Mean-field theory for consensus-based optimization and extensions to constrained and multi-objective problems

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH Aachen University zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Giacomo Borghi, M. Sc.

aus Correggio, Italy

Berichter: Prof. Dr. Michael Herty
Prof. Dr. Lorenzo Pareschi

Tag der mündlichen Prüfung: 01.02.2024

Diese Dissertation ist auf den Internetseiten der Universitätsbibliothek verfügbar.

RWTH AACHEN UNIVERSITY

UNIVERSITÀ DI PARMA

Acknowledgements

I would like to thank my supervisors, Prof. Michael Herty and Prof. Lorenzo Pareschi, for supporting and advising me throughout my doctoral studies and in writing this thesis. This work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Projektnummer 320021702/GRK2326 Energy, Entropy, and Dissipative Dynamics (EDDy).

The thesis has been written whithin an agreement of joint thesis supervision between RWTH Aachen University and Università degli Studi di Parma.

Abstract

Stochastic particle methods in optimization constitute a popular class of heuristic techniques where a set of possible solutions is iteratively updated according to deterministic and stochastic mechanisms. These algorithms are oftentimes inspired by natural phenomena such as the collective motion of birds in a swarm or the reproduction of genes in biology. Recent works apply kinetic tools developed for modeling such phenomena to the analysis of particle-based optimization methods to develop a full mathematical understanding of their convergence properties.

In this thesis, we follow this line of research and propose a semi-discrete mean-field model for the analysis of Consensus-Based Optimization (CBO) methods. In the literature, CBO methods are typically analyzed after a twofold approximation: first, the algorithm dynamics are approximated by a system of time-continuous processes, then, the system is approximated by a mono-particle process of McKean type. Here, we directly consider a mean-field approximation of the algorithm's update and derive a mono-particle difference equation. We also adapt the convergence results for the time-continuous mean-field model to these semi-discrete settings, claiming that this modeling procedure avoids the introduction of an unnecessary additional approximation error.

The second part of the thesis extends the class of CBO methods to solve constrained and multi-objective optimization problems. For constrained optimization, we couple the CBO update rule with an exact penalization technique. By adding a penalty term, we reformulate the constrained problem as an unconstrained one and further propose an algorithmic technique to tune, during the computation, the penalization strength to its optimal value. For multi-objective problems, we modify the CBO dynamics so that every particle aims to optimize a different, parameterized scalar sub-problem. In this way, we are able to compute an approximation of the entire Pareto front with a single run of the algorithm. The parameters of the scalarized sub-problems are further adapted following binary repulsive dynamics to improve the diversity of the computed front.

Via mean-field approximation, we show that the proposed algorithms are able to converge into a neighborhood of the solutions under mild assumptions, both in the case of constrained and multi-objective optimization problems. Numerical experiments over benchmark problems illustrate the algorithms' performance in different settings and for different problem types.

Zusammenfassung

In der Optimierung stellen stochastische Partikelmethoden eine beliebte Klasse heuristischer Techniken dar, bei denen eine Menge möglicher Lösungen iterativ nach deterministischen und stochastischen Mechanismen aktualisiert wird. Diese Algorithmen sind oft von natürlichen Phänomenen inspiriert, wie der kollektiven Bewegung von Vögeln in einem Schwarm oder der Reproduktion von Genen in der Biologie. Aktuelle Arbeiten wenden Werkzeuge aus der Kinetik, die für die Modellierung solcher Phänomene entwickelt wurden, auf die Analyse partikelbasierter Optimierungsverfahren an, um ein umfassendes mathematisches Verständnis ihrer Konvergenzeigenschaften zu erhalten.

In dieser Arbeit folgen wir dieser Forschungsrichtung und schlagen ein semidiskretes Mean-Field-Modell für die Analyse von konsensbasierten Optimierungsverfahren (CBO) vor. In der Literatur werden CBO-Methoden typischerweise nach einer zweifachen Approximation analysiert: Zunächst wird die Dynamik des Algorithmus durch ein System zeitkontinuierlicher Prozesse approximiert, dann wird das System durch einen Ein-Teilchen-Prozess vom McKean-Typ approximiert. Hier betrachten wir direkt eine Mean-Field-Approximation der Aktualisierung des Algorithmus und leiten eine Mono-Teilchen-Differenzgleichung her. Weiterhin passen wir auch die Konvergenzergebnisse für das zeitkontinuierliche Mean-Field-Modell an diese semidiskreten Einstellungen an und behaupten, dass dieses Modellierungsverfahren die Einführung eines unnötigen zusätzlichen Approximationsfehlers vermeidet.

Der zweite Teil der Arbeit erweitert die Klasse der CBO-Methoden, um eingeschränkte sowie Mehrzieloptimierungsprobleme zu lösen. Für die eingeschränkte Optimierung koppeln wir die CBO-Aktualisierungsregel mit einer exakten Bestrafungsregel. Durch Hinzufügen eines Bestrafungsterms formulieren wir das eingeschränkte Problem als ein nicht eingeschränktes Problem um und stellen eine algorithmische Technik vor, die die Stärke der Bestrafung während der Berechnung auf ihren optimalen Wert abstimmt. Bei Mehrzielproblemen modifizieren wir die CBO-Dynamik so, dass jeder Partikel ein anderes, parametrisiertes skalares Teilproblem optimiert. Auf diese Weise sind wir in der Lage, eine Approximation der gesamtem Pareto-Front mit einem einzigen Durchlauf des Algorithmus zu berechnen. Die Parameter der skalaren Teilprobleme werden nach der binären repulsiven Dynamik weiter angepasst, um die Diversität der berechneten Front zu verbessern.

Mittels Mean-Field-Approximation zeigen wir, dass die vorgeschlagenen Algorithmen in der Lage sind, unter milden Annahmen in eine Nachbarschaft der Lösungen zu konvergieren, sowohl bei eingeschränkten als auch bei Mehrzieloptimierungsproblemen. Numerische Experimente von Benchmark-Problemen zeigen die Leistung der Algorithmen in verschiedenen Situationen und für verschiedene Problemtypen.

Sommario

I metodi particellari stocastici nell'ottimizzazione costituiscono una popolare classe di tecniche euristiche in cui un insieme di possibili soluzioni viene aggiornato iterativamente secondo meccanismi deterministici e stocastici. Questi algoritmi sono spesso ispirati a fenomeni naturali come il moto collettivo degli uccelli in uno stormo o la riproduzione dei geni in biologia. Lavori di ricerca recenti applicano strumenti cinetici sviluppati per la modellazione di tali fenomeni all'analisi dei metodi di ottimizzazione basati sulle particelle per sviluppare una piena comprensione matematica delle loro proprietà di convergenza.

In questa tesi, seguiamo questa linea di ricerca e proponiamo un modello di campo medio semi-discreto per l'analisi dei metodi di ottimizzazione basati sul consenso (CBO). In letteratura, i metodi CBO sono tipicamente analizzati dopo una duplice approssimazione: dapprima, la dinamica dell'algoritmo è approssimata da un sistema di processi tempo-continui, quindi, il sistema è approssimato da un processo mono-particellare di tipo McKean. In questa tesi, consideriamo direttamente un'approssimazione di campo medio dell'iterazione dell'algoritmo e deriviamo un'equazione alle differenze mono-particellare. Adattiamo inoltre i risultati di convergenza per il modello di campo medio continuo nel tempo a queste impostazioni semidiscrete, notando come questa procedura di modellazione eviti l'introduzione di un inutile errore di approssimazione aggiuntivo.

La seconda parte della tesi estende la classe dei metodi CBO per risolvere problemi di ottimizzazione vincolata e multi-obiettivo. Per l'ottimizzazione vincolata, accoppiamo la regola di aggiornamento CBO con una tecnica di penalizzazione esatta. Aggiungendo un termine di penalizzazione, riformuliamo il problema vincolato come un problema non vincolato e proponiamo una tecnica algoritmica per regolare, durante la computazione, la forza della penalizzazione al suo valore ottimale. Per i problemi multi-obiettivo, modifichiamo la dinamica del CBO in modo che ogni particella miri ad ottimizzare un diverso problema scalare parametrizzato. In questo modo, siamo in grado di calcolare un'approssimazione dell'intero fronte di Pareto con una singola esecuzione dell'algoritmo. I parametri dei problemi scalari vengono ulteriormente adattati attraverso una dinamica repulsiva binaria per migliorare la diversità del fronte calcolato.

Tramite un'approssimazione di campo medio, dimostriamo che gli algoritmi proposti sono in grado di convergere in un intorno delle soluzioni in presenza di lievi ipotesi, sia nel caso di problemi di ottimizzazione vincolati che multi-obiettivo. Esperimenti numerici su test di riferimento illustrano le prestazioni degli algoritmi in diversi contesti e per diversi tipi di problemi.

Contents

Chapter 1

Introduction

Optimization problems can be found everywhere. What is the shortest route to go home? How can I schedule my meetings in the most efficient way? Already in such common settings, we often rely on mobile applications to get answers to our optimization problems. What mobile applications do on our phones is simply run an optimization algorithm that returns the optimal solution possible. For the route planning problem, we insert a destination and the algorithm computes the best itinerary for us. In engineering and applied sciences, the complexity of optimization problems can increase exponentially. What is the optimal balance between ingredients in a chemical solution? What is the car design which minimizes air friction? What is the best Neural Network architecture for a given Natural Language Processing task? To solve these kinds of optimization problems having efficient and reliable algorithms is of paramount importance. In this thesis, we study and extend a class of optimization algorithms, the Consensus-Based Optimization methods, which have proven, in recent years, to be both effective and theoretically sound.

Historically, optimization algorithms are divided into exact methods and heuristics. Most exact methods are designed to provably converge to local solutions to the problems, eventually taking an infinite computational time to do so. Heuristics are usually designed with a specific application in mind and aim to find the best solution possible, given a certain computational budget. Exact methods are theoretically analyzed and are popular among the mathematical community, while heuristics are more popular among engineers and practitioners. Quoting [76], "*[...] these two communities are almost completely disjoint: they have different journals, different conferences, and different test functions.*"

One of the features which makes heuristics popular in the applied community is that their core mechanism has oftentimes an intuitive, nature-inspired interpretation. To name a few, the Particle Swarm Optimization method [64] mimics the behavior of a swarm of birds; the Genetic Algorithm [79] is inspired by Darwin's principle of natural selection; Simulated Annealing [65] resembles the annealing technique in metallurgy. Names aside, the main reason why the mathematical community has traditionally had no interest in

studying heuristics is that, to put it simply, they are too difficult to analyze. Being the result of several trials and errors, a heuristic algorithm is a stratified collection of mechanisms involving stochastic choices and complex recombination of possible solutions. The usual mathematical programming techniques developed for exact methods are not suitable for the analysis of such algorithms. This is where another branch of mathematics, namely mathematical physics, enters the story.

Statistical mechanics was first developed by L. Boltzmann (1844–1906) to study the macroscopic properties of rarefied gases. Even in a rarefied state, a gas contains an enormous number of particles, colliding with each other at every instant. One cannot study how the gas behaves, let us say, how its pressure evolves, by tracking the trajectory of every single particle: this would be computationally prohibitive. To deal with this issue, Boltzmann suggested describing the gas in probabilistic terms. We should not try to describe the state of every single particle but rather model how they are *likely* to behave, that is, we should look at probability distributions rather than particle systems. At the cost of sacrificing a certain amount of precision, we may obtain a useful mathematical model that we can employ to predict how the gas pressure or its temperature evolves, for instance.

We can find the same situation in many other contexts. In biology, we know how cells interact with each other at the microscopic level, and we would like to know how a certain tissue grows. In economics, from a model of the traders' behavior, we would like to infer the long-time evolution of stock prices. In social sciences, can we guess the election output if we know how people exchange opinions with one another? In all these cases, the aim is to analytically study the collective behavior of a system from the laws that regulate the interaction between the individuals (to be intended in a broad sense). There is a large body of literature in which Boltzmann's idea of switching from a discrete, individual-based model to a continuum mathematical model dealing with probability distributions has been successfully applied in these contexts, see, for instance, [24, 28, 33, 70, 71, 80] and the references therein.

Some heuristic algorithms can also be described as individual-based models where the individuals are given possible solutions to the optimization problems. It should come as no surprise that this is the case with the previously mentioned Particle Swarm Optimization method and Genetic Algorithm. The heuristic's mechanism which combines and modifies the solutions plays here the role of the interaction mechanisms in physics, biology, or economics. The only difference is that it is the algorithm's designer who makes the rules, and not nature or society. The most common heuristics optimization paradigms were developed in the '80s and '90s [64, 65, 79]. Thanks to they broad popularity and flexibility, these methods are the subject of an impressive number of publications in which the algorithms are constantly improved, see for instance [60] for a literature review on the time span 1983–2016. Up until recent years, though, parallelism between heuristics and individual-based models has been long overseen.

In [72], authors proposed the Consensus-Based Optimization (CBO) heuristics where

a set of possible solutions interact through dynamics inspired by the ones used to study opinion formation among social individuals. CBO shares many features with Particle Swarm Optimization, but it was designed specifically to be amenable to theoretical analysis via tools from mathematical physics. Indeed, the algorithm's mechanism can be studied in probabilistic terms to theoretically analyze its collective behavior.

From [72], many authors are now developing statistical mechanics of heuristics which aim to analytically study the convergence properties of popular heuristics with the new mathematical tools developed in the analysis of collective dynamics [2, 5, 16, 19, 41, 50, 54, 59]. Can we still consider it to be non-exact a heuristic that is mathematically proven to converge to a solution with a high probability? Is it still a heuristic if the algorithm's dynamics were the result of a theoretical investigation, rather than several trials and errors against benchmark problems? This body of work thinned the gap between heuristic algorithms and exact methods and calls for a new type of classification in the optimization literature.

1.1 Contribution

This thesis deals with CBO methods and their theoretical analysis. The contribution is twofold: we propose a novel mean-field mathematical analysis that does not rely on continuous-in-time approximations, and we extend the class of CBO methods to solve constrained and multi-objective optimization problems.

Mean-field analysis. One of the possible ways of deriving a statistical description starting from an individual-based model, is to take the so-called *mean-field limit*. Let us briefly give an intuition on how to perform such an approximation. Given our set of N individuals, which we call *particles* from now on, we focus on a particular one (anyone) and study how it interacts with the remaining $N-1$. Assuming N large, we approximate the rest of the ensemble with a probability distribution. In this way, we now need to model only the interaction between a single particle and a probability measure describing the average behavior of the rest. Due to the arbitrariness of the initial choice of particle, it turns out that such a measure describes the probability distribution of our particle as well. We are left, therefore, with a mono-particle process that aims to approximate a N particle system. It is called *mean-field limit* as its accuracy typically increases as the number of particles N increases, and it becomes exact as $N \to \infty$. Going from a model of N equations to a model of one single equation greatly simplifies the theoretical analysis of the system.

The starting point for the derivation of mean-field limits for swarming models, biological models, or opinion models, is, most of the time, a set of N, time-continuous, ordinary (or stochastic) differential equations. This is because, in these contexts, events are considered to happen seamlessly, at every instant of time. On the contrary, optimization algorithms, and CBO methods in particular, are intrinsically discrete. The particles are updated at each iteration and there is no concept of time continuity in the physics of

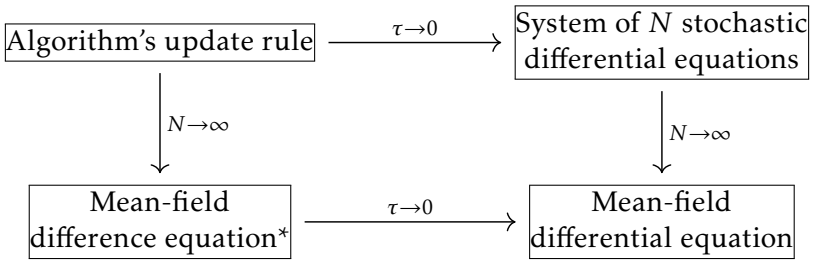

Figure 1.1: In the diagram, τ is step-size of the particles' update in the CBO algorithm, while N is the number of particles. In this thesis, we study the convergence of CBO methods via the mathematical model (*).

the algorithm. Yet, the most frequent approach in the analysis of CBO methods is to first perform a continuous-in-time approximation of the algorithm's update rule, see, for instance, [16, 41, 72]. The consensus dynamics is, so, approximated by a system N Stochastic Differential Equations (SDEs). From this point onwards, one can apply the classical tools in mean-field theory to derive the mono-particle process, which can either be described by a SDE of McKean type, or a Partial Differential Equation (PDE). After a twofold approximation, the mathematical model of the heuristic can be theoretically analyzed.

In this thesis, we suggest a different approach that requires to perform a single approximation. We achieve that by directly taking the mean-field limit $N \to \infty$, without relying on a time-continuous approximation of the dynamics, see Figure 1.1. We argue that this is sufficient to prove the convergence of the CBO dynamics towards global minima, thanks to the simplicity of the consensus mechanism. Moreover, it avoids introducing an unnecessary approximation error in the theoretical investigation. Our analysis is built upon the recent results on the convergence of CBO methods [41, 42], but takes a novel point of view. The convergence results we are able to prove for the classical CBO method, we also prove for a variant that includes memory effects, and for a projected version of the CBO dynamics. For the former, we also provide quantitative error estimates of the mean-field approximation, in the case of convex, bounded domains.

The second part of the thesis concerns the design of novel heuristic algorithms.

Constrained optimization problems. We start by adapting CBO to solve constrained optimization problems. In such problems, a solution does not only have to minimize an objective but should also satisfy certain constraints to be admissible. To include constraints, we add an exact penalization strategy to CBO. Penalization techniques aim to rule out non-admissible solutions by adding a penalization term to those that do not satisfy the constraints. The strength of the penalization is determined by the so-called penalty parameter. For a penalization to be exact, it means it ensures correspondence

between the penalized problem and the original constrained one, provided the penalty parameter is sufficiently large. Exact penalization typically makes the objective function non-smooth and, therefore, it is non-suitable for gradient-based methods. As CBO methods are gradient-free, this does not represent an issue.

In the proposed algorithm, we couple the CBO dynamics with a dynamical update of the penalty parameter. Whenever the constraint violation of the particles is larger than a certain tolerance, the penalty parameter is increased. If the violation remains within the tolerance, instead, we decrease the tolerance for the next iteration. We perform a mean-field analysis of the novel method and derive conditions under which the particle system is expected to converge toward an admissible minimizer of the objective problem. Several numerical experiments are performed to assess the performance of the algorithm against benchmark problems. We also compare the experimental results with the theoretical analysis and investigate the role of the different parameters involved in the method. This application shows how the original CBO dynamics can be easily enriched with further mechanisms, without affecting its convergence properties and theoretical understanding.

Part of this work on CBO for constrained optimization has already been published in [7][1]. We note that the algorithm analysis in the semi-discrete settings is novel.

Multi-objective optimization problems. The second algorithm we propose is designed to solve multi-objective optimization problems, but it can also be applied to any task where several problems need to be solved in parallel.

In multi-objective optimization, two or more objectives need to be taken into account at the same time. Since objectives are typically in conflict with each other, we expect to be no solution that minimizes all of them. Therefore, the classical notion of optimality is relaxed to the notion of Edgeworth-Pareto (EP) optimality. To find the possibly many EP optimal points, we employ a scalarization strategy that allows us to translate the multi-objective problem into several parameterized single-objective sub-problems. The CBO algorithm we propose distributes the optimization tasks among the particles so that every particle is in charge of solving a different sub-problem. During the computation, the multi-objective values of the particles are shared among the ensemble so that we can compute a different consensus point for each problem.

As in the case of exact penalization, the choice of the sub-problems' parameters is essential to guarantee a good quality of the solution. In multi-objective optimization, one should find EP solutions that are as *diverse* as possible to one another, meaning that their images should not be close from one another in the objective space. To this end, we couple the CBO dynamics with a short-range repulsive evolution in the space of the parameters. Through several numerical experiments with bi- and tri-objective benchmark problems,

[1] **CRediT authorship contribution statement for [10–12]:** Giacomo Borghi: Conceptualization, Software, Validation, Formal analysis, Data curation, Writing – original draft, Visualization. Michael Herty & Lorenzo Pareschi: Conceptualization, Methodology, Formal analysis, Writing - Review & Editing, Supervision, Project administration, Funding acquisition.

we show the validity of the proposed algorithm. We also derive a mean-field description of the coupled dynamics and theoretically prove convergence towards EP optimal points under mild assumptions on the parameterized sub-problems.

This chapter on CBO methods for multi-objective optimization is built upon the published works [7,7,7][1]. As for the penalized CBO algorithm, we remark the novelty of the converge analysis in the semi-discrete mean-field settings.

1.2 Outline

The thesis is organized as follows. Chapter 2 is devoted to the presentation of the CBO algorithm for single-objective, unconstrained problems and the derivation of the corresponding mean-field difference equation. After presenting the original CBO method, we show how it can be derived as a simplified version of the popular PSO algorithm. Then, we consider the CBO with anisotropic diffusion and derive its mean-field approximation via the propagation of chaos assumption. We also present and discuss a CBO variant where the personal best of each particle is included in the dynamics, as in PSO methods.

Chapter 3 is devoted to the theoretical analysis of the mean-field model, both in the case of dynamics with and without memory effects. Under suitable growth conditions of the objective function and choice of the algorithm parameters, we show that the mean-field process concentrates in an arbitrary small neighborhood of the minimizer. We also analyze a further variant where particles are constrained, via projection, in a convex and bounded domain. For the projected dynamics, we not only show convergence in mean-field law but also provide a quantitative error of the mean-field approximation in terms of number of particles.

In Chapter 4, we show how to adapt the CBO method to solve constrained optimization problems via exact penalization. After presenting the algorithm and, in particular, the adaptive strategy for the penalty parameters, we analyze the corresponding mean-field description and its convergence properties. Then, we extensively test the method on benchmark problems to validate the theoretical analysis and further understand the role of the different algorithm's parameters.

Chapter 5 is dedicated to multi-objective optimization. We propose here a CBO method that makes use of the scalarization technique and that aims to minimize several parameterized sub-problems simultaneously. We start by presenting the CBO dynamics in case of fixed parameters and then propose different adaptive strategies for bi-objective problems and general multi-objective ones. After deriving their mean-field approximations, we analytically study their convergence properties towards the Pareto front. Several numerical experiments are performed for both bi- and tri-obective problems to validate the CBO dynamics and the different adaptive strategies.

Final remarks and possible future research directions are discussed in Chapter 6.

Chapter 2

Consensus-Based Optimization methods and mean-field theory

Consensus-Based Optimization (CBO) algorithms make use of a set of possible solutions, which we will call *particles*, to explore the search space and find a global minimum of the objective function. The update rule is the result of an interaction of consensus type between the particles where convergence towards a consensus point is coupled with random steps favoring exploration of the search space.

The collective behavior of interacting particle systems has been subject of extensive mathematical investigation in last decades, see, for instance, [52,71] and references therein. Following this large body of literature, convergence analysis of CBO algorithms typically makes use of a twofold approximation of the dynamics [16, 41, 72]. First, the particles' update rule is approximated by a continuous-in-time dynamics. Then, the particle system is approximated by a mono-particle process of mean-field type which greatly simplifies the convergence analysis [16, 41]. Quantitative estimates of the two approximations allow eventually to recover error estimates for the original CBO algorithm [41].

In this chapter, we propose a different strategy to study the collective behavior of the CBO particle system. In particular, we directly perform a mean-field approximation of the CBO algorithm. We simplify in this way the modeling procedure and avoid the introduction of a second approximation error (the one which depends on the step size). This also simplifies the well-posedness analysis of the mean-field approximation.

Before delving into the mathematical modeling, we introduce in Section 2.1 the CBO update rule. Following its genesis [72], we present the CBO interaction as a simplified version of the popular Particle Swarm Optimization (PSO) heuristic. We recall from [43] the main steps on how to derive CBO dynamics from PSO ones. Next, we show how to leverage such simplified dynamics to provide a statistical description of the particle evolution. We will first present in Section 2.2 the typical twofold approximation used in the CBO literature. In Section 2.3 we then present the novel mean-field approximation of the

CBO dynamics. We also illustrate in Section 2.4 how to re-introduce the typical memory effects of PSO methods, while keeping the particle system amenable to mathematical description. Additional remarks and details regarding the literature are provided at the end of the chapter.

2.1 Derivation of CBO from PSO

Inspired by the intelligent behavior of flocks of birds, the first PSO algorithm was proposed in [64], and has now been improved in many aspects to serve a wide range of use cases [56]. Following the steps proposed in [43], we introduce in this section the CBO dynamics as a modification of the PSO algorithm.

Let $\mathcal{E} : \mathbb{R}^d \to \mathbb{R}$ be the objective function the swarm aims to optimize, the problem reads

$$\text{minimize } \mathcal{E}(x). \tag{2.1}$$

In the original PSO, $N \in \mathbb{N}$ particles are characterized by their locations $\{X^i\}_{i=1}^N \subset \mathbb{R}^d$ and their velocities $\{V^i\}_{i=1}^N \subset \mathbb{R}^d$. At every iteration k of the algorithm, each particle modifies its velocity to move towards two different points in the search space. The first point is given by the particle own best position found so far up to step k. We indicate the best position of the i-th particle at step k as $Y^i_{(k)}$:

$$Y^i_{(k)} \quad \text{such that} \quad \mathcal{E}(Y^i_{(k)}) = \min_{h,\, 0\le h\le k} \mathcal{E}(X^i_{(h)}).$$

In case the above minimum is not unique, the first one reached during the computation is typically set to be the particle's personal best [56].

The second point considered in the velocity update is the global best location among the all personal bests. Let $f^{N,y}_{(k)} = (1/N)\sum_{i=1}^N \delta_{Y^i_{(k)}}$ be the empirical measure of the personal bests, where with δ_y we denote the Dirac measure at $y \in \mathbb{R}^d$. We will indicate the global best location as $m^*[f^{N,y}_{(k)}]$:

$$m^*[f^{N,y}_{(k)}] \quad \text{such that} \quad \mathcal{E}(m^*[f^{N,y}_{(k)}]) = \min_{j,\, 1\le j\le N} \mathcal{E}(Y^j_{(k)}).$$

To be more precise, for any f belonging to the set $\mathcal{P}(\mathbb{R}^d)$ of Borel probability measure over $\mathbb{R}^d$ we consider a map $m^* : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}^d$ satisfying

$$m^*[f] \in \operatorname*{argmin}_{x\in \mathrm{supp}(f)} \mathcal{E}(x).$$

To include stochasticity in the dynamics, two vectors for each particle are randomly sampled from the uniform distribution over $[0,1]^d$, $\theta^{i,1}_{(k)}, \theta^{i,2}_{(k)} \sim \mathrm{Unif}([0,1]^d)$. Let $c_1, c_2 > 0$

be positive constants and $\odot$ denote the component-wise (Hadamard) product, the PSO algorithm iteration then reads

$$\begin{cases} X^i_{(k+1)} & = X^i_{(k)} + V^i_{(k+1)} \\ V^i_{(k+1)} & = V^i_{(k)} + c_1 \left(Y^i_{(k)} - X^i_{(k)} \right) \odot \theta^{i,1}_{(k)} + c_2 \left(m^*[f^{N,y}_{(k)}] - X^i_{(k)} \right) \odot \theta^{i,2}_{(k)} \end{cases} \tag{2.2}$$

for all $i = 1, \dots, N$.

To be able to analytically analyze the PSO dynamics, authors in [72] suggested to: consider a first-order dynamics, avoid the use of personal bests, and avoid the use of any memory mechanism at all. Also, for modeling purposes, the global best position is regularized by considering a weighted average of the particles locations.

Let $f \in \mathcal{P}(\mathbb{R}^d)$ be a probability measure describing a particle distribution. The corresponding weighted average $m^\alpha[f] \in \mathbb{R}^d$ is computed through the Boltzmann-Gibbs distribution associated with the objective function $\mathcal{E}$ and thermodynamic temperature $1/\alpha$, for $\alpha > 0$:

$$m^\alpha[f] := \frac{\int x \exp(-\alpha \mathcal{E}(x))\, df(x)}{\int \exp(-\alpha \mathcal{E}(x))\, df(x)}. \tag{2.3}$$

One can verify that such weighted average is an approximation of the global best position for empirical particle distributions. Indeed, if the global best position is unique among the personal bests, direct computations leads to

$$m^\alpha[f^{N,y}_{(k)}] \longrightarrow m^*[f^{N,y}_{(k)}], \quad \text{as} \quad \alpha \to \infty.$$

We remark that proving the above limit for generic, non-empirical, probability measures is a non-trivial task which requires additional assumptions on the objective function $\mathcal{E}$, see the quantitative version of the so-called Laplace principle [42, Proposition 1].

Let $f^N_{(k)}$ be the empirical distribution associated with the particle locations $\{X^i_{(k)}\}_{i=1}^N$. After disregarding personal bests and velocities, and by regularizing m^*, update rule (2.2) simplifies to

$$X^i_{(k+1)} = X^i_{(k)} + c_2 \left(m^\alpha[f^N_{(k)}] - X^i_{(k)} \right) \odot \theta^{i,2}_{(k)}$$

which is not yet the CBO dynamics, though. As suggested in [43], the CBO update rule can be obtained after splitting the variation $c_2 \left(m^*[f^N_{(k)}] - X^i_{(k)} \right) \odot \theta^{i,2}_{(k)}$ into a deterministic term and a zero-mean term. Such components are also tuned separately through two different parameters, $\lambda > 0$ and $\sigma > 0$ respectively for greater flexibility. At the cost of making the system over-parameterized, an additional parameter $\tau \in (0, 1)$ regulating the step-size is also introduced.

The obtained CBO update rule finally reads

$$X^i_{(k+1)} = X^i_{(k)} + \lambda \tau \left(m^\alpha[f^N_{(k)}] - X^i_{(k)} \right) + \sigma \sqrt{\tau} \left(m^\alpha[f^N_{(k)}] - X^i_{(k)} \right) \odot \theta^i_{(k)}, \tag{2.4}$$

where $\theta^i_{(k)} \sim \mathcal{N}(0, I_\mathrm{d})$ is now sampled form the standard normal distribution.

We note that, while the CBO dynamics differs significantly from the PSO one (2.2), the main features of PSO, that is,

- flock convergence towards a (regularized) global best position
- exploration via random component

are preserved in (2.4). On the other hand, we will see in the following how this simpler update rule is more amenable to mathematical analysis, thanks to its first-order nature, the lack of personal bests, and, most importantly, the regularization of the global best.

Remark 2.1. *Equation* (2.4) *corresponds to the anisotropic CBO dynamics proposed in [19], where the noise acts differently on each coordinate of the search space. The algorithm iteration was originally proposed with isotropic exploration:*

$$X^i_{(k+1)} = X^i_{(k)} + \lambda\tau\left(m^\alpha[f^N_{(k)}] - X^i_{(k)}\right) + \sigma\sqrt{\tau}\,|m^\alpha[f^N_{(k)}] - X^i_{(k)}|\,\theta^i_{(k)}.$$

While the mean-field modeling procedure proposed in this work can be applied to both iso- and aniso-tropic CBO dynamics, we will focus on the latter as it has been proved to be more robust against large-scale optimization problems ($\mathrm{d} \gg 1$*) [5, 19].*

2.2 Literature review on continuous mean-field approximations

Most works in the CBO literature analyze the numerical particle iteration through two approximations: first, the algorithm is approximated by a continuous-in-time dynamics of N particles, then, a mean-field approximation of the particle interaction is performed [16, 41, 42, 72]. Loosely speaking, this corresponds to performing two different limits, $\tau \to 0$ and $N \to \infty$. For completeness, we recall in this section the main steps of this modeling procedure.

To derive the first approximation, we note that (2.4) corresponds to an Euler-Maruyama discretization [55, 73] of the following system of N Itô's Stochastic Differential Equations (SDEs)

$$dX^i_t = \lambda\left(m^\alpha[f^N_t] - X^i_t\right)dt + \sigma\left(m^\alpha[f^N_t] - X^i_t\right) \odot dB^i_t \tag{2.5}$$

for $i = 1, \dots, N$, where $\{(B^i_t)_{t\geq 0}\}_{i=1}^N$ are independent Wiener processes and $f^N_t = (1/N)\sum_{i=1}^N \delta_{X^i_t}$ is the (random) empirical particle measure. Such approximation introduces an error which is polynomial in the step size τ [73]. Well-possesses of the above system of SDEs is guaranteed by classical arguments in SDE analysis, provided the objective function $\mathcal{E}$ is locally Lipschitz continuous [16].

We note that this first approximation does not address the main mathematical challenge of the CBO particle system, that, is the (possibly) large number N of equations involved.

To do so, authors in [72] propose a mean-field approximation of the particles' interaction by making the so-called *propagation of chaos* assumption on the marginals [78]. Let $F_t^N = \mathrm{Law}(X_t^1, \dots, X_t^N) \in \mathcal{P}((\mathbb{R}^{\mathrm{d}})^N)$ be the probability distribution of the particle system at time $t \geq 0$. We assume that for $N \gg 1$ it factorizes as $F_t^N \approx f_t^{\otimes N}$ for some $f_t \in \mathcal{P}(\mathbb{R}^{\mathrm{d}})$. Under such assumption, particles $X_t^1, \dots, X_t^N$ are identically, independently distributed at all times $t \geq 0$ and evolve according to the McKean mono-particle process

$$d\overline{X}_t = \lambda\Big(m^\alpha[f_t] - \overline{X}_t\Big)dt + \sigma\Big(m^\alpha[f_t] - \overline{X}_t\Big) \odot dB_t \tag{2.6}$$

with $f_t = \mathrm{Law}(\overline{X}_t)$. Thanks to Itô-Doeblin formula, f_t weakly satisfies the Fokker-Planck equation

$$\frac{d}{dt}\int \phi(x) df_t(x) = \lambda \int \nabla\phi(x) \cdot (m^\alpha[f_t] - x) df_t(x) + \frac{\sigma^2}{2}\int \sum_{\ell=1}^{\mathrm{d}} \partial_{\ell\ell}\phi(x)(m^\alpha[f_t] - x)_\ell^2 df_t(x) \tag{2.7}$$

for any $\phi \in C_b^\infty(\mathbb{R}^{\mathrm{d}})$. Well-posedness of the mean-field dynamics (2.6) for given initial data $\overline{X}_0 \sim f_0$ can be proved under mild assumptions on the objective function by means of the Leray-Schauder fixed point theorem [16].

A rigorous analysis of the mean-field limit $N \to \infty$ has been carried out in [57] where authors justify the propagation of chaos assumption. Under suitable initial data, authors prove that the empirical random measures f^N and the solution $f \in \mathcal{P}(C([0,T], \mathbb{R}^{\mathrm{d}}))$ to (2.7) for a time horizon $T > 0$ satisfies

$$f^N \rightharpoonup f \quad \text{in law as} \quad N \to \infty.$$

This is proven under mild assumptions on the objective $\mathcal{E}$, that is, local Lipschitz continuity, quadratic lower bound, and quadratic growth at infinity. Quantitative error estimates for the mean-field approximation are provided in [41] under the same assumptions for $\mathcal{E}$, provided the dynamics remains bounded (and this is shown to happen with high probability). Let $T > 0$ be a given time horizon. By coupling the N particle system (2.5) with N solutions $\{\overline{X}_t^i\}_{i=1}^N$ to the mean-field dynamics (2.6), authors in [41] prove, in particular, that

$$\max_{i=1,\dots,N} \sup_{t\in[0,T]} \mathbb{E}\Big[|X_t^i - \overline{X}_t^i|^2 \,\Big|\, \Omega_M\Big] \leq C_{\mathrm{MFA}} N^{-1},$$

where Ω_M is the set of realizations satisfying $(1/N)\sum_{i=1}^N \max\{|X_t^i|^4, |\overline{X}_t^i|^4\} \leq M$ for all $t \in [0,T]$, and where $C_{\mathrm{MFA}} > 0$ is a constant that depends on T, M.

By studying the mean-field description (2.7) one can investigate the large-time behavior of the particles' system, such as the creation of consensus and the convergence towards a global minimizer. While we will review the different approaches to the analysis of the mean-field model in Chapter 3, we remark now that any convergence analysis of this type

necessarily includes an error of order $\mathcal{O}(\tau + N^{-1})$ introduced by the twofold approximation. In the following section, we propose a mean-field approximation of the dynamics which only consider the limit $N \to \infty$ and keep $\tau > 0$ fixed. Therefore, this modeling procedure is expected to introduce an error of order $\mathcal{O}(N^{-1})$ only.

2.3 A semi-discrete mean-field model

We start by considering the discrete CBO update rule (2.4) which generates for $k = 0, 1, \dots$ a sequence of random variables $\{X_{(k)}\}_{i=1}^N$. Then, we directly assume propagation of chaos of the marginals. Let $F_{(k)}^N \in \mathcal{P}((\mathbb{R}^d)^N)$ be the probability distribution associated with $\{X_{(k)}\}_{i=1}^N$. We assume that for large particle systems, $N \gg 1$, $F_{(k)}^N \approx f_{(k)}^{\otimes N}$ for some $f_{(k)} \in \mathcal{P}(\mathbb{R}^d)$. Under such assumption, the consensus point satisfies

$$m^\alpha[f_{(k)}^N] = \frac{\int x \exp(-\alpha \mathcal{E}(x)) d f_{(k)}^N(x)}{\int \exp(-\alpha \mathcal{E}(x)) d f_{(k)}^N(x)} \approx \frac{\int x \exp(-\alpha \mathcal{E}(x)) d f_{(k)}(x)}{\int \exp(-\alpha \mathcal{E}(x)) d f_{(k)}(x)} = m^\alpha[f_{(k)}]$$

and so the update rule (2.4) becomes, as before, independent on the particle index i:

$$\overline{X}_{(k+1)} = \overline{X}_{(k)} + \lambda \tau \left(m^\alpha[f_{(k)}] - \overline{X}_{(k)} \right) + \sigma \sqrt{\tau} \left(m^\alpha[f_{(k)}] - \overline{X}_{(k)} \right) \odot \theta_{(k)} \tag{2.8}$$

where $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$ and $\theta_{(k)} \sim \mathcal{N}(0, I_d)$. The obtained difference equation is of McKean type and is the semi-discrete equivalent of model (2.6).

For simplicity, let us forget for a moment about the iterative step k and introduce the following notation

$$\begin{aligned} x' &= x + (\lambda \tau + \sigma \sqrt{\tau} \theta) \odot (m^\alpha[f] - x) \\ &=: \mathcal{C}^{CBO}(x, f, \theta) \end{aligned} \tag{2.9}$$

for a given $f \in \mathcal{P}(\mathbb{R}^d)$ and $\theta \in \mathbb{R}^d$. For any measurable test function $\phi : \mathbb{R}^d \to \mathbb{R}$, we have that the mono-particle process satisfies

$$\mathbb{E}\left[\phi(\overline{X}_{(k+1)})\right] = \mathbb{E}\left[\phi(\mathcal{C}^{CBO}(\overline{X}_{(k)}, f_{(k)}, \theta_{(k)})\right] = \mathbb{E}\left[\phi(\overline{X}'_{(k)})\right]$$

or, equivalently,

$$\int \phi(x)\, d f_{(k+1)}(x) = \iint \phi(x')\, d f_{(k)}(x) d\mu(\theta) \tag{2.10}$$

where, here, $x' = \mathcal{C}^{CBO}(x, f_{(k)}, \theta)$ is given by (2.9) and $\mu = \mathrm{Law}(\theta_{(k)})$. We will favor the notation x' over the more precise $\mathcal{C}^{CBO}(x, f, \theta)$ for the sake of readability, whenever it is clear from the context the arguments we consider.

Again, we note that equation (2.10) describes a dynamics of mean-field type which is discrete time and that can therefore be interpreted as the discrete counterpart of (2.7). We prove now the well-posedness of the above iterative scheme under the following assumptions.

Assumption 2.2. *The objective function $\mathcal{E}$ satisfies:*

- $\underline{\mathcal{E}} := \inf \mathcal{E} > -\infty$;
- *there exists $L_{\mathcal{E}}$, $c_u, c_l, R_l > 0$ such that*

$$\begin{cases} |\mathcal{E}(x) - \mathcal{E}(y)| \le L_{\mathcal{E}}(1 + |x| + |y|)|x - y| & \forall\, x, y \in \mathbb{R}^{\mathrm{d}}, \\ \mathcal{E}(x) - \underline{\mathcal{E}} \le c_u(1 + |x|^2) & \forall\, x \in \mathbb{R}^{\mathrm{d}}, \\ \mathcal{E}(x) - \underline{\mathcal{E}} \ge c_l|x|^2 & \forall\, x : |x| > R_l\,. \end{cases} \tag{2.11}$$

Let $\mathcal{P}_q(\mathbb{R}^{\mathrm{d}})$ be the set of Borel probability measures with bounded q-moment, $\mathcal{P}_q(\mathbb{R}^{\mathrm{d}}) = \{f \in \mathcal{P}(\mathbb{R}^{\mathrm{d}}) : \int |x|^q df(x) < \infty\}$.

Proposition 2.3. *Under Assumption 2.2, if $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^{\mathrm{d}})$, then $f_{(k)} \in \mathcal{P}_2(\mathbb{R}^{\mathrm{d}})$ for all $k \ge 1$.*

Before presenting the proof, we collect the following auxiliary results.

Lemma 2.4 ([16, Lemma 3.3])**.** *Let $\mathcal{E}$ satisfy Assumption 2.2 and $f \in \mathcal{P}_2(\mathbb{R}^{\mathrm{d}})$. Then*

$$\int |x|^2 \frac{\exp(-\alpha\mathcal{E}(x))}{\int \exp(-\alpha\mathcal{E}(x))df(x)} df(x) \le b_1 + b_2 \int |x|^2 df(x)$$

with constants

$$b_1 = R_l^2 + b_2, \quad b_2 = 2\frac{c_u}{c_l}\left(1 + \frac{1}{\alpha c_l}\frac{1}{R_l^2}\right)$$

depending only on R_l, c_u, c_l

Lemma 2.5. *Let θ_ℓ be a random variable taking values in $\mathbb{R}$ with $\mathbb{E}[\theta_\ell] = 0$ and $\mathbb{E}[\theta_\ell^2] = 1$, and define*

$$S_1(\theta_\ell) := 1 - (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell), \qquad S_2(\theta_\ell) := \lambda\tau + \sigma\sqrt{\tau}\theta_\ell\,.$$

It holds

$$\begin{aligned} \mathbb{E}[S_1^2(\theta_\ell)] &= (1 - (2\lambda\tau - \sigma^2\tau - \lambda^2\tau^2)), \\ \mathbb{E}[S_2^2(\theta_\ell)] &= \sigma^2\tau + \lambda^2\tau^2, \\ \mathbb{E}[S_1(\theta_\ell)S_2(\theta_\ell)] &:= \lambda\tau - (\sigma^2\tau + \lambda^2\tau^2). \end{aligned}$$

Proof. Let us start by computing $\mathbb{E}[S_2^2(\theta_\ell)]$. By computing $S_2^2(\theta_\ell)$, we obtain

$$S_2^2(\theta_\ell) = \left(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell\right)^2 = \lambda^2\tau^2 + 2\lambda\tau\sigma\sqrt{\tau}\theta_\ell + \sigma^2\tau\theta_\ell^2.$$

When taking the expectation, the second term vanishes as $\mathbb{E}[\theta_\ell] = 0$. Using assumption $\mathbb{E}[\theta_\ell^2] = 1$ one obtains the desired value

$$\mathbb{E}[S_2^2(\theta_\ell)] = \lambda^2\tau^2 + \sigma^2\tau.$$

Next, we note that $S_1(\theta_\ell) = 1 - S_2(\theta_\ell)$ and so it holds

$$S_1^2(\theta_\ell) = (1 - S_2(\theta_\ell))^2 = 1 - 2S_2(\theta_\ell) + S_2^2(\theta_\ell).$$

By taking the expectation and using the computed value for $\mathbb{E}[S_2^2(\theta_\ell)]$, we get

$$\mathbb{E}[S_1^2(\theta_\ell)] = 1 - 2\lambda\tau + \sigma^2\tau + \lambda^2\tau^2.$$

For the mixed term, it holds

$$S_1(\theta_\ell)S_2(\theta_\ell) = (1 - S_2(\theta_\ell))S_2(\theta_\ell) = S_2(\theta_\ell) - S_2^2(\theta_\ell).$$

Again, by re-using the computed values we obtain

$$\mathbb{E}[S_1(\theta_\ell)S_2(\theta_\ell)] = \lambda\tau - \sigma^2\tau - \lambda^2\tau^2.$$

□

Proof of Proposition 2.3. By applying (2.10) with $\phi(x) = |x|^2$, we obtain

$$\begin{aligned}
\int |x|^2 d f_{(k+1)}(x) &= \iint |x'|^2 d f_{(k)}(x) d\mu(\theta) \\
&= \mathbb{E} \int |(1 - (\lambda\tau + \sigma\sqrt{\tau}\theta)) \odot x + (\lambda\tau + \sigma\sqrt{\tau}\theta) \odot m^\alpha[f_{(k)}]|^2 d f_{(k)}(x) \\
&= \sum_{\ell=1}^{\mathrm{d}} \mathbb{E} \int \left((1 - (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell))x_\ell + (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)m_\ell^\alpha[f_{(k)}]\right)^2 d f_{(k)}(x),
\end{aligned}$$

where the expectation is taken with respect to the sampling of θ. Using the notations S_1, S_2 introduced in Lemma 2.5, the argument of the integral can be written as

$$\begin{aligned}
(x'_\ell)^2 &= |(1 - (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)x_\ell + (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)m_\ell^\alpha[f_{(k)}]|^2 \\
&= (1 - (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell))^2 x_\ell^2 + 2(1 - (\lambda\tau + \sigma\sqrt{\theta}_\ell))(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)x_\ell m_\ell^\alpha[f_{(k)}] \\
&\quad + (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)^2 m_\ell^\alpha[f_{(k)}]^2 \\
&= S_1^2(\theta_\ell)x_\ell^2 + 2S_1(\theta_\ell)S_2(\theta_\ell)x_\ell m_\ell^\alpha[f_{(k)}] + S_2(\theta_\ell)^2 m_\ell^\alpha[f_{(k)}]^2.
\end{aligned}$$

By applying Lemma 2.5 we have that

$$\begin{aligned}\mathbb{E}\int (x'_\ell)^2 df_{(k)}(x) &= \mathbb{E}[S_1^2(\theta_\ell)]\int x_\ell^2 df_{(k)}(x) + \mathbb{E}[2S_1(\theta_\ell)S_2(\theta_\ell)]\int x_\ell m_\ell^\alpha[f_{(k)}]df_{(k)}(x)\\ &\quad + \mathbb{E}[S_2^2(\theta_\ell)]m_\ell^\alpha[f_{(k)}]^2\\ &= (1-(2\lambda\tau-\sigma^2\tau-\lambda^2\tau^2))\int x_\ell^2 df_{(k)}(x)\\ &\quad + 2(\lambda\tau-(\sigma^2\tau+\lambda^2\tau^2))\int x_\ell m_\ell^\alpha[f_{(k)}]df_{(k)}(x) + (\sigma^2\tau+\lambda^2\tau^2)m_\ell^\alpha[f_{(k)}]^2\\ &\le (1+\tau c_1)\int x_\ell^2 df_{(k)}(x) + \tau c_2 m_\ell^\alpha[f_{(k)}]^2\end{aligned}$$

for some constants $c_1, c_2 > 0$ independent on τ for $\tau \in (0,1)$, but that depend on λ, σ. By summing all components $\ell = 1, \dots, \mathrm{d}$ we obtain the estimate

$$\int |x|^2 df_{(k+1)}(x) \le (1+\tau c_1)\int |x|^2 df_{(k)}(x) + \tau c_2 |m^\alpha[f_{(k)}]|^2.$$

Thanks to Assumptions 2.2, we are able to apply Lemma 2.4 and Jensen's inequality to bound $|m^\alpha[f_{(k)}]|^2$ as

$$|m^\alpha[f_{(k)}]|^2 \le \frac{\int |x|^2 \exp(-\alpha\mathcal{E}(x))df_{(k)}(x)}{\int \exp(-\alpha\mathcal{E}(x))df_{(k)}(x)} \le b_1 + b_2\int |x|^2 df_{(k)}(x)$$

for positive constants b_1, b_2 depending on c_u, c_l, R_l, α. This further leads to

$$\int |x|^2 df_{(k+1)}(x) \le (1+\tau c_3)\int |x|^2 df_{(k)}(x) + \tau c_4$$

for some $c_3, c_4 > 0$. By iterating the argument for all $h = 0, 1, \dots, k$, we can conclude

$$\begin{aligned}\int |x|^2 df_{(k)}(x) &\le (1+\tau c_3)^k \int |x|^2 df_{(0)}(x) + \tau c_4 \sum_{h=0}^{k-1}(1+\tau c_3)^h\\ &= (1+\tau c_3)^k \int |x|^2 df_{(0)}(x) + \tau c_4 \frac{(1+\tau c_3)^k - 1}{\tau c_3}\\ &\le e^{c_3 k\tau}\int |x|^2 df_{(0)}(x) + \frac{c_4}{c_3}(e^{c_3 k\tau} - 1),\end{aligned}$$

due to inequality $1 + s \le e^s$ for any $s \in \mathbb{R}$. Therefore, if $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^\mathrm{d})$, the above estimates guarantees that $f_{(k)} \in \mathcal{P}_2(\mathbb{R}^\mathrm{d})$ for all $k \ge 1$. □

We note that the obtained bound on the second moments is analogous to the one obtained in [16, Theorem 3.2] for the continuous mean-field model (2.6) under the same assumptions on $\mathcal{E}$. Indeed, we will see how most of the arguments used in the continuous settings can be adapted to the semi-discrete model, thanks to the simplicity of the CBO dynamics.

Remark 2.6. *By formally taking the limit $\tau \to 0$ in (2.10) with $\phi \in C_b^\infty(\mathbb{R}^d)$, we obtain the Fokker-Planck equation (2.7). For any $k \geq 1$, by Fubini's theorem it holds as $\tau \to 0$*

$$\begin{aligned}\int \phi(x)\, df_{(k+1)}(x) &= \iint \phi(x')\, df_{(k)}(x) d\mu(\theta) \\ &= \int \mathbb{E}\Big[\phi(x) + \nabla\phi(x)\cdot(x'-x) \\ &\qquad + \frac{1}{2}\sum_{\ell_1,\ell_2=1}^{d} \partial_{\ell_1\ell_2}\phi(x)(x'-x)_{\ell_1}(x'-x)_{\ell_2}\Big] df_{(k)}(x) + o(\tau).\end{aligned}$$

Now, we recall from the proof of Proposition 2.3 that $\mathbb{E}[x'-x] = \lambda\tau(m^\alpha[f_{(k)}]-x)$, $\mathbb{E}[(x'-x)_{\ell_1}(x'-x)_{\ell_2}] = o(\tau)$ for all $\ell_1 \neq \ell_2$, and $\mathbb{E}[(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)^2] = \lambda\tau^2 + \sigma^2\tau$. This yields to

$$\int \phi(x)\, df_{(k+1)}(x) = \int \Big(\phi(x) + \lambda\tau\nabla\phi(x)\cdot(m^\alpha[f_{(k)}]-x) + \tau\frac{\sigma^2}{2}\sum_{\ell=1}^{d}\partial_{\ell\ell}\phi(x)(x'-x)_\ell^2\Big) df_{(k)}(x) + o(\tau)$$

which we rewrite as

$$\begin{aligned}&\frac{\int \phi(x)\, df_{(k+1)}(x) - \int \phi(x) df_{(k)}(x)}{\tau} \\ &\qquad = \int \Big(\phi(x) + \lambda\nabla\phi(x)\cdot(m^\alpha[f_{(k)}]-x) + \frac{\sigma^2}{2}\sum_{\ell=1}^{d}\partial_{\ell\ell}\phi(x)(x'-x)_\ell^2\Big) df_{(k)}(x) + o(1).\end{aligned}$$

As $\tau \to 0$, the above equation formally converges to the time-continuous mean-field model (2.7).

2.4 A variant with memory effects

We have shown in Section 2.1 how the CBO particle interaction can be seen as a simpler variant of the PSO dynamics. Inspired by recent analysis of PSO algorithms [43, 59] and CBO variants [74], we propose in the following a mean-field CBO model where the memory mechanism of PSO methods is reintroduced. This is achieved by coupling the particle location $X^i_{(k)}$ with the corresponding personal best $Y^i_{(k)}$ and by computing the regularized

global best position through the measure $f_{(k)}^{N,y} = (1/N)\sum_{i=1}^{N} \delta_{Y_{(k)}^i}$:

$$\begin{aligned} X_{(k+1)}^i &= X_{(k)}^i + \lambda\tau\left(m^\alpha[f_{(k)}^{N,y}] - X_{(k)}^i\right) + \sigma\sqrt{\tau}\left(m^\alpha[f_{(k)}^{N,y}] - X_{(k)}^i\right)\odot\theta_{(k)}^i, \\ Y_{(k+1)}^i &= \begin{cases} X_{(k+1)}^i & \text{if } \mathcal{E}(X_{(k+1)}^i) < \mathcal{E}(Y_{(k)}^i), \\ Y_{(k)}^i & \text{else}. \end{cases} \end{aligned} \tag{2.12}$$

Let $f_{(k)}^N \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ be the particles' empirical measure. Under the same propagation of chaos assumption we used for the CBO dynamics, $f_{(k)}^N \approx f_{(k)} \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ for $N \gg 1$, we obtain the mean-field process

$$\begin{aligned} \overline{X}_{(k+1)} &= \overline{X}_{(k)} + \lambda\tau\left(m^\alpha[f_{(k)}^y] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha[f_{(k)}^y] - \overline{X}_{(k)}\right)\odot\theta_{(k)}, \\ \overline{Y}_{(k+1)} &= \begin{cases} \overline{X}_{(k+1)} & \text{if } \mathcal{E}(\overline{X}_{(k+1)}) < \mathcal{E}(\overline{Y}_{(k)}), \\ \overline{Y}_{(k)} & \text{else}, \end{cases} \end{aligned} \tag{2.13}$$

where $f_{(k)}^y = \mathrm{Law}(\overline{Y}_{(k)})$.

For any test measurable function $\phi : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$, it holds

$$\begin{aligned} \mathbb{E}\left[\phi(\overline{X}_{(k+1)}, \overline{Y}_{(k+1)})\right] &= \mathbb{E}\left[\phi(\overline{X}_{(k+1)}, \overline{X}_{(k+1)}) \,\middle|\, \mathcal{E}(\overline{X}_{(k+1)}) < \mathcal{E}(\overline{Y}_{(k)})\right] \\ &\quad + \mathbb{E}\left[\phi(\overline{X}_{(k+1)}, \overline{Y}_{(k)}) \,\middle|\, \mathcal{E}(\overline{X}_{(k+1)}) \geq \mathcal{E}(\overline{Y}_{(k)})\right] \\ &= \mathbb{E}\left[\phi(\overline{X}_{(k+1)}, \overline{X}_{(k+1)})\mathbf{1}_{\{\mathcal{E}(\overline{X}_{(k+1)}) < \mathcal{E}(\overline{Y}_{(k)})\}}\right] \\ &\quad + \mathbb{E}\left[\phi(\overline{X}_{(k+1)}, \overline{Y}_{(k)})\mathbf{1}_{\{\mathcal{E}(\overline{X}_{(k+1)}) \geq \mathcal{E}(\overline{Y}_{(k)})\}}\right] \end{aligned}$$

in which the indicator function is defined as $\mathbf{1}_{\{\mathtt{True}\}} = 1$ and $\mathbf{1}_{\{\mathtt{False}\}} = 0$. As a consequence, $f_{(k+1)}$ satisfies

$$\begin{aligned} &\int \phi(x,y)\, df_{(k+1)}(x,y) \\ &\qquad = \iint \left(\phi(x',x')\mathbf{1}_{\{\mathcal{E}(x') < \mathcal{E}(y)\}} + \phi(x',y)\mathbf{1}_{\{\mathcal{E}(x') \geq \mathcal{E}(y)\}}\right) df_{(k)}(x,y)\, d\mu(\theta) \end{aligned} \tag{2.14}$$

where x' is again given by the CBO dynamics (2.9), but with $f = f_{(k)}^y$.

It is interesting to notice how the model (2.14) exactly implements the memory mechanism, unlike similar time continuous models in which the personal bests update rule is regularized [43, 59, 74]. As we will see in the next chapter, it is possible to prove convergence towards the global minimizer directly for (2.14), but a regularized update rule may be required to obtain quantitative error bounds for the mean-field approximation. Moments' estimates for (2.13) can be derived similarly to the CBO dynamics without memory effects.

Proposition 2.7. *Under Assumption 2.2, if $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$, then $f_{(k)} \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$ for all $k \geq 1$.*

Proof. Assume $f_{(k)} \in \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d)$. Following the computations done in the proof of Proposition 2.3 we obtain

$$\begin{aligned}\int |x|^2 df^x_{(k+1)}(x) &\leq (1-\tau c_1)\int |x|^2 df^x_{(k)}(x) + \tau c_2 |m^\alpha[f^y_{(k)}]|^2 \\ &\leq (1-\tau c_1)\int |x|^2 df^x_{(k)}(x) + \tau c_2 \left(b_1 + b_2 \int |y|^2 df^y_{(k)}\right) < \infty\end{aligned}$$

where we applied, again, Lemma 2.4 to bound $|m^\alpha[f^y_{(k)}]|^2$. For the particle best, a rough estimate can be given by noting that

$$\begin{aligned}\int |y|^2 df^y_{(k+1)}(y) &= \int |x'|^2 \mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} df_{(k)}(x,y) + \int |y|^2 \mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} df_{(k)}(x,y) \\ &\leq \int |x'|^2 df^x_{(k)}(x) + \int |y|^2 df^y_{(k)}(y) < \infty,\end{aligned}$$

due to the previously computed bound. This proves the claim by induction.

As it will be useful later for the convergence analysis, let us derive a sharper estimate by exploiting the memory mechanism. Thanks to Assumption 2.2, we have

$$\begin{aligned}\int |y|^2 df^y_{(k)}(y) &= \int_{B^2_{R_l}(0)} |y|^2 df^y_{(k)}(y) + \int_{(B^2_{R_l}(0))^c} |y|^2 df^y_{(k)}(y) \\ &\leq R_l^2 + \frac{1}{c_l}\int_{\mathbb{R}^d} (\mathcal{E}(y) - \underline{\mathcal{E}})\, df^y_{(k)}(y) \\ &\leq R_l^2 + \frac{1}{c_l}\left(\int_{\mathbb{R}^d} \mathcal{E}(y) df^y_{(k)}(y) - \int_{\mathbb{R}^d} \mathcal{E}(y_0) df^y_{(0)}(y_0) + \int_{\mathbb{R}^d} (\mathcal{E}(y_0) - \underline{\mathcal{E}})\, df^y_{(0)}(y_0)\right).\end{aligned}$$

Now, since $\mathcal{E}(\overline{Y}_{(k)}) \leq \mathcal{E}(\overline{Y}_{(0)})$ with probability 1, as the personal best can improve its objective value, it holds

$$\int_{\mathbb{R}^d} \mathcal{E}(y) df^y_{(k)}(y) \leq \int_{\mathbb{R}^d} \mathcal{E}(y_0) df^y_{(0)}(y_0).$$

Therefore, we can drop their difference in the previous estimate and further bound the second moment as

$$\int |y|^2 df^y_{(k)}(y) \leq R_l^2 + \frac{1}{c_l}\int (\mathcal{E}(y_0) - \underline{\mathcal{E}})\, df^y_{(0)}(y_0) < +\infty. \tag{2.15}$$

We note that the right-hand side of (2.15) is independent on both k and τ. □

2.5 Discussion

In this chapter we introduced the CBO algorithm and proposed a novel mean-field approximation by directly considering the many-particle limit $N \to \infty$ of the CBO update rule. We also recalled the most common modeling approach in the CBO literature which consist of deriving the mean-field model from a time-continuous system of SDEs. To conclude the chapter we mention other approaches used in the literature.

In [48], authors approximate the CBO particle system as $\tau \to 0$, but do not perform a mean-field approximation of the particle interaction. This approach is taken further in [49, 67], where the convergence analysis is directly performed at the fully discrete CBO particle system. As the theoretical analysis does not rely on mean-field averaging, convergence is proved under rather restrictive assumptions on the initial data (see, for instance, [49, Theorem 3.2]). We note that a time-discrete model which is similar to the one proposed here, has been employed in [18], where authors propose a variant of CBO, the Consensus-Based Sampling method, for inverse and optimization problems.

Chapter 3

Convergence analysis of consensus-based optimization algorithms

In this chapter, we investigate the large-time behavior of the CBO algorithm iteration by exploiting the mean-field approximation of the particle dynamics. We recall that for an initial particle distribution $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^d)$ the proposed mean-field model consists of a sequence of probability measures $(f_{(k)})_{k\in\mathbb{Z}_+} \subset \mathcal{P}_2(\mathbb{R}^d)$ (thanks to Proposition (2.3)) where $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$, and $\overline{X}_{(k)}$ is given by the difference equation

$$\overline{X}_{(k+1)} = \overline{X}_{(k)} + \lambda\tau\left(m^\alpha[f_{(k)}] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha[f_{(k)}] - \overline{X}_{(k)}\right)\odot\theta_{(k)} \tag{3.1}$$

with $\theta_{(k)} \sim \mathcal{N}(0, I_d)$.

Among other assumptions, we will restrict our analysis to objective functions $\mathcal{E}$ attaining a unique global minimizer $x^* := \mathrm{argmin}_{x\in\mathbb{R}^d}\mathcal{E}(x)$. Let $f \in \mathcal{P}_2(\mathbb{R}^d)$ be a particle distribution, we investigate convergence by studying the expected mean-field error of f with respect to x^*:

$$\mathrm{Err}[f] := \int |x - x^*|^2 df(x). \tag{3.2}$$

In particular, we show that $\mathrm{Err}[f_{(k)}]$ can reach, during the computation, any desired accuracy $\varepsilon > 0$, provided the objective function $\mathcal{E}$ and the CBO parameters $\{\lambda, \sigma, \tau, \alpha\}$ satisfy certain conditions. We note that the above error coincide with the squared Wasserstein-2 distance between f and the probability measure concentrated at the solution x^*, $\mathrm{Err}[f] = W_2^2(f, \delta_{x^*})$. We refer to [75] for the definition and more details on Wasserstein distances.

In this chapter, we start by reviewing in Section 3.1 the main strategies used in the literature to prove convergence of the continuous-in-time mean-field model. In Section 3.2 we then adapt such techniques to study (3.1) and prove convergence in mean-field

law without relying on time-continuous approximations. We will do that for the CBO algorithm both with and without memory effects. Next, in Section 3.3 we consider particle dynamics with bounded search space and provide quantitative error estimates of the mean-field approximation in terms of number of particles N. By coupling the mean-field approximation error with the mean-field convergence analysis, we are able to prove convergence of the CBO algorithm for bounded, convex domains in expectation. Final remarks are discussed in Section 3.4, while Section 3.5 collects the proof of auxiliary lemmas.

3.1 Literature review on convergence of CBO dynamics

The key ingredient for the convergence analysis of CBO methods is the Boltzmann-Gibbs distribution used to compute the consensus point. In particular, this choice of weights allows to use a classical result of large-deviation analysis: the *Laplace principle* [27]. This principle states that for any $f \in \mathcal{P}(\mathbb{R}^d)$ absolutely continuous with respect to the Lebesgue measure it holds

$$\lim_{\alpha\to\infty}\left(-\frac{1}{\alpha}\log\left(\int e^{-\alpha\mathcal{E}(x)}df(x)\right)\right) = \inf_{x\in\mathrm{supp}(f)}\mathcal{E}(x).$$

The approaches available in the literature can be summarized into two main ones, which we briefly recall for completeness. In the following, $f \in C([0,T],\mathcal{P}_2(\mathbb{R}^d))$ is assumed to be a solution to the mean-field continuous model (2.7).

The first technique was proposed in [16] and further applied in many other works [19, 38, 39, 48–50]. To follow this strategy, one needs to study the consensus formation through the evolution of the system variance $\mathrm{Var}[f_t] = (1/2)\int |x - m^0[f]|^2 df_t(x)$. Under suitable assumptions, it is typically possible to prove exponential decay of the variance, leading to $W_2(f_t,\delta_{\tilde{x}}) \to 0$ as $t\to\infty$ for some point $\tilde{x}\in\mathbb{R}^d$. Next, provided $\mathrm{Var}[f_0]$ is sufficiently small, $\mathcal{E}\in C^2(\mathbb{R}^d)$ is sufficiently regular, and the interaction parameters are well-chosen, one can prove the upper bound as $\alpha\to\infty$

$$\mathcal{E}(\tilde{x}) \le -\frac{1}{\alpha}\log\Big(\int e^{-\alpha\mathcal{E}(x)}df_0(x)\Big) + \mathcal{O}(\alpha^{-1}),$$

see, for instance, [19, Theorem 3.2]. Finally, thanks to the Laplace principle, the above upper bound can be made arbitrary close to the objective value global minimum $\mathcal{E}(x^*)$, assuming large α and $x^*\in\mathrm{supp}(f_0)$.

This variance-based approach, has shown to be flexible enough to be successfully applied to many variants of CBO algorithms [5, 58]. On the other hand, differentiability of the objective function is required together with a strict assumption on the initial variance, as $\mathrm{Var}[f_0]$ and α need to be inversely proportional. This gives a local flavor to the result.

The second approach has been proposed in [41] for isotropic CBO dynamics, and adapted in [42] for anisotropic ones. The main intuition behind these techniques is to show that the mean-field dynamics corresponds to a gradient flow with respect to the mean squared error. In this case, $\mathrm{Err}[f_t]$ itself plays the role of a Lyapunov functional, rather than the variance. In the analysis of the proposed semi-discrete mean-field model (3.1) we are going to take this same point of view. Therefore, we recall in the following the main steps of the convergence analysis carried out in [41, 42].

As before, the key result used for the analysis is the Laplace principle. In [41], authors derive a quantitative version of this asymptotic result to prove that the consensus point $m^\alpha[f_t]$ can be made arbitrary close to the global minimum x^* if the mass around the minimizer does not vanish during the evolution. To be more precise, the quantitative version of the Laplace principle for anisotropic CBO can be obtained under the assumptions below.

Assumption 3.1 (Growth conditions around minimizer)**.** *There exists a unique minimizer x^* of $\mathcal{E}$ and $c_1, p_1 > 0, R > 0$ such that*

$$c_1 \|x - x^*\|_\infty^{p_1} \leq \mathcal{E}(x) - \mathcal{E}(x^*) \qquad \forall x,\ \|x - x^*\|_\infty \leq R, \tag{3.3}$$

and lower bound $\mathcal{E}_\infty > 0$ such that

$$\mathcal{E}_\infty < \mathcal{E}(x) - \mathcal{E}(x^*) \qquad \forall x,\ \|x - x^*\|_\infty > R. \tag{3.4}$$

In the following, we denote with $B_r^\infty(x)$ the closed ℓ_∞-ball centered in $x \in \mathbb{R}^\mathrm{d}$ of radius $r > 0$. The quantitative version of the Laplace principle then reads:

Proposition 3.2. *[42, Proposition 1] Let $\inf \mathcal{E} = 0$, $f \in \mathcal{P}(\mathbb{R}^\mathrm{d})$ and fix $\alpha > 0$. For any $r > 0$ we define $\mathcal{E}_r := \sup_{x \in B_r^\infty(x^*)} \mathcal{E}(x)$. Then, under Assumption 3.1, for any $r \in (0, R]$ and $q > 0$ such that $q + \mathcal{E}_r < \mathcal{E}_\infty$, we have*

$$|m^\alpha[f] - x^*| \leq c_1 \sqrt{\mathrm{d}} (q + \mathcal{E}_r)^{1/p_1} + \frac{\sqrt{\mathrm{d}} \exp(-\alpha q)}{f(B_r^\infty(x^*))} \int |x - x^*| df(x).$$

Thanks to this quantitative estimate, authors in [41, 42] are able to identify, for a given accuracy $\varepsilon > 0$ and initial data f_0, a sufficiently large time horizon T^* and α such that

$$\min_{t \in [0, T^*]} \mathrm{Err}[f_t] \leq \varepsilon.$$

Also, the error $\mathrm{Err}[f_t]$ exponentially decays until the desired accuracy is reached. We remark that assumption $\inf \mathcal{E} = 0$ can be dropped without loss of generality, as $m^\alpha[f]$ is invariant under translations of $\mathcal{E}$.

In the next sections, we show to obtain the same convergence result for the finite-difference mean-field model.

Remark 3.3. *The dependence on the dimension* d *in Proposition 3.2 is a consequence of the use of ℓ_∞-norms in the growth conditions* (3.3) *and* (3.4)*. Bounds in terms of ℓ_2-norms do not lead to such an explicit dependence on the problem dimension, see [41, Proposition 21]. We include the above result as it is more suitable to analyze the CBO dynamics with anisotropic exploration.*

3.2 Convergence analysis in mean-field law

To study the convergence properties of the mean-field model (3.1) (or (2.13) if memory effects are considered) we follow the proof technique introduced in [41, 42]. The proof for anisotropic CBO requires the following steps:

1. study time evolution of the error $\mathrm{Err}[f_{(k)}]$ for $k \geq 0$;
2. provide a lower-bound on the mass around the neighborhood of the minimizer $f_{(k)}(B_r^\infty(x^*))$, for some small radius $r > 0$;
3. apply the quantitative Laplace principle (Proposition 3.2) to show that $|m^\alpha[f_{(k)}] - x^*| = \mathcal{O}(\varepsilon)$ for large $\alpha > 0$;
4. use the collected estimates to prove $\mathrm{Err}[f_{(K)}] = \mathcal{O}(\varepsilon)$ for some large $K > 0$.

As we will see, steps 1, 3 and 4 can be adapted with no effort to the semi-discrete settings. In case of CBO without memory effects, we provide a lower bound on the probability mass around the minimizer (step 2) by relying on computations performed for the time continuous mean-field model. As a result, we require a rather strict upper bound on the step size τ. With the introduction of memory effects, though, we follow a new approach to address step 3 which does not require any strict assumption on τ.

3.2.1 CBO without memory effects

In this subsection, we consider the sequence of probability measures $(f_{(k)})_{k\in\mathbb{Z}_+}$ where $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)}) \in \mathcal{P}_2(\mathbb{R}^\mathrm{d})$ with $\overline{X}_{(k)}$ defined by (3.1). We start the analysis by studying the evolution of the mean-squared error.

Proposition 3.4 (Error evolution)**.** *For all $k = 0, 1, \dots$ and $\lambda\tau < 1$, it holds*

$$\mathrm{Err}[f_{(k+1)}] \leq \left(1 - \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{2}\right)\mathrm{Err}[f_{(k)}] + \tau\frac{2\lambda + \sigma^2 + \lambda^2\tau}{2}|m^\alpha[f_{(k)}] - x^*|^2. \tag{3.5}$$

Proof. By definition of $\mathrm{Err}[f_{(k)}]$ and the weak formulation (2.10), it holds

$$\mathrm{Err}[f_{(k+1)}] = \int |x - x^*|^2 df_{(k+1)}(x) = \mathbb{E}_{\theta\sim\mu}\int |x' - x^*|^2 df_{(k)}(x),$$

with $x' = \mathcal{C}^{CBO}(x, f_{(k)}, \theta)$, as in (2.9). For the sake of notational semplicity, we will simply use $\mathbb{E}[\cdot]$ to denote $\mathbb{E}_{\theta\sim\mu}[\cdot]$ in the following computations.

We rewrite $x' - x^*$ as

$$\begin{aligned} x' - x^* &= x + (\lambda\tau + \sigma\sqrt{\tau}\theta) \odot (m^\alpha[f_{(k)}] - x) - x^* \\ &= (1 - (\lambda\tau + \sigma\sqrt{\tau}\theta)) \odot (x - x^*) + (\lambda\tau + \sigma\sqrt{\tau}\theta) \odot (m^\alpha[f_{(k)}] - x^*) \end{aligned}$$

to obtain

$$\begin{aligned} \mathbb{E}\Big[|x' - x^*|^2\Big] &= \mathbb{E}\Big[|(1 - (\lambda\tau + \sigma\sqrt{\tau}\theta)) \odot (x - x^*) + (\lambda\tau + \sigma\sqrt{\tau}\theta) \odot (m^\alpha[f_{(k)}] - x^*)|^2\Big] \\ &= \sum_{\ell=1}^{\mathrm{d}} \mathbb{E}\Big[\Big((1 - (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell))(x - x^*)_\ell + (\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)(m^\alpha[f_{(k)}] - x^*)\Big)_\ell^2\Big] \\ &= \sum_{\ell=1}^{\mathrm{d}} \mathbb{E}\Big[\Big(S_1(\theta_\ell)(x - x^*)_\ell + S_2(\theta_\ell)(m^\alpha[f_{(k)}] - x^*)_\ell\Big)^2\Big], \end{aligned}$$

where we used the definition of S_1, S_2 given in Lemma 2.5. By applying now Lemma 2.5 one further obtains for all coordinates $\ell = 1, \dots, \mathrm{d}$

$$\begin{aligned} &\mathbb{E}\Big[\Big(S_1(\theta_\ell)(x - x^*)_\ell + S_2(\theta_\ell)(m^\alpha[f_{(k)}] - x^*)_\ell\Big)^2\Big] \\ &\quad= \mathbb{E}\Big[S_1^2(\theta_\ell)(x - x^*)_\ell^2 + 2S_1(\theta_\ell)S_2(\theta_\ell)(x - x^*)_\ell(m^\alpha[f_{(k)}] - x^*)_\ell \\ &\qquad + S_2^2(\theta_\ell)(m^\alpha[f_{(k)}] - x^*)_\ell^2\Big] \\ &\quad= \mathbb{E}[S_1^2(\theta_\ell)](x - x^*)_\ell^2 + 2\mathbb{E}[S_1(\theta_\ell)S_2(\theta_\ell)](x - x^*)_\ell(m^\alpha[f_{(k)}] - x^*)_\ell \\ &\qquad + \mathbb{E}[S_2^2(\theta_\ell)](m^\alpha[f_{(k)}] - x^*)_\ell^2 \\ &\quad= (1 - (2\lambda\tau - \sigma^2\tau - \lambda^2\tau^2))(x - x^*)_\ell^2 + 2\Big(\lambda\tau - (\sigma^2\tau + \lambda^2\tau^2)\Big)(x - x^*)_\ell(m^\alpha[f_{(k)}] - x^*)_\ell \\ &\qquad + (\sigma^2\tau + \lambda^2\tau^2)(m^\alpha[f_{(k)}] - x^*)_\ell^2 \\ &\quad\le \left(1 - \left(\lambda\tau - \frac{\sigma^2\tau + \lambda^2\tau^2}{2}\right)\right)(x - x^*)_\ell^2 + \left(\lambda\tau + \frac{\sigma^2\tau + \lambda^2\tau^2}{2}\right)(m^\alpha[f_{(k)}] - x^*)_\ell^2, \end{aligned}$$

where in the last step we used inequality $ab \le \frac{1}{2}a^2 + \frac{1}{2}b^2$ for any $a, b \in \mathbb{R}$ to bound the second term.

Finally, we sum over all coordinates $\ell = 1, \dots, \mathrm{d}$ and integrate x with respect to $f_{(k)}$ to get the desired estimate

$$\begin{aligned} \mathbb{E}\int |x' - x^*|^2 d f_{(k)}(x) \le \left(1 - \tau\left(\lambda - \frac{\sigma^2 + \lambda^2\tau}{2}\right)\right)\int |x - x^*|^2 d f_{(k)}(x) \\ + \tau\left(\lambda + \frac{\sigma^2 + \lambda^2\tau}{2}\right)|m^\alpha[f_{(k)}] - x^*|^2. \end{aligned}$$

□

It is clear from Proposition 3.4 that if one can bound the term $|m^\alpha[f_{(k)}] - x^*|$, then a suitable choice of λ, σ, τ leads to exponential decay of $\mathrm{Err}[f_{(k)}]$. We will do that by applying the quantitative Laplace principle (Proposition 3.2) over a sufficiently large time window. This requires to provide a lower bound on

$$f_{(k)}(B_r^\infty(x^*)) = f_{(k)}\Big(\Big\{x : \|x - x^*\|_\infty < r\Big\}\Big)$$

for any $r > 0$.

To do so, we follow [42] and consider the mollifier

$$\phi_r(x) = \begin{cases} \prod_{\ell=1}^{\mathrm{d}} \exp\Big(1 - \frac{r^2}{r^2 - (x - x^*)_\ell^2}\Big), & \text{if } \|x - x\|_\infty < r, \\ 0, & \text{else,} \end{cases} \tag{3.6}$$

which satisfies $\mathrm{Im}(\phi_r) = [0, 1]$, $\mathrm{supp}(\phi_r) = B_r^\infty(x^*)$, $\phi \in C_c^\infty(\mathbb{R}^{\mathrm{d}})$. Hence, the expected value of ϕ_r gives us a lower bound on the mass around the minimizer

$$f_{(k)}(B_r^\infty(x^*)) \geq \int \phi_r(x) d f_{(k)}(x).$$

Lemma 3.5 (Lower bound on mass around x^*)**.** *For a given $T > 0$ and $r > 0$, let*

$$\max_{k,\, k\tau \leq T} |m^\alpha[f_{(k)}] - x^*| \leq B$$

for some $B > 0$.

Provided $\tau > 0$ is sufficiently small, there exists a positive constant $a = a(r, \mathrm{d}, B, \lambda, \sigma)$ (independent on τ) such that

$$f_{(k)}(B_r^\infty(x^*)) \geq \left(\frac{1}{2} \int \phi_r(x) d f_{(0)}(x)\right) \exp(-ak\tau) \quad \forall k : k\tau \leq T.$$

The proof of the lemma is rather technical and it is given at the end of the chapter in Section 3.5. Now, we are ready to present and tackle the main convergence theorem.

Theorem 3.6. *Fix an accuracy $\varepsilon > 0$ and assume the objective function $\mathcal{E}$ satisfies Assumption 3.1. Consider an initial distribution $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^{\mathrm{d}})$ such that $x^* \in \mathrm{supp}(f_{(0)})$ and parameters $\{\lambda, \sigma\} \subset \mathbb{R}_+$ satisfying $\lambda > \sigma^2$. Define the time horizon $T^* > 0$*

$$T^* := \frac{4}{\lambda - \sigma^2} \log\left(\frac{\mathrm{Err}[f_{(0)}]}{\varepsilon}\right).$$

Let $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$ with $\overline{X}_{(k)}$ updated according to (3.1). Provided $\alpha > 0$ is sufficiently large and the step-size $\tau > 0$ is sufficiently small, it holds

$$\min_{k : k\tau \leq T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon. \tag{3.7}$$

Moreover, as long as $\mathrm{Err}[f_{(k-1)}] > \varepsilon$, *it holds*

$$\mathrm{Err}[f_{(k)}] \leq \exp\left(-k\tau\frac{\lambda-\sigma^2}{4}\right)\mathrm{Err}[f_{(0)}]. \tag{3.8}$$

Proof. We start by defining suitable $\tau_\varepsilon, \alpha_\varepsilon$ such that for all $\tau \leq \tau_\varepsilon$ and $\alpha \geq \alpha_\varepsilon$ we have convergence in the sense of (3.7). To this end, we introduce

$$C_{\lambda,\sigma} := \sqrt{\frac{\lambda-\sigma^2}{2(3\lambda+\sigma^2)}}$$

and $q_\varepsilon, r_\varepsilon$ given by

$$q_\varepsilon := \frac{1}{2}\min\left\{\frac{1}{c_1\sqrt{\mathrm{d}}}\left(\sqrt{\varepsilon}C_{\lambda,\sigma}\right)^{p_1}, \mathcal{E}_\infty\right\}, \quad r_\varepsilon := \max_{s\in[0,R]}\left\{\max_{x\in B_s^\infty(x^*)}\mathcal{E}(x) \leq q_\varepsilon\right\},$$

such that, thanks to this choice, it holds

$$c_1\sqrt{\mathrm{d}}(q_\varepsilon + \mathcal{E}_{r_\varepsilon})^{1/p_1} < \frac{\sqrt{\varepsilon}C_{\lambda,\sigma}}{2} \tag{3.9}$$

(recall $\mathcal{E}_r$ is defined as $\mathcal{E}_r := \sup_{x\in B_r^\infty(x^*)}\mathcal{E}(x)$ in Proposition 3.2).

Then, we pick a step-size $\tau_\varepsilon > 0$ such that $\lambda\tau_\varepsilon < 1$ and such that Lemma 3.5 holds with $r = r_\varepsilon$, $B := C_{\lambda,\sigma}\sqrt{\mathrm{Err}[f_{(0)}]}$, and $T = T^*$. In this way, we obtain an exponential decay of $f_{(k)}(B_{r_\varepsilon}^\infty(x^*))$ with exponent $a > 0$. We remark that the definition of τ_ε, ultimately, only depends on $\varepsilon, f_{(0)}\lambda, \sigma$ and it does not depend on any choice of α.

Next, we consider any α_ε satisfying

$$\frac{\sqrt{\mathrm{d}}\exp(-\alpha_\varepsilon q_\varepsilon + aT^*)}{\int \phi_{r_\varepsilon}(x)df_{(0)}(x)}\sqrt{\mathrm{Err}[f_{(0)}]} < \frac{C_{\lambda,\sigma}}{2}\sqrt{\varepsilon}.$$

These particular choices of $\tau_\varepsilon, \alpha_\varepsilon$ will allow us to apply both Lemma 3.5 and Proposition 3.2 as long as $|m^\alpha[f_{(k)}] - x^*| \leq B$.

Now, for any fixed $\tau \leq \tau_\varepsilon$ and $\alpha \geq \alpha_\varepsilon$, we consider the iterative step K

$$K := \sup\left\{k \,:\, \mathrm{Err}[f_{(h)}] > \varepsilon \text{ and } |m^\alpha[f_{(h)}] - x^*| < C_{\lambda,\sigma}\sqrt{\mathrm{Err}[f_{(h)}]} \;\; \forall h \leq k\right\}. \tag{3.10}$$

By applying Proposition 3.4 and since $\lambda\tau \le \lambda\tau_\varepsilon < 1$, we obtain that for all $k \le K$

$$\begin{aligned}
\mathrm{Err}[f_{(k)}] &\le \left(1 - \frac{2\lambda\tau - \sigma^2\tau - \lambda^2\tau^2}{2}\right)\mathrm{Err}[f_{(k-1)}] + \frac{2\lambda\tau + \sigma^2\tau + \lambda^2\tau^2}{2}|m^\alpha[f_{(k-1)}] - x^*|^2 \\
&\le \left(1 - \tau\frac{\lambda - \sigma^2}{2}\right)\mathrm{Err}[f_{(k-1)}] + \tau\frac{3\lambda + \sigma^2}{2}|m^\alpha[f_{(k-1)}] - x^*|^2 \\
&\le \left(1 - \tau\frac{\lambda - \sigma^2}{2}\right)\mathrm{Err}[f_{(k-1)}] + \tau\frac{3\lambda + \sigma^2}{2}C_{\lambda,\sigma}^2\,\mathrm{Err}[f_{(k-1)}] \\
&\le \left(1 - \tau\frac{\lambda - \sigma^2}{4}\right)\mathrm{Err}[f_{(k-1)}],
\end{aligned}$$

where the last two inequalities follow from the definition of K and $C_{\lambda,\sigma}$, respectively. By iterating the above argument, we obtain for all $k \le K$

$$\mathrm{Err}[f_{(k)}] \le \left(1 - \tau\frac{\lambda - \sigma^2}{4}\right)^k \mathrm{Err}[f_{(0)}] \le \exp\left(-k\tau\frac{\lambda - \sigma^2}{4}\right)\mathrm{Err}[f_{(0)}].$$

We note that the right-hand-side above is decreasing due to the assumption $\lambda > \sigma^2$. Therefore, $\mathrm{Err}[f_{(k)}]$ and $|m^\alpha[f_{(k)}] - x^*|$ can be bounded as following:

$$\begin{aligned}
\max_{0\le k\le K}\mathrm{Err}[f_{(k)}] &\le \mathrm{Err}[f_{(0)}] \\
\max_{0\le k\le K}|m^\alpha[f_{(k)}] - x^*| &\le C_{\lambda,\sigma}\max_{0\le k\le \bar{k}}\sqrt{\mathrm{Err}[f_{(k)}]} \le C_{\lambda,\sigma}\sqrt{\mathrm{Err}[f_{(0)}]} = B.
\end{aligned}$$

where, again, we used that $|m^\alpha[\rho_{(k)}] - x^*| < C_{\lambda,\sigma}\mathrm{Err}[f_{(k)}]$ as long as $k \le K$.

To show $\mathrm{Err}[f_{(k)}] \le \varepsilon$ for some k, we check three different cases.

Case $K \ge T^*/\tau$. Thanks to the definition of T^* and the error exponential decay up to iteration K, we have

$$\mathrm{Err}[f_{(K)}] \le \exp\left(-K\tau\frac{\lambda - \sigma^2}{4}\right)\mathrm{Err}[f_{(0)}] \le \exp\left(-T^*\frac{\lambda - \sigma^2}{4}\right)\mathrm{Err}[f_{(0)}] \le \varepsilon\,.$$

Case $K < T^*/\tau$ and $\mathrm{Err}[f_{(K)}] \le \varepsilon$. Nothing to prove here.

Case $K < T^*/\tau$, $\mathrm{Err}[f_{(K)}] > \varepsilon$ and $|m^\alpha[f_{(K)}] - x^*| \ge C_{\lambda,\sigma}\sqrt{\mathrm{Err}[f_{(K)}]}$.

We prove that, due to our choice $\alpha \ge \alpha_\varepsilon$ and $\tau \le \tau_\varepsilon$ it holds $|m^\alpha[f_{(K)}] - x^*| < C_{\lambda,\sigma}\sqrt{\mathrm{Err}[f_{(K)}]}$ and therefore we are led to a contradiction.

By applying Proposition 3.2 with $q = q_\varepsilon$, $r = r_\varepsilon$, and Jensen's inequality we obtain

$$\begin{aligned}
|m^\alpha[f_{(K)}] - x^*| &\le c_1\sqrt{\mathrm{d}}(q_\varepsilon + \mathcal{E}_{r_\varepsilon})^{1/p_1} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon)}{f(B_{r_\varepsilon}^\infty(x^*))}\int |x - x^*|\,df_{(K)}(x) \\
&< \frac{\sqrt{\varepsilon}C_{\lambda,\sigma}}{2} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon)}{f_{(K)}(B_{r_\varepsilon}^\infty(x^*))}\sqrt{\mathrm{Err}[f_{(K)}]},
\end{aligned}$$

thanks to our choice of $q_\varepsilon, r_\varepsilon$ satisfying (3.9).

To bound the second term, we recall that $\mathrm{Err}[f_{(K)}] \leq \mathrm{Err}[f_{(0)}]$. Thanks to the choice $\tau \leq \tau_\varepsilon$, we can also apply Lemma 3.5 with r_ε to obtain a lower bound on the mass around the minimizer which decays exponentially with the exponent $a > 0$:

$$|m^\alpha[f_{(K)}] - x^*| \leq \frac{\sqrt{\varepsilon} C_{\lambda,\sigma}}{2} + \frac{\sqrt{\mathrm{d}} \exp(-\alpha q_\varepsilon + aT^*)}{\int \phi_{r_\varepsilon}(x) d f_{(0)}(x)} \sqrt{\mathrm{Err}[f_{(0)}]}.$$

Now, we note that α_ε was chosen exactly to bound the second term on the right-hand side above. In particular, for any $\alpha \geq \alpha_\varepsilon$ it holds

$$\frac{\sqrt{\mathrm{d}} \exp(-\alpha q_\varepsilon + aT^*)}{\int \phi_{r_\varepsilon}(x) d f_{(0)}(x)} \sqrt{\mathrm{Err}[f_{(0)}]} \leq \frac{C_{\lambda,\sigma}}{2} \sqrt{\varepsilon}.$$

Altogether, since the considered case assumes $\varepsilon \leq \mathrm{Err}[f_{(K)}]$, it holds

$$|m^\alpha[f_{(K)}] - x^*| < C_{\lambda,\sigma} \sqrt{\varepsilon} \leq C_{\lambda,\sigma} \sqrt{\mathrm{Err}[f_{(K)}]},$$

which is the desired contradiction.

□

3.2.2 CBO with memory effects

We saw how one of the key requirements in the application of the Laplace principle is ensuring that the mass around the minimizer does not vanish completely. The proof has required a strict condition on the step-size τ and the use of rather technical arguments. This may sound paradoxical as we expect, on the contrary, dynamics where particles do not leave promising areas of the search space.

Such expected behavior can be prescribed by including memory effects in the CBO dynamics, as illustrated in Section 2.4. This is also interesting for modeling purposes, as it allows to recover a key mechanism of PSO methods. We recall here the corresponding mean-field dynamics of the particle $\overline{X}_{(k)}$ and its personal best $\overline{Y}_{(k)}$

$$\begin{aligned} \overline{X}_{(k+1)} &= \overline{X}_{(k)} + \lambda\tau\left(m^\alpha[f^y_{(k)}] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha[f^y_{(k)}] - \overline{X}_{(k)}\right) \odot \theta_{(k)}, \\ \overline{Y}_{(k+1)} &= \begin{cases} \overline{X}_{(k+1)} & \text{if } \mathcal{E}(\overline{X}_{(k+1)}) < \mathcal{E}(\overline{Y}_{(k)}), \\ \overline{Y}_{(k)} & \text{else}, \end{cases} \end{aligned} \tag{3.11}$$

together with its weak formulation in terms of $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)}), \overline{Y}_{(k)}) \in \mathcal{P}(\mathbb{R}^\mathrm{d} \times \mathbb{R}^\mathrm{d})$

$$\begin{aligned} &\int \phi(x,y)\, d f_{(k+1)}(x,y) \\ &\qquad = \iint \left(\phi(x',x')\mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} + \phi(x',y)\mathbf{1}_{\{\mathcal{E}(x')\geq\mathcal{E}(y)\}}\right) d f_{(k)}(x,y) d\mu(\theta) \end{aligned} \tag{3.12}$$

with $x' = x + \lambda\tau(m^\alpha[f^y_{(k)}] - x) + \sigma\sqrt{\tau}(m^\alpha[f^y_{(k)}] - x) \odot \theta = \mathcal{C}^{CBO}(x, f^y_{(k)}, \theta)$, as before.

Thanks to the personal bests, it becomes easier to provide a lower bound on the mass around the minimizer x^*.

Lemma 3.7. *Assume $\mathcal{E}$ to be continuous and to attain a unique global minimum x^*. At every step k and for any radius $r > 0$, there exists $r_0 \in (0, r]$ such that*

$$f^y_{(k)}(B^\infty_r(x^*)) \geq f^y_{(0)}\left(B^\infty_{r_0}(x^*)\right).$$

Proof. For $\delta > 0$, we consider the objective function level set

$$L_\delta := \{x : \mathcal{E}(x) \leq \mathcal{E}(x^*) + \delta\}.$$

As we would expect due to the personal best mechanism, the probability mass of every level set is non-decreasing during the computation. Indeed, by considering the indicator function on the level set $\phi(x, y) = \mathbf{1}_{\{\mathcal{E}(y) \leq \mathcal{E}(x^*)+\delta\}}$ in (3.12) we have

$$\begin{aligned} f^y_{(k+1)}(L_\delta) &= \int \mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\, df_{(k+1)}(x) \\ &= \iint \left(\mathbf{1}_{\{\mathcal{E}(x')\leq\mathcal{E}(x^*)+\delta\}}\mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} + \mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\mathbf{1}_{\{\mathcal{E}(x')\geq\mathcal{E}(y)\}}\right) df_{(k)}(x, y)d\mu(\theta). \end{aligned}$$

Thanks to the inclusion

$$\{\mathcal{E}(y) \leq \mathcal{E}(x^*) + \delta\} \cap \{\mathcal{E}(x') < \mathcal{E}(y)\} \subseteq \{\mathcal{E}(x') \leq \mathcal{E}(x^*) + \delta\} \cap \{\mathcal{E}(x') < \mathcal{E}(y)\},$$

we can bound the above as

$$\begin{aligned} f^y_{(k+1)}(L_\delta) &\geq \iint \left(\mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} + \mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\mathbf{1}_{\{\mathcal{E}(x')\geq\mathcal{E}(y)\}}\right) df_{(k)}(x, y)d\mu(\theta) \\ &= \iint \left(\mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\left(\mathbf{1}_{\{\mathcal{E}(x')<\mathcal{E}(y)\}} + \mathbf{1}_{\{\mathcal{E}(x')\geq\mathcal{E}(y)\}}\right)\right) df_{(k)}(x, y)d\mu(\theta) \\ &= \iint \left(\mathbf{1}_{\{\mathcal{E}(y)\leq\mathcal{E}(x^*)+\delta\}}\right) df_{(k)}(x, y)d\mu(\theta) = f^y_{(k)}(L_\delta). \end{aligned}$$

By continuity of $\mathcal{E}$ and by uniqueness of the global minimizer, for all $r > 0$ there exists $\delta > 0$ such that $L_\delta \subseteq B^\infty_r(x^*)$. Similarly, for any $\delta > 0$ there exists a $r_0 > 0$ such that $B^\infty_{r_0}(x^*) \subseteq L_\delta$, leading to

$$B^\infty_{r_0}(x^*) \subseteq L_\delta \subseteq B^\infty_r(x^*).$$

Therefore, thanks to preservation of mass property of the level sets, we can conclude by noting that

$$f^y_{(k)}(B^\infty_r(x^*)) \geq f^y_{(k)}(L_\delta) \geq f^y_{(0)}(L_\delta) \geq f^y_{(0)}\left(B^\infty_{r_0}(x^*)\right).$$

□

As for the case without personal bests, we study convergence by looking at the expected squared error (3.2) of the mean-field particle location $\overline{X}_{(k)}$, that is, $\mathrm{Err}[f^x_{(k)}]$. We note that the estimates we provided in Proposition 3.4 still hold when the consensus point is computed among the personal bests. Therefore, we have

$$\mathrm{Err}[f^x_{(k+1)}] \leq \left(1 - \tau \frac{2\lambda - \sigma^2 - \lambda^2\tau}{2}\right)\mathrm{Err}[f^x_{(k)}] + \tau \frac{2\lambda + \sigma^2 + \lambda^2\tau}{2}|m^\alpha[f^y_{(k)}] - x^*|^2 . \tag{3.13}$$

Next, we briefly adapt the convergence proof for CBO with memory effects

Theorem 3.8. *Fix an accuracy $\varepsilon > 0$ and assume the objective function $\mathcal{E}$ satisfies Assumptions 2.2 and 3.1. Consider an initial distribution $f^x_{(0)} \in \mathcal{P}_2(\mathbb{R}^\mathrm{d})$ such that $x^* \in \mathrm{supp}(f^x_{(0)})$ and parameters $\{\lambda, \sigma, \tau\} \subset \mathbb{R}_+$ satisfying $2\lambda - \sigma^2 - \lambda^2\tau > 0$. Define the time horizon $T^* > 0$*

$$T^* = \frac{4}{2\lambda - \sigma^2 - \lambda^2\tau} \log\left(\frac{\mathrm{Err}[f^x_{(0)}]}{\varepsilon}\right).$$

Let $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)}, \overline{Y}_{(k)})$ with $\overline{X}_{(k)}, \overline{Y}_{(k)}$ updated according to (3.11). Provided α is sufficiently large, it holds

$$\min_{k\,:\,k\tau \leq T^*} \mathrm{Err}[f^x_{(k)}] \leq \varepsilon . \tag{3.14}$$

Moreover, as long as $\mathrm{Err}[f^x_{(k-1)}] > \varepsilon$, it holds

$$\mathrm{Err}[f^x_{(k)}] \leq \exp\left(-k\tau \frac{2\lambda - \sigma^2 - \lambda^2\tau}{4}\right)\mathrm{Err}[f^x_{(0)}]. \tag{3.15}$$

Proof. We start by defining suitable α_ε such that for all $\alpha \geq \alpha_\varepsilon$ we have convergence in the sense of (3.14). Let $C_{\lambda,\sigma,\tau}$ be given by

$$C_{\lambda,\sigma,\tau} := \sqrt{\frac{2\lambda - \sigma^2 - \lambda^2\tau}{2(2\lambda + \sigma^2 + \lambda^2\tau)}}.$$

We start by applying the quantitative Laplace principle (Proposition 3.2) with $f = f^y_{(k)}$ and pick, as in the proof of Theorem 3.6, q_ε and r_ε given by

$$q_\varepsilon := \frac{1}{2}\min\left\{\frac{1}{c_1\sqrt{\mathrm{d}}}\left(\sqrt{\varepsilon}C_{\lambda,\sigma,\tau}\right)^{p_1}, \mathcal{E}_\infty\right\}, \quad r_\varepsilon := \max_{s\in[0,R]}\left\{\max_{x\in B^\infty_s(x^*)} \mathcal{E}(x) \leq q_\varepsilon\right\}.$$

We obtain

$$\begin{aligned}|m^\alpha[f^y_{(k)}] - x^*| &\leq c_1\sqrt{\mathrm{d}}(q_\varepsilon + \mathcal{E}_{r_\varepsilon})^{1/p_1} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon)}{f^y_{(k)}(B^\infty_{r_\varepsilon}(x^*))}\int |y - x^*| d f^y_{(k)}(y) \\ &\leq \frac{\sqrt{\varepsilon}C_{\lambda,\sigma,\tau}}{2} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon)}{f^y_{(k)}(B^\infty_{r_\varepsilon}(x^*))}\int |y - x^*| d f^y_{(k)}(y). \end{aligned} \tag{3.16}$$

Next, we recall from Lemma 3.7 that there exists $r_{\varepsilon,0} > 0$ such that $f^y_{(k)}(B^\infty_{r_\varepsilon}(x^*)) \geq f^y_{(0)}(B^\infty_{r_{\varepsilon,0}}(x^*))$. Thanks to Jensen's inequality and the bound (2.15) for the second moments derived in Proposition 2.7, it also holds

$$\begin{aligned}\int |y - x^*| d f^y_{(k)}(y) &\leq |x^*| + \int |y| d f^y_{(k)}(y) \\ &\leq |x^*| + \left(\int |y|^2 d f^y_{(k)}(y) \right)^{1/2} \\ &\leq |x^*| + R_l + \frac{1}{c_l} \left(\int (\mathcal{E}(y) - \mathcal{E}(x^*)) d f^y_{(0)}(y) \right)^{1/2} =: C_{\mathcal{E},0}\,.\end{aligned}$$

Therefore, from (3.16) we obtain

$$|m^\alpha[f^y_{(k)}] - x^*| \leq \frac{\sqrt{\varepsilon} C_{\lambda,\sigma,\tau}}{2} + \frac{\sqrt{d}\exp(-\alpha q_\varepsilon)}{f^y_{(0)}(B^\infty_{r_{\varepsilon,0}}(x^*))} C_{\mathcal{E},0}\,.$$

We can now pick α_ε large enough such for all $\alpha \geq \alpha_\varepsilon$ it holds

$$|m^\alpha[f^y_{(k)}] - x^*| < C_{\lambda,\sigma,\tau}\sqrt{\varepsilon}\,.$$

We remark that this choice of α_ε is independent on k and on time horizon T^*. This is different from the proof of the analogous convergence result for CBO without memory effects.

Let us couple the above estimate with the error evolution (3.13) to note that, as long as $\mathrm{Err}[f^x_{(k)}] > \varepsilon$, the error decays exponentially:

$$\begin{aligned}\mathrm{Err}[f^x_{(k+1)}] &\leq \left(1 - \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{2}\right)\mathrm{Err}[f^x_{(k)}] + \tau\frac{2\lambda + \sigma^2 + \lambda^2\tau}{2}|m^\alpha[f^y_{(k)}] - x^*|^2 \\ &\leq \left(1 - \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{2}\right)\mathrm{Err}[f^x_{(k)}] + \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{4}\varepsilon \\ &\leq \left(1 - \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{4}\right)\mathrm{Err}[f^x_{(k)}] \\ &\leq \exp\left(-\frac{2\lambda - \sigma^2 - \lambda^2\tau}{4}k\tau\right)\mathrm{Err}[f^x_{(0)}],\end{aligned}$$

Since the upper bound is smaller than ε for $k\tau > T^*$, the result follows. □

Corollary 3.9. *Under the settings of Theorem 3.8, assume the parameter α is increased during the computation such that $\alpha_{(k)} \to \infty$ as $k \to \infty$. It holds*

$$\inf_{k\geq 0} \mathrm{Err}[f_{(k)}] = 0\,.$$

Proof. As noted, the choice of α_ε in the proof of Theorem 3.8 is independent on k and T^*. Therefore, for all $\varepsilon > 0$ there exists k large enough such that $\alpha_{(k)} \geq \alpha_\varepsilon$. This leads to

$$\inf_{k\geq 0} \mathrm{Err}[f_{(k)}] \leq \min_{k, k\tau<T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon\,.$$

□

3.3 Mean-field approximation error for bounded domains

In this section, we consider CBO dynamics where the particles are constrained into a set $\mathcal{D} \subset \mathbb{R}^{\mathrm{d}}$, which is assumed to be closed, bounded, and convex. For such constrained dynamics, we provide quantitative error estimates of the mean-field approximation. We also couple this result with the mean-field convergence analysis to prove convergence of the CBO algorithm with a finite number N of particles.

Let $\Pi_{\mathcal{D}} : \mathbb{R}^{\mathrm{d}} \to \mathcal{D}$ be the projection operator into $\mathcal{D}$

$$\Pi_{\mathcal{D}}(x) = \operatorname*{argmin}_{z\in\mathcal{D}} |z - x|^2\,.$$

The projection is well-defined because the minimum is attained uniquely due to the convexity of $\mathcal{D}$. The projected CBO dynamics (without memory effects) simply adds a projection step after the particles update to ensure that the system remains confined over $\mathcal{D}$. The update rule reads for all $i = 1, \dots, N$,

$$X^i_{(k+1)} = \Pi_{\mathcal{D}}\Big(X^i_{(k)} + \lambda\tau\left(m^\alpha[f^N_{(k)}] - X^i_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha[f^N_{(k)}] - X^i_{(k)}\right)\odot\theta^i_{(k)}\Big)\,. \tag{3.17}$$

The corresponding mean-field difference equation is then given by

$$\overline{X}_{(k+1)} = \Pi_{\mathcal{D}}\Big(\overline{X}_{(k)} + \lambda\tau\left(m^\alpha[f_{(k)}] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha[f_{(k)}] - \overline{X}_{(k)}\right)\odot\overline{\theta}_{(k)}\Big) \tag{3.18}$$

with, as before, $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$ and $\overline{\theta}_{(k)} \sim \mathcal{N}(0, I_{\mathrm{d}})$.

Theorem 3.10. *Assume $\mathcal{E}$ satisfies Assumption 2.2. Let $\{X^i_{(k)}\}_{i=1}^N$ be given by* (3.17) *and consider N copies of the mean-field dynamics* (3.18) *$\{\overline{X}^i_{(k)}\}_{i=1}^N$ with same initial data $X^i_{(0)} = \overline{X}^i_{(0)}$ and random vectors $\theta^i_{(h)} = \overline{\theta}^i_{(h)}$ for all $h \leq k$.*

It holds

$$\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] \leq C_{\mathrm{MFA}} N^{-1}$$

with $C_{\mathrm{MFA}} = C_{\mathrm{MFA}}(k, \mathrm{diam}(\mathcal{D}), \alpha, \mathcal{E}, \lambda, \sigma, \tau)$.

Before providing a proof, we first recall some auxiliary results.

Lemma 3.11. *[16, Lemma 3.2] Let $\mathcal{E}$ satisfy Assumption 2.2 and $f, \hat{f} \in \mathcal{P}_2(\mathbb{R}^d)$ with*

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

Then the following estimates holds

$$|m^\alpha[f] - m^\alpha[\hat{f}]| \leq C_0 W_2(f, \hat{f}),$$

for a constant C_0 depending only on $\alpha, L_\mathcal{E}, K$.

Lemma 3.12. *Let $\mathcal{E}$ satisfy Assumption 2.2. Let $\{\overline{X}^i_{(k)}\}_{i=1}^N$, for $k = 0, 1, \ldots$, be N copies of the mean-field process (3.17), which are i.i.d. with common distribution $f_{(k)} \in \mathcal{P}(\mathcal{D})$. Then, there exists a constant C_1 depending only on $\operatorname{diam}(\mathcal{D})$ and $C_\alpha := \exp(\alpha(\sup_{x\in\mathcal{D}} \mathcal{E} - \inf_{x\in\mathcal{D}} \mathcal{E}))$ such that*

$$\sup_{k\in\mathbb{Z}_+} \mathbb{E}\left[|m^\alpha[\overline{f}^N_{(k)}] - m^\alpha[f_{(k)}]|^2\right] \leq C_1 N^{-1}.$$

where $\overline{f}^N_{(k)}$ is the empirical random measure given by $\overline{f}^N_{(k)} := (1/N)\sum_{i=1}^N \delta_{\overline{X}^i_{(k)}}$.

Thanks to the boundedness of $\mathcal{D}$, a proof can be given by following step-by-step the proof of Lemma 3.1 in [40] for CBO mean-field dynamics over hypersurfaces in time-continuous settings.

Proof of Theorem 3.10. For the sake of notational simplicity, let $X^i_{(k+1/2)}, \overline{X}^i_{(k+1/2)}$ be the random variables before the projection step in (3.17) and (3.18), respectively. We note that $\Pi_\mathcal{D}$ is non-expansive thanks to the convexity of $\mathcal{D}$. Namely, for any $x, y \in \mathbb{R}^d$, it holds

$$|\Pi_\mathcal{D}(x) - \Pi_\mathcal{D}(y)| \leq |x - y|.$$

Therefore, we have

$$|X^i_{(k+1)} - \overline{X}^i_{(k+1)}|^2 = \left|\Pi_\mathcal{D}\left(X^i_{(k+1/2)}\right) - \Pi_\mathcal{D}\left(\overline{X}^i_{(k+1/2)}\right)\right|^2 \leq |X^i_{(k+1/2)} - \overline{X}^i_{(k+1/2)}|^2$$

for all $i = 1, \ldots, N$. Next, we bound the expected value as following

$$\begin{aligned}
\mathbb{E}\left[|X^i_{(k+1/2)} - \overline{X}^i_{(k+1/2)}|^2\right] &\leq 2\mathbb{E}\left[|(1 - \lambda\tau - \sigma\sqrt{\tau}\theta^i_{(k)}) \odot (X^i_{(k)} - \overline{X}^i_{(k)})|^2\right] \\
&\quad + 2\mathbb{E}\left[|(\lambda\tau + \sigma\sqrt{\tau}\theta^i_{(k)}) \odot (m^\alpha[f^N_{(k)}] - m^\alpha[f_{(k)}])|^2\right] \\
&= 2\left(1 - \tau(2\lambda - \sigma^2 - \lambda^2\tau)\right)\mathbb{E}\left[|X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] \\
&\quad + 2\tau(\lambda^2\tau + \sigma^2)\mathbb{E}\left[|m^\alpha[f^N_{(k)}] - m^\alpha[f_{(k)}]|^2\right],
\end{aligned}$$

where we used Lemma 2.5 to compute the expected values with respect to $\theta^i_{(k)}$. To bound the second term, we consider $\overline{f}^N_{(k)} = (1/N)\sum_{i=1}^N \delta_{\overline{X}^i_{(k)}}$ and apply the triangular-like inequality

$$|m^\alpha[f^N_{(k)}] - m^\alpha[f_{(k)}]|^2 \le 2|m^\alpha[f^N_{(k)}] - m^\alpha[\overline{f}^N_{(k)}]|^2 + 2|m^\alpha[\overline{f}^N_{(k)}] - m^\alpha[f_{(k)}]|^2. \tag{3.19}$$

Thanks to Assumption 2.2, we are able to apply Lemma 3.11 with $K = \mathrm{diam}(\mathcal{D})^4$ to get for some constant $C_0 = C_0(\mathrm{diam}(\mathcal{D}), \alpha, L_{\mathcal{E}}) > 0$

$$|m^\alpha[f^N_{(k)}] - m^\alpha[\overline{f}^N_{(k)}]|^2 \le C_0^2 W_2^2\left(f^N_{(k)}, \overline{f}^N_{(k)}\right) \le \frac{C_0^2}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2.$$

For the second term in (3.19), we apply Lemma 3.12 to obtain

$$\mathbb{E}\left[|m^\alpha[\overline{f}^N_{(k)}] - m^\alpha[f_{(k)}]|^2\right] \le C_1 N^{-1}$$

for some constant $C_1 > 0$ depending on $M = \mathrm{diam}(\mathcal{D})$ and $C_\alpha = \exp(\alpha(\sup\mathcal{E} - \inf\mathcal{E}))$, but independent on k.

We collect the computed estimates and sum over all particles $i = 1, \dots, N$ to obtain

$$\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k+1)} - \overline{X}^i_{(k+1)}|^2\right] \le 2(1+\tau C_2)\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] + \tau C_1 C_3 N^{-1}$$

for some constants C_2, C_3 depending on λ, σ, τ as well as on $\mathrm{diam}(\mathcal{D}), \alpha, L_{\mathcal{E}}$. By iterating the argument for all $h = 1, \dots, k$ one obtains

$$\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] \le 2^k(1+\tau C_2)^k \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(0)} - \overline{X}^i_{(0)}|^2\right] + \frac{\tau C_1 C_3}{N}\sum_{h=1}^k 2^h(1+\tau C_2)^h.$$

Thanks to assumption $X^i_{(0)} = \overline{X}^i_{(0)}$ for all $i = 1, \dots, N$ we conclude that

$$\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(0)} - \overline{X}^i_{(k)}|^2\right] \le C_{\mathrm{MFA}} N^{-1}$$

for some positive constant $C_{\mathrm{MFA}} = C_{\mathrm{MFA}}(k, C_1, C_2, C_3)$. □

Theorem 3.13. *Assume the objective $\mathcal{E}$ satisfies Assumption 3.1 with the unique global minimum belonging to the bounded search space, $x^* \in \mathcal{D}$.*

Statements of Theorem 3.6 also hold for the projected mean-field dynamics (3.18).

Proof. We recall that the key ingredients to prove convergence are given by:

- estimates on evolution of $\mathrm{Err}[f_{(k)}]$ (Proposition 3.4);
- lower bound on probability mass around the minimizer x^* (Lemma 3.5);
- application of the quantitative Laplace principle (Proposition 3.2).

First, we note that the Laplace principle holds independently on the projection strategy if $x^* \in \mathcal{D}$. Therefore, we only need to ensure that estimates provided in Proposition 3.4 and Lemma 3.5 still hold for the projected dynamics. As we will see, this is the case thanks to the non-expansive property of the projector operator $\Pi_{\mathcal{D}}$.

As above, let $\overline{X}_{(k+1/2)}$ be the particle before the projection step, that is, the argument of $\Pi_{\mathcal{D}}$ in (3.18), and let $f_{(k+1/2)}$ be its law. We note that the same estimates we derived for the non-projected dynamics still hold for $f_{(k+1/2)}$, which is the law before the projection step. Consider the expected mean squared error $\mathrm{Err}[f_{(k+1)}]$ after the projection step. Since $x^* = \Pi_{\mathcal{D}}(x^*)$ it holds

$$\mathrm{Err}[f_{(k+1)}] = \int |\Pi_{\mathcal{D}}(x') - x^*|^2 \, d f_{(k)}(x) \leq \int |x' - x^*|^2 \, d f_{(k)}(x) = \mathrm{Err}[f_{(k+1/2)}].$$

Therefore, the bound on the error evolution given by Proposition 3.4 still holds when projection is performed.

Similarly, when looking at the mass around the global minimizer, $f_{(k+1)}(B_r^\infty(x^*))$, due to $x^* \in \mathcal{D}$ it holds

$$\Pi_{\mathcal{D}}(B_r^\infty(x^*)) \subseteq B_r^\infty(x^*),$$

from which follows $B_r^\infty(x^*) \subseteq \Pi_{\mathcal{D}}^{-1}\left(B_r^\infty(x^*)\right)$. This leads to

$$f_{(k+1)}\left(B_r^\infty(x^*)\right) = f_{(k+1/2)}\left(\Pi_{\mathcal{D}}^{-1}\left(B_r^\infty(x^*)\right)\right) \geq f_{(k+1/2)}\left(B_r^\infty(x^*)\right).$$

Given that we have a lower bound for $f_{(k+1/2)}\left(B_r^\infty(x^*)\right)$ thanks to Lemma 3.5, the same estimates hold for the projected CBO dynamics. □

Now that we have proved quantitative error estimates of the mean-field approximation and convergence of the mean-field dynamics, we can couple the two results to obtain an error estimate for the N particle system.

Theorem 3.14. *Let the objective function $\mathcal{E}$ satisfy Assumptions 2.2 and 3.1 with global minimum $x^* \in \mathcal{D}$. Consider a particle system $\{X_{(k)}^i\}_{i=1}^N$ updated according to the projected CBO dynamics* (3.17) *with $X_0^i \sim f_{(0)}$ such that $x^* \in \mathrm{supp}(f_{(0)})$. Let the parameters $\{\tau, \lambda, \sigma\} \subset \mathbb{R}_+$ satisfy $\lambda > \sigma^2$.*

Fix an accuracy $\varepsilon > 0$ and define the time horizon $T^* > 0$

$$T^* = \frac{4}{\lambda - \sigma^2} \log\left(\frac{2\mathrm{Err}[f_{(0)}]}{\varepsilon}\right).$$

Provided α is sufficiently large and the step-size τ is sufficiently small, it holds

$$\min_{k:k\tau\le T^*} \mathrm{Err}[f^N_{(k)}] \le C_{\mathrm{MFA}} N^{-1} + \varepsilon\,, \tag{3.20}$$

with $C_{\mathrm{MFA}} = C_{\mathrm{MFA}}(T^*/\tau, \mathrm{diam}(\mathcal{D}), \alpha, \mathcal{E}, \lambda, \sigma, \tau)$.

Proof. Let $\{\overline{X}^i_{(k)}\}_{i=1}^N$ be independent mean-field processes evolving according to (3.18) with $X^i_{(0)} = \overline{X}^i_{(0)}$ and $\theta^i_{(k)} = \overline{\theta}^i_{(k)}$ for all $k \ge 1$, $i = 1,\dots,N$. We decompose the mean squared error of the particle system as

$$\begin{aligned}\mathrm{Err}[f^N_{(k)}] &= \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - x^*|^2\right] \\ &\le 2\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] + 2\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |\overline{X}^i_{(k)} - x^*|^2\right].\end{aligned} \tag{3.21}$$

Given that the mean-field processes are independent we have

$$\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |\overline{X}^i_{(k)} - x^*|^2\right] = \mathrm{Err}[f_{(k)}],$$

with $f_{(k)}$ being their common law. By applying Theorem 3.13 with accuracy $\varepsilon/2$ we get

$$\min_{k:k\tau\le T^*} \mathrm{Err}[f^N_{(k)}] \le \varepsilon/2\,,$$

provided α is sufficiently large and τ is sufficiently small. Next, we apply Theorem 3.10 to bound the first term as

$$2\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N |X^i_{(k)} - \overline{X}^i_{(k)}|^2\right] \le C_{\mathrm{MFA}} N^{-1}$$

for some $C_{\mathrm{MFA}} = C_{\mathrm{MFA}}(T^*/\tau, \mathrm{diam}(\mathcal{D}), \alpha, \mathcal{E}, \lambda, \sigma, \tau)$. By plugging the collected estimates in (3.21) we obtain (3.20). □

Remark 3.15. *One could wonder if a similar mean-field error estimates holds when projection is applied in presence of memory effects. To follow the proof strategy of Theorem 3.10, we should be able to provide stability estimates for the dynamics of the personal bests. This is a non-trivial task as the update rule for* $\{Y^i_{(k)}\}_{i=1}^N$ *and* $\{\overline{Y}^i_{(k)}\}_{i=1}^N$ *is non-Lipschitz, see, for instance,* (3.11).

A workaround suggested in [43] consist of relaxing the memory mechanism as following. Let $S^\beta(x,y)$ *be a Lipschitz approximation of* $\mathbf{1}_{\{\mathcal{E}(x)<\mathcal{E}(y)\}}$ *as* $\beta \to \infty$ *with* $\mathrm{Im}(S^\beta) \subseteq [0,1]$ *and* $\nu > 0$ *be an additional parameter satisfying* $\nu\tau \in [0,1]$. *The relaxed personal bests update reads*

$$Y^i_{(k+1)} = Y^i_{(k)} + \nu\tau S^\beta(X^i_{(k+1)}, Y^i_{(k)})(X^i_{(k+1)} - Y^i_{(k)}). \tag{3.22}$$

Due to the convexity of $\mathcal{D}$, no projection is needed in the above update rule as $Y^i_{(k+1)}$ is a convex combination of random variables taking values in $\mathcal{D}$. We note that as $\beta \to \infty$ and for $\nu = 1/\tau$, the above dynamics converges to the original memory mechanism (2.12)*. Mean-field approximation errors for* (3.22) *can than be derived following the steps of Theorem 3.10. Convergence towards the global solution x^* in mean-field law for* (3.22) *can also be proved under stricter assumptions on the objective function. This is done in [7] by assuming local convexity of $\mathcal{E}$ around x^* and in [74] by adding a stochastic component to the update rule* (3.22)*.*

3.4 Discussion

We have seen in this chapter that the quantitative Laplace principle is a powerful tool to study CBO dynamics. In mean-field law, this result allows to virtually place $m^\alpha[f_{(k)}]$ as close as one wants to the global minimizer x^* by only assuming large values of α. Considering then an iterative update, rather than time-continuous dynamics, does not make much of a difference when it comes to tracking the evolution of the error, or of the moments.

Assuming boundedness of the search space further allows us to provide quantitative estimates of the mean-field approximation error in terms of number of particles N. We note that such approximation error can be provided with high probability also in the case of unbounded search domains, see [41].

Finally, it is important to remark that the above convergence analysis only focuses on the *expected* outcome of the CBO algorithm. By definition of $\mathrm{Err}[f]$, indeed, we are averaging the error among all the possible realizations. One could wonder, how far does a single CBO run deviate from the expected performance? While many numerical experiments show that CBO algorithms are reliable in this sense, a theoretical answer to this question is still missing.

3.5 Proof of Lemma 3.5

Proof of Lemma 3.5. Taylor's expansion around x' with Lagrange remainder, gives us for some $\xi \in B_r^\infty(x^*)$

$$\begin{aligned}\phi_r(x') - \phi_r(x) &= \nabla\phi_r(x)(x'-x) + \frac{1}{2}(x'-x)^\top \nabla^2\phi_r(x)(x'-x) \\ &\quad + \sum_{\ell,i,j=1}^{\mathrm{d}} \frac{\partial^3_{\ell ij}\phi_r(\xi)}{3!}(x'-x)_\ell (x'-x)_i (x'-x)_j \\ &=: T^{(1)}(x',x) + T^{(2)}(x',x) + T^{(3)}(x',x).\end{aligned}$$

To bound $T^{(1)}, T^{(2)}, T^{(3)}$ we exploit computations done for the time-continuous dynamics for the terms of order $\mathcal{O}(\tau)$, while we derive new bounds for the higher order terms.

Let $T_\ell^{(1)}(x',x) := \partial_\ell \phi_r(x)(x'-x)_\ell$ such that $T^{(1)} = \sum_{\ell=1}^{\mathrm{d}} T_\ell^{(1)}$. We take the expectation with respect to θ_ℓ to obtain

$$\begin{aligned}\mathbb{E}[T_{1\ell}(x',x)] &= \mathbb{E}[\partial_\ell \phi_r(x)(x'-x)_\ell] \\ &= \mathbb{E}\Big[\partial_\ell \phi_r(x)(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)(m^\alpha[f_{(k)}]-x)_\ell\Big] \\ &= \lambda\tau(m^\alpha[f_{(k)}]-x)_\ell \partial_\ell \phi_r(x),\end{aligned}$$

since $\mathbb{E}[\theta_\ell] = 0$ for all $\ell = 1,\dots,\mathrm{d}$.

Now, let $T^{(2)} = \sum_{\ell,i=1}^{\mathrm{d}} T_{\ell,i}^{(2)}$ with $T_{\ell,i}^{(2)}(x',x) = (1/2)\partial_{\ell i}^2 \phi_r(x)(x'-x)_\ell (x'-x)_i$. For the diagonal terms it holds

$$\begin{aligned}\mathbb{E}\Big[T_{\ell\ell}^{(2)}(x',x)\Big] &= \mathbb{E}\left[\frac{1}{2}\partial_{\ell\ell}^2 \phi_r(x)(x'-x)_\ell^2\right] \\ &= \mathbb{E}\left[\frac{1}{2}\partial_{\ell\ell}^2 \phi_r(x)(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)^2(m^\alpha[f_{(k)}]-x)_\ell^2\right] \\ &= \frac{1}{2}\partial_{\ell\ell}^2 \phi_r(x)(\lambda^2\tau^2 + \sigma^2\tau)(m^\alpha[f_{(k)}]-x)_\ell^2\end{aligned}$$

given that

$$\mathbb{E}[(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)^2] = \mathbb{E}[\lambda^2\tau^2 + 2\lambda\tau\sigma\sqrt{\tau}\theta_\ell + \sigma^2\tau\theta_\ell^2] = \lambda^2\tau^2 + \sigma^2\tau$$

due to $\mathbb{E}[\theta_\ell] = 0$, $\mathbb{E}[\theta_\ell^2] = 1$. Non-diagonal terms, instead, are purely of order $\mathcal{O}(\tau^2)$:

$$\begin{aligned}\mathbb{E}\Big[T_{\ell i}^{(2)}(x',x)\Big] &= \mathbb{E}\left[\frac{1}{2}\partial_{\ell i}^2 \phi_r(x)(x'-x)_\ell (x'-x)_i\right] \\ &= \mathbb{E}\left[\frac{1}{2}\partial_{\ell i}^2 \phi_r(x)(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)(\lambda\tau + \sigma\sqrt{\tau}\theta_i)(m^\alpha[f_{(k)}]-x)_\ell (m^\alpha[f_{(k)}]-x)_i\right] \\ &= \frac{1}{2}\partial_{\ell i}^2 \phi_r(x)\lambda^2\tau^2(m^\alpha[f_{(k)}]-x)_\ell (m^\alpha[f_{(k)}]-x)_i,\end{aligned}$$

due to $\mathbb{E}[\theta_\ell] = 0$, and $\mathbb{E}[\theta_\ell\theta_i] = 0$ whenever $\ell \neq i$.

Before having a look at $T^{(3)}(x',x)$, we apply the estimates derived in [42, Proposition 2] for the terms of order $\mathcal{O}(\tau)$ of the expansion to obtain

$$\int \mathbb{E}\Big[T_\ell^{(1)}(x',x) + T_{\ell\ell}^{(2)}(x',x)\Big] df_{(k)}(x) \geq -\tau\frac{a}{\mathrm{d}}\int \phi_r(x) df_{(k)}(x) \tag{3.23}$$

where $a > 0$ is given by

$$\begin{aligned}&a = 2\mathrm{d}\max\{a_1 + a_2, a_3\} \\ &\text{with} \quad a_1 = \frac{\lambda(cr + B\sqrt{c})}{(1-c^2)r}, \quad a_2 = \frac{(\lambda^2\tau + \sigma^2)(cr^2 + B^2)(2c+1)}{(1-c)^4 r^2}, \quad a_3 = \frac{2\lambda^2}{(2c-1)\sigma^2}\end{aligned} \tag{3.24}$$

with $c \in (1/2, 1)$ being any constant satisfying $(1-c^2) \le (2c-1)c$. We note that, differently from [42, Proposition 2], the term $\lambda^2\tau$ appears in the definition of a_2, making the above definition dependent on τ. The decay rate a, though, can be made independent on τ by simply assuming $\tau \le 1$.

To bound the higher order terms, we note that

$$\phi_r(x) = \phi\left(\frac{x-x^*}{r}\right)$$

with $\phi \in C_0^\infty(\mathbb{R}^{\mathrm{d}})$ independent on r. Therefore, there exists $C_1 > 0$ independent on r such that

$$\partial^2_{\ell i}\phi_r(x) = \frac{1}{r^2}\partial_{\ell i}\phi\left(\frac{x-x^*}{r}\right) \ge -\frac{C_1}{r^2}$$
$$\partial^3_{\ell ij}\phi_r(x) = \frac{1}{r^3}\partial^3_{\ell ij}\phi\left(\frac{x-x^*}{r}\right) \ge -\frac{C_1}{r^3}$$

for all $x \in \mathbb{R}^{\mathrm{d}}$.

For $\ell \neq i$, therefore, we have

$$\begin{aligned}
\mathbb{E}\left[T^{(2)}_{\ell i}(x', x)\right] &= \frac{1}{2}\partial^2_{\ell i}\phi_r(x)\lambda^2\tau^2(m^\alpha[f_{(k)}] - x)_\ell(m^\alpha[f_{(k)}] - x)_i \\
&\ge -\frac{1}{2}\frac{C_1}{r^2}\lambda\tau^2|(m^\alpha[f_{(k)}] - x)_\ell||(m^\alpha[f_{(k)}] - x)_i| \\
&\ge -\frac{1}{2}\frac{C_1}{r^2}\lambda\tau^2(B+r)^2 \\
&\ge -C_2\frac{\tau^2}{\mathrm{d}(\mathrm{d}-1)r^2}
\end{aligned} \tag{3.25}$$

with $C_2 = C_2(\mathrm{d}, \lambda, B)$ for $r \le 1$. Similarly, for the remainder $T^{(3)} =: \sum_{\ell,i,j=1}^{\mathrm{d}} T^{(3)}_{\ell ij}$ we have

$$\begin{aligned}
\mathbb{E}[T^{(3)}_{\ell ij}(x', x)] &= \mathbb{E}\left[\frac{\partial^3_{\ell\ell\ell}\phi_r(\xi)}{3!}(x'-x)_\ell(x'-x)_i(x'-x)_j\right] \\
&\ge -\frac{C_1}{3!r^3}|(m^\alpha[f_{(k)}] - x)_\ell||(m^\alpha[f_{(k)}] - x)_i||(m^\alpha[f_{(k)}] - x)_j| \cdot \\
&\quad \cdot \mathbb{E}\left[|(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)(\lambda\tau + \sigma\sqrt{\tau}\theta_i)(\lambda\tau + \sigma\sqrt{\tau}\theta_j)|\right]
\end{aligned}$$

We bound the expectation term above as following

$$\begin{aligned}
&\mathbb{E}\left[|(\lambda\tau + \sigma\sqrt{\tau}\theta_\ell)(\lambda\tau + \sigma\sqrt{\tau}\theta_i)(\lambda\tau + \sigma\sqrt{\tau}\theta_j)|\right] \\
&\qquad \le \mathbb{E}\left[(\lambda\tau + \sigma\sqrt{\tau}|\theta_\ell|)(\lambda\tau + \sigma\sqrt{\tau}|\theta_i|)(\lambda\tau + \sigma\sqrt{\tau}|\theta_j|)\right] \\
&\qquad \le \lambda^3\tau^3 + 3\lambda\tau\sigma^2\tau + 3\lambda^2\tau^2\sigma\sqrt{\tau} + \sigma^3\tau\sqrt{\tau},
\end{aligned}$$

leading to

$$\mathbb{E}[T^{(3)}_{\ell ij}(x',x)] \geq -\frac{C_1}{3!r^3}(B+r)^3(\lambda^3\tau^3+3\lambda\tau\sigma^2\tau+3\lambda^2\tau^2\sigma\sqrt{\tau}+\sigma^3\tau\sqrt{\tau})$$
$$\geq -C_3\frac{\tau\sqrt{\tau}}{\mathrm{d}^3 r^3} \tag{3.26}$$

for some constants C_3, with $C_3 = C_3(\mathrm{d},\lambda,\sigma,B)$ independent of τ for $\tau, r \leq 1$.

By putting together estimates (3.23), (3.25), and (3.26), we obtain

$$\begin{aligned}
\int \phi_r(x) df_{(k+1)}(x) - \int \phi_r(x) df_{(k)}(x) &= \int \mathbb{E}\Big[T^{(1)}(x',x)+T^{(2)}(x',x)+T^{(3)}(x',x)\Big] df_{(k)}(x) \\
&= \int \sum_{\ell=1}^{\mathrm{d}} \mathbb{E}\Big[T^{(1)}_{\ell}(x',x)+T^{(2)}_{\ell\ell}(x',x)\Big] df_{(k)}(x) \\
&\quad + \int \mathbb{E}\left[\sum_{\ell,i=1,\ell\neq i}^{\mathrm{d}} T^{(2)}_{\ell i}(x'-x) + \sum_{\ell,i,j=1}^{\mathrm{d}} T^{(3)}_{\ell ij}(x',x)\right] df_{(k)}(x) \\
&\geq -\tau a \int \phi_r(x) df_{(k)}(x) - \frac{\tau^2}{r^2}C_2 - \frac{\tau\sqrt{\tau}}{r^3}C_3 \\
&\geq -\tau a \int \phi_r(x) df_{(k)}(x) - \frac{\tau\sqrt{\tau}}{r^3}C
\end{aligned}$$

for some $C = C(\mathrm{d},\lambda,\sigma,B)$ independent of τ, r for $\tau, r \leq 1$.

Now, we can iterate the above estimate for all $h = 1,\dots,k$ to get

$$\begin{aligned}
\int \phi_r(x) f_{(k)}(x) &\geq (1-\tau a)\int \phi_r(x) df_{(k-1)}(x) - \frac{\tau\sqrt{\tau}}{r^3}C \\
&\geq (1-\tau a)^k \int \phi_r(x) df_{(0)}(x) - \frac{\tau\sqrt{\tau}}{r^3}C\sum_{h=0}^{k-1}(1-\tau a)^{k-h} \\
&\geq (1-\tau a)^k \int \phi_r(x) df_{(0)}(x) - \frac{k\tau\sqrt{\tau}}{r^3}C.
\end{aligned}$$

Let us consider τ sufficiently small such that for all $k\tau < T$

$$\frac{k\tau\sqrt{\tau}}{r^3}C \leq \frac{1}{2}e^{-aT}\int \phi_r(x) df_{(0)}(x).$$

This leads to the desired lower bound

$$\begin{aligned}\int \phi_r(x) d f_{(k)}(x) &\geq (1-\tau a)^k \int \phi_r(x) d f_{(0)}(x) - \frac{k\tau\sqrt{\tau}}{r^3} C \\ &\geq e^{-ak\tau} \int \phi_r(x) d f_{(0)}(x) - \frac{1}{2} e^{-aT} \int \phi_r(x) d f_{(0)}(x) \\ &\geq \frac{1}{2} e^{-ak\tau} \int \phi_r(x) d f_{(0)}(x)\end{aligned}$$

where we used that $(1-\tau a)^k \leq e^{-ak\tau}$. □

Remark 3.16. *The fact that x^* is the global minimizer of the objective function $\mathcal{E}$ does not play a role in the above proof. Therefore, Lemma 3.5 holds for any point $\hat{x} \in \mathbb{R}^d$. In particular, this means that if $\hat{x} \in \operatorname{supp}(f_{(0)})$, then $\hat{x} \in \operatorname{supp}(f_{(k)})$ for all $k \geq 0$. This is an important property that we will use for the analysis of CBO algorithms in constrained and multi-objective settings.*

Chapter 4

Constrained optimization problems

Optimization problems in applications typically require solutions to satisfy certain admissibility criteria. Such criteria are mathematically modeled as equality or inequality constraints, leading to a problem formulation of the type

$$\begin{aligned} &\text{minimize} \quad \mathcal{E}(x) \\ &\text{subject to} \quad \begin{cases} g_\ell(x) \leq 0, & \ell = 1, \dots, m, \\ h_\jmath(x) = 0, & \jmath = 1, \dots, n, \end{cases} \end{aligned} \tag{4.1}$$

for an objective function $\mathcal{E} \in C(\mathbb{R}^d)$ and constraint functions $g_\ell, h_\jmath \in C(\mathbb{R}^d)$.

In this chapter, we discuss a variant of the Consensus-Based Optimization (CBO) algorithm which takes into account such additional requirements and ensures that the computed solution belongs to the feasible set

$$\mathcal{M} = \{x \in \mathbb{R}^d \,:\, g_\ell(x) \leq 0,\ h_\jmath(x) = 0,\ \forall \ell = 1, \dots, m,\ \forall \jmath = 1, \dots, n\}.$$

In Section 3.3 we already discussed a modification of the CBO dynamics where particles are constrained within a given convex search space through a projection step. Such simple mechanism can clearly be extended to any feasible set $\mathcal{M}$ for which a projection operator is well-posed and explicitly available, at least in a neighborhood of $\mathcal{M}$. Authors in [38–40] extended CBO methods to implicitly defined hyper-surfaces, with a focus on the case of the sphere which is particularly relevant for applications.

It may happen, though, that the projection map towards the feasible set $\mathcal{M}$ is not explicitly known, or it is difficult to compute. In these cases, therefore, it is not possible to constrain the dynamics via projection. A common approach to deal with this scenario is to allow the algorithm to search for solutions on the entire space $\mathbb{R}^d$ but to modify the objective function in way that unfeasible points are, loosely speaking, less likely to be picked by the method. This is known in the literature as *penalization* approach and it

allows to reformulate the original constrained optimization problem into an unconstrained one:

$$\text{minimize} \quad P_\beta(x) := \mathcal{E}(x) + \beta r(x), \tag{4.2}$$

where the added penalty term βr satisfies

$$\beta \in \mathbb{R}_+ \qquad \text{and} \qquad r(x) \begin{cases} = 0, & \text{if } x \in \mathcal{M}, \\ > 0, & \text{else}. \end{cases}$$

The new objective function P_β, therefore, coincides with $\mathcal{E}$ over the feasible set $\mathcal{M}$, and the mechanism is regulated via a positive parameter β. Choosing $r(x) = +\infty$ for all $x \notin \mathcal{M}$, for instance, makes the penalized problem (4.2) equivalent to (4.1), but this is typically of little use for designing algorithms. On the other hand, penalization terms which continuously depend on the constraint functions g_ℓ, h_j such as ℓ_p-penalization for some $p > 0$

$$r(x) = \sum_{\ell=1}^{m} (\max\{0, g_\ell(x)\})^p + \sum_{j=1}^{n} |h_j(x)|^p,$$

allow to preserve the continuity of the objective function and are commonly used by optimizing routines. Continuity, though, is paid at a price: when ℓ_p-penalization is used, the penalized problem (4.2) may not be equivalent to (4.1), in the sense that the respective set of global solutions may not coincide anymore.

Let us consider for instance a problem with smooth, convex objective and constraints. When ℓ_2-penalization is employed, it is well-known that the global solution to (4.2) does not satisfy for finite β the constraints if they are active, see, for instance, [81]. In this case, the solution to (4.2) converges to the true solution to (4.1) only asymptotically as $\beta \to \infty$. On the other hand, in these settings, the non-smooth ℓ_1-penalization ensures equivalence between the problems provided the penalization parameter β is sufficiently large: $\beta \geq \overline{\beta}$ for some finite $\overline{\beta} \in \mathbb{R}_+$ [81]. The threshold value $\overline{\beta}$ depends on the corresponding Lagrangian multiplier and, so, it is problem-dependent and unknown.

Penalization techniques which allow to reach equivalence between of the penalized problem for finite values of β, are called *exact*. In the following, we show how to couple the CBO dynamics with exact penalization in order to find the optimal penalization parameter $\overline{\beta}$ and, hence, to solve (4.1). We note that a CBO method which makes use of non-exact penalization has been proposed in [21].

The remaining part of the chapter is organized as follows. In Section 4.1, we present a CBO algorithm for exact penalization. Via mean-field approximation of the dynamics, we analyze the convergence properties of the proposed algorithm in Section 4.2. In Section 4.3, we illustrate the results of numerical experiments where the algorithm is tested against several constrained benchmark problems in different dimensions of the search space. Finally, Section 4.4 collects additional remarks.

4.1 Adaptive exact penalization strategy

In this section, we propose and algorithm which aims to solve the constrained optimization problem (4.1) by considering the penalized sub-problems (4.2). In particular, we assume the penalization strategy to be exact in the following sense.

Assumption 4.1. *The penalty term r is exact, that is, the set of global solutions to problems* (4.2) *and* (4.1) *coincide if β is sufficiently large. Moreover, we set*

$$\overline{\beta} := \inf\Big\{\beta > 0 \ : \ \text{problem (4.2) is equivalent to (4.1)}\Big\}.$$

Without loss of generality, we assume that for $\beta = \overline{\beta}$, the two problems are also not equivalent.

Exact penalization is a well-known strategy in constrained optimization and we refer to [6, 7, 15, 81] for more details on the topic. As already mentioned, ℓ_1-penalization is a common choice. As a further example, consider a constrained problem with $\mathcal{E}, g_\ell$ being twice differentiable and without equality constraints. Assume the Karush-Kuhn-Tucker (KKT) conditions hold at the global solution x^* with Lagrangian multiplier λ^*. If the weak second-order sufficient optimality conditions also hold, then ℓ_1-penalization is exact for $\beta \geq \|\lambda^*\|_\infty$ [7].

The simplest way to couple the CBO solver with exact penalization would consist in solving, iteratively, the penalty problem (4.2) for different values of β. If we increase β each time we solve the penalty problem, eventually we reach the threshold value $\overline{\beta}$ and obtain a feasible solution. This procedure therefore may require to solve many similar sub-problems before finding a feasible solution. To save computational time, we suggest an algorithmic strategy where the penalty parameter β is adapted during the computation depending on the constraint violation of the current particle system.

To do so, we introduce in the CBO dynamics the following evolving parameters:

- $\beta_{(k)} > 0$: penalty parameter which determines the objective function $P_{\beta_{(k)}}$;
- $\vartheta_{(k)} > 0$: parameter which determines the tolerance for the constraint violation;
- $\eta_\beta, \eta_\vartheta > 1$: update rate for $\beta_{(k)}$ and $\vartheta_{(k)}$, respectively.

As for the particle dynamics, we consider for simplicity the anisotropic CBO particle system

$$X^i_{(k+1)} = X^i_{(k)} + \lambda\tau\left(m^\alpha_{\beta_{(k)}}[f^N_{(k)}] - X^i_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha_{\beta_{(k)}}[f^N_{(k)}] - X^i_{(k)}\right)\odot\theta^i_{(k)}, \tag{4.3}$$

where, as in the previous chapters, $\tau, \lambda, \sigma > 0$ are fixed during the computation and $\theta^i_{(k)} \sim \mathcal{N}(0, I_\mathrm{d})$ is sampled from the standard normal distribution. We recall that the objective function enters the dynamics via the computation of the consensus point. As the objective $P_{\beta_{(k)}}$ now depends on $\beta_{(k)}$, we underlined this dependency in (4.3) by writing $m^\alpha_{\beta_{(k)}}[f^N_{(k)}]$.

To measure the constraint violation of the particle system, we introduce the functional $R : \mathcal{P}(\mathbb{R}^d) \to [0, \infty]$

$$R[f] := \int r(x) df(x), \tag{4.4}$$

where r is the exact penalty term given by Assumption 4.1.

Let $f^N_{(k)}$ be the empirical measure associated with the particle system (4.3). At each step k, we then check the following feasibility condition

$$R[f^N_{(k+1)}] \leq \frac{1}{\sqrt{\vartheta_{(k)}}}. \tag{4.5}$$

If the condition is satisfied, then the tolerance is decreased (by increasing $\vartheta_{(k)}$). If the condition is not satisfied, instead, we increase both the penalty parameter $\beta_{(k)}$ and $\vartheta_{(k)}$. The use of the square root in the feasibility check (4.5) is arbitrary and not necessary for the algorithm implementation, but it simplifies the computations in the theoretical analysis. This is because we will assume the constrained violation to be the distance from the feasible set (see Assumption 4.3 later), while the error will be given by the *squared* distance from the global minimizer.

The overall optimization method is described by Algorithm 1.

Set $\lambda, \sigma, \tau > 0, \eta_\beta, \eta_\vartheta > 1$;
Initialize $\{X^i_{(0)}\}_{i=1}^N$ and $\beta_{(0)}, \vartheta_{(0)} > 0$;
for $k = 0, 1, \ldots$ **do**
  Compute $X^i_{(k+1)}$ according to (4.3) for all $i = 1, \ldots, N$;
  if $R[f^N_{(k+1)}] \leq 1/\sqrt{\vartheta_{(k)}}$ **then**
    #feasibility check satisfied#;
    $\vartheta_{(k+1)} = \eta_\vartheta \vartheta_{(k)}$;
    $\beta_{(k+1)} = \beta_{(k)}$
  else
    #feasibility check violated#;
    $\vartheta_{(k+1)} = \min\{\vartheta_{(k)}/\eta_\vartheta, \vartheta_{(0)}\}$;
    $\beta_{(k+1)} = \eta_\beta \beta_{(k)}$
  end
end

Algorithm 1: CBO adaptive algorithm for exact penalization.

In CBO methods, the Boltzmann-Gibbs distribution is used to compute the consensus point and, in general, to guide the particle system. We also propose to check the feasibility condition with

$$R^\alpha_\beta[f] := \frac{\int r(x) \exp(-\alpha P_\beta(x)) df(x)}{\int \exp(-\alpha P_\beta(x)) df(x)}, \tag{4.6}$$

for $\beta, \alpha > 0$, that is, the weighted counter-part of (4.4). For an empirical distributions f^N, $R^\alpha_\beta[f^N]$ converges as $\alpha \to \infty$ to the constraint violation of the best particle of the ensemble in terms of the objective function P_β. Therefore, in the limit case $\alpha = +\infty$, the feasibility check is performed directly at the algorithm's candidate solution. The numerical experiments presented in Section 4.3 will show that using the Boltzmann-Gibbs distribution typically makes the adaptive mechanism more accurate.

Similar strategies have been introduced in [53, 77] for quadratic programming problems in the context of finite-dimensional optimization. Methods suggested in [53, 77] additionally include smooth approximation of the ℓ_1-penalty term. This is not needed in our settings because CBO methods do not require any smoothness of the objective function. We also note that similar adaptive strategies have been extended to infinite-dimensional settings [45, 46].

In the next section, we analyze the convergence properties of Algorithm 1 by relying on the mean-field approximation of the CBO dynamics.

4.2 Convergence analysis in mean-field law

The idea behind adaptive strategy can be summarized as follows. In a first stage of the computation, we expect the penalty parameter $\beta_{(k)}$ to be smaller than the threshold value $\bar{\beta}$. Therefore, being the penalization strength not sufficient, particles are likely to concentrate in an unfeasible area of the search space and, as a consequence, to violate the feasibility check (4.5). Eventually, $\beta_{(k)}$ becomes larger than $\bar{\beta}$, leading to the second stage of the computation where the penalized problem (4.2) is equivalent to the constrained one (4.1). Ideally, $\beta_{(k)}$ should stop increasing at this point in order to avoid the term $\beta_{(k)} r$ to overwhelm $\mathcal{E}$ in the penalized objective. In this section, we analytically investigate under which conditions Algorithm 1 follows such expected behavior.

To do so, we recall the mean-field approximation of (4.3) which we already introduced in the previous chapters:

$$\overline{X}_{(k+1)} = \overline{X}_{(k)} + \lambda\tau\left(m^\alpha_{\beta_{(k)}}[f_{(k)}] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha_{\beta_{(k)}}[f_{(k)}] - \overline{X}_{(k)}\right)\odot\overline{\theta}_{(k)}\,, \tag{4.7}$$

where $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$. The feasibility check, in the mean-field dynamics, is then performed by using $R[f_{(k+1)}]$, see Algorithm 2.

As for the case where the objective is fixed (see Proposition 2.3), $f_{(k)} \in \mathcal{P}_2(\mathbb{R}^d)$ follows from $f_{(0)} \in \mathcal{P}_2(\mathbb{R}^d)$ provided the penalty sub-problems satisfy the following assumptions.

Assumption 4.2. *The penalized objective function P_β satisfies*

- $\underline{P}_\beta := \inf P_\beta > -\infty$;

Set $\lambda, \sigma, \tau > 0, \eta_\beta, \eta_\vartheta > 1$;
Initialize $\overline{X}_{(0)}, \beta_{(0)}, \vartheta_{(0)}$;
for $k = 0, 1, \dots$ **do**
  Compute $\overline{X}_{(k+1)}$ according to (4.7);
  if $R[f_{(k+1)}] \leq 1/\sqrt{\vartheta_{(k)}}$ **then**
    #feasibility check satisfied#;
    $\vartheta_{(k+1)} = \eta_\vartheta \vartheta_{(k)}$;
    $\beta_{(k+1)} = \beta_{(k)}$
  else
    #feasibility check violated#;
    $\vartheta_{(k+1)} = \min\{\vartheta_{(k)}/\eta_\vartheta, \vartheta_{(0)}\}$;
    $\beta_{(k+1)} = \eta_\beta \beta_{(k)}$
  end
end

Algorithm 2: Mean-field CBO adaptive algorithm for exact penalization.

- *there exists* $L_{P_\beta}, c_u, c_l, R_l > 0$ *such that*

$$\begin{cases} |P_\beta(x) - P_\beta(y)| \leq L_{P_\beta}(1 + |x| + |y|)|x - y| & \forall\, x, y \in \mathbb{R}^{\mathrm{d}}, \\ P_\beta(x) - \underline{P}_\beta \leq c_u(1 + |x|^2) & \forall\, x \in \mathbb{R}^{\mathrm{d}}, \\ P_\beta(x) - \underline{P}_\beta \geq c_l |x|^2 & \forall\, x : |x| > R_l\,. \end{cases} \tag{4.8}$$

We note that the constants in the above estimates may depend on β, but this does not affect the computation of moment estimates for $f_{(k)}$. Indeed, during a run of Algorithm 2, only a finite number of penalized sub-problems corresponding to $\beta_{(h)}, h \leq k$ are relevant to the CBO dynamics up to time k.

Furthermore, we simplify the analytical investigation by considering a specific penalty term. For a given set $A \subset \mathbb{R}^{\mathrm{d}}$, let $\mathrm{dist}(\cdot, A)$ denote the set distance, that is, $\mathrm{dist}(x, A) = \inf_{y \in A} |x - y|$, for any $x \in \mathbb{R}^{\mathrm{d}}$. We assume:

Assumption 4.3. *The penalty term is given by* $r(x) = \mathrm{dist}(x, \mathcal{M})$, *with* $\mathcal{M}$ *being the feasible set.*

As previously discussed, this choice r is not strictly necessary for the penalized sub-problems to be exact, and other penalization terms, like ℓ_1-penalization, can be used. On the other hand, this assumption makes the constraint violation $r(x) = \mathrm{dist}(x, \mathcal{M})$ a lower bound for the error $|x - x^*|$ for any feasible solution $x^* \in \mathcal{M}$. This property will be useful to theoretically control the feasibility check in terms of expected mean-squared error.

To analyze the algorithm, we rely on the convergence properties of the CBO dynamics and, in particular, on Theorem 3.6. The convergence result needs to be applied to every

sub-problem (4.2) in order for us to be able to study Algorithm 2. To this end, we introduce the following inverse continuity assumption on P_β :

Assumption 4.4 (Growth conditions around minimizer of P_β)**.** *There exists a unique global minimizer x_β^* of P_β and $c_1, p_1, > 0, R > 0$ such that*

$$c_1 \|x - x_\beta^*\|_\infty^{p_1} \leq P_\beta(x) - P_\beta(x_\beta^*) \qquad \forall x,\ \|x - x_\beta^*\|_\infty \leq R \tag{4.9}$$

and lower bound $C_\infty > 0$ such that

$$C_\infty < P_\beta(x) - P_\beta(x_\beta^*) \qquad \forall x,\ \|x - x_\beta^*\|_\infty > R. \tag{4.10}$$

Similarly to the analysis for the unconstrained CBO method, we track the convergence towards the minimizer x_β^*, by studying the evolution of the expected ℓ_2-error

$$\mathrm{Err}_\beta[f] := \int |x - x_\beta^*|^2 df(x)$$

for any $f \in \mathcal{P}_2(\mathbb{R}^d)$. We note that, whenever penalization is exact ($\beta > \overline{\beta}$), x_β^* corresponds to the solution to the constrained problem (4.1) which we simply indicate as x^*. We indicate the error with respect to x^* as

$$\mathrm{Err}[f] := \int |x - x^*|^2 df(x). \tag{4.11}$$

We are now ready to enunciate the main convergence result.

Theorem 4.5. *Let Assumptions 4.1, 4.3 hold for r and let P_β satisfy Assumptions 4.2, 4.4 for all $\beta > 0$. Consider any $\overline{X}_{(0)}$ with $\mathrm{supp}(f_{(0)}) = \mathbb{R}^d$, $\beta_{(0)} > 0$ and parameters $\eta_\beta, \eta_\vartheta > 1$, $\lambda, \sigma > 0$ with $\lambda > \sigma^2$.*

Let $\overline{X}_{(k)}$ be constructed according to Algorithm 2, and $f_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$. For any accuracy $\varepsilon > 0$, there exists a choice of parameters $\{\alpha, \tau, \vartheta_{(0)}\}$ such that the mean-field mono-particle process constructed via Algorithm 2 satisfies

$$\min_{k : k\tau \leq T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon$$

for some time horizon $T^* > 0$ sufficiently large.

To prove the above convergence result, we consider two different scenarios. When penalization is not exact $\beta_{(k)} \leq \overline{\beta}$, we show that the feasibility check will be necessary violated until the threshold value is reached, provided $\vartheta_{(0)}$ is sufficiently large. The second scenario corresponds to the exact penalization case, where $\beta_{(k)} > \overline{\beta}$ at some iteration k. If this happens, we show that the CBO mean-field process converges to the true solution up to the given accuracy ε.

4.2.1 Convergence for exact penalization (case $\beta_{(k)} > \overline{\beta}$)

We start by analyzing the algorithm behavior in the case where, at some step k_0 of the computation, $\beta_{(k_0)} > \overline{\beta}$. First, we show that the mean-field distribution concentrates around the minimizer x^* to (4.1), provided $x^* \in \mathrm{supp}(f_{(k_0)})$. We recall that for all $\beta > \overline{\beta}$, $x^*_\beta = x^*$ and so $\mathrm{Err}[\cdot] = \mathrm{Err}_\beta[\cdot]$.

Proposition 4.6. *Under the settings of Theorem 4.5, assume at some algorithmic step k_0 it holds $\beta_{(k_0)} > \overline{\beta}$ and $x^* \in \mathrm{supp}(f_{(k_0)})$.*

Then, there exists α, T^* sufficiently large and τ sufficiently small such that

$$\min_{k \geq k_0, k\tau \leq T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon.$$

Proof. If $\mathrm{Err}[f_{(k_0)}] \leq \varepsilon$ the result is trivially true, so we assume $\mathrm{Err}[f_{(k_0)}] > \varepsilon$. Thanks to the assumptions considered, we can apply Theorem 3.6 for the objective $P_{\beta_{(k_0)}}$ and obtain that

$$\mathrm{Err}[f_{(k_0+1)}] \leq \exp\left(-\tau \frac{\lambda - \sigma^2}{4}\right) \mathrm{Err}[f_{(k_0)}]$$

due to the exponential decay estimates (3.8), provided $\alpha > \alpha_{k_0}, \tau < \tau_{k_0}$ for some $\alpha_{k_0}, \tau_{k_0} > 0$. Next, as $\beta_{(k)}$ is non-decreasing, it holds $\beta_{(k_0+1)} > \overline{\beta}$, and so the global minimum of the penalized sub-problem corresponding to $\beta_{(k_0+1)}$ is still given by x^*. We also note that, thanks to Lemma 3.5, $x^* \in \mathrm{supp}(f_{(k_0+1)})$. As consequence, we can iteratively apply Theorem 3.6 for $k \geq k_0$ to obtain

$$\mathrm{Err}[f_{(k)}] \leq \exp\left(-(k-k_0)\tau \frac{\lambda - \sigma^2}{4}\right) \mathrm{Err}[f_{(k_0)}] \tag{4.12}$$

until the desired accuracy ε is reached.

We remark that the choice of α, τ may become more and more restrictive at each iteration. Since the decay estimate is independent on α and τ, though, only a finite number of iterations is required to reach the desired accuracy ε. Therefore, there exists a suitable choice of parameters $\alpha, T^* < \infty$ large enough and $\tau > 0$ sufficiently small such that

$$\min_{k \geq k_0, k\tau \leq T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon.$$

□

It is clear from the above proof that the feasibility condition does not play a role in the convergence analysis once $\beta_{(k)}$ has reached the threshold value $\overline{\beta}$. In practice, though, a penalty parameter which is too large with respect to $\overline{\beta}$ makes the optimization problem more challenging for the CBO particle system (see numerical experiments in Section 4.3). Next, we derive conditions under which $\beta_{(k_0)}$ stop increasing after reaching $\overline{\beta}$.

Proposition 4.7. *Under the settings of Proposition 4.6, if* $\eta_\vartheta, \vartheta_{(k_0-1)}$ *satisfy*

$$\eta_\vartheta \leq \exp\left(\tau\frac{\lambda-\sigma^2}{4}\right) \qquad \textit{and} \qquad \mathrm{Err}[f_{(k_0)}] \leq \frac{1}{\vartheta_{(k_0-1)}},$$

then the feasibility condition is satisfied until $\mathrm{Err}[f_{(k)}] \leq \varepsilon$.

Proof. Given that x^* is solution to the constrained problem, it necessarily belongs to the feasible space $\mathcal{M}$. This means that the error is an upper bound to the constraint violation: $r(x) = \mathrm{dist}(x, \mathcal{M}) \leq |x - x^*|$ for any $x \in \mathbb{R}^d$. As a consequence, for any $f \in \mathcal{P}(\mathbb{R}^d)$, we can bound the constraint violation with the expected mean-squared error as following:

$$R[f] = \int \mathrm{dist}(x, M) df(x) \leq \int |x - x^*| df \leq \left(\int |x - x^*|^2 df\right)^{\frac{1}{2}} = \sqrt{\mathrm{Err}[f]} \tag{4.13}$$

where we also used Jensen's inequality.

Intuitively, the assumption on η_ϑ ensures the that the decay rate of the tolerance is faster than the decay of the error square root. Thus, the tolerance will continue to be larger than $\sqrt{\mathrm{Err}[f_{(k)}]}$ and, in particular, of $R[f_{(k)}]$ due to (4.13). Thanks to the choice of η_ϑ, it holds indeed

$$\exp\left(-(k+1-k_0)\tau\frac{\lambda-\sigma^2}{4}\right) \leq \eta_\vartheta^{-(k+1-k_0)}.$$

We now plug the above estimate into the decay rate (4.12) provided in the proof of Proposition 4.6. Together with the assumption $\mathrm{Err}[f_{(k_0)}] < 1/\vartheta_{(k_0-1)}$, we obtain

$$\mathrm{Err}[f_{(k+1)}] \leq \exp\left(-(k+1-k_0)\tau\frac{\lambda-\sigma^2}{4}\right)\mathrm{Err}[f_{(k_0)}] \leq \eta_\vartheta^{-(k+1-k_0)}\mathrm{Err}[f_{(k_0)}] \leq \frac{1}{\eta_\vartheta^{(k+1-k_0)}\vartheta_{(k_0-1)}}$$

for all $k \geq k_0$, up until the desired accuracy is reached. Given that $\vartheta_{(k)} \leq \eta_\vartheta^{(k+1-k_0)}\vartheta_{(k_0-1)}$, we further obtain

$$R[f_{(k+1)}] \leq \sqrt{\mathrm{Err}[f_{(k+1)}]} \leq \frac{1}{\sqrt{\theta_{(k)}}},$$

that is, the feasibility check is satisfied for any $k \geq k_0$ until $\mathrm{Err}[f_{(k)}] \leq \varepsilon$. □

In the next section, we will analyze instead the algorithm's behavior when the penalty parameter is not sufficiently large for the penalization technique to be exact.

4.2.2 Violation of feasibility check (case $\beta_{(k)} \leq \overline{\beta}$) and proof of Theorem 4.5

As we assumed all penalized objectives P_β to satisfy the inverse continuity assumption (Assumption (4.4)) required by Theorem 3.6, we can use the convergence estimates provided by the theorem to study the CBO dynamics. Intuitively, if $\beta_{(k)} \leq \overline{\beta}$ for all k, the particle system will eventually concentrate in a ε-neighborhood of an infeasible solution. If the feasibility tolerance (which depends on $\vartheta_{(0)}$) is sufficiently small, then the feasibility check will be violated, leading to a contradiction. In the following, we make these statements more precise and finally provide a proof to Theorem 4.5.

Proposition 4.8. *Under the settings of Theorem 4.5, assume at some algorithmic step $k_0 > 1$ to hold $\beta_{(k_0)} \leq \overline{\beta}$ and $x^*_{\beta_{k_0}} \in \operatorname{supp}(f_{(k_0)})$. If condition*

$$r\left(x^*_{\beta_{k_0}}\right) > \frac{1}{\sqrt{\vartheta_0}}, \tag{4.14}$$

is satisfied, then the feasibility condition will be violated at some finite $k > k_0$, provided α, T^* are sufficiently large and τ is sufficiently small.

Proof. For notational simplicity, in the following we set $x^*_0 = x^*_{\beta_{(k_0)}}$. Seeking a contradiction, suppose the feasibility check is always satisfied and, so, $\beta_{(k)} = \beta_{(k_0)}$ for all $k \geq k_0$. Given that the objective function is fixed and $x^*_0 \in \operatorname{supp}(f_{(k_0)})$, we can directly apply Theorem 3.6 for some arbitrary accuracy $\varepsilon' > 0$.

We obtain that there exists α_0, T^*_0, τ_0 such that

$$\min_{k \geq k_0 \,:\, k\tau \leq T^*_0} \operatorname{Err}_{\beta_{(k_0)}}[f_{(k)}] \leq \varepsilon'. \tag{4.15}$$

We show that for ε' sufficiently small, we are lead to a contradiction.

By the choice of r (Assumption 4.3) and the triangular inequality, we have

$$r(x^*_0) = \operatorname{dist}(x^*_0, \mathcal{M}) \leq \operatorname{dist}(x, \mathcal{M}) + |x - x^*_0| = r(x) + |x - x^*_0|$$

and we can bound for any $f \in \mathcal{P}(\mathbb{R}^d)$ the constraint violation from below as

$$R[f] = \int r(x) df(x) \geq r(x^*_0) - \int |x - x^*_0| df(x) \geq r(x^*_0) - \sqrt{\operatorname{Err}_{\beta_{(k_0)}}[f]}$$

where we also used Jensen's inequality in the last step. As a consequence of (4.15), at some $k^* \geq k_0$ it holds

$$R[f_{(k^*)}] \geq r(x^*_0) - \sqrt{\varepsilon'}.$$

At the same time, since we supposed the feasibility check to be satisfied at all $k \geq k_0$, we also have

$$R[f_{(k^*)}] \leq \frac{1}{\sqrt{\vartheta_{(k^*-1)}}} \leq \frac{1}{\sqrt{\vartheta_{(0)}}}.$$

By combining the inequalities together, we obtain

$$r(x_0^*) \leq \frac{1}{\sqrt{\vartheta_{(0)}}} + \sqrt{\varepsilon'},$$

which contradicts the given assumption $r(x_0^*) > 1/\sqrt{\vartheta_{(0)}}$ for ε' sufficiently small. □

Proof of Theorem 4.5. If the initial penalty parameter already satisfies $\beta_{(0)} > \overline{\beta}$, we can directly apply Proposition 4.6 and conclude that there exists α, T^* sufficiently large and τ sufficiently small such that the desired error estimate holds.

If this is not the case, let us consider the set of global solutions $x^*_{\beta_{(k)}}$ for the penalized sub-problems $\beta_{(k)}, k \geq 0$

$$S^* := \{x^*_{\beta_{(k)}} : k \geq 0\}.$$

We note that S^* is a finite collection of points since $x^*_\beta = x^*$ for all $\beta > \overline{\beta}$. Given that the mass around any point can be bounded from below as in Lemma 3.5 (as noted in Remark 3.16), we have that

$$S^* \subset \operatorname{supp}(f_{(k)}) \qquad \text{for all} \quad k \geq 0 \tag{4.16}$$

thanks to $\operatorname{supp}(f_{(0)}) = \mathbb{R}^d$.

To apply Proposition 4.8 for all $\beta_{(k)} \geq \overline{\beta}$ we additionally assume $\vartheta_{(0)}$ to satisfy

$$r(x^*_{\beta_{(k)}}) > \frac{1}{\sqrt{\vartheta_{(0)}}} \qquad \text{for all} \quad k \geq 0. \tag{4.17}$$

Thanks to (4.16) and (4.17), we can iteratively apply Proposition 4.8 (a finite number of times) to obtain that there exists an iteration K' such that $\beta_{(K')} > \overline{\beta}$ for some α', T' sufficiently large and τ' sufficiently small. Finally, since $x^* \in \operatorname{supp}(f_{(K')})$ and we are in the exact penalization case, we can conclude by applying Proposition 4.6: there exists $T^* > T'$, $\alpha > \alpha'$ and $\tau < \tau'$ such that

$$\min_{k : k\tau \leq T^*} \operatorname{Err}[f_{(k)}] \leq \varepsilon.$$

□

Remark 4.9.

- *Contrary to the unconstrained settings, the above analysis does not provide an explicit definition for T^* in terms of initial conditions $f_{(0)}$ and accuracy ε. This is because when applying Proposition* (4.8) *we do not know when exactly the feasibility check will be violated, but only that it will be violated at some time.*
- *The use of a discrete-in-time description of the CBO dynamics greatly simplifies the analytical investigation with respect to [7], where the CBO dynamics is continuous, but the parameters tuning happens at discrete times. In particular, we can avoid proving the stability of the constraint violation through the mollification procedure (see [7, Section 3.4]).*

4.3 Numerical experiments

In this section, we test Algorithm 1 against different benchmark problems. We start by performing numerical experiments with numerous particles to validate the mean-field analysis performed in Section 4.2. Then, we show how the algorithm performs with a relatively small number of particles by testing it against benchmark problems in search spaces of dimension up to d = 20. We will focus, in particular, on understanding how sensible the algorithm is with respect to its parameters and if the suggested adaptive strategy is able to identify the optimal penalty parameter value. As suggested by the analysis, we set the parameter α to $\alpha = 10^6$: a relatively large value with respect to the objective functions considered.

4.3.1 Simulation of mean-field regime

Being the theoretical analysis based on a mean-field approximation of the CBO dynamics, we start by running the algorithm with a large number of particles N (large with respect to the amount needed for a comparable performance). Therefore, we employ in the following $N = 10^6$ particles.

Let us consider the constrained problem given by:

$$\begin{aligned} &\min_{x\in\mathbb{R}^2} \mathcal{E}(x) := \frac{1}{2}\sum_{\ell=1}^{2}\left(\frac{x_\ell^4}{5} - 2x_\ell^2 + x_\ell\right) + 10 \\ &\text{subject to} \quad g(x) = R(z) = \frac{1}{2}\sum_{i=1}^{2} z_\ell^2 - 10\cos(2\pi z_\ell) + 5 \le 0\,, \end{aligned} \tag{4.18}$$

where the affine change of variable $x \mapsto z$ is determined by

$$z = \begin{pmatrix} \cos(\pi/6) & -\sin(\pi/6) \\ \sin(\pi/6) & \cos(\pi/6) \end{pmatrix} (x - (1,1)^\top).$$

The problem is illustrated in Figure 4.1. We note that $\mathcal{E}$ attains a non-feasible global minimum $\hat{x}$ and that the feasible set is made of disjoint sets of different areas. The solution to the constrained problem x^* is not a local minimum of $\mathcal{E}$ and so the constraint is active. To penalize infeasible points, we use ℓ_1-penalization, $r(x) = |g(x)|$. Being the constraint function g an affine transformation of the Rastrigin function [63], the penalized functions P_β are highly non-convex and attain several local minima, see for instance Figure 4.1b. The precise threshold value $\overline{\beta}$ is unknown, but numerical experiments suggest that $\overline{\beta} \in [4, 5]$.

Because of the low dimensionality of the search space, we use isotropic exploration in the CBO particles update rule. We set the initial penalty parameter to $\beta_{(0)} = 0.1$, so that is holds $\beta_{(0)} < \overline{\beta}$. The remaining algorithm parameters are set to $\lambda = 1, \sigma = 0.5, \tau = 0.1$. Figure 4.2 shows the evolution of the particle distribution $f_{(k)}$ with parameters $\eta_\vartheta = 1.1, \eta_\beta = 1.1$

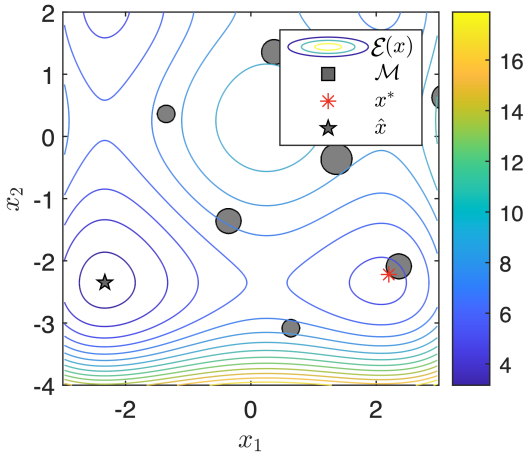

(a) Contour values of the objective function $\mathcal{E}$ and the feasible set $\mathcal{M}$ (in gray)

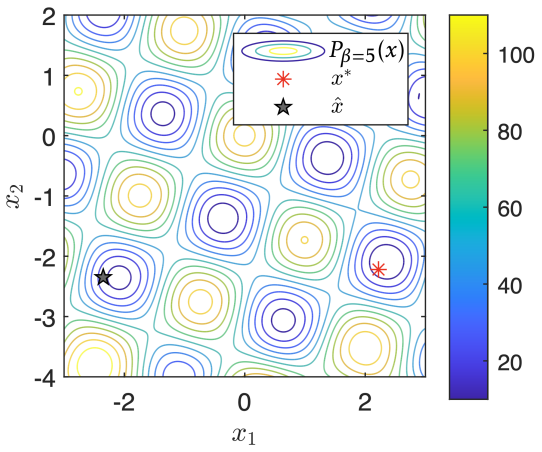

(b) Contour values of the penalty function P_β with $\beta = 5$.

Figure 4.1: Illustration of the constrained problem (4.18) (left) and an equivalent penalty sub-problem (right). The global solution to the constrained problem x^* in shown in red, while the infeasible minimum $\hat{x}$ of $\mathcal{E}$ is shown in black.

and initial tolerance $\vartheta_{(0)} = 2$. Given that the initial penalty parameter is not sufficiently large, $\beta_{(0)} < \overline{\beta}$, the particles initially converge around the infeasible minimum of the objective function, see Figures 4.2a, 4.2b, and 4.2c. After a sufficient number of iterations, it holds $\beta_{(k)} > \overline{\beta}$ due to the adaptive mechanism and the particles move towards the solution of the constrained problem, see Figures 4.2e, 4.2f. We note that, between these two phases, there is an intermediate step (Figure 4.2d) where particles are still located around the infeasible minimum, but the mass is partially redistributed over the search space.

Figure 4.3a shows the evolution of three quantities: the constraint violation $R[f_{(k)}]$ (4.4), the feasibility check tolerance $1/\sqrt{\vartheta_{(k)}}$, and the penalty parameter $\beta_{(k)}$. We note that, even though $\beta_{(k)}$ reaches the threshold value $\overline{\beta}$ already around time $t = k\tau = 5$, the feasibility check is violated up after $t = 12$. This choice of the initial tolerance $\vartheta_{(0)} = 2$ satisfies requirement (4.14) of Proposition 4.8, that is, the constraint violation of infeasible minima (only $\hat{x}$ in this case) is larger than the initial tolerance. Indeed, we have $r(\hat{x}) \approx 0.88$, while $1/\sqrt{\vartheta_{(0)}} = 1/\sqrt{2} \approx 0.71$. If we increase the initial tolerance by setting $\vartheta_{(0)} = 0.25$, on the other hand, condition (4.14) is not satisfied anymore and the particles concentrate around the infeasible minimum, see Figure 4.3b.

Proposition 4.7 suggests that, in order for $\beta_{(k)}$ not to increase during the entire computation, η_ϑ should not exceed the convergence rate of the CBO dynamics. This is the case of the first experiment performed, with $\eta_\vartheta = 1.1$, where indeed the tolerance decreases faster than then constraint violation, see Figure 4.3a. If we increase the update rate for

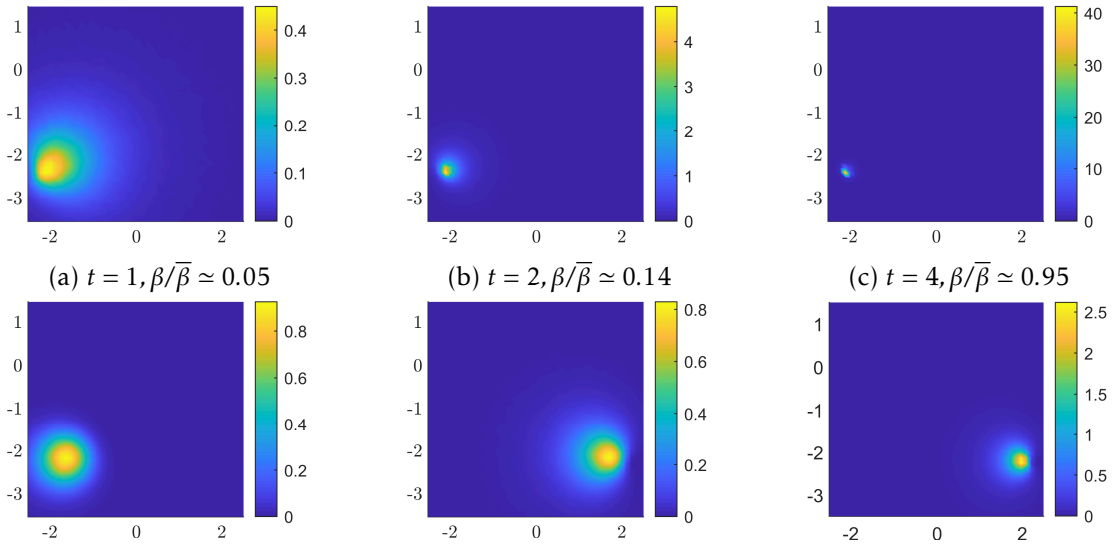

(a) $t = 1, \beta/\overline{\beta} \simeq 0.05$ (b) $t = 2, \beta/\overline{\beta} \simeq 0.14$ (c) $t = 4, \beta/\overline{\beta} \simeq 0.95$

(d) $t = 5, \beta/\overline{\beta} \simeq 2.47$ (e) $t = 8, \beta/\overline{\beta} \simeq 43$ (f) $t = 10, \beta/\overline{\beta} \simeq 289$

Figure 4.2: Evolution of the particle distribution at different times. Plot shows a density constructed via Gaussian kernel, starting from the empirical distribution made of $N = 10^6$ particles. CBO dynamics with isotropic diffusion is used. Algorithm parameters are set to $\lambda = 1, \sigma = 0.5, \tau = 0.1, \vartheta_{(0)} = 2, \beta_{(0)} = 0.1, \eta_\vartheta = 1.1, \eta_\beta = 1.1$. The time psudo-time t is given by $t = k\tau$.

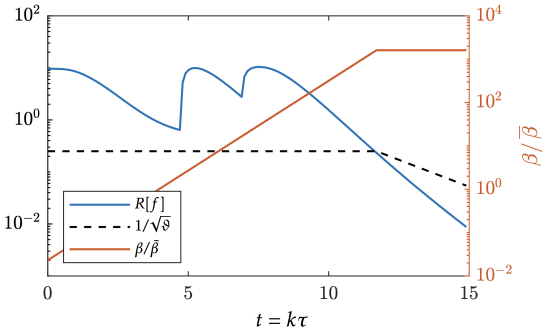

(a) $1/\sqrt{\vartheta_{(0)}} = 0.25$, $\eta_\vartheta = 1.1$, unweighted constraint violation R. $\mathrm{Err}[f^N_{(k^*)}] \simeq 1.66 \times 10^{-4}$

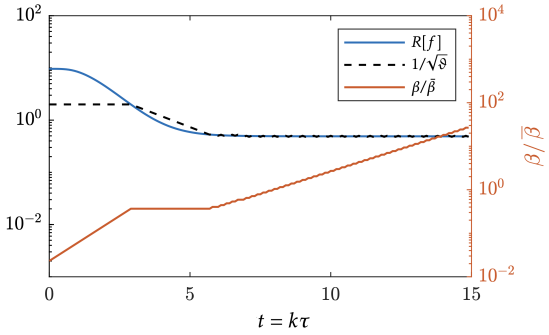

(b) $1/\sqrt{\vartheta_{(0)}} = 2$, $\eta_\vartheta = 1.1$, unweighted constraint violation R. $\mathrm{Err}[f^N_{(k^*)}] \simeq 9.26$

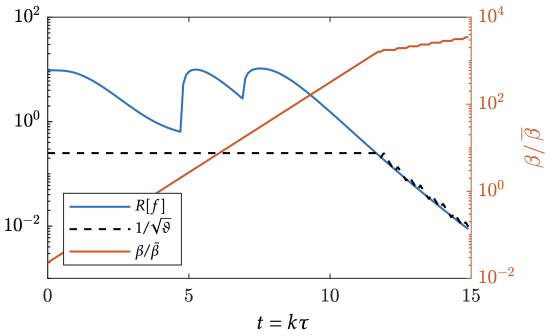

(c) $1/\sqrt{\vartheta_{(0)}} = 0.25$, $\eta_\vartheta = 1.5$, unweighted constraint violation R. $\mathrm{Err}[f^N_{(k^*)}] \simeq 1.66 \times 10^{-4}$

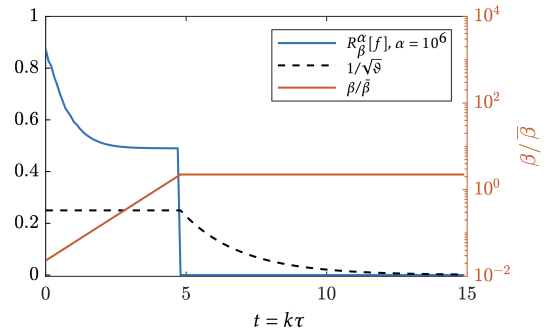

(d) $1/\sqrt{\vartheta_{(0)}} = 0.25$, $\eta_\vartheta = 1.1$, weighted constraint violation R^α_β. $\mathrm{Err}[f^N_{(k^*)}] \simeq 1.66 \times 10^{-4}$

Figure 4.3: Evolution of the constraint violation (in blue), the tolerance $1/\sqrt{\vartheta_{(k)}}$ and the penalty parameter $\beta_{(k)}$ in the simulation of the mean-field algorithm applied to problem (4.18) in four different settings. Only in Figure 4.3d the constraint violation is calculated using R^α_β which uses the Boltzmann-Gibbs distribution, see (4.6). The average ℓ_2-error $\mathrm{Err}[f^N_{(k^*)}]$ (4.11) shows the reached accuracy at the final step $k^* = 15/\tau$ of the computation.

$\vartheta_{(k)}$ to $\eta_\vartheta = 1.5$, we have that the feasibility check is violated also when the particles are concentrating at the true solution of the constrained problem, see Figure 4.3c.

Finally, we perform a last experiment where we use $\vartheta_{(0)} = 2$ and $\eta_\vartheta = 1.1$, as in the first one, but compute the feasibility check with the weighted constraint violation R^α_β (given by (4.6)) instead of R. While the accuracy obtained is comparable with the first experiment, we note that the feasibility check is satisfied soon after the threshold value $\overline{\beta}$ is reached. As a consequence, the final penalty parameter computed is much closer the optimal value, see Figure 4.3d. We will show in the next sections that using the weighted constraint violation R^α_β has a strong impact on the algorithm performance when a smaller number of particles is used.

4.3.2 Benchmark problems in $\mathbf{d} = 5$

We now perform experiments in a higher dimension, $\mathrm{d} = 5$, where we employ a smaller number of particles, that is, $N = 200$.

The problems we consider are given by combinations between two objective functions,

$$\mathcal{E}_1(x) = \frac{1}{\mathrm{d}} \sum_{i=1}^{\mathrm{d}} \left(\frac{x_i^4}{5} - 2x_i^2 + x_i \right) + 10$$

$$\mathcal{E}_2(x) = -20 \exp\left(-0.2 \sqrt{\frac{1}{\mathrm{d}} \sum_{i=1}^{\mathrm{d}} (x-o)_i^2} \right) - \exp\left(\frac{1}{\mathrm{d}} \sum_{i=1}^{\mathrm{d}} \cos(2\pi(x-o)_i) \right) + 20 + e$$

$$\text{where} \quad o = (1.7\bar{6}, 1.5\bar{3}, 1.\bar{3}, 1.0\bar{6}, 0.8\bar{3}),$$

and two admissible sets, the sphere $\mathcal{M}_1 = \mathbb{S}^4$, and the torus, $\mathcal{M}_2 = \mathbb{T}^4$. Therefore, the four constrained optimization problems are

$$\min_{x \in \mathbb{R}^5} \mathcal{E}_l(x) \quad \text{subject to} \quad x \in \mathcal{M}_i \quad \text{for} \quad l = 1,2, \quad i = 1,2. \tag{4.19}$$

As assumed in the theoretical analysis (Assumption 4.3), we use the corresponding distance functions to penalize the objectives:

$$r_1(x) = \mathrm{dist}(x, \mathbb{S}^4) = \big| |x| - 1 \big|$$

$$r_2(x) = \mathrm{dist}(x, \mathbb{T}^4) = \left| \sqrt{\left(\sqrt{|x|^2 - x_{\mathrm{d}}^2} - 1 \right)^2 + x_{\mathrm{d}}^2} - 0.5 \right|,$$

where x_{d} indicates the d-th component of x. Both objectives are non-convex and attain several local minima. We note that the first one coincides with the objective function of problem (4.18), while $\mathcal{E}_2$ is the Ackley function, a well-known benchmark function in global optimization [63].

In the following experiments, we set the algorithm parameters to: $\lambda = 1, \sigma = 0.6, \vartheta_{(0)} = 4, \eta_\vartheta = 1.1, \eta_\beta = 1.1$. Particles evolve with isotropic diffusion. We consider different values of $\beta_{(0)}$ to simulate the scenario where the magnitude of the optimal penalty parameter $\overline{\beta}$ is unknown. For the considered problems, it holds $\overline{\beta} \in [1, 10]$, and so we initialize β in a range between $10^{-5}\overline{\beta}$ and $10^3\overline{\beta}$, approximately. To measure the algorithm performance, we consider a run successful if

$$\|m^\alpha[f^N_{(k^*)}] - x^*\|_\infty < 0.1\,,$$

where $k^* = 300$ is the last iteration, and x^* is, as before, the global solution to the constrained problem. The 200 particles are initially sampled from the uniform distribution over $[-2, 2]^{\mathrm{d}}$.

In the first two experiments, we compare the use of the constraint violation R (4.4) and its weighted counterpart R^α_β in the evaluation of the feasibility check. As shown in Figures 4.4a and 4.4b, the use of R^α_β drastically improves the algorithm success rate when $\beta_{(0)} < \overline{\beta}$. When the initial guess overshoots the optimal value, $\beta_{(0)} > \overline{\beta}$, the algorithm performance is rather poor in both cases. To improve the algorithm performance in this scenario, we propose an additional adaptive strategy: $\beta_{(k)}$ is decreased with rate η_β up until the first feasibility check is violated. With this additional heuristic strategy, the algorithm performance considerably improves, see Figure 4.4c. To remark, once again, the importance of the correct tuning of β, we also perform experiments with no adaptive strategy at all. As expected, particles concentrate around the solution to the constrained problem only if $\beta_{(0)} > \overline{\beta}$, see Figure 4.4d.

As already observed in the simulation of the mean-field regime, the use of the weighted constraint violation R^α_β speeds up the adaptive mechanisms, in the sense that the feasibility check is satisfied sooner after the $\beta_{(k)}$ reaches the threshold $\overline{\beta}$. This can be shown by looking at the final value β reached at the end of the computation, see Figures 4.5a and 4.5b. Decreasing the penalty parameter at the beginning of the computation makes the final value $\beta_{(k^*)}$ reached by the algorithm almost independent on the initialization of the penalty parameter, see Figure 4.5c.

Finally, Figure 4.6 shows the accuracy reached in the third experiment, where R^α_β and decreasing strategy are used. When a relatively small number of particles is used, $N = 200$, the average ℓ_2-error seems to be sensitive to outliers, whereas the performance with respect to the ℓ_∞-error $\|m^\alpha[f^N_{(k^*)}] - x^*\|_\infty$ seems to be more stable with respect to the choice of $\beta_{(0)}$. The best performance in terms of accuracy are obtained when $\beta_{(0)}$ is approximately equal to $\overline{\beta}$, or slightly larger.

4.3.3 Benchmark problems in higher dimensions

In this section, we consider benchmark constrained problems scalable to any search space dimension d. So far, we have employed isotropic diffusion in the CBO dynamics for

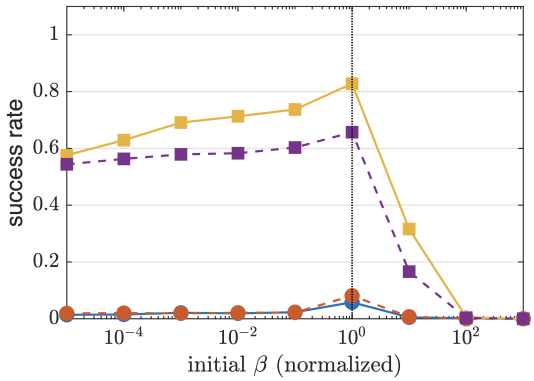

(a) Feasibility check with R

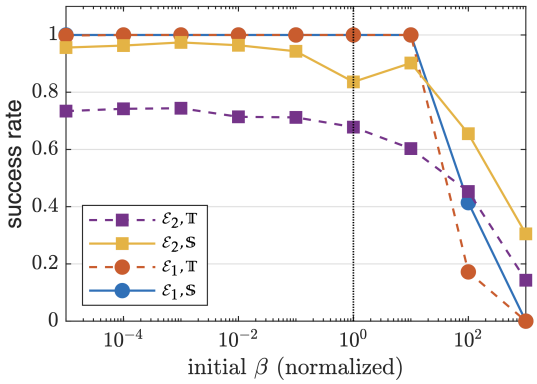

(b) Feasibility check with R_β^α

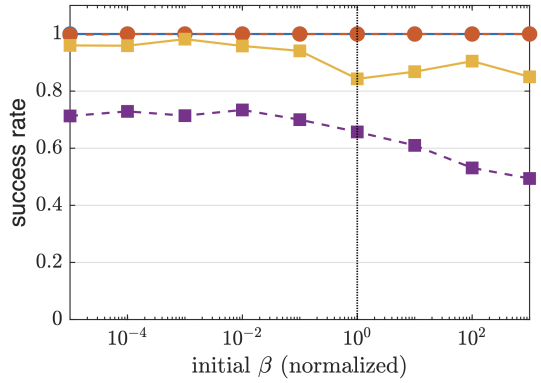

(c) Feasibility check with R_β^α and initial decreasing strategy for β

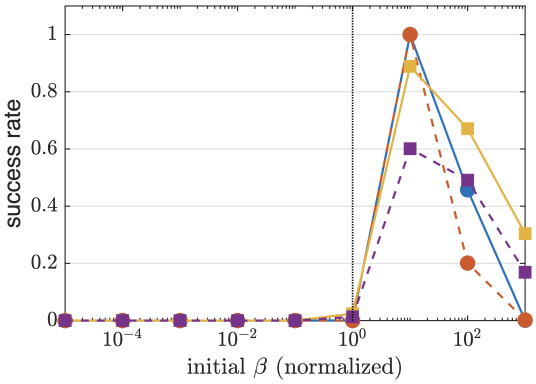

(d) No adaptive strategy

Figure 4.4: Success rate obtained for different initialization of $\beta_{(0)}$ normalized with respect to the threshold value $\overline{\beta}$. The proposed method is tested against problems (4.19) and results are averaged over 1000 runs.

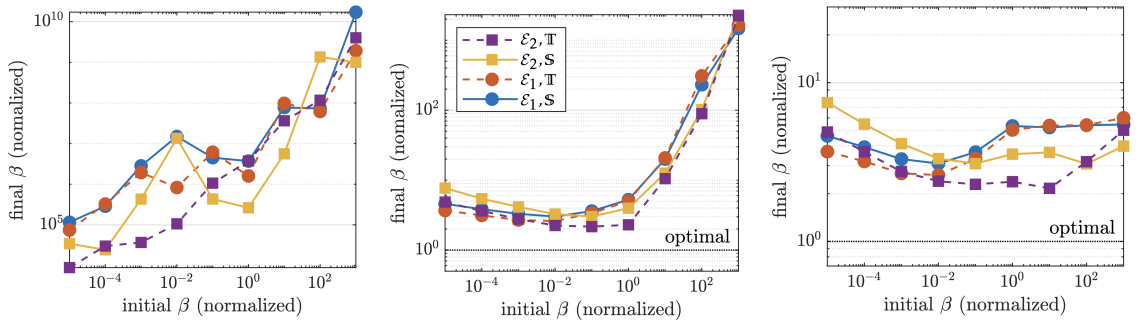

(a) Feasibility check with R (b) Feasibility check with R^α_β (c) Feasibility check with R^α_β, decreasing strategy for β

Figure 4.5: Final value of β obtained for different initialization of $\beta_{(0)}$, normalized with respect to the threshold value $\overline{\beta}$. The proposed method is tested against problems (4.19) and results are averaged over 1000 runs.

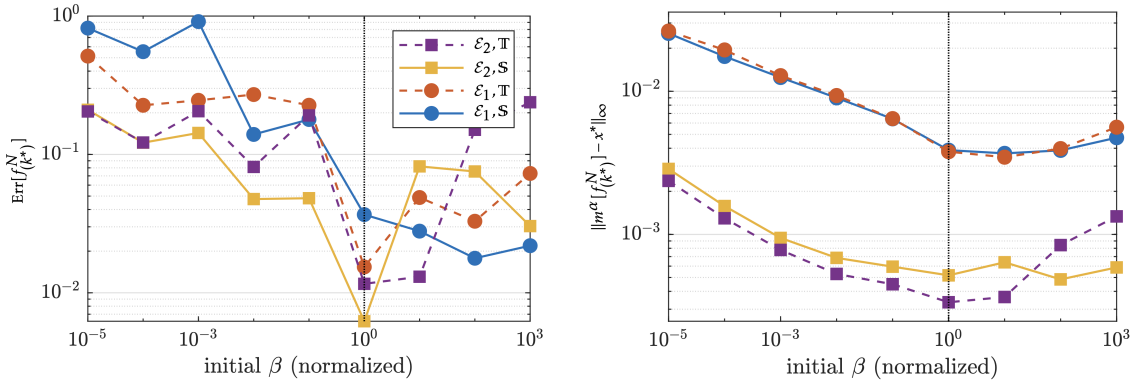

(a) Final average ℓ_2-error $\mathrm{Err}[f^N_{(k^*)}]$ (b) Final ℓ_∞-error $\|m^\alpha[f^N_{(k^*)}]-x^*\|_\infty$

Figure 4.6: Accuracy obtained for different initialization of $\beta_{(0)}$ normalized with respect to the threshold value $\overline{\beta}$. Feasibility check with R^α_β and decreasing strategy for β are used in the computation. The proposed method is tested against problems (4.19) and results are averaged over 1000 runs, but only successful runs are considered in the statistics.

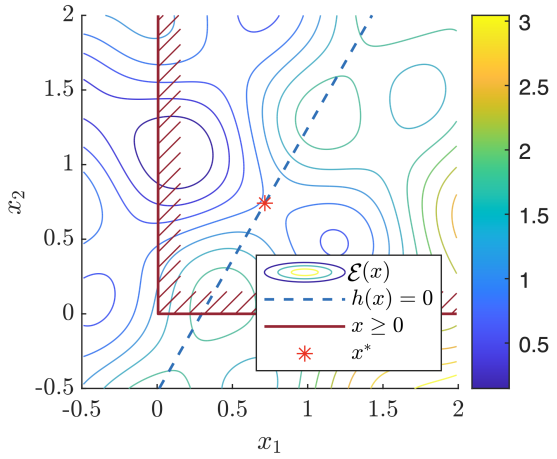

Figure 4.7: Illustration of a constrained optimization problem of type (4.21) with d = 2. The figure shows: the objective function $\mathcal{E}$, the solution x^*, the set of points satisfying the equality constraint $h(x) = H^\top x + h^0 = 0$ and the inequality constraint $x > 0$.

which the convergence rate is known to be dimension dependent [16,41]. Therefore, in this section, we also employ anisotropic exploration and compare the algorithm performance between the two strategies. We investigate, in particular, how sensitive the algorithm performance is with respect to the choice of the parameter σ, which regulates the exploration strength in the CBO dynamics.

The constrained optimization problems we consider are non-convex and randomly generated. We follow the procedure suggested in [77] and start by constructing a quadratic optimization problem of the form

$$\min_{x\in\mathbb{R}^d} \frac{1}{2}x^\top Ax - b^\top x \quad \text{subject to} \quad H^\top x + h^0 = 0,\ x \geq 0, \tag{4.20}$$

with $A \in \mathbb{R}^{d\times d}, H \in \mathbb{R}^{d\times p}, h \in \mathbb{R}^p$, for which the solution x^* is known by construction. The number p of equality constraint is given by $p = \lfloor d/2 \rfloor$. In order to make the problem non-convex, we add local minima belonging to the feasible set. Let $y^\perp$ and $y^\parallel$ be the projection of $y \in \mathbb{R}^d$ to $\ker(H^\top)$ and $\mathrm{span}(H^\top)$ respectively. We introduce

$$\gamma(x) = \frac{\cos(2\pi|(x-x^*)^\perp|)}{2} + \frac{\cos(\pi + 2\pi|(x-x^*)^\parallel|)}{2}$$

and consider the new constrained problem with same global solution x^*

$$\min_{x\in\mathbb{R}^d} \frac{1}{2}x^\top Ax - b^\top x + \gamma(x) \quad \text{subject to} \quad H^\top x + h^0 = 0,\ x \geq 0. \tag{4.21}$$

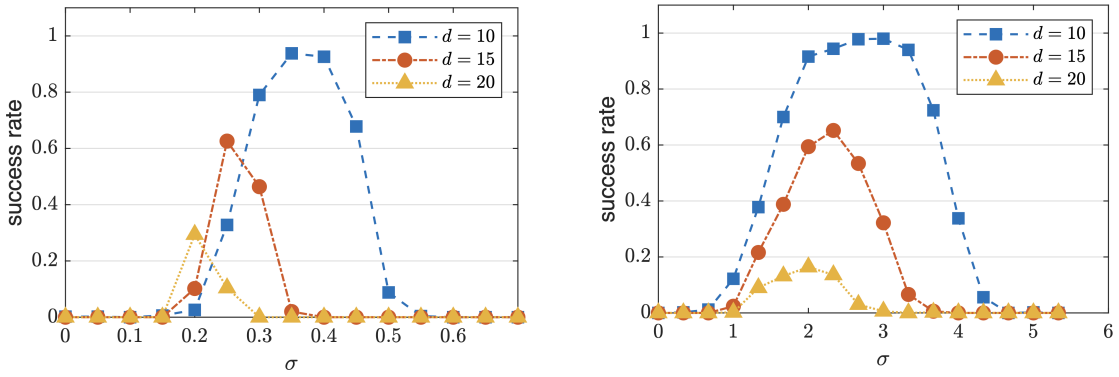

(a) Isotropic exploration

(b) Anisotropic exploration

Figure 4.8: Success rate as a function of the diffusion parameter σ for the two different exploration processes. The tests are run 500 times to solve problems (4.21).

We refer to Figure 4.7 for an illustration of a problem generated with the above procedure for d = 2 and $p = 1$.

For the penalization strategy we use ℓ_1-penalization $r = \|H^\top x + h^0\|_1$, which is known to be exact with threshold value $\overline{\beta} \approx 1$. We use the same algorithmic parameters as in the previous section, but test the algorithm for different values of σ. We initialize $N = 500$ particles uniformly sampled over $[0,2]^{\mathrm{d}}$ and consider a run successful whenever $\|m^\alpha[f^N_{(k^*)}] - x^*\|_\infty \leq 0.25$, $k^* = 300$ being the total number of iterations. In the following, we fix $\beta_{(0)} = 10^{-2}$.

In the first experiment, we consider three different space dimensions, d = 5, 10, 15 and test the algorithm with different values for σ: $\sigma \in [0, 0.7]$ for anisotropic diffusion and $\sigma \in [0, 5.25]$ for anisotropic diffusion. As already noted in [19], the optimal value for σ depends on the space dimension when isotropic diffusion is used, whereas employing anisotropic diffusion in CBO dynamics allows for a less restrictive choice of σ, see Figure 4.8. We note that anisotropic diffusion also leads to a better performance in terms of success rate for the problem considered.

Finally, we observe that the algorithm performance decreases as the problem dimension increases. A final experiment with anisotropic diffusion and $\sigma = 2$, shows that the success rate decreases sensibly for d > 10, see Figure 4.9. In particular, as shown in Figure 4.9b, the adaptive mechanism for the penalty parameter β becomes less effective.

4.4 Discussion

In this chapter, we showed how the CBO dynamics can be coupled to an additional adaptive strategy for an algorithm parameter in the context of exact penalization. Being

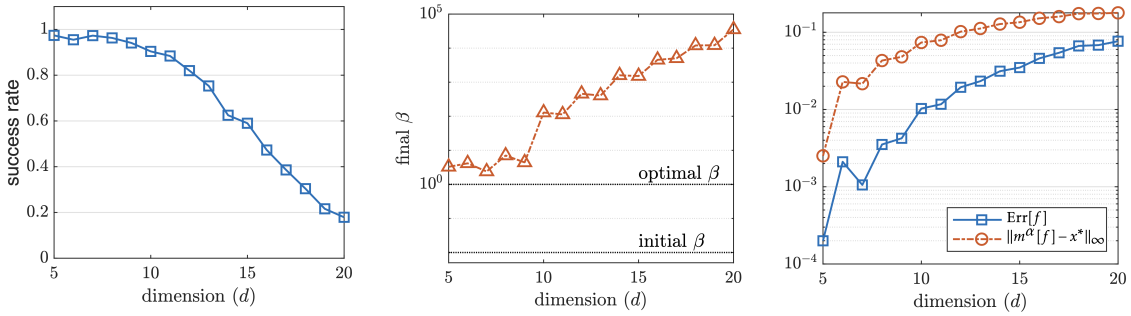

(a) Success rate (b) Final penalty parameter (c) Accuracy for successful runs

Figure 4.9: Algorithm performance obtained for different dimensions d of the search space for problem (4.21). Anisotropic exploration is used and $\sigma = 2$. Results are averaged over 1000 runs.

the optimal value of the penalty parameter unknown, the adaptive strategy leverages the information given by the whole particle ensemble to tune the penalty parameter during a single algorithm run.

We adapted the mean-field convergence analysis performed for the CBO dynamics to the proposed algorithm. Thanks to the CBO convergence rate, we have also been able to analytically study the properties of the adaptive strategy. We validated the performance of the suggested algorithm and the theoretical analysis by testing the algorithm against several benchmark problems. For moderate search space dimension $d \leq 10$, the algorithm performs well and is able to adapt the penalty parameter to a suitable value, without overshooting the threshold $\overline{\beta}$.

In the next chapter, we show how to leverage the ensemble information not to adapt a single parameter, but to solve many optimization problems in parallel in the context of multi-objective optimization.

Chapter 5

Multi-objective optimization problems

In real-life applications, optimization problems often require considering two or more objectives at the same time, which may be in conflict with each other. While designing an engine, for instance, maximizing performances and minimizing fuel consumption are two opposing aims. The optimization procedure must balance these two objectives and return a solution that is a compromise between best performance and engine efficiency. Also, such solution is typically not unique as it depends on how important we consider one objective with respect to the other (or others).

To formulate a multi-objective problem mathematically, we introduce the notion of Edgeworth-Pareto optimality. Assume $\overline{\mathcal{E}}_\ell : \mathbb{R}^{\mathrm{d}} \to \mathbb{R}$, $\ell = 1, \dots, m$ to be the $m \geq 2$ objective functions given, and that we aim to minimize. A solution to

$$\text{minimize } \overline{\mathcal{E}}(x) := (\overline{\mathcal{E}}_1(x), \dots, \overline{\mathcal{E}}_m(x))^\top \tag{5.1}$$

is intended to be the set of points satisfying the following notion of optimality.

Definition 5.1 (Edgeworth-Pareto optimality)**.** *A point x^* is (strongly) Edgeworth-Pareto (EP) optimal if and only if there is no other point x such that*

$$\overline{\mathcal{E}}_\ell(x) \leq \overline{\mathcal{E}}_\ell(x^*) \qquad \text{for all} \quad \ell = 1, \dots, m.$$

A multi-objective problem may have an infinite number of EP optimal point. In practice, solving (5.1) only requires to find a finite subset of EP optimal points which is representative of the entire solution set. Going back to our initial example, the aim is to find a finite number of different optimal designs, covering the spectrum between best engine performance and smallest fuel consumption. To qualitative estimate how well a finite set of computed solutions describes the entire set of EP optimal points, we will introduce later a concept known in the literature as *diversity* [32, 61]. It is essential to bear

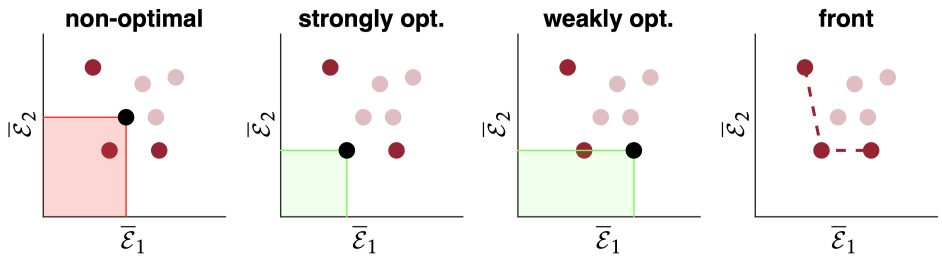

Figure 5.1: Illustrative example of Edgeworth-Pareto optimality among a set of points on a two-dimensional objective space. In the first three pictures, from left to right, the black dot indicates a point which is, respectively, non EP optimal, strongly EP optimal, and weakly EP optimal (see Definitions 5.1 and 5.2). On the last illustration, the set of three EP optimal points forming the so-called Pareto front is highlighted.

in mind that, from the multi-objective prospective, all EP optimal points are equally good and that, ultimately, we do not seek to find just one single solution.

Optimization routines for (5.1) either compute one EP solution at the time, or attempt to compute an entire set of optimal points at once. The latter is the case of popular heuristics in multi-objective optimization like NSGA-II [26] and MOEA/D [83]. We refer to the recent survey [23] for more details on such particle-based heuristics and, in particular, on Evolutionary Algorithms. Many classical mathematical programming routines such as descend methods, Newton's method, and trust region methods have been adapted in the settings of multi-objective problems, and they typically compute one single EP optimal point per run [37, 44]. For more details on deterministic mathematical programming algorithms in multi-objective optimization we refer to the survey [30].

Following the strategy of the mentioned heuristics, we adapt in this chapter the CBO iteration to compute a set of EP optimal points with a single run. This is done by partially distributing the optimization task among the particle system. Through a scalarization procedure, we first translate (5.1) into a set of parameterized, scalar sub-problems which are then solved simultaneously by the ensemble. An efficient communication strategy makes the proposed method computationally more efficient than performing several independent runs of the single-objective CBO optimizer. To seek diversity, we also couple the CBO dynamics with an additional adaptive strategy in the parameters space.

The chapter is organized as follows. In Section 5.1 we illustrate the scalarization strategy and introduce a quantitative notion of diversity based on two-body energy potentials. Next, we propose in Section 5.2 different CBO dynamics and parameters adaptive strategies to solve (5.1). Section 5.3 is devoted to the theoretical analysis of the methods via mean-field approximation. Numerical experiments for bi- and tri-objective benchmark optimization problems are illustrated in Section 5.4. Finally, in Section 5.5 we collect additional remarks and recall the most important ideas and results of the chapter.

5.1 Scalarization strategy and energy-based diversity measures

Scalarization is a popular technique in multi-objective optimization because it allows to trace back the problem to the well-studied single-objective settings. Among the different scalarization techniques available in the literature, see, for instance, [62], we consider for simplicity the weighted semi-norm approach. This is because, unlike other strategies such as the ε-Constraint Method [29], scalarization by means of weighted semi-norms leads to unconstrained single objective sub-problems.

Let Ω_m denote the m-dimensional unit simplex,

$$\Omega_m = \left\{ w \in \mathbb{R}^m \,:\, w_\ell \geq 0 \text{ for all } \ell = 1, \dots, m, \text{ and } \sum_{\ell=1}^{m} w_\ell = 1 \right\}.$$

We will assume throughout the chapter to know a lower bound of the objective functions. Without loss of generality, in particular, we assume $\overline{\mathcal{E}}(x) > 0$ component-wise for all $x \in \mathbb{R}^{\mathrm{d}}$. Let $w \in \Omega_m$ be a vector of weights. For any $p \in [1, \infty]$ we define the scalarized (by means of weighted semi-norm) objective $\mathcal{E}^p(w;\cdot) : \mathbb{R}^{\mathrm{d}} \to \mathbb{R}$ as

$$\mathcal{E}_p(w;x) := \begin{cases} \left(\sum_{\ell=1}^m w_\ell |\overline{\mathcal{E}}_\ell(x)|^p \right)^{1/p} & \text{if } p \in [1, \infty), \\ \sup_{\ell=1,\dots,m} w_\ell |\overline{\mathcal{E}}_\ell(x)| & \text{if } p = \infty. \end{cases} \tag{5.2}$$

For a given $p \in [0, \infty]$, the scalar sub-problem corresponding to $w \in \Omega_m$ then reads

$$\text{minimize } \mathcal{E}_p(w;x) \quad \text{subject to} \quad x \in \mathbb{R}^{\mathrm{d}}. \tag{5.3}$$

Before showing the relation between (5.1) and (5.3) we introduce a weaker version of EP optimality.

Definition 5.2 (Weak Edgeworth-Pareto optimality)**.** *A point x^* is weakly Edgeworth-Pareto (EP) optimal if and only if there is no other point x such that*

$$\overline{\mathcal{E}}_\ell(x) < \overline{\mathcal{E}}_\ell(x^*) \qquad \text{for all} \quad \ell = 1, \dots, m.$$

Clearly, weakly EP optimality follows from EP optimality. We recall from [62] the following classical results.

Theorem 5.3 ([62, Theorem 5.25, Corollary 11.21])**.** *Assume $\overline{\mathcal{E}}$ to be component-wise positive.*

a) Let $p \in [1, \infty]$. If x^* is a solution to (5.3) *for some $w \in \Omega_m$, then x^* is a weakly EP optimal point.*

b) Let $p = \infty$. A point x^* is weakly EP optimal if and only if x^* is a solution to (5.3) *for some $w \in \Omega_m$.*

c) *Let $p = \infty$ and assume all sub-problems* (5.3) *attain a unique minimum. Then, x^* is EP optimal if and only if x^* is the solution to* (5.3) *for some $w \in \Omega_m$.*

Therefore, by solving (5.3) for different vectors of weights we can find weakly EP optimal points. Since the aim is to compute a finite set of EP optimal points which approximates the entire (possibly uncountable) set of solutions, a natural choice is to pick $p = \infty$ due to the if-and-only-if statement in Theorem 5.3 (also known and *Chebyschev* scalarization strategy). Despite this, oftentimes this is not the preferred choice as $\mathcal{E}_\infty(w;\cdot)$ is in general non-differentiable, see (5.2). Exactly as in the context of exact penalization, being CBO methods gradient-free allows us to choose $p = \infty$ and, so, to use the most convenient scalarization strategies. We remark that the common choice $p = 1$ does not allow to recover the full set of EP optimal points for general non-convex multi-objective problems.

Assume we are allowed to solve $N \in \mathbb{N}$ scalarized sub-problems (5.3). Which sub-problems should we choose? Or, equivalently, which vectors of weights $\{w^i\}_{i=1}^N \subset \Omega_m$ should we pick? To answer this question, let us go back to the concept of *diversity*. The aim of the multi-objective algorithm consist of providing the decision maker a set of EP optimal solutions which represents at best all the possible optimal objective values one can reach. To be more precise, let $F \subset \mathbb{R}^m$ be the so-called *Pareto* front

$$F := \left\{ \overline{\mathcal{E}}(x^*) \,:\, x^* \text{ is EP optimal} \right\}. \tag{5.4}$$

The best choice of parameters w^i would than be the one such that the image of the obtained solutions well approximates the Pareto front. Qualitatively, this choice corresponds to the points $\{x^{*,i}\}_{i=1}^N$ whose images $\{\overline{\mathcal{E}}(x^{*,i})\}_{i=1}^N$ are uniformly distributed over F. The optimal choice of parameters is therefore problem-dependent as it depends on the Pareto front geometry, see Figure 5.2.

To quantify how well the computed solutions cover the front, or, are *diverse*, many metrics have been proposed in the literature and there is no unique agreement on which one is the most convenient. Among the many, we mention *hypervolume contribution* [84], *crowding distance* [26] and the *Riesz s-energy* [32, 68]. The choice of a diversity metric is important not only to assess the algorithm performance in benchmark problems, but also to design the algorithm itself. Recently, particular attention has been drawn to energy-based diversity measures thanks to their flexibility and theoretical properties [23, 32]. Since energy functionals naturally fit into the CBO particle framework, and, in particular, into its mean-field analysis, we mathematically quantify the concept of diversity in terms of energy-based diversity metrics.

Assume the computed solution to (5.1) comprises N points $\{x^{*,i}\}_{i=1}^N \subset \mathbb{R}^d$ with $\rho^{*,N} = (1/N)\sum_{i=1}^N \delta_{x^{*,i}} \in \mathcal{P}(\mathbb{R}^d)$ being the associated empirical measure. For a given interaction kernel $U : \mathbb{R}^m \to (-\infty, \infty]$ we define the functional $\mathcal{U} : \mathcal{P}(\mathbb{R}^m) \to (-\infty, \infty]$ as

$$\mathcal{U}[\zeta] := \frac{1}{2} \iint U(z-y)\, d\zeta(z)\, d\zeta(y). \tag{5.5}$$

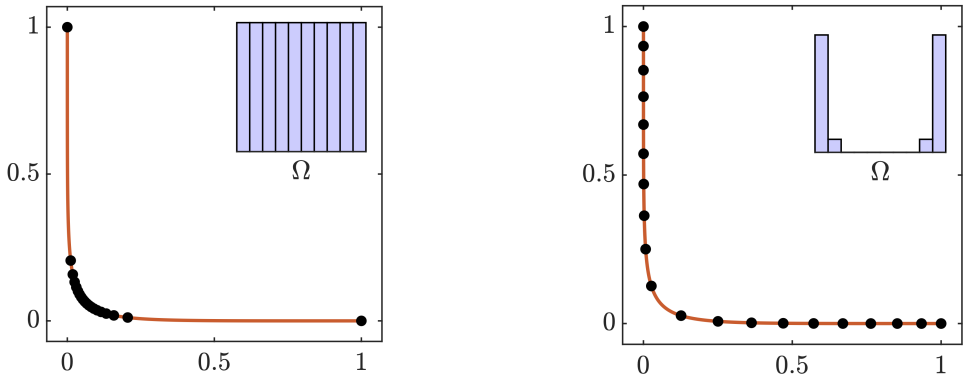

(a) High-energy configuration

(b) Low-energy configuration

Figure 5.2: Illustration of two particle systems in the image space. Each of the systems comprises solutions of $N = 20$ scalarized problems (5.3). On the left, the problems' parameters are taken uniformly over the simplex, while, on the right, the parameters are taken such that the Riesz 1-energy is minimized. The histograms illustrate the parameters distribution on the simplex. As expected, the low-energy configuration better describes the Pareto front (in red), that is, the low-energy configuration is more diverse.

Provided continuity of $\overline{\mathcal{E}}$, we consider the push-forward measure $\overline{\mathcal{E}}_{\#}\rho^{*,N} \in \mathcal{P}(\mathbb{R}^m)$ of $\rho^{*,N}$, which is defined by:

$$\overline{\mathcal{E}}_{\#}\rho^{*,N}(A) := \rho^{*,N}\left(\overline{\mathcal{E}}^{-1}(A)\right)$$

for any Borel set $A \subset \mathbb{R}^m$.

We note that, if U models a repulsive potential, the total energy of the particle system in the image space, that is, $\mathcal{U}[\overline{\mathcal{E}}_{\#}\rho^{*,N}]$, will be larger if particles are close to each other in the image space. On the other hand, the more the particle are distanced from one another, the smaller the total energy is. Therefore, quantity $\mathcal{U}[\overline{\mathcal{E}}_{\#}\rho^{*,N}]$ can be considered as a metric for the diversity of sets of points $\{x^{*,i}\}_{i=1}^{N}$, see also Figure 5.2 for an illustrative example. The already mentioned Riesz s-energy

$$U_R(z) = \frac{1}{|z|^s}, \quad \text{for } s > 0$$

is an example of a short-range repulsive potential. Another choice, which is non-singular, is given for instance by the Morse potential

$$U_M(z) = e^{-C|z|} \quad \text{for some } C > 0.$$

Restricting the probability measure to be supported on the set of EP optimal points

$F_x := \overline{\mathcal{E}}^{-1}(F)$, the best set of $N \in \mathbb{N}$ EP optimal points is the solution of the problem

$$\text{minimize } \mathcal{U}[\overline{\mathcal{E}}_{\#}\rho^{N,*}] \quad \text{subject to} \quad \rho^{*,N} = (1/N)\sum_{i=1}^{N} \delta_{x^{*,i}}, \text{ with } \{x^{*,i}\}_{i=1}^{N} \subset F_x. \tag{5.6}$$

In the following, we propose a strategy to solve (5.6) via consensus-based optimization.

5.2 Consensus-based optimization and adaptive strategies

In this section, we present several CBO dynamics that leverage the particle system to solve N different sub-problems simultaneously. We will first assume to the vector of weights of the scalar sub-problems to be fixed, and then couple the suggested dynamics with an adaptive strategy for the parameters to improve the diversity of the system. The aim, in particular, is to select a set of scalarized sub-problems whose corresponding set of solutions attains low energy $\mathcal{U}$ value. We derive an adaptive strategy which approximates the energy gradient flow for bi-objective problems and, then, we suggest a second adaptive strategy which can be applied to general multi-objective problems.

5.2.1 CBO for parameterized sub-problems

Assume we want to solve $N \in \mathbb{N}$ scalar sub-problems (5.3) determined by the set of parameters $\{w^i\}_{i=1}^N \subset \Omega_m$. We aim of doing that with a single system of N particles $\{X^i_{(k)}\}_{i=1}^N \subset \mathbb{R}^d$. To distribute the task among the particles, we couple each of them with a problem and consider the tuples $(X^i_{(k)}, w^i)$ so that every particle has its own scalarized objective $\mathcal{E}_p(w^i;\cdot)$ to minimize. In single-objective CBO methods, the update rule drives each particle towards a consensus point which is considered to be in a good location for the objective function. Now, since every particle targets a different scalarized objective, the consensus point should differ for each particle (talking about *consensus* seems not appropriate anymore, but we keep this term to stress the parallelism with the single-objective settings).

Let $\rho \in \mathcal{P}(\mathbb{R}^d)$ be a particle probability measure. For any problem $w \in \Omega_m$, we define the corresponding consensus point as

$$m^\alpha_w[\rho] := \frac{\int x \exp(-\alpha \mathcal{E}_p(w;x))\, d\rho(x)}{\int \exp(-\alpha \mathcal{E}_p(w;x))\, d\rho(x)}, \tag{5.7}$$

for some parameter $\alpha > 0$. At step $k \geq 0$, we denote the empirical measure associated with the particle system $\{X^i_{(k)}\}_{i=1}^N$ as $\rho^N_{(k)}$. The Multi-objective Consensus Based Optimization (M-CBO) update rule with anisotropic diffusion reads

$$X^i_{(k+1)} = X^i_{(k)} + \lambda\tau\left(m^\alpha_{w^i}[\rho^N_{(k)}] - X^i_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha_{w^i}[\rho^N_{(k)}] - X^i_{(k)}\right) \odot \theta^i_{(k)}, \tag{5.8}$$

with $\theta^i_{(k)} \sim \mathcal{N}(0, I_d)$ for all $i = 1, \dots, N$, as in the single-objective CBO dynamics.

According to (5.8), every particle moves towards a possibly different area of the search space. Communication between the particles happens via the computation of the consensus points. Indeed, the consensus-point for the i-th particle is computed by taking into account the whole particle system. Thanks to the Boltzmann-Gibbs weights in (5.7), particles attaining large values of the i-th objective $\overline{\mathcal{E}}(w^i; \cdot)$ have little impact in the computation of the i-th objective point.

Remark 5.4. *It is interesting to note that the proposed M-CBO update rule can be applied to any task where it is required to solve many parameterized sub-problems at the same time. Indeed, update rule* (5.8) *does not make use of the particular structure of sub-problems* (5.3)*. In practice, though, the heuristic derivation implicitly assumes the following scenarios.*

- *For updating the N particles, we are required to compute $\mathcal{E}_p(w^i; X^j_{(k)})$ for all $i, j = 1, \dots, N$, possibly leading to an update step of complexity $\mathcal{O}(N^2)$. Therefore, the evaluation of the sub-objective should be computationally cheap. This is the case of the considered sub-objectives* (5.2)*, as we only need to evaluate the multi-objective function $\overline{\mathcal{E}}$ N times per iteration, once per particle.*

- *Loosely speaking, the sub-problems should not differ too much from one another. By using* (5.7)*, we implicitly assume that the i-th particle can obtain valuable information from the rest of the ensemble. For the scalarized sub-objectives* (5.2)*, we can expect continuity with respect to the parameter w. Therefore, particles with similar parameters can share valuable information regarding their objective function landscapes.*

5.2.2 Parameters adaptation for bi-objective problems

A classical strategy to reduce the energy of a particle system is to let the particles evolve according to the associated vector field. We now derive an adaptive strategy for the parameters which aims to approximate a gradient flow dynamics in the image space. When it comes to gradient-based dynamics, working with continuous-in-time systems is the most convenient choice. Therefore, during the following formal derivation, we assume the vector of weights to be continuous-in-time processes. To this end, we introduce $W^i_t \in \Omega_m$ for all $i = 1, \dots, N$ and $t \geq 0$. The algorithm's update rule for $W^i_{(k)}$ will be given by a suitable discretization of the continuous dynamics.

Let $\{Z^i_t\}_{i=1}^N \subset F$ be a set of particles in the objective space evolving over the Pareto front, with $\zeta^N_t = (1/N) \sum_{i=1}^N \delta_{Z^i_t}$. For the sake of the formal derivation, let us assume U to be smooth and F sufficiently regular such that the projection $\Pi_{T(z,F)}$ towards the tangential space $T(z, F)$ at $z \in F$ is well-defined. Recall that, to obtain a diverse approximation of the Pareto front F, particles should reach a low-energy configuration over F. In theory, following the particles' descent directions, we can lower the energy of the system. Recall

that $\mathcal{U}[\zeta_t^N]$ takes the form of

$$\mathcal{U}[\zeta_t^N] = \frac{1}{2}\iint U(z-y)\,d\zeta_t^N(z)d\zeta_t^N(y) = \frac{1}{2N^2}\sum_{i,j=1}^{N} U(Z_t^i - Z_t^j).$$

The point $Z_t^i \in \mathbb{R}^m$, in particular, is subject to the potential

$$\int U(Z_t^i - y)\,d\zeta_t^N(y) = \frac{1}{N}\sum_{j=1}^{N} U(Z_t^i - Z_t^j),$$

with associated vector field at Z_t^i given by

$$\int \nabla U(Z_t^i - y)\,d\zeta_t^N(y) = \frac{1}{N}\sum_{j=1}^{N} \nabla U(Z_t^i - Z_t^j).$$

We obtain in this way the (projected) descent dynamics

$$\frac{d}{dt}Z_t^i = \Pi_{T(Z_t^i,F)}\left(-\frac{1}{N}\sum_{j=1}^{N}\nabla U(Z_t^i - Z_t^j)\right). \tag{5.9}$$

For Chebyshev scalarization ($p = \infty$) and assuming solution uniqueness, Theorem 5.3 (c) allows us to associate to any scalarized sub-problem $w \in \Omega_m$ an EP optimal point $x^*(w)$. This gives a parameterization of the Pareto front via $z^* : \Omega_m \to F$, with $z^*(w) := \overline{\mathcal{E}}(x^*(w))$. In our settings, Z_t^i is to be interpreted as the image of the exact solution of the i-th problem W_t^i, that is, $Z_t^i = z^*(W_t^i)$. In a situation where, again, the solution operator z^* is sufficiently regular, we can finally perform a change of variable and obtain from (5.9) an adaptive strategy for W_t^i:

$$\frac{d}{dt}W_t^i = \left(Dz^*(W_t^i)\right)^+ \Pi_{T(Z_t^i,F)}\left(-\frac{1}{N}\sum_{j=1}^{N}\nabla U(Z_t^i - Z_t^j)\right) \tag{5.10}$$

where M^+ is the pseudo inverse of M and $Dz^*(w)$ is the differential at $w \in \Omega_m$. We refer to [51, Ch. IV, Sec. 5] for more details on differential equations and numerical integration on manifolds via projection.

Since the parameterization z^* as the well the tangential spaces $T(z,F)$ are unknown, in practice (5.10) needs to be approximated. We do that by considering the simple case where $m = 2$ and the front is a linear segment coinciding with the unit simplex Ω_2 itself. For Chebyshev scalarization, one can show that the parameterization mapping the simplex into the front is linear and given by

$$z_L^*(w) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} w$$

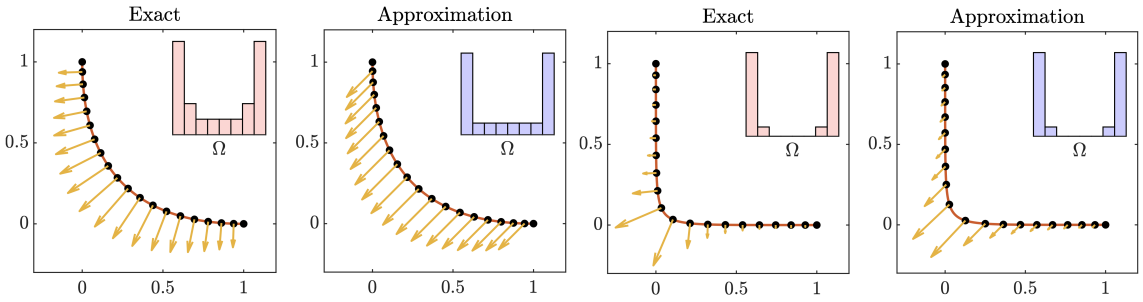

(a) Problem Lamé $\gamma = 0.5$ (b) Problem Lamé $\gamma = 0.25$

Figure 5.3: Final configuration over the objective space $Z_t^i = z^*(W_t^i)$ of $N = 20$ particles evolved according to (5.10) (exact dynamics) and (5.11) (approximated dynamics) under 1-Riesz binary potential. The arrows indicate the total potential forces acting on the particles. Two different types of fronts are taken into account. Numerical experiments are performed by means of an explicit Euler scheme with $\Delta t = 10^{-8}, T = 0.01$. Histograms show the final distribution over the unit simplex Ω_2. Despite the many approximations, dynamics (5.11) allows us to recover good descriptions of the Pareto fronts.

Similarly, if $F = \Omega_2$, the projection to the tangential space for $z \in \Omega_2, z \neq (1,0)^\top, (0,1)^\top$ is explicitly known and given for any $u \in \mathbb{R}^2$ by

$$\Pi_{T(z,\Omega_2)}(u) = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} u\,.$$

We note that for $m > 2$, under assumption $F = \Omega_m$ the mapping z^* is non-injective as for the case $m = 2$. Therefore, we cannot apply for $m > 2$ the change of variable which we used to obtain (5.10) from (5.9). For bi-objective problems, performing approximations $z^* \approx z_L^*$ and $\Pi_{T(z,F)} \approx \Pi_{T(w,\Omega_2)}$ then leads to

$$\frac{d}{dt} W_t^i = \Pi_{T(W_t^i,\Omega_2)} \left(\frac{1}{N} \sum_{j=1}^{N} \nabla U(Z_t^i - Z_t^j) \right) \tag{5.11}$$

where we used that for $u \in \mathbb{R}^2$

$$(Dz_L^*(w))^\top \Pi_{T(z,\Omega_2)}(-u) = -\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}^\top \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} u = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} u = \Pi_{T(z,\Omega_2)}(u).$$

Numerical experiments performed in the settings where the parameterization $z^* : \Omega_2 \to F$ is explicitly known show that (5.11) can be a valid approximation to obtain diverse distributions over the Pareto front, see Figure 5.3.

Finally, we discretize (5.11) to derive an update rule for $\{W^i_{(k)}\}_{i=1}^N$. We do that via explicit, projected Euler scheme with time step τ. Moreover, to compute the vector fields in the image space, we substitute the image of the true solution Z^i_t of the i-th problem with the current multi-objective value of the i-th particle $\overline{\mathcal{E}}(X^i_{(k)})$. The update rule derived with the above procedure reads

$$W^i_{(k+1)} = \Pi_{\Omega_2}\left(W^i_{(k)} + \nu \frac{\tau}{N} \sum_{j=1}^{N} \nabla U\left(\overline{\mathcal{E}}(X^i_{(k)}) - \overline{\mathcal{E}}(X^j_{(k)}) \right) \right) \tag{5.12}$$

where the additional parameter $\nu > 0$ is introduced to tune the velocity of the adaptive procedure with respect to the CBO dynamics (5.8).

Remark 5.5. *We note that the projection step may in principle be avoided in the update rule* (5.12) *and used only in the evaluation of the scalarized objective $\mathcal{E}(w;\cdot)$. This would allow, though, $\{W^i_{(k)}\}_{i=1}^N$ to evolve far from the simplex, with potentially a negative effect on the dynamics due to long-range interactions. We leave this alternative strategy for future investigations.*

5.2.3 Parameters adaptation for general multi-objective problems

In the previous section, we showed how to simulate a gradient flow dynamics in the objective space by making use of the Chebyshev parameterization given by Theorem 5.3 (c). This was possible for the case $m = 2$ where a suitable linearization allowed to approximate the gradient flow evolution. To derive an update rule for general multi-objective problems $m \geq 2$, we now take a different perspective and focus on repulsive dynamics taking place, directly, on the unit simplex Ω_m. We follow here the literature on active particle systems and consider a time-continuous adaptation of the parameters $\{W^i_t\}_{i=1}^N, t \geq 0$. The update rule for $\{W^i_{(k)}\}_{i=1}^N$ is then obtained through an explicit numerical method applied to the continuous dynamics.

As before, let $z^* : \Omega_m \to F$ be the map associating to any problem w the image of the corresponding EP optimal point $z^*(w) = \overline{\mathcal{E}}(x^*(w))$. The map existence is ensured under settings of Theorem 5.3 (c), but it is typically not bijective for $m > 2$. From the definition of the scalarized sub-problems (5.3) we can expect $z^*(w)$ to depend continuously on w. To reach a uniform distribution over the front, we derive binary repulsive dynamics which distance two particles' parameters W^i_t, W^j_t whenever the corresponding solutions $Z^i_t = z^*(W^i_t)$ and $Z^j_t = z^*(W^j_t)$ (or an approximation of them) are too close in the objective space. We recall the dynamics is constrained to the unit simplex Ω_m, and so the vector field needs to be projected into the tangential space $T(W^i_t, \Omega_m)$.

For an interaction kernel $\mathcal{K} : \Omega_m \times \Omega_m \times \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}^m$ (which we will define explicitly

later), the paradigmatic model then reads for all $i = 1, \dots, N$

$$\frac{d}{dt} W_t^i = \Pi_{T(W_t^i, \Omega_m)} \left(\frac{1}{N} \sum_{j=1}^{N} \mathcal{K}\left(W_t^i, W_t^j, Z_t^i, Z_t^j\right)(W_t^i - W_t^j) \right). \tag{5.13}$$

System (5.13) is a first-order microscopic model sharing many features with swarming and opinion models which have been extensively studied in the last decades [1, 3, 17, 25, 71]. For more references in swarming models in bounded domains we refer in particular to [4, 20, 34]. To describe short-range repulsion that depends on the binary potential U, we suggest an interaction kernel of the form

$$\mathcal{K}\left(W_t^i, W_t^j, Z_t^i, Z_j^i\right) = -|\nabla U(Z_t^i - Z_t^j)| \|W_t^i - W_t^j|^{-1}. \tag{5.14}$$

Due to the normalizing factor $|W_t^i - W_t^j|^{-1}$, the strength of the displacement between the i-th and the j-th particle has exactly magnitude $|\nabla U(z^*(W_t^i) - z^*(W_t^j))|$. This corresponds to the gradient flow-like dynamics (5.11) derived in the previous section for bi-objective problems. Here, though, the displacement direction is simply determined by $W_t^i - W_t^j$. We note that a generalization of the well-known Cucker-Smale model [25] with nonlinear velocity coupling has been investigated in [47]. In particular, settings of [47] cover the case where the velocities are normalized, similarly to (5.13) equipped with (5.14). The suggested choice of $\mathcal{K}$ is made under the assumption of U being purely repulsive. If U models, for instance, long-range attraction effects, (5.14) can be modified to possibly assume negative values when particles should attract each other.

Finally, we derive from (5.13), (5.14) an update rule for $W_{(k)}^i$ by means of an explicit Euler discretization with step size $\tau > 0$. As we did for the case $m = 2$, the unknown exact solution $Z_t = z^*(W_t^i)$ of the i-th scalarized sub-problem is substituted in the algorithm by the i-th particle location in the objective space $\overline{\mathcal{E}}(X_{(k)}^i)$. The adaptive strategy for general multi-objective problems then reads

$$W_{(k+1)}^i = \Pi_{\Omega_m} \left(W_{(k)}^i - \nu \frac{\tau}{N} \sum_{j=1}^{N} \left| \nabla U\left(\overline{\mathcal{E}}(X_{(k)}^i) - \overline{\mathcal{E}}(X_{(k)}^j)\right) \right| \frac{W_{(k)}^i - W_{(k)}^j}{|W_{(k)}^i - W_{(k)}^j|} \right) \tag{5.15}$$

where, $\nu > 0$, again, is an additional parameter which will allow us to regulate the time scale of the adaptive procedure (5.15) with respect to the CBO dynamics.

Remark 5.6. *In [35], authors perform a mean-field analysis of particle systems subject to potentials in bounded domains, i.e. settings similar to* (5.12) *and* (5.15)*. The authors show that the system can evolve towards a non-optimal equilibrium in which a certain amount of particles, or mass, is concentrated at the boundary of the domain. Inspired by [36], where a diffusive*

component is added to overcome this effect, we may modify the proposed dynamics to include a stochastic component.

Let $\Phi = \Phi(X^i_{(k)}, X^j_{(k)}, W^i_{(k)}, W^j_{(k)})$ *denote the right-hand side of either of the suggested update rule* (5.12) *and* (5.15)*, the perturbed adaptive strategy then reads*

$$\begin{cases} W^i_{(k+1/2)} & = \Phi(X^i_{(k)}, X^j_{(k)}, W^i_{(k)}, W^j_{(k)}) \\ W^i_{(k+1)} & = \Pi_{\Omega_m}\left(W^i_{(k+1/2)} + \hat{\sigma}\,\xi^i_{(k)}\right) \end{cases} \tag{5.16}$$

where $\hat{\sigma} > 0$ *is a parameter controlling the variance and* $\xi^i_{(k)} \sim \mathcal{N}(0, I_m)$*. In Section* (5.4) *we will see how this stochastic behavior allows, indeed, to reach lower energy configurations in the performed numerical experiments.*

5.3 Convergence analysis in mean-field law

In multi-objective optimization, a common metric to assess the performance of an optimizer is given by the Generation Distance (*GD*) [82]. When the exact Pareto front F is available, *GD* measures the squared distance between the image of every EP optimal point computed and F. Let $\rho \in \mathcal{P}(\mathbb{R}^d)$ encode the computed set of solutions, the metric $GD : \mathcal{P}(\mathbb{R}^d) \to [0, \infty]$ is defined as

$$GD[\rho] := \left(\int \operatorname{dist}(\overline{\mathcal{E}}(x), F)^2 \, d\rho(x) \right)^{\frac{1}{2}}. \tag{5.17}$$

In this section, we perform a mean-field analysis of the proposed algorithms to investigate their convergence properties. In particular, we derive conditions on the multi-objective function $\overline{\mathcal{E}}$ and on the scalarized sub-problems (5.3) which lead to an exponential decay of *GD* up to a certain pre-fixed accuracy $\varepsilon > 0$.

5.3.1 Mean-field approximation of the coupled dynamics

Let us start by deriving mean-field approximations of the dynamics introduced in the previous sections. We recall each particle is characterized by its location $X^i_{(k)}$ on the search space $\mathbb{R}^d$ and the vector of weights $W^i_{(k)}$ taking values in the simplex Ω_m. We denote the (random) empirical probability measure associated to $\{X^i_{(k)}\}_{i=1}^N$ with $\rho^N_{(k)}$, while the one associated to the parameters $\{W^i_{(k)}\}_{i=1}^N$ with $\mu^N_{(k)}$. Also, we introduce the empirical distribution $f^N_{(k)} = (1/N)\sum_{i=1}^N \delta_{X^i_{(k)}} \otimes \delta_{W^i_{(k)}}$ of the coupled system $\{(X^i_{(k)}, W^i_{(k)})\}_{i=1}^N$.

Following the procedure illustrated in Chapter 2, we derive the mean-field model under the *propagation of chaos* assumption. For large ensemble sizes, $N \gg 1$, we assume factorization of the marginals, leading to $f^N_{(k)} \approx f_{(k)}$ for some $f \in \mathcal{P}(\mathbb{R}^d \times \Omega_m)$. As a consequence,

it also holds $\rho^N_{(k)} \approx \rho_{(k)} \in \mathcal{P}(\mathbb{R}^d)$ and $\mu^N_{(k)} \approx \mu_{(k)} \in \mathcal{P}(\Omega_m)$ with $\rho_{(k)}, \mu_{(k)}$ being the marginal distributions of $f_{(k)}$ with respect to the first and second component, respectively. Let the initial particle distribution be given by $f_{(0)} \in \mathcal{P}(\mathbb{R}^d \times \Omega_m)$. Under such approximation, the multi-objective CBO update rule (5.8) becomes independent on the index i and reads

$$\overline{X}_{(k+1)} = \overline{X}_{(k)} + \lambda\tau\left(m^\alpha_{\overline{W}_{(k)}}[\rho_{(k)}] - \overline{X}_{(k)}\right) + \sigma\sqrt{\tau}\left(m^\alpha_{\overline{W}_{(k)}}[\rho_{(k)}] - \overline{X}_{(k)}\right) \odot \overline{\theta}_{(k)}, \tag{5.18}$$

with $\overline{\theta}_{(k)} \sim \mathcal{N}(0, I_d)$ for all $i = 1, \dots, N$, as in the single-objective CBO dynamics.

If there is no interaction between the parameters $\{W^i_{(k)}\}_{i=1}^N$, the mean-field dynamics over the simplex is simply given by

$$\overline{W}_{(k+1)} = \overline{W}_{(k)} \tag{5.19}$$

for all $k \geq 0$. When the adaptive mechanism is employed, the binary interaction between the particles is also approximated by an interaction of mean-field type.

Consider update rule (5.12) for bi-objective problems. In this case, if we assume $\rho^N_{(k)} \approx \rho_{(k)}$ we obtain

$$\frac{1}{N}\sum_{j=1}^N \nabla U\left(\overline{\mathcal{E}}(X^i_{(k)}) - \overline{\mathcal{E}}(X^j_{(k)})\right) \approx \int \nabla U\left(\overline{\mathcal{E}}(X^i_{(k)}) - \overline{\mathcal{E}}(x)\right) d\rho_{(k)}(x)$$

for all $i = 1, \dots, N$, leading to the mean-field dynamics

$$\overline{W}_{(k+1)} = \Pi_{\Omega_m}\left(\overline{W}_{(k)} + \nu\tau \int \nabla U\left(\overline{\mathcal{E}}(\overline{X}_{(k)}) - \overline{\mathcal{E}}(x)\right) d\rho_{(k)}(x)\right). \tag{5.20}$$

Similarly, by inserting $f^N_{(k)} \approx f_{(k)}$ in the update rule (5.15) for general multi-objective problems we obtain

$$\frac{1}{N}\sum_{j=1}^N \left|\nabla U\left(\overline{\mathcal{E}}(X^i_{(k)}) - \overline{\mathcal{E}}(X^j_{(k)})\right)\right| \frac{W^i_{(k)} - W^j_{(k)}}{|W^i_{(k)} - W^j_{(k)}|} \approx \int \left|\nabla U\left(\overline{\mathcal{E}}(X^i_{(k)}) - \overline{\mathcal{E}}(x)\right)\right| \frac{W^i_{(k)} - w}{|W^i_{(k)} - w|} d f_{(k)}(x, w).$$

Under the above mean-field approximation, the adaptive dynamics (5.15) finally reads

$$\overline{W}_{(k+1)} = \Pi_{\Omega_m}\left(\overline{W}_{(k)} - \nu\tau \int \left|\nabla U\left(\overline{\mathcal{E}}(\overline{X}_{(k)}) - \overline{\mathcal{E}}(x)\right)\right| \frac{\overline{W}_{(k)} - w}{|\overline{W}_{(k)} - w|} d f_{(k)}(x, w)\right). \tag{5.21}$$

5.3.2 Convergence to the Pareto front

We consider now a mean-field approximation $(\overline{X}_{(k)}, \overline{W}_{(k)})$ of the M-CBO particles system, where the parameter $\overline{W}_{(k)}$ is iteratively defined according to either (5.19), (5.20) or (5.21). With $f_{(k)} \in \mathcal{P}(\mathbb{R}^d \times \Omega_m)$ we indicate the law of $(\overline{X}_{(k)}, \overline{W}_{(k)})$, while $\rho_{(k)} \in \mathcal{P}(\mathbb{R}^d)$, $\mu_{(k)} \in \mathcal{P}(\Omega_m)$ indicate the law of $\overline{X}_{(k)}$ and $\overline{W}_{(k)}$, respectively. In the following, we collect the necessary assumptions to prove convergence towards the Pareto front with respect to the GD metric (5.17).

The first assumption concerns the solvability of the scalarized sub-problems (5.3) via the CBO dynamics. From the analysis of single-objective CBO methods, we recall that well-conditioning around the minimizer is required to apply the quantitative Laplace principle (see Chapter 3). In the context of multi-objective optimization, we need to apply the quantitative Laplace principle with same parameter $\alpha > 0$ for all $w \in \Omega_m$. Therefore, we require a uniform behavior around the minimizers of the different sub-problems by assuming common lower and upper bounds.

Assumption 5.7 (Inverse continuity around minimizer)**.** *Let $p \in [1, \infty]$ be fixed. For all $w \in \Omega_m$, the scalarized function $\mathcal{E}_p(w; \cdot)$ is continuous and admits a unique minimizer $x^*(w)$.*

There exist constants $c_1, c_2, c_3, p_1, p_2 > 0$ and $R > 0$ independent on w, such that

$$c_1 \|x - x^*(w)\|_\infty^{p_1} \leq \mathcal{E}_p(w; x) - \mathcal{E}_p(w; x^*(w)) \leq c_2 \|x - x^*(w)\|_\infty^{p_2} \qquad \forall x,\ \|x - x^*(w)\|_\infty \leq R, \quad (5.22)$$

and such that

$$c_3 < \mathcal{E}_p(w; x) - \mathcal{E}_p(w; x^*(w)) \qquad \forall x,\ \|x - x^*(w)\|_\infty > R. \quad (5.23)$$

Another necessary assumption for the application of the quantitative Laplace principle is having a lower bound on the mass around the minimizer. Again, in the multi-objective settings, we require a lower bound for the different minimizers $x^*(w)$ which is uniform with respect to $w \in \Omega$. This will be achieved through the following assumption on the set of minimizers and on the initial data.

Assumption 5.8 (Boundedness and initial datum)**.** *The set of EP optimal points is bounded.*

Furthermore, let $H \subset \mathbb{R}^d$ be a bounded, open set containing them. The initial particle distribution $\rho_{(0)}$ is given by the uniform distribution over H, that is, $\rho_{(0)} = \mathrm{Unif}(H)$.

We note that, under Assumption 5.7 and for $\overline{\mathcal{E}}$ component-wise positive, the minimizers of the scalarized sub-problems $x^*(\Omega_m) = \{x^*(w) : w \in \Omega_m\}$ are also EP optimal points [62]. Therefore, by definition of the Pareto front F, we have that the pre-image of F is contained by the set of minimizers of the scalarized sub-problems:

$$x^*(\Omega_m) \subseteq \overline{\mathcal{E}}^{-1}(F).$$

The two set coincides only when Chebyshev scalarization is used (choice $p = \infty$), see Theorem 5.3 (c).

Finally, we assume Lipschitz regularity of the minimizers with respect to parameter $w \in \Omega_m$.

Assumption 5.9. *The map $x^* : \Omega_m \to \mathbb{R}^d$ is globally Lipschitz continuous.*

Before checking convergence in the objective space, we show that the mean-field particle system converges to EP optimal points in the search space $\mathbb{R}^d$. Similarly to the single-objective case, we consider to this end the mean-squared error defined now as $\mathrm{Err} : \mathcal{P}(\mathbb{R}^d \times \Omega_m) \to [0, \infty]$ with

$$\mathrm{Err}[f] := \int |x - x^*(w)|^2 \, df(x, w). \tag{5.24}$$

Theorem 5.10 (Convergence towards EP optimal points)**.** *Fix an accuracy $\varepsilon > 0$ and let Assumptions 5.7, 5.8, and 5.9 hold. Assume U to be differentiable with ∇U bounded. Consider parameters $\lambda, \sigma > 0$, $\nu \geq 0$ satisfying*

$$\delta := \lambda - \sigma^2 - 4C\frac{\nu}{\sqrt{\varepsilon}} - 4C^2\frac{\nu^2}{\varepsilon} > 0, \quad \text{where} \quad C = \mathrm{Lip}(x^*) \sup_{z \in \mathbb{R}^m} |\nabla U(z)|. \tag{5.25}$$

Define the time horizon $T^* > 0$

$$T^* = \frac{4}{\delta} \log\left(\frac{\mathrm{Err}[f_{(0)}]}{\varepsilon}\right).$$

Provided $\alpha > 0$ is sufficiently large and the step-size $\tau > 0$ is sufficiently small, it holds

$$\min_{k : k\tau \leq T^*} \mathrm{Err}[f_{(k)}] \leq \varepsilon. \tag{5.26}$$

Moreover, as long as $\mathrm{Err}[f_{(k-1)}] > \varepsilon$, it holds

$$\mathrm{Err}[f_{(k)}] \leq \exp\left(-k\tau\frac{\delta}{4}\right)\mathrm{Err}[f_{(0)}]. \tag{5.27}$$

Before providing a proof in the next section, we show how to apply this result to prove decay of the *GD* metric. Recall $\rho_{(k)} = \mathrm{Law}(\overline{X}_{(k)})$.

Corollary 5.11. *Under the settings of Theorem 3.6, further assume that the multi-objective function $\overline{\mathcal{E}}$ is globally Lipschitz continuous and component-wise positive.*

It holds

$$\min_{k : k\tau \leq T^*} GD[\rho_{(k)}] \leq \mathrm{Lip}(\overline{\mathcal{E}})\sqrt{\varepsilon}.$$

Moreover, as long as $\mathrm{Err}[f_{(k-1)}] > \varepsilon$, it holds

$$GD[\rho_{(k)}] \leq \exp\left(-k\tau\frac{\delta}{4}\right)\mathrm{Lip}(\overline{\mathcal{E}})\sqrt{\mathrm{Err}[f_{(0)}]}. \tag{5.28}$$

Remark 5.12. *In order to be satisfied, condition* (5.25) *requires a choice of the parameter* ν *such that* $\nu = \mathcal{O}(\sqrt{\varepsilon})$*. Therefore, convergence towards the Pareto front is guaranteed only if the sub-problems parameters are adapted slowly with respect to the CBO dynamics. This suggests that the two evolutionary dynamics should have a different time scale. Numerically, one can simulate such behavior by either choosing a relatively small parameter* ν*, or by adapting the sub-problems parameters* $W^i_{(k)}$ *less frequently with respect to the particles' positions* $X^i_{(k)}$*. We will experiment both approaches in the numerical tests performed in Section 5.4.*

5.3.3 Proof of Theorem 5.10

For the sake of notational simplicity, we introduce the following notation to indicate the particles after the update rule. For a given $x \in \mathbb{R}^d, w \in \Omega_m$, $f_{(k)} \in \mathcal{P}(\mathbb{R}^d \times \Omega_m)$ with first marginal $\rho_{(k)} \in \mathcal{P}(\mathbb{R}^d)$, and $\overline{\theta}_{(k)} \sim \mathcal{N}(0, I_d)$ we define x', w' as

$$x' = \mathcal{C}^{CBOw}(x, w, \rho_{(k)}, \overline{\theta}_{(k)}) := x + (\lambda\tau + \sigma\sqrt{\tau}\overline{\theta}_{(k)}) \odot (m^\alpha_w[\rho_{(k)}] - x) \tag{5.29}$$

$$w' = \mathcal{C}^A(w, x, f_{(k)}) := \begin{cases} w & \text{if (5.19) is used,} \\ \Pi_{\Omega_m}\left(w + \nu\tau \int \nabla U(g(x) - g(y)) d\rho_{(k)}(y)\right) & \text{if (5.20) is used,} \\ \Pi_{\Omega_m}\left(w - \nu\tau \int |\nabla U(g(x) - g(y))| \frac{w-v}{|w-v|} d f_{(k)}(y, v)\right) & \text{if (5.21) is used.} \end{cases}$$

With this notation, the law $f_{(k)}$ of $(\overline{X}_{(k)}, \overline{W}_{(k)})$ satisfies for any $\phi \in \mathcal{C}(\mathbb{R}^d \times \Omega_m)$

$$\int \phi(x, w) d f_{(k+1)}(x, w) = \mathbb{E} \int \phi(x', w') d f_{(k)}(x, w). \tag{5.30}$$

where the expectation is taken with respect to the sampling procedure $\overline{\theta}_{(k)} \sim \mathcal{N}(0, I_d)$.

To prove Theorem 5.10, we follow the proof strategy used for the single-objective CBO algorithm without memory effects in Chapter 3. Therefore, we start by studying the error evolution.

Proposition 5.13 (Error evolution)**.** *Under the settings of Theorem 5.10, let* λ, τ *be such that* $\lambda\tau \leq 1$*. Until the desired accuracy* ε *is reached, it holds*

$$\begin{aligned} \mathrm{Err}[f_{(k+1)}] \leq \left(1 - \tau\frac{\delta}{2}\right)\mathrm{Err}[f_{(k)}] + 2\tau C \frac{\nu}{\sqrt{\varepsilon}} \sqrt{\mathrm{Err}[f_{(k)}]} \sup_{v \in \Omega_m} |m^\alpha_v[\rho_{(k)}] - x^*(v)| \\ + \tau \frac{3\lambda + \sigma^2}{2} \sup_{v \in \Omega_m} |m^\alpha_v[\rho_{(k)}] - x^*(v)|^2. \end{aligned}$$

Proof. We apply (5.30) with $\phi(x,w) = |x - x^*(w)|^2$ to obtain

$$\begin{aligned}\int |x - x^*(w)|^2 f_{(k+1)}(x,w) &= \mathbb{E}\int |x' - x^*(w')|^2 d f_{(k)}(x,w)\\ &= \mathbb{E}\int |x' - x^*(w) + x^*(w) - x^*(w')|^2 d f_{(k)}(x,w)\\ &= \mathbb{E}\int |x' - x^*(w)|^2 d f_{(k)}(x,w)\\ &\quad - \mathbb{E}\int 2(x' - x^*(w))\cdot(x^*(w) - x^*(w'))d f_{(k)}(x,w)\\ &\quad + \mathbb{E}\int |x^*(w) - x^*(w')|^2 d f_{(k)}(x,w) =: I_1 + I_2 + I_3\,.\end{aligned}$$

We note that, in the computation of x', the consensus point corresponding to problem w is used. Therefore, in I_1, there is agreement between the consensus $m_w^\alpha[\rho_{(k)}]$ point and $x^*(w)$, as they both depend on the pre-update parameter w. This is the same situation we had in the single-objective case, and so the same computations done in the proof of Proposition 3.4 hold here too. This leads to the estimate

$$I_1 \leq \left(1 - \tau\frac{2\lambda - \sigma^2 - \lambda^2\tau}{2}\right)\mathrm{Err}[f_{(k)}] + \tau\frac{2\lambda + \sigma^2 + \lambda^2\tau}{2}\int |m_w^\alpha[\rho_{(k)}] - x^*(w)|^2 d\mu_{(k)}(w).$$

By using assumption $\lambda\tau < 1$, the above further simplifies to

$$I_1 \leq \left(1 - \tau\frac{\lambda - \sigma^2}{2}\right)\mathrm{Err}[f_{(k)}] + \tau\frac{3\lambda + \sigma^2}{2}\sup_{v\in\Omega_m} |m_v^\alpha[\rho_{(k)}] - x^*(v)|^2. \tag{5.31}$$

Next, let us bound I_3. Thanks to the Lipschitz continuity of x^* (Assumption 5.9), boundedness of ∇U, and the non-contractivity of the projection Π_{Ω_m}, it holds

$$|x^*(w) - x^*(w')| \leq \mathrm{Lip}(x^*)|w - w'| \leq \nu\tau\,\mathrm{Lip}(x^*)\sup_{z\in\mathbb{R}^m}|\nabla U(z)| = \nu\tau C \tag{5.32}$$

leading to

$$I_3 \leq \nu^2\tau^2 C^2.$$

By multiplying by $\mathrm{Err}[f_{(k)}]/\varepsilon \geq 1$, we further get an estimate in terms of the error

$$I_3 \leq \frac{\nu^2}{\varepsilon}\tau^2 C^2 \mathrm{Err}[f_{(k)}]. \tag{5.33}$$

To bound I_2 we use will use similar arguments. First of all, we note that the integrand in I_2 is linear in $\overline{\theta}_{(k)}$. By using $\mathbb{E}[\overline{\theta}_{(k)}] = 0$, we obtain

$$\mathbb{E}[x'] = x + \lambda\tau(m_w^\alpha[\rho_{(k)}] - x) = (1 - \lambda\tau)x + \lambda\tau m_w^\alpha[\rho_{(k)}],$$

which allows us to rewrite I_2 as

$$\begin{aligned} I_2 &= -\mathbb{E}\int 2(x' - x^*(w))\cdot(x^*(w) - x^*(w'))df_{(k)}(x,w) \\ &= -(1-\lambda\tau)\int 2(x - x^*(w))\cdot(x^*(w) - x^*(w'))df_{(k)}(x,w) \\ &\quad - \lambda\tau\int 2(m_w^\alpha[\rho_{(k)}] - x^*(w))\cdot(x^*(w) - x^*(w'))df_{(k)}(x,w). \end{aligned}$$

An application of Chauchy-Schwarz inequality and of (5.32) further leads to the estimate

$$\begin{aligned} I_2 &\le (1-\lambda\tau)2\sqrt{\mathrm{Err}[f_{(k)}]}\left(\int |x^*(w) - x^*(w')|^2 d\mu_{(k)}(w)\right)^{\frac{1}{2}} \\ &\quad + 2\lambda\tau\sup_{v\in\Omega_m}|m_v^\alpha[\rho_{(k)}] - x^*(v)|\int |x^*(w) - x^*(w')|d\mu_{(k)}(w) \\ &\le (1-\lambda\tau)2\nu\tau C\sqrt{\mathrm{Err}[f_{(k)}]} + 2\lambda\tau\nu\tau C\sup_{v\in\Omega_m}|m_v^\alpha[\rho_{(k)}] - x^*(v)|. \end{aligned}$$

As before, we derive an estimate in terms of $\mathrm{Err}[f_{(k)}]$ by multiplying by $\sqrt{\mathrm{Err}[f_{(k)}]}/\sqrt{\varepsilon} > 1$, and by using $\lambda\tau \in (0,1)$

$$I_2 \le 2\tau C\frac{\nu}{\sqrt{\varepsilon}}\mathrm{Err}[f_{(k)}] + 2\tau C\frac{\nu}{\sqrt{\varepsilon}}\sqrt{\mathrm{Err}[f_{(k)}]}\sup_{v\in\Omega_m}|m_v^\alpha[\rho_{(k)}] - x^*(v)|. \tag{5.34}$$

The desired upper bound for $\mathrm{Err}[f_{(k+1)}]$ follows after collecting estimates (5.31), (5.33), (5.34), and the definition of δ given in (5.25). □

Next, we aim to bound $|m_w^\alpha[\rho_{(k)}] - x^*(w)|$ uniformly in $w \in \Omega_m$ by means of the Laplace principle. Thanks to the assumptions on the sub-problems, we are able, in particular, to adapt Proposition 3.2 such that the estimate provided by the quantitative Laplace principle is independent on the sub-problem $w \in \Omega_m$ considered.

Proposition 5.14 (Quantitative Laplace principle for parameterized sub-problems)**.** *Let Assumptions 5.7, 5.8, 5.9 be satisfied and let $f \in \mathcal{P}(\mathbb{R}^{\mathrm{d}}\times\Omega)$ be such that its first marginal $\rho \in \mathcal{P}(\mathbb{R}^{\mathrm{d}})$ satisfies $H \subseteq \mathrm{supp}(\rho)$.*

For any $r \in [0,R]$, $q > 0$ such that $q + c_2 r^{p_2} \le c_3$ (constants introduced in Assumption 5.7) it holds

$$\sup_{v\in\Omega_m}|m_v^\alpha[\rho] - x^*(v)| \le c_1\sqrt{\mathrm{d}}(q + c_2 r^{p_2})^{1/p_1} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q)}{\inf_{v\in\Omega_m}\rho(B_r^\infty(x^*(v)))}\left(\sqrt{\mathrm{Err}[f]} + \mathrm{diam}(H)\right).$$

Proof. Without loss of generalization, we can assume $\inf_{x\in\mathbb{R}^{\mathrm{d}}} \mathcal{E}_p(w;x) = 0$ for all $w \in \Omega_m$. Since $x^*(w) \in H \subset \mathrm{supp}(\rho)$, we apply Proposition 3.2 for all $w \in \Omega_m$. As a consequence, for any $r \in (0, R]$ and $q > 0$ such that $q + \sup_{x\in B_r^\infty(x^*(w))} \mathcal{E}_p(w;x) < c_3$ the following upper bound holds

$$|m_w^\alpha[\rho] - x^*(w)| \leq c_1 \sqrt{\mathrm{d}}(q + \sup_{x\in B_r^\infty(x^*(w))} \mathcal{E}_p(w;x))^{1/p_1} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q)}{\rho(B_r^\infty(x^*(w)))} \int |x - x^*(w)| d\rho(x). \quad (5.35)$$

We aim now to make the above bound independent on $w \in \Omega_m$. Thanks to Assumption 5.7, we note that $\sup_{x\in B_r^\infty(x^*(w))} \mathcal{E}_p(w;x) \leq c_2 r^{p_2}$ for all $r \in (0, R]$ and therefore we can make a feasible choice of q, r which is common among all $w \in \Omega_m$. By boundedness of H (Assumption 5.8) and continuity of x^* (Assumption 5.9) it holds $\inf_{w\in\Omega_m} \rho(x^*(w)) > 0$ since $H \subseteq \mathrm{supp}(\rho)$. Finally, the last term we need to bound in (5.35) is $\int |x - x^*(w)| d\rho(x)$. By triangular and Jensen's inequalities we have

$$\int |x - x^*(w)| d\rho(x) \leq \int (|x - x^*(v)| + |x^*(v) - x^*(w)|)\, df(x,v) \leq \sqrt{\mathrm{Err}[f]} - \mathrm{diam}(H).$$

The estimates we collected lead to the desired upper bound for $|m_w^\alpha[\rho] - x^*(w)|$ which is independent on the choice of $w \in \Omega_m$. □

With the purpose of applying Proposition 5.14 with $f = f_{(k)}$, we require to prove that $H \subseteq \mathrm{supp}(\rho_{(k)})$ for $k \geq 0$. Moreover, it is desirable to have a lower estimate on $\inf_{v\in\Omega_m} \rho_{(k)}(B_r^\infty(x^*(v)))$ in terms of the initial data $\rho_{(0)}$. Analogously to the single-objective case, we study the mass around a EP optimal point $x^*(w)$ via the mollifier $\phi_{r,w} \in C_c^\infty(\mathbb{R}^{\mathrm{d}})$

$$\phi_{r,w}(x) := \begin{cases} \prod_{\ell=1}^{\mathrm{d}} \exp\left(1 - \frac{r^2}{r^2 - (x - x^*(w))_\ell^2}\right), & \text{if } \|x - x\|_\infty < r, \\ 0, & \text{else,} \end{cases} \quad (5.36)$$

for which holds

$$\rho_{(k)}(B_r^\infty(x^*(w))) \geq \int \phi_{r,w}(x) d\rho_{(k)}(x).$$

We do that through the following lemma, which is the multi-objective equivalent of Lemma 3.5.

Lemma 5.15 (Lower bound on mass around EP optimal points.)**.** *For a given $T > 0$ and $r > 0$, let*

$$\max_{k, k\tau\leq T} \sup_{w\in\Omega_m} |m_w^\alpha[\rho_{(k)}] - x^*(w)| \leq B$$

for some $B > 0$.

Provided $\tau > 0$ is sufficiently small, there exists a positive constant $a = a(r, \mathrm{d}, B, \lambda, \sigma, \mathrm{diam}(H))$ (independent on τ) such that

$$\inf_{w\in\Omega} \rho_{(k)}\left(B_r^\infty(x^*(w))\right) \geq m_{r,0} \exp(-ak\tau) \quad \forall k \,:\, k\tau \leq T\,,$$

with $m_{r,0} := (1/2)\inf_{w\in\Omega_m} \int \phi_{r,w}(x) d\rho_{(0)}(x)$.

Proof. For any arbitrary, but fixed $v \in \Omega_m$, the proof can be carried out exactly as the proof of Lemma 3.5, by considering the Taylor expansion of $\phi_{r,v} = \phi_{r,v}(x)$ around x at $x' = \mathcal{C}^{CBOw}(x, w, \rho_{(k)}, \theta)$ and the estimates provided for the time-continuous dynamics in [7, Lemma 4.2].

We also remark the importance, here, of Assumption 5.8. The boundedness of the set of EP optimal points, in particular, allows us to use the following triangular estimate

$$|m_w^\alpha[\rho_{(k)}] - x^*(v)| \leq |m_w^\alpha[\rho_{(k)}] - x^*(w)| + |x^*(w) - x^*(v)| \leq |m_w^\alpha[\rho_{(k)}] - x^*(w)| + \mathrm{diam}(H).$$

for any $w, v \in \Omega_m$, which is essential to bound the displacement $|x' - x^*(v)|$ for the particle (x, w). □

Finally, we provide a proof of the main theorem.

Proof of Theorem 5.10. We start by defining suitable $\tau_\varepsilon, \alpha_\varepsilon$ such that for all $\tau \leq \tau_\varepsilon$ and $\alpha \geq \alpha_\varepsilon$ we have convergence in the sense of (5.26). To this end, we introduce

$$C_\delta := \min\left\{\frac{1}{8}\frac{\delta}{2C(\nu/\sqrt{\varepsilon})}, \sqrt{\frac{1}{8}\frac{2\delta}{3\lambda + \sigma^2}}\right\}$$

and $q_\varepsilon, r_\varepsilon$ given by

$$q_\varepsilon := \frac{1}{2}\min\left\{\frac{1}{c_1\sqrt{\mathrm{d}}}\left(\sqrt{\varepsilon}C_\delta\right)^{p_1}, c_3\right\}, \quad r_\varepsilon := \frac{1}{c_2} q_\varepsilon^{1/p_2}.$$

Then, we pick a step size $\tau_\varepsilon > 0$ such that $\lambda\tau_\varepsilon < 1$ and such that Lemma 5.15 holds with $r = r_\varepsilon$, $B = C_\delta\sqrt{\mathrm{Err}[f_{(0)}]}$ and $T = T^*$. We remark that the definition of τ_ε, ultimately, only depends on $\varepsilon, f_{(0)}, \lambda, \sigma, \nu$ and it does not depend on any choice of α.

Next, we consider any α_ε satisfying

$$\frac{\sqrt{\mathrm{d}}\exp(-\alpha_\varepsilon q_\varepsilon + aT^*)}{m_{r_\varepsilon,0}}\left(\sqrt{\mathrm{Err}[f_{(0)}]} + \mathrm{diam}(H)\right) < C_\delta\sqrt{\varepsilon}\,, \tag{5.37}$$

with $a > 0$ being the exponent determined by Lemma 5.15.

This particular choices of $\tau_\varepsilon, \alpha_\varepsilon$ will allow us to apply Lemma 5.15 and Proposition 5.14 as long as $\sup_{w\in\Omega_m} |m_w^\alpha[\rho_{(k)}] - x^*(w)| \leq C_\delta\sqrt{\mathrm{Err}[f_{(0)}]}$.

Now, for any fixed $\tau \le \tau_\varepsilon$ and $\alpha \ge \alpha_\varepsilon$, we consider the iterative step K

$$K := \sup\left\{ k \,:\, \mathrm{Err}[f_{(h)}] > \varepsilon \text{ and } \sup_{w\in\Omega_m} |m_w^\alpha[\rho_{(h)}] - x^*(w)| < C_\delta \sqrt{\mathrm{Err}[f_{(h)}]} \;\; \forall h \le k \right\}. \tag{5.38}$$

By applying Proposition 5.13 and since $\lambda\tau \le \lambda\tau_\varepsilon < 1$, we obtain that for all $k \le K$

$$\begin{aligned}
\mathrm{Err}[f_{(k+1)}] &\le \left(1 - \tau\frac{\delta}{2}\right)\mathrm{Err}[f_{(k)}] + 2\tau C\frac{\nu}{\sqrt{\varepsilon}}\sqrt{\mathrm{Err}[f_{(k)}]}\sup_{v\in\Omega_m}|m_v^\alpha[\rho_{(k)}] - x^*(v)| \\
&\quad + \tau\frac{3\lambda+\sigma^2}{2}\sup_{v\in\Omega_m}|m_v^\alpha[\rho_{(k)}] - x^*(v)|^2 \\
&\le \left(1 - \tau\frac{\delta}{2}\right)\mathrm{Err}[f_{(k)}] + 2\tau C\frac{\nu}{\sqrt{\varepsilon}}\sqrt{\mathrm{Err}[f_{(k)}]}C_\delta\sqrt{\mathrm{Err}[f_{(h)}]} + \tau\frac{3\lambda+\sigma^2}{2}C_\delta^2\mathrm{Err}[f_{(h)}] \\
&\le \left(1 - \tau\frac{\delta}{2}\right)\mathrm{Err}[f_{(k)}] + \tau\frac{\delta}{8}\mathrm{Err}[f_{(h)}] + \tau\frac{\delta}{8}\mathrm{Err}[f_{(h)}] \le \left(1 - \tau\frac{\delta}{4}\right)\mathrm{Err}[f_{(k)}]
\end{aligned}$$

where we also applied the definition of K and C_δ. By iterating the above argument, we obtain for all $k \le K$

$$\mathrm{Err}[f_{(k)}] \le \left(1 - \tau\frac{\delta}{4}\right)^k \mathrm{Err}[f_{(0)}] \le \exp\left(-k\tau\frac{\delta}{4}\right)\mathrm{Err}[f_{(0)}].$$

We note that the right-hand-side above is decreasing due to the assumption $\delta > 0$. Therefore, $\mathrm{Err}[f_{(k)}]$ and $\sup_{w\in\Omega_m}|m_w^\alpha[\rho_{(k)}] - x^*(w)|$ can be bounded as following:

$$\begin{aligned}
\max_{0\le k\le K}\mathrm{Err}[f_{(k)}] &\le \mathrm{Err}[f_{(0)}] \\
\max_{0\le k\le K}\sup_{w\in\Omega_m}|m_w^\alpha[\rho_{(k)}] - x^*(w)| &\le C_\delta \max_{0\le k\le \bar{k}}\sqrt{\mathrm{Err}[f_{(k)}]} \le C_\delta\sqrt{\mathrm{Err}[f_{(0)}]} = B.
\end{aligned}$$

where, again, we used that $\sup_{w\in\Omega_m}|m_w^\alpha[\rho_{(k)}] - x^*(w)| < C_\delta\mathrm{Err}[f_{(k)}]$ as long as $k \le K$.

To show $\mathrm{Err}[f_{(k)}] \le \varepsilon$ for some k, we check three different cases.

Case $K \ge T^*/\tau$. Thanks to the definition of T^* and the error exponential decay up to iteration K, we have

$$\mathrm{Err}[f_{(K)}] \le \exp\left(-K\tau\frac{\delta}{4}\right)\mathrm{Err}[f_{(0)}] \le \exp\left(-T^*\frac{\delta}{4}\right)\mathrm{Err}[f_{(0)}] \le \varepsilon.$$

Case $K < T^*/\tau$ and $\mathrm{Err}[f_{(K)}] \le \varepsilon$. Nothing to prove here.

Case $K < T^*/\tau$, $\mathrm{Err}[f_{(K)}] > \varepsilon$ and $\sup_{w\in\Omega_m}|m_w^\alpha[f_{(K)}] - x^*(w)| \ge C_\delta\sqrt{\mathrm{Err}[f_{(K)}]}$.

We prove that due to our choice of α, τ it holds $\sup_{w\in\Omega_m}|m_w^\alpha[f_{(K)}] - x^*(w)| < C_\delta\sqrt{\mathrm{Err}[f_{(K)}]}$ and therefore we are led to a contradiction.

By applying Proposition 5.14 with $q = q_\varepsilon$, $r = r_\varepsilon$, we obtain thanks to our previous choice of $q_\varepsilon, r_\varepsilon$

$$\sup_{v\in\Omega_m} |m_v^\alpha[\rho_{(K)}] - x^*(v)| \le c_1\sqrt{\mathrm{d}}(q + c_2 r^{p_2})^{1/p_1} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q)}{\inf_{v\in\Omega_m}\rho_{(K)}(B_r^\infty(x^*(v)))}\left(\sqrt{\mathrm{Err}[f_{(K)}]} + \mathrm{diam}(H)\right)$$
$$< \frac{\sqrt{\varepsilon}C_\delta}{2} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon)}{\inf_{v\in\Omega_m}\rho_{(K)}(B_r^\infty(x^*(v)))}\left(\sqrt{\mathrm{Err}[f_{(K)}]} + \mathrm{diam}(H)\right).$$

To bound the second term, we use that $\mathrm{Err}[f_{(K)}] \le \mathrm{Err}[f_{(0)}]$. Thanks to the choice $\tau \le \tau_\varepsilon$, we can also apply Lemma 5.15 with $r = r_\varepsilon$ and obtain a lower bound on the mass around the minimizers with the exponent $a > 0$

$$\inf_{v\in\Omega_m} \rho_{(K)}(B_r^\infty(x^*(v))) \ge m_{r,0}\exp(-aK\tau).$$

This leads to

$$\sup_{v\in\Omega_m} |m_v^\alpha[\rho_{(K)}] - x^*(v)| \le \frac{\sqrt{\varepsilon}C_\delta}{2} + \frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon + aT^*)}{m_{r_\varepsilon,0}}\left(\sqrt{\mathrm{Err}[f_{(0)}]} + \mathrm{diam}(H)\right).$$

Now, we note that α_ε was chosen exactly to bound the second term on the right-hand side above. In particular, due to the choice (5.37) and $\alpha \ge \alpha_\varepsilon$, we have

$$\frac{\sqrt{\mathrm{d}}\exp(-\alpha q_\varepsilon + aT^*)}{m_{r_\varepsilon,0}}\left(\sqrt{\mathrm{Err}[f_{(0)}]} + \mathrm{diam}(H)\right) \le \frac{\sqrt{\varepsilon}C_\delta}{2}.$$

Altogether, since the considered case assumes $\varepsilon \le \mathrm{Err}[f_{(K)}]$, it holds

$$|m^\alpha[f_{(k)}] - x^*| < C_\delta\sqrt{\varepsilon} \le C_\delta\sqrt{\mathrm{Err}[f_{(K)}]},$$

which is the desired contradiction. □

5.4 Numerical experiments

In this section, we test the proposed methods against bi-objective and tri-objective optimization problems using Chebyshev scalarization ($p = \infty$ in (5.3)). We are particularly interested in verifying if the additional adaptive strategies allow the algorithm to reach a low-energy description of the Pareto front and in validating the theoretical analysis performed in Section 5.3.

To quantitatively assess the algorithms' performance, we will use different metrics. We recall from Section 5.1 that short-range repulsive potentials can be used to measure the

diversity of a computed solution. We consider in the following three different potentials: the Riesz $(m-1)$-energy, the Morse potential and the Newtonian one, given respectively by

$$U_R(z) = \frac{1}{|z|^{m-1}}, \qquad U_M(z) = e^{-C|z|}, \qquad U_N(z) = \begin{cases} -\log|z| & \text{if } m = 2 \\ 1/|z|^{m-2} & \text{if } m > 2 \end{cases}$$

for some $C > 0$. To measure how close the computed solution is to the Pareto front, we use the Generational Distance (GD) which was already introduced and studied in Section 5.2. The metric GD is numerically computed via a reference approximation F^M the Pareto front made of a collection of $M = 100$ points $F^M = \{z_j\}_{j=1}^M$. We use the same notation for simplicity:

$$GD[\rho_{(k)}^N] := \left(\frac{1}{N} \sum_{i=1}^N \operatorname{dist}(\overline{\mathcal{E}}(X_{(k)}^i), F^M)^2 \right)^{\frac{1}{2}}. \tag{5.39}$$

We further introduce two common performance metrics: the Inverted Generational Distance (IGD) and the hyper-volume contribution ($\mathcal{S}$) [22, 84]. The metric IGD is defined similarly to GD, but the roles of the computed and the reference solutions are inverted:

$$IGD[\rho_{(k)}^N] := \left(\frac{1}{M} \sum_{j=1}^M \operatorname{dist}(y^j, E_{(k)})^2 \right)^{\frac{1}{2}} \quad \text{with} \quad E_{(k)} := \{\overline{\mathcal{E}}(X_{(k)}^i)\}_{i=1}^N. \tag{5.40}$$

While GD only measures how close the computed solution is to the Pareto front, IGD also takes into account whether the output is well-spread over the front (assuming the reference solution is itself well-spread). For this reason, IGD is considered to be a good metric to quantitative assess the overall performance of a multi-objective algorithm [22]. The hyper-volume contribution, instead, measures the amount of points that are dominated (a dominates b, written $a \prec b$, if $a < b$ component-wise) by the computed solution. Let $\overline{\mathcal{E}}^*$ be an upper bound for the Pareto front and $\mathcal{L}^m$ denote the Lebesgue measure over $\mathbb{R}^m$. The hyper-volume contribution is defined by

$$\mathcal{S}[\rho_{(k)}^N] := \mathcal{L}^m \left(\bigcup_{i=1}^N \left\{ y \in \mathbb{R}^m \mid \overline{\mathcal{E}}(X_{(k)}^i) \prec y \prec \overline{\mathcal{E}}^* \right\} \right). \tag{5.41}$$

We note that, differently from GD and IGD, the computation of $\mathcal{S}$ does not require to have a reference solution, but only an upper bound of the front.

In the following, we qualitatively and quantitatively test the proposed adaptive CBO dynamics against several bi- and tri-objective benchmark problems. The problems' definitions are provided in Section 5.4.3.

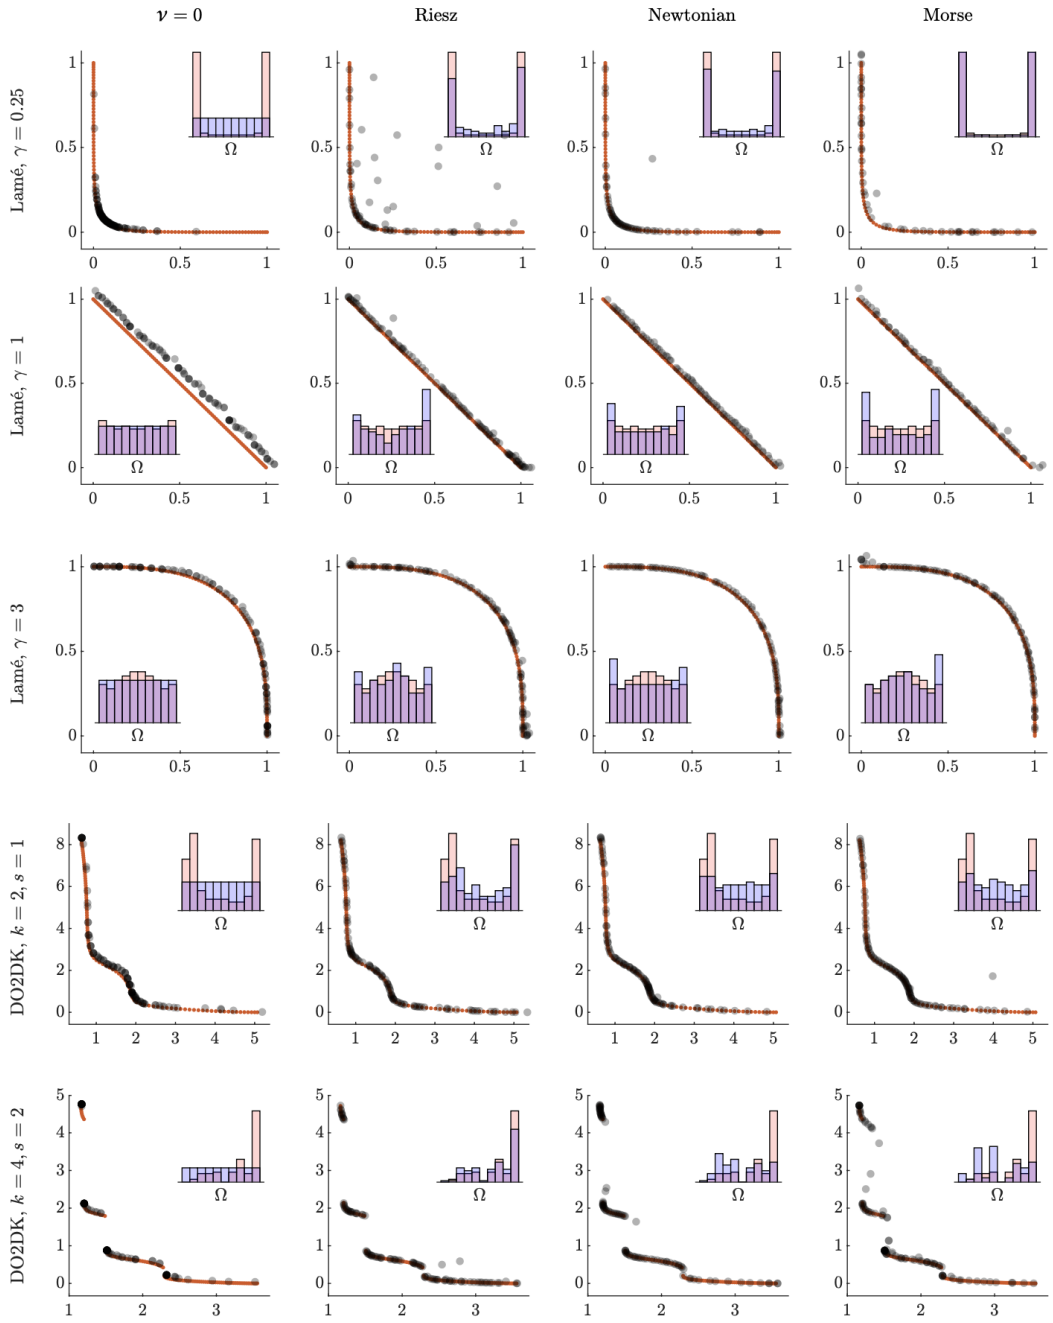

Figure 5.4: Shown in gray are the positions of the particles in the image space after a single run. The reference solution is shown in red. Four different adaptive strategies are used: no interaction, interaction with Riesz, Newtonian, and Morse potential. Histograms show the final distribution over the simplex Ω_2 (blue) and the optimal one (red).

5.4.1 Bi-objective problems

In this section, we focus on the adaptive strategy (5.12) inspired by the gradient-flow dynamics. We start by qualitatively and quantitative investigating the algorithm performance for different shapes of the Pareto front. Then, we check how sensitive the suggested strategy is with respect to the parameters' choice and to the search space dimension d.

Test problems. We test the algorithms in 4 different settings:

1. no parameters' adaptation ($\nu = 0$);
2. Adaptation through Riesz potential with $\nu = 10^{-5}$;
3. Adaptation through Newtonian potential with $\nu = 10^{-3}$;
4. Adaptation through Morse potential with $\nu = 10^{-1}, C = 20$.

The choice of ν has been optimized for each settings. We note how, the stronger the singularity at the origin is, the slower the adaptive process should be. The remaining parameters are chosen as following: $\lambda = 1, \sigma = 4, \alpha = 10^6, \tau = 10^{-2}$. We use $N = 100$ particles. Their initial positions $X^i_{(0)}$ are uniformly sampled from $[0,1]^{\mathrm{d}}$, $X^i_{(0)} \sim \mathrm{Unif}([0,1]^{\mathrm{d}})$, while the initial sub-problems parameters $W^i_{(0)}$ are uniformly and deterministically taken over Ω_2. The algorithm runs for a total of $k_{\max} = 5000$ iterations. The search-space dimension is set to $\mathrm{d} = 10$.

Figure 5.4 illustrates the particle configuration over the objective space for the different settings considered and for 5 benchmark problems. The plots also include a histogram illustrating the final distribution of the sub-problems' parameters $\{W^i_{(k)}\}_{i=1}^N$ over the simplex Ω_2.

When tested against problem Lamé $\gamma = 0.25$, the CBO dynamics alone (stategy $\nu = 0$) is not able to recover a good approximation of the front. This was expected as a uniform distribution over Ω_2 is not suitable for multi-objective problems with Pareto fronts much different from the simplex Ω_2 itself. For this problem, we note that the best front approximation is reached by the algorithm run using Morse potential. The Riesz potential seems to be too strong here, as many particles did not converge towards the front. For the problem Lamé $\gamma = 1$, the different adaptive strategies return similar results. Even though the choice of uniform distribution is optimal in this case, since $F = \Omega_2$, it is interesting to note that strategy $\nu = 0$ returns a worse approximation of the front. We conjecture that the additional adaptive dynamics improves the CBO algorithm (even when it is not needed) by further adding stochasticity to the evolution. The adaptive mechanism seems to improve the algorithm output also when the Pareto front is discontinuous, see problem DO2DK $k = 4, s = 2$ in Figure 5.4. In this case, though, we note that a certain number of particles keep moving from one disconnected set to the other during the

computation. While this behavior cannot be avoided due to the repulsive dynamics, such sub-optimal points can be easily discarded by adding a post-processing step to the algorithm.

The algorithm quantitative performance in the different settings is provided in Table 5.1. Results obtained via the well-known NSGA-II [26] algorithm are added for comparison. Algorithm NSGA-II is tested with the implementation provided in [69] with same number of iterations $k_{\max} = 5000$, same number of particles $N = 100$ and default choice of parameters (mutation strength $m_s = 0.05$, mutation probability $p_m = 0.5$, crossover probability $p_c = 0.9$).

As expected from the theoretical analysis, in particular from Theorem 5.10 and Corollary 5.11, the lowest value of GD are usually attained when no adaptive dynamics is used ($\nu = 0$). On the other hand, we obtain lower energy configurations when there is short-range interaction between the particles. As already noticed from Figure 5.4, Morse potential is oftentimes the best choice for the benchmark problems considered. The adaptive mechanisms improve the CBO method also with respect to the IGD metric.

Figure 5.5 shows how the different performance metrics evolve during the computation for two of the benchmark problems considered. We can recognize two different stages in the evolution. During the first stage, GD decreases quickly, indicating that the particles converge towards EP optimal points, and so towards the Pareto front in the image space. As a consequence, the energy of the particles' configuration increases too. The second stage is characterized, instead, by a slow decrease of the potential energies. This is particularly evident for the Morse energy. The presence of two-different time scales was already noticed during the theoretical analysis (see Remark 5.12) and it is further confirmed here by the numerical experiments. The decrease of the potential energies was also expected due to the formal derivation of the adaptive strategy (5.12) from a gradient flow dynamics. It is interesting to note that the multi-swarm CBO approach proposed in [66] leads to the opposite behavior, where the sub-problems' parameters adapt first and the particles reach a consensus afterwards.

Effect of the parameter ν and scalability. We test the algorithm for different values of ν, keeping the other parameters fixed, in order to experimentally investigate the importance of ν. The remaining algorithm parameters are taken as in the previous tests.

Figures 5.6a and 5.6b show the final values of GD and IGD when different binary potentials are used in the computation. As can be expected, relatively large values of ν lead to a strong interaction in the parameter space, which affects the CBO mechanism. Therefore, the GD metric increases for large values of ν. Interestingly, the lowest GD values are not always achieved for the smallest values of ν. This is an indication that the additional dynamics of the weight vectors may help the CBO mechanism in the optimization of the sub-problems.

The optimal value of ν is different for each test case, as shown by the IGD metrics in Figure 5.6b. In particular, a strong interaction in the parameter space benefits the DO2DK

Problem	Interaction	GD	$\mathcal{U}_R$	$\mathcal{U}_N$	$\mathcal{U}_M$	$\mathcal{S}$	IGD
Lamé	$\nu = 0$	2.33e-02	1.00e+10	2.41e+00	4.86e-01	9.69e-01	1.31e-01
$\gamma = 0.25$	Riesz	8.74e+00	5.65e+00	-1.94e-01	9.62e-02	7.77e-01	4.06e-02
	Newtonian	1.11e+01	8.23e+00	-3.53e-01	1.14e-01	8.38e-01	4.25e-02
	Morse	1.49e+01	1.81e+04	-1.36e+00	3.40e-02	7.45e-01	2.64e-02
	NSGA-II	6.99e-03	7.98+00	9.00e-01	7.45e-01	9.84e-01	4.19e-03
Lamé	$\nu = 0$	9.88e-02	9.60e+09	9.96e-01	1.26e-01	3.74e-01	8.28e-02
$\gamma = 1$	Riesz	1.63e-01	6.54e+00	8.57e-01	1.23e-01	4.59e-01	1.56e-02
	Newtonian	9.81e-01	8.39e+00	5.77e-01	9.47e-02	4.62e-01	1.91e-02
	Morse	6.83e-01	8.97e+05	4.41e-01	7.95e-02	4.48e-01	1.78e-02
	NSGA-II	2.48e-02	5.86e+00	1.03e+00	1.06e-01	4.71e-01	4.33e-03
Lamé	$\nu = 0$	1.93e-02	8.40e+09	9.56e-01	1.30e-01	8.45e-02	2.18e-02
$\gamma = 3$	Riesz	5.64e-02	7.06e+00	7.68e-01	1.14e-01	1.01e-01	1.32e-02
	Newtonian	2.34e-01	6.33e+00	5.74e-01	9.10e-02	1.03e-01	1.11e-02
	Morse	3.02e-01	9.57e+06	5.04e-01	7.84e-02	1.02e-01	1.29e-02
	NSGA-II	3.64e-02	4.65e+00	9.16e-01	8.10e-02	9.68e-02	5.28e-03
DO2DK	$\nu = 0$	1.80e-01	1.00e+10	-3.30e-01	6.94e-02	8.84e+01	2.82e-01
k=2,s=1	Riesz	5.03e-02	1.58e+00	-6.04e-01	4.18e-02	8.94e+01	1.18e-01
	Newtonian	6.48e-02	1.77e+01	-5.98e-01	3.85e-02	8.94e+01	1.07e-01
	Morse	9.59e-02	7.67e+08	-5.64e-01	3.75e-02	8.94e+01	9.33e-02
	NSGA-II	9.97e-02	9.44e-01	-9.18e-01	1.85e-02	8.91e+01	3.42e-02
DO2DK	$\nu = 0$	6.60e-02	1.00e+10	2.69e+00	2.64e-01	8.66e+01	1.36e-01
k=4, s=2	Riesz	8.95e-01	3.34e+00	-1.27e-01	7.08e-02	8.40e+01	2.61e-02
	Newtonian	1.50e+00	2.18e+01	-2.26e-01	7.52e-02	8.44e+01	3.61e-02
	Morse	9.85e+00	1.94e+09	-1.56e-01	9.09e-02	7.63e+01	3.45e-02
	NSGA-II	7.80e-02	2.36e+00	9.85e-02	4.19e-02	8.66e+01	1.11e-01

Table 5.1: Performance of the algorithm for different settings and problems. The performance of the NSGA II algorithm is also shown as a reference for comparison with well-established heuristics. Green boxes show the proposed algorithm's best results. The results are averaged over a set of 25 runs.

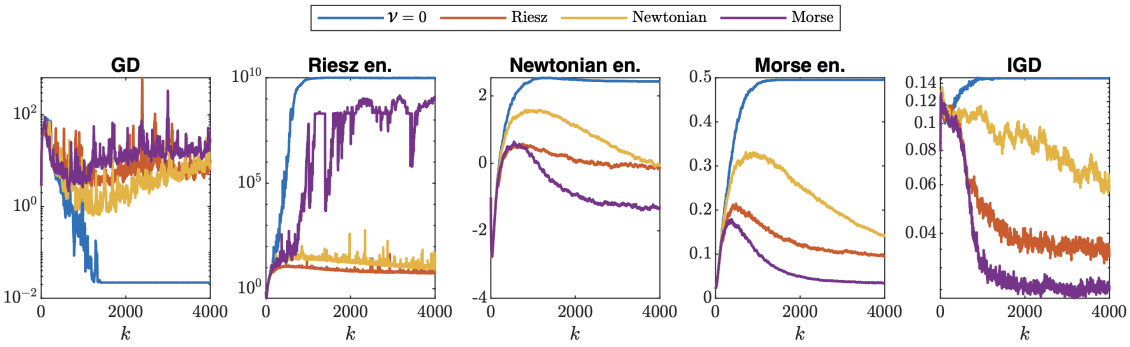

(a) Problem Lamé $\gamma = 0.25$

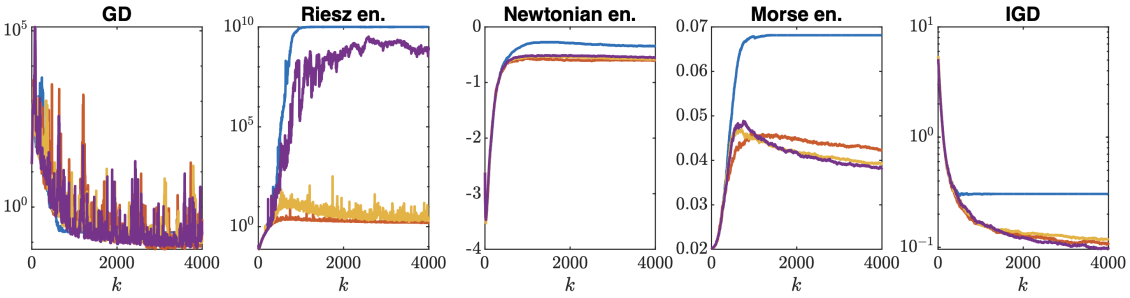

(b) Problem DO2DK $k = 2, s = 1$

Figure 5.5: Performance metrics evolution. The results are averaged over a set of 25 runs.

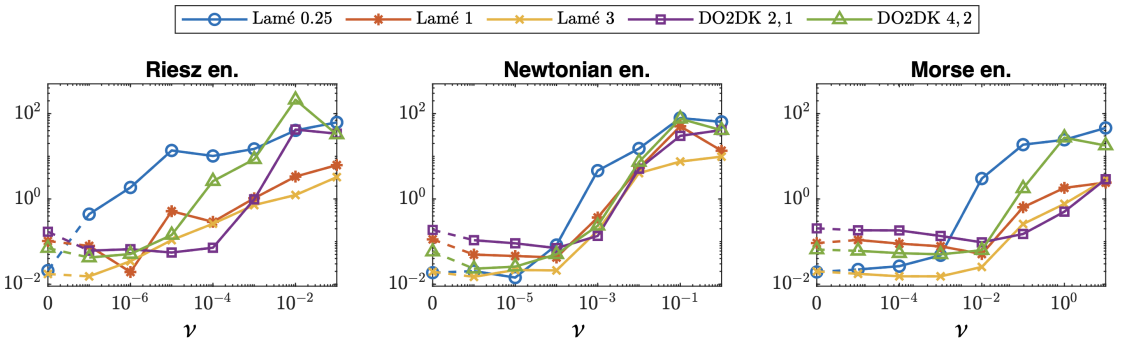

(a) Generational Distance (GD) (5.39)

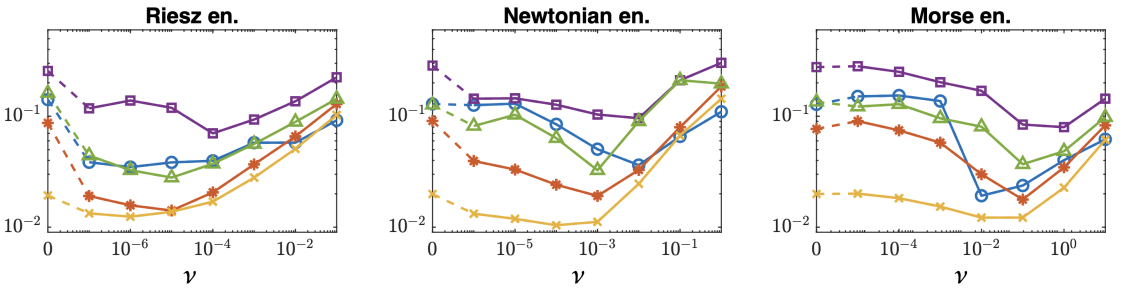

(b) Inverted Generational Distance (IGD) (5.40)

Figure 5.6: Performance metrics as a function of ν for all problems under consideration. The results are averaged over a set of 10 runs.

problems. This can be explained by the geometry of the front (Figure 5.4): the length of the front is longer and, as a consequence, the particles tend to be further apart in the image space, which makes the interaction of the binary potential weaker. Larger values of ν mitigate this effect and lead to a better performance of the algorithm. This issue could be addressed by estimating the length of the front and choosing the parameter ν accordingly, if the extrema of the Pareto front are known in advance. We also note that the proposed algorithm seems to perform better when the Morse potential is used to compute the adaptive mechanism.

As mentioned above, the adaptive strategy in Ω_m adds stochasticity to the particle evolution in $\mathbb{R}^d$. Therefore, the stochastic component in the position dynamics (5.8) may not be necessary. However, assuming $\sigma = 0$ leads to bad approximations of the Pareto front, see Figure 5.7b. This is an indication that the diffusion term is still of paramount importance for the exploratory behavior of the particles and their statistical independence. From Figure 5.7a it is clear that the optimal diffusion parameter σ is larger the smaller ν is. Specifically, for $\nu = 0$ the particles do not deviate from the EP optimal points until $\sigma > 10$. Also, for some problems, when σ is too small, larger values for ν improve convergence to

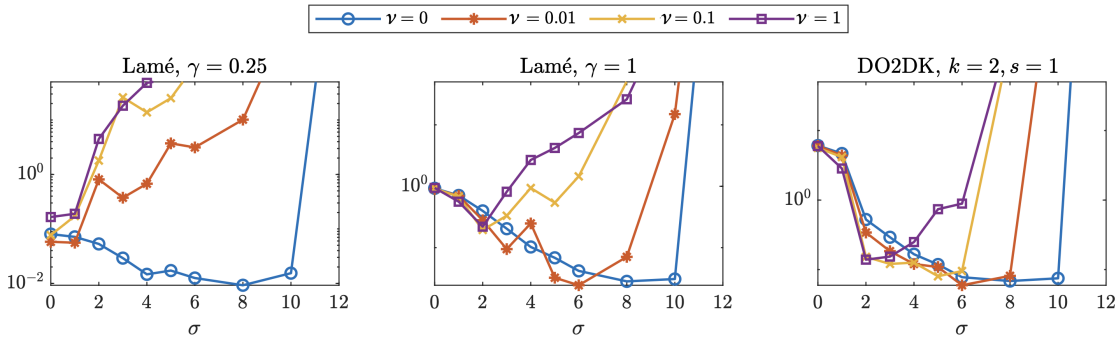

(a) *GD*

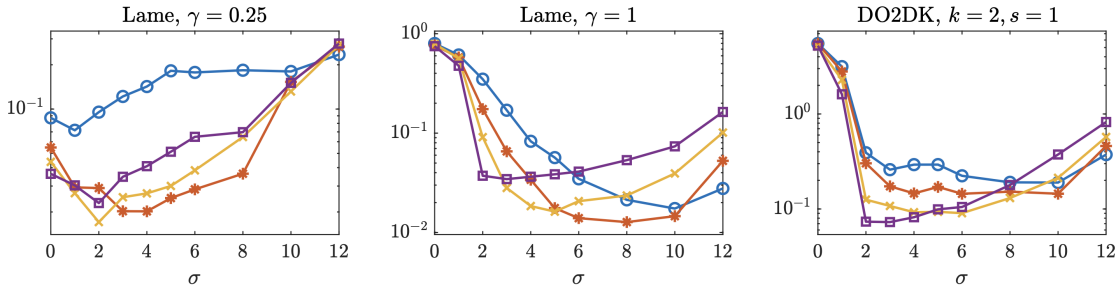

(b) *IGD*

Figure 5.7: Metrics as functions of σ, for different values of ν. Results are averaged over a set of 5 runs using Morse interaction.

optimal points.

Finally, with the same choice of parameters, we test the performance of the algorithm for different dimensions d of the search space. If the same number $N = 100$ of particles is used, the *IGD* of the computed solutions increases with the increase of the dimension d of the space, see Figure 5.8 (left). This effect can be mitigated by simply increasing the number of particles linearly with the space dimension, see Figure 5.8 (right).

5.4.2 Tri-objective problems

For tri-objective problems, $m = 3$, we have suggested the adaptive dynamics (5.15), with, eventually, the addition of noise as in (5.16). Being tri-objective objective optimization problems more challenging than bi-objective ones, we experiment in this section a different implementation of the particle systems (5.8) and (5.15) to improve performance.

First of all, the two-scale behavior of the system is obtained by updating the sub-problem parameters $W^i_{(k)}$ less frequently than the particles' positions $X^i_{(k)}$. Specifically, we update the parameters every 50 updates of the positions. Secondly, we assign to

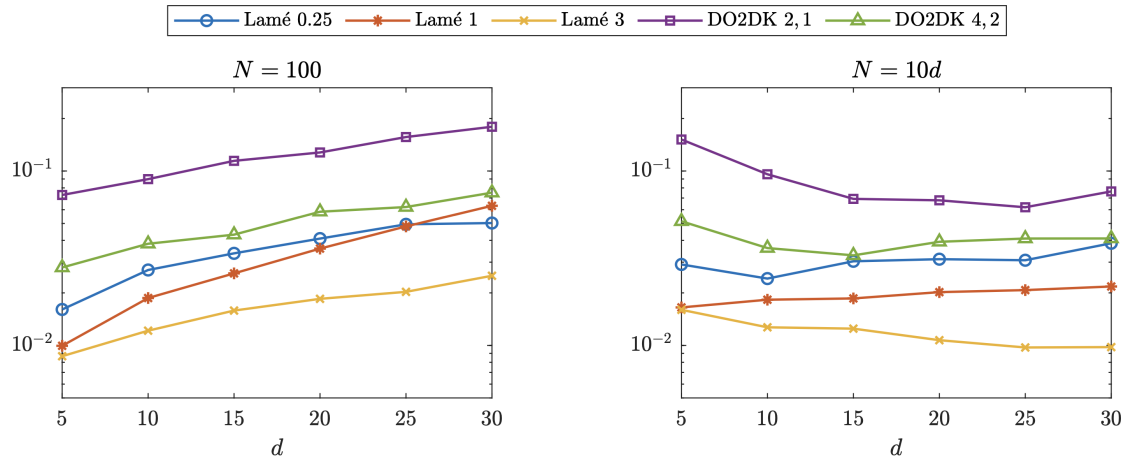

Figure 5.8: Metric IGD as a function of the search space dimension d. Morse interaction is used and results are averaged over a set of 20 runs.

each sub-problem $n \in \mathbb{N}$ particles' positions. Therefore, we do not have a one-to-one correspondence between sub-problems and particles. With these changes, we are allocating more computational resources to each scalarized sub-problem and, so, we expect the CBO dynamics to obtain more accurate solutions. The detailed algorithmic strategy is outlined in Algorithm 3.

We test Algorithm 3 against the tri-objective problems: Lamé with $\gamma = 0.5, 2$, and the inverted DTLZ1 problem. We use $N_w = 66$ sub-problems and $n = 10$ particles per sub-problem. The adaptive dynamics is computed via Morse binary potential with $C = 30$ and parameter $\nu = 10^{-2}$. The remaining algorithm parameters are set to $\lambda = 1, \sigma = 1, \tau = 10^{-2}, \alpha = 10^5, k_{\max} = 1000$. Search space dimension is set to d = 3. When used, the stochastic component in the adaptive dynamics has strength $\hat{\sigma} = 10^{-6}$.

Figures 5.9, 5.10, and 5.11 show the image of the computed EP optimal points and the final configuration over the simplex Ω_3. As for the bi-objective case, we note that without adaptive dynamics the algorithm is able to compute EP optimal points, but that their image is not well-distributed among the Pareto front. Remarkably, the adaptive dynamics allow to reach, in the case of the inverted DTLZ1 problem, an almost uniform distribution over the front, even though this corresponds to a non-trivial distribution over the simplex, see Figures 5.9b and 5.9c. The addition of stochasticity typically reduces the total energy of the system in the image space, as expected.

The evolution of the potential energies and of the IGD metric are illustrated in Figure 5.12. For the considered Lamé problems, the addition of noise significantly improves the quality of the solution with respect to both measures, while for the inverted DTLZ1 problem, the noise does not play a significant role. The adaptive strategies allow better

Set $\lambda, \sigma, \nu, \tau > 0, k_{\max} \in \mathbb{N}$;
Set number of sub-problems $N_w \in \mathbb{N}$ and number of particles per sub-problem $n \in \mathbb{N}$;
Initialize $X^{i,j}_{(0)}$ for all $i = 1, \dots, N_w$, $j = 1, \dots, n$ uniformly over $[0,1]^{\mathrm{d}}$;
Initialize $W^i_{(0)}$ for all $i = 1, \dots, N_w$, equispaced over Ω_m;
for $k = 0, 1, \dots, k_{\max}$ **do**
  $\rho^N_{(k)} = 1/(nN_w) \sum_{i,j} \delta_{X^{i,j}_{(k)}}$;
  $\theta^{i,j}_{(k)} \sim \mathcal{N}(0, \mathbf{I}_{\mathrm{d}})$ for all i, j ;
  $X^{i,j}_{(k+1)} = \mathcal{C}^{CBOw}(X^{i,j}_{(k)}, W^i_{(k)}, \rho^N_{(k)}, \theta^{i,j}_{(k)})$ for all i, j (see (5.29))
  if mod $(k, 50) = 0$ **then**
    $Z^i_{(k)} = \mathcal{E}(m^{\alpha}_{W^i_{(k)}}[\rho^N_{(k)}])$ for all $i = 1, \dots, N_w$;
    $W^i_{(k+1/2)} = W^i_{(k)} - \nu(\tau/N_w) \sum_{j=1}^{N_w} \left| \nabla U\left(Z^i_{(k)} - Z^j_{(k)}\right) \right| \frac{W^i_{(k)} - W^j_{(k)}}{|W^i_{(k)} - W^j_{(k)}|}$ for all $i = 1, \dots, N_w$;
    $W^i_{(k+1)} = \Pi_{\Omega_m}(W^i_{(k+1/2)})$ for all $i = 1, \dots, N_w$
  else
    $W^i_{(k+1)} = W^i_{(k)}$ for all $i = 1, \dots, N_w$;
  end
end

Algorithm 3: Adaptive M-CBO, modified. The two-scale behavior of the method is reached by updating the sub-problems parameters only every 50 CBO iterations. For every sub-problem, n different particles are used. Since there is not one-to-to correspondence between particles' positions and sub-problems, we compute the potential vector field in the objective space via the consensus points $m^{\alpha}_{W^i_{(k)}}[\rho^N_{(k)}]$, $i = 1, \dots, N_w$ (cf. (5.15)). The total number of particles is given by $N = n \cdot N_w$. Noise may be also added in the update of $\{W^i_{(k)}\}_{i=1}^{N_w}$, as in (5.16). We note that the use of many particles per sub-problem makes the algorithm similar to the multi-swarm CBO approach suggested in [66].

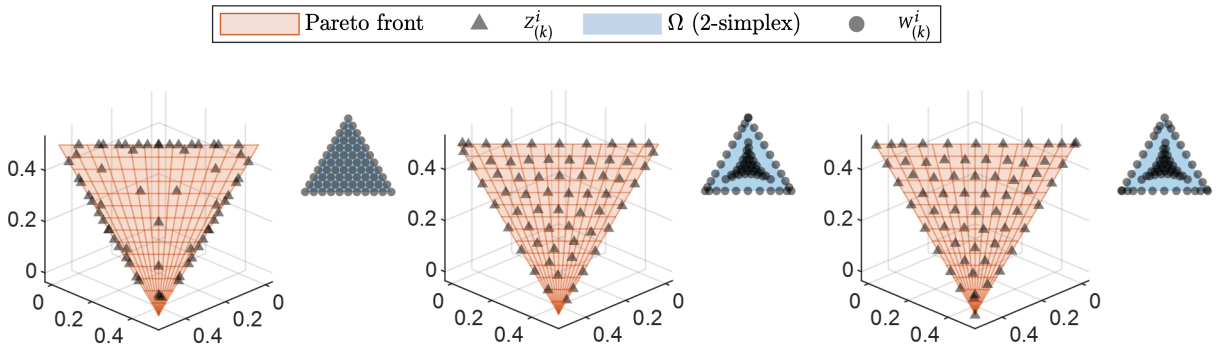

(a) No update, $\mathcal{U}[\rho^N_{(k)}] = 1.92e-02$

(b) Update (5.15), $\mathcal{U}[\rho^N_{(k)}] = 1.23e-02$

(c) Update (5.15) with noise (5.16), $\mathcal{U}[\rho^N_{(k)}] = 1.26e-02$

Figure 5.9: Tri-objective inverted DTLZ1 problem.

(a) No update, $\mathcal{U}[\rho^N_{(k)}] = 6.31e-02$

(b) Update (5.15), $\mathcal{U}[\rho^N_{(k)}] = 6.20e-02$

(c) Update (5.15) with noise (5.16), $\mathcal{U}[\rho^N_{(k)}] = 1.30e-02$

Figure 5.10: Tri-objective Lamé problem with $\gamma = 0.5$.

(a) No update, $\mathcal{U}[\rho^N_{(k)}] = 4.83e-02$

(b) Update (5.15), $\mathcal{U}[\rho^N_{(k)}] = 4.39e-02$

(c) Update (5.15) with noise (5.16), $\mathcal{U}[\rho^N_{(k)}] = 3.89e-03$

Figure 5.11: Tri-objective Lamé problem with $\gamma = 2$.

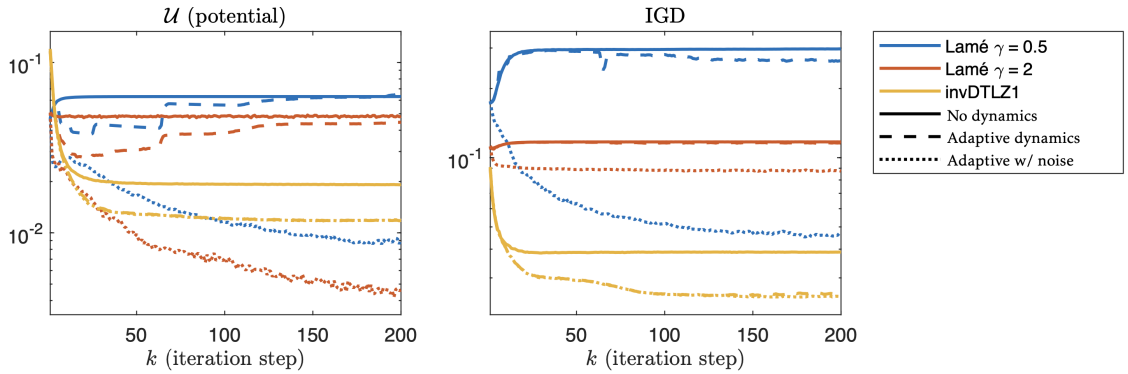

Figure 5.12: Evolution of two performance metrics: the Morse potential and the IGD as functions of the iteration step k with different repulsive dynamics for the tri-objective problems considered.

solutions in terms of IGD in all cases.

5.4.3 Problems definition

We report here the definition of the benchmark problems, together with the penalization strategy and a known parameterization of F. The Lamé [31] and the DO2DK [14] problems were originally formulated as constrained multi-objective optimization problems where the search space is given by $H = [0,1]^{\mathrm{d}}$. The set of EP optimal points corresponds to $[0,1]\times\{0\}^{\mathrm{d}-1}$. Adding a projection step to H has a significant impact on the dynamics of the algorithm, as any point belonging to $\mathbb{R}\times\mathbb{R}^{\mathrm{d}-1}_{\leq 0}$ is projected to an EP optimal point. Therefore, we make use of an exact penalization strategy to ensure that particles stay in the feasible region. We use a ℓ_1-penalty term of the form $\beta\mathrm{dist}(x,H)$, $\beta > 0$, which we add the original objective functions.

Let $x \in \mathbb{R}^{\mathrm{d}}, x = (x_1,\dots,x_{\mathrm{d}})$ for $\mathrm{d} \geq 1$ the objective functions are given by

- Lamé [31] with $\gamma \in \mathbb{R}_{>0}$. Let $r(x) = \sqrt{\sum_{i=2}^{\mathrm{d}} x_i^2}$, for $m = 2$ the objectives read

$$\overline{\mathcal{E}}_1(x) = \left|\cos\left(\frac{\pi}{2}x_1\right)\right|^{\frac{2}{\gamma}}(1+r(x)) + \frac{\pi}{\gamma}\mathrm{dist}(x,H)$$
$$\overline{\mathcal{E}}_2(x) = \left|\sin\left(\frac{\pi}{2}x_1\right)\right|^{\frac{2}{\gamma}}(1+r(x)) + \frac{\pi}{\gamma}\mathrm{dist}(x,H),$$

while, for $m = 3$,

$$\overline{\mathcal{E}}_1(x) = \left|\cos\left(\frac{\pi}{2}x_1\right)\right|^{\frac{2}{\gamma}}(1 + r(x)) + \frac{\pi}{\gamma}\mathrm{dist}(x, H)$$
$$\overline{\mathcal{E}}_2(x) = \left|\sin\left(\frac{\pi}{2}x_1\right)\cos\left(\frac{\pi}{2}x_2\right)\right|^{\frac{2}{\gamma}}(1 + r(x)) + \frac{\pi}{\gamma}\mathrm{dist}(x, H)$$
$$\overline{\mathcal{E}}_2(x) = \left|\sin\left(\frac{\pi}{2}x_1\right)\sin\left(\frac{\pi}{2}x_2\right)\right|^{\frac{2}{\gamma}}(1 + r(x)) + \frac{\pi}{\gamma}\mathrm{dist}(x, H).$$

- DO2DK [14] with $k \in \mathbb{N}, s \in \mathbb{R}_{>0}$, for $m = 2$ the objectives reads

$$\overline{\mathcal{E}}_1(x) = \sin\left(\frac{\pi}{2}x_1 + \left(1 + \frac{2^s - 1}{2^{s+2}}\right)\pi + 1\right)r_a(x)r_b(x) + 10\mathrm{dist}(x, H)$$
$$\overline{\mathcal{E}}_2(x) = \left(\cos\left(\frac{\pi}{2}x_1 + \pi\right) + 1\right)r_a(x)r_b(x) + 10\mathrm{dist}(x, H)$$

  with

$$r_a(x) = 1 + \frac{9}{\mathrm{d} - 1}\sum_{i=2}^{\mathrm{d}} x_i$$
$$r_b(x) = 5 + 10\left(x_1 - \frac{1}{2}\right)^2 + \frac{2^{\frac{s}{2}}\cos(2k\pi x_1)}{k}.$$

- Inverted DTLZ1 [61]. For $m = 3$, the objectives read

$$\overline{\mathcal{E}}_1(x) = (1 - x_1x_2)\frac{1 + r(x)}{2} + 10\mathrm{dist}(x, H)$$
$$\overline{\mathcal{E}}_2(x) = (1 - x_1(1 - x_2))\frac{1 + r(x)}{2} + 10\mathrm{dist}(x, H)$$
$$\overline{\mathcal{E}}_3(x) = -x_1\frac{1 + r(x)}{2} + 10\mathrm{dist}(x, H)$$

  with

$$r(x) = 100(\mathrm{d} - m + (x_{\mathrm{d}} - 0.5)^2 - \cos(20\pi x_{\mathrm{d}})).$$

The parameterization used to construct reference solutions is given by

$$h : [0, 1] \to F, \quad h(r) = g(r, 0, \ldots, 0).$$

5.5 Discussion

In this chapter, we proposed a CBO algorithm for multi-objective optimization problems. Via scalarization, the multi-objective problems is translated into several parameterized sub-problems which are efficiently solved in parallel by the system. This is done by assigning to every particle (or to a small group of particles) a different sub-problem. The particles exchange information through the computation of the consensus points. Via mean-field analysis, we proved that the particles converge towards EP optimal solutions, provided the scalarized objectives share common lower and upper bounds around their global minimizers.

We also suggested different adaptive strategies to choose the scalarized sub-problems leading to a uniform approximation of the Pareto front. Such adaptive dynamics make use of short-range repulsive potentials to distance the particles over the objective space so that they are well-spread over the front. Several numerical experiments validated the suggested strategy both in case of bi- and tri-objective problems. As indicated by the theoretical mean-field analysis, convergence to EP optimal points is ensured only if the adaptive dynamics have a slower time-scale with respect to the CBO one.

We remark that the suggested strategy where the work is distributed among the particles can be applied to any optimization problem where one needs to solve two or more minimization problems in parallel.

Chapter 6

Conclusion and Outlook

6.1 Conclusion

In this thesis, we proposed a new mean-field model to study the convergence properties of the time-discrete CBO particle dynamics. The model consists of a difference equation of McKean type which approximates the algorithm update rule for a large number of particles N. Thanks to the simplicity of the consensus dynamics, we adapted the analysis developed for the time-continuous mean-field model to such semi-discrete settings. With this approach, we have been able to simplify the theoretical analysis of CBO by avoiding an unnecessary approximation step. In the case of bounded domains, we have been able to quantity in terms number of particles N the error introduced by the mean-field approximation.

Next, we suggested a novel strategy to solve constrained optimization problems. Via exact penalization and a suitable adaptive strategy for the penalty parameter, we proved convergence towards the constrained global minima in mean-field law. In this context, avoiding a time-continuous description has allowed us to include the additional mechanism in the theoretical analysis.

With the same methodology, we designed a CBO algorithm for multi-objective optimization problems. In this context, we proposed to distribute the computational effort among the particles, leading to a highly efficient algorithm where a different parameterized scalar sub-problem is associated with every particle. Inspired by many recent works on repulsive dynamics over bounded domains, we also suggested an additional adaptive mechanism to select the parameters of the sub-problems. The aim here was to provide a good approximation of the Pareto front by computing low-energy configurations of the particle system. During the computation, the energy is decreased by updating the parameters with gradient-flow-like dynamics. Therefore, taking the point of view of mathematical physicists has not only allowed us to theoretically analyze CBO algorithms, but also served as an inspiration to design novel, mathematically sound, optimizing dynamics.

6.2 Outlook

As shown in Section 2.1, CBO is closely related to Particle Swarm Optimization (PSO) methods. Time-continuous mean-field models for PSO have been proposed in [43] and further studied in [59]. The approach suggested in this thesis may also be applied in this context to derive a semi-discrete description of the PSO dynamics. What about other heuristics? In the work [13] we focus on deriving a kinetic description of the popular Genetic Algorithm (GA). We note that preliminary work has been conducted in [2]. In [2], though, authors suggest a single-parent approach that does not model the crossover mechanism which is typical of GA methods. Thus, further investigations are needed to fully explain the success of such heuristics.

Another promising research direction consists of extending stochastic particle methods, like CBO, to infinite-dimensional settings. Robust, global optimization algorithms in general Hilbert, Banach, or even just metric spaces could be applied in any different areas such as optimal control problems or learning problems. As we have seen, the core idea of the CBO methods is made of a rather simple consensus dynamics, which may be generalized to spaces with a weaker structure than the Euclidean one. Ongoing work investigates multi-agent consensus dynamics in Wasserstein spaces of probability measures with possible application in the training of Generative Models, Neural Networks or in the numerical simulation of Gradient Flows.

Bibliography

[1] G. Albi, N. Bellomo, L. Fermo, S.-Y. Ha, J. Kim, L. Pareschi, D. Poyato, and J. Soler. Vehicular traffic, crowds, and swarms: From kinetic theory and multiscale methods to applications and research perspectives. *Mathematical Models and Methods in Applied Sciences*, 29(10):1901–2005, 2019.

[2] G. Albi, F. Ferrarese, and C. Totzeck. Kinetic-based optimization enhanced by genetic dynamics. *Mathematical Models and Methods in Applied Sciences*, 33(14):2905–2933, 2023.

[3] G. Albi, L. Pareschi, G. Toscani, and M. Zanella. Recent advances in opinion modeling: control and social influence. In N. Bellomo, D. Pierre, and T. Eitan, editors, *Active Particles, Volume 1*, Modeling and Simulation in Science, Engineering and Technology, pages 49–98. Birkhäuser, Cham, 2017.

[4] D. Armbruster, S. Motsch, and A. Thatcher. Swarming in bounded domains. *Physica D: Nonlinear Phenomena*, 344:58–67, 2017.

[5] A. Benfenati, G. Borghi, and L. Pareschi. Binary interaction methods for high dimensional global optimization and machine learning. *Applied Mathematics & Optimization*, 86(1):9, June 2022.

[6] D. Bertsekas. *Constrained Optimization and Lagrange Multiplier Methods*. Academic Press, 1982.

[7] J. Bonnans, J. Gilbert, C. Lemarechal, and C. Sagastizábal. *Numerical Optimization: Theoretical and Practical Aspects*. Universitext. Springer Berlin Heidelberg, 2013.

[8] G. Borghi. Repulsion dynamics for uniform Pareto front approximation in multi-objective optimization problems. *PAMM*, 23(1):e202200285, 2023.

[9] G. Borghi, S. Grassi, and L. Pareschi. Consensus based optimization with memory effects: Random selection and applications. *Chaos, Solitons & Fractals*, 174:113859, 2023.

[10] G. Borghi, M. Herty, and L. Pareschi. A consensus-based algorithm for multi-objective optimization and its mean-field description. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4131–4136, 2022.

[11] G. Borghi, M. Herty, and L. Pareschi. An adaptive consensus based method for multi-objective optimization with uniform Pareto front approximation. *Applied Mathematics & Optimization*, 88(2):58, Aug 2023.

[12] G. Borghi, M. Herty, and L. Pareschi. Constrained consensus-based optimization. *SIAM Journal on Optimization*, 33(1):211–236, 2023.

[13] G. Borghi and L. Pareschi. Kinetic description and convergence analysis of genetic algorithms for global optimization. *arXiv preprint arXiv:2310.08562*, 2023.

[14] J. Branke, K. Deb, H. Dierolf, and M. Osswald. Finding knees in multi-objective optimization. In X. Yao, E. K. Burke, J. A. Lozano, J. Smith, J. J. Merelo-Guervós, J. A. Bullinaria, J. E. Rowe, P. Tiňo, A. Kabán, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature - PPSN VIII*, pages 722–731, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[15] J. V. Burke. An exact penalization viewpoint of constrained optimization. *SIAM J. Control Optim.*, 4(29):968–998, 1991.

[16] J. A. Carrillo, Y.-P. Choi, C. Totzeck, and O. Tse. An analytical framework for consensus-based global optimization method. *Math. Models Methods Appl. Sci.*, 28(6):1037–1066, 2018.

[17] J. A. Carrillo, M. Fornasier, J. Rosado, and G. Toscani. Asymptotic flocking dynamics for the kinetic Cucker–Smale model. *SIAM Journal on Mathematical Analysis*, 42(1):218–236, 2010.

[18] J. A. Carrillo, F. Hoffmann, A. M. Stuart, and U. Vaes. Consensus-based sampling. *Studies in Applied Mathematics*, 148(3):1069–1140, 2022.

[19] J. A. Carrillo, S. Jin, L. Li, and Y. Zhu. A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM: COCV*, 27:S5, 2021.

[20] J. A. Carrillo, D. Slepčev, and L. Wu. Nonlocal-interaction equations on uniformly prox-regular sets. *Discrete and Continuous Dynamical Systems*, 36(3):1209–1247, 2016.

[21] J. A. Carrillo, C. Totzeck, and U. Vaes. Consensus-based optimization and ensemble kalman inversion for global optimization problems with constraints. In *Modeling and Simulation for Collective Dynamics*, pages 195–230, 2023.

[22] S. Cheng, Y. Shi, and Q. Qin. On the performance metrics of multiobjective optimization. In *International Conference in Swarm Intelligence*, pages 504–512. Springer, 2012.

[23] C. A. Coello Coello, S. González Brambila, J. Figueroa Gamboa, M. G. Castillo Tapia, and R. Hernández Gómez. Evolutionary multiobjective optimization: open research areas and some challenges lying ahead. *Complex & Intelligent Systems*, 6(2):221–236, Jul 2020.

[24] S. Cordier, L. Pareschi, and G. Toscani. On a kinetic model for a simple market economy. *Journal of Statistical Physics*, 120:253–277, 2005.

[25] F. Cucker and S. Smale. Emergent behavior in flocks. *IEEE Transactions on automatic control*, 52(5):852–862, 2007.

[26] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Trans. Evol. Comp*, 6(2):182–197, apr 2002.

[27] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer-Verlag Berlin Heidelberg, 2010.

[28] R. Eftimie, G. de Vries, M. A. Lewis, and F. Lutscher. Modeling group formation and activity patterns in self-organizing collectives of individuals. *Bulletin of Mathematical Biology*, 69(5):1537–1565, Jul 2007.

[29] M. Ehrgott. *Multicriteria Optimization*. Springer, Berlin, Heidelberg, 2005.

[30] G. Eichfelder. Twenty years of continuous multiobjective optimization in the twenty-first century. *EURO Journal on Computational Optimization*, 9:100014, 2021.

[31] M. T. M. Emmerich and A. H. Deutz. Test problems based on Lamé superspheres. In *Proceedings of the 4th International Conference on Evolutionary Multi-Criterion Optimization*, EMO'07, page 922–936, Berlin, Heidelberg, 2007. Springer-Verlag.

[32] J. G. Falcón-Cardona, E. Covantes Osuna, and C. A. Coello Coello. An overview of pair-potential functions for multi-objective optimization. In H. Ishibuchi, Q. Zhang, R. Cheng, K. Li, H. Li, H. Wang, and A. Zhou, editors, *Evolutionary Multi-Criterion Optimization*, pages 401–412, Cham, 2021. Springer International Publishing.

[33] R. C. Fetecau. Collective behavior of biological aggregations in two dimensions: A nonlocal kinetic model. *Mathematical Models and Methods in Applied Sciences*, 21(07):1539–1569, 2011.

[34] R. C. Fetecau and M. Kovacic. Swarm equilibria in domains with boundaries. *SIAM Journal on Applied Dynamical Systems*, 16, 08 2016.

[35] R. C. Fetecau and M. Kovacic. Swarm equilibria in domains with boundaries. *SIAM J. Appl. Dyn. Syst.*, 16:1260–1308, 2017.

[36] R. C. Fetecau, M. Kovacic, and I. Topaloglu. Swarming in domains with boundaries: Approximation and regularization by nonlinear diffusion. *Discrete and Continuous Dynamical Systems - B*, 24(4):1815–1842, 2019.

[37] J. Fliege, L. M. G. n. Drummond, and B. F. Svaiter. Newton's method for multiobjective optimization. *SIAM Journal on Optimization*, 20(2):602–626, 2009.

[38] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based optimization on the sphere: Convergence to global minimizers and machine learning. *J. Machine Learning Research*, 22(237):1–55, 2021.

[39] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Anisotropic diffusion in consensus-based optimization on the sphere. *SIAM Journal on Optimization*, 32(3):1984–2012, 2022.

[40] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based optimization on hypersurfaces: Well-posedness and mean-field limit. *Mathematical Models and Methods in Applied Sciences*, 30(14):2725–2751, 2020.

[41] M. Fornasier, T. Klock, and K. Riedl. Consensus-based optimization methods converge globally. *arXiv:2103.15130*, 2021.

[42] M. Fornasier, T. Klock, and K. Riedl. Convergence of anisotropic consensus-based optimization in mean-field law. In J. L. Jiménez Laredo, J. I. Hidalgo, and K. O. Babaagba, editors, *Applications of Evolutionary Computation*, pages 738–754, Cham, 2022. Springer International Publishing.

[43] S. Grassi and L. Pareschi. From particle swarm optimization to consensus based optimization: Stochastic modeling and mean-field limit. *Mathematical Models and Methods in Applied Sciences*, 31(08):1625–1657, 2021.

[44] L. Graña Drummond and B. Svaiter. A steepest descent method for vector optimization. *Journal of Computational and Applied Mathematics*, 175(2):395–414, 2005.

[45] M. Gugat and M. Herty. The smoothed-penalty algorithm for state constrained optimal control problems for partial differential equations. *Optim. Methods Softw.*, 25(4-6):573–599, 2010.

[46] M. Gugat and M. Herty. A smoothed penalty iteration for state constrained optimal control problems for partial differential equations. *Optimization*, 62(3):379–395, 2013.

[47] S.-Y. Ha, T. Ha, and J.-H. Kim. Emergent behavior of a Cucker-Smale type particle model with nonlinear velocity couplings. *IEEE Transactions on Automatic Control*, 55(7):1679–1683, 2010.

[48] S.-Y. Ha, S. Jin, and D. Kim. Convergence of a first-order consensus-based global optimization algorithm. *Mathematical Models and Methods in Applied Sciences*, 30(12):2417–2444, 2020.

[49] S.-Y. Ha, S. Jin, and D. Kim. Convergence and error estimates for time-discrete consensus-based optimization algorithms. *Numerische Mathematik*, 147(2):255–282, Feb 2021.

[50] S.-Y. Ha, M. Kang, D. Kim, J. Kim, and I. Yang. Stochastic consensus dynamics for nonconvex optimization on the Stiefel manifold: Mean-field limit and convergence. *Mathematical Models and Methods in Applied Sciences*, 0(0):1–85, 2022.

[51] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration*, volume 31 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, second edition, 2006. Structure-preserving algorithms for ordinary differential equations.

[52] D. Helbing. *Quantitative Sociodynamics: Stochastic Methods and Models of Social Interaction Processes*. Springer Berlin Heidelberg, 2010.

[53] M. Herty, A. Klar, A. K. Singh, and P. Spellucci. Smoothed penalty algorithms for optimization of nonlinear models. *Comput. Optim. Appl.*, 37(2):157–176, 2007.

[54] M. Herty and G. Visconti. Kinetic methods for inverse problems. *Kinetic and Related Models*, 12(5):1109–1130, 2019.

[55] D. J. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3):525–546, 2001.

[56] E. H. Houssein, A. G. Gad, K. Hussain, and P. N. Suganthan. Major advances in particle swarm optimization: Theory, analysis, and application. *Swarm and Evolutionary Computation*, 63:100868, 2021.

[57] H. Huang and J. Qiu. On the mean-field limit for the consensus-based optimization. *Mathematical Methods in the Applied Sciences*, 45(12):7814–7831, 2022.

[58] H. Huang, J. Qiu, and K. Riedl. Consensus-based optimization for saddle point problems, 2022.

[59] H. Huang, J. Qiu, and K. Riedl. On the global convergence of particle swarm optimization methods. *Applied Mathematics & Optimization*, 88(2):30, 2023.

[60] K. Hussain, M. N. Mohd Salleh, S. Cheng, and Y. Shi. Metaheuristic research: a comprehensive survey. *Artificial Intelligence Review*, 52(4):2191–2233, Dec 2019.

[61] H. Ishibuchi, Y. Setoguchi, H. Masuda, and Y. Nojima. Performance of decomposition-based many-objective algorithms strongly depends on pareto front shapes. *IEEE Transactions on Evolutionary Computation*, 21(2):169–190, 2017.

[62] J. Jahn. *Vector optimization - theory, applications, and extensions*. Springer, Berlin, Heidelberg, 2004.

[63] M. Jamil and X.-S. Yang. A literature survey of benchmark functions for global optimization problems. *Int. Journal of Mathematical Modelling and Numerical Optimisation*, 2(4):150–194, 2013.

[64] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks*, volume 4, pages 1942–1948 vol.4, 1995.

[65] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.

[66] K. Klamroth, M. Stiglmayr, and C. Totzeck. Consensus-based optimization for multi-objective problems: A multi-swarm approach. *arXiv:2103.15130*, 2022.

[67] D. Ko, S.-Y. Ha, S. Jin, and D. Kim. Convergence analysis of the discrete consensus-based optimization algorithm with random batch interactions and heterogeneous noises. *Mathematical Models and Methods in Applied Sciences*, 32(06):1071–1107, 2022.

[68] L. A. Márquez-Vega, J. G. Falcón-Cardona, and E. Covantes Osuna. Towards a Pareto front shape invariant multi-objective evolutionary algorithm using pair-potential functions. In I. Batyrshin, A. Gelbukh, and G. Sidorov, editors, *Advances in Computational Intelligence*, pages 369–382, Cham, 2021. Springer International Publishing.

[69] V. Martínez-Cagigal. Non Sorting Genetic Algorithm II (NSGA-II). `https://www.mathworks.com/matlabcentral/fileexchange/65494-non-sorting-\genetic-algorithm-ii-nsga-ii`, 2023. [Online; Retrieved March 23, 2023].

[70] G. Naldi, L. Pareschi, and G. Toscani. *Mathematical modeling of collective behavior in socio-economic and life sciences*. Springer Science & Business Media, 2010.

[71] L. Pareschi and G. Toscani. *Interacting Multiagent Systems: Kinetic equations and Monte Carlo methods*. Oxford University Press, 2013.

[72] R. Pinnau, C. Totzeck, O. Tse, and S. Martin. A consensus-based model for global optimization and its mean-field limit. *Math. Models Methods Appl. Sci.*, 27(1):183–204, 2017.

[73] E. Platen. An introduction to numerical methods for stochastic differential equations. *Acta Numerica*, 8:197–246, 1999.

[74] K. Riedl. Leveraging memory effects and gradient information in consensus-based optimisation: On global convergence in mean-field law. *European Journal of Applied Mathematics*, page 1–32, 2023.

[75] F. Santambrogio. *Optimal Transport for Applied Mathematicians*. Birkhäuser, 2015.

[76] Y. D. Sergeyev, D. E. Kvasov, and M. S. Mukhametzhanov. On the efficiency of nature-inspired metaheuristics in expensive global optimization with limited budget. *Scientific Reports*, 8(1):453, 2018.

[77] P. Spellucci. Solving QP problems by penalization and smoothing. *preprint, TU Darmstadt*, 2002.

[78] A.-S. Sznitman. Topics in propagation of chaos. In P.-L. Hennequin, editor, *Ecole d'Eté de Probabilités de Saint-Flour XIX — 1989*, pages 165–251, Berlin, Heidelberg, 1991. Springer Berlin Heidelberg.

[79] K.-S. Tang, K.-F. Man, S. Kwong, and Q. He. Genetic algorithms and their applications. *IEEE signal processing magazine*, 13(6):22–37, 1996.

[80] C. M. Topaz and A. L. Bertozzi. Swarming patterns in a two-dimensional kinematic model for biological groups. *SIAM Journal on Applied Mathematics*, 65(1):152–174, 2004.

[81] M. Ulbrich and S. Ulbrich. *Nichtlineare Optimierung*. Springer-Verlag, 2012.

[82] D. A. Van Veldhuizen, G. B. Lamont, et al. Evolutionary computation and convergence to a Pareto front. In *Late breaking papers at the genetic programming 1998 conference*, pages 221–228. Citeseer, 1998.

[83] Q. Zhang and H. Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *Evolutionary Computation, IEEE Transactions on*, 11:712 – 731, 01 2008.

[84] E. Zitzler and L. Thiele. Multiobjective optimization using evolutionary algorithms - a comparative case study. In *Proceedings of the 5th International Conference on Parallel Problem Solving from Nature*, PPSN V, page 292–304, Berlin, Heidelberg, 1998. Springer-Verlag.

Eidesstattliche Erklärung - Declaration of Authorship

I, Giacomo Borghi, declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

Hiermit erkläre ich an Eides statt / I do solemnely swear that:

1. This work was done wholly or mainly while in candidature for the doctoral degree at this faculty and university;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this university or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others or myself, this is always clearly attributed;
4. Where I have quoted from the work of others or myself, the source is always given. This thesis is entirely my own work, with the exception of such quotations;
5. I have acknowledged all major sources of assistance;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published before as:
    - G. Borghi. Repulsion dynamics for uniform Pareto front approximation in multi-objective optimization problems. *PAMM*, 23(1):e202200285, 2023.
    - G. Borghi, S. Grassi, and L. Pareschi. Consensus based optimization with memory effects: Random selection and applications. *Chaos, Solitons & Fractals*, 174:113859, 2023.
    - G. Borghi, M. Herty, and L. Pareschi. A consensus-based algorithm for multi-objective optimization and its mean-field description. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4131–4136, 2022.
    - G. Borghi, M. Herty, and L. Pareschi. An adaptive consensus based method for multi-objective optimization with uniform Pareto front approximation. *Applied Mathematics & Optimization*, 88(2):58, Aug 2023.
    - G. Borghi, M. Herty, and L. Pareschi. Constrained consensus-based optimization. *SIAM Journal on Optimization*, 33(1):211–236, 2023.

February 20, 2024, Aachen

_____

(Giacomo Borghi)