



Constrained consensus-based optimization and numerical heuristics for the few particle regime

Jonas Beddrich¹ · Enis Chenchene² · Massimo Fornasier^{1,3,4}  · Hui Huang⁵ · Barbara Wohlmuth¹

Received: 20 December 2024 / Accepted: 24 March 2026

© The Author(s) 2026

Abstract

Consensus-based optimization (CBO) is a versatile multi-particle optimization method for performing nonconvex and nonsmooth global optimizations in high dimensions. Proofs of global convergence in probability have been achieved for a broad class of objective functions in unconstrained optimizations. In this work we adapt the algorithm for solving constrained optimizations on compact and unbounded domains with boundary by leveraging emerging reflective boundary conditions. In particular, we close a relevant gap in the literature by providing a global convergence proof for the many-particle regime comprehensive of convergence rates. On the one hand, for the sake of minimizing running cost, it is desirable to keep the number of particles small. On the other hand, reducing the number of particles implies a diminished capability of exploration of the algorithm. Hence numerical heuristics are needed to ensure convergence of CBO in the few-particle regime. In this work, we

✉ Massimo Fornasier
massimo.fornasier@cit.tum.de

Jonas Beddrich
beddrich@ma.tum.de

Enis Chenchene
enis.chenchene@univie.ac.at

Hui Huang
huihuang1@hnu.edu.cn

Barbara Wohlmuth
wohlmuth@cit.tum.de

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

² Faculty of Mathematics, University of Vienna, Vienna, Austria

³ Munich Data Science Institute, Technical University of Munich, Garching by Munich, Germany

⁴ Munich Center for Machine Learning, Munich, Germany

⁵ School of Mathematics, Hunan University, Changsha, China

also significantly improve the convergence and complexity of CBO by utilizing an adaptive region control mechanism and by choosing geometry-specific random noise. In particular, by combining a hierarchical noise structure with a multigrid finite element method, we are able to compute global minimizers for a constrained p-Allen–Cahn problem with obstacles, a very challenging variational problem.

Keywords Constrained global optimization · Consensus-based optimization · Heuristics

Mathematics Subject Classification 90C26 · 90C56 · 90C59 · 65K10 · 35Q90 · 60H10

Contents

1	Introduction	
2	Architecture of the proof of the main result	
2.1	Step 1: convergence of the numerical scheme for $\Delta t \rightarrow 0$	
2.2	Step 2: mean-field limit for $N \rightarrow \infty$	
2.3	Step 3: large time analysis of the mean-field model	
3	Main technical challenges	
4	Well-posedness and mean-field limit	
4.1	Well-posedness	
4.2	Mean-field limit	
5	Global convergence to the minimizer	
5.1	Global convergence proof when $\overline{\mathcal{D}}$ is bounded	
5.2	Global convergence proof when $\overline{\mathcal{D}}$ is unbounded	
5.3	Quantitative convergence result for the numerical scheme	
6	Numerical experiments and applications	
6.1	Effect of the parameters	
6.1.1	The effect of λ and σ	
6.1.2	Ablation study on the heuristics	
6.1.3	Quantitative comparisons with respect to N	
6.2	Performance over nonconvex domains	
6.3	Solving a PDE problem with an obstacle	
6.3.1	Hierarchical structured noise	
6.3.2	Obstacle problem	
	Appendix	
	Consensus-based optimization within the global optimization landscape	
	Proof of Proposition 5.2	
	References	

1 Introduction

In this paper we are concerned with theoretical guarantees for global optimization, namely ensuring computability of global minimizers

$$x^* \in \arg \min_{x \in \overline{\mathcal{D}}} \mathcal{E}(x) \quad (1.1)$$

of a possibly nonconvex and nonsmooth objective function $\mathcal{E} : \overline{\mathcal{D}} \subset \mathbb{R}^d \rightarrow \mathbb{R}$. Here $\mathcal{D}$ is a domain in $\mathbb{R}^d$ and $\overline{\mathcal{D}}$ its closure. This problem can be rightfully considered the ultimate challenge of mathematical optimization, because there are no clear directional or geometrical principles to guide the search for global minima [1, 2]. Local optimization techniques like gradient descent or (quasi-)Newton's methods struggle with nonconvex problems because they often get trapped in possibly numerous local minima. Even solving a simple polynomial equation can highlight these limitations, requiring multiple independent runs and possibly randomizing the initial conditions to be able to find all solutions.

Very-many-runs first or second order methods and metaheuristics, such as Simulated Annealing, Genetic Algorithms, Particle Swarm Optimization, Ant Colony Optimization, and other evolution strategies [3, 4], have achieved extraordinary empirical successes in global optimization and have been extensively benchmarked, as evidenced by platforms like COCO [5]. For Simulated Annealing the convergence in probability can be guaranteed by the large time analysis of the corresponding Fokker–Planck equation with vanishing diffusion by leveraging log-Sobolev inequalities as forms of Polyak–Łojasiewicz conditions, see [6] and the older results [7–9]. Unfortunately, the theoretically guaranteed rate of convergence in time $\mathcal{O}(\log \log(t) / \log(t))$ of Simulated Annealing is rather poor, see [6, Theorem 4.1]. Instead, proving the global convergence of the widely used multi-agent methods Genetic Algorithms [10] and Particle Swarm Optimization [11] with quantitative rates has remained an elusive challenge for decades, until recent breakthroughs did connect these to multi-agent dynamics [12–15]. More specifically, the most recent groundbreaking paradigm of Consensus-Based Optimization (CBO) [16] introduced by Pinnau et al. and its first analytical description [15] by Carrillo et al. have introduced the right mathematical tools for attacking the long-standing problem of rigorously proving global convergence for such methods. CBO aims at fusing the cooling strategy of Simulated Annealing [7–9, 17, 18] towards Gibbs equilibria with the space exploration by multiple particles/explorers as in Particle Swarm Optimization [11, 19] by taking advantage of a consensus mechanism as in the Cucker–Smale alignment model [20] or the Hegselmann–Krause opinion formation model [21], in which an average orientation or opinion is obtained from the individual observations. For the case $\mathcal{D} = \mathbb{R}^d$ of unconstrained optimization as first introduced and addressed in [16], the equations defining the iterates X_k^i of the CBO algorithm read ($i = 1, \dots, N$ labelling the particles, and $k = 0, \dots, K$ denoting the iterates)

$$X_{(k+1)\Delta t}^i = X_{k\Delta t}^i - \Delta t \lambda (X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)) + \sqrt{\Delta t} \sigma |X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)| N_k^i(0, 1), \quad (1.2)$$

where $\mathcal{E}$ is the objective function to be minimized, $X_\alpha(\rho_{k\Delta t}^N) = \frac{1}{\sum_{i=1}^N \omega_\alpha^\mathcal{E}(X_{k\Delta t}^i)} \sum_{i=1}^N X_{k\Delta t}^i \omega_\alpha^\mathcal{E}(X_{k\Delta t}^i)$ is called the instantaneous consensus-point, $\omega_\alpha^\mathcal{E}(x) = e^{-\alpha \mathcal{E}(x)}$ is the Gibbs weight, and $N_k^i(0, 1)$ is an independent standard Gaussian random vector. The other constants λ , σ , Δt , α are all positive parameters of the algorithm, describing drift strength, volatility, time discretization, and inverse temperature respectively. The initial optimizers X_0^i are drawn i.i.d. at random according to a given probability distribution ρ_0 . The choice of the weight function




Fig. 1 Collocation of Consensus-Based Optimization (CBO) within the global optimization landscape. Here SA=“Simulated Annealing”, PSO=“Particle Swarm Optimization”, CH=“Consensus Hopping”, ES=“Evolution Strategies”, MMS=“Minimizing Movement Scheme” (or Proximal Point Method), TR=“Trust-Region”, SGD=“Stochastic Gradient Descent”. A full explanation of the connections is reported in the Appendix

$\omega_\alpha^\mathcal{E}(x)$ comes from the well-known Laplace’s principle [22, 23], which states that for any probability measure $\mu \in \mathcal{P}(\mathbb{R}^d)$, there holds

$$\lim_{\alpha \rightarrow \infty} \left(-\frac{1}{\alpha} \log \left(\int_{\mathbb{R}^d} \omega_\alpha^\mathcal{E}(x) \mu(dx) \right) \right) = \inf_{x \in \text{supp}(\mu)} \mathcal{E}(x). \quad (1.3)$$

The algorithm (1.2) has a very simple formulation (it can be implemented in couple of lines of code) with a computational cost of $\mathcal{O}(N)$ and is easily parallelizable, see, e.g., [24, 25]. Moreover, it is based on sole pointwise evaluations of the objective function $\mathcal{E}$, hence, it does not require any higher order information.

Since 2017 CBO has attracted enormous attention in virtue of its elegant formulation, its ability of breaking locality, and its astounding potential of providing a blueprint for proving convergence of other evolution strategies such as local Monte Carlo, Metropolis–Hastings, CMA-ES, see [26], and Particle Swarm Optimization, see [13, 14, 27]. In Figure 1 we illustrate how CBO is collocated within the global optimization landscape, with a full explanation of the relationships with different relevant methods in the Appendix.

Several proofs of global convergence of CBO have been proposed [15, 31–34], prominently the most comprehensive one of one of us and collaborators [28], which provides global convergence in probability with *quantitative* convergence rates depending on $\Delta t \rightarrow 0$, $N \rightarrow \infty$, $\alpha \rightarrow \infty$ under mild conditions on $\mathcal{E}$, namely local Lipschitz continuity and local growth at minimizers. This convergence result is based first on proving that the dynamics for $\Delta t \rightarrow 0$ and $N \rightarrow \infty$ in (1.2) can be described by its mean-field approximation [35–37]

$$d\bar{X}_t = -\lambda (\bar{X}_t - X_\alpha(\rho_t)) dt + \sigma |\bar{X}_t - X_\alpha(\rho_t)| dB_t. \quad (1.4)$$

where $X_\alpha(\rho) = \frac{\int x \omega_\alpha^\mathcal{E}(x) d\rho(x)}{\int \omega_\alpha^\mathcal{E}(x) d\rho(x)}$ and $\rho_t = \text{Law}(\bar{X}_t)$, which fulfills the nonlinear Fokker–Planck equation

$$\partial_t \rho_t = \lambda \operatorname{div} \left((x - X_\alpha(\rho_t)) \rho_t \right) + \frac{\sigma^2}{2} \Delta (|x - X_\alpha(\rho_t)|^2 \rho_t). \quad (1.5)$$

The large time behavior of (1.5) that finally explains the convergence of (1.2) is a particularly challenging problem, because (1.5) does not fulfill a recognizable gradient flow structure and it is not prone to more common techniques such as entropy dissipation methods. For this reason, in [28] an *ad hoc* technique based on a quantitative version of Laplace’s principle, see [28, Proposition 21 and Proposition 23], had to be devised resulting in proving that, for α large enough, the squared Wasserstein distance $W_2^2(\rho_t, \delta_{x^*})$ (here x^* is the assumed unique global minimizer of $\mathcal{E}$) is a natural Lyapunov functional for (1.5) at finite time, with exponential decay

$$W_2^2(\rho_t, \delta_{x^*}) \leq W_2^2(\rho_0, \delta_{x^*}) e^{-(1-\theta)(2\lambda - d\sigma^2)t}, \quad (1.6)$$

for a suitable $\theta \in (0, 1)$ for $t \in [0, T^*]$ and $W_2^2(\rho_{T^*}, \delta_{x^*}) \leq \epsilon$. The meaning of (1.6) is that, remarkably, in the many-particle limit the fundamental mechanism of CBO is to transform a nonconvex optimization problem into the canonical problem of minimizing the squared distance to the global minimizer. Hence, the average trajectories of the particles move precisely in straight lines towards the global minimizer, demonstrating robust convergence even in the presence of a nonconvexity. Standard results of numerical approximation of solutions of SDEs [38] one obtains that the numerical scheme (1.2) converges for $\Delta t \rightarrow 0$ to a system of SDE; a quantitative mean-field limit [36, Theorem 2.6] for $N \rightarrow \infty$ shows that the SDE system converges in suitable sense to the mean-field model (1.4); finally the large time analysis of (1.5) allows to conclude (1.6). The combination of these three approximation levels yields [28, Theorem 3.8], showing convergence in probability of (1.2)

$$\begin{aligned} \mathbb{P} \left(\left| \frac{1}{N} \sum_{i=1}^N X_{T^*}^i - x^* \right|^2 \leq \epsilon \right) &\geq \\ 1 - \left[\epsilon^{-1} (C_{\text{NA}} \Delta t + C_{\text{MFA}} N^{-1} + C_0 e^{-(1-\theta)(2\lambda - d\sigma^2)T^*}) + \delta \right], \end{aligned} \quad (1.7)$$

where ϵ, δ, θ are arbitrarily small positive constants and $T^* = \Delta t K$ is a sufficiently large running time depending on ϵ .

In the meanwhile, the original CBO model [16] has been adapted to address a multitude of different problems, namely solving for multiple minimizers, saddle point problems, equilibria of games, multiobjective optimizations and more. It has been also employed with success in a large spectrum of real-life applications. The related literature is quite vast by now. Rather than attempting to include a necessarily incomplete account of this very fast growing field, we refer to the review paper [39], and to [40] for a more recent and relatively comprehensive report. Below we limit ourselves to recall the relevant results, which are more focused on the specific subject of the present work.

In the following let $\mathcal{D} \subsetneq \mathbb{R}^d$ be an open convex set with the boundary $\partial \mathcal{D}$ and, by $\bar{\mathcal{D}}$ we denote again the closure of $\mathcal{D}$. The scope of the present paper is to close one of a relevant gaps

left in the literature, that is providing a suitable formulation for CBO and related theoretical guarantees to solve global optimization problems of the type (1.1) constrained on convex domains with boundary.

It should be mentioned that the results included in this paper are not the only recent contribution aiming at formulating a CBO method for constrained optimization.

The first appearing work in this direction is [41], where the authors analyze a fully discrete CBO scheme like (1.2) with a correction implemented by an orthogonal projection onto $\overline{\mathcal{D}}$, see also (1.8) below, for the case where the domain is convex. While in the paper the authors offer guarantees of convergence to a consensus point, they do not go so far to provide convergence rates to global minimizers in terms of number N of particles as in (1.7). Successive contributions [42, 43] formulated the constrained optimization problem as a unconstrained optimization with penalty. While [43] is focused on the formulation of the method and numerical results, the convergence proof in mean-field law in [42] is based on the analysis of large time behavior of a Fokker–Planck equation as in (1.5) and makes use of [28, Theorem 3.7]. While in this paper the domain $\mathcal{D}$ does not need to be convex, the result relies on the fundamental assumption that the penalization is exact ([42, Assumption A1]) and, again, it does not achieve convergence rates in terms of number N of particles as in (1.7). The more recent paper [44] contributes with a CBO model for equality constrained optimizations (hence the domain $\mathcal{D}$ is a manifold described globally as zero set of a smooth function) and provides a proof of convergence in mean-field law, following the blueprint of [28, Theorem 3.7]. Finally, in the series of papers [45–47], two of us introduced and analyzed formulations of CBO for optimizations over compact hypersurfaces with no boundary, providing proof of *local* convergence with quantitative estimates for problems on hyperspheres in [46, 47].

In the present paper we analyze the following CBO scheme, which was already introduced in [41]:

Algorithm: CBO on convex domains with boundary

Given a time horizon $T > 0$ and a time discretization $t_0 = 0 < \Delta t < \dots < K \Delta t = T$ of $[0, T]$, we define

$$X_{(k+1)\Delta t}^i = \Pi_{\overline{\mathcal{D}}} \left(X_{k\Delta t}^i - \Delta t \lambda \left(X_{k\Delta t}^i - X_{\alpha}(\rho_{k\Delta t}^N) \right) \right) \quad (1.8)$$

$$+ \sigma D(X_{k\Delta t}^i - X_{\alpha}(\rho_{k\Delta t}^N)) N^i(0, \Delta t)$$

$$X_0^i \sim \rho_0, \quad i = 1, \dots, N, \quad (1.9)$$

where $\{N^i(0, \Delta t)\}_{i=1}^N$ are independent Gaussian random vectors with zero mean and covariance matrix $\Delta t \mathbf{Id}_d$, $D(\cdot)$ denotes the diagonal matrix produced by its argument, and $\Pi_{\overline{\mathcal{D}}}$ is the orthogonal projection onto the closed convex set $\overline{\mathcal{D}}$.

Let us stress at this point that the only relevant assumption on the initial distribution ρ_0 is that it has the 16-th moment bounded (Lemma 4.4). Hence, except for this concentration assumption, we can choose quite freely how the initial particles are picked.

For the main results we need the following assumptions, here summarized informally (for details see Assumptions 1-3):

Assumptions

- (A1) $\mathcal{D}$ is a sufficiently regular *convex* domain (Assumption 1).
- (A2) The objective function $\mathcal{E}$ is locally Lipschitz continuous with at most linear growth of the Lipschitz constant and it enjoys a quadratic growth (Assumption 2).
- (A3) The objective function $\mathcal{E}$ has a polynomial growth in the vicinity of the *unique* minimizer x^* (Assumption 3).

A few comments on the assumptions are in order: condition (A1), in particular, the convexity of the domain, is useful for ensuring well-posedness of the SDE model underlying Algorithm (1.8) and for certain inequalities such as (5.8) to promote convergence. The convexity cannot be easily relaxed because of the need for well-definiteness of orthogonal projections $\Pi_{\overline{\mathcal{D}}}$ in (1.8). In fact, the analysis of SDEs with reflecting boundary conditions in non-convex domains remains a significant mathematical challenge, as standard techniques based on convex projection (the Skorokhod Problem) are not directly applicable. In assumption (A2) we require at most local Lipschitz continuity of $\mathcal{E}$ and its quadratic growth in order to ensure the well-posedness of certain continuous models (2.1) and (2.4) and for proving their mean-field relationships. This latter assumption can be relaxed to other polynomial growth, see, e.g., [36]. Instead, assumption (A3) is crucial to prove global convergence to x^* . While we currently restrict our analysis to the case of a unique global minimizer x^* , owing to the quantitative Laplace principle presented in Proposition 5.1, this assumption can be relaxed as done in recent papers [34, Theorem 4.9], [48, Theorem 3.22] by utilizing the qualitative Laplace principle instead. Note here that all these models can find only one of the global minimizers, depending on the initial particle distribution. To simultaneously identify multiple global minimizers, specially designed CBO models with localizing kernels have been introduced in [40, 49]. Examples of functions fulfilling Assumption (A2) are the objectives of lasso and ridge regression, or empirical risk functions with, for instance, the least squares loss and weight decay. Moreover, several standard benchmark functions in optimization [50], such as the nonconvex Rastrigin or Ackley function obey also Assumption (A3).

This paper contributes with the following advances:

- First theoretical guarantees for the scheme (1.8) for constrained CBO on convex domains;
- New constrained mean-field formulation enabling Wasserstein-type contractivity with reflective boundaries;
- Clarification that boundaries do not destroy exponential convergence, even when minimizers lie on the boundary;

- Extension of the quantitative Laplace principle to bounded domains, allowing treatment of boundary minimizers;
- Simpler convergence proofs than existing ones of the unconstrained theory.

More details about novelties are provided in Section 3. We summarize our main result as follows:

Theorem 1.1 *Assume (A1)–(A3) and fix $\varepsilon > 0$, $2\lambda > \sigma^2$, and $\alpha > 0$ large enough (depending on ε and ρ_0). Let $\{(X_{k\Delta t}^i)_{k=1,\dots,K}\}_{i=1}^N$ be the iterations generated by the Euler–Maruyama scheme (1.8), where $K\Delta t = T_*$, and T_* is large enough (depending on $\varepsilon > 0$ and $\alpha > 0$). Then the final iterations fulfill the following quantitative error estimate*

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N X_{K\Delta t}^i - x^* \right|^2 \right] \leq 3C_{NA}\Delta t \log(1/\Delta t) + 3C_{MFA} \frac{1}{N} + 3\varepsilon, \quad (1.10)$$

where C_{MFA} comes from Theorem 4.3, and C_{NA} depends on λ , σ , α , d , T_* , N and ε .

Remark 1.1 A few comments about the assumptions and quantitative error bound are in order:

- The stability assumption $2\lambda > \sigma^2$ is essentially sharp for unconstrained minimizations and the large particle regime and not an artifact of the theory; on bounded domains numerical experiments show that it can be slightly violated, see Figure 4 in Section 6.1.1;
- The time T_* depends on ε and α and increases as ε vanishes and α diverges; these dependencies are determined in Theorem 5.4;
- Given the generality of the objective functions we consider (Assumption (A2)), which includes problems intrinsically affected by the curse of dimension, we cannot exclude a super-polynomial dependency of the constants C_{NA} and C_{MFA} on the dimension d . However, in multiple optimization problems of practical relevance in high dimension, CBO performs far better than other state of the art methods such as CMA-ES (see for instance recent breakthrough experiments in humanoid robots [51]);
- Connected to the previous point, how large N needs to be for the quantitative bound to be meaningful (1.10) is problem/objective dependent, through the constant C_{MFA} , for which we can provide only a worst case analysis in as Theorem 4.3.

This result goes beyond the above mentioned previous contributions [41, 42, 44] as it provides a *quantitative rate of convergence* with respect to relevant parameters of the algorithm, namely number of iterations K , time discretization step Δt , and number of particles N . Moreover, differently from [42], we do not need to impose exact penalizations. Moreover, while our previous work [45–47] for optimizations on hyperspheres obtained results of local convergence, Theorem 1.1 ensures global convergence in probability for problems on convex domains with boundary. Indeed, a simple application of Markov inequality allows to obtain from (1.10) convergence in probability with a quantitative estimate of the type (1.7), hence we do not stress it here further.

The quantitative estimate (1.10) does explain well the observed behavior of the algorithm, as we will show in Section 6.1 with numerical experiments. Nevertheless, it is particularly meaningful for the many-particle regime $N \gg 1$, especially because the constant C_{MFA} shows exponential dependence on α and, depending on the objective function $\mathcal{E}$, this may be reflected in an exponential dependence on the dimension d , see, e.g., [28].

On the one hand, for the sake of minimizing running cost, it is certainly desirable to keep the number of particles quite small instead. On the other hand, reducing the number of particles implies a diminished capability of the algorithm of exploration the optimization domain and forming a consensus on the location of the global minimizer. Hence numerical heuristics are needed to ensure that CBO in the few-particle regime mimics and behaves as the many-particle one.

As the second relevant contribution of this paper, in this work we also significantly improve the convergence and complexity of CBO in the few-particle regime by utilizing an adaptive region control mechanism and by choosing geometry-specific random noise. In particular, we show that these new heuristics are extremely powerful, allowing to solve the global optimization of the Rastrigin function in dimension $d = 100$ with or without constraints, a problem that was so far completely out of reach of CBO methods, as it was becoming extremely challenging for CBO already for moderated dimensions ($d = 20$), see, e.g., [16, 42, 52, 53]. It remains an open and very challenging problem to provide theoretical guarantees for such heuristics for the few-particle regime.

Moreover, in order to test our novel approach beyond standard benchmark optimizations much tested in previous work, see, e.g., [16, 42, 52, 53], in this paper we apply for the first time CBO for solving a challenging problem in scientific computing. In particular, by integrating a hierarchical noise structure in a multigrid finite element method, we demonstrate how the CBO scheme (1.8) can be effectively applied to compute global minimizers for a constrained p -Allen–Cahn problem, both with and without obstacles. Given the existence of continua of (local) minimizers of p -Allen–Cahn energies, such global variational problems are extremely challenging. We consider this successful test a breakthrough that paves the way for further applications of CBO in scientific computing, highlighting once more its versatility and ease of implementation. A complete theoretical study of these numerical results is left for future research.

Summary of results

The results of this paper are split into two specific directions:

- (i) A theoretical analysis of the Algorithm (1.8) as fully described in the Theorem 1.1 under the Assumptions (A1)–(A3). The results are meaningful for the number of particles N large.
- ii) A set of heuristics for the few particle regime, which are not covered by the theory. Namely enhanced volatility to promote exploration, and projection onto shrinking balls to promote concentration.

The paper is organized as follows. In Section 2 we provide a concise explanation of the architecture of the proof of Theorem 1.1, which is based on the asymptotic approximation by auxiliary mean-field models for $\Delta t \rightarrow 0$ and $N \rightarrow \infty$. In Section 4 we address the well-posedness of such models and we quantify the approximation rate for $N \rightarrow \infty$, by establishing a quantitative mean-field limit. Section 5 is dedicated then to the proof of Theorem 1.1. In Section 6 we introduce new heuristics to improve the complexity for the few-particle regime and we illustrate their efficacy in solving constrained optimizations for well-known benchmark cases and for a constrained p -Allen–Cahn problem with obstacles.

2 Architecture of the proof of the main result

While the proof of Theorem 1.1 builds on the blueprint of [28], it requires nevertheless major and nontrivial technical adaptations to deal with the boundary of $\mathcal{D}$. In particular, thanks to the convexity of $\mathcal{D}$ and the emergence of reflecting boundary conditions (see below), we obtain inequalities, such as (5.8), which play at our favor in crucial estimates. Moreover, we need to devise new observations, such as Proposition 5.2 and Proposition 5.5, to deal with the case where the minimizer x^* may lay precisely on the boundary, a situation, which is actually quite common in practical scenarios. See more details in Section 3 on the relevance of the analytic advances of this paper and their insights.

Now, to make sense of Theorem 1.1 and how it is obtained, we need to introduce further notations. For $x \in \partial\mathcal{D}$, we denote by $\mathcal{H}_x$ the set of all supporting hyperplanes of $\mathcal{D}$ at x . By an outward normal vector $n(x)$ at $x \in \partial\mathcal{D}$ one means any outward unit vector perpendicular to some $H \in \mathcal{H}_x$. Moreover, we denote by $\mathcal{N}_x$ the set of all outward unit normal vectors at $x \in \partial\mathcal{D}$. Indeed, there is a possibility that $\#(\mathcal{N}_x) = \infty$ if the boundary $\partial\mathcal{D}$ lacks smoothness in the vicinity of x . In our current study, we do not require the boundary smoothness condition, thanks to the results established by Tanaka [54]. For later use let us denote $[N] := \{1, \dots, N\}$. In the following, we outline the architecture of the proof of our main result Theorem 1.1. The proof is divided into three steps, and the results obtained in each step are summarized in Fig. 2.

2.1 Step 1: convergence of the numerical scheme for $\Delta t \rightarrow 0$

The proof of Theorem 1.1 is based on observing first that the iterations (1.8) converge in suitable sense for $\Delta t \rightarrow 0$ to the solution of the following Skorokhod stochastic differential equations (SDEs)

$$\begin{cases} dX_t^i = -\lambda(X_t^i - X_\alpha(\rho_t^N))dt + \sigma D(X_t^i - X_\alpha(\rho_t^N))dB_t^i - dL_t^i, & (2.1a) \\ L_t^i = \int_0^t n(X_s^i)d|L^i|_s, \quad |L^i|_t = \int_0^t \mathbf{1}_{\partial\mathcal{D}}(X_s^i)d|L^i|_s, \quad i \in [N] & (2.1b) \end{cases}$$




Fig. 2 Scheme of results: in Section 4 we address the well-posedness of (2.1) and (2.4) and their relationship via mean-field limit $N \rightarrow \infty$. In Section 5 we analyze the large time behavior and we derive the global convergence of the numerical scheme (1.8)

where $n(X_s^i) \in \mathcal{N}_{X_s^i}$ if $X_s^i \in \partial\mathcal{D}$, $\{B_{t \geq 0}^i\}_{i=1}^N$ are N independent Brownian motions and $\{L_{t \geq 0}^i\}_{i=1}^N$ are continuous reflecting processes associated to $\{X_{t \geq 0}^i\}_{i=1}^N$ with bounded total variation, which prevent particles leaving the domain. Moreover, $|L|_t^i$ denotes the total variation of L_s^i on $[0, t]$, namely

$$|L|_t^i = \sup \sum_k |L_{t_k}^i - L_{t_{k-1}}^i|, \quad (2.2)$$

where the supremum is taken over all partitions such that $0 = t_0 < t_1 < \dots < t_n = t$. Here, again the consensus point is defined as

$$X_\alpha(\rho_t^N) = \frac{\int_{\mathcal{D}} x \omega_\alpha^\mathcal{E}(x) \rho^N(t, dx)}{\int_{\mathcal{D}} \omega_\alpha^\mathcal{E}(x) \rho^N(t, dx)}, \quad \rho^N(t, dx) = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i} dx. \quad (2.3)$$

2.2 Step 2: mean-field limit for $N \rightarrow \infty$

As $N \rightarrow \infty$, we further prove that the CBO dynamics (2.1) will well approximate the solution of the following mean-field kinetic McKean–Vlasov type equation

$$\begin{cases} d\bar{X}_t = -\lambda(\bar{X}_t - X_\alpha(\rho_t))dt + \sigma D(\bar{X}_t - X_\alpha(\rho_t))dB_t - d\bar{L}_t, \\ \bar{L}_t = \int_0^t n(\bar{X}_s) d|\bar{L}|_s, \quad |\bar{L}|_t = \int_0^t \mathbf{I}_{\partial\mathcal{D}}(\bar{X}_s) d|\bar{L}|_s, \end{cases} \quad (2.4a)$$

$$\quad (2.4b)$$

where

$$X_\alpha(\rho_t) = \frac{\int_{\mathcal{D}} x \omega_\alpha^\mathcal{E}(x) \rho(t, dx)}{\int_{\mathcal{D}} \omega_\alpha^\mathcal{E}(x) \rho(t, dx)}, \quad (2.5)$$

with $\rho(t, x)$ being required to be the law of $\bar{X}_t$, which makes the set of equations (2.4) nonlinear and self-consistent.

Equations (2.1) and (2.4) are known as Skorokhod SDEs, analogous to the one-dimensional case initially studied by Skorokhod in [55]. The extension of Skorokhod SDEs to domains beyond just a half-line or half-space was first addressed by Tanaka in [54], where the domain $\mathcal{D}$ is assumed to be convex. This convexity condition was subsequently relaxed by Lions and Sznitman in [56] together with an admissibility condition that essentially requires $\mathcal{D}$ to be approximated in some way by smooth domains. This admissibility condition was later

eliminated by Saisho in [57]. Finally, for a comprehensive review of stochastic differential equations with reflection, we refer the reader to [58].

A direct application of Itô's formula, the law $\rho_t := \rho(t, \cdot)$ of $\bar{X}$ at time t is a weak solution to the following nonlinear Vlasov–Fokker–Planck equation with non-flux boundary condition

$$\left\{ \begin{array}{l} \partial_t \rho_t = \lambda \nabla \cdot ((x - X_\alpha(\rho_t)) \rho_t) + \frac{\sigma^2}{2} \sum_{k=1}^d \partial_{x_k, x_k}^2 ((x - X_\alpha(\rho_t))_k^2 \rho_t), \quad x \in \mathcal{D}, t > 0 \\ \langle (x - X_\alpha(\rho_t)) \rho_t, n(x) \rangle + \sum_{k=1}^d \partial_{x_k} ((x - X_\alpha(\rho_t))_k^2 \rho_t) n_k(x) = 0, \quad x \in \partial \mathcal{D}. \end{array} \right. \quad (2.6a)$$

with the initial data $\rho_0(x) = \text{Law}(\bar{X}_0)$. Indeed for any test function $\phi \in C^2(\bar{\mathcal{D}})$ with $\partial_{x_k} \phi(x) n_k(x) = 0$ on $\partial \mathcal{D}$ for all $k = 1, \dots, d$, using Itô's lemma deduces

$$\begin{aligned} d\phi(\bar{X}_t) = & \nabla \phi(\bar{X}_t) \cdot [-\lambda(\bar{X}_t - X_\alpha(\rho_t))dt + \sigma D(X_t - X_\alpha(\rho_t))dB_t] \\ & - \nabla \phi(\bar{X}_t) \cdot n(\bar{X}_t) \mathbf{I}_{\partial \mathcal{D}}(\bar{X}_t) d|\bar{L}|_t + \frac{\sigma^2}{2} \partial_{x_k x_k}^2 \phi(\bar{X}_t) (X_t^i - X_\alpha(\rho_t))_k^2 dt. \end{aligned} \quad (2.7)$$

Then taking expectation on both sides implies that $\rho_t(x) = \text{Law}(\bar{X}_t)$ is a weak solution to (2.6).

2.3 Step 3: large time analysis of the mean-field model

The combination of the above mentioned quantitative approximations for $\Delta t \rightarrow 0$, $N \rightarrow \infty$, and the analysis of the large time behavior of the solution ρ_t of (2.6) or its underlying process $\bar{X}_t$ of (2.4) yields the estimate (1.10) of Theorem 1.1. The details of the proof follow in Section 5.

During the revision of this paper, we got aware of the independently developed work [59], which addressed the same class of mean-field models as (2.4). While well-posedness and particle approximation of the model were developed also in this latter paper, the proof of convergence to minimizers for CBO is based on a variance analysis, which requires bounded objectives f , bounded domain $\mathcal{D}$, and have a suboptimal scaling of the parameters λ, σ depending on α (compare the condition $2\lambda > \sigma^2(1 + e^{2\alpha(f_{\max} - f_{\min})})$ in [59, Theorem 5.2] with the optimal one $2\lambda > \sigma^2$ as in our Theorem 1.1). Moreover, in contrast to Theorem 1.1, in [59] no proof of convergence of the numerical scheme (1.8) was achieved.

3 Main technical challenges

This paper provides the first rigorous proof of *global* convergence for CBO in convex domains with boundaries by developing a new mean-field formulation and analytic framework that extends Wasserstein contractivity and the quantitative Laplace principle to constrained set-

tings. The main insights of our novel analysis is that the presence of a boundary does not destroy the exponential-in-time contractivity of CBO towards global minimizers and that the algorithm (1.8) is precisely the right one to discretize the dynamics. In particular, we aim to address a significant gap in the literature: although the numerical scheme (1.8) has been known since its initial formulation in [41], no proof of global convergence has been provided to date. This gap arises for two main reasons. First, a proper analytical approach was developed only very recently in [28] for unconstrained optimization. Second, the approach in [28] relies on contractivity estimates involving the Wasserstein distance, as in (1.6). While this argument works well for unconstrained optimization, it has remained elusive how to adapt it to domains with boundaries, and how the presence of boundaries would affect estimates involving Wasserstein-like distances. To emphasize such difficulty, let us reiterate that in all existing CBO literature on constrained optimization, authors have consistently sought to avoid dealing directly with boundaries [41–47]. (The only exception is the recent work [59] already discussed above at the end of the previous section.) One of the main contributions of this paper is the identification of a proper formulation of the mean-field approximation (2.4), which enables inequalities such as (5.8) that are crucial in estimates involving Wasserstein-like distances, such as (5.4). Without finding the right formulation and carefully developing the entire argument, it is by no means obvious that adapting the system to reflective boundary conditions would permit contractivity of the functional (5.4), making this a central and non-trivial achievement of our work. Indeed, for minimizers precisely at the boundary, one could have been concerned that reflecting boundary conditions could contribute to counteract the convergence, an effect that the theoretical results of the present paper eventually excludes, at least for convex domains (for nonconvex domains, this issue remains open).

Another central technical tool used in the paper [28] is the so-called *quantitative Laplace principle*, which was again formulated to work for problems defined on the entire Euclidean space. Proposition 5.1, Proposition 5.2, Proposition 5.5, and Corollary 5.3 finally allow for the adaptation to bounded domains. Without such a proper formulation, it would have not been possible to deal with the analysis for the global minimizers laying precisely at the boundary of the domain, a situation very common in optimization.

As a further contribution, the proofs of convergence as in Theorem 5.4 and Theorem 5.7 are simpler compared to the formulation in [28], offering a more straightforward approach to further developments.

Although these theoretical results are already significant advances for the reliable use of CBO for constrained optimization, we felt interesting to explore how flexible and effective CBO can be for constrained optimization, by pushing a bit further the boundaries of numerical experiments beyond classical benchmarks [50] as reported so many times in other previous studies [16, 42, 52, 53, 59]. Therefore, we introduce novel heuristics tailored to few-particle regimes, including adaptive region control and geometry-dependent noise. The numerical experiments we report demonstrate that, even with a small number of particles, with such heuristics CBO can successfully solve problems that previous studies could not address. We also explore new challenging applications, such as the reliable computation of ground states

of nonconvex energies, and show for the first time in the CBO literature that the p -Allen–Cahn variational problem with obstacles can be solved efficiently and robustly, highlighting the adaptability of CBO to various geometries and problems.

4 Well-posedness and mean-field limit

Let us now study in details the well-posedness of the models (2.3)–(2.4) and their mean-field relationship. First, we introduce the Condition (B) from [54] for a convex domain $\mathcal{D}$.

Assumption 1 There exists $\varepsilon > 0$ and $\delta > 0$ such that for any $x \in \partial D$ we can find an open ball $B_\varepsilon(x_0) = \{y \in \mathbb{R}^d : |y - x_0| < \varepsilon\}$ satisfying $B_\varepsilon(x_0) \subset \mathcal{D}$ and $|x - x_0| \leq \delta$.

One can see that the above assumption always holds if $\mathcal{D}$ is bounded or if $d = 2$.

Assumption 2 Throughout this section we are interested in the objective function $\mathcal{E} \in \mathcal{C}(\overline{\mathcal{D}})$, for which

1. it holds that $\underline{\mathcal{E}} \leq \mathcal{E}(x)$ for all $x \in \overline{\mathcal{D}}$ and satisfies

$$\forall x, y \in \overline{\mathcal{D}}, \quad |\mathcal{E}(x) - \mathcal{E}(y)| \leq L_{\mathcal{E}}(1 + |x| + |y|)|x - y|; \quad (4.1)$$

2. there exist some constants $c_u, c_\ell, C_u, C_\ell > 0$ such that

$$\begin{aligned} \forall x \in \overline{\mathcal{D}}, \quad c_\ell |x|^2 - C_\ell &\leq \mathcal{E}(x) - \underline{\mathcal{E}} \text{ and} \\ \forall x \in \overline{\mathcal{D}}, \quad \mathcal{E}(x) - \underline{\mathcal{E}} &\leq c_u |x|^2 + C_u. \end{aligned} \quad (4.2)$$

4.1 Well-posedness

To keep the notation concise in what follows, let us denote the state vector of the entire particle system (2.1) by $\mathbf{X} \in \mathcal{C}([0, \infty), \mathbb{R}^{Nd})$ with $\mathbf{X}(t) = \mathbf{X}_t = ((X_t^1)^T, \dots, (X_t^N)^T)^T$ for every $t \geq 0$. Equation (2.1) can then be reformulated as

$$d\mathbf{X}_t = -\text{diag}(\lambda)\mathbf{F}(\mathbf{X}_t)dt + \text{diag}(\sigma)\mathbf{M}(\mathbf{X}_t)d\mathbf{B}_t - d\mathbf{L}_t \quad (4.3)$$

with $(\mathbf{B}_t)_{t \geq 0}$ being a standard Brownian motion in $\mathbb{R}^{Nd}$ and definitions

$$\begin{aligned} \mathbf{L}_t &:= ((L_t^1)^T, \dots, (L_t^N)^T)^T \\ \mathbf{F}(\mathbf{X}_t) &:= (F^1(\mathbf{X}_t)^T, \dots, F^N(\mathbf{X}_t)^T)^T \quad \text{with} \quad F^i(\mathbf{X}_t) = (X_t^i - X_\alpha(\rho_t^N)) \in \mathbb{R}^d \\ \mathbf{M}(\mathbf{X}_t) &:= \text{diag}(M^1(\mathbf{X}_t), \dots, M^N(\mathbf{X}_t)) \quad \text{with} \quad M^i(\mathbf{X}_t) = D(X_t^i - X_\alpha(\rho_t^N)) \in \mathbb{R}^d. \end{aligned}$$

The diag operator in the definition of $\mathbf{M}$ maps the input matrices onto a block-diagonal matrix with them as its diagonal, and $\text{diag}(\lambda)$ and $\text{diag}(\sigma)$ are $Nd \times Nd$ -dimensional diagonal matrices, whose entries are λ and σ .

Lemma 4.1 Assume the cost function $\mathcal{E}$ satisfies Assumption 2-(1). Let $N \in \mathbb{N}$, $\alpha > 0$ be arbitrary. Then for any $\mathbf{X}, \widehat{\mathbf{X}} \in \overline{\mathcal{D}}^N$ satisfying $|\mathbf{X}|, |\widehat{\mathbf{X}}| \leq R$, it holds that

$$|F^i(\mathbf{X})| \leq 2|X^i| + |\mathbf{X}| \leq 2|\mathbf{X}| \quad (4.4)$$

and

$$|F^i(\mathbf{X}) - F^i(\widehat{\mathbf{X}})| \leq 2(2 + RL_{\mathcal{E}}(1 + 2R)\alpha \exp(\alpha(\bar{\mathcal{E}}_R - \underline{\mathcal{E}}_R))) |\mathbf{X} - \widehat{\mathbf{X}}|_{\infty}, \quad (4.5)$$

all $i \in [N]$, where $L_{\mathcal{E}}$ comes from (4.1), $\bar{\mathcal{E}}_R := \max_{|x| \leq R} \mathcal{E}(x)$, $\underline{\mathcal{E}}_R := \min_{|x| \leq R} \mathcal{E}(x)$, $|\mathbf{X}|_{\infty} := \max_{i=1, \dots, N} |X^i|$ for any $\mathbf{X} \in \mathbb{R}^{Nd}$, and $|\cdot|$ denotes the standard Euclidean norm.

Proof Recall that,

$$F^i(\mathbf{X}) = X^i - \frac{\sum_{j=1}^N (X^i - X^j) \omega_{\alpha}^{\mathcal{E}}(X^j)}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(X^j)}, \quad (4.6)$$

then the estimate (4.4) follows directly. Now we prove (4.5). For any $i \in [N]$, we have

$$\begin{aligned} F^i(\mathbf{X}) - F^i(\widehat{\mathbf{X}}) &= X^i - \widehat{X}^i + \frac{\sum_{j=1}^N (X^i - X^j) \omega_{\alpha}^{\mathcal{E}}(X^j)}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(X^j)} - \frac{\sum_{j=1}^N (\widehat{X}^i - \widehat{X}^j) \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j)}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j)} \\ &= X^i - \widehat{X}^i + \frac{\sum_{j=1}^N (X^i - \widehat{X}^i + \widehat{X}^j - X^j) \omega_{\alpha}^{\mathcal{E}}(X^j)}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(X^j)} + \frac{\sum_{j=1}^N (\widehat{X}^i - \widehat{X}^j) (\omega_{\alpha}^{\mathcal{E}}(X^j) - \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j))}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(X^j)} \\ &\quad + \sum_{j=1}^N (\widehat{X}^i - \widehat{X}^j) \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j) \frac{\sum_{j=1}^N (\omega_{\alpha}^{\mathcal{E}}(X^j) - \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j))}{\sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(X^j) \sum_{j=1}^N \omega_{\alpha}^{\mathcal{E}}(\widehat{X}^j)} \\ &=: X^i - \widehat{X}^i + I_1 + I_2 + I_3. \end{aligned} \quad (4.7)$$

It is not difficult to see that

$$|I_1| \leq |X^i - \widehat{X}^i| + |\mathbf{X} - \widehat{\mathbf{X}}|_{\infty} \leq 2|\mathbf{X} - \widehat{\mathbf{X}}|_{\infty}. \quad (4.8)$$

Using assumption (4.1) one can further compute that

$$\begin{aligned} |I_2| &\leq \frac{\sum_{j=1}^N |\widehat{X}^i - \widehat{X}^j| \alpha \exp(-\alpha((1-c)\mathcal{E}(X^j) + c\mathcal{E}(\widehat{X}^j))) L_{\mathcal{E}}(1 + |X^j| + |\widehat{X}^j|) |X^j - \widehat{X}^j|}{N \exp(-\alpha \bar{\mathcal{E}}_R)} \\ &\leq 2RL_{\mathcal{E}}(1 + 2R)\alpha \exp(\alpha(\bar{\mathcal{E}}_R - \underline{\mathcal{E}}_R)) \frac{\sum_{j=1}^N |X^j - \widehat{X}^j|}{N} \\ &\leq 2RL_{\mathcal{E}}(1 + 2R)\alpha \exp(\alpha(\bar{\mathcal{E}}_R - \underline{\mathcal{E}}_R)) |\mathbf{X} - \widehat{\mathbf{X}}|_{\infty}, \end{aligned} \quad (4.9)$$

where $c \in [0, 1]$ comes from the mean-value theorem. Similarly, we also have

$$|I_3| \leq 2RL_{\mathcal{E}}(1 + 2R)\alpha \exp(\alpha(\bar{\mathcal{E}}_R - \underline{\mathcal{E}}_R)) |\mathbf{X} - \widehat{\mathbf{X}}|_{\infty}. \quad (4.10)$$

Putting all the above terms together yields the required estimate. $\square$

Now for any fixed number N of particles, one obtains the well-posedness of the particle system (4.3).

Theorem 4.1 Let $\mathcal{D}$ satisfy Assumption 1 and $\mathcal{E}$ satisfy Assumption 2-(1). Then for each $N \in \mathbb{N}$ and any initial data $\{X_0^i\}_{i=1}^N \in \overline{\mathcal{D}}^N$ satisfying $\mathbb{E}[|X_0|^2] < \infty$, there exists a pathwise unique solution $(\mathbf{X}_t, \mathbf{L}_t)_{t \in [0, T]}$ to the particle system (2.1) or (4.3) for any $T > 0$, and it satisfies $\mathbb{E}[\sup_{0 \leq t \leq T} |\mathbf{X}_t|^2] < \infty$.

Proof Lemma 4.1 deduce that $F^i, i \in [N]$ is locally Lipschitz continuous and has sublinear growth. Consequently, both $\mathbf{F}$ and $\mathbf{M}$ are locally Lipschitz continuous and have sublinear growth. If additionally we have the non-explosion criterion [60, Theorem 3.5] or [61, Theorem 3.1], then it guarantees the well-posedness of the particle system (2.1) or (4.3) according to [54, Theorem 4.1]. Indeed, applying Itô's formula leads to

$$d|X_t^i - X_0^i|^2 = 2(X_t^i - X_0^i) \cdot [-\lambda(X_t^i - X_\alpha(\rho_t^N))dt + \sigma D(X_t^i - X_\alpha(\rho_t^N))dB_t^i - dL_t^i] + \sigma^2 |X_t^i - X_\alpha(\rho_t^N)|^2 dt. \quad (4.11)$$

Using the convexity of $\mathcal{D}$ one has

$$(X_t^i - X_0^i) \cdot dL_t^i = (X_t^i - X_0^i) \cdot n(X_t^i) \mathbf{I}_{\partial \mathcal{D}}(X_t^i) d|L_t^i| \geq 0. \quad (4.12)$$

Then (4.4) implies that for all $\tau > 0$

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq t \leq \tau} |X_t^i - X_0^i|^2 \right] &\leq \lambda \int_0^\tau \mathbb{E}[|X_s^i - X_0^i|^2] + \mathbb{E}[|F^i(\mathbf{X}_s)|^2] ds + \sigma^2 \int_0^\tau \mathbb{E}[|F^i(\mathbf{X}_s)|^2] ds \\ &\quad + 2\sigma \mathbb{E} \left[\sup_{0 \leq t \leq \tau} \left| \int_0^t (X_s^i - X_0^i) \cdot D(X_s^i - X_\alpha(\rho_s^N)) dB_s^i \right| \right] \\ &\leq C \int_0^\tau \mathbb{E}[|X_s^i - X_0^i|^2] ds + C \int_0^\tau \mathbb{E}[|\mathbf{X}_s|^2] ds \\ &\quad + 2\sigma \mathbb{E} \left[\sup_{0 \leq t \leq \tau} \left| \int_0^t (X_s^i - X_0^i) \cdot D(X_s^i - X_\alpha(\rho_s^N)) dB_s^i \right| \right]. \end{aligned}$$

It follows from Burkholder–Davis–Gundy inequality that

$$\begin{aligned} &\mathbb{E} \left[\sup_{0 \leq t \leq \tau} \left| \int_0^t (X_s^i - X_0^i) \cdot D(X_s^i - X_\alpha(\rho_s^N)) dB_s^i \right| \right] \\ &\leq C \mathbb{E} \left[\left(\int_0^\tau |X_s^i - X_0^i|^2 |F^i(\mathbf{X}_s)|^2 ds \right)^{1/2} \right] \\ &\leq C \mathbb{E} \left[\left(\sup_{0 \leq t \leq \tau} |X_t^i - X_0^i|^2 \int_0^\tau |F^i(\mathbf{X}_s)|^2 ds \right)^{1/2} \right] \\ &\leq \frac{1}{2} \mathbb{E} \left[\sup_{0 \leq t \leq \tau} |X_t^i - X_0^i|^2 \right] + C \int_0^\tau \mathbb{E}[|F^i(\mathbf{X}_s)|^2] ds \\ &\leq \frac{1}{2} \mathbb{E} \left[\sup_{0 \leq t \leq \tau} |X_t^i - X_0^i|^2 \right] + C \int_0^\tau \mathbb{E}[|\mathbf{X}_s|^2] ds. \end{aligned}$$

This yields that

$$\mathbb{E} \left[\sup_{0 \leq t \leq \tau} |X_t^i - X_0^i|^2 \right] \leq C \int_0^\tau \mathbb{E}[|X_s^i - X_0^i|^2] ds + C \int_0^\tau \mathbb{E}[|\mathbf{X}_s - \mathbf{X}_0|^2] ds + C \mathbb{E}[|\mathbf{X}_0|^2],$$

which leads to

$$\mathbb{E} \left[\sup_{0 \leq t \leq \tau} |\mathbf{X}_t - \mathbf{X}_0|^2 \right] \leq C \int_0^\tau \mathbb{E}[|\mathbf{X}_s - \mathbf{X}_0|^2] ds + C \mathbb{E}[|\mathbf{X}_0|^2]. \quad (4.13)$$

Therefore Gronwall's inequality concludes that

$$\mathbb{E} \left[\sup_{0 \leq t \leq \tau} |\mathbf{X}_t - \mathbf{X}_0|^2 \right] \leq C \mathbb{E}[|\mathbf{X}_0|^2] \exp(C\tau) \quad \forall \tau > 0 \quad (4.14)$$

where $C > 0$ depend only on λ, σ and N . Namely, the solution exists globally in time for each fixed N . $\square$

In order to prove the well-posedness of the mean-field dynamics (2.4), we need the following stability estimate from [15, Lemma 3.2]:

Lemma 4.2 *Let $\mathcal{E}$ satisfy Assumption 2-(1) and $\rho, \hat{\rho} \in \mathcal{P}_4(\overline{\mathcal{D}})$ with*

$$\int |x|^4 d\rho(dx), \quad \int |x|^4 d\hat{\rho}(dx) \leq K.$$

Then it holds that

$$|X_\alpha(\rho) - X_\alpha(\hat{\rho})| \leq c_0 W_2(\rho, \hat{\rho}), \quad (4.15)$$

where c_0 depends only on $\alpha, L_{\mathcal{E}}$ and K and W_2 is the 2-Wasserstein distance defined in Table 1.

We will also need the following estimate on $X_\alpha(\rho)$ [36, Proposition A.3]:

Lemma 4.3 *Let $\mathcal{E}$ satisfy Assumption 2-(2) and $\rho \in \mathcal{P}_p(\overline{\mathcal{D}})$. Then there exists some constant $c_1 > 0$ depending only on $c_u, c_\ell, C_u, C_\ell, \alpha$ such that it holds*

$$|X_\alpha(\rho)| \leq \left(c_1 \int_{\overline{\mathcal{D}}} |x|^p \rho(dx) \right)^{\frac{1}{p}} \quad \text{for } p \geq 1. \quad (4.16)$$

Theorem 4.2 *Let $\mathcal{D}$ satisfy Assumption 1, $\mathcal{E}$ satisfy Assumption 2, and the initial data $\rho_0 \in \mathcal{P}_4(\overline{\mathcal{D}})$. Then there exists a unique process $\bar{X} \in \mathcal{C}([0, T]; \overline{\mathcal{D}})$, $T > 0$ satisfying the mean-field dynamic (2.4) in strong sense with $\rho \in \mathcal{C}([0, T]; \mathcal{P}_2(\overline{\mathcal{D}}))$.*

Proof We follow the proof of [20, Theorem 3.1, Theorem 3.2]. For any given $u \in \mathcal{C}([0, T]; \mathbb{R}^d)$, using [54, Theorem 4.1] we may uniquely solve the following linear SDE

$$\begin{cases} d\bar{X}_t = -\lambda(\bar{X}_t - u_t)dt + \sigma D(\bar{X}_t - u_t)dB_t - dL_t, \\ L_t = \int_0^t n(\bar{X}_s)d|L|_s, \quad |L|_t = \int_0^t \mathbf{I}_{\partial\mathcal{D}}(\bar{X}_s)d|L|_s, \end{cases} \quad (4.17a)$$

$$(4.17b)$$

Table 1 Notation Table. Further notations will be recalled in the text.

Notation	Definition
$\mathcal{E}$	The cost function one wishes to minimize.
x^*	The unique global minimizer of $\mathcal{E}$.
$\overline{\mathcal{D}}$	The closure of a convex domain $\mathcal{D}$.
$\partial\mathcal{D}$	The boundary of $\mathcal{D}$.
$ \mathcal{D} $	The diameter of $\mathcal{D}$.
$n(x)$	An outward normal vector at $x \in \partial\mathcal{D}$.
$\Pi_{\overline{\mathcal{D}}}(\cdot)$	The orthogonal projection onto $\overline{\mathcal{D}}$.
$ \cdot $	The standard Euclidean distance.
$\ \cdot\ _\infty$	The L^∞ norm.
$\mathcal{P}_p(\mathbb{R}^d)$	The set of probability measures on $\mathbb{R}^d$ with a finite p th-moment.
$W_2(\mu, \nu)$	The 2-Wasserstein distance between probability measures $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$. Namely, $W_2^2(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu, \nu)} \iint_{\mathbb{R}^d \times \mathbb{R}^d} x - y ^2 \gamma(dx, dy)$, where $\Gamma(\mu, \nu)$ is the set of all couplings between μ and ν .
$B_r(x)$	An open ball with origin x and radius r .

with the initial data $\overline{X}_0$ distributed according to $\rho_0 \in \mathcal{P}_4(\overline{\mathcal{D}})$. In particular $\sup_{t \in [0, T]} \mathbb{E}[|\overline{X}_t|^4] \leq K$ for some $K < \infty$ depends only on $\mathbb{E}[|\overline{X}_0|^4]$, T , λ and σ . Let us denote by $g_t = \text{Law}(\overline{X}_t) \in \mathcal{P}(\overline{\mathcal{D}})$ then $g \in \mathcal{C}([0, T]; \mathcal{P}_2(\overline{\mathcal{D}}))$. Indeed, for any $t < s$, $t, s \in (0, T)$, it follows from [54, Theorem 3.1] we have

$$\begin{aligned} |\overline{X}_t - \overline{X}_s|^2 &\leq 2 \int_s^t (X_\tau - X_s) \cdot [-\lambda(\overline{X}_\tau - u_\tau) d\tau + \sigma D(\overline{X}_\tau - u_\tau) dB_\tau] \\ &\quad + \int_s^t \sigma^2 |\overline{X}_\tau - u_\tau|^2 d\tau. \end{aligned} \quad (4.18)$$

Then taking expectation on both sides implies

$$\begin{aligned} \mathbb{E}[|\overline{X}_t - \overline{X}_s|^2] &\leq 2\lambda \int_s^t \mathbb{E}[|X_\tau - X_s| \cdot |\overline{X}_\tau - u_\tau|] d\tau + \sigma^2 \int_s^t \mathbb{E}[|\overline{X}_\tau - u_\tau|^2] d\tau \\ &\leq \lambda \int_s^t \mathbb{E}[|\overline{X}_\tau - u_\tau|^2] d\tau + (\sigma^2 + \lambda) \int_s^t \mathbb{E}[|\overline{X}_\tau - u_\tau^1|^2] d\tau \\ &\leq C|t - s|, \end{aligned} \quad (4.19)$$

for some $C > 0$ depending only on λ , σ , K and $\|u\|_\infty$. Thus we have $W_2(g_t, g_s) \leq C|t - s|^{\frac{1}{2}}$. Set

$$X_\alpha(g_t) = \frac{\int_{\overline{\mathcal{D}}} x \omega_\alpha^\mathcal{E}(x) g(t, dx)}{\int_{\overline{\mathcal{D}}} \omega_\alpha^\mathcal{E}(x) g(t, dx)}, \quad (4.20)$$

which provides the self-mapping property of the map

$$\mathcal{T} : \mathcal{C}([0, T]; \mathbb{R}^d) \rightarrow \mathcal{C}([0, T]; \mathbb{R}^d) \text{ with } u \mapsto \mathcal{T}(u) = X_\alpha(g) \quad (4.21)$$

for which we show to be compact later.

Applying Lemma 4.2 we obtain

$$|X_\alpha(g_t) - X_\alpha(g_s)| \leq c_0 W_2(g_t, g_s) \leq c_0 C |t - s|^{\frac{1}{2}}, \quad (4.22)$$

which indicates the Hölder continuity of $t \mapsto X_\alpha(g_t)$ with the exponent $1/2$. This implies the compactness of $\mathcal{T}$ because of the compact embedding $\mathcal{C}^{1/2}([0, T]; \mathbb{R}^d) \hookrightarrow \mathcal{C}([0, T]; \mathbb{R}^d)$. Now let $u \in \mathcal{C}([0, T]; \mathbb{R}^d)$ satisfy $u = \xi \mathcal{T}(u)$ for some $\xi \in [0, 1]$. In particular, there exists $g \in \mathcal{C}([0, T]; \mathcal{P}_2(\overline{\mathcal{D}}))$, i.e. the law of solution $\overline{X}$ to SDE (4.17) such that $u = \xi X_\alpha(g)$. Then according to Lemma 4.3, we have for all $t \in (0, T)$, it holds $|u_t|^2 = \xi^2 |X_\alpha(g_t)|^2 \leq \tau^2 c_1 \int_{\overline{\mathcal{D}}} |x|^2 g_t(dx)$. Following similar computations as in (4.11)-(4.13) it yields an estimate on $\sup_{t \in [0, T]} \int_{\overline{\mathcal{D}}} |x|^2 g_t(dx) = \sup_{t \in [0, T]} \mathbb{E}[|\overline{X}_t|^2] < \infty$. This implies $\|u\|_\infty < \infty$. Finally, applying the Leray–Schauder fixed point theorem provides a fixed point u for $\mathcal{T}$ thereby a solution to (2.4). As for the uniqueness, we can follow similar arguments in [20, Theorem 3.1] by using Lemma 4.2. The details are omitted here. $\square$

4.2 Mean-field limit

Now let $\{(\overline{X}_t^i)_{t \geq 0}\}_{i=1}^N$ be N independent copies of solutions to the mean-field dynamics (2.4), so they are i.i.d. with the same distribution ρ . First, one can prove the following moment estimates for empirical measures.

Lemma 4.4 *Let $\mathcal{E}$ satisfy Assumption 2, and suppose that $\rho_0 \in \mathcal{P}_{2p}(\overline{\mathcal{D}})$ for any $p \geq 2$.*

1. *Consider particle system (2.1) with $\rho_0^{\otimes N}$ -distributed initial data, and let ρ_t^N be the corresponding empirical measure. Then there exists some constant $c_2 > 0$ independent of N such that it holds*

$$\sup_{t \in [0, T]} \left\{ \sup_{i \in [N]} \mathbb{E}[|X_t^i|^p] + \mathbb{E}[\rho_t^N[|x|^p]] + \mathbb{E}[|X_\alpha(\rho_t^N)|^p] \right\} < c_2. \quad (4.23)$$

2. *Consider particles $\{(\overline{X}_t^i)_{t \geq 0}\}_{i=1}^N$ with $\rho_0^{\otimes N}$ -distributed initial data, and let $\overline{\rho}_t^N$ be the corresponding empirical measure. Then there exists some constant $c_2 > 0$ independent of N such that it holds*

$$\sup_{t \in [0, T]} \left\{ \sup_{i \in [N]} \mathbb{E}[|\overline{X}_t^i|^p] + \mathbb{E}[\overline{\rho}_t^N[|x|^p]] + \mathbb{E}[|X_\alpha(\overline{\rho}_t^N)|^p] \right\} < c_2. \quad (4.24)$$

Proof Arguments for (1) and (2) are parallel. Here we only prove (1) for the case $p = 4$. Indeed, for any fixed $i \in [N]$ applying Itô's formula leads to

$$\begin{aligned} d|X_t^i - X_0^i|^4 &= 4|X_t^i - X_0^i|^2 (X_t^i - X_0^i) \cdot [-\lambda(X_t^i - X_\alpha(\rho_t^N))dt \\ &\quad + \sigma D(X_t^i - X_\alpha(\rho_t^N))dB_t^i - dL_t^i] + 6\sigma^2 |X_t^i - X_0^i|^2 |X_t^i - X_\alpha(\rho_t^N)|^2 dt. \end{aligned} \quad (4.25)$$

Using the convexity of $\mathcal{D}$ and taking expectation on both sides one has

$$\begin{aligned} d\mathbb{E}[|X_t^i - X_0^i|^4] &\leq 3(\lambda + \sigma^2)\mathbb{E}[|X_t^i - X_0^i|^4]dt + (\lambda + 3\sigma^2)\mathbb{E}[|X_t^i - X_\alpha(\rho_t^N)|^4]dt \\ &\leq 3(\lambda + \sigma^2)\mathbb{E}[|X_t^i - X_0^i|^4]dt + 8(\lambda + 3\sigma^2)\mathbb{E}[|X_t^i|^4]dt + 8(\lambda + 3\sigma^2)\mathbb{E}[|X_\alpha(\rho_t^N)|^4]dt \\ &\leq 3(\lambda + \sigma^2)\mathbb{E}[|X_t^i - X_0^i|^4]dt + 8(\lambda + 3\sigma^2)\mathbb{E}[|X_t^i|^2]dt + 8(\lambda + 3\sigma^2)c_1\mathbb{E}[(\rho_t^N[|x|^4])]dt, \end{aligned}$$

where in the last inequality we used Lemma 4.3. Notice that here Itô's integral term disappears, namely

$$\mathbb{E}\left[4|X_t^i - X_0^i|^2(X_t^i - X_0^i) \cdot \sigma D(X_t^i - X_\alpha(\rho_t^N))dB_t^i\right] = 0 \quad (4.26)$$

because it satisfies,

$$\begin{aligned} \mathbb{E}\left[\int_0^t \left|4|X_s^i - X_0^i|^2(X_s^i - X_0^i) \cdot \sigma D(X_s^i - X_\alpha(\rho_s^N))\right|^2 ds\right] \\ \leq C \int_0^t \mathbb{E}[|X_s^i|^8 + |X_0^i|^8] ds < \infty, \end{aligned} \quad (4.27)$$

which can be guaranteed by the assumption $\rho_0 \in \mathcal{P}_8(\overline{\mathcal{D}})$.

Since particles are exchangeable, it implies that

$$\mathbb{E}[\rho_t^N[|x|^4]] = \mathbb{E}[|X_t^i|^4]. \quad (4.28)$$

Thus we have

$$\begin{aligned} d\mathbb{E}[|X_t^i - X_0^i|^4] &\leq 3(\lambda + \sigma^2)\mathbb{E}[|X_t^i - X_0^i|^4]dt + 8(\lambda + 3\sigma^2)(1 + c_1)\mathbb{E}[|X_t^i|^4]dt \\ &\leq 3(\lambda + \sigma^2)\mathbb{E}[|X_t^i - X_0^i|^4]dt + 64(\lambda + 3\sigma^2)(1 + c_1)\mathbb{E}[|X_t^i - X_0^i|^4]dt \\ &\quad + 64(\lambda + 3\sigma^2)(1 + c_1)\mathbb{E}[|X_0^i|^4]dt \end{aligned}$$

Gronwall's inequality gives that

$$\sup_{t \in [0, T], i \in [N]} \mathbb{E}[|X_t^i|^4] \leq C \quad (4.29)$$

for some $C > 0$ independent of N . Then estimate (4.23) follows from (4.28) and Lemma 4.3. $\square$

Then we recall a large deviation bound estimate from [36, Lemma 3.7]:

Lemma 4.5 *Let $\mathcal{E}$ satisfy Assumption 2, and $\bar{\rho}_t^N$ be the empirical measure associated to the particles $\{(\bar{X}_t^i)_{t \in [0, T]}\}_{i=1}^N$ satisfying (2.4) up to any time $T > 0$, which are i.i.d. with common distribution ρ satisfying $\sup_{t \in [0, T]} \int_{\overline{\mathcal{D}}} |x|^4 \rho_t(dx) < \infty$. Then there exists some constant $c_3 > 0$ depending only on $\mathcal{E}, \alpha, T$ and $\sup_{t \in [0, T]} \int_{\overline{\mathcal{D}}} |x|^4 \rho_t(dx)$ such that*

$$\sup_{t \in [0, T]} \mathbb{E}[|X_\alpha(\bar{\rho}_t^N) - X_\alpha(\rho_t)|^2] \leq c_3 N^{-1}. \quad (4.30)$$

Moreover, we will also need the following improved stability estimate [36, Corollary 3.3] on $X_\alpha(\rho)$:

Lemma 4.6 Suppose $\mathcal{E}$ satisfy Assumption 2. Then for all $R > 0$ there exists some constant $L > 0$ depending only on $L_{\mathcal{E}}, R, \alpha$ such that

$$|X_{\alpha}(\rho) - X_{\alpha}(\widehat{\rho})| \leq L W_2(\rho, \widehat{\rho}), \quad \forall (\rho, \widehat{\rho}) \in \mathcal{P}_{2,R}(\overline{\mathcal{D}}) \times \mathcal{P}_2(\overline{\mathcal{D}}), \quad (4.31)$$

where $\mathcal{P}_{2,R}(\overline{\mathcal{D}}) := \{\rho \in \mathcal{P}_2(\overline{\mathcal{D}}) : \int_{\overline{\mathcal{D}}} |x|^2 \rho(dx) \leq R\}$.

Compared to Lemma 4.2, here L depends only on the second moment bound of measure ρ . This lemma leads to the following mean-field limit estimate:

Theorem 4.3 Assume that $\mathcal{D}$ satisfies Assumption 1 and $\mathcal{E}$ satisfies Assumption 2. For any $T > 0$, let $\{(X_t^i)_{t \in [0, T]}\}_{i=1}^N$ and $\{(\overline{X}_t^i)_{t \in [0, T]}\}_{i=1}^N$ be the solutions to the interacting particle system (2.1) and the mean-field dynamics (2.4) respectively up to time T with the same initial data $\{(X_0^i)\}_{i=1}^N$ (i.i.d. distributed according to $\rho_0 \in \mathcal{P}_{16}(\overline{\mathcal{D}})$) and Brownian motions $\{(B_t^i)_{t \in [0, T]}\}_{i=1}^N$. Then there exists some $C_{\text{MFA}} > 0$ depending only on $\lambda, \sigma, \alpha, T, \mathcal{E}$ and $\sup_{t \in [0, T]} \mathbb{E}[|\overline{X}_t^1|^8]$ such that

$$\sup_{t \in [0, T]} \sup_{i=1, \dots, N} \mathbb{E}[|X_t^i - \overline{X}_t^i|^2] \leq C_{\text{MFA}} N^{-1}. \quad (4.32)$$

Proof Notice that

$$\begin{aligned} d(X_t^i - \overline{X}_t^i) &= -\lambda(X_t^i - \overline{X}_t^i)dt + \lambda(X_{\alpha}(\rho_t^N) - X_{\alpha}(\rho_t))dt \\ &\quad + \sigma D(X_t^i - \overline{X}_t^i)dB_t^i + \sigma D(X_{\alpha}(\rho_t) - X_{\alpha}(\rho_t^N))dB_t^i - dL_t^i + d\overline{L}_t^i, \end{aligned} \quad (4.33)$$

which by applying Itô's formula leads to

$$\begin{aligned} d(X_t^i - \overline{X}_t^i)^2 &= 2(X_t^i - \overline{X}_t^i) \cdot [-\lambda(X_t^i - \overline{X}_t^i)dt + \lambda(X_{\alpha}(\rho_t^N) - X_{\alpha}(\rho_t))dt] \\ &\quad + 2(X_t^i - \overline{X}_t^i) \cdot [\sigma D(X_t^i - \overline{X}_t^i)dB_t^i + \sigma D(X_{\alpha}(\rho_t) - X_{\alpha}(\rho_t^N))dB_t^i - dL_t^i + d\overline{L}_t^i] \\ &\quad + \sigma^2[(X_t^i - \overline{X}_t^i + (X_{\alpha}(\rho_t) - X_{\alpha}(\rho_t^N))^2]dt. \end{aligned} \quad (4.34)$$

Using the fact that

$$(X_t^i - \overline{X}_t^i) \cdot dL_t^i = (X_t^i - \overline{X}_t^i) \cdot n(X_t^i) \mathbf{I}_{\partial \mathcal{D}}(X_t^i) d|L_t^i|_t \geq 0, \quad (4.35)$$

$$(X_t^i - \overline{X}_t^i) \cdot d\overline{L}_t^i = (X_t^i - \overline{X}_t^i) \cdot n(\overline{X}_t^i) \mathbf{I}_{\partial \mathcal{D}}(\overline{X}_t^i) d|\overline{L}_t^i|_t \leq 0, \quad (4.36)$$

and taking expectation on both sides of (4.34) deduce

$$\begin{aligned} d\mathbb{E}[|X_t^i - \overline{X}_t^i|^2] &\leq -(\lambda - 2\sigma^2)\mathbb{E}[|X_t^i - \overline{X}_t^i|^2]dt + (\lambda + 2\sigma^2)\mathbb{E}[|X_{\alpha}(\rho_t^N) - X_{\alpha}(\rho_t)|^2]dt \\ &\leq -(\lambda - 2\sigma^2)\mathbb{E}[|X_t^i - \overline{X}_t^i|^2]dt + 2(\lambda + 2\sigma^2)\mathbb{E}[|X_{\alpha}(\rho_t^N) - X_{\alpha}(\overline{\rho}_t^N)|^2]dt + 2(\lambda + 2\sigma^2)\mathbb{E}[|X_{\alpha}(\overline{\rho}_t^N) \\ &\quad - X_{\alpha}(\rho_t)|^2]dt \\ &\leq -(\lambda - 2\sigma^2)\mathbb{E}[|X_t^i - \overline{X}_t^i|^2]dt + 2(\lambda + 2\sigma^2)\mathbb{E}[|X_{\alpha}(\rho_t^N) - X_{\alpha}(\overline{\rho}_t^N)|^2]dt + 2(\lambda + 2\sigma^2)c_3 \frac{1}{N}dt, \end{aligned} \quad (4.37)$$

where in the third inequality we have used Lemma 4.5. We now bound the term

$$\mathbb{E}[|X_{\alpha}(\rho_t^N) - X_{\alpha}(\overline{\rho}_t^N)|^4] \leq 8\mathbb{E}[|X_{\alpha}(\rho_t^N)|^4] + 8\mathbb{E}[|X_{\alpha}(\overline{\rho}_t^N)|^4] \leq 16c_2.$$

In order to deal with the non-Lipschitz property of the weighted mean $X_\alpha(\rho)$, we introduce for each $t \in [0, T]$ the event

$$A_{N,t} := \left\{ \omega \in \Omega : \frac{1}{N} \sum_{i=1}^N |\bar{X}_t^i(\omega)|^2 \geq R \right\}, \quad (4.38)$$

where $R > \sup_{t \in [0, T]} \mathbb{E}[|\bar{X}_t^1|^2]$ is fixed. Then according to [36, Lemma 2.5] we have the following bound

$$\mathbb{P}(A_{N,t}) \leq CN^{-2}, \quad \forall t \in [0, T] \quad (4.39)$$

holds if $\sup_{t \in [0, T]} \mathbb{E}[|\bar{X}_t^1|^8] < \infty$, which is guaranteed by Lemma 4.4 under the assumption that $\rho_0 \in \mathcal{P}_{16}(\mathcal{D})$. Then one splits the term

$$\begin{aligned} & \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2] \\ &= \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2 \mathbf{I}_{\Omega/A_{N,t}}] + \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2 \mathbf{I}_{A_{N,t}}] \\ &\leq \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2 \mathbf{I}_{\Omega/A_{N,t}}] + (\mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^4])^{1/2} (\mathbb{P}(A_{N,t}))^{1/2} \\ &\leq \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2 \mathbf{I}_{\Omega/A_{N,t}}] + 4c_2^{1/2} C \frac{1}{N}. \end{aligned} \quad (4.40)$$

Meanwhile, the improved stability Lemma 4.6 concludes that

$$\begin{aligned} \mathbb{E}[|X_\alpha(\rho_t^N) - X_\alpha(\bar{\rho}_t^N)|^2 \mathbf{I}_{\Omega/A_{N,t}}] &\leq C \mathbb{E}[W_2(\rho_t^N, \bar{\rho}_t^N)^2] \leq C \mathbb{E}\left[\frac{1}{N} \sum_{i \in [N]} |X_t^i - \bar{X}_t^i|^2\right] \\ &= C \mathbb{E}[|X_t^i - \bar{X}_t^i|^2], \end{aligned} \quad (4.41)$$

which leads to

$$d\mathbb{E}[|X_t^i - \bar{X}_t^i|^2] \leq C \mathbb{E}[|X_t^i - \bar{X}_t^i|^2] dt + 4c_2^{1/2} C \frac{1}{N} dt. \quad (4.42)$$

Thus applying Gronwall's inequality concludes the desired estimate. $\square$

5 Global convergence to the minimizer

In this section, we present our main result about the global convergence in mean-field law for cost functions satisfying the following conditions.

Assumption 3 Throughout this section we are interested in the objective function $\mathcal{E} \in \mathcal{C}(\bar{\mathcal{D}})$, for which

1. there exists a unique $x^* \in \bar{\mathcal{D}}$ such that $\mathcal{E}(x^*) = \min_{x \in \bar{\mathcal{D}}} \mathcal{E}(x) =: \underline{\mathcal{E}}$.
2. there exists some $\mathcal{E}_\infty, R_0, \eta, \nu > 0$ such that

$$\eta |x - x^*| \leq |\mathcal{E}(x) - \underline{\mathcal{E}}|^\nu \text{ for all } x \in B_{R_0}(x^*) \cap \bar{\mathcal{D}}, \quad (5.1)$$

$$\mathcal{E}_\infty < \mathcal{E}(x) - \underline{\mathcal{E}} \text{ for all } x \in (B_{R_0}(x^*))^c \cap \bar{\mathcal{D}}. \quad (5.2)$$

3. For any $q > 0$, there exists some $r \in (0, R_0]$ such that

$$|\mathcal{E}(x) - \underline{\mathcal{E}}| \leq q, \quad \text{for all } x \in B_r(x^*) \cap \overline{\mathcal{D}}. \quad (5.3)$$

In this section we define

$$\mathcal{V}(t) := \mathbb{E}[|\overline{X}_t - x^*|^2], \quad (5.4)$$

then one has the following lemma

Lemma 5.1 *The functional $\mathcal{V}(t)$ defined in (5.4) satisfies*

$$\frac{d\mathcal{V}(t)}{dt} \leq -(2\lambda - \sigma^2)\mathcal{V}(t) + 2(\lambda + \sigma^2)\mathcal{V}(t)^{\frac{1}{2}}|x^* - X_\alpha(\rho_t)| + \sigma^2|x^* - X_\alpha(\rho_t)|^2 \quad (5.5)$$

Proof It is easy to compute that

$$d|\overline{X}_t - x^*|^2 = 2(\overline{X}_t - x^*)dX_t + \sigma^2|\overline{X}_t - X_\alpha(\rho_t)|^2 dt. \quad (5.6)$$

which implies

$$\begin{aligned} d\mathcal{V}(t) &= -2\lambda\mathbb{E}[(\overline{X}_t - x^*) \cdot (\overline{X}_t - X_\alpha(\rho_t))]dt - 2\mathbb{E}[(\overline{X}_t - x^*) \cdot n(\overline{X}_t)\mathbf{I}_{\partial\mathcal{D}}(\overline{X}_t)d|\overline{L}|_t] \\ &\quad + \sigma^2\mathbb{E}[|\overline{X}_t - X_\alpha(\rho_t)|^2]dt \\ &\leq -2\lambda\mathbb{E}[(\overline{X}_t - x^*) \cdot (\overline{X}_t - X_\alpha(\rho_t))]dt + \sigma^2\mathbb{E}[|\overline{X}_t - X_\alpha(\rho_t)|^2]dt \\ &= -2\lambda\mathcal{V}(t) - 2\lambda\mathbb{E}[(\overline{X}_t - x^*) \cdot (x^* - X_\alpha(\rho_t))] + \sigma^2\mathbb{E}[|\overline{X}_t - X_\alpha(\rho_t)|^2] \\ &\leq -(2\lambda - \sigma^2)\mathcal{V}(t) + 2(\lambda + \sigma^2)\mathcal{V}(t)^{\frac{1}{2}}|x^* - X_\alpha(\rho_t)| + \sigma^2|x^* - X_\alpha(\rho_t)|^2, \end{aligned} \quad (5.7)$$

where in the second inequality we have used the fact that

$$-2(\overline{X}_t - x^*) \cdot n(\overline{X}_t) \leq 0 \quad (5.8)$$

for all $\overline{X}_t \in \partial\mathcal{D}$ and $x^* \in \overline{\mathcal{D}}$, since $\mathcal{D}$ is convex and $n(\overline{X}_t)$ is the outward normal vector at $\overline{X}_t$. This completes the proof. $\square$

In the following, we shall use the notation for a ℓ^∞ -ball with a radius r and the center x^* , i.e. $B_r(x^*) := \{x : |x - x^*|_\infty \leq r\}$. Let us quantify the maximum discrepancy for the objective function $\mathcal{E}$ around the minimizer x^* , namely for $s > 0$,

$$\mathcal{E}_s := \sup_{x \in B_s(x^*) \cap \overline{\mathcal{D}}} |\mathcal{E}(x) - \underline{\mathcal{E}}|. \quad (5.9)$$

Proposition 5.1 *Assume that $\mathcal{E}$ satisfies Assumption 3. For any $t > 0$ and $s \in (0, R_0]$ let $\mathcal{E}_s$ be defined as in (5.9), $0 < q \leq \frac{\mathcal{E}_\infty}{2}$, and $r := \max\{s \in (0, R_0] : \mathcal{E}_s \leq q\}$ then we have*

$$|x^* - X_\alpha(\rho_t)| \leq \frac{(2q)^\nu}{\eta} + \frac{\exp(-\alpha q)}{\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int_{\overline{\mathcal{D}}} |x - x^*| \rho_t(dx). \quad (5.10)$$

Proof Let $\tilde{r} \geq r > 0$, and using Jensen's inequality one can deduce

$$|x^* - X_\alpha(\rho_t)| \leq \int_{B_{\tilde{r}}(x^*) \cap \overline{\mathcal{D}}} |x - x^*| \frac{\omega_\alpha(x) \rho_t(dx)}{\|\omega_\alpha(\cdot)\|_{L_1(\rho_t)}} + \int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*| \frac{\omega_\alpha(x) \rho_t(dx)}{\|\omega_\alpha(\cdot)\|_{L_1(\rho_t)}}. \quad (5.11)$$

The first term is bounded by $\tilde{r}$ since $|x - x^*| \leq \tilde{r}$ for all $x \in B_{\tilde{r}}(x^*)$. Moreover, it follows from Markov's inequality that

$$\begin{aligned} \|\omega_\alpha(\cdot)\|_{L_1(\rho_t)} &\geq \exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\rho_t(\{x \in \overline{\mathcal{D}} : \exp(-\alpha\mathcal{E}(x)) \geq \exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\}) \\ &= \exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\rho_t(\{x \in \overline{\mathcal{D}} : \mathcal{E}(x) \leq (\mathcal{E}_r + \underline{\mathcal{E}})\}) \\ &\geq \exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\}). \end{aligned} \quad (5.12)$$

Then for the second term we have

$$\begin{aligned} &\int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*| \frac{\omega_\alpha(x)\rho_t(dx)}{\|\omega_\alpha(\cdot)\|_{L_1(\rho_t)}} \\ &\leq \frac{1}{\exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*|\omega_\alpha(x)\rho_t(dx) \\ &\leq \frac{\exp(-\alpha \min_{x \in (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} \mathcal{E}(x))}{\exp(-\alpha(\mathcal{E}_r + \underline{\mathcal{E}}))\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*|\rho_t(dx) \\ &= \frac{\exp(-\alpha(\min_{x \in (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} \mathcal{E}(x) - \mathcal{E}_r - \underline{\mathcal{E}}))}{\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*|\rho_t(dx). \end{aligned} \quad (5.13)$$

Thus for any $\tilde{r} \geq r > 0$ we obtain

$$\begin{aligned} &|x^* - X_\alpha(\rho_t)| \\ &\leq \tilde{r} + \frac{\exp(-\alpha(\min_{x \in (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} \mathcal{E}(x) - \mathcal{E}_r - \underline{\mathcal{E}}))}{\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int_{(B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} |x - x^*|\rho_t(dx). \end{aligned} \quad (5.14)$$

Next we choose $\tilde{r} = \frac{(q + \mathcal{E}_r)^\nu}{\eta}$, then it holds that

$$\tilde{r} = \frac{(q + \mathcal{E}_r)^\nu}{\eta} \geq \frac{\mathcal{E}_r^\nu}{\eta} = \frac{\left(\sup_{x \in B_r(x^*) \cap \overline{\mathcal{D}}} |\mathcal{E}(x) - \underline{\mathcal{E}}|\right)^\nu}{\eta} \geq \sup_{x \in B_r(x^*) \cap \overline{\mathcal{D}}} |x - x^*| = r \quad (5.15)$$

by using Assumption 3. Additionally one notice that

$$\min_{x \in (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} \mathcal{E}(x) - \underline{\mathcal{E}} \geq \begin{cases} \mathcal{E}_\infty, & \text{if } x \in (B_{R_0}(x^*))^c \cap (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}, \\ (\tilde{r}\eta)^{\frac{1}{\nu}}, & \text{if } x \in B_{R_0}(x^*) \cap (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}. \end{cases} \quad (5.16)$$

Thus, since $(\tilde{r}\eta)^{\frac{1}{\nu}} = q + \mathcal{E}_r \leq 2q \leq \mathcal{E}_\infty$, we have

$$\min_{x \in (B_{\tilde{r}}(x^*))^c \cap \overline{\mathcal{D}}} \mathcal{E}(x) - \mathcal{E}_r - \underline{\mathcal{E}} \geq (\tilde{r}\eta)^{\frac{1}{\nu}} - \mathcal{E}_r = q + \mathcal{E}_r - \mathcal{E}_r = q. \quad (5.17)$$

Inserting this and the definition of $\tilde{r}$ into (5.14), we conclude the result. $\square$

To eventually apply the above proposition, one needs to ensure that $\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\}) = \mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}})$ is bounded away from 0 for a finite time horizon T . To do so, we employ a rather technical argument inspired from [28, Proposition 23], and introduce

the mollifier $\phi_r^{x_0} : \mathbb{R}^d \rightarrow \mathbb{R}$, with $r > 0$ and $x_0 \in \mathcal{D}$ defined by

$$\phi_r^{x_0}(x) := \begin{cases} \prod_{k=1}^d \exp\left(1 - \frac{r^2}{r^2 - (x - x_0)_k^2}\right), & \text{if } x \in B_r(x_0) \cap \mathcal{D} \\ 0, & \text{else.} \end{cases} \quad (5.18)$$

So we have $\text{Im}(\phi_r) = [0, 1]$, $\text{supp}(\phi_r) = B_r(x_0) \cap \mathcal{D}$, $\phi_r \in \mathcal{C}_c^\infty(\mathbb{R}^d)$ and

$$\begin{aligned} \partial_{x_k} \phi_r^{x_0}(x) &= -2r^2 \frac{(x - x_0)_k}{(r^2 - (x - x_0)_k^2)^2} \phi_r(x), \\ \partial_{x_k}^2 \phi_r^{x_0}(x) &= 2r^2 \left(\frac{2(2(x - x_0)_k^2 - r^2)(x - x_0)_k^2 - (r^2 - (x - x_0)_k^2)^2}{(r^2 - (x - x_0)_k^2)^4} \right) \phi_r(x). \end{aligned} \quad (5.19)$$

The proof will be split into two cases:

5.1 Global convergence proof when $\overline{\mathcal{D}}$ is bounded

In this subsection we assume that $\overline{\mathcal{D}}$ is bounded. Namely, there exists some diameter $0 < |\mathcal{D}| < \infty$ such that $|\mathcal{D}| = \max_{x, y \in \overline{\mathcal{D}}} |x - y|$. First, we get the following result, whose proof is postponed to the Appendix.

Proposition 5.2 *Let $(\overline{X}_t)_{0 \leq t \leq T}$ be the solution to (2.4) up to any time $T > 0$. Then for any $x_0 \in \mathcal{D}$ and $r > 0$, there exists some constant $\vartheta > 0$ depends only on d, r, λ, σ and $|\mathcal{D}|$ such that*

$$\mathbb{P}(\overline{X}_t \in B_r(x_0) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_r^{x_0}(\overline{X}_0)] \exp(-\vartheta t) \quad (5.20)$$

holds for all $t \in [0, T]$, where $\phi_r^{x_0}$ is defined as in (5.18).

This result directly implies

Corollary 5.3 *Let $(\overline{X}_t)_{0 \leq t \leq T}$ be the solution to (2.4) up to any time $T > 0$. Then for any $r > 0$, let $\tilde{x}^* \in \mathcal{D}$ and $0 < \tilde{r} \leq r$ be choosen such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Then there exists some constant $\vartheta > 0$ depends only on d, λ, σ and $|\mathcal{D}|$ such that*

$$\mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)] \exp(-\vartheta t) \quad (5.21)$$

holds for all $t \in [0, T]$. Especially when $x^* \in \mathcal{D}$ one can choose $\tilde{x}^* = x^*$ and $\tilde{r} = r$.

Proof *Case 1:* if $x^* \in \mathcal{D}$, then one may let $x_0 = x^*$ in Proposition 5.2 and immediately gets

$$\mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_r^{x^*}(\overline{X}_0)] \exp(-\vartheta t). \quad (5.22)$$

Case 2: if $x^* \in \partial\mathcal{D}$, then we can always find some $x_0 = \tilde{x}^* \in \mathcal{D}$ and $0 < \tilde{r} < r$ such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Using Proposition 5.2 again one yields

$$\mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{P}(\overline{X}_t \in B_{\tilde{r}}(\tilde{x}^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)] \exp(-\vartheta t). \quad (5.23)$$

□

Now we are ready to prove the global convergence result, which provides a rate for the variance function $\mathcal{V}(t) = \mathbb{E}[|\bar{X}_t - x^*|^2]$ within a prescribed time-range.

Theorem 5.4 Assume that $\mathcal{E}$ satisfies Assumption 3, and λ, σ satisfy $2\lambda > \sigma^2$. Furthermore assume the initial data satisfies $\mathcal{V}(0) > 2\varepsilon$ for any prescribed accuracy $\varepsilon > 0$, and $\mathbb{E}[\phi_r^{\tilde{x}^*}(\bar{X}_0)] > 0$ with $\tilde{x}^*, \tilde{r}$ to be determined later. Let $\tau \in (0, 1)$ and choose α to be sufficiently large satisfying

$$\alpha \geq \frac{\vartheta T_\varepsilon - \log(\frac{1}{2}c_5\mathbb{E}[\phi_r^{\tilde{x}^*}(\bar{X}_0)])}{c_4}, \quad (5.24)$$

where ϑ comes from Corollary 5.3,

$$T_\varepsilon := \frac{1}{(1-\tau)(2\lambda-\sigma^2)} \log\left(\frac{\mathcal{V}(0)}{\varepsilon}\right), \quad c_5 := \min\left\{\frac{\tau(2\lambda-\sigma^2)}{4(\lambda+\sigma^2)}, \sqrt{\frac{\tau(2\lambda-\sigma^2)}{2\sigma^2}}\right\},$$

and

$$c_4 := \frac{1}{2} \min\left\{\left(\eta \frac{c_5\sqrt{\varepsilon}}{2}\right)^{\frac{1}{\nu}}, \mathcal{E}_\infty\right\}, \quad r := \left\{\max_{s \in (0, R_0]} : \mathcal{E}_s \leq c_4\right\},$$

and $\tilde{x}^* \in D$ and $0 < \tilde{r} < r$ are chosen such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Then there exists some $0 < T_* \leq T_\varepsilon$ such that $\mathcal{V}(t)$ satisfies

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1-\tau)(2\lambda-\sigma^2)t) \quad \text{for all } t \in [0, T_*]. \quad (5.25)$$

and it reaches the prescribed accuracy at time T_* , namely $\mathcal{V}(T_*) = \varepsilon$.

Proof We define

$$T_\alpha := \inf\{t \geq 0 : \mathcal{V}(t) = \varepsilon\}, \quad (5.26)$$

and for all $t \in [0, T_\alpha]$ let

$$C_\alpha(t) := c_5\sqrt{\mathcal{V}(t)} = \min\left\{\frac{\tau(2\lambda-\sigma^2)}{4(\lambda+\sigma^2)}, \sqrt{\frac{\tau(2\lambda-\sigma^2)}{2\sigma^2}}\right\}\sqrt{\mathcal{V}(t)}. \quad (5.27)$$

Here the time T_α represents the first time when the variance $\mathcal{V}(t)$ reaches the prescribed accuracy ε , and it depends on α because $\mathcal{V}$ does. It is obvious that $T_\alpha > 0$ is because of the assumption on the initial data that $\mathcal{V}(0) > 2\varepsilon$. According to the definition of T_α one has $\mathcal{V}(t) > \varepsilon$ for all $t \in [0, T_\alpha)$ and $\mathcal{V}(T_\alpha) = \varepsilon$. Next we will prove that $T_\alpha \leq T_\varepsilon$ and $\mathcal{V}(t)$ decreases exponentially on $[0, T_\alpha)$.

Case $T_\alpha \leq T_\varepsilon$: It follows from Proposition 5.1 that for all $t \in [0, T_\alpha)$

$$\begin{aligned} |x^* - X_\alpha(\rho_t)| &\leq \frac{(2q_t^\alpha)^\nu}{\eta} + \frac{\exp(-\alpha q_t^\alpha)}{\rho_t(\{x \in \bar{D} : x \in B_r(x^*)\})} \int |x - x^*| \rho_t(dx) \\ &\leq \frac{(2q_t^\alpha)^\nu}{\eta} + \frac{\exp(-\alpha q_t^\alpha)}{\rho_t(\{x \in \bar{D} : x \in B_r(x^*)\})} \sqrt{\mathcal{V}(t)}, \end{aligned} \quad (5.28)$$

where we have constructed

$$q_t^\alpha := \frac{1}{2} \min \left\{ \left(\eta \frac{C_\alpha(t)}{2} \right)^{\frac{1}{v}}, \mathcal{E}_\infty \right\}. \quad (5.29)$$

Then it holds

$$c_4 = \frac{1}{2} \min \left\{ \left(\eta \frac{c_5 \sqrt{\varepsilon}}{2} \right)^{\frac{1}{v}}, \mathcal{E}_\infty \right\} < q_t^\alpha \leq \frac{1}{2} \mathcal{E}_\infty \quad (5.30)$$

because $C_\alpha(t) > c_5 \sqrt{\varepsilon}$. Moreover, here we let

$$r = \left\{ \max_{s \in [0, R_0]} : \mathcal{E}_s \leq c_4 \right\}. \quad (5.31)$$

By construction, these choices satisfy $r \leq R_0$ and $\frac{(2q_t^\alpha)^v}{\eta} \leq \frac{C_\alpha(t)}{2}$.

Moreover, $\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\}) = \mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_r^{\tilde{x}^*}(\overline{X}_0)] \exp(-\vartheta t)$ holds according to (5.21), where ϑ depends only on d, r, λ, σ and $|\mathcal{D}|$. This concludes that for all $t \in [0, T_\alpha)$

$$|x^* - X_\alpha(\rho_t)| \leq \frac{C_\alpha(t)}{2} + \frac{\exp(-\alpha c_4) \exp(\vartheta T_\varepsilon)}{\mathbb{E}[\phi_r^{\tilde{x}^*}(\overline{X}_0)]} \sqrt{\mathcal{V}(t)} \leq C_\alpha(t), \quad (5.32)$$

where we choose $\alpha \geq \alpha_0$ with

$$\alpha_0 = \frac{\vartheta T_\varepsilon - \log(\frac{1}{2} c_5 \mathbb{E}[\phi_r^{\tilde{x}^*}(\overline{X}_0)])}{c_4}. \quad (5.33)$$

Let us recall the upper bound for the time derivative of $\mathcal{V}(t)$ given in Lemma 5.1:

$$\frac{d\mathcal{V}(t)}{dt} \leq -(2\lambda - \sigma^2)\mathcal{V}(t) + 2(\lambda + \sigma^2)\mathcal{V}(t)^{\frac{1}{2}} |x^* - X_\alpha(\rho_t)| + \sigma^2 |x^* - X_\alpha(\rho_t)|^2. \quad (5.34)$$

Then using the definition of c_5 and the estimate from (5.32) one can deduce that

$$\frac{d\mathcal{V}(t)}{dt} \leq -(1 - \tau)(2\lambda - \sigma^2)\mathcal{V}(t), \quad (5.35)$$

which by Gronwall's inequality leads to

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1 - \tau)(2\lambda - \sigma^2)t) \text{ for } t \in [0, T_\alpha). \quad (5.36)$$

Case $T_\varepsilon < T_\alpha$: By the definition of T_α we know $\mathcal{V}(t) > \varepsilon$ for all $t \in [0, T_\varepsilon]$. Then following the same argument as in the first case, one can conclude

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1 - \tau)(2\lambda - \sigma^2)t) \text{ for } t \in [0, T_\varepsilon]. \quad (5.37)$$

Using this estimate the fact that $T_\varepsilon = \frac{1}{(1-\tau)(2\lambda-\sigma^2)} \log(\frac{\mathcal{V}(0)}{\varepsilon})$ implies $\mathcal{V}(T_\varepsilon) \leq \varepsilon$, which is a contradiction. Thus this case can never happen. $\square$

5.2 Global convergence proof when $\overline{\mathcal{D}}$ is unbounded

In this subsection we consider the case where $\mathcal{D}$ is unbounded, namely a finite diameter $|\mathcal{D}|$ does not exist. Therefore, Proposition 5.2 and Corollary 5.3 need some adjustments. Firstly, following the proof of Proposition 5.2 it is easy to get

Proposition 5.5 *Let $(\overline{X}_t)_{0 \leq t \leq T}$ be the solution to (2.4) up to any time $T > 0$. For any fixed $x_0 \in \mathcal{D}$, assume that $\sup_{t \in [0, T]} |x_0 - X_\alpha(\rho_t)| \leq B$ for some $B > 0$. Then for any $r > 0$, there exists some constant $\vartheta > 0$ depends only on d, r, λ, σ and B such that*

$$\mathbb{P}(\overline{X}_t \in B_r(x_0) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_r^{x_0}(\overline{X}_0)] \exp(-\vartheta t) \quad (5.38)$$

holds for all $t \in [0, T]$, where $\phi_r^{x_0}$ is defined as in (5.18).

Proof Compared to Proposition 5.2, the additional assumption on the upper bound B compensates for the lack of an upper bound on $|\mathcal{D}|$. The proof is basically identical to the one of Proposition 5.2 except we shall bound $|(\overline{X}_t - X_\alpha(\rho_t))_k|$ in (6.25) differently as follows

$$|(\overline{X}_t - X_\alpha(\rho_t))_k| \leq |(\overline{X}_t - x_0)_k| + |(x_0 - X_\alpha(\rho_t))_k| \leq \sqrt{c}r + B.$$

□

This implies the following result immediately similar to the bounded domain case.

Corollary 5.6 *Let $(\overline{X}_t)_{0 \leq t \leq T}$ be the solution to (2.4) up to any time $T > 0$. For any $r > 0$, let $\tilde{x}^* \in \mathcal{D}$ and $0 < \tilde{r} < r$ be choosen such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Assume that $\sup_{t \in [0, T]} |\tilde{x}^* - X_\alpha(\rho_t)| \leq B$ for some $B > 0$. Then there exists some constant $\vartheta > 0$ depends only on d, r, λ, σ and B such that*

$$\mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)] \exp(-\vartheta t) \quad (5.39)$$

holds for all $t \in [0, T]$. Especially when $x^* \in \mathcal{D}$ one can choose $\tilde{x}^* = x^*$ and $\tilde{r} = r$.

Now we are ready to prove the global convergence result as in Theorem 5.4

Theorem 5.7 *Assume that $\mathcal{E}$ satisfies Assumption 3, and λ, σ satisfy $2\lambda > \sigma^2$. Furthermore, assume the initial data satisfies $\mathcal{V}(0) > 2\varepsilon$ for any prescribed accuracy $\varepsilon > 0$, and $\rho_0(\{x \in \overline{\mathcal{D}} : x \in B_{r_0}(x^*)\}) > 0$, $\mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)] > 0$ with $\tilde{x}^*, \tilde{r}$ to be determined later. Let $\tau \in (0, 1)$ and choose α to be sufficiently large satisfying*

$$\alpha \geq \max \left\{ \frac{-\log(\frac{1}{2}c_5\rho_0(\{x \in \overline{\mathcal{D}} : x \in B_{r_0}(x^*)\}))}{q_0}, \frac{\vartheta T_\varepsilon - \log(\frac{1}{2}c_5\mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)])}{c_4} \right\}, \quad (5.40)$$

where ϑ comes from Proposition 5.6,

$$T_\varepsilon := \frac{1}{(1-\tau)(2\lambda - \sigma^2)} \log \left(\frac{\mathcal{V}(0)}{\varepsilon} \right), \quad c_5 := \min \left\{ \frac{\tau(2\lambda - \sigma^2)}{4(\lambda + \sigma^2)}, \sqrt{\frac{\tau(2\lambda - \sigma^2)}{2\sigma^2}} \right\},$$

$$q_0 := \frac{1}{2} \min \left\{ \left(\eta \frac{c_5 \sqrt{\mathcal{V}(0)}}{2} \right)^{\frac{1}{\nu}}, \mathcal{E}_\infty \right\}, \quad c_4 := \frac{1}{2} \min \left\{ \left(\eta \frac{c_5 \sqrt{\varepsilon}}{2} \right)^{\frac{1}{\nu}}, \mathcal{E}_\infty \right\},$$

$$r_0 := \left\{ \max_{s \in (0, R_0]} : \mathcal{E}_s \leq q_0 \right\}, \quad r := \left\{ \max_{s \in (0, R_0]} : \mathcal{E}_s \leq c_4 \right\},$$

and $\tilde{x}^* \in \mathcal{D}$ and $0 < \tilde{r} < r$ are chosen such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Then there exists some $0 < T_* \leq T_\varepsilon$ such that $\mathcal{V}(t)$ satisfies

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1 - \tau)(2\lambda - \sigma^2)t) \quad \text{for all } t \in [0, T_*]. \quad (5.41)$$

and it reaches the prescribed accuracy at time T_* , namely $\mathcal{V}(T_*) = \varepsilon$.

Proof We define

$$T_\alpha := \inf \left\{ t \geq 0 : \mathcal{V}(t) = \varepsilon \quad \text{or} \quad |x^* - X_\alpha(\rho_t)| = 2c_5 \sqrt{\mathcal{V}(0)} \right\}, \quad (5.42)$$

and for all $t \in [0, T_\alpha]$ define

$$C_\alpha(t) := c_5 \sqrt{\mathcal{V}(t)} = \min \left\{ \frac{\tau(2\lambda - \sigma^2)}{4(\lambda + \sigma^2)}, \sqrt{\frac{\tau(2\lambda - \sigma^2)}{2\sigma^2}} \right\} \sqrt{\mathcal{V}(t)}. \quad (5.43)$$

Here the time T_α represents the first time when the variance $\mathcal{V}(t)$ reaches the prescribed accuracy ε or $|x^* - X_\alpha(\rho_t)| = 2c_5 \sqrt{\mathcal{V}(0)}$, and it depends on α because $\mathcal{V}$ does.

Now we show that $T_\alpha > 0$. It follows from Proposition 5.1 that

$$\begin{aligned} |x^* - X_\alpha(\rho_0)| &\leq \frac{(2q_0)^\nu}{\eta} + \frac{\exp(-\alpha q_0)}{\rho_0(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \int |x - x^*| \rho_0(dx) \\ &\leq \frac{(2q_0)^\nu}{\eta} + \frac{\exp(-\alpha q_0)}{\rho_0(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \sqrt{\mathcal{V}(0)}, \end{aligned} \quad (5.44)$$

where we have constructed

$$q_0 := \frac{1}{2} \min \left\{ \left(\eta \frac{c_5 \sqrt{\mathcal{V}(0)}}{2} \right)^{\frac{1}{\nu}}, \mathcal{E}_\infty \right\}. \quad (5.45)$$

Then it holds $q_0 \leq \frac{1}{2} \mathcal{E}_\infty$. Moreover, here we let

$$r_0 = \left\{ \max_{s \in [0, R_0]} : \mathcal{E}_s \leq q_0 \right\}. \quad (5.46)$$

By construction, these choices satisfy $r_0 \leq R_0$ and $\frac{(2q_0)^\nu}{\eta} \leq \frac{c_5 \sqrt{\mathcal{V}(0)}}{2}$. This means that

$$|x^* - X_\alpha(\rho_0)| \leq \frac{c_5 \sqrt{\mathcal{V}(0)}}{2} + \frac{\exp(-\alpha q_0)}{\rho_0(\{x \in \overline{\mathcal{D}} : x \in B_{r_0}(x^*)\})} \sqrt{\mathcal{V}(0)}. \quad (5.47)$$

If now we choose α sufficiently large, e.g., $\alpha \geq \alpha_0$ with

$$\alpha_0 = \frac{-\log\left(\frac{1}{2} c_5 \rho_0(\{x \in \overline{\mathcal{D}} : x \in B_{r_0}(x^*)\})\right)}{q_0}$$

then one has

$$|x^* - X_\alpha(\rho_0)| \leq c_5 \sqrt{\mathcal{V}(0)}.$$

This together with $\mathcal{V}(0) \geq 2\varepsilon$ implies $T_\alpha > 0$.

According to the definition of T_α one has

$$\mathcal{V}(t) > \varepsilon \quad \text{and} \quad |x^* - X_\alpha(\rho_t)| < 2c_5 \sqrt{\mathcal{V}(0)}, \quad \text{for all } t \in [0, T_\alpha),$$

and at $t = T_\alpha$, it holds $\mathcal{V}(T_\alpha) = \varepsilon$ or $|x^* - X_\alpha(\rho_{T_\alpha})| = 2c_5 \sqrt{\mathcal{V}(0)}$. Next we will prove that $T_\alpha \leq T_\varepsilon$ and $\mathcal{V}(t)$ decreases exponentially on $[0, T_\alpha)$.

Case $T_\alpha \leq T_\varepsilon$: It follows from Proposition 5.1 that for all $t \in [0, T_\alpha)$

$$|x^* - X_\alpha(\rho_t)| \leq \frac{(2q_t^\alpha)^\nu}{\eta} + \frac{\exp(-\alpha q_t^\alpha)}{\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\})} \sqrt{\mathcal{V}(t)}, \quad (5.48)$$

where we have constructed

$$q_t^\alpha := \frac{1}{2} \min \left\{ \left(\eta \frac{C_\alpha(t)}{2} \right)^{\frac{1}{\nu}}, \mathcal{E}_\infty \right\}. \quad (5.49)$$

Then it holds

$$c_4 = \frac{1}{2} \min \left\{ \left(\eta \frac{c_5 \sqrt{\varepsilon}}{2} \right)^{\frac{1}{\nu}}, \mathcal{E}_\infty \right\} < q_t^\alpha \leq \frac{1}{2} \mathcal{E}_\infty \quad (5.50)$$

because $C_\alpha(t) > c_5 \sqrt{\varepsilon}$. Moreover here we let

$$r = \left\{ \max_{s \in [0, R_0]} : \mathcal{E}_s \leq c_4 \right\}. \quad (5.51)$$

By construction, these choices satisfy $r \leq R_0$ and $\frac{(2q_t^\alpha)^\nu}{\eta} \leq \frac{C_\alpha(t)}{2}$.

Next we choose $\tilde{x}^* \in D$ and $0 < \tilde{r} < r$ such that $B_{\tilde{r}}(\tilde{x}^*) \subset B_r(x^*)$. Then we have

$$\sup_{t \in [0, T_\alpha)} |\tilde{x}^* - X_\alpha(\rho_t)| \leq |\tilde{x}^* - x^*| + \sup_{t \in [0, T_\alpha)} |x^* - X_\alpha(\rho_t)| < r + 2c_5 \sqrt{\mathcal{V}(0)} =: B.$$

Moreover according to (5.39) it holds that

$$\rho_t(\{x \in \overline{\mathcal{D}} : x \in B_r(x^*)\}) = \mathbb{P}(\overline{X}_t \in B_r(x^*) \cap \overline{\mathcal{D}}) \geq \mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)] \exp(-\vartheta t),$$

where ϑ depends only on d, r, λ, σ and B . This concludes that for all $t \in [0, T_\alpha)$

$$|x^* - X_\alpha(\rho_t)| \leq \frac{C_\alpha(t)}{2} + \frac{\exp(-\alpha c_4) \exp(\vartheta T_\varepsilon)}{\mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)]} \sqrt{\mathcal{V}(t)} \leq C_\alpha(t), \quad (5.52)$$

where we choose $\alpha \geq \alpha_0$ with

$$\alpha_0 = \frac{\vartheta T_\varepsilon - \log(\frac{1}{2} c_5 \mathbb{E}[\phi_{\tilde{r}}^{\tilde{x}^*}(\overline{X}_0)])}{c_4}. \quad (5.53)$$

Let us recall the upper bound for the time derivative of $\mathcal{V}(t)$ given in Lemma 5.1:

$$\frac{d\mathcal{V}(t)}{dt} \leq -(2\lambda - \sigma^2)\mathcal{V}(t) + 2(\lambda + \sigma^2)\mathcal{V}(t)^{\frac{1}{2}} |x^* - X_\alpha(\rho_t)| + \sigma^2 |x^* - X_\alpha(\rho_t)|^2. \quad (5.54)$$

Then using the definition of c_5 and the estimate from (5.52) one can deduce that

$$\frac{d\mathcal{V}(t)}{dt} \leq -(1-\tau)(2\lambda - \sigma^2)\mathcal{V}(t), \quad (5.55)$$

which by Gronwall's inequality leads to

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1-\tau)(2\lambda - \sigma^2)t), \quad \text{for } t \in [0, T_\alpha]. \quad (5.56)$$

This implies

$$|x^* - X_\alpha(\rho_{T_\alpha})| \leq C_\alpha(T_\alpha) = c_5 \sqrt{\mathcal{V}(T_\alpha)} \leq c_5 \sqrt{\mathcal{V}(0)}.$$

Then, by the definition of T_α we must have $\mathcal{V}(T_\alpha) = \varepsilon$.

Case $T_\varepsilon < T_\alpha$: By the definition of T_α we know $\mathcal{V}(t) > \varepsilon$ and $|x^* - X_\alpha(\rho_t)| < 2c_5\sqrt{\mathcal{V}(0)}$ for all $t \in [0, T_\varepsilon]$. Then following the same argument as in the first case, one can conclude

$$\mathcal{V}(t) \leq \mathcal{V}(0) \exp(-(1-\tau)(2\lambda - \sigma^2)t), \quad \text{for } t \in [0, T_\varepsilon]. \quad (5.57)$$

Using the this estimate the fact that $T_\varepsilon = \frac{1}{(1-\tau)(2\lambda - \sigma^2)} \log(\frac{\mathcal{V}(0)}{\varepsilon})$ implies $\mathcal{V}(T_\varepsilon) \leq \varepsilon$, which is a contradiction. Thus this case can never happen. $\square$

5.3 Quantitative convergence result for the numerical scheme

We will use the Euler–Maruyama scheme from [62] to solve the particle system (2.1) numerically. For this purpose we define the orthogonal projection operator

$$\Pi_{\overline{\mathcal{D}}}(x) := \arg \min_{z \in \overline{\mathcal{D}}} |x - z|, \quad \text{for } x \in \mathbb{R}^d. \quad (5.58)$$

Note that for a given $x \in \mathbb{R}^d$, the projection $\Pi_{\overline{\mathcal{D}}}(x)$ exists and is uniquely determined due to the convexity of $\overline{\mathcal{D}}$. Given a time horizon $T > 0$ and a time discretization $t_0 = 0 < \Delta t < \dots < K\Delta t = T$ of $[0, T]$. Then the Euler–Maruyama scheme amounts to:

$$\begin{aligned} X_{(k+1)\Delta t}^i &= \Pi_{\overline{\mathcal{D}}}\left(X_{k\Delta t}^i - \Delta t \lambda (X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)) + \sigma D(X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)) N^i(0, \Delta t)\right) \\ X_0^i &\sim \rho_0, \quad i = 1, \dots, N, \end{aligned} \quad (5.59)$$

where $\{N^i(0, \Delta t)\}_{i=1}^N$ are independent Gaussian random vectors with zero mean and covariance matrix $\Delta t \mathbf{Id}_d$. Now collecting results from Theorem 4.3 and Theorem 5.4 we can establish a quantitative convergence result for the numerical scheme (5.59). We do so by controlling the following discrete counterpart of (5.4):

$$\mathcal{V}_N(t) := W_2^2(\rho_t^N, \delta_{x^*}) = \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N X_t^i - x^* \right|^2 \right], \quad t = k\Delta t, k \in [K]. \quad (5.60)$$

We are now ready to prove the main result of this paper (cf. Theorem 1.1):

Theorem 5.8 *Under the assumptions of Theorem 4.3 and Theorem 5.4 (or Theorem 5.7), let $\{(X_{k\Delta t}^i)_{k=1, \dots, K}\}_{i=1}^N$ be the iterations generated by Euler–Maruyama scheme (5.59) with $K\Delta t = T_*$, where T_* comes from Theorem 5.4 (or Theorem 5.7) such that $\mathcal{V}(T_*) = \varepsilon$ for any*

prescribed accuracy $\varepsilon > 0$. Then the final iterations fulfill the following quantitative error estimate

$$\mathcal{V}_N(T_*) \leq 3C_{\text{NA}} \Delta t \log(1/\Delta t) + 3C_{\text{MFA}} \frac{1}{N} + 3\varepsilon, \quad (5.61)$$

where C_{MFA} comes from Theorem 4.3, and C_{NA} depends on $\lambda, \sigma, \alpha, d, T_*, N$ and ε .

Proof Recall that $\{(X_t^i)_{t \in [0, T_*]}\}_{i=1}^N$ and $\{(\bar{X}_t^i)_{t \in [0, T_*]}\}_{i=1}^N$ are solutions to the CBO particle system (2.1) and N independent copies of solutions to the mean-field dynamics (2.4) up to time $T_* = K \Delta t$ respectively. Then we split the error

$$\begin{aligned} \mathcal{V}_N(T_*) \leq & 3\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (X_{K\Delta t}^i - X_{T_*}^i) \right|^2 \right] + 3\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (X_{T_*}^i - \bar{X}_{T_*}^i) \right|^2 \right] \\ & + 3\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \bar{X}_{T_*}^i - x^* \right|^2 \right], \end{aligned} \quad (5.62)$$

which divide the overall error into an approximation error of the Euler scheme, the mean-field limit estimate error and the optimization error in mean-field law. The first term on the right hand side of (5.62) can be estimated by applying the result from [62, Theorem 3.2], which yields

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (X_{K\Delta t}^i - X_{T_*}^i) \right|^2 \right] \leq C_{\text{NA}} \Delta t \log(1/\Delta t).$$

The second term can be bounded by using estimate from Theorem 4.3, which establishes

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N (X_{T_*}^i - \bar{X}_{T_*}^i) \right|^2 \right] \leq C_{\text{MFA}} \frac{1}{N}.$$

Finally the third term follows from Theorem 5.4 or Theorem 5.7, and it holds

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \bar{X}_{T_*}^i - x^* \right|^2 \right] \leq \varepsilon.$$

□

6 Numerical experiments and applications

In this section, we present numerical experiments, which are performed in Python, partially based on CBXpy [63], and are available for reproducibility at <https://github.com/echnen/CBO-with-boundaries>. Let us introduce now the rationale of these experiments.

The success of the CBO algorithm lies in its unique blend of exploration and exploitation. It achieves exploration through the diffusion of a large number of particles, while exploitation is facilitated by harnessing collective information to form a consensus point that approximates the global minimizer. The synergy of these effects—exploration through particle diffusion

and exploitation via the consensus point—becomes increasingly powerful as the number of particles grows. Theorem 5.8 guarantees that, as the number of particles N increases, the algorithm converges in expectation at a rate of $\mathcal{O}(N^{-1})$. This indicates that the accuracy improves significantly with more particles. However, the constants in the error bound (5.62), especially C_{MFA} , can increase exponentially with certain problem parameters, such as dimensionality. Consequently, it is challenging to directly apply these convergence results when using the CBO algorithm (5.59) with a limited number of particles N . In the few-particle regime, the effectiveness of the CBO algorithm exploration and exploitation mechanisms diminishes. To ensure successful convergence even with fewer particles, we must augment these mechanisms with specialized heuristics. These enhancements should enable the algorithm to keep dynamics similar to those observed with a large number of particles, thereby preserving its effectiveness despite the reduced particle count.

In this section, we introduce two key strategies to enhance the exploration and exploitation capabilities of the CBO algorithm in a few-particle regime.

1. **Enhanced exploitation:** To improve exploitation, we leverage the insight that the consensus point typically remains within an exponentially shrinking region around the global minimizer. Indeed, combining (5.32) or (5.52), for which

$$|x^* - X_\alpha(\rho_t)| \lesssim \sqrt{\mathcal{V}(t)}$$

with the exponential decay in (5.25) or (5.41)

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

we conclude that there is an exponentially shrinking *trust-region* $\{x \in \mathbb{R}^d : |x - X_\alpha(\rho_t)| \lesssim \exp(-(1 - \tau)(\lambda - \sigma^2/2)t)\}$ in which we can assume that particles are moving with high-probability. Based on this observation, even with fewer particles, we will—optimistically—constrain them to stay within progressively shrinking balls centered around the consensus-point, following a carefully planned shrinking schedule. Admittedly, such optimism would be misplaced if few particles would not explore enough, but this is counteracted by the Heuristic 2. which enforces at the same time an enhanced exploration¹.

2. **Enhanced exploration:** We propose a heuristic that boosts exploration by increasing the volatility in the noise. Specifically, we allow the noise variance σ^2 to exceed 2λ , contrary to the conditions set by Theorem 5.4. This adjustment aims to compensate for the reduced number of particles by providing greater variability in their movements.

¹ Another reasoning supporting Heuristic 1. is that trust-region methods [64], which are connected to CBO as explained in the Appendix, are known to converge (at least to local minimizers). While in trust-region methods the search is based on a quadratic interpolation and a guess-and-accept/reject mechanism for the trust radius, in CBO the search is based on stochastic exploration and the definition of consensus-point, which determine the trust-region. We refer to the Appendix for more insights.

These enhancements are designed to ensure that the dynamics of the algorithm remain effective and similar to those observed with a larger number of particles, even when operating in a few-particle regime.

To our knowledge, these two enhancement mechanisms—boosting exploration through increased noise volatility and enhancing exploitation by constraining particles to shrinking regions around the global minimizer—have not been explored in the existing literature. However, they enable us to obtain numerical results that were previously unobserved in previous numerical experiments, see, e.g., [16, 42, 52, 53].

For example, it has been well-documented that the standard CBO algorithm struggles to minimize the Rastrigin function in higher dimensions (e.g., $d \geq 20$), regardless of how reasonably large is the number of particles used. Remarkably, our enhanced CBO implementation surpasses this dimensionality barrier, demonstrating consistent convergence with a moderate number of particles. This breakthrough suggests that our approach can significantly improve the algorithm performance in high-dimensional optimization problems. We illustrate these findings in Section 6.1 and Section 6.2.

Moreover, as noted in [31, 65, 66], incorporating anisotropic noise or jump process noise can significantly enhance exploration in the few-particle regime, especially for high-dimensional optimization problems. This observation implies that adaptively shaping the noise based on the specific application is a crucial heuristic for optimizing the performance of the CBO algorithm. Tailoring the noise to suit the problem at hand can provide additional benefits, particularly when dealing with complex, high-dimensional landscapes.

Below, we will demonstrate how these adaptive noise strategies, combined with our previously discussed enhancements, contribute to improved convergence and robustness of the CBO algorithm even with a reduced number of particles. In particular, in Section 6.3 we explore the efficient computation of global minimizers of 1D p -Allen–Cahn energies with boundary conditions and additional obstacle constraints. The problem is solved by combining a hierarchical approximation and noise shaping within a multigrid finite element method. Given the existence of continua of (local) minimizers of the p -Allen–Cahn energy, the variational problems is extremely challenging. Yet, our hierarchical CBO solver does converge robustly to the global minimizer by using very few particles, despite the high-dimensionality of the optimization problem and, very importantly, without starting from initial data in the vicinity to the sought solution. To our knowledge, this is the first example in the literature of use of CBO for solving problems in scientific computing. These very promising results suggest that CBO could be employed to solve other challenging nonlinear equations, certainly a very interesting direction for future research.




Fig. 3 Domains and initial particles considered in Section 6.1 in dimension $d = 2$. The contoured orange shadow represents the Rastrigin function

6.1 Effect of the parameters

In this section, we explore the effect of the parameters for the convergence of (5.59) considering a standard nonlinear multimodal function for testing, i.e., the Rastrigin function:

$$R(x) := 10d + \sum_{i=1}^d (x_i^2 - 10 \cos(2\pi x_i)), \quad \text{for } x := (x_1, \dots, x_d) \in \mathbb{R}^d, \quad (6.1)$$

which we aim to minimize over the convex and nonconvex domains depicted in Figure 3.

In our numerical implementation, three key heuristics have been applied. The first, which is quite common, involves using an increasing sequence of parameters α generated by the update rule:

$$\alpha_k := \alpha_0 + \frac{k}{K}(\alpha_K - \alpha_0), \quad \text{for all } k \in \{0, \dots, K\}, \quad (6.2)$$

where $K \in \mathbb{N}$ is the final iteration of the method, and $\alpha_0, \alpha_K > 0$ are starting and final parameters defined by the user. If not specified otherwise, we set $\alpha_0 = 10^6$ and $\alpha_K = 10^9$. Similarly, we use default values of $\lambda = 1$, $N = 1000$, $\Delta t = 10^{-2}$, and run each simulation 1000 for iterations.

The next two are the two novel enhancements introduced in this paper, which we detail below.

Heuristic 1 (Enhanced Exploitation) At each iteration number $k \in \mathbb{N}$, we consider a ball centered on the consensus point $X_\alpha(\rho_{k\Delta t}^N)$ with radius $R := \gamma \max_i \{|X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)|\}$ with $\gamma \in (0, 1]$ and project every particle $X_{(k+1)\Delta t}^i$ onto this ball.

It is important to note that the parameter γ should be considered as an additional hyperparameter and should be tuned by the user. The second heuristic enhances particle exploration by allowing the noise variance σ^2 to exceed the theoretical bound 2λ :

Heuristic 2 (Enhanced Exploration) Given a drift parameter $\lambda > 0$, we exceed the admissible noise level $\sigma = \sqrt{2\lambda}$ given by Theorem 5.4 by setting $\sigma = S\sqrt{2\lambda}$ for $S > 1$.

We anticipate that, based on our experiments, a good choice for S is 5, independent of the dimension of the problem. Higher values of S can also be considered, especially when used




Fig. 4 Logarithm of the residual to the solution at final iteration, i.e., $\ln(\widehat{V}_N(T))$ for $\widehat{V}_N(t)$ defined as in (6.3) for different parameter's choices. The green line indicates the values of λ and σ for which the final residual coincides with the initial residual to the solution $\ln(\widehat{V}_N(0))$ and therefore separates the (numerical) converging regime from the non-converging one. For this parameter study, we consider dimension $d = 2$, $\alpha_0 = \alpha_K = 10^6$, and run each simulation for 100 iterations

in conjunction with Heuristic 1, as this combination helps mitigate excessive exploration away from the consensus point.

6.1.1 The effect of λ and σ

In this section, we investigate the effect of the parameters λ and σ in the presence of constraints. We consider² $\Omega_1 = [-6.12, 5.12]^2$, fix $\alpha_0 = \alpha_K = 10^6$, and pick 100 choices of λ and σ spaced evenly in a logarithmic scale from 10^{-1} to $10^{2.5}$ and from 10^{-1} to $10^{1.1}$, respectively. We initialize the particles sampling from a Gaussian distribution with mean $(5.12, 5.12)/\sqrt{2}$ and variance 10, in this way the initial discrete distribution is likely to cover the whole domain while not being centered on the optimal solution, but actually on a local (non global) minimizer. We apply a further projection to the domain for those particles that fell outside Ω . Once the initial measure is fixed, for each λ and σ we let the algorithm run for 100 iterations and keep track of the error at the last iteration. This is computed in each iteration t by an empirical approximation of $\mathcal{V}_N(t)$ according to (5.60), namely,

$$\widehat{V}_N(t) := \frac{1}{N} \sum_{i=1}^N |X_t^i - x^*|^2, \quad (6.3)$$

where $x^* \in \mathbb{R}^d$ is the optimal solution, in this case $x^* = 0$. For the sake of comparison, we run exactly the same experiment with $\Omega = \mathbb{R}^2$, i.e., simply dropping the additional projection onto the bounded domain and by considering the same initialization without the further projection onto Ω . Additionally, we repeat the same experiment with isotropic noise, for which the critical line (at least for unbounded domains) is $2\lambda - d\sigma^2 = 0$, compare, e.g., [28]. We plot all the logarithmic values of these errors in Figure 4.

We learn that utilizing anisotropic noise and introducing constraints do in fact influence the convergence of the CBO scheme allowing for larger diffusion parameters. Indeed, in

² Note that the tiltedness of the domain is intentional.

Figure 4 we observe that in these cases, the critical line $2\lambda - \sigma^2 = 0$ is no longer sharp. This interesting behaviour, which gives a first numerical evidence of the effectiveness of Heuristic 2, is completely unexpected and suggests that the analysis in these settings could be further refined. We leave these intriguing questions for future work.

6.1.2 Ablation study on the heuristics

In this ablation study, we show that the combination of Heuristics 1 and 2 can indeed be beneficial in practice, and only with both heuristics being active, solving problems in high dimensions becomes feasible. Specifically, we consider dimensions ranging in $d = 2, 15, 20$ and consider Heuristics 2 with $S = 5$, i.e., we make the non-feasible parameter choice:

$$\lambda = 1, \quad \text{and} \quad \sigma = 5\sqrt{2\lambda}. \quad (6.4)$$

Note that these values fall outside the range of admissible parameters in Theorem 5.4. To test the influence of the parameter γ in Heuristics 1, we considered two possible scenarios:

$$\Omega_1 := [-6.12, 5.12]^d, \quad \text{and} \quad \Omega_2 := [0, 11.24]^d.$$

Note that in the latter, the global minimizer lies on the boundary of Ω . We initialize the particles sampling from a Gaussian with mean $(5.12, \dots, 5.12)/\sqrt{d}$, and variance 10. Some initial particles might fall out of Ω : We fix this by projecting all of these onto Ω . We run the algorithm 20 times in each case and show the results in the two top rows of Figure 5. For the sake of comparison, we also repeat the experiment with a parameter choice that is indeed admissible in the sense of Theorem 5.4, namely, $\sigma = \sqrt{2\lambda}$ and $\lambda = 1$, and show the results in the two bottom rows of Figure 5.

In this example, we observe that Heuristic 1 has a beneficial effect on the performance of the proposed method. Specifically, as demonstrated in the top two rows of Figure 5, it can significantly attenuate oscillations, which is a particularly desirable effect in optimization [67]. Additionally, in other cases, such as in the bottom row of Figure 5 (left graphic), it can markedly speed up convergence. It is also worth noting that we did not encounter any instance where Heuristics 1 with $\gamma \geq 1/2$ deteriorated the convergence performance of the method. However, we did not present cases with $\gamma < 1/2$ as those indeed performed poorly. The influence of the heuristics is further quantified in Table 2. We display the success rate for the scenarios run on Ω_2 , i.e., the second and fourth row of Figure 5, whereby a run is considered successful if $W_2^2(\rho_T^N, \delta_{x^*}) < 10^{-10}$. It becomes apparent that both heuristics have to be active for optimizations in high dimensions to become feasible.

To further corroborate our numerical heuristics, we run Algorithm (1.8) in the same setting, fixing the dimension to $d = 5$ and choosing $\Omega = \mathbb{R}^d$. We consider 20 values of S evenly spaced in $[0.5, 10]$, first applying Heuristic 1 ($\gamma = 0.75$) and then disabling it ($\gamma = 1$). For each value of S , we repeat the experiment five times and report the final residual at time $T = 10$ in Figure 6. The figure illustrates the positive impact of Heuristic 2, which achieves its best performance for $S \simeq 5$, in agreement with the choice made in (6.4).




Fig. 5 Error decrease as a function of time for the dimensions: $d = 2, 15, 20$. In each case, we run the algorithm 20 times and display the average in thicker lines. In the first and second row, Heuristics 1. and 2. are active for $\sigma = 5\sqrt{2\lambda}$ and $\gamma < 1$, while in the third and fourth row only Heuristics 1. is active and we show the behavior for the critical value $\sigma = \sqrt{2\lambda}$, as in standard CBO

Table 2 Success rate depending on Heuristic 1 ($\checkmark$: $\gamma = 0.75$, $\times$: $\gamma = 1$) and Heuristic 2 ($\checkmark$: $\sigma = 5\sqrt{2\lambda}$, $\times$: $\sigma = \sqrt{2\lambda}$) being on or off. We define success as $W_2^2(\rho_t^N, \delta_{x^*}) < 10^{-10}$ and, in each scenario, run the algorithm 20 times on $\Omega_2 = [0, 11.24]^d$

Heuristic 1 ($\gamma = 0.75$)	Heuristic 2 ($\sigma = 5\sqrt{2\lambda}$)	$d = 2$	$d = 15$	$d = 20$
$\times$	$\times$	0%	0%	0%
$\checkmark$	$\times$	100%	0%	0%
$\times$	$\checkmark$	90%	35%	25%
$\checkmark$	$\checkmark$	100%	95%	80%




Fig. 6 Final residual as a function of the exploration parameter S in Heuristic 2, first with Heuristic 1 enabled and then without it. The vertical red line marks the critical stability threshold at $S = 1$. We consider $\sigma = S\sqrt{2\lambda}$ with $\lambda = 1$




Fig. 7 Error decrease as a function of time for $d = 100$ with parameters (6.4). In each case, we run the algorithm 20 times and display the average with a thicker line. Curves (a) and (b) represent the $\Omega = \Omega_2$ and $\Omega = \mathbb{R}^d$ cases employing Heuristics 1 ($\gamma = 0.95$) and 2 ($\sigma = 5\sqrt{2\lambda}$, $\lambda = 1$), while (c) represents the standard CBO ($\sigma = \sqrt{2\lambda}$, $\lambda = 1$) for $\Omega = \mathbb{R}^d$

Eventually, we consider a relatively difficult high dimensional problem with $d = 100$. We consider the constraint $\Omega := \Omega_2$ and initialize the particles as above by sampling from a Gaussian distribution with mean $(5.12, \dots, 5.12)/\sqrt{d}$ and variance 10 and projecting onto the domain. We run 20 independent experiments by considering all enhancements: Increasing parameters α as in (6.2), Heuristic 1 with $\gamma = 0.95$ and Heuristic 2 with $S = 5$, and show the residual decrease as a function of the iterations in Figure 7. For comparison, we also show the performance in the unconstrained setting but employing Heuristic 1 and Heuristic 2, and in the unconstrained setting without employing any heuristic. While the latter does not converge to an optimal solution, in the first two cases the parameter choice (6.4) allows us to optimize the Rastrigin function in $d = 100$ dimensions, a result never obtained so far in the literature with CBO.

6.1.3 Quantitative comparisons with respect to N

We compare the performance of the standard CBO method ($\lambda = 1$, $\sigma = \sqrt{2\lambda}$), CBO with enhanced exploration ($\lambda = 1$, $\sigma = 5\sqrt{2\lambda}$, $\gamma = 1$), and CBO with both enhanced exploration




Fig. 8 Quantitative comparison of CBO variants and PSO. In each case, we run the algorithm 20 times and display the average of the discrete Wasserstein distance after 1000 iterations

and enhanced exploitation ($\lambda = 1$, $\sigma = 5\sqrt{2\lambda}$, $\gamma = 0.75$). As a benchmark, we include a particle swarm optimization (PSO) algorithm based on [11], which is implemented as part of the PySwarms package [68]. Unless specified otherwise, parameters are as chosen in Section 6.1.1. We study the performance of all methods using increasing particle numbers $N \in \{1, 2, 4, \dots, 2^{14}\}$, applied to the Rastrigin problem in dimensions $d = 5, 10, 15, 20$ on the domain Ω_1 . The particles of each method are initialized and constrained to remain within the domain. Each simulation is run 20 times, and performance is measured by the mean discrete Wasserstein distance of all runs evaluated after 1000 iterations. A low score on this measure indicates that all runs perform well, while a single poor result skews the performance measure towards a high score. The results are displayed in Figure 8.

While the overall performance of all methods improves with the number of particles, none of the algorithms reliably solves the optimization problem with $N \leq 128$ particles. The CBO variants with enhanced exploration effectively handle the optimization problem when using more than $N = 1024$ particles, with the variant with enhanced exploitation ($\gamma = 0.75$) outperforming all other methods. The standard CBO performs well for $d = 5$ and $N \geq 2048$, but its performance deteriorates in higher dimensions. The PSO algorithm, on the other hand, struggles to reliably solve the optimization problem across all scenarios, while all considered CBO variants yield better or at least similar results.

6.2 Performance over nonconvex domains

Motivated by the excellent results achieved, we decided to test the method further by exploring nonconvex domains, despite the current lack of theoretical backing. Specifically, for each dimension $d = 2, 15, 20$ we consider two cases:

$$\Omega_{i+2} := B_r(0) \setminus (B_1(\bar{x}_i) \cup B_1(-\bar{x}_i)) \quad , \quad i = 1, 2,$$

For the following experiment, we consider the default parameters outlined in Section 6.1 and choose $\lambda = 1$ as well as $\sigma = 5\sqrt{2\lambda}$. As shown in Figure 9, the proposed method is able to deliver excellent performances also in the case of a nonconvex domain with shape Ω_3 . How-




Fig. 9 Residual decrease as a function of the iterations over two nonconvex domains, for details refer to Section 6.2

ever, we believe that the anisotropic bias of the noise along axes makes this case particularly lucky and likely explains the significant difference in Figure 9 between optimizing on Ω_3 or Ω_4 . This intriguing result suggests that optimal performance might be achievable only with an *adaptive covariance* of the noise. Unfortunately, we have yet to model, implement, and analyze a version of CBO with adaptive covariance, but some techniques in this direction, such as those in [69], seem to open the door to such variations also for CBO diffusion.

In the next section, we move away from standard benchmarks, such as the Rastrigin function, and present results on the use of CBO for solving partial differential equations with constraints.

6.3 Solving a PDE problem with an obstacle

Partial differential equations and optimization problems are inherently intertwined, as many PDE systems emerge as gradient flows of free energy functionals. Noticeably, the Allen–Cahn and Cahn–Hilliard equations can be derived as the L^2 and H^{-1} gradient flows of the Ginzburg–Landau (GL) free energy functional, respectively, see, e.g., [70, 71]. By employing CBO with hierarchically structured noise, we efficiently compute the global minimizer of a variational p -Allen–Cahn-type problem within a finite element framework.

Let $\mathcal{E}_{GL}(v)$ be the p -GL free energy functional over the interval $[0, 1]$

$$\mathcal{E}_{GL}(v) = \int_{[0,1]} \frac{1}{p} |\nabla v(s)|^p + \frac{1}{\epsilon^2} F(v(s)) ds, \quad (6.5)$$

where $F(v) = (v - w_1)^2(v - w_2)^2$ is a so-called double well potential. The minimization problem reads as

$$v^* = \arg \min_{v \in V} \mathcal{E}_{GL}(v), \quad (6.6)$$

where $V = \{v \in W^{1,p}[0, 1] \mid v(0) = v_0, v(1) = v_1\}$. We discretize the continuous minimization problem using the finite element method. The interval $[0, 1]$ is discretized into M one-dimensional elements $T_j = [(j-1)h, jh]$, $j = 1, \dots, M$ of equal length $h = M^{-1}$, and we define the finite element space V_M

$$V_M = \{v \in V \mid v|_{T_j} \in \mathcal{P}_1(T_j), \forall j \in \{1, \dots, M\}\}. \quad (6.7)$$

Expanding $v \in V_M$ in terms of first-order Lagrange basis functions $\{\phi_j\}_{j=0}^M$, such that $v = \sum_{j=0}^M \phi_j x_j$, whereby $x_0 = v_0, x_M = v_1$, we obtain the discrete formulation

$$x^* = \arg \min_{x \in \mathbb{R}^{M-1}} \sum_{j=1}^M \left(\frac{h}{p} \left| \frac{x_j - x_{j-1}}{h} \right|^p + \frac{1}{\epsilon^2} \tilde{F}_j \right). \quad (6.8)$$

The integral over the double-well potential $\tilde{F}_j = \int_{(j-1)h}^{jh} F(v(t))dt$ can be calculated numerically or even analytically. The optimization problem is high-dimensional with $d = M - 1$ and can be associated with $V_{M,0} = \text{span} \{\phi_j, j = 1, \dots, M - 1\}$. Due to the double-well potential, numerous local minima do exist, which, depending on ϵ , retain arbitrarily steep basins of attraction. Since the GL free energy function depends on the gradient of v , Gaussian noise is badly suited for the optimization, as it does not reflect the smoothness of the solution. Therefore, naively applying CBO for fine discretizations requires a large number of particles, small time steps, and large times, and thus, becomes unfeasible due to the computational cost.

6.3.1 Hierarchical structured noise

We overcome these limitations using a hierarchical approach similar to a multi-grid method. We set $M = 2^m$ and solve subsequently the optimization problem over a nested sequence of finite element spaces $V_2 \subset V_4 \subset V_8 \subset \dots \subset V_{2^m}$ using the final particle distribution of the previous run as the initial particle distribution of the following. Thereby, we exploit that CBO solves the coarse problem efficiently and can find the solution for fine resolutions, provided that the initial particle distribution is sufficiently close to the solution. Similar hierarchical approaches are well studied in the context of solving obstacle problems; see [72–74].

Alternatively, the hierarchical approach can be implemented using the highest resolution and solely changing the noise of the algorithm. During the i th run, at each time step and for every particle, we replace the standard anisotropic noise vector $\beta = \sigma D(X_{k\Delta t}^j - X_\alpha(\rho_{k\Delta t}^N))N^j(0, \Delta t) \in \mathbb{R}^d$ defined in (1.8) with $\beta^i \in \mathbb{R}^d$, such that

$$\beta_l^i = \int_0^1 \phi_l^m(s) \sum_{k=1}^{2^i-1} \phi_k^i(s) \beta_{k2^{m-i}} ds, \quad l = 1, \dots, 2^m - 1, \quad (6.9)$$

where $\{\phi_l^m\}_{l=1}^{2^m-1}$ and $\{\phi_k^i\}_{k=1}^{2^i-1}$ are first-order Lagrangian bases of $V_{2^m,0}$ and $V_{2^i,0}$, respectively. Accordingly, the initial particle distribution of the first CBO run is drawn with respect to the coarsest resolution. We note that by this construction, we automatically decrease the




Fig. 10 Particle distributions and consensus points at the end of the CBO runs for different $M = 2^i$




Fig. 11 Development of the energy 11a and the consensus points at the end of CBO runs with dimensions $d = 2^i - 1$, $i = 2, \dots, 7$ and $w_2 = 0.75$ (gray dotted line)

correlation length of our noise with each refinement step. The resulting particle distributions are visualized in Figure 10c.

For the numerical experiment, we choose $v_0 = 0.5$, $v_1 = 1$, $w_1 = 0.25$, $w_2 = 0.75$, $\epsilon^{-2} = 500$, $p = 1.5$ and consider spatial resolutions with $M = 2^i$, $i = 2, \dots, 7$ elements. The parameters of the algorithm are chosen as $\Delta t = 10^{-2}$, $N = 20d = 2540$, $\lambda = 1$, $\sigma = 7$, and $\alpha = 10^6$. We use 100 iterations at every resolution except the finest, and 1000 iterations at the finest level. As visualized in Figure 11a, the energy of the consensus of the CBO run quickly decreases with every refinement step, requiring only a few iterations before reaching the optimizer of the current resolution, and thus, stagnating until the next refinement step. Figure 11b shows the convergence of the consensus points at the end of every refinement step, whereby the optimal solutions of coarse resolutions, e.g., $i = 2$, may qualitatively differ from those of fine resolutions. By incorporating hierarchically structured noise and tailoring the algorithm's randomness to the specific optimization problem, we effectively solve a 127-dimensional problem using only 2540 particles and 1500 iterations.




Fig. 12 Various obstacles and consensus points for the p -Allen–Cahn problem described in Section 6.3.1 with $p = 1.5$, $\epsilon^{-2} = 500$. The obstacle constraints influence the consensus point despite not being active

6.3.2 Obstacle problem

Bridging the gap between global optimization on bounded domains and the finite element formulation previously discussed, we impose an obstacle constraint to the p -Allen–Cahn problem. Specifically, we enforce that the solution $v(x)$ satisfies $g(x) \leq v(x) \leq f(x)$ at all grid points of the finest level $x \in \{0, h, 2h, \dots, 1\}$. Thus, the resulting discrete optimization problem is defined within a convex 127-dimensional hypercuboid. To reduce the complexity, we impose the constraints only at the grid points of the current resolution, beginning with a coarse obstacle and progressively refining it in tandem with the noise resolution. Thus, the solution obtained under coarse obstacle constraints can exhibit qualitative differences compared to those obtained with finer resolutions. Different obstacle constraints and the corresponding consensus points are visualized in Figure 12, emphasizing the influence of the constraints on the consensus despite not being active.

We apply the CBO to different parameter combinations for p and ϵ with 1000 and 10000 iterations for the coarser and the finest resolution, respectively. Note that for $p = 1$, $\epsilon^{-2} = 2000$ and $p = 1.5$, $\epsilon^{-2} = 10000$, there exist multiple global optimizers. The development of the energy of the best particle and the final consensus point are visualized in Figure 13. As the obstacle evolves across varying refinement levels, a solution that was previously optimal may become infeasible or suboptimal following the refinement step, resulting in a non-monotonic energy curve. The progression of the energy associated with the best-performing particle is plotted relative to the minimal energy achieved at the current resolution. CBO converges fast for coarse resolution yet requires more iterations as the resolution of the obstacle and noise becomes finer. In Figure 13a, for $p = 1$ and $p = 1.25$, the energy curve exhibits a plateau before eventually declining as the algorithm overcomes a local minimizer. We obtain qualitatively good solutions, even for parameter settings characterized by local minima with steep basins of attraction.

By incorporating hierarchically structured noise, we obtain efficiently the global minimizer of the 1D p -Allen–Cahn energy discretized by low order H^1 -conforming elements with a reasonably small number of particles. This advancement extends the applicability of CBO to a broader class of PDE-related problems. Furthermore, it highlights the potential of





(a) Energy development, $\epsilon^{-2} = 2000$ (b) Consensus points, $\epsilon^{-2} = 2000$ (c) Energy development, $p = 1.5$ (d) Consensus points, $p = 1.5$

Fig. 13 Development of the energy of the best particle and the consensus point of different p -Allen-Cahn problems with an obstacle. The gray dotted lines in (B) and (D) show the double-well potential with $w_1 = 0.25$ and $w_2 = 0.75$

problem-specific noise to enhance algorithmic performance, to significantly decrease computational costs by reducing the number of particles, and to identify global optimizers in even higher-dimensional spaces.

Acknowledgements E.C. acknowledges the support of the project P 34922-N of the Austrian Science Fund (FWF). M.F. acknowledges the support of the Munich Center for Machine Learning and the ERC Advanced Grant NEITALG, grant agreement No. 101198055.

Funding Open Access funding enabled and organized by Projekt DEAL. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This project has received funding from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement No. 101198055, project acronym NEITALG).

Data availability All data and numerical experiments presented in this paper are partially based on CBXpy [63] and are available for reproducibility at <https://github.com/echnen/CBO-with-boundaries>.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

Consensus-based optimization within the global optimization landscape

Here, we refer to Figure 1, which illustrates how Consensus-Based Optimization (CBO) fits within the broader landscape of global optimization. We begin with the original 2017 paper [16], where the authors introduced CBO as a hybrid approach, combining elements of Particle Swarm Optimization (PSO) and Simulated Annealing (SA). While the connections in [16] were initially based on arguments by analogy, subsequent works have rigorously established them. In particular, [14, 27] showed that CBO can be derived as the zero-inertia limit of a suitably formulated version of PSO. This insight enabled the first rigorous analysis of global convergence for PSO, as proposed in [13].

The fact that CBO is a form of SA was made rigorous in the paper [28], where one of us proved that as the number of particles $N \rightarrow +\infty$ and for $\alpha \rightarrow +\infty$, each particle tends to the solution of the following annealed Langevin dynamics:

$$dX(t) = -\nabla \mathcal{F}(X(t)) + \sigma |X(t) - x^*| dB_t,$$

where $\mathcal{F}(x) = \frac{\lambda}{2} |x - x^*|^2$.

Further surprising connections of CBO with other relevant global optimization methods have been derived in [26] through the Euler-Maruyama approximation of the CBO dynamics (1.2):

$$X_{(k+1)\Delta t}^i = X_{k\Delta t}^i - \Delta t \lambda (X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)) + \sqrt{\Delta t} \sigma |X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)| N_k^i(0, 1), \quad (6.10)$$

By scaling $\lambda = \frac{1}{\Delta t}$ one obtains

$$X_{(k+1)\Delta t}^i = X_\alpha(\rho_{k\Delta t}^N) + \sqrt{\Delta t} \sigma |X_{k\Delta t}^i - X_\alpha(\rho_{k\Delta t}^N)| N_k^i(0, 1), \quad (6.11)$$

and, consequently, the so-called *Consensus Hopping scheme (CH)* as introduced in [26]

$$X_{k+1}^{CH} = X_\alpha(\mathcal{N}(X_k^{CH}, \sigma_k^2)), \quad (6.12)$$

which is the consensus point of a Gaussian centered in the previous consensus point with a suitable variance σ_k^2 . Let us first notice that the CH scheme is an instance of $(1, \lambda)$ Evolution

Strategies (ES), with $\mu = 1$ parent and $\lambda = N$ offspring³. Proofs of global convergence for ES were obtained long ago, for instance in [76, 77], using arguments based on Markov chains. However, most classical proofs of global convergence for ES do not provide explicit convergence rates. Based on the connection between CBO and ES obtained through (6.10)–(6.12), it is now possible to also establish exponential convergence rates, as discovered in the forthcoming paper [29]. Connections between CBO and ES have been further investigated in [30].

By the Laplace principle, the explicit expression for the consensus point in (6.12) can be approximated, for a suitable scaling of $\alpha \rightarrow \infty$ with respect to $\sigma_k \rightarrow 0$, by

$$\begin{aligned} X_{k+1}^{CH} &= X_\alpha(\mathcal{N}(X_k^{CH}, \sigma_k^2)) \\ &= \int_{\mathbb{R}^d} x e^{-\alpha \left(\mathcal{E}(x) + \frac{1}{2\sigma_k^2 \alpha} |x - X_k^{CH}|^2 \right)} dx \\ &\approx \arg \min_{x \in \mathbb{R}^d} \mathcal{E}(x) + \frac{1}{2\tau} |x - X_k^{CH}|^2, \end{aligned} \quad (6.13)$$

which can be readily recognized as a *Minimizing Movement Scheme (MMS)* or a *Proximal Point Method*. In other words, the iteration of CH follows approximately a Gradient Descent (GD), a surprising connection that was recently made rigorous in [26], where it was shown that CBO can be interpreted as a suitable stochastic relaxation of GD.

Furthermore, the MMS is just one instance of a broader class of proximal point schemes of the form

$$X_{k+1} = \arg \min_{x \in \mathbb{R}^d} \mathcal{E}(x) + \frac{1}{p\tau^{p-1}} |x - X_k|^p, \quad (6.14)$$

for any $1 < p < \infty$. As $p \rightarrow +\infty$, such schemes converge to

$$X_{k+1} = \arg \min_{x \in B(X_k, R_\tau)} \mathcal{E}(x), \quad (6.15)$$

which restricts the search for the next approximation to an optimal point within a *Trust Region (TR)*. Notice that for $\tau \rightarrow 0$ suitable piecewise interpolations of both iterations in (6.14) and (6.15) converge to a gradient flow trajectory with different time reparametrizations, a well-known insight since the 1993 groundbreaking work of De Giorgi [78]. Hence, the conceptual path from (6.12) to (6.15) illustrates how CBO is also connected to both GD and TR methods [64]. We conclude mentioning the work [79] where *Genetic Algorithms (GA)* are also put in connection to CBO through kinetic models [80].

In conclusion, CBO occupies a very central position (Figure 1) within a composite landscape of related methods for global optimization. In light of these connections, the powerful mathematical techniques very recently developed for analyzing the global convergence of CBO (starting with [15, 28, 31–34]), as further developed in this paper, have also enabled

³ We recall that a (μ, λ) -Evolution Strategy [75] maintains a population of μ candidate solutions, called parents. In each iteration, the algorithm generates λ new candidate solutions, called offspring, by applying a “mutation” to the parents. The mutation operator typically adds random perturbations drawn from a normal distribution to the parent solutions. (Compare (6.12).)

further breakthroughs, such as describing the global convergence of PSO [13] and deriving convergence rates for ES [29]. We believe that, given their broad applicability, the significance of these mathematical techniques cannot be overstated, and they have the potential to inspire further advances in global optimization; this paper wants to represent a meaningful contribution in that direction.

Proof of Proposition 5.2

Proof To simplify the notation, in the following, we drop the subscript x_0 in $\phi_r^{x_0}$, and it is easy to compute

$$\frac{d\mathbb{E}[\phi_r(\bar{X}_t)]}{dt} = \sum_{k=1}^d \mathbb{E}[T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t)] \quad (6.16)$$

with

$$T_{1,k} := -\lambda(\bar{X}_t - X_\alpha(\rho_t))_k \partial_{x_k} \phi_r(\bar{X}_t) \text{ and } T_{2,k} := \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 \partial_{x_k}^2 \phi_r(\bar{X}_t). \quad (6.17)$$

Here the boundary term disappears because $\text{supp}(\phi_r) = B_r(x_0) \cap \mathcal{D} \subset \mathcal{D}$. For each $k = 1, \dots, d$, let us define the subsets

$$K_{1,k} := \{x \in \bar{\mathcal{D}} : |(x - x_0)_k| > \sqrt{cr}\} \quad (6.18)$$

and

$$K_{2,k} := \left\{x \in \bar{\mathcal{D}} : -\lambda(x - X_\alpha(\rho_t))_k (x - x_0)_k (r^2 - (x - x_0)_k^2)^2 \right. \quad (6.19)$$

$$\left. > \tilde{c}r^2 \frac{\sigma^2}{2} (x - X_\alpha(\rho_t))_k^2 (x - x_0)_k^2 \right\} \quad (6.20)$$

where $\tilde{c} := 2c - 1 \in (0, 1)$. Then we decompose $B_r(x_0)$ according to

$$B_r(x_0) = (K_{1,k}^c \cap B_r(x_0)) \cup (K_{1,k} \cap K_{2,k}^c \cap B_r(x_0)) \cup (K_{1,k} \cap K_{2,k} \cap B_r(x_0)) \quad (6.21)$$

and one considers each of these subsets separately.

Subset $K_{1,k}^c \cap B_r(x_0)$: We have $|(\bar{X}_t - x_0)_k| \leq \sqrt{cr}$ for each $\bar{X}_t \in K_{1,k}^c$. For $T_{1,k}$, we use the expression of $\partial_{x_k} \phi_r$ and get

$$T_{1,k}(\bar{X}_t) = 2r^2 \lambda (\bar{X}_t - X_\alpha(\rho_t))_k \frac{(\bar{X}_t - x_0)_k}{(r^2 - (\bar{X}_t - x_0)_k^2)^2} \phi_r(\bar{X}_t) \quad (6.22)$$

$$\geq -2r^2 \lambda \frac{|(\bar{X}_t - X_\alpha(\rho_t))_k| |(\bar{X}_t - x_0)_k|}{(r^2 - (\bar{X}_t - x_0)_k^2)^2} \phi_r(\bar{X}_t) \geq -\frac{2\lambda(\sqrt{cr} + |\mathcal{D}|)\sqrt{c}}{(1-c)^2 r} \phi_r(\bar{X}_t) \quad (6.23)$$

$$=: -\vartheta_1 \phi_r(\bar{X}_t), \quad (6.24)$$

where in the last inequality we have used

$$|(\bar{X}_t - X_\alpha(\rho_t))_k| \leq |(\bar{X}_t - x_0)_k| + |(X_\alpha(\rho_t) - x_0)_k| \leq \sqrt{c}r + |\mathcal{D}|. \quad (6.25)$$

Similarly for $T_{2,k}$ one obtains

$$T_{2,k}(\bar{X}_t) \quad (6.26)$$

$$\begin{aligned} &= \sigma^2 r^2 (\bar{X}_t - X_\alpha(\rho_t))_k^2 \frac{2 \left(2 (\bar{X}_t - x_0)_k^2 - r^2 \right) (\bar{X}_t - x_0)_k^2 - \left(r^2 - (\bar{X}_t - x_0)_k^2 \right)^2}{\left(r^2 - (\bar{X}_t - x_0)_k^2 \right)^4} \phi_r(\bar{X}_t) \\ &\geq -\frac{\sigma^2 (\sqrt{c}r + |\mathcal{D}|)^2 (2c + 1)}{(1 - c)^4 r^2} \phi_r(\bar{X}_t) =: -\vartheta_2 \phi_r(\bar{X}_t). \end{aligned} \quad (6.27)$$

Subset $K_{1,k} \cap K_{2,k}^c \cap B_r(x_0)$: As $\bar{X}_t \in K_{1,k}$ we have $|(\bar{X}_t - x_0)_k| > \sqrt{c}r$. We observe that $T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t) \geq 0$ for all $\bar{X}_t$ satisfying

$$\left(-\lambda (\bar{X}_t - X_\alpha(\rho_t))_k (\bar{X}_t - x_0)_k + \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 \right) \left(r^2 - (\bar{X}_t - x_0)_k^2 \right)^2 \quad (6.28)$$

$$\leq \sigma^2 (\bar{X}_t - X_\alpha(\rho_t))_k^2 \left(2 (\bar{X}_t - x_0)_k^2 - r^2 \right) (\bar{X}_t - x_0)_k^2. \quad (6.29)$$

Indeed, this can be verified by first showing that

$$-\lambda (\bar{X}_t - X_\alpha(\rho_t))_k (\bar{X}_t - x_0)_k \left(r^2 - (\bar{X}_t - x_0)_k^2 \right)^2 \quad (6.30)$$

$$\leq \tilde{c} r^2 \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 (\bar{X}_t - x_0)_k^2 \quad (6.31)$$

$$= (2c - 1) r^2 \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 (\bar{X}_t - x_0)_k^2 \quad (6.32)$$

$$\leq \left(2 (\bar{X}_t - x_0)_k^2 - r^2 \right) \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 (\bar{X}_t - x_0)_k^2, \quad (6.33)$$

where we have used the fact that $\bar{X}_t \in K_{1,k} \cap K_{2,k}^c$ and $\tilde{c} = 2c - 1$. One also notice that

$$\frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 \left(r^2 - (\bar{X}_t - x_0)_k^2 \right)^2 \leq \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 (1 - c)^2 r^4 \quad (6.34)$$

$$\leq \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 (2c - 1) r^2 c r^2 \quad (6.35)$$

$$\leq \frac{\sigma^2}{2} (\bar{X}_t - X_\alpha(\rho_t))_k^2 \left(2 (\bar{X}_t - x_0)_k^2 - r^2 \right) (\bar{X}_t - x_0)_k^2 \quad (6.36)$$

by using $(1 - c)^2 \leq (2c - 1)c$. Hence (6.28) holds and we have $T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t) \geq 0$.

Subset $K_{1,k} \cap K_{2,k} \cap B_r(x_0)$: Notice that when $(\bar{X}_t)_k = (X_\alpha(\rho_t))_k$, we have $T_{1,k} = T_{2,k} = 0$, so in this case there is nothing to prove. If $(\bar{X}_t)_k \neq (X_\alpha(\rho_t))_k$, or $\sigma^2 (\bar{X}_t - X_\alpha(\rho_t))_k^2 > 0$

($\sigma > 0$), we exploit $\bar{X}_t \in K_{2,k}$ to get

$$\frac{(\bar{X}_t - X_\alpha(\rho_t))_k (\bar{X}_t - x_0)_k}{(r^2 - (\bar{X}_t - x_0)_k^2)^2} \geq \frac{-|(\bar{X}_t - X_\alpha(\rho_t))_k| |(\bar{X}_t - x_0)_k|}{(r^2 - (\bar{X}_t - x_0)_k^2)^2} \quad (6.37)$$

$$> \frac{2\lambda (\bar{X}_t - X_\alpha(\rho_t))_k (\bar{X}_t - x_0)_k}{\tilde{c}r^2\sigma^2 |(\bar{X}_t - X_\alpha(\rho_t))_k| |(\bar{X}_t - x_0)_k|} \geq -\frac{2\lambda}{\tilde{c}r^2\sigma^2}. \quad (6.38)$$

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

$$T_{1,k}(\bar{X}_t) = 2r^2\lambda (\bar{X}_t - X_\alpha(\rho_t))_k \frac{(\bar{X}_t - x_0)_k}{(r^2 - (\bar{X}_t - x_0)_k^2)^2} \phi_r(\bar{X}_t) \quad (6.39)$$

$$\geq -\frac{4\lambda^2}{\tilde{c}\sigma^2} \phi_r(\bar{X}_t) =: -\vartheta_3 \phi_r(\bar{X}_t). \quad (6.40)$$

Moreover since $\bar{X}_t \in K_{1,k}$ and $2(2c-1)c \geq (1-c)^2$ implied by the assumption, one has

$$2 \left(2(\bar{X}_t - x_0)_k^2 - r^2 \right) (\bar{X}_t - x_0)_k \geq (r^2 - (\bar{X}_t - x_0)_k^2)^2, \quad (6.41)$$

which yields that $T_{2,k}(\bar{X}_t) \geq 0$.

Concluding the proof: Collecting estimates from above, we get

$$\begin{aligned} \frac{d\mathbb{E}[\phi_r(\bar{X}_t)]}{dt} &= \sum_{k=1}^d \mathbb{E}[T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t)] \\ &= \sum_{k=1}^d \left(\mathbb{E}[(T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t)) \mathbf{I}_{K_{1,k} \cap K_{2,k} \cap B_r(x_0)}] \right. \\ &\quad + \mathbb{E}[(T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t)) \mathbf{I}_{K_{1,k} \cap K_{2,k}^c \cap B_r(x_0)}] \\ &\quad \left. + \mathbb{E}[(T_{1,k}(\bar{X}_t) + T_{2,k}(\bar{X}_t)) \mathbf{I}_{K_{1,k}^c \cap B_r(x_0)}] \right) \\ &\geq -d(\vartheta_1 + \vartheta_2 + \vartheta_3) \mathbb{E}[\phi_r(\bar{X}_t)] =: -\vartheta \mathbb{E}[\phi_r(\bar{X}_t)]. \end{aligned} \quad (6.42)$$

An application of Gronwall's inequality concludes that

$$\mathbb{P}(\bar{X}_t \in B_r(x_0)) \geq \mathbb{E}[\phi_r(\bar{X}_t)] \geq \mathbb{E}[\phi_r(\bar{X}_0)] \exp(-\vartheta t). \quad (6.43)$$

□

References

1. Banholzer, D., Fliege, J., Werner, R.: A radial basis function method for noisy global optimisation. *Mathematical Programming*, 1–44 (2024)
2. Vinod, A.P., Israel, A., Topcu, U.: Constrained, global optimization of unknown functions with Lipschitz continuous gradients. *SIAM J. Optim.* **32**(2), 1239–1264 (2022)

3. Blum, C., Roli, A.: Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Comput. Surv.* **35**(3), 268–308 (2003)
4. Gendreau, M., Potvin, J.-Y. (eds.): *Handbook of Metaheuristics*, 2nd edn. Springer, New York, NY, USA (2013)
5. Hansen, N., Auger, A., Ros, R., Mersmann, O., Tušar, T., Brockhoff, D.: COCO: A platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* **36**(1), 114–144 (2021)
6. Chizat, L.: Mean-field Langevin dynamics: Exponential convergence and annealing. *Transactions on Machine Learning Research* (2022)
7. Chiang, T.-S., Hwang, C.-R., Sheu, S.J.: Diffusion for global optimization in $\mathbb{R}^n$. *SIAM J. Control. Optim.* **25**(3), 737–753 (1987)
8. Geman, S., Hwang, C.-R.: Diffusions for global optimization. *SIAM J. Control. Optim.* **24**(5), 1031–1043 (1986)
9. Holley, R., Stroock, D.: Simulated annealing via Sobolev inequalities. *Commun. Math. Phys.* **115**(4), 553–569 (1988)
10. Reeves, C.R.: Genetic algorithms. *Handbook of metaheuristics*, 109–139 (2010)
11. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: *Proceedings of 1995 IEEE International Conference on Neural Networks*, 4, 1942–1948 (1995). <https://doi.org/10.1109/ICNN.1995.488968>. IEEE
12. Borghi, G., Pareschi, L.: Kinetic description and convergence analysis of genetic algorithms for global optimization. [arXiv:2310.08562](https://arxiv.org/abs/2310.08562) (2023)
13. Huang, H., Qiu, J., Riedl, K.: On the global convergence of particle swarm optimization methods. *Applied Mathematics & Optimization* **88**(2), 30 (2023)
14. Grassi, S., Pareschi, L.: From particle swarm optimization to consensus based optimization: Stochastic modeling and mean-field limit. *Math. Models Methods Appl. Sci.* **31**(08), 1625–1657 (2021). <https://doi.org/10.1142/S0218202521500342>
15. Carrillo, J.A., Choi, Y.-P., Totzeck, C., Tse, O.: An analytical framework for consensus-based global optimization method. *Math. Models Methods Appl. Sci.* **28**(06), 1037–1066 (2018). <https://doi.org/10.1142/s0218202518500276>
16. Pinnau, R., Totzeck, C., Tse, O., Martin, S.: A consensus-based model for global optimization and its mean-field limit. *Math. Models Methods Appl. Sci.* **27**(01), 183–204 (2017)
17. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**(4598), 671–680 (1983)
18. Kirkpatrick, S.: Optimization by simulated annealing: Quantitative studies. *J. Stat. Phys.* **34**, 975–986 (1984)
19. Zhang, Y., Wang, S., Ji, G.: A comprehensive survey on particle swarm optimization algorithm and its applications. *Math. Probl. Eng.* **2015**(1), 931256 (2015)
20. Carrillo, J.A., Fornasier, M., Rosado, J., Toscani, G.: Asymptotic flocking dynamics for the kinetic Cucker-Smale model. *SIAM J. Math. Anal.* **42**(1), 218–236 (2010)
21. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation* **5**(3) (2002)
22. Miller, P.D.: *Applied Asymptotic Analysis* vol. 75. American Mathematical Soc. (2006). <https://doi.org/10.1090/gsm/075>
23. Dembo, A., Zeitouni, O.: *Large Deviations Techniques and Applications*. Springer (2010). <https://doi.org/10.1007/978-3-642-03311-7>
24. Benfenati, A., Borghi, G., Pareschi, L.: Binary interaction methods for high dimensional global optimization and machine learning. *Applied Mathematics & Optimization* **86**(1), 9 (2022). <https://doi.org/10.1007/s00245-022-09836-5>
25. Carrillo, J.A., Trillos, N.G., Li, S., Zhu, Y.: FedCBO: Reaching group consensus in clustered federated learning through consensus-based optimization. *Journal of Machine Learning Research (JMLR)* **25**(214), 1–51 (2024)

26. Riedl, K., Klock, T., Geldhauser, C., Fornasier, M.: Gradient is all you need? how consensus-based optimization can be interpreted as a stochastic relaxation of gradient descent. [arXiv:2306.09778](https://arxiv.org/abs/2306.09778) (2023)
27. Cipriani, C., Huang, H., Qiu, J.: Zero-inertia limit: from particle swarm optimization to consensus-based optimization. *SIAM J. Math. Anal.* **54**(3), 3091–3121 (2022)
28. Fornasier, M., Klock, T., Riedl, K.: Consensus-based optimization methods converge globally. *SIAM J. Optim.* **34**(3), 2973–3004 (2024). <https://doi.org/10.1137/22M1527805>
29. Fornasier, M., Huang, H., Klemenc, J., Malaspina, G.: From consensus-based optimization to evolution strategies:1 proof of global convergence2. in final preparation (2025)
30. Roith, T., Bungert, L., Wacker, P.: Consensus-based optimization for closed-box adversarial attacks and a connection to evolution strategies. [arXiv:2506.24048](https://arxiv.org/abs/2506.24048) (2025)
31. Carrillo, J.A., Jin, S., Li, L., Zhu, Y.: A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM: Control, Optimisation and Calculus of Variations* **27**, 5 (2021)
32. Ha, S.-Y., Jin, S., Kim, D.: Convergence and error estimates for time-discrete consensus-based optimization algorithms. *Numer. Math.* **147**, 255–282 (2021)
33. Ko, D., Ha, S.-Y., Jin, S., Kim, D.: Convergence analysis of the discrete consensus-based optimization algorithm with random batch interactions and heterogeneous noises. *Math. Models Methods Appl. Sci.* **32**(06), 1071–1107 (2022)
34. Huang, H., Kouhkouh, H., Sun, L.: Faithful global convergence for the rescaled consensus-based optimization. *arXiv preprint* [arXiv:2503.08578](https://arxiv.org/abs/2503.08578) (2025)
35. Huang, H., Qiu, J.: On the mean-field limit for the consensus-based optimization. *Math. Meth. Appl. Sci.* **45**(12), 7814–7831 (2022)
36. Gerber, N.J., Hoffmann, F., Vaes, U.: Mean-field limits for consensus-based optimization and sampling. *ESAIM: Control, Optimisation and Calculus of Variations* **31**, 74 (2025)
37. Huang, H., Kouhkouh, H.: Uniform-in-time mean-field limit estimate for the consensus-based optimization. *ESAIM: Control, Optimisation and Calculus of Variations* **31**, 69 (2025)
38. Platen, E.: An introduction to numerical methods for stochastic differential equations. *Acta Numer* **8**, 197–246 (1999)
39. Totzeck, C.: Trends in consensus-based optimization. In: *Active Particles, Volume 3: Advances in Theory, Models, and Applications*, pp. 201–226. Springer (2021)
40. Fornasier, M., Sun, L.: A pde framework of consensus-based optimization for objectives with multiple global minimizers. *Comm. Partial Differ. Equ.* **50**(4), 493–541 (2025)
41. Bae, H.-O., Ha, S.-Y., Kang, M., Lim, H., Min, C., Yoo, J.: A constrained consensus based optimization algorithm and its application to finance. *Appl. Math. Comput.* **416**, 126726–10 (2022). <https://doi.org/10.1016/j.amc.2021.126726>
42. Borghi, G., Herty, M., Pareschi, L.: Constrained consensus-based optimization. *SIAM J. Optim.* **33**(1), 211–236 (2023). <https://doi.org/10.1137/22M1471304>
43. Carrillo, J.A., Totzeck, C., Vaes, U.: Consensus-based optimization and ensemble Kalman inversion for global optimization problems with constraints. In: *Modeling and Simulation for Collective Dynamics*, pp. 195–230. World Scientific (2023)
44. Carrillo, J.A., Jin, S., Zhang, H., Zhu, Y.: An interacting particle consensus method for constrained global optimization. *arXiv preprint* [arXiv:2405.00891](https://arxiv.org/abs/2405.00891) (2024)
45. Fornasier, M., Huang, H., Pareschi, L., Sünnen, P.: Consensus-based optimization on hypersurfaces: Well-posedness and mean-field limit. *Math. Models Methods Appl. Sci.* **30**(14), 2725–2751 (2020). <https://doi.org/10.1142/S0218202520500530>
46. Fornasier, M., Huang, H., Pareschi, L., Sünnen, P.: Consensus-based optimization on the sphere: convergence to global minimizers and machine learning. *J. Mach. Learn. Res. (JMLR)* **22**, 237–55 (2021)
47. Fornasier, M., Huang, H., Pareschi, L., Sünnen, P.: Anisotropic diffusion in consensus-based optimization on the sphere. *SIAM J. Optim.* **32**(3), 1984–2012 (2022). <https://doi.org/10.1137/21M140941X>
48. Fornasier, M., Huang, H., Klemenc, J., Malaspina, G.: From consensus-based optimization to evolution strategies: Proof of global convergence. *arXiv preprint* [arXiv:2602.11677](https://arxiv.org/abs/2602.11677) (2026)

49. Bungert, L., Roith, T., Wacker, P.: Polarized consensus-based dynamics for optimization and sampling. *Mathematical Programming*, 1–31 (2024)
50. Momin Jamil, X.-S.Y.: A literature survey of benchmark functions for global optimization problems. *Int. Journal of Mathematical Modelling and Numerical Optimisation* 4(2), 150–194 (2013)
51. Sun, X., Jordana, A., Fornasier, M., Etesami, J., Khadiv, M.: Consensus-based optimization (cbo): Towards global optimality in robotics. [arXiv:2602.06868](https://arxiv.org/abs/2602.06868) (2026)
52. Borghi, G., Grassi, S., Pareschi, L.: Consensus based optimization with memory effects: Random selection and applications. *Chaos, Solitons & Fractals* 174, 113859 (2023) <https://doi.org/10.1016/j.chaos.2023.113859>
53. Totzeck, C., Wolfram, M.-T.: Consensus-based global optimization with personal best. *Mathematical Biosciences and Engineering*. MBE 17(5), 6026–6044 (2020). <https://doi.org/10.3934/mbe.2020320>
54. Tanaka, H.: Stochastic differential equations with reflecting boundary condition in convex regions. *Stochastic Processes: Selected Papers of Hiroshi Tanaka* 9, 157 (1979)
55. Skorokhod, A.V.: Stochastic equations for diffusion processes in a bounded region. I–II. *Theory of Probability & Its Applications* 7(1), 3–23 (1962)
56. Lions, P.-L., Sznitman, A.-S.: Stochastic differential equations with reflecting boundary conditions. *Commun. Pure Appl. Math.* 37(4), 511–537 (1984)
57. Saisho, Y.: Stochastic differential equations for multi-dimensional domain with reflecting boundary. *Probab. Theory Relat. Fields* 74(3), 455–477 (1987)
58. Pilipenko, A.: *An Introduction to Stochastic Differential Equations with Reflection* vol. 1. Universitaetsverlag Potsdam (2014)
59. Hinds, P.D., Sharma, A., Tretyakov, M.V.: Well-posedness and approximation of reflected McKean–Vlasov SDEs with applications. *Math. Models Methods Appl. Sci.* 35(08), 1845–1887 (2025). <https://doi.org/10.1142/S0218202525500241>
60. Khasminskii, R.: *Stochastic Stability of Differential Equations*. Springer (2012)
61. Durrett, R.: *Stochastic Calculus: a Practical Introduction*. CRC press (2018)
62. Słomiński, L.: Euler’s approximations of solutions of sdes with reflecting boundary. *Stochastic processes and their applications* 94(2), 317–337 (2001)
63. Bailo, R., Barbaro, A., Gomes, S.N., Riedl, K., Roith, T., Totzeck, C., Vaes, U.: CBX: Python and Julia Packages for Consensus-Based Interacting Particle Methods. *Journal of Open Source Software* 9(98), 6611 (2024) <https://doi.org/10.21105/joss.06611>
64. Conn, A.R., Gould, N.I.M., Toint, P.L.: *Trust-Region Methods*. MPS/SIAM Series on Optimization, p. 959. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA (2000). Volume 1
65. Kalise, D., Sharma, A., Tretyakov, M.V.: Consensus-based optimization via jump-diffusion stochastic differential equations. *Math. Models Methods Appl. Sci.* 33(02), 289–339 (2023)
66. Fornasier, M., Klock, T., Riedl, K.: Convergence of anisotropic consensus-based optimization in mean-field law. In: *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*, pp. 738–754 (2022). Springer
67. Attouch, H., Chbani, Z., Fadili, J., Riahi, H.: First-order optimization algorithms via inertial systems with hessian driven damping. *Math. Program.* 193(1), 113–155 (2022). <https://doi.org/10.1007/s10107-020-01591-1>
68. Miranda, L.J.: Pyswarms: a research toolkit for particle swarm optimization in python. *Journal of Open Source Software* 3(21), 433 (2018)
69. Burger, M., Erbar, M., Hoffmann, F., Matthes, D., Schlichting, A.: Covariance-modulated optimal transport and gradient flows. *Arch. Ration. Mech. Anal.* 249(1), 7 (2025)
70. Feng, X., Prohl, A.: Numerical analysis of the Allen–Cahn equation and approximation for mean curvature flows. *Numer. Math.* 94, 33–65 (2003)
71. Lee, D., Huh, J.-Y., Jeong, D., Shin, J., Yun, A., Kim, J.: Physical, mathematical, and numerical derivations of the cahn–hilliard equation. *Comput. Mater. Sci.* 81, 216–225 (2014)

72. Hübner, S., Wohlmuth, B.: A primal–dual active set strategy for non-linear multibody contact problems. *Comput. Methods Appl. Mech. Eng.* **194**(27), 3147–3166 (2005). <https://doi.org/10.1016/j.cma.2004.08.006>
73. Gräser, C., Kornhuber, R.: Multigrid methods for obstacle problems. *Journal of Computational Mathematics*, 1–44 (2009)
74. Wohlmuth, B.: Variationally consistent discretization schemes and numerical algorithms for contact problems. *Acta Numer* **20**, 569–734 (2011). <https://doi.org/10.1017/S0962492911000079>
75. Auger, A., Hansen, N.: The theory of evolution strategies: A tutorial. In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) Tutorial Track*, Washington, DC, USA (2005). ACM Press
76. Rudolph, G.: *Convergence Analysis of Stochastic Global Optimization Algorithms*. Springer, Berlin / Heidelberg / New York (1997)
77. Rudolph, G.: Global convergence of evolutionary algorithms for continuous search spaces. *Stochastic Processes and their Applications* **76**(2), 171–179 (1998). [https://doi.org/10.1016/S0304-4149\(98\)00047-X](https://doi.org/10.1016/S0304-4149(98)00047-X)
78. De Giorgi, E.: New problems on minimizing movements. In: *Boundary Value Problems for Partial Differential Equations and Applications*. RMA Res. Notes Appl. Math., vol. 29, pp. 81–98. Masson, Paris (1993)
79. Borghi, G., Pareschi, L.: Kinetic description and convergence analysis of genetic algorithms for global optimization. *Commun. Math. Sci.* **23**(3), 641–668 (2025). <https://doi.org/10.4310/cms.250208214404>
80. Borghi, G., Herty, M., Pareschi, L.: Kinetic models for optimization: a unified mathematical framework for metaheuristics. [arxiv:2410.10369](https://arxiv.org/abs/2410.10369) (2024)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.