimate holds*

$$|E(\rho_{T^*}) - v^*| \leq \epsilon, \quad (3.7)$$

for

$$\epsilon := C(C_0, c_1, \beta) \left((1 + C_{\alpha^*, \mathcal{E}}^\beta) \left(\frac{\lambda C_{T^*}}{\lambda \theta - 4C_{\alpha^*, \mathcal{E}} C_{\sigma, d}} \delta^{\frac{d-2}{4}} + \varepsilon_1 \right)^{\min\{1, \frac{\beta}{2}\}} + \varepsilon^\beta \right). \quad (3.8)$$

Remark 7 *The conditions of well-preparation (3.4) require that the initial datum ρ_0 is both well-concentrated and at the same time $v_{\alpha^*, \mathcal{E}}(\rho_0)$ already approximates well a global minimizer. Technically this is enforced by requiring that the product $C_{\alpha^*, \mathcal{E}}^{2\max\{1, \beta\}} V(\rho_0)^{1/2}$ is small for α^* large. Of course, this condition is fulfilled for any initial density ρ_0 , which is well-concentrated in the near of a global minimizer. Hence, the conditions (3.4) of well-preparation of ρ_0 may have a locality flavour. However, in the case the function \mathcal{E} is symmetric, i.e., $\mathcal{E}(v) = \mathcal{E}(-v)$ (as it happens in numerous applications, in particular the ones we present in this paper), then the condition is generically/practically satisfied at least for one of the two global minimizers $\pm v^*$, yielding essentially a global result.*

The proof of Theorem 6 is based on showing the monotone decay of the variance $V(\rho_t)$ under the assumption of well-preparation (Definition 5) and simultaneously by using the Laplace principle (1.7) and the inverse continuity property 4. of Assumptions 3.1 to derive the quantitative estimate

$$\left| \frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right| \leq C(C_0, c_1, \beta) \left((C_{\alpha, \mathcal{E}})^\beta V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^\beta \right). \quad (3.9)$$

The monotone decay of the variance is deduced by computing and estimating explicitly its derivative:

$$\begin{aligned} \frac{d}{dt} V(\rho_t) &= -\lambda V(\rho_t) \langle E(\rho_t), v_{\alpha, \mathcal{E}} \rangle - \frac{\lambda v_{\alpha, \mathcal{E}}^2 + 1}{2} 2V(\rho_t) + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \mathcal{E}})^2 d\rho_t \\ &\quad + C_{\sigma, d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha, \mathcal{E}})^2 \langle E(\rho_t), v \rangle d\rho_t \\ &\leq -\lambda V(\rho_t) \left(\langle E(\rho_t), v_{\alpha, \mathcal{E}} \rangle + \frac{v_{\alpha, \mathcal{E}}^2 + 1}{2} \right) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \mathcal{E}})^2 d\rho_t + 4C_{\alpha, \mathcal{E}} C_{\sigma, d} V(\rho_t). \end{aligned}$$

The idea is to balance all the terms on the right-hand side by using the parameters λ, σ in such a way of obtaining a negative sign. This also requires to show that, as soon as $V(\rho_t)$ is small enough, $|E(\rho_t)| \approx \langle E, v_{\alpha, \mathcal{E}} \rangle \approx |v_{\alpha, \mathcal{E}}| \approx 1$ and these estimates are worked out in Lemma 8. For ease of notation, for any vector $v \in \mathbb{R}^d$ we may write v^2 to mean $|v|^2$.

3.2 Auxiliary lemmas

A simple computation yields $2V(\rho_t) = 1 - E(\rho_t)^2$. In particular, as soon as $V(\rho_t)$ is small $E(\rho_t)^2 \approx 1$ and below we will silently apply the assignment $\mathcal{E}(E(\rho_t)) := \mathcal{E}\left(\frac{E(\rho_t)}{|E(\rho_t)|}\right)$ by normal extension. Since $E(\rho_t) = \mathbb{E}[\bar{V}_t]$, it follows from (3.2) that

$$\frac{d}{dt}E(\rho_t) = - \int_{\mathbb{S}^{d-1}} \eta_t d\rho_t - \int_{\mathbb{S}^{d-1}} \frac{(d-1)\sigma^2}{2} (v - v_{\alpha,\mathcal{E}})^2 v d\rho_t. \quad (3.10)$$

In the following lemma, we summarize some useful estimates of $v_{\alpha,\mathcal{E}}(\rho_t)$, $E(\rho_t)$ and $V(\rho_t)$. Here we recall the definition

$$v_{\alpha,\mathcal{E}}(\rho_t) := \frac{\int_{\mathbb{S}^{d-1}} v \omega_{\alpha}^{\mathcal{E}}(v) d\rho_t(v)}{\|\omega_{\alpha}^{\mathcal{E}}\|_{L^1(\rho_t)}} = \frac{\int_{\mathbb{S}^{d-1}} v e^{-\alpha\mathcal{E}(v)} d\rho_t(v)}{\|e^{-\alpha\mathcal{E}}\|_{L^1(\rho_t)}}. \quad (3.11)$$

Lemma 8 *Let $v_{\alpha,\mathcal{E}}(\rho_t)$ be defined as above. It holds that*

1. $\int_{\mathbb{S}^{d-1}} |v - v_{\alpha,\mathcal{E}}(\rho_t)|^2 d\rho_t \leq 4C_{\alpha,\mathcal{E}} V(\rho_t)$ and $\int_{\mathbb{S}^{d-1}} |v - v_{\alpha,\mathcal{E}}(\rho_t)| d\rho_t \leq 2C_{\alpha,\mathcal{E}} V(\rho_t)^{\frac{1}{2}}$;
2. $v_{\alpha,\mathcal{E}}(\rho_t)^2 \geq 1 - 4C_{\alpha,\mathcal{E}}^2 V(\rho_t)$;
3. $|v_{\alpha,\mathcal{E}}(\rho_t) - E(\rho_t)|^2 \leq (4C_{\alpha,\mathcal{E}}^2 - 2)V(\rho_t)$;

where $C_{\alpha,\mathcal{E}} = e^{\alpha(\bar{\mathcal{E}} - \underline{\mathcal{E}})}$.

Before proving the key estimate (3.9), we need a lower bound on the norm of the weights $\|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_t)}$, which is ensured by the following auxiliary result.

Lemma 9 *Let c_1, c_2 be the constants from the assumptions on \mathcal{E} . Then we have*

$$\frac{d}{dt} \|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_t)}^2 \geq -b_1(\sigma, d, \alpha, c_1, c_2, \underline{\mathcal{E}}) V(\rho_t) - b_2(d, \alpha, c_1, \underline{\mathcal{E}}) \lambda V(\rho_t)^{\frac{1}{2}} \quad (3.12)$$

with $0 \leq b_1, b_2 \leq 1$ and $b_1, b_2 \rightarrow 0$ as $\alpha \rightarrow \infty$.

3.3 Proof of the large time asymptotic result

Proposition 10 *For any fixed $T > 0$, assume that*

$$\bar{V}_T := \sup_{0 \leq t \leq T} V(\rho_t) \leq \min \left\{ T^{-1} \|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_0)}^2, T^{-1} \lambda^{-2} \|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_0)}^4, \frac{3}{8} \right\}.$$

Then for any $\varepsilon > 0$, there exists a minimizer v^ of \mathcal{E} such that*

$$\left| \frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right| \leq C(C_0, c_1, \beta) \left((C_{\alpha,\mathcal{E}})^{\beta} V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^{\beta} \right) \quad \text{for all } t \in [0, T] \quad (3.13)$$

holds for any $\alpha > \alpha_0$ with some $\alpha_0 \gg 1$, where $C_{\alpha,\mathcal{E}} = e^{\alpha(\bar{\mathcal{E}} - \underline{\mathcal{E}})}$, and C_0, c_1, β are used in Assumption 3.1. Moreover, as soon as $|E(\rho_t)| \geq 1/2$

$$\left| v_{\alpha,\mathcal{E}}(\rho_t) - \frac{E(\rho_t)}{|E(\rho_t)|} \right|^2 \leq \left(8C_{\alpha,\mathcal{E}}^2 - \frac{4}{3} \right) V(\rho_t). \quad (3.14)$$

As it is needed in the proof of this proposition, for readers' convenience, we give a brief introduction of the Wasserstein metric in the following definition, we refer, e.g., to (Ambrosio et al., 2008) for more details.

Definition 11 (Wasserstein Metric) For any $1 \leq p < \infty$, let $\mathcal{P}_p(\mathbb{R}^d)$ be the space of Borel probability measures on \mathbb{R}^d with finite p moment. We equip this space with the Wasserstein distance

$$W_p^p(\mu, \nu) := \inf \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} |z - \hat{z}|^p d\pi(\mu, \nu) \mid \pi \in \Pi(\mu, \nu) \right\} \quad (3.15)$$

where $\Pi(\mu, \nu)$ denotes the collection of all Borel probability measures on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals μ and ν in the first and second component respectively. If $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ have bounded support, then the 1-Wasserstein distance can be equivalently expressed in terms of the dual formulation

$$W_1(\mu, \nu) := \sup \left\{ \int_{\mathbb{R}^d} f(v) d(\mu - \nu)(v) \mid f \in \text{Lip}(\mathbb{R}^d), \text{Lip}(f) \leq 1 \right\} \quad (3.16)$$

Proof (Proposition 10) It follows from Lemma 9 that

$$\begin{aligned} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)}^2 &\geq \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2 - b_1(\alpha) \int_0^t V(\rho_s) ds - b_2(\alpha) \lambda \int_0^t V(\rho_s)^{\frac{1}{2}} ds \\ &\geq \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2 - b_1(\alpha) \bar{\mathcal{V}}_T T - b_2(\alpha) \lambda \bar{\mathcal{V}}_T^{\frac{1}{2}} T \\ &\geq \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2 - b_1(\alpha) \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2 - b_2(\alpha) \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2, \end{aligned}$$

where we have used the assumption

$$\bar{\mathcal{V}}_T := \sup_{0 \leq t \leq T} V(\rho_t) \leq \min \left\{ T^{-1} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2, T^{-1} \lambda^{-2} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^4 \right\}. \quad (3.17)$$

The above inequality implies

$$-\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} \leq -\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)} - \frac{1}{2\alpha} \log(1 - b_1(\alpha) - b_2(\alpha)).$$

The Laplace principle states

$$\lim_{\alpha \rightarrow \infty} -\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)} = \underline{\mathcal{E}}, \quad (3.18)$$

which implies the existence of an $\alpha_1 \gg 1$ such that any $\alpha > \alpha_1$ it holds

$$-\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)} - \underline{\mathcal{E}} < \frac{\varepsilon}{2} \quad (3.19)$$

for any $\varepsilon > 0$. Together with the fact that $b_1(\alpha), b_2(\alpha) \rightarrow 0$ as $\alpha \rightarrow \infty$, it yields that

$$-\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \underline{\mathcal{E}} \leq -\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)} - \underline{\mathcal{E}} - \frac{1}{2\alpha} \log(1 - b_1(\alpha) - b_2(\alpha)) \leq \varepsilon,$$

for any $\alpha > \alpha_2$ with some $\alpha_2 \gg 1$. Let us assume that $\bar{V}_T \leq \frac{3}{8}$, then

$$\frac{1}{2} \leq |E(\rho_t)| \leq 1.$$

By the dual representation of 1-Wasserstein distance W_1 , we know that

$$\begin{aligned} & \left| \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \omega_{\mathcal{E}}^\alpha \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) \right| = \left| \int_{\mathbb{R}^d} e^{-\alpha \mathcal{E}(v)} d\rho_t(v) - \delta_{\frac{E(\rho_t)}{|E(\rho_t)|}}(v) \right| \\ & \leq \alpha e^{-\alpha \underline{\mathcal{E}}} \|\nabla \mathcal{E}\|_\infty W_1(\rho_t, \delta_{\frac{E(\rho_t)}{|E(\rho_t)|}}) \leq \alpha c_1 e^{-\alpha \underline{\mathcal{E}}} W_2(\rho_t, \delta_{\frac{E(\rho_t)}{|E(\rho_t)|}}) \leq 2\sqrt{\frac{2}{3}} \alpha c_1 e^{-\alpha \underline{\mathcal{E}}} V(\rho_t)^{\frac{1}{2}}. \end{aligned} \quad (3.20)$$

Here we have used the fact that

$$W_2(\rho_t, \delta_{\frac{E(\rho_t)}{|E(\rho_t)|}})^2 \leq \int_{\mathbb{S}^{d-1}} \left| v - \frac{E(\rho_t)}{|E(\rho_t)|} \right|^2 d\rho_t = 2 - 2|E(\rho_t)| = \frac{4V(\rho_t)}{1 + |E(\rho_t)|} \leq \frac{8}{3} V(\rho_t). \quad (3.21)$$

Above (3.20) leads to

$$\begin{aligned} & \left| -\frac{1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \mathcal{E} \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) \right| = \left| -\frac{1}{\alpha} \left(\log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \log \omega_{\mathcal{E}}^\alpha \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) \right) \right| \\ & \leq \frac{e^{\alpha \bar{\mathcal{E}}}}{\alpha} \left| \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \omega_{\mathcal{E}}^\alpha \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) \right| \leq 2\sqrt{\frac{2}{3}} c_1 C_{\alpha, \mathcal{E}} V(\rho_t)^{\frac{1}{2}}. \end{aligned}$$

Hence we have

$$\begin{aligned} 0 \leq \mathcal{E} \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) - \underline{\mathcal{E}} & \leq \mathcal{E} \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) - \frac{-1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} + \frac{-1}{\alpha} \log \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_t)} - \underline{\mathcal{E}} \\ & \leq 2\sqrt{\frac{2}{3}} c_1 C_{\alpha, \mathcal{E}} V(\rho_t)^{\frac{1}{2}} + \varepsilon, \end{aligned}$$

which yields that

$$\left| \frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right| \leq C_0 \left| \mathcal{E} \left(\frac{E(\rho_t)}{|E(\rho_t)|} \right) - \underline{\mathcal{E}} \right|^\beta \leq C(C_0, c_1, \beta) \left((C_{\alpha, \mathcal{E}})^\beta V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^\beta \right).$$

by the inverse continuity 4. in Assumption 3.1, where v^* is a minimizer of \mathcal{E} . Next we compute

$$\begin{aligned} & \left| v_{\alpha, \mathcal{E}}(\rho_t) - \frac{E(\rho_t)}{|E(\rho_t)|} \right|^2 = \int_{\mathbb{S}^{d-1}} \left| v_{\alpha, \mathcal{E}}(\rho_t) - v + v - \frac{E(\rho_t)}{|E(\rho_t)|} \right|^2 d\rho_t(v) \\ & = \int_{\mathbb{S}^{d-1}} |v_{\alpha, \mathcal{E}}(\rho_t) - v|^2 d\rho_t + \int_{\mathbb{S}^{d-1}} \left| v - \frac{E(\rho_t)}{|E(\rho_t)|} \right|^2 d\rho_t + 2 \int_{\mathbb{S}^{d-1}} \left\langle v_{\alpha, \mathcal{E}}(\rho_t) - v, v - \frac{E(\rho_t)}{|E(\rho_t)|} \right\rangle d\rho_t \\ & \leq 4C_{\alpha, \mathcal{E}}^2 V(\rho_t) + \frac{8}{3} V(\rho_t) + 2|E(\rho_t)| - 2 + \left(2 - \frac{2}{|E(\rho_t)|} \right) \langle v_{\alpha, \mathcal{E}}(\rho_t), E(\rho_t) \rangle \\ & \leq 4C_{\alpha, \mathcal{E}}^2 V(\rho_t) + \frac{8}{3} V(\rho_t) - 2V(\rho_t) + \left(2 - \frac{2}{|E(\rho_t)|} \right) \langle v_{\alpha, \mathcal{E}}(\rho_t), E(\rho_t) \rangle \end{aligned}$$

$$= (4C_{\alpha,\varepsilon}^2 + \frac{2}{3})V(\rho_t) + \left(2 - \frac{2}{|E(\rho_t)|}\right) \langle v_{\alpha,\varepsilon}(\rho_t), E(\rho_t) \rangle,$$

where we have used (3.21) and $\frac{1}{2} \leq |E(\rho_t)| \leq 1$. Notice that

$$\begin{aligned} \left(2 - \frac{2}{|E(\rho_t)|}\right) \langle v_{\alpha,\varepsilon}(\rho_t), E(\rho_t) \rangle &= \left(2 - \frac{2}{|E(\rho_t)|}\right) \frac{v_{\alpha,\varepsilon}(\rho_t)^2 + E(\rho_t)^2 - |v_{\alpha,\varepsilon}(\rho_t) - E(\rho_t)|^2}{2} \\ &\leq \left(\frac{2}{|E(\rho_t)|} - 2\right) \frac{|v_{\alpha,\varepsilon}(\rho_t) - E(\rho_t)|^2}{2} \leq (4C_{\alpha,\varepsilon}^2 - 2)V(\rho_t). \end{aligned}$$

Thus we have

$$\left|v_{\alpha,\varepsilon}(\rho_t) - \frac{E(\rho_t)}{|E(\rho_t)|}\right|^2 \leq \left(8C_{\alpha,\varepsilon}^2 - \frac{4}{3}\right) V(\rho_t).$$

Hence we complete the proof. \blacksquare

The next ingredient is proving the monotone decay of the variance $V(\rho_t)$ under assumptions of well-preparation (see Definition 5).

Proposition 12 *Let us fix any $T > 0$ and choose α large enough and assume that the parameters and the initial datum are well-prepared in the sense of Definition 5. Then it holds*

$$V(\rho_t) \leq V(\rho_0)e^{-(\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d})t} + \frac{\lambda C_T}{\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} \quad \text{for all } t \in [0, T]. \quad (3.22)$$

Proof Let us compute the derivative of the variance (where $C_{\sigma,d} = \frac{(d-1)\sigma^2}{2}$)

$$\begin{aligned} \frac{d}{dt}V(\rho_t) &= \frac{1}{2} \frac{d}{dt} \left(\int_{\mathbb{S}^{d-1}} v^2 d\rho_t - E(\rho_t)^2 \right) = \frac{1}{2} \frac{d}{dt} \left(1 - E(\rho_t)^2 \right) = -E(\rho_t) \frac{d}{dt} E(\rho_t) \\ &= E(\rho_t) \int_{\mathbb{S}^{d-1}} \eta_t d\rho_t + C_{\sigma,d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t \\ &= \lambda \int_{\mathbb{S}^{d-1}} \langle v_{\alpha,\varepsilon}, v \rangle \langle E(\rho_t), v \rangle - \langle E(\rho_t), v_{\alpha,\varepsilon} \rangle d\rho_t + C_{\sigma,d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t. \end{aligned}$$

Notice that

$$\langle E(\rho_t), v \rangle = \frac{1}{2}(E(\rho_t)^2 + v^2 - |E(\rho_t) - v|^2) = \frac{1}{2}(E(\rho_t)^2 + 1 - (E(\rho_t) - v)^2).$$

Then one has

$$\begin{aligned} \frac{d}{dt}V(\rho_t) &= \lambda \left(\frac{E(\rho_t)^2 + 1}{2} - 1 \right) \langle E(\rho_t), v_{\alpha,\varepsilon}(\rho_t) \rangle - \frac{\lambda}{2} \int_{\mathbb{S}^{d-1}} \langle v_{\alpha,\varepsilon}, v \rangle (E(\rho_t) - v)^2 d\rho_t \\ &\quad + C_{\sigma,d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t \\ &= -\lambda V(\rho_t) \langle E(\rho_t), v_{\alpha,\varepsilon}(\rho_t) \rangle - \frac{\lambda}{2} \int_{\mathbb{S}^{d-1}} \langle v_{\alpha,\varepsilon}, v \rangle (E(\rho_t) - v)^2 d\rho_t \\ &\quad + C_{\sigma,d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t, \end{aligned}$$

where we have used the fact that $2V(\rho_t) = 1 - E(\rho_t)^2$. Moreover, since

$$\langle v_{\alpha,\varepsilon}, v \rangle = \frac{1}{2}(v_{\alpha,\varepsilon}^2 + v^2 - |v_{\alpha,\varepsilon} - v|^2) = \frac{1}{2}(v_{\alpha,\varepsilon}^2 + 1 - (v_{\alpha,\varepsilon}(\rho_t) - v)^2)$$

and $\int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 d\rho_t = 2V(\rho_t)$, we have

$$\begin{aligned} \frac{d}{dt}V(\rho_t) &= -\lambda V(\rho_t)\langle E(\rho_t), v_{\alpha,\varepsilon}(\rho_t) \rangle - \frac{\lambda}{2} \frac{v_{\alpha,\varepsilon}^2 + 1}{2} 2V(\rho_t) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha,\varepsilon})^2 d\rho_t + C_{\sigma,d} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 \langle E(\rho_t), v \rangle d\rho_t \\ &\leq -\lambda V(\rho_t) \left(\langle E(\rho_t), v_{\alpha,\varepsilon}(\rho_t) \rangle + \frac{v_{\alpha,\varepsilon}^2 + 1}{2} \right) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha,\varepsilon})^2 d\rho_t + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t), \end{aligned}$$

where we have used estimate (4.3) in the last inequality.

Next we observe that

$$\begin{aligned} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 d\rho_t &= \int_{\mathbb{S}^{d-1}} (v - E(\rho_t) + E(\rho_t) - v_{\alpha,\varepsilon})^2 d\rho_t \\ &= \int_{\mathbb{S}^{d-1}} (v - E(\rho_t))^2 d\rho_t + (E(\rho_t) - v_{\alpha,\varepsilon})^2 = 2V(\rho_t) + E(\rho_t)^2 + v_{\alpha,\varepsilon}^2 - 2\langle E(\rho_t), v_{\alpha,\varepsilon} \rangle. \end{aligned}$$

So it holds

$$\begin{aligned} \langle E(\rho_t), v_{\alpha,\varepsilon} \rangle &= V(\rho_t) + \frac{E(\rho_t)^2 + v_{\alpha,\varepsilon}^2}{2} - \frac{1}{2} \int_{\mathbb{S}^{d-1}} (v - v_{\alpha,\varepsilon})^2 d\rho_t \\ &\geq V(\rho_t) + \frac{E(\rho_t)^2 + v_{\alpha,\varepsilon}^2}{2} - 2C_{\alpha,\varepsilon}V(\rho_t), \end{aligned} \tag{3.23}$$

where we have used (4.3) again. Thus we obtain that

$$\begin{aligned} \frac{d}{dt}V(\rho_t) &\leq -\lambda V(\rho_t) \left(V(\rho_t) + \frac{2v_{\alpha,\varepsilon}^2 + 1 + E(\rho_t)^2}{2} - 2C_{\alpha,\varepsilon}V(\rho_t) \right) + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha,\varepsilon})^2 d\rho_t \\ &= -\lambda V(\rho_t) (v_{\alpha,\varepsilon}^2 + 1 - 2C_{\alpha,\varepsilon}V(\rho_t)) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha,\varepsilon})^2 d\rho_t + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t) \\ &\leq -\lambda V(\rho_t) (2 - 2C_{\alpha,\varepsilon}V(\rho_t) - 4C_{\alpha,\varepsilon}^2V(\rho_t)) \\ &\quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha,\varepsilon})^2 d\rho_t + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t), \end{aligned}$$

where we have used $2V(\rho_t) = 1 - E(\rho_t)^2$ in the second equality and 2) from Lemma 8 in the last inequality.

Let v^* be the minimizer used in Proposition 10, and one has

$$\begin{aligned}
 & \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v_{\alpha, \varepsilon})^2 d\rho_t \\
 = & \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v_{\alpha, \varepsilon} - v^*)^2 d\rho_t \\
 & + 2 \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 \langle v - v^*, v^* - v_{\alpha, \varepsilon} \rangle d\rho_t \\
 \leq & \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + 2 \left(v_{\alpha, \varepsilon} - \frac{E(\rho_t)}{|E(\rho_t)|} \right)^2 \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 d\rho_t \\
 & + 2 \left(\frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right)^2 \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 d\rho_t \\
 & + 4 \left| v_{\alpha, \varepsilon} - \frac{E(\rho_t)}{|E(\rho_t)|} \right| \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 d\rho_t + 4 \left| \frac{E(\rho_t)}{|E(\rho_t)|} - v^* \right| \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 d\rho_t \\
 \leq & \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + 2(8C_{\alpha, \varepsilon}^2 - \frac{4}{3})V(\rho_t)^2 \\
 & + 2C(C_0, c_1, \beta) \left((C_{\alpha, \varepsilon})^{2\beta} V(\rho_t)^\beta + \varepsilon^{2\beta} \right) V(\rho_t) \\
 & + 4(8C_{\alpha, \varepsilon}^2 - \frac{4}{3})^{\frac{1}{2}} V(\rho_t)^{\frac{1}{2}} V(\rho_t) + 4C(C_0, c_1, \beta) \left((C_{\alpha, \varepsilon})^\beta V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^\beta \right) V(\rho_t),
 \end{aligned}$$

where we have used estimate (3.14) and Proposition 10 for $\alpha > \alpha_0$. This implies that

$$\begin{aligned}
 \frac{d}{dt} V(\rho_t) & \leq -\lambda V(\rho_t) \left(2 - 2C_{\alpha, \varepsilon} V(\rho_t) - 4C_{\alpha, \varepsilon}^2 V(\rho_t) - \frac{1}{2} (8C_{\alpha, \varepsilon}^2 - \frac{4}{3}) V(\rho_t) \right. \\
 & \quad - \frac{1}{2} C(C_0, c_1, \beta) \left((C_{\alpha, \varepsilon})^{2\beta} V(\rho_t)^\beta + \varepsilon^{2\beta} \right) - (8C_{\alpha, \varepsilon}^2 - \frac{4}{3})^{\frac{1}{2}} V(\rho_t)^{\frac{1}{2}} \\
 & \quad \left. - C(C_0, c_1, \beta) \left((C_{\alpha, \varepsilon})^\beta V(\rho_t)^{\frac{\beta}{2}} + \varepsilon^\beta \right) \right) \\
 & \quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + 4C_{\alpha, \varepsilon} C_{\sigma, d} V(\rho_t) \\
 & \leq -\lambda V(\rho_t) \left(2 - C^* \left(C_{\alpha, \varepsilon}^{2 \max\{1, \beta\}} \bar{V}_T^{\min\{1/2, \beta\}} + \varepsilon^\beta \right) \right) \\
 & \quad + \frac{\lambda}{4} \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + 4C_{\alpha, \varepsilon} C_{\sigma, d} V(\rho_t),
 \end{aligned}$$

where $\bar{V}_T := \sup_{0 \leq t \leq T} V(\rho_t) \leq \frac{1}{2}$, and $C^* > 0$ is a constant depending only on c_1, β and C_0 .

Now we treat the term $\int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t$, which can be split into two parts

$$\begin{aligned}
 & \int_{\mathbb{S}^{d-1}} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t \\
 = & \int_{\mathcal{D}_\delta} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + \int_{\mathbb{S}^{d-1}/\mathcal{D}_\delta} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t,
 \end{aligned}$$

for some $\delta > 0$, where

$$\mathcal{D}_\delta := \left\{ v \in \mathbb{S}^{d-1} \mid -1 \leq \langle v, v^* \rangle \leq -1 + \delta \right\}.$$

This means that

$$\frac{\lambda}{4} \int_{\mathbb{S}^{d-1}/\mathcal{D}_\delta} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t \leq \lambda(2 - \delta)V(\rho_t). \quad (3.24)$$

Hence one can conclude

$$\begin{aligned} \frac{d}{dt}V(\rho_t) &\leq -\lambda V(\rho_t) \left(\delta - C^* \left(C_{\alpha,\varepsilon}^{2\max\{1,\beta\}} \bar{V}_T^{\min 1/2\{1,\beta\}} + \varepsilon^\beta \right) \right) \\ &\quad + \frac{\lambda}{4} \int_{\mathcal{D}_\delta} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t), \end{aligned} \quad (3.25)$$

where we emphasize that $\delta > 0$.

Notice that \mathcal{D}_δ can be understood as a small cap on the sphere that is on the opposite side of the minimizer v^* . By the assumption that $\rho_0 \in L^2(\mathbb{S}^{d-1})$ (see Definition 5), we have the solution ρ_t is not just a measure but it is a function, and for any given $T > 0$ it satisfies $\rho \in L^\infty([0, T]; L^2(\mathbb{S}^{d-1}))$. This can be proved through a standard argument of PDE theory, which we provide in Theorem 13. Thus we have

$$\int_{\mathcal{D}_\delta} d\rho_t = \int_{\mathcal{D}_\delta} \rho_t(v) dv \leq \|\rho_t\|_2 |\mathcal{D}_\delta|^{\frac{1}{2}} \leq C(T)(A_\delta)^{\frac{1}{2}}, \quad (3.26)$$

where A_δ denotes the area of the hyperspherical cap \mathcal{D}_δ , which satisfies the formula

$$A_\delta = \frac{1}{2} a_d I_{2\delta-\delta^2} \left(\frac{d-1}{2}, \frac{1}{2} \right) \leq C \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2})} \frac{(d-1)^{\frac{1}{2}}}{d-2} \delta^{\frac{d-2}{2}}, \quad (3.27)$$

where a_d represents the area of a unit ball and $I_x(a, b)$ is the regularized incomplete beta function. Note that

$$A_\delta \rightarrow 0 \text{ as } \delta \rightarrow 0. \quad (3.28)$$

This means that for d sufficiently large it holds

$$\int_{\mathcal{D}_\delta} (E(\rho_t) - v)^2 (v - v^*)^2 d\rho_t \leq 16 \int_{\mathcal{D}_\delta} d\rho_t(v) \leq C(\lambda, \sigma, T, \|\rho_0\|_2)(A_\delta)^{\frac{1}{2}} \leq 4C_T \delta^{\frac{d-2}{4}}.$$

Therefore we have

$$\frac{d}{dt}V(\rho_t) \leq -\lambda V(\rho_t) \left(\delta - C^* \left(C_{\alpha,\varepsilon}^{2\max\{1,\beta\}} \bar{V}_T^{\min 1/2\{1,\beta\}} + \varepsilon^\beta \right) \right) + 4C_{\alpha,\varepsilon}C_{\sigma,d}V(\rho_t) + \lambda C_T \delta^{\frac{d-2}{4}} \quad (3.29)$$

for all $t \in [0, T]$. Let us assume that

$$\delta - C^* \left(C_{\alpha,\varepsilon}^{2\max\{1,\beta\}} \bar{V}_T^{\min 1/2\{1,\beta\}} + \varepsilon^\beta \right) \geq \theta > 0, \quad \text{i.e. } 0 \leq C_{\alpha,\varepsilon}^{2\max\{1,\beta\}} \bar{V}_T^{\min 1/2\{1,\beta\}} + \varepsilon^\beta \leq \frac{\delta - \theta}{C^*}. \quad (3.30)$$

Then we have

$$\frac{d}{dt}V(\rho_t) \leq -(\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d})V(\rho_t) + \lambda C_T \delta^{\frac{d-2}{4}},$$

which leads to

$$V(\rho_t) \leq V(\rho_0)e^{-(\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d})t} + \frac{\lambda C_T}{\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} \text{ for all } t \in [0, T],$$

which is contractive as soon as $\lambda\theta > 4C_{\alpha,\varepsilon}C_{\sigma,d}$. We are left to verify the assumptions that $\bar{V}_T \leq \min \left\{ T^{-1} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2, T^{-1} \lambda^{-2} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^4, \frac{3}{8} \right\}$ and (3.30), which hold if we assume that

$$\begin{aligned} C_{\alpha,\varepsilon}^{2 \max\{1,\beta\}} \left(V(\rho_0) + \frac{\lambda C_T}{\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} \right)^{\frac{1}{2} \min\{1,\beta\}} + \varepsilon^\beta &< \frac{\delta - \theta}{C^*}, \\ V(\rho_0) + \frac{\lambda C_T}{\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} &\leq \min \left\{ T^{-1} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^2, T^{-1} \lambda^{-2} \|\omega_{\mathcal{E}}^\alpha\|_{L^1(\rho_0)}^4, \frac{3}{8} \right\}. \end{aligned}$$

Hence we complete the proof. \blacksquare

Proof (Theorem 6) Proposition 12 implies that for any $\varepsilon_1 > 0$, there exists some T^* large enough such that

$$V(\rho_{T^*}) \leq \varepsilon_0 := \frac{\lambda C_{T^*}}{\lambda\theta - 4C_{\alpha,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} + \varepsilon_1.$$

Moreover $1 \geq |E(\rho_{T^*})| = \sqrt{1 - 2V(\rho_{T^*})} \geq \sqrt{1 - 2\varepsilon_0}$ and

$$\left| E(\rho_{T^*}) - \frac{E(\rho_{T^*})}{|E(\rho_{T^*})|} \right| \leq \frac{1 - |E(\rho_{T^*})|}{|E(\rho_{T^*})|} \leq \frac{1 - \sqrt{1 - 2\varepsilon_0}}{\sqrt{1 - 2\varepsilon_0}} \leq 2\varepsilon_0,$$

as soon as $0 \leq \varepsilon_0 \leq \frac{1}{4}(\sqrt{5} - 1)$, which is fulfilled as soon as δ, ε_1 are chosen small enough. These estimates, triangle inequality and Proposition 10 lead to the quantitative estimate

$$|E(\rho_{T^*}) - v^*| \leq C(C_0, c_1, \beta) \left((1 + C_{\alpha^*,\varepsilon}^\beta) \left(\frac{\lambda C_{T^*}}{\lambda\theta - 4C_{\alpha^*,\varepsilon}C_{\sigma,d}} \delta^{\frac{d-2}{4}} + \varepsilon_1 \right)^{\min\{1, \frac{\beta}{2}\}} + \varepsilon^\beta \right).$$

Note once again here that ε, δ , and ε_1 can be all chosen to be sufficiently small. \blacksquare

3.4 Proof of the main result

Let us finally address the proof of the main theorem of this paper.

Proof (Theorem 1) In order to show a concrete instance of the result, we develop the proof for the case where $\{V_n^i := V_{\Delta t, n}^i : n = 0, \dots, n_{T^*}; i = 1 \dots N\}$ are generated by the iterative algorithm (1.8). However, any other numerical scheme of order m can be considered (Platen, 1999). The SDE system (1.2) is well-posed by (Fornasier et al., 2020, Theorem 2.1) and it admits a pathwise strong solution $V_t^i, i = 1, \dots, N$. The iterative algorithm (1.8) is the discrete-time (projected Euler-Maruyama) approximation of the SDE system (1.2)

with order of approximation $m = 1/2$ by classical results, e.g., see (Higham et al., 2002, Theorem 2.2)

$$\mathbb{E} \left[\sup_{n=0, \dots, n_{T^*}} |V_{\Delta t, n} - V_{t_n}|^2 \right] \leq \bar{C}_1 (\Delta t)^{2m}, \quad (3.31)$$

for \bar{C}_1 which depends linearly on d and N , and possibly exponentially on T^* , λ , and σ (see in particular the estimates before (2.11) in the proof of (Higham et al., 2002, Theorem 2.2)). Let us stress that the introduction of the post-projection $V_{n+1}^i \leftarrow \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i|$ to enforce the dynamics on the sphere may produce an additional error of at most order Δt because

$$\left| \tilde{V}_{n+1}^i - \tilde{V}_{n+1}^i / |\tilde{V}_{n+1}^i| \right|^2 = |\tilde{V}_{n+1}^i|^2 + 1 - 2|\tilde{V}_{n+1}^i| = (|\tilde{V}_{n+1}^i| - 1)^2$$

and, in view of

$$\tilde{V}_{n+1}^i = V_n^i + \Delta t P(V_n^i) V_n^{\alpha, \mathcal{E}} + \sigma |V_n^i - V_n^{\alpha, \mathcal{E}}| P(V_n^i) \Delta B_n^i - \Delta t \frac{\sigma^2}{2} (V_n^i - V_n^{\alpha, \mathcal{E}})^2 (d-1) V_n^i,$$

we obtain

$$\begin{aligned} \mathbb{E} \left[(|\tilde{V}_{n+1}^i| - 1)^2 \right] &= \mathbb{E} \left[(|\tilde{V}_{n+1}^i| - |V_n^i|)^2 \right] \\ &\leq \mathbb{E} \left[((1 + 2(d-1)\sigma^2)\Delta t + 2\sigma|\Delta B_n^i|)^2 \right] \\ &\leq (1 + 2(d-1)\sigma^2)^2 \Delta t^2 + 4\sigma^2 d \Delta t + 2\sigma(1 + 2(d-1)\sigma^2) \Delta t \sqrt{d\Delta t} \\ &\leq \bar{C}'_1 \Delta t. \end{aligned} \quad (3.32)$$

By (Fornasier et al., 2020, Theorem 2.2) we have also well-posedness of (1.11) with pathwise strong solution \bar{V}_t . For \bar{V}_0^i drawn i.i.d. according to ρ_0 , $i = 1 \dots, N$, an application of (Fornasier et al., 2020, Theorem 3.1) yields

$$\sup_{t \in [0, T]} \sup_{i=1, \dots, N} \mathbb{E} \left[|V_t^i - \bar{V}_t^i|^2 \right] \leq \bar{C}_2 N^{-1}, \quad (3.33)$$

for any $T > 0$ time horizon. As clarified in (Fornasier et al., 2020, Remark 3.2 and Lemma 3.1), the constant \bar{C}_2 depends at most linearly on d , and, as a worst case analysis, polynomially on $C_{\alpha^*, \mathcal{E}}$, and exponentially on T . By law of large numbers, for $\rho_t = \text{law}(\bar{V}_t)$ it holds

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N \bar{V}_{T^*}^i - E(\rho_{T^*}) \right|^2 \leq \bar{C}'_2 N^{-1}. \quad (3.34)$$

Under the assumptions of well-preparation, Theorem 6 yields

$$|E(\rho_{T^*}) - v^*|^2 \leq \bar{C}_3 \epsilon, \quad (3.35)$$

for \bar{C}_3 that depends polynomially on $C_{\alpha^*, \mathcal{E}}$. By combining the strong convergence (3.31), the mean-field limit (3.33), the law of large numbers (3.34), and the large time asymptotics (3.35) we conclude by multiple applications of Jensen inequality the final error estimate

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N V_{\Delta t, n_{T^*}}^i - v^* \right|^2 \right]$$

$$\begin{aligned}
 &\leq 8 \left(\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (V_{\Delta t, n_{T^*}}^i - V_{T^*}^i) \right|^2 \right] + \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (V_{T^*}^i - \bar{V}_{T^*}^i) \right|^2 \right] \right. \\
 &\quad \left. + \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \bar{V}_{T^*}^i - E(\rho_{T^*}) \right|^2 \right] + |E(\rho_{T^*}) - v^*|^2 \right) \\
 &\leq 8\bar{C}_1(\Delta t)^{2m} + 8(\bar{C}_2 + \bar{C}'_2)N^{-1} + 8\bar{C}_3\epsilon^2.
 \end{aligned} \tag{3.36}$$

■

4. Auxiliary Results and Proofs

4.1 Proofs of auxiliary lemmas

Proof (Lemma 2)

From (2.2) we get

$$\begin{aligned}
 \langle V_{n+1}^i, V_{n+1}^i \rangle &= \langle V_n^i, V_n^i \rangle + \langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i) \rangle \\
 &\quad + 2\langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), V_n^i \rangle.
 \end{aligned}$$

Assuming $\langle V_{n+1}^i, V_{n+1}^i \rangle = \langle V_n^i, V_n^i \rangle$ implies

$$\begin{aligned}
 0 &= \langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i) \rangle + 2\langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), V_n^i \rangle \\
 &= \langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i) + 2V_n^i \rangle \\
 &= \langle \Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i), V_{n+1}^i + V_n^i \rangle
 \end{aligned}$$

where we used the fact that $\Phi(\Delta t, V_n^i, V_{n+1}^i, \xi_n^i) = V_{n+1}^i - V_n^i$. ■

Proof (Lemma 8) Using Jensen's inequality, one concludes that

$$\int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 d\rho_t \leq \frac{1}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}} \int_{\mathbb{S}^{d-1}} \int_{\mathbb{S}^{d-1}} |v - u|^2 e^{-\alpha \mathcal{E}(u)} d\rho_t(v) d\rho_t(u). \tag{4.1}$$

The expression on the right can be further estimated as follows

$$\int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 d\rho_t \leq 4 \frac{e^{-\alpha \mathcal{E}}}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}} V(\rho_t) \tag{4.2}$$

$$\leq 4C_{\alpha, \mathcal{E}} V(\rho_t), \tag{4.3}$$

where $C_{\alpha, \mathcal{E}} = e^{\alpha(\bar{\mathcal{E}} - \mathcal{E})}$. Similarly one has

$$\begin{aligned}
 \int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t &\leq \frac{1}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}} \int \int |v - u| e^{-\alpha \mathcal{E}(u)} d\rho_t(v) d\rho_t(u) \leq 2 \frac{e^{-\alpha \mathcal{E}}}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}} V(\rho_t)^{\frac{1}{2}} \\
 &\leq 2C_{\alpha, \mathcal{E}} V(\rho_t)^{\frac{1}{2}}.
 \end{aligned} \tag{4.4}$$

Next we notice that

$$1 - v_{\alpha, \mathcal{E}}(\rho_t)^2 = \frac{\int_{\mathbb{S}^{d-1}} (v - v_{\alpha, \mathcal{E}}(\rho_t))^2 \omega_{\mathcal{E}}^{\alpha}(u) d\rho_t(u)}{\|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_t)}} \leq 4C_{\alpha, \mathcal{E}}^2 V(\rho_t), \quad (4.5)$$

where we have used (4.3) in the last inequality. This implies estimate 2).

To obtain 3), we compute

$$\begin{aligned} |v_{\alpha, \mathcal{E}}(\rho_t) - E(\rho_t)|^2 &= \int_{\mathbb{S}^{d-1}} |v_{\alpha, \mathcal{E}}(\rho_t) - v + v - E|^2 d\rho_t(v) \\ &= \int_{\mathbb{S}^{d-1}} |v_{\alpha, \mathcal{E}}(\rho_t) - v|^2 d\rho_t + \int_{\mathbb{S}^{d-1}} |v - E|^2 d\rho_t + 2 \int_{\mathbb{S}^{d-1}} \langle v_{\alpha, \mathcal{E}}(\rho_t) - v, v - E \rangle d\rho_t \\ &\leq 4C_{\alpha, \mathcal{E}}^2 V(\rho_t) + 2V(\rho_t) + 2E^2 - 2 = (4C_{\alpha, \mathcal{E}}^2 - 2)V(\rho_t), \end{aligned}$$

which completes the proof. \blacksquare

Proof (Lemma 9) The derivative of $\|\omega_{\mathcal{E}}^{\alpha}\|_{L^1(\rho_t)}$ is given by

$$\begin{aligned} \frac{d}{dt} \int_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha}(v) d\rho_t &= \int_{\mathbb{S}^{d-1}} \frac{\sigma^2}{2} |v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 \Delta_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha} \\ &\quad - \lambda (\langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - v_{\alpha, \mathcal{E}}(\rho_t)) \cdot \nabla_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha} d\rho_t \\ &= \int_{\mathbb{S}^{d-1}} \frac{\sigma^2}{2} |v - v_{\alpha, \mathcal{E}}(\rho_t)|^2 \Delta_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha} + \lambda P(v) v_{\alpha, \mathcal{E}}(\rho_t) \cdot \nabla_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha} d\rho_t \\ &=: \mathbf{I} + \mathbf{II}. \end{aligned} \quad (4.6)$$

The gradient and the Laplacian of the weight function can be computed as

$$\nabla_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha}(v) = \nabla \omega_{\mathcal{E}}^{\alpha} \left(\frac{v}{|v|} \right) \Big|_{|v|=1} = \frac{1}{|v|} \left(I - \frac{vv^T}{|v|^2} \right) \nabla \omega_{\mathcal{E}}^{\alpha} \Big|_{|v|=1} = -\alpha e^{-\alpha \mathcal{E}} (I - vv^T) \nabla \mathcal{E} \Big|_{|v|=1} \quad (4.7)$$

and

$$\Delta_{\mathbb{S}^{d-1}} \omega_{\mathcal{E}}^{\alpha}(v) = \Delta \omega_{\mathcal{E}}^{\alpha} \left(\frac{v}{|v|} \right) \Big|_{|v|=1} = \frac{\Delta \omega_{\mathcal{E}}^{\alpha}}{|v|} - (d-1) \frac{v}{|v|^3} \cdot \nabla \omega_{\mathcal{E}}^{\alpha} - \frac{vv^T}{|v|^3} : \nabla^2 \omega_{\mathcal{E}}^{\alpha} \Big|_{|v|=1}. \quad (4.8)$$

We further have

$$\nabla \omega_{\mathcal{E}}^{\alpha} = -\alpha e^{-\alpha \mathcal{E}} \nabla \mathcal{E} \in \mathbb{R}^d; \quad (4.9)$$

$$\nabla^2 \omega_{\mathcal{E}}^{\alpha} = -\alpha e^{-\alpha \mathcal{E}} (-\alpha \nabla \mathcal{E} \otimes \nabla \mathcal{E} + \nabla^2 \mathcal{E}) \in \mathbb{R}^{d \times d}; \quad (4.10)$$

$$\Delta \omega_{\mathcal{E}}^{\alpha} = \alpha^2 e^{-\alpha \mathcal{E}} |\nabla \mathcal{E}|^2 - \alpha e^{-\alpha \mathcal{E}} \Delta \mathcal{E} \in \mathbb{R}. \quad (4.11)$$

We estimate the term \mathbf{I} as follows

$$\mathbf{I} = \frac{\sigma^2}{2} \int |v - v_{\alpha, \mathcal{E}}|^2 \left(\Delta \omega_{\mathcal{E}}^{\alpha} - (d-1) v \cdot \nabla \omega_{\mathcal{E}}^{\alpha} - v \otimes v : \nabla^2 \omega_{\mathcal{E}}^{\alpha} \right) d\rho_t(v)$$

$$\begin{aligned}
 &= \frac{\sigma^2}{2} \int |v - v_{\alpha, \mathcal{E}}|^2 \left[\alpha^2 |\nabla \mathcal{E}|^2 - \alpha \Delta \mathcal{E} + \alpha(d-1)v \cdot \nabla \mathcal{E} \right. \\
 &\quad \left. + \alpha \left(v \otimes v : (-\alpha \nabla \mathcal{E} \otimes \nabla \mathcal{E}) + v \otimes v : \nabla^2 \mathcal{E} \right) \right] e^{-\alpha \mathcal{E}} d\rho_t(v) \\
 &\geq \frac{\sigma^2}{2} \int |v - v_{\alpha, \mathcal{E}}|^2 \left[-\alpha \Delta \mathcal{E} + \alpha(d-1)\nabla \mathcal{E} \cdot v - \alpha^2 |\nabla \mathcal{E}|^2 - \alpha |\nabla^2 \mathcal{E}| \right] e^{-\alpha \mathcal{E}} d\rho_t(v) \\
 &\geq \frac{\sigma^2}{2} \int |v - v_{\alpha, \mathcal{E}}|^2 e^{-\alpha \mathcal{E}} \left[-\alpha c_2 - \alpha(d-1)c_1 - \alpha^2 c_1^2 - \alpha c_2 \right] d\rho_t(v) \\
 &\geq -2\sigma^2 \alpha e^{-2\alpha \underline{\mathcal{E}}} (2c_2 + (d-1)c_1 + \alpha c_1^2) \frac{V(\rho_t)}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}}, \tag{4.12}
 \end{aligned}$$

where we have used that $|\nabla \mathcal{E}| \leq c_1$; $|\Delta \mathcal{E}|, |\nabla^2 \mathcal{E}| \leq c_2$, estimate (4.2) and the property

$$v \otimes v : \nabla \mathcal{E} \otimes \nabla \mathcal{E} = \sum_{i,j} v_i v_j \partial_i \mathcal{E} \partial_j \mathcal{E} \leq \left(\sum_i \partial_i \mathcal{E} \right)^2 \leq |\nabla \mathcal{E}|^2. \tag{4.13}$$

For the term **II** we get

$$\begin{aligned}
 \mathbf{II} &= -\alpha \lambda \int_{\mathbb{S}^{d-1}} e^{-\alpha \mathcal{E}} P(v) v_{\alpha, \mathcal{E}}(\rho_t) \cdot (\nabla \mathcal{E} - v v^T \nabla \mathcal{E}) d\rho_t = -\alpha \lambda \int_{\mathbb{S}^{d-1}} e^{-\alpha \mathcal{E}} P(v) v_{\alpha, \mathcal{E}}(\rho_t) \cdot \nabla \mathcal{E} d\rho_t \\
 &= \alpha \lambda \int_{\mathbb{S}^{d-1}} e^{-\alpha \mathcal{E}} (\langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - v_{\alpha, \mathcal{E}}(\rho_t)) \cdot \nabla \mathcal{E} d\rho_t \\
 &\geq -\alpha \lambda c_1 e^{-\alpha \underline{\mathcal{E}}} \int_{\mathbb{S}^{d-1}} |\langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t, \tag{4.14}
 \end{aligned}$$

where in the second equality we have used the fact that $v \cdot P(v) v_{\alpha, \mathcal{E}}(\rho_t) = 0$. We observe that

$$\begin{aligned}
 \int_{\mathbb{S}^{d-1}} |\langle v_{\alpha, \mathcal{E}}(\rho_t), v \rangle v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t &= \int_{\mathbb{S}^{d-1}} |\langle v_{\alpha, \mathcal{E}}(\rho_t) - v, v \rangle v + v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t \\
 &\leq \int_{\mathbb{S}^{d-1}} |\langle v_{\alpha, \mathcal{E}}(\rho_t) - v, v \rangle v| d\rho_t + \int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t \\
 &\leq 2 \int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t \leq 4 \frac{e^{-\alpha \underline{\mathcal{E}}}}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}} V(\rho_t)^{\frac{1}{2}}, \tag{4.15}
 \end{aligned}$$

where we have used (4.4) in the last inequality. Thus we have

$$\mathbf{II} \geq -2\alpha \lambda c_1 e^{-\alpha \underline{\mathcal{E}}} \int_{\mathbb{S}^{d-1}} |v - v_{\alpha, \mathcal{E}}(\rho_t)| d\rho_t(v) \geq -4\alpha \lambda c_1 e^{-2\alpha \underline{\mathcal{E}}} \frac{V(\rho_t)^{\frac{1}{2}}}{\|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}}. \tag{4.16}$$

Combining the inequalities (4.12) and (4.16) yields

$$\begin{aligned}
 \frac{1}{2} \frac{d}{dt} \|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)}^2 &= \|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)} \frac{d}{dt} \|\omega_\alpha^\mathcal{E}\|_{L^1(\rho_t)} \\
 &\geq -2\sigma^2 \alpha e^{-2\alpha \underline{\mathcal{E}}} (2c_2 + (d-1)c_1 + \alpha c_1^2) V(\rho_t) - 4\alpha \lambda c_1 e^{-2\alpha \underline{\mathcal{E}}} V(\rho_t)^{\frac{1}{2}} \\
 &=: -b_1(d, \sigma, \alpha, c_1, c_2, \underline{\mathcal{E}}) V(\rho_t) - b_2(\alpha, c_1, \underline{\mathcal{E}}) \lambda V(\rho_t)^{\frac{1}{2}}, \tag{4.17}
 \end{aligned}$$

where $b_1, b_2 \rightarrow 0$ as $\alpha \rightarrow \infty$. ■

4.2 Well-posedness and regularity result

Theorem 13 *For any given $T > 0$, let $\rho_0 \in L^2(\mathbb{S}^{d-1})$. Then there exists a unique weak solution ρ to equation (1.10). Moreover it has the following regularity*

$$\rho \in L^\infty([0, T]; L^2(\mathbb{S}^{d-1})) \cap L^2([0, T]; H^1(\mathbb{S}^{d-1})) \text{ and } \partial_t \rho \in L^2([0, T]; H(\mathbb{S}^{d-1})). \quad (4.18)$$

Proof The proof is standard and based on Picard's iteration. We sketch below the details. Let $\rho^0(x, t) \equiv \rho_0(x)$. For $n \geq 0$, let ρ^{n+1} be the unique weak solution to following linear equation

$$\partial_t \rho_t^{n+1} = \lambda \nabla_{\mathbb{S}^{d-1}} \cdot ((\langle v_{\alpha, \varepsilon}(\rho_t^n), v \rangle v - v_{\alpha, \varepsilon}(\rho_t^n)) \rho_t^{n+1}) + \frac{\sigma^2}{2} \Delta_{\mathbb{S}^{d-1}} (|v - v_{\alpha, \varepsilon}(\rho_t^n)|^2 \rho_t^{n+1}), \quad t > 0, \quad (4.19)$$

with the initial data $\rho^{n+1}(x, 0) = \rho_0(x)$ for any given $\rho^n \in L^\infty([0, T]; L^2(\mathbb{S}^{d-1})) \cap L^2([0, T]; H^1(\mathbb{S}^{d-1}))$. For any given $T > 0$ and $t \in [0, T]$, it is easy to compute that

$$\begin{aligned} & \frac{1}{2} \frac{d}{dt} \|\rho_t^{n+1}\|_2^2 + \frac{\sigma^2}{2} \int_{\mathbb{S}^{d-1}} \nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1} \cdot \nabla_{\mathbb{S}^{d-1}} (|v - v_{\alpha, \varepsilon}(\rho_t^n)|^2 \rho_t^{n+1}) dv \\ &= -\lambda \int_{\mathbb{S}^{d-1}} \nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1} \cdot (\langle v_{\alpha, \varepsilon}(\rho_t^n), v \rangle v - v_{\alpha, \varepsilon}(\rho_t^n)) \rho_t^{n+1} dv \leq \lambda \int_{\mathbb{S}^{d-1}} |\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}| |\rho_t^{n+1}| dv. \end{aligned}$$

This lead to

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|\rho_t^{n+1}\|_2^2 &\leq \lambda \int_{\mathbb{S}^{d-1}} |\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}| |\rho_t^{n+1}| dv - \frac{\sigma^2}{2} \int_{\mathbb{S}^{d-1}} |\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}|^2 |v - v_{\alpha, \varepsilon}(\rho_t^n)|^2 dv \\ &\quad - \sigma^2 \int_{\mathbb{S}^{d-1}} \nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1} \cdot (v - v_{\alpha, \varepsilon}(\rho_t^n)) \rho_t^{n+1} dv \\ &\leq -\frac{\sigma^2}{2} \min_{t \in [0, T]} \operatorname{ess\,inf}_{v \in \mathbb{S}^{d-1}} |v - v_{\alpha, \varepsilon}(\rho_t^n)|^2 \|\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}\|_2^2 \\ &\quad + \varepsilon \|\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}\|_2^2 + C(\varepsilon, \sigma, \lambda) \|\rho_t^{n+1}\|_2^2 \\ &\leq C(\varepsilon, \sigma, \lambda) \|\rho_t^{n+1}\|_2^2, \end{aligned}$$

where we have used Hölder's inequality in the second inequality. Applying Gronwall's inequality it yields that

$$\|\rho_t^{n+1}\|_2^2 + \int_0^t \|\nabla_{\mathbb{S}^{d-1}} \rho_t^{n+1}\|_2^2 dt \leq C(T, \sigma, \lambda, \|\rho_0\|_2). \quad (4.20)$$

We also get that for all $\psi \in H^1(\mathbb{S}^{d-1})$

$$\begin{aligned} & \|\partial_t \rho_t^{n+1}\|_{H^1(\mathbb{S}^{d-1})'} = \sup_{\|\psi\|_{H^1} \leq 1} |\langle \partial_t \rho_t^{n+1}, \psi \rangle| \\ & \leq \sup_{\|\psi\|_{H^1} \leq 1} \left| \langle \nabla_{\mathbb{S}^{d-1}} \psi, \lambda (\langle v_{\alpha, \varepsilon}(\rho_t^n), v \rangle v - v_{\alpha, \varepsilon}(\rho_t^n)) \rho_t^{n+1} + \frac{\sigma^2}{2} \nabla_{\mathbb{S}^{d-1}} (|v - v_{\alpha, \varepsilon}(\rho_t^n)|^2 \rho_t^{n+1}) \rangle \right| \\ & \leq C(\lambda, \sigma) \|\rho_t^{n+1}\|_{H^1}. \end{aligned}$$

Thus we obtain $\partial_t \rho^{n+1} \in L^2([0, T]; H(\mathbb{S}^{d-1})')$. Note that this also implies that $\rho^{n+1} \in \mathcal{C}([0, T]; L^2(\mathbb{S}^{d-1}))$ due to the fact that

$$\max_{0 \leq t \leq T} \|\rho^{n+1}\|_2 \leq C(\|\rho^{n+1}\|_{L^2([0, T], H^1)} + \|\partial_t \rho^{n+1}\|_{L^2([0, T]; H(\mathbb{S}^{d-1})'}),$$

where C depends only T . Then by Aubin-Lions lemma, there exists a subsequence ρ^{n_k} and a function $\rho \in L^2([0, T] \times \mathbb{S}^{d-1})$ such that

$$\rho^{n_k} \rightarrow \rho \text{ in } L^2([0, T] \times \mathbb{S}^{d-1}) \text{ as } k \rightarrow \infty. \quad (4.21)$$

To finish the proof of existence we are left to pass the limit and verify ρ is the solution, we omit the details here of this very standard concluding step (see, e.g., (Albi et al., 2017, Theorem 2.4) for similar arguments).

As for the uniqueness, it has been obtained in (Fornasier et al., 2020, Section 2.2 and Section 2.3) by using the uniqueness of the corresponding nonlinear SDE (3.2). \blacksquare

5. Conclusions

We presented the numerical implementation of a new consensus-based model for global optimization on the sphere, which is inspired by the kinetic Kolmogorov-Kuramoto-Vicsek equation. The main result of this paper is about the first and currently unique proof of the convergence of consensus-based optimization to global minimizers provided conditions of well-preparation of the initial datum. We present several numerical experiments in low dimension and synthetic examples in order to illustrate the behavior of the method and we tested the algorithms in high dimension against state of the art methods in a couple of challenging problems in signal processing and machine learning, namely the phase retrieval problem and the robust subspace detection. These experiments show that the algorithm proposed in the present paper scales well with the dimension and is very versatile (one just needs to modify the definition of the function \mathcal{E} and the rest goes with the same code⁵!). The algorithm is able to perform essentially as good as *ad hoc* state of the art methods and in some instances it obtains quantifiably better results. The theoretical rate of convergence is of order N^{-1} in the particle number N and it does not depend on the dimension. Multiplicative constants may depend at most linearly on the dimension d and, as worst case scenario, exponentially in the parameter α . The rate of convergence is exponential and explicitly computable from the parameters of the method, i.e., $\lambda\theta - 2(d-1)e^{\alpha(\bar{\mathcal{E}}-\mathcal{E})}\sigma^2$. The numerical experiments in high dimension ($d \approx 3000$) confirm that the method is in general not affected by curse of dimensionality. Moreover, the requirement of well-preparation of the initial datum (Definition 5) is due to the proving technique we are using based on the monotone decay of the variance. In the case of symmetric cost functions $\mathcal{E}(v) = \mathcal{E}(-v)$, the well-preparation is by no means a severe restriction. We conjecture that with other proving techniques the conditions of well-preparation can be removed, since in the numerical experiments the initialization by uniform distribution yields to global convergence consistently. In our view, this work represents a fundamental theoretical contribution to CBO methods

5. <https://github.com/PhilippeSu/KV-CBO>

on the sphere, on which to build variations of the algorithm with the aim of further improving its complexity and convergence towards the global minimum. A promising perspective in this direction is to consider the introduction of anisotropic noise in order to reduce dependence of the parameters from the dimension and to better explore the search space in case of very high dimensional problems (Carrillo et al., 2021). This and other algorithmic improvements are left to future research.

Acknowledgments

Massimo Fornasier and Hui Huang acknowledge the support of the DFG Project "Identification of Energies from Observation of Evolutions" and the DFG SPP 1962 "Non-smooth and Complementarity-based Distributed Parameter Systems: Simulation and Hierarchical Optimization". The present project and Philippe Sünnen are supported by the National Research Fund, Luxembourg (AFR PhD Project Idea "Mathematical Analysis of Training Neural Networks" 12434809). Lorenzo Pareschi acknowledges the support of the John Von Neumann guest Professorship program of the Technical University of Munich during the preparation of this work. The authors acknowledge the support and the facilities of the LRZ Compute Cloud of the Leibniz Supercomputing Center of the Bavarian Academy of Sciences, on which the numerical experiments of this paper have been tested.

References

- Emile Aarts and Jan Korst. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. 1989.
- Giacomo Albi and Lorenzo Pareschi. Binary interaction algorithms for the simulation of flocking and swarming dynamics. *Multiscale Modeling & Simulation*, 11(1):1–29, 2013.
- Giacomo Albi, Young-Pil Choi, Massimo Fornasier, and Dante Kalise. Mean field control hierarchy. *Applied Mathematics & Optimization*, 76(1):93–135, 2017.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2008.
- Thomas Back, David B. Fogel, and Zbigniew Michalewicz, editors. *Handbook of Evolutionary Computation*. 1997.
- Bubacarr Bah, Holger Rauhut, Ulrich Terstiege, and Michael Westdickenberg. Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. arXiv:1910.05505, 2019.
- Wilma. A. Bainbridge, Philipp Isola, and Aude Oliva. *The Intrinsic Memorability of Face Photographs*. *Journal of Experimental Psychology: General*, 142(4), 1323–1334., 2013.
- Radu Balan, Pete Casazza, and Dan Edidin. On signal reconstruction without phase. *Applied and Computational Harmonic Analysis*, 20(3):345–356, 2006.

- Afonso S Bandeira, Jameson Cahill, Dustin G Mixon, and Aaron A Nelson. Saving phase: Injectivity and stability for phase retrieval. *Applied and Computational Harmonic Analysis*, 37(1):106–125, 2014.
- Alessandro Benfenati, Giacomo Borghi, and Lorenzo Pareschi. Binary interaction methods for high dimensional global optimization and machine learning. arxiv:2105.02695, 2021.
- Yoshua Bengio, Patrice Simard, Paolo Frasconi, et al. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
- Christian Blum and Andrea Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Comput. Surv.*, 35(3):268–308, September 2003. ISSN 0360-0300.
- Aleksandar Mijatović, Veno Mramor, and Gerónimo U. Bravo. A note on the exact simulation of spherical Brownian motion. *Statistics and Probability Letters*, 165.108836, 2020.
- Emmanuel J. Candés, Yonina C. Eldar, Thomas. Strohmer, and Vladislav. Voroninski. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, 2013.
- Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- Yuan Cao and Quanquan Gu. Generalization error bounds of gradient descent for learning over-parameterized deep relu networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3349–3356, Apr. 2020.
- José A Carrillo, Young-Pil Choi, Claudia Totzeck, and Oliver Tse. An analytical framework for consensus-based global optimization method. *Mathematical Models and Methods in Applied Sciences*, 28(06):1037–1066, 2018.
- José A Carrillo, Shi Jin, Lei Li, and Yuhua Zhu. A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM: Control, Optimisation and Calculus of Variations*, 27:S5, 2021.
- Rohan Chandra, Ziyuan Zhong, Justin Hontz, Val McCulloch, Christoph Studer, and Tom Goldstein. Phasepack: A phase retrieval library. pages 1617–1621, 2017.
- Jingrun Chen, Shi Jin, and Liyao Lyu. A consensus-based global optimization method with adaptive momentum estimation. arxiv:2012.04827, 2020.
- Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, 176(1):5–37, 2019.
- Cristina Cipriani, Hui Huang, and Jinniao Qiu. Zero-inertia limit: from particle swarm optimization to consensus based optimization. arXiv:2104.06939, 2021.

- Amir Dembo and Ofer Zeitouni. *Large Deviations Techniques and Applications*. Springer-Verlag Berlin Heidelberg, 2010.
- Marco Dorigo and Christian Blum. Ant colony optimization theory: A survey. *Theoretical computer science*, 344(2-3):243–278, 2005.
- Yonina C Eldar and Shahar Mendelson. Phase retrieval: Stability and recovery guarantees. *Applied and Computational Harmonic Analysis*, 36(3):473–494, 2014.
- Veit Elser, Ti-Yen Lan, and Tamir Bendory. Benchmark problems for phase retrieval. *SIAM Journal on Imaging Sciences*, 11(4):2429–2455, 2018.
- Hugo Jair Escalante, Manuel Montes, and Luis Enrique Sucar. Particle swarm model selection. *Journal of Machine Learning Research*, 10(2), 2009.
- Razvan C Fetecau, Hui Huang, and Weiran Sun. Propagation of chaos for the Keller–Segel equation over bounded domains. *Journal of Differential Equations*, 266(4):2142–2174, 2019.
- James R Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.
- David B. Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence (IEEE Press Series on Computational Intelligence)*. 2006.
- Massimo Fornasier, Hui Huang, Lorenzo Pareschi, and Philippe Sünnen. Consensus-based optimization on hypersurfaces well-posedness and mean-field limit. *Mathematical Models and Methods in Applied Sciences*, 30(14):2725–2751, 2020.
- Massimo Fornasier, Timo Klock, and Konstantin Riedl. Consensus-based optimization methods converge globally in mean-field law. arxiv:2103.15130, 2021.
- Ralph W Gerchberg. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237–246, 1972.
- Fred W Glover and Gary A Kochenberger. *Handbook of metaheuristics*, volume 57. Springer Science & Business Media, 2006.
- Sara Grassi and Lorenzo Pareschi. From particle swarm optimization to consensus based optimization: Stochastic modeling and mean-field limit. *Mathematical Models and Methods in Applied Sciences*, pages 1–33, 2021.
- Robert Großmann, Fernando Peruani, and Markus Bär. A geometric approach to self-propelled motion in isotropic & anisotropic environments. *The European Physical Journal Special Topics*, 224(7):1377–1394, 2015.
- Seung-Yeal Ha, Shi Jin, and Doheon Kim. Convergence of a first-order consensus-based global optimization algorithm. *Mathematical Models and Methods in Applied Sciences*, 30(12):2417–2444, 2020.

- Ernst Hairer, Christian Lubich, and Gerhard Wanner. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*, volume 31. Springer Science & Business Media, 2006.
- Robert W Harrison. Phase problem in crystallography. *JOSA a*, 10(5):1046–1055, 1993.
- Wilfred K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- Desmond J. Higham, Xuerong Mao, and Andrew M. Stuart. Strong convergence of euler-type methods for nonlinear stochastic differential equations. *SIAM Journal on Numerical Analysis*, 40(3):1041–1063, 2002.
- John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. 1992.
- Richard Holley and Daniel Stroock. Simulated annealing via Sobolev inequalities. *Communications in Mathematical Physics*, 115(4):553–569, 1988.
- Hui Huang and Jian-Guo Liu. Error estimate of a random particle blob method for the Keller–Segel equation. *Mathematics of Computation*, 86(308):2719–2744, 2017.
- Hui Huang and Jinniao Qiu. On the mean-field limit for the consensus-based optimization. arxiv:2105.12919, 2021.
- Norman E. Hurt. *Phase Retrieval and Zero Crossings: Mathematical Methods in Image Reconstruction*.
- Adel Javanmard, Marco Mondelli, and Andrea Montanari. Analysis of a two-layer neural network via displacement convexity. *The Annals of Statistics*, 48(6):3619–3642, 2020.
- Shi Jin, Lei Li, and Jian-Guo Liu. Random batch methods (rbm) for interacting particle systems. *Journal of Computational Physics*, 400:108877, 2020.
- James Kennedy. Particle swarm optimization. *Encyclopedia of machine learning*, pages 760–766, 2010.
- Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid strict saddle points. *Mathematical programming*, 176(1):311–337, 2019.
- Gilad Lerman and Tyler Maunu. Fast, robust and non-convex subspace recovery. *Information and Inference: A Journal of the IMA*, 7(2):277–336, 2018.
- Gilad Lerman, Michael B McCoy, Joel A Tropp, and Teng Zhang. Robust computation of linear models by convex relaxation. *Foundations of Computational Mathematics*, 15(2):363–410, 2015.

- Shengchao Liu, Dimitris Papailiopoulos, and Dimitris Achlioptas. Bad global minima exist and sgd can reach them. *Advances in Neural Information Processing Systems*, 33, 2020.
- Michael Herty Lorenzo Pareschi and Giuseppe Visconti. Mean field models for large data-clustering problems. *Network and Heterogeneous Media*, 15(3):463–487, 2020.
- George Marsaglia. Choosing a point from the surface of a sphere. *The Annals of Mathematical Statistics*, 43(2):645–646, 1972.
- Tyler Maunu, Teng Zhang, and Gilad Lerman. A well-tempered landscape for non-convex robust subspace recovery. *Journal of Machine Learning Research*, 20(37), 2019.
- 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.
- Peter David Miller. *Applied Asymptotic Analysis*, volume 75. American Mathematical Soc., 2006.
- Marco Mondelli and Andrea Montanari. Fundamental limits of weak recovery with applications to phase retrieval. *Foundations of Computational Mathematics*, 19(3):703–773, 2019.
- Mervin E Muller. A note on a method for generating points uniformly on n-dimensional spheres. *Communications of the ACM*, 2(4):19–20, 1959.
- John A. Nelder and Roger Mead. A simplex method for function minimization. *Computer Journal*, 7:308–313, 1965.
- René Pinnau, Claudia Totzeck, Oliver Tse, and Stephan Martin. A consensus-based model for global optimization and its mean-field limit. *Mathematical Models and Methods in Applied Sciences*, 27(01):183–204, 2017.
- Eckhard Platen. An introduction to numerical methods for stochastic differential equations. *Acta numerica*, 8:197–246, 1999.
- Riccardo Poli, James Kennedy, and Tim Blackwell. Particle swarm optimization. *Swarm intelligence*, 1(1):33–57, 2007.
- Harry M. Quiney. Coherent diffractive imaging using short wavelength light sources. *Journal of modern optics*, 57(13):1109–1149, 2010.
- L. A. Rastrigin. The convergence of the random search method in the external control of many-parameter system. *Automation and Remote Control*, 24:1337–1342, 1963.
- Alain-Sol Sznitman. Topics in propagation of chaos. In *Ecole d’été de probabilités de Saint-Flour XIX—1989*, pages 165–251. Springer, 1991.
- Adriaan Walther. The question of phase retrieval in optics. *Optica Acta: International Journal of Optics*, 10(1):41–49, 1963.

Guo-zhen Yang, Bi-zhen Dong, Ben-yuan Gu, Jie-yao Zhuang, and Okan K Ersoy.
Gerchberg–Saxton and Yang–Gu algorithms for phase retrieval in a nonunitary transform system: a comparison. *Applied optics*, 33(2):209–218, 1994.