Consensus Based Optimization with Finite Range Interaction

Brünnette, Tim

Award date:
2021

Awarding institution:
Universität Stuttgart

TU/e

University of Stuttgart
Germany

MASTER'S THESIS

Consensus Based Optimization with Finite Range Interaction

submitted by

Author Tim BRÜNNETTE

Supervisor
Prof. Dr. Oliver TSE

Supervisor
Prof. Dr. Andrea BARTH

Contents

Chapter 1

Introduction

Interacting particle systems are models used in research to gain an understanding of the often complex behavior of dynamic groups and populations. The phenomenons covered by this can stem from physics, where the particles represent physical particles, but also frequently from fields such as biology or social sciences where they emulate individuals or states of individuals. Well known examples include bird flocking behavior([MT11]) and opinion formation([Lig99]).

In the field of optimization many meta-heuristics have been developed over the last decades, such as Simulated Annealing, Genetic Programming or Random Search ([GKR94]). Some of these draw their motivation from the aforementioned field of particle systems or can at least be interpreted in such ways. Notable examples of such particle based optimization meta-heuristics are Ant Colony Optimization, the Bees Algorithm or Particle Swarm Optimization (see [CHS03], [YPM$^+$13], [BM17]).

These methods have been a staple of optimization for a while now and they are applied to both, continuous and discrete problems. The surge of machine learning techniques and especially neural networks has only increased the need to reliably find good optima, even more so for extremely high dimensional problems. Stochastic Gradient Langevin Dynamics (SGLD) are one of the most popular optimization methods to train neural nets and e.g. [KPP19] proposes an interacting version of it. All these methods have a certain degree of randomness in common, which causes the fact that many of the mathematical tools that are used for the analysis come from the area of probability theory.

The **Consensus Based Optimization** (CBO), first introduced in [PTTM17], is such a particle based optimization algorithm with possible applications for machine learning tasks. The method is based on N particles roaming and evaluating the function landscape. Most importantly, the method does not use gradient information and is therefore suited to almost any continuous optimization problem. Further, given its particle nature, the parallelization of the method is very natural.

The function evaluations at the particle positions are used to compute a common attraction point. The attraction point is calculated as a weighted center of mass of the particles, with the weights according to the respective evaluation of the transformed cost function. Motivated by the Laplace principle this transformation is $\omega_f^\alpha(x) = \exp\big(-\alpha f(x)\big)$ for $\alpha > 0$. The bigger the parameter α, the more distinguished the optima of the transformed cost function. Compared to the often seen choice of using the best particles position (i.e. an arg min in the formulation) as the attraction point, this approach allows for a more rewarding analytical investigation.

For each particle, the drift term steers it towards the attraction point while the volatility term guarantees the explorative nature of the method. The magnitude of the distance from the attraction point determines the amount of randomness in each particles movement. The further away the particles, the bigger the volatility and thus the random exploration.

While the authors do show well-posedness of this particle system under reasonable assumptions for the target function, most of the analysis is carried out on the postulated corresponding mean-field limit process. Here some convergence towards unique minima can indeed be guaranteed.

In Chapter 2 we give a more thorough mathematical overview over this system and its properties. We also take a look at the refinements and additional results that have been introduced since in the literature. The article [CCTT18] gives the aforementioned convergence guarantee, [HJK20] focuses on the particle system instead of the mean-field SDE in its analysis regarding convergence, while [CJLZ21] proposes to cancel the Brownian motion and thus the randomness of the system dimension wise. The series [FHPS20], [FHPS21b], [FHPS21a] gives a full analysis of the method on compact hyperspheres, while [TW20] introduces the use of particle history.

Subsequently, in the main parts of this work in Chapter 3 and Chapter 4 we first present one possible modification to the CBO method where the communication be-

tween the particles is no longer uniform and specifically the interaction only happens with limited range. Chapter 3 also includes a well-posedness proof of both, the particle system and its resulting postulated mean-field limit. Especially the latter differs noticeably from the original well-posedness proof in [CCTT18]. It uses a fixed point theorem on the space of measure-valued curves. In Chapter 4 we tackle one gap that remains so far in the literature. The convergence of the particle system towards the postulated mean field limit has not yet been shown rigorously. We explore the use of Large Deviations to close that gap. Lastly, Chapter 5 contains numerical studies regarding the applicability and utility of the previously given modification with the use of simple examples.

Chapter 2

Previously on CBO

We refer to the formulations from [CCTT18] as the standard CBO particle system and standard CBO mean field limit respectively. To introduce this, we prescribe the behavior, for $i = 1, \dots, N$, $N \in \mathbb{N}$ different particles, in the following way:

$$\mathrm{d}X_t^i = -\lambda \left(X_t^i - m_t \right) \mathrm{d}t + \sigma \left| X_t^i - m_t \right| \mathrm{d}W_t^i \,, \tag{2.1a}$$

$$m_t = \frac{1}{\sum_{j=1}^N \omega_f^\alpha(X_t^j)} \sum_{j=1}^N X_t^j \omega_f^\alpha(X_t^j) \,, \tag{2.1b}$$

$$\omega_f^\alpha(x) = \exp\left(-\alpha f(x)\right) \,. \tag{2.1c}$$

This is a weakly interacting system of stochastic differential equations(SDEs) in the sense of Itô, where each of the N coupled equations describes the movement of a particle in a d-dimensional landscape, with W_t^i representing N independent Brownian motions of dimension $d \in \mathbb{N}$ on a Borel probability space. Lastly, $\lambda, \sigma, \alpha > 0$ are parameters of the system. Existence and uniqueness of a strong solution for each $N \in \mathbb{N}$, all times $t > 0$ and for initial conditions satisfying $\mathbb{E}\left|\mathbf{X}_0^{(N)}\right|^2 < \infty$ is shown in [CCTT18].

In [PTTM17] the CBO method is presented for the first time. Here, the drift term still includes a smoothed Heaviside function that is multiplied with the drift term. Equation (2.1a) is thus instead formulated as

$$\mathrm{d}X_t^i = -\lambda \left(X_t^i - m_t \right) H_\epsilon \left(f(X_t^i) - f(m_t) \right) \mathrm{d}t + \sigma \left| X_t^i - m_t \right| \mathrm{d}W_t^i \,. \tag{2.2}$$

In this version, the Heaviside function ensures that an attraction only happens towards a better solution. However, this term is subsequently dropped in order to make the system more accessible for analysis. The core mathematical properties are still preserved in this updated version and the numerical experiments are showing no significant advantages or disadvantages regarding the performance for either version.

Excursion: Solutions to SDEs

This excursion is supposed to give merely an overview to fill a potential knowledge gap and enable an understanding of the following passages. The formal definitions include some additional concepts such as filtrations and martingales, which we can mostly sidestep in this context. A very complete introduction to the topic can be found in [KS98], Chapter 1-3,5 and a shorter and thus possibly more accessible variant in [Dur96], Chapter 1,2,5.

Generally speaking, stating a solution to an SDE requires the existence of a complete probability space with a family of σ-algebras $(\Omega, \mathscr{F}, P)$ as a basic requirement. That family $F = \{F_t\}_{t\geq 0}$ is called the filtration. The filtration is often said to encode the available information and it is understandable that the available information must suffice to describe the behavior of the process. This is then phrased as "the process is adapted to the filtration". In many cases, one is not limited to a specific filtration beforehand and is thus free to pick exactly the filtration with the minimal set of information which enables one to describe the process. This is called the "filtration generated by the process" or "natural filtration". In fact, it is so natural that it is often not explicitly talked about anymore. This is also true in our context, where the σ-algebra is generally the Borel-σ-algebra and the probability measures are Borel-measures, possibly with some additional requirements at times.

The equation

$$\mathrm{d}X_t = b(t, X_t)\,\mathrm{d}t + \sigma(t, X_t)\,\mathrm{d}W_t$$

is shorthand for

$$X_t = X_0 + \int_0^t b(t, X_t)\,\mathrm{d}t + \int_0^t \sigma(t, X_t)\,\mathrm{d}W_t\,,$$

which is calculable with the definition of the Itô integral

$$\int_a^b f(t)\,\mathrm{d}W_t = \lim_{n\to\infty} \sum_{[t_{i-1},t_i]\in\pi_n} f(t_{i-1})\left(W_{t_i} - W_{t_{i-1}}\right)$$

where the π_n are a sequence of refining partitions. Thus, it resembles the definition of the Riemann integral, but the evaluations at the mesh points are weighted by the stochastic process against which we are integrating. The evaluation of the function happens at the left side of the interval at t-1 and therefore yields the Itô integral, which is the chosen formulation in this work. A function evaluation at the midpoint of the interval would lead to convergence to a different object, the Stratonovich integral - an important distinction.

This works for a whole class of processes, but most importantly also for the Brownian motion W_t. In that case, the difference $W_{t_i} - W_{t_{i-1}}$ is a random variable with a normal distribution. As usual, there is much more to this definition, its construction and its consequences (basically the whole field of stochastic calculus) as can be experienced by looking into the aforementioned references.

Now since the the solution is a stochastic process, it does not surprise that there are different notions of a solution and with these solutions come of course also different notions of uniqueness. For one, there is the **strong** solution. Here, given a probability space, a Brownian motion and an initial condition X_0, one can give a process X_t that fulfills the integral formulation of the SDE. Two solution processes X_t, X'_t are indistinguishable if

$$P(\sup_{t\geq 0}\left|X_t - X'_t\right| = 0) = 1\,,$$

so with full probability two paths for a shared given Brownian motion do not deviate.

For the **weak** solution however, the order of operations is different. Here, the process still has to fulfill the integral formulation, but one is allowed to pick the process and the Brownian motion (and the underlying probability space and filtration) together. Concretely, it allows for example for the Brownian motion to be formulated in terms of the solution. The pathwise notion of uniqueness is not applicable anymore since the solutions do not even necessarily exist on the same probability space and therefore cannot necessarily be compared directly by taking the difference. The

correct notion of weak uniqueness is uniqueness in distribution. Two solutions X, X' are the same, if given the same initial distribution the processes as random variables on $C([0,T],\mathbb{R})$ have the same distribution.

As the naming already suggests, a strong solution implies a weak solution and pathwise uniqueness (strong uniqueness) implies uniqueness in distribution (weak uniqueness).

The Attraction Point

The point m_t is the attraction point of the system. To determine this point, the particles X_t^i are weighted by the value of their function evaluation $f(X_t^i)$. However, the function landscape is first stretched with a negative exponential. This is motivated by the Laplace method, stating that

$$\lim_{\alpha\to\infty} -\frac{1}{\alpha}\log\int_{\mathbb{R}^d} e^{-\alpha f}\,\mathrm{d}\rho(x) = f(x^*) > 0 \tag{2.3}$$

for any compactly supported Borel probability measure $\rho \in \mathscr{P}(\mathbb{R}^d)$ with $x^* \in \mathrm{supp}(\rho)$. We can see that $e^{-\alpha f}\rho$ accumulates its mass around x^* because of the exponential stretching dominating the probability measure. Using the empirical measure, i.e. formulating the particle state in terms of the Dirac measure $\delta_{X_t^i}$ at the positions $X_t^i \in \mathbb{R}^d$,

$$\rho_t^N = \frac{1}{N}\sum_{i=1}^N \delta_{X_t^i} \tag{2.4}$$

we can rewrite m_t as

$$m_t = \frac{1}{\left\|\omega_f^\alpha\right\|_{L^1(\rho_t^N)}}\int x\omega_f^\alpha\,\mathrm{d}\rho_t^N\,. \tag{2.5}$$

If ρ_t^N converges in some sense to the measure ρ for $\mathbb{N}\to\infty$ then this last term (2.5) is an approximation of the expectation of the Laplace principle integral from (2.3), which we know to tend towards a point evaluation at the minimum for $\alpha\to\infty$ for a large class of measures. Therefore we can reasonably hope that it gives an approxi-

mation of said minimum x^*.

In Particle Swarm Optimization the attraction point is usually chosen to be $m_t = \operatorname{argmin}_i f(X_t^i)$. The choice given here approximates this term for $\alpha \to \infty$ without singling out a specific particle. This enables us to pass to a mean-field limit, a feat that is not possible when one particle amongst many fulfills a special role.

The distance between the particle and the attraction point also determines the magnitude of its random movement. This can be seen as another episode of the "exploitation vs. exploration" argument. Whenever the particle is still far away form the suspected minimum m_t it can afford more randomness and thus exploration, whereas close particles should really take a more thorough look at the environment around that point. By doing so, they exploit the so far accumulated knowledge, which is used to compute a suspected optimum.

Convergence Guarantees: Mean Field Equation

For the particle system Eq. (2.1) the mean-field limit is postulated as

$$\mathrm{d}\overline{X}_t = -\lambda\left(\overline{X}_t - m[\overline{\mu}_t]\right)\mathrm{d}t + \sigma\left|\overline{X}_t - m[\overline{\mu}_t]\right|\mathrm{d}W_t\,, \tag{2.6a}$$

$$m[\overline{\mu}_t] = \frac{1}{\int \omega_\alpha^f(y)\,\overline{\mu}_t(\mathrm{d}y)}\int y\omega_\alpha^f(y)\,\overline{\mu}_t(\mathrm{d}y)\,, \tag{2.6b}$$

$$\overline{\mu}_t = \mathrm{law}(\overline{X}_t)\,, \tag{2.6c}$$

a d-dimensional stochastic differential equation with the initial condition

$$\mathrm{law}(\overline{X}_0) = \overline{\mu}_0\,. \tag{2.6d}$$

Notice, that the law $\overline{\mu}_t$ of the process $\overline{X}_t$ itself shows up on the right hand side. This is called a non-linearity in the sense of McKean ([McK66]). That fact makes the analysis of the SDE nonstandard. The Fokker-Planck equation corresponding to Equation (2.6) is then

$$\partial_t\overline{\mu}_t = \Delta(\kappa_t\overline{\mu}_t) + \nabla\cdot(\gamma_t\overline{\mu}_t)\,, \tag{2.7}$$

with the drift and diffusion coefficients

$$\kappa_t = -\lambda(x - m[\overline{\mu}_t]) \,, \quad \gamma_t = \frac{\sigma^2}{2} \left| x - m[\overline{\mu}_t] \right|^2 \,. \tag{2.8}$$

This PDE is non-linear and non-local, so the analysis is equally nontrivial here. Nevertheless some convergence guarantees have been determined in [CCTT18]. In order to appropriately talk about the conditions posed on those measures, we define the spaces of Borel probability measures with finite j-th moment.

Definition 2.1 (Measure spaces with finite measure)**.** We define the space

$$\mathscr{P}_j(\mathbb{R}^d) = \left\{ \mu \in \mathscr{P}(\mathbb{R}^d) \,\middle\|\, \int_{\mathbb{R}^d} |x|^j \, \mathrm{d}\mu < \infty \right\} \quad \text{for } j = 1, 2, 4 \tag{2.9}$$

of Borel probability measures with finite j-th moment equipped with the Wasserstein metric

$$W_j(\mu, \hat{\mu}) = \sqrt[j]{\inf \mathbb{E}[\left| Z - \hat{Z} \right|^j]} \tag{2.10}$$

where the infimum is taken over all joint distributions of $Z, \hat{Z}$ with marginals $\mu, \hat{\mu}$.

This metric space is Polish. Convergence in this metric is equivalent to weak convergence of the measures plus convergence of the first j moments.

Further, some restrictions have to be introduced regarding the regularity of the cost function. The assumptions made in [CCTT18] regarding the cost function are:

Assumptions A1. The cost function is bounded from below, so

$$f : \mathbb{R}^d \to \mathbb{R} \,, \quad \inf f = f(x^*) =: \underline{f} \quad . \tag{2.11a}$$

Assumptions A2. A Lipschitz-like regularity and a quadratic growth condition hold for the cost function, i.e.

$$\begin{aligned} \left| f(x) - f(y) \right| &\leq L_f \left(|x| + \left| y \right| \right) \left| x - y \right| \,, \\ f(x) - \underline{f} &\leq c_f (1 + |x|^2) \,, \end{aligned}$$

for all $x, y \in \mathbb{R}^d$ and for $L_f, c_f > 0$.

Assumptions A3. Either the function assumes its maximum or it grows at most quadratically outside of a ball,

$$\overline{f} = \sup f, \qquad \text{or} \qquad f(x) - \underline{f} \geq c_l(1 + |x|^2) \quad \text{for} \quad |x| \geq M\,.$$

Under assumptions A1, A2, A3 the authors prove well-posedness of the mean-field SDE (2.6) for initial conditions $\overline{\mu}_0 \in \mathscr{P}_4(\mathbb{R}^d)$. That is, the SDE has a unique strong solution $\overline{\mu}^*$. This is achieved by defining a mapping from an attraction point curve to another and finding a fixed-point thereof.

Further, the authors prove a concentration estimate, adding the following assumptions on the cost function.

Assumptions A4.

$$\begin{aligned} &f \in C^2(\mathbb{R}^d)\,, \\ &\left\|\nabla^2 f\right\|_\infty \leq c_d \\ &\inf f \geq 0\,, \\ &\Delta f \leq c_0 + c_1 \left|\nabla f\right|^2 \end{aligned}$$

for some constants $c_0, c_1, c_d \leq 0$.

For example, it is possible to smoothly transition the function to a constant outside a large ball and thus fulfill assumption A4, as long as the global minimum is inside the ball and accordingly not changed.

With the assumptions A1 - A4 in place, the authors prove a convergence estimate. For some non-empty set of parameters λ, σ, α there exists a consensus point $\tilde{x}$ with the property

$$\mathbb{E}[\overline{\mu}_t] \to \tilde{x} \quad \text{and} \quad m_t[\overline{\mu}_t] \to \tilde{x} \quad \text{as} \quad t \to \infty\,, \tag{2.12}$$

with the convergence happening in an exponential manner. Further, if $x^* \in \operatorname{supp}(\overline{\mu}_0)$, then $\tilde{x} = x^*$. One condition on the parameters is $\alpha \geq \alpha_0$, the landscape parameter has to pass a certain threshold. Generally speaking one can say that an increase of α always improves the convergence behavior. The intuition behind that is that this

stretches the landscape more radically and thus provides more distinguished optima. While this effect of the α parameter is also reflected in the analysis, the situation regarding the actual computation is not that clear cut. An increase of α can pose numerical problems as the evaluations of $\exp(-\alpha f)$ come very close, potentially computationally indistinguishably close to zero. Further the choice of the parameters λ, σ also depends on the dimension, sometimes making it hard to find suitable setups in real applications.

In [FKR21] an alternative angle at the convergence analysis of the mean-field system is given. Instead of comparing the variance of the particles to their expectation, it is noted that averaged over many instantiations, particles behave similar to the convexification $\|x - x^*\|^2$ of the function. This functional enables a relaxation of the conditions posed on the target function, notably not requiring any differentiability (but expectedly still some form of Lipschitz continuity).

Dimension-Wise Diffusion

In [CJLZ21] there are two modifications proposed for the method. The first one, batch-wise computation, is mostly of computational nature and is discussed in Chapter 5. The second, more fundamental modification concerns the computation of the random component in the particle movement. In the standard scheme (Equation (2.1)), the random movement has the same magnitude in every dimension, depending on the overall distance between the particle and the attraction point.

This is changed and replaced by a dimension-wise assessment of the distance and then a following arrangement of the magnitude of the randomness in that specific dimension. Specifically, equation (2.1a) is replaced by

$$\mathrm{d}X_t^i = -\lambda\left(X_t^i - m_t\right)\mathrm{d}t + \sigma \sum_{k=1}^{d} (X_t^i - m_t)_k \,\mathrm{d}W_t^{i,k}\vec{e}_k\,, \tag{2.13}$$

with $\vec{e}_k$ the k-th unit vector, $W^{i,k}$ independent standard Brownian motions and $(X^i - m_t)_k$ the k-th component of $(X^i - m_t)$. Now, if a particle is close to the attraction point in one dimension, its randomness in that direction will be small. Essentially, the stochasticity of the particles is cancelled dimension-wise. The thereby gained potential to proceed in a more step-by-step manner can understandably lead to a

faster consensus, especially in higher dimensions.

The authors give a proof that, under similar assumptions to A4, the parameter selection is not sensitive to the dimensionality anymore. This is in contrast to the standard CBO. While these results are achieved for the mean-field description, the numerical studies conducted in the paper also seem to confirm this for the particle system.

Convergence Guarantees: Particle System

We shift our focus from the mean-field equation (2.6) back onto the particle system (2.1). In [CCTT18] the well-posedness of this coupled SDE system is shown. This is achieved by showing Lipschitz continuity and sub-linear growth for the drift and diffusion coefficients and then applying standard existence results from SDE theory (see e.g. [Øks03]).

In [HJK20] it is shown that the particle system generally is contractive and thus converges to a consensual final state and that this final state under some circumstances is not too far away from the minimum of the target function. This analysis is applied to both, the continuous formulation as well as the discrete description (the de-facto computational algorithm) of the dimension-wise variant (2.13) with the small change of a shared Brownian motion between the particles.

Concretely, in addition to the requirement of a target function $f \in C_b^2$ with finite second derivatives (similar to Assumptions A1 and A4) we pose another assumption:

Assumptions A5.

$$2\lambda > \sigma^2$$
$$x^* \in \text{supp}(\text{law}(X_0^i))$$

for starting spots X_0^i that are sampled i.i.d.

Here it is required that the initial distribution $\text{law}(X_0^i)$ is essentially "not too bad" in a Laplace principle sense; a very technical and hard to track assumption that is given based on the target function, its second derivatives and the expected diameter of the initial condition. The dominance of the drift over the diffusion, as well as the containment of the optimum in the initialization area are recurring requirements. So

with Assumptions A1, A4, A5 and a good initial condition it holds that

$$\operatorname*{ess\,inf}_{\omega\in\Omega} f(X^{\infty}(\omega)) \leq \underline{f} + \mathcal{O}\left(\frac{1}{\alpha}\right), \tag{2.14}$$

where X^{∞} is the almost sure consensual limit of the particles X_t^i, which gets arbitrarily close to the minimum for growing α.

This establishes thus that both, the particle system as well as the mean field limit exhibit convergence behavior towards a global optimum under suitable conditions. Furthermore, their behavior in numerical experiments is, by eye-check, the same. Nevertheless, despite these indications, a rigorous proof that the postulated system (2.6) is indeed the limiting process for the particle system (2.1) remains to be done.

CBO on a Sphere and other Hypersurfaces

In [FHPS20] a version of the Consensus Based Optimization on connected smooth compact hypersurfaces is introduced. Technically, the process takes place in the space $\mathbb{R}^d$, which the hypersurface is embedded into. The drift is chosen such that particles move on the surface and a third term is introduced with the purpose of compensating for the random component that is normal to the surface, ensuring the hypersurface is never left. Under these circumstances, the authors prove not only the well-posedness of the particle and mean field system, but also prove the convergence of the particle system towards the mean field limit. Crucial for that is the condition of a **compact** hypersurface. This allows for a straightforward comparison of N copies of the mean field process with the N-particle process, or as the authors put it: "by the classical Sznitman's theory". Moreover, they reason that the weak convergence of the empirical measure towards the mean field measure has rate N^{-1}. In the accompanying pre-print [FHPS21b], the same authors show further results on hyperspheres. On the sphere, a minimization property is proven, similarly to those in the other works mentioned already. Together with the mean-field limit result, which includes hyperspheres in its scope, this gives a complete characterization of a CBO scheme. Their main result characterizes the convergence of the discretized-in-time particle system

as

$$\text{Total error} = \text{discret. error} + \text{mean field approx.} + \text{Laplace principle},$$

a combination of three errors. The total error is measured as the expected squared distance of the mean of the discretized N-particle system at time T with timestep Δt described by the term

$$\mathbb{E}\left[\left|\frac{1}{N}\sum_{i=1}^{N} X_{\Delta t,T}^{i,(N)} - x^*\right|^2\right].$$

The discretization error, in case of an Euler-Maruyama scheme, usually has weak convergence in Δt of order 1. The mean field result for the general hypersurfaces tranfers to this case, as does the convergence of rate N^{-1}. Lastly, similar to the minimization result for the standard CBO, exponential convergence of the mean field process towards the minimizer in time has been shown. This holds, like before, for a suitable, provably non-empty set of parameters λ, σ, α in relation to the dimension d. In [FHPS21a], the same authors deploy the idea of the dimension-wise diffusion — they call it anisotropic diffusion — to the CBO on the sphere. This works out, and in the same way as for the optimisation on $\mathbb{R}^d$, the dimension disappears from the convergence rate of the mean field process, i.e. the error component "Laplace principle" is free from the curse of dimensionality.

Use of Particle History

In [TW20] a variant is proposed wherein a personal best per particle trajectory is computed. This is in line with the Particle Swarm Approach, where a personal best is stored in all common variants. Similarly motivated as for the inter-particle consensus, not the argmax is used but rather a weighted mean over the trajectory inspired by the Laplace principle. That is

$$p^i(x) = \frac{1}{\int_0^t \omega_f^\beta(x(t))\,\mathrm{d}t}\int_0^t x(t)\omega_f^\beta(x(t))\,\mathrm{d}t$$

with the corresponding dynamics for $i = 1, \dots, N$

$$\mathrm{d}X_t^i = \left(-\lambda\left(X_t^i - m_t\right) - \gamma\left(X_t^i - p_t^i\right)\right)\mathrm{d}t + \sigma\left|X_t^i - m_t\right|\mathrm{d}W_t^i\,,$$

where β and γ are newly introduced parameters who function analogously to α and λ. For this now non-Markovian particle system a well-posedness result has been given alongside some interesting numerical results.

Chapter 3

CBO with Finite Range Interaction

We are proposing the following scheme as a variation of the standard particle scheme (2.1). The SDE system reads

$$\mathrm{d}X_t^i = -\lambda\left(X_t^i - m_t^i\right)\mathrm{d}t + \sigma\left(\left|X_t^i - m_t^i\right|\right)\mathrm{d}W_t^i\,, \tag{3.1a}$$

$$m_t^i = \frac{1}{\varepsilon + \sum_{j=1}^N \phi(X_t^i - X_t^j)\omega_f^\alpha(X_t^i)} \sum_{j=1}^N X_t^j \omega_f^\alpha(X_t^i)\phi(X_t^i - X_t^j)\,, \tag{3.1b}$$

with $\lambda, \alpha > 0$ the already known parameters for the drift and the scaling of the landscape. Further, there are two noteworthy differences from the standard model introduced in Chapter 2. Firstly, the application of the diffusion function σ with the following properties:

Assumptions A6.

$$\begin{aligned} &\left|\sigma(x) - \sigma(y)\right| \leq L_\sigma \left|x - y\right| \quad \text{for } L_\sigma > 0 \\ &\sigma(0) = 0 \end{aligned}$$

The Lipschitz continuity of σ is unsurprisingly a technical necessity for the well-posedness of the system. The second condition is not imperative for the analysis carried out in this work, but certainly sensible as one wants a steady state whenever $X_t^i = m_t^i$ for all $i = 1, \ldots, N$, i.e. whenever the particles have reached a consensus. It is thus on the one hand useful as a mental model and on the other hand streamlines the notation, enough reason to make this assumption already at this point. Replacing

the parameter (or linear function) by this function σ clearly yields a more general optimization model. Further, this set-up allows for an upper bound on the diffusion, a property that is necessary for proving the large deviation results in Chapter 4.

It is mostly convenient to imagine the function as $\sigma(x) = \max\{\sigma_c x, K_\sigma\}$ with $K_\sigma > 0$, a capped off linear function. Alternatively, the definition also allows for a logistic function or similar.

The presence of the $\varepsilon > 0$ ensures that the computation of the center of mass is also sensibly defined in the case where the sum (or even more importantly the integral after taking the limit) in the denominator is zero. This small disturbance purely facilitates an easier analysis of the system and is assumed to be very small, even negligible in the context of the application of the optimization method. The full significance of that small disturbance is discussed at a later point.

Secondly, ϕ is the newly introduced communication kernel. The only way in which the N different particles influence each other is via the effect they have on the common attraction point. With the communication kernel in place, each particle now has its own attraction point m_t^i. Not all other particles have to have the same uniform influence on this individual attraction point, in contrast to the unmodified and thus uniform communication.

Introducing this kernel is interesting from a modelling perspective. The development of the system can be understood as opinion dynamics, and therefore making the influence the particles have on each other non-uniform is an intuitive enhancement. Particles that are closer by some metric can have a bigger impact on the opinion formation for example. This can be represented by radial functions taking on the form $\tilde{\phi}(x, y) = \phi(\|x - y\|)$. Prominent examples possible in this setup are the Gaussian kernel $\phi(x) = \exp(-x^2)$ or an indicator function $\phi(x) = I_{[0,R]}(x)$. For the kernel $\tilde{\phi}(x, y) \equiv 1$, the system reduces to the standard case, albeit with a disturbance ε.

Additionally, computational aspects are also a driving factor. With real world applications in mind, it can be hugely beneficial if not every particle has to communicate its status to every other particle. It is easy to imagine such sparsity being computationally relevant, especially in cases with actual physically different particles/nodes.

With these applications in mind, we make the followings assumption on the communication kernel:

Assumptions A7.

$$\begin{aligned}&\phi \in C_c^1(B_r(0)),\\ &\phi(x) \geq 0, \quad \text{for all } x \in \mathbb{R}^d \text{with}\\ &\tilde{\phi}(x,y) = \phi(x-y)\,.\end{aligned}$$

This assumption clearly limits the area of influence of each particle to its surrounding ball with radius r, the interaction happens only within a finite range. One example of a kernel fitting this description would be

$$\tilde{\phi}(x) = \begin{cases} \exp(-1/1-x) & \text{for } x < 1 \\ 0 & \text{otherwise} \end{cases}$$

here with $r = 1$. While that assumption prohibits the reversion to the standard case for now, in Chapter 4 about large deviations and limiting processes the standard CBO is briefly discussed again. For the following results, the Assumptions A1, A6, A7 are assumed to hold as well as continuous differentiability of the target function f.

3.1 Well-Posedness of the Particle System

The kernel CBO particle system is written in vector notation as

$$\mathrm{d}\mathbf{X}_t^{(N)} = -\lambda \mathbf{F}_N(\mathbf{X}_t^{(N)})\,\mathrm{d}t + \mathbf{M}_N(\mathbf{X}_t^{(N)})\,\mathrm{d}\mathbf{W}_t^{(N)}\,, \tag{3.2}$$

with $\mathbf{W}$ the standard Brownian motion in $\mathbb{R}^{Nd}$ and

$$\mathbf{F}_N^i(\mathbf{X}) = X_i - \frac{\sum_i X_j \phi(X_i - X_j)\omega_f^\alpha(X_j)}{\varepsilon + \sum_i \phi(X_i - X_j)\omega_f^\alpha(X_j)}\,, \tag{3.3}$$

$$\mathbf{M}_N(\mathbf{X}) = \mathrm{diag}(\sigma(\left|F_N^1(\mathbf{X})\right|)\mathbb{I}_d, \ldots, \sigma(\left|F_N^N(\mathbf{X})\right|)\mathbb{I}_d) \quad \in \mathbb{R}^{Nd\times Nd}\,. \tag{3.4}$$

Theorem 3.1 (Well-posedness of the kernel CBO particle system)**.** *The particle system* (2.1) *or equivalently in vector notation* (3.2) *has a strong solution and path-wise uniqueness holds.*

Proof. We aim to show linear growth and local Lipschitz continuity of the drift and diffusion coefficients, in order to then apply a standard result for the well-posedness. Starting with linear growth, we can derive for the drift

$$\begin{aligned}\left|\mathbf{F}_N^i(\mathbf{X})\right| &\leq \left|\frac{\varepsilon X_i + \sum_i (X_i - X_j)\phi(X_i - X_j)\omega_f^\alpha(X_j)}{\varepsilon + \sum_i \phi(X_i - X_j)\omega_f^\alpha(X_j)}\right| \\ &\leq |X_i| + \left|\frac{\sum_i (X_i - X_j)\phi(X_i - X_j)\omega_f^\alpha(X_j)}{\sum_i \phi(X_i - X_j)\omega_f^\alpha(X_j)}\right| \leq |X_i| + r \end{aligned} \tag{3.5}$$

that a simple linear growth restriction holds for every $i = 1, \dots, N$, as the second term is a convex combination of the particles within the interaction radius. Since the function σ is also a Lipschitz continuous function and finite at $x = 0$, this result easily extends to the diffusion $\mathbf{M}_N$.

Under the assumptions A4, A7 we posed on f and ϕ, both are continuously differentiable. That is enough for F_N^i to also be continuously differentiable. We then deduce from that that F_N^i is locally Lipschitz for each $i = 1, \dots, N$. Again, given the benevolent properties of the function σ, this extends to $\mathbf{M}_N$.

To conclude this section, we refer to a result from [Dur96], Chapter 5, Theorem (3.2). Analogous to [CCTT18], we show to this end, that

$$-2\lambda \mathbf{X} \cdot \mathbf{F}_N(\mathbf{X}) + \operatorname{trace}(\mathbf{M}_N \mathbf{M}_N^\top)(\mathbf{X}) \leq b_N(1 + |\mathbf{X}|^2) \tag{3.6}$$

holds for some $b_N > 0$ which may depend on the number of particles. Similar to the linear growth property (3.5) the following inequalities hold:

$$\begin{aligned} -X_i \cdot \mathbf{F}_N^i(\mathbf{X}) &\leq -|X_i|^2 + |X_i|\, r\,, \\ \sigma\left(\left|\mathbf{F}_N^i(\mathbf{X})\right|\right) &\leq L_\sigma\ (|X_i| + r)\ , \end{aligned}$$

and thus, keeping in mind that $\mathbf{M}_N$ is a diagonal matrix,

$$\begin{aligned} -2\lambda \mathbf{X} \cdot \mathbf{F}_N(\mathbf{X}) + \operatorname{trace}(\mathbf{M}_N \mathbf{M}_N^\top)(\mathbf{X}) &\leq \sum_i 2\lambda(-|X_i|^2 + |X_i|\, r) + d L_\sigma^2\ (|X_i| + r)^2 \\ &\leq 2N\left(\lambda r + d L_\sigma^2\right)|\mathbf{X}|^2 + 2Nr\left(\lambda + d L_\sigma^2 r\right)\ , \end{aligned}$$

which proves the assertion. □

Remark. The cited result from [Dur96] also offers moment bounds for the solutions. However, in this case, they increase with $N \to \infty$ and are not helpful for the (transition to the) mean field system.

3.2 Well-Posedness of the Mean-Field Limit SDE

The formally resulting mean-field SDE for the kernel CBO reads as

$$\mathrm{d}\overline{X}_t = -\lambda\left(\overline{X}_t - m[\overline{\mu}_t](\overline{X}_t)\right)\mathrm{d}t + \sigma\left(\left|\overline{X}_t - m[\overline{\mu}_t](\overline{X}_t)\right|\right)\mathrm{d}W_t\,, \tag{3.7a}$$

$$m[\overline{\mu}_t](\overline{X}_t) = \frac{1}{\varepsilon + \int \phi(\overline{X}_t, y)\omega_f^\alpha(y)\,\overline{\mu}_t(\mathrm{d}y)} \int y\phi(\overline{X}_t, y)\omega_f^\alpha(y)\,\overline{\mu}_t(\mathrm{d}y)\,, \tag{3.7b}$$

$$\overline{\mu}_t = \mathrm{law}(\overline{X}_t)\,, \quad \mathrm{law}(\overline{X}_0) = \overline{\mu}_0\,. \tag{3.7c}$$

This especially includes the fact, that the coefficients in this formulation yield the particle dynamics as well if given the empirical measure as input.

Theorem 3.2 (Well-posedness of the kernel CBO mean-field SDE)**.** *For $\overline{\mu}_0 \in \mathscr{P}_2$ the mean-field system* (3.7) *has a unique continuous solution curve $\overline{\mu}^* \in C([0,T], \mathscr{P}_2(\mathbb{R}^d))$.*

Remark. In contrast to the particle version of the kernel CBO, the methods from [CCTT18] cannot be fully applied to the mean field version. There, the strategy is to define a mapping from one trajectory of the center of mass m_t to the next one and then find a fixed point of this mapping in $C([0,T],\mathbb{R}^d) \to C([0,T],\mathbb{R}^d)$. However, the fact that the center of mass now also depends on the instantiation of the process in addition to the measure describing the distribution makes this strategy infeasible for the kernel CBO model. While a fixed point stratetgy is still applied, the spaces of the mappings are changed significantly.

Remark. After this continuous measure-valued curve is obtained as a solution in Theorem 3.2, this can be fed into the SDE as a fixed input. By a fairly standard argument, also to be seen in Step 1 of the proof, this then also guarantees a unique strong solution to the SDE system.

Proof. This proof consists of three main steps. Firstly, a mapping is defined using an auxiliary SDE system and its well-definedness and continuity is shown. Secondly, the conditions of a fixed point theorem are checked for this mapping. And lastly, it is shown that this fixed point is the unique solution of the kernel CBO mean-field system.

Step 1 (Auxiliary Mapping): We denote by $\mathscr{A} = (C([0,T], \mathscr{P}_2(\mathbb{R}^d) \parallel \eta(0) = \eta_0)$ the space of measure valued curves with finite second moment with an arbitrary initial condition $\eta_0 \in \mathscr{P}_2(\mathbb{R}^d)$. We define the mapping

$$\mathscr{T} : \mathscr{A} \to \mathscr{A} , \qquad \left(\mathscr{T}\eta\right)_t = \mathrm{law}(Z_t) \tag{3.8a}$$

with Z_t the solution to the auxiliary SDE

$$\mathrm{d}Z_t = -\lambda\left(Z_t - m[\eta_t](Z_t)\right)\mathrm{d}t + \sigma\left(\left|Z_t - m[\eta_t](Z_t)\right|\right)\mathrm{d}W_t , \tag{3.8b}$$

with a given measure valued curve $\eta \in \mathscr{A}$. In order to make the mapping well-defined, Eq. (3.8b) needs to have a unique solution Z_t. For that, we show that

$$\tilde{b}(x,t) = b(x,\eta_t) := -\lambda\left(x - m[\eta_t](x)\right) , \tag{3.9}$$

$$\tilde{\sigma}(x,t) = \sigma\left(\left|x - m[\eta_t](x)\right|\right) , \tag{3.10}$$

fulfill the standard conditions from literature (e.g [Øks03], §5.2). That is, a linear growth and Lipschitz condition hold for the coefficients $\tilde{b}$ and $\tilde{\sigma}$. Given the limited area of influence of the communication kernel ϕ, we can conclude

$$\begin{aligned}\left|x - m[\eta_t](x)\right| &= \left|x - \frac{\int_{\mathbb{R}^d} y\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}{\varepsilon + \int_{\mathbb{R}^d}\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}\right| \\ &= \left|\frac{x\varepsilon + \int_{B_r(x)}(x-y)\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}{\varepsilon + \int_{B_r(x)}\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}\right| \leq |x| + r ,\end{aligned} \tag{3.11}$$

immediately showing the linear growth property of $\tilde{b}$ and $\tilde{\sigma}$. At the core of this lies the fact that $m[\eta_t](x)$ is a weighted average over the measure η_t. However, due to ϕ, this

average is only evaluated within a ball centered around x. The additional summand ε in the denominator can at worst completely dominate the normalization, therefore causing the additive x in the final result.

For the local Lipschitz property of the coefficients, we take a look at the derivative of the center of mass m, and see that

$$\begin{aligned}(\partial_j m[\eta_t])(x) &= \frac{\int_{\mathbb{R}^d}(y-x)(\partial_j\phi)(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}{\varepsilon+\int_{\mathbb{R}^d}\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)} \\ &\quad + (x-m[\eta_t](x))\frac{\int_{\mathbb{R}^d}(\partial_j\phi)(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)}{\varepsilon+\int_{\mathbb{R}^d}\phi(x-y)\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y)} \\ &\leq \frac{(2r+|x|)}{\varepsilon}\left\|\nabla\phi\right\|_{L^\infty}\int_{\mathbb{R}^d}\omega_f^\alpha(y)\,\eta_t(\mathrm{d}y) \end{aligned} \tag{3.12}$$

holds for all $\eta_t \in \mathscr{P}(\mathbb{R}^d)$ and $j = 1,\dots,d$. With the target function f bounded from above, this is bounded on a compact ball. We conclude that the derivative of $m[\eta_t](x)$ and thus consequently also those of $\tilde{b}(x,t),\tilde{\sigma}(x,t)$ are locally bounded and also locally Lipschitz continuous.

The mapping $\mathscr{T}$ is therefore well-defined. We are also interested in the continuity of the mapping. We interpret the mapping as a composition of two mappings

$$\mathscr{A} = C([0,T],\mathscr{P}_2(\mathbb{R}^d)) \to L^2((\Omega,\mathbb{P}),C([0,T],\mathbb{R}^d)) \to C([0,T],\mathscr{P}_2(\mathbb{R}^d)) = \mathscr{A}\,, \tag{3.13}$$

and cite [Arn74, Theorem 7.3.1] for the continuity of the first arrow. For that purpose, $\mathscr{A}$ is interpreted as the parameter space of the coefficients of the auxiliary SDE. The theorem requires the linear growth to be independent of the parameter, as proven in (3.11) and the continuity of the coefficients in the parameter space, which we show later in (3.19). Any solution to the SDE has time marginals with finite second moment and the growth of the second moment in time is restricted. That justifies the mapping implied by the second arrow. The continuity of the second arrow is obtained by pulling the supremum into the integral, giving

$$\sup_{t\in[0,T]} W_2^2(\mathrm{law}(X(t)),\mathrm{law}(Y(t))) \leq \sup_{t\in[0,T]}\mathbb{E}\left[|X(t)-Y(t)|^2\right] \leq \mathbb{E}\left[\sup_{t\in[0,T]}|X(t)-Y(t)|^2\right].$$

Step 2 (Fixed Point): With the operator $\mathscr{T} : \mathscr{A} \to \mathscr{A}$ in place, one can look for a fixed point. We cite a fixed point theorem from [MN75]:

Theorem 3.3 (Schauder-Tychonoff)**.** *Let $\mathscr{A}$ be a convex subset of a locally convex space and f a continuous map of $\mathscr{A}$ into a compact subset of $\mathscr{A}$. Then f has a fixed point.*

Step 2a (Locally Convex Structure): For the locally convex space, we equip the space of all finite measures $\mathscr{M}^b(\mathbb{R}^d)$ with the narrow topology.

Definition 3.4 (Convergence in the narrow topology)**.** The measure μ_n is said to converge narrowly to μ in $\mathbb{R}^d$ if and only if

$$\int f \, \mathrm{d}\mu_n \to \int f \, \mathrm{d}\mu$$

for all continuous and bounded functions $f \in C_b^0(\mathbb{R}^d)$.

The space of continuous curves with a certain starting point

$$\mathscr{L} := \left\{ \eta \in C\left([0,T], \mathscr{M}^b(\mathbb{R}^d)\right) \;\middle\|\; \eta(0) = \eta_0 \right\}$$

constitutes a locally convex space, as the family of semi-norms can be transferred from $\mathscr{M}^b(\mathbb{R}^d)$ via the supremum. Then, due to the definition of $\mathscr{P}_2(\mathbb{R}^d)$ given in Definition 2.1, we have that $\mathscr{A}$ is a convex subset of $\mathscr{L}$.

Step 2b (Continuity): The continuity of $\mathscr{T} : C\left([0,T], \mathscr{M}^b(\mathbb{R}^d)\right) \hookrightarrow C\left([0,T], \mathscr{M}^b(\mathbb{R}^d)\right)$ is implied by the Wasserstein-2 continuity as shown in Step 1 (3.13).

Step 2c (Compactness): Next, we take a look at the set $\mathscr{T}(A)$ and prove its embedding into a compact subset. We make use of the fact that for $\eta_0 \in \mathscr{P}_2(\mathbb{R}^d)$ we have a second-order moment estimate for $\left(\mathscr{T}\eta\right)_t$. This is obtained after squaring both sides of the auxiliary equation (3.8b), taking the expectation and using Itô's isometry,

which gives

$$\begin{aligned}\mathbb{E}\,|Z_t|^2 &\le \mathbb{E}\,|Z_0|^2 + 2\lambda^2\mathbb{E}\left|\int_0^t (x - m[\eta_t](x))\,\mathrm{d}t\right|^2 + 2\mathbb{E}\left|\int_0^t \sigma(\left|x - m[\eta_t](x)\right|)\,\mathrm{d}W_t\right|^2 \\ &\le \mathbb{E}\,|Z_0|^2 + 2(\lambda^2 t + L_\sigma^2)\mathbb{E}\left|\int_0^t (x - m[\eta_t](x))^2\,\mathrm{d}t\right|^2,\end{aligned}$$

then subsequently applying the linear estimate from (3.11) and Gronwall's inequality yields

$$\mathbb{E}\,|Z_t|^2 \le \left(\mathbb{E}\,|Z_0|^2 + 4Tr^2(\lambda^2 T + L_\sigma^2)\right)\exp\left(4(\lambda^2 T + L_\sigma^2)\right) := C_2\,.$$

This specifically also provides

$$\sup_{\eta\in A}\sup_{t\in[0,T]}\int |x|^2\,\mathrm{d}(\mathscr{T}\eta)_t \le C_2 \tag{3.14}$$

for some positive constant $C_2 < \infty$.

By the integral characterization of tightness and Prokhorov's theorem, as given in [AGS05, Remark 5.1.5. and Theorem 5.1.3] respectively, the set of measures

$$\left\{(\mathscr{T}\eta)_t \;\middle|\; t\in[0,T], \eta\in A\right\} \quad \text{is relatively compact w.r.t. the narrow topology in } \mathscr{P}_2(\mathbb{R}^d).$$

We can use this result to apply a "refined version of Arzelà-Ascoli" [AGS05, Proposition 3.3.1] in order to also show compactness in the space of measure valued curves $\mathscr{A}$. The second ingredient to this is

$$\begin{aligned}\mathbb{E}\,|Z_t - Z_s|^2 &\le 2\lambda^2\mathbb{E}\left|\int_s^t (Z_\tau - m[\eta_\tau](Z_\tau))\,\mathrm{d}\tau\right|^2 + 2\mathbb{E}\left|\int_s^t \sigma(\left|Z_\tau - m[\eta_\tau](Z_\tau)\right|)\,\mathrm{d}W_\tau\right|^2 \\ &\le 4(r + C_2)(\lambda^2 T + L_\sigma^2)\,|t - s|\,,\end{aligned}$$

which yields some form of (equi-)continuity of the curves in time. We used the linearity of the coefficients and the uniformly bounded second moment of the process

for this. Concretely, this provides the necessary result

$$W_2\left((\mathscr{T}\eta)_t, (\mathscr{T}\eta)_s\right) \leq c_0 |t-s|^{\frac{1}{2}}\ , \tag{3.15}$$

where the constant $c_0 > 0$ is independent of $\eta \in A$. This compactness argument then yields a W_2-continuous limit curve η^*, an accumulation point for a sequence in $\mathscr{T}(A)$. Since the moment operator $M_2 : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}$, $M_2(\mu) = \int |x|^2 \,\mathrm{d}\mu$ is lower semi-continuous w.r.t. the narrow topology(see [AGS05, Section 5.1.1]), we can locate this accumulation point within the set $\mathscr{A}$ again,

$$\int |x|^2 \,\mathrm{d}\eta_t^* \leq \liminf_{n\in\mathbb{N}} \int |x|^2 \,\mathrm{d}\eta_t^n, \qquad \forall\, t \in [0,T]\,. \tag{3.16}$$

With this, we have shown the image of $\mathscr{A}$ under $\mathscr{T}$ to be relatively compact and the accumulation points to be within $\mathscr{A}$ again. Thus, the function $\mathscr{T}$ maps into a compact subset of $\mathscr{A}$ and fulfills the condition from the fixed point theorem, Theorem 3.3.

Step 3 (Uniqueness): For this step, we take a look at the difference process $Z_t - \hat{Z}_t$ for the same Brownian motion, and to this end also introduce a stopping time, at which the process is aborted. This stopping time is defined as

$$\tau_R = \inf\Big\{ t \geq 0 \mid \max\{\|Z_t\|, \|\hat{Z}_t\|\} > R \Big\}, \qquad R > 0,$$

making sure that both processes take place in the ball $B_R(0)$. With the formulation of the difference process given by

$$\begin{aligned} Z_{t\wedge\tau_R} - \hat{Z}_{t\wedge\tau_R} = Z_0 - \hat{Z}_0 &+ \lambda \int_0^{t\wedge\tau_R} (Z_s - \hat{Z}_s)\,\mathrm{d}s + \lambda \int_0^{t\wedge\tau_R} (m[\rho](Z_s) - m[\hat{\rho}](\hat{Z}_s))\,\mathrm{d}s \\ &+ \int_0^{t\wedge\tau_R} \sigma(\big|Z_s - m[\rho](Z_s)\big| - \sigma(\big|\hat{Z}_s - m[\hat{\rho}](\hat{Z}_s)\big|)\,\mathrm{d}W_s\,, \end{aligned}$$

we decide to take a closer look at

$$m[\rho](Z_s) - m[\hat{\rho}](\hat{Z}_s). \tag{3.17}$$

Lemma 3.5. *For $x, \hat{x} \in B_{R(0)} \subset \mathbb{R}^d$ and $\rho, \hat{\rho} \in \mathscr{P}_2(\mathbb{R}^d)$ it holds*

$$\left|m[\rho](x) - m[\hat{\rho}](\hat{x})\right| \leq \left\|\nabla m[\rho]\right\|_{L^\infty_{B_{R(0)}}} |x - \hat{x}| + C(a, L_f, \varepsilon, R, r) W_2(\rho, \hat{\rho}) \,, \tag{3.18}$$

that is the difference can be bounded in terms of the location of the centers and the Wasserstein-2 norm of the measures. The proof is given at the end of this section.

Now, after squaring the difference process and taking the expectation on both sides and using Itô's isometry, we obtain

$$\begin{aligned}\mathbb{E}\left|Z_{t\wedge\tau_R} - \hat{Z}_{t\wedge\tau_R}\right|^2 \leq 2\mathbb{E}\left|Z_0 - \hat{Z}_0\right|^2 + (4\lambda^2 t 8 L_\sigma^2) \int_0^t \mathbb{E}\left|Z_s - \hat{Z}_s\right|^2 \mathrm{d}s \\ + (4\lambda^2 t + 8L_\sigma^2) \int_0^t \mathbb{E}\left|m[\rho](x) - m[\hat{\rho}](\hat{x})\right|^2 \mathrm{d}s \,.\end{aligned}$$

Using the fact that $W_2^2(\rho_t, \hat{\rho}_t) \leq \mathbb{E}|Z_t - \hat{Z}_t|^2$, we can apply Gronwall's inequality to

$$\mathbb{E}\left|Z_{t\wedge\tau_R} - \hat{Z}_{t\wedge\tau_R}\right|^2 \leq (4\lambda t + 8L_\sigma^2)\left(1 + \left\|\nabla m[\rho]\right\|^2_{L^\infty_{B_R(0)}} + C(a, L_f, \varepsilon, R, r)\right) \int \mathbb{E}\left|Z_s - \hat{Z}_s\right|^2 \mathrm{d}s$$

to obtain that $\mathbb{E}|Z_{t\wedge\tau_R} - \hat{Z}_{t\wedge\tau_R}|^2 = 0$ for all times $t \geq 0$. To characterise the time τ_R, we observe that the processes have a uniformly bounded second moment. Markov's inequality therefore gives

$$\mathbb{P}(\tau_R < T) \leq \mathbb{P}(\max\{|Z_T|, |\hat{Z}_T|\} > R) \leq \frac{1}{R^2}\mathbb{E}\max\{|Z_T|^2, |\hat{Z}_T|^2\} \leq \frac{C_2}{R^2},$$

where $C_2 > 0$ is the constant appearing in (3.14). Therefore, τ_R converges in measure, which combined with the monotonicity of τ_R in $R > 0$ yields the almost sure convergence $\tau_R \to +\infty$ as $R \to \infty$. Owing to the pointwise convergence and the continuity of $t \mapsto Z_t, \hat{Z}_t$ almost surely, the dominated convergence theorem then yields

$$\mathbb{E}\left|Z_t - \hat{Z}_t\right|^2 = \lim_{R\to\infty} \mathbb{E}\left|Z_{t\wedge\tau_R} - \hat{Z}_{t\wedge\tau_R}\right|^2 = 0 \qquad \text{for every } t \geq 0,$$

i.e. $Z_t = \hat{Z}_t$ almost surely for every $t \geq 0$. This concludes the proof of Theorem 3.2. □

Proof of Lemma 3.5. We begin by separating the term and already revealing the two-fold structure of the estimate

$$\left|m[\rho](x) - m[\hat{\rho}](\hat{x})\right| \leq \left|m[\rho](x) - m[\rho](\hat{x})\right| + \left|m[\rho](\hat{x}) - m[\hat{\rho}](\hat{x})\right| ,$$

tackling those terms separately. Given that everything is happening inside a ball, the first term can be estimated via the (local) bound on the derivative given in Equation (3.12), yielding

$$\left|m[\rho](x) - m[\rho](\hat{x})\right| \leq \left\|\nabla m[\rho]\right\|_{L^\infty_{B_{R(0)}}} |x - \hat{x}| \quad \text{for } x, \hat{x} \in B_R(0).$$

For the estimation of the second term, we take inspiration from the playbook of [CCTT18, Lemma 3.2], thus formulating

$$m[\rho](\hat{x}) - m[\hat{\rho}](\hat{x}) = \iint \left[\frac{y\phi(y-\hat{x})\omega_f^\alpha(y)}{\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\rho)} + \varepsilon} - \frac{\hat{y}\phi(\hat{y}-\hat{x})\omega_f^\alpha(\hat{y})}{\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\hat{\rho})} + \varepsilon} \right] \mathrm{d}\pi(y,\hat{y})$$

$$= \iint \frac{\left(\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\hat{\rho})} + \varepsilon\right)\left(y\phi(y-\hat{x})\omega_f^\alpha(y)\right) - \left(\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\rho)} + \varepsilon\right)\left(\hat{y}\phi(\hat{y}-\hat{x})\omega_f^\alpha(\hat{y})\right)}{\left(\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\hat{\rho})} + \varepsilon\right)\left(\left\|\phi(\cdot,\hat{x})\omega_f^\alpha\right\|_{L^1(\rho)} + \varepsilon\right)} \mathrm{d}\pi(y,\hat{y})$$

where $\pi \in \Pi(\rho,\hat{\rho})$ is an arbitrary coupling between the measures ρ and $\hat{\rho}$. We take a look at the derivative of the function in the numerator

$$y \mapsto y\phi(y-\hat{x})\omega_f^\alpha(y) =: \Sigma(y,\hat{x}) ,$$

and see that for

$$(D_y\Sigma)(y,\hat{x}) = \left[y(D\phi)(y-\hat{x}) - \alpha\, y\phi(y-\hat{x})(Df)(y) + \phi(y-\hat{x})\right]\omega_f^\alpha(y) ,$$

both ϕ and $(D\phi)$ are zero whenever $\left|y-\hat{x}\right| > r$. The whole derivative is compactly supported and by the assumptions also continuous, therefore bounded.

Given that one of the normalization factors is bigger than the other and $\varepsilon > 0$, we

obtain

$$\begin{aligned}\left|m[\rho](\hat{x}) - m[\hat{\rho}](\hat{x})\right| &\leq \iint \frac{L_\Sigma}{\varepsilon^2}\left|y-\hat{y}\right| \mathrm{d}\pi \\ &\leq C(a, L_f, \varepsilon, R, r)\left(\iint \left|y-\hat{y}\right|^2 \mathrm{d}\pi\right)^{\frac{1}{2}},\end{aligned}$$

upon which optimizing over all couplings π gives the desired result:

$$\left|m[\rho](\hat{x}) - m[\hat{\rho}](\hat{x})\right| \leq C(a, L_f, \varepsilon, R, r) W_2(\rho, \hat{\rho})\,. \tag{3.19}$$

□

Chapter 4

Large Deviation Principle for the Mean-Field Limit

In this chapter, no prior knowledge about Large Deviations (LDs) shall be assumed. We give a short and not rigorous introduction inspired by [Tou11], which aims to bring the reader to a level to adequately assess the results presented later on.

4.1 What Are Large Deviations?

Principally, we are interested in the normalized sum of random variables which is denoted as $S_N = \frac{1}{N}\sum_{i=1}^{N} X^i$. In its simplest form, those random variables are real valued. While in our set up the random variables are paths, it is still a very helpful simplification to gain an understanding. Specifically in this work, the sum we are interested in is denoted as

$$\mu_\omega^N := \frac{1}{N}\sum_{i=1}^{N} \delta_{X^{i,N}(\cdot,\omega)}\,, \tag{4.1}$$

where $X^{i,N}(t,\omega)$ are the particles of the N-particle model at time t. This means, the variable $\mu_\omega^N \in \mathscr{P}(C([0,T],\mathbb{R}^d))$ is now a probability measure on the trajectories in $\mathbb{R}^d$, a path measure. This is different from the viewpoint we took in the proof of the well-posedness of the kernel CBO presented in Section 3.2, where we studied trajectories on the space of probability measures, measure-valued curves.

The aim is to characterize the principal behavior of the logarithmized probability density function (PDF) of the sum, i.e. find a formulation such that

$$p_{S_N}(s) \approx e^{-NI(s)} . \tag{4.2}$$

The function $I(s) \geq 0$ is called the rate function. It controls the rate of decay for $N \to \infty$. The only points where the PDF does not tend to zero, are the points where s is such that $I(s) = 0$. Consequently the PDF concentrates around these points. The large deviation formulation also gives quantitative estimates for the occurrence of events that are far away from these concentration points, so called rare events or tail events. In this sense, one can say that LD theory is another take on formalizing the convergence of probability measures.

Formally, a few intricacies are involved, but the ideas remain the same. We give the exact definitions for completeness.

Definition 4.1 (Rate Function)**.** A rate function is a lower semi-continuous function $I : S \to [0, \infty]$. A rate function is called **good**, if its sublevel sets are not only closed due to the semi-continuity, but also compact. In general, S is the target space of the random variables one wants to characterize.

Definition 4.2 (Laplace Principle)**.** The sequence of S-valued random variables $(X^n)_{n\in\mathbb{N}}$ satisfies the Laplace principle with rate function I if for all bounded continuous functions $G \in C_b(S)$

$$\lim_{n\to\infty} -\frac{1}{n} \log \mathbb{E}\left[\exp(-n \cdot G(X^n))\right] = \inf_{s\in S}(I(s) + G(s)) .$$

Definition 4.3 (Large Deviation Principle)**.** The sequence of S-valued random variables $(X^n)_{n\in\mathbb{N}}$ satisfies the large deviation principle with rate function I if for all Borel-measurable sets $B \in \mathscr{B}(S)$

$$-\inf_{s\in B^o} I(s) \leq \liminf_{n\to\infty} \frac{1}{n} \log P(X^n \in B) \leq \limsup_{n\to\infty} \frac{1}{n} \log P(X^n \in B) \leq -\inf_{s\in\overline{B}} I(s) ,$$

where B^o denotes the interior and $\overline{B}$ denotes the closure of the set B.

For good rate functions, the Laplace Principle from Definition 4.2 and the Large Deviation Principle from 4.3 are equivalent.

Establishing an LDP for the variable μ_ω^N means we can identify the limiting process for the increasing number of particles $N \to \infty$ as the root of the rate function.

4.2 The Large Deviation Property for Interacting Particle Systems

In [BDF12] a set of sufficient conditions is devised and subsequently streamlined in [Fis14] to establish a large deviation principle for weakly interacting processes of a certain form. Unfortunately, this does not immediately include the application to the process given by the equations (2.6). We present again the conditions here for the convenience of the reader and then discuss how to modify them for our purposes.
The analysis in the source is based on a controlled version of the process, in our case for the particle version

$$\mathrm{d}\hat{X}_t^{i,N} = -\lambda\left(\hat{X}_t^{i,N} - m_t^{i,N}\right)\mathrm{d}t + \sigma\left(\left|\hat{X}_t^{i,N} - m_t^{i,N}\right|\right)u_t^i\,\mathrm{d}t + \sigma\left(\left|\hat{X}_t^{i,N} - m_t^{i,N}\right|\right)\mathrm{d}W_t^{i,N} \tag{4.3}$$

for all $i = 1,\dots,N$, with a measurable stochastic control function $u_t = (u_t^1, u_t^2, \dots, u_t^N) : [0,T]\times\Omega \to \mathbb{R}^{N\times d}$ such that

$$\mathbb{E}\left[\int_0^T |u_t|^2\,\mathrm{d}t\right] < \infty \tag{4.4}$$

the control has finite energy. In the mean field version the controlled system reads as

$$\mathrm{d}\hat{\overline{X}}_t = -\lambda\left(\hat{\overline{X}}_t - m_t\right)\mathrm{d}t + \sigma\left(\left|\hat{\overline{X}}_t - m_t\right|\right)u_t\,\mathrm{d}t + \sigma\left(\left|\hat{\overline{X}}_t - m_t\right|\right)\mathrm{d}W_t \tag{4.5}$$

The hat is used to mark controlled variables.
With this in place, we cite the proposed set of sufficient conditions:

(H1) For some $\rho_0 \in \mathscr{P}(\mathbb{R}^d)$, $\frac{1}{N}\sum_{i=1}^N \delta_{X^{i,N}} \to \rho_0$ for $N \to \infty$.

(H2) The coefficients b, σ are continuous (on $\mathbb{R}^d \times \mathscr{P}(\mathbb{R}^d)$).

(H3) For all $N \in \mathbb{N}$ the particle system (3.1) has a unique strong solution.

(H4) Weak uniqueness holds for the controlled mean field system (4.5).

(H5) If the controls u^N are such that

$$\sup_{N\in\mathbb{N}}\left(\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^{N}\int_0^T \left|u^{i,N}\right|^2 \mathrm{d}t\right]\right)<\infty\,, \tag{4.6}$$

then the controlled empirical measures $\hat{\mu}^N \in \mathscr{P}(C([0,T],\mathbb{R}^d))$, solutions to the controlled particle system (4.3) under u^N, form a tight family.

In [Fis14] an additional condition is introduced to determine the form of the rate function.

(H6) Given any $\eta \in \mathscr{P}(C([0,T],\mathbb{R}^d))$ weak existence and path-wise uniqueness holds for the auxiliary SDE (3.8b).

Then, according to [Fis14], [BDF12] the following theorem holds.

Theorem 4.4 (Laplace Property for Interacting Particle Systems)**.** *Given (H1) - (H6), the Laplace principle holds for the sequence* $\left(\mu^N\right)_{N\in\mathbb{N}}$, $\mu^N \in \mathscr{P}(C([0,T],\mathbb{R}^d))$ *with the rate function*

$$I(\eta) = R(\eta\|\,\mathrm{law}(Z^\eta))\,, \quad \eta \in \mathscr{P}(C([0,T],\mathbb{R}^d))\,, \tag{4.7}$$

where Z^η *is the solution to the auxiliary SDE*

$$\mathrm{d}Z_t = -\lambda\left(Z_t - m[\eta_t](Z_t)\right)\mathrm{d}t + \sigma\left(\left|Z_t - m[\eta_t](Z_t)\right|\right)\mathrm{d}W_t\,,$$

(see Eq (3.8b) *for more context) given the corresponding measure valued curve* η *and* $R(\cdot\|\cdot)$ *denotes the relative entropy between the measures.*

Remark. If this theorem applies to the kernel CBO we can characterize the limit process as the fixed point of the auxiliary SDE (3.8b), since the relative entropy is zero only there.

Definition 4.5 (Relative Entropy)**.** The relative entropy between two path measures

$\eta, \rho \in \mathscr{P}(C([0,T],\mathbb{R}^d))$ is defined as

$$R(\eta\|\rho)) = \int_{C([0,T],\mathbb{R}^d)} \log\left(\frac{\mathrm{d}\eta}{\mathrm{d}\rho}(\gamma)\right)\mathrm{d}\eta(\gamma)\,.$$

4.3 The Large Deviation Property for Finite-Range CBO

So far, we have introduced the Theorem 4.4 on the underlying space $\mathbb{R}^d \times \mathscr{P}(\mathbb{R}^d)$, however that does not suit the scenario of the kernel CBO. Clearly, it would be non-trivial to even define the coefficients of the particle system especially with regards to the computation of the center of mass m_t, the intuition of which is based on finite expectation.For a way around this, [BDF12], Remark 7.1 states that the proof can also be exercised on the the space $\mathbb{R}^d \times \mathscr{P}_2(\mathbb{R}^d)$ since $\left(W_2, \mathscr{P}_2(\mathbb{R}^d)\right)$, as given in Definition (2.1) is Polish. However, this then causes additional conditions to be introduced. These conditions, labeled (H7) and (H8) are later stated rigorously and dealt with. This then means that the theorem holds for path measures on $C([0,T],\mathbb{R}^d)$ whose time marginals additionally are in $\mathscr{P}_2(\mathbb{R}^d)$.
In the previous chapter we have established measure-valued curves as solutions. With the work we have done there, it is clear that they can be reinserted into the well-defined SDE to then provide a path measure which is additionally guaranteed to have marginals in $\mathscr{P}_2(\mathbb{R}^d)$.

About Condition (H1)

This condition requires the starting distributions to converge towards a probability measure. Since this is up to the applicant of the method, this does not pose a problem. One can however not expect to find the optimum if it lies outside the support of the starting measure.

About Condition (H2)

This condition is satisfied by the result from Lemma 3.5.

About Condition (H3)

Showing the strong existence and uniqueness of solutions to the particle system has already been done in Section 3.1.

About Condition (H4)

Even though it is stated in [BDF12] that "It is typical that such weak uniqueness holds if it holds for the uncontrolled system.", some work is required for this condition. In [Fis14], Appendix C sufficient conditions for (H1) - (H5) are given. Going back, one notices that the requirements and the proof that is provided there for condition (H4) can be replicated here in a slightly adapted form. Similarly to the technique we applied in Step 3: Uniqueness of the long proof of Theorem 3.2, we define a stopping time: the process is stopped when either process or now additionally also the energy introduced by the control variable leaves a certain bounded area. Up until this time, many other properties of the system are bounded as well, such that we can derive path-wise uniqueness in Proposition (C.1). It is then argued in [Fis14, Proposition (C.2)] that uniqueness in law follows from the path-wise uniqueness.

For the first part, Proposition (C.1) in [Fis14], we plug Lemma 3.5, which holds until the stopping time. The usefulness of that lemma lies in the fact that the estimate in terms of the position of the process x and the Wasserstein-2-metric of the measure also give us an estimate in terms of the path of the process φ and the Wasserstein-1-metric:

$$\left|b(t,\varphi,\eta)-b(t,\tilde{\varphi},\tilde{\eta})\right|+\left|\sigma(t,\varphi,\eta)-\sigma(t,\tilde{\varphi},\tilde{\eta})\right|\leq L_G\left(\sup_{s\in[0,T]}\left|\varphi(s)-\tilde{\varphi}(s)\right|+W_1(\eta,\tilde{\eta})\right) \tag{4.8}$$

holds for the kernel CBO coefficients with $\varphi,\tilde{\varphi}\in C([0,T],\mathbb{R}^d)$ denoting the trajectories. The stopping time guarantees that $\sup_{s\in[0,T]}\left|\varphi(s)\right|\vee\left|\tilde{\varphi}(s)\right|\leq G$ for a $G>0$. This is an adapted version of (L) in the proof by Fischer. Further, due to the definition of the Wasserstein-metric,

$$W_1(\mathrm{law}(X),\mathrm{law}(Y))\leq\mathbb{E}[|X-Y|]\,, \tag{4.9}$$

holds. This is then used in the replication of Proposition (C.1) briefly presented in its adapted form in the following. Let $\hat{X}, \hat{Y}$ be solutions to (4.5) with control u. Define the stopping time

$$\tau_G := \inf_{t\in[0,T]} \left\{ \left|\hat{X}(t)\right| \vee \left|\hat{Y}(t)\right| \vee \int_0^t |u(s)|^2 \,\mathrm{d}s \geq G \right\}. \tag{4.10}$$

Using the Hölder inequaltiy, Doob's inequality, inequalities (4.8) and (4.9) we derive

$$\begin{aligned}
&\mathbb{E}\left[\sup_{s\in[0,t]} \left|\hat{X}(t\wedge\tau_G) - \hat{Y}(t\wedge\tau_G)\right|^2 \right] \\
&\leq 8L_G^2(T+G+4)\mathbb{E}\left[\int_0^{\tau_G\wedge t} \sup_{r\in[0,s]} \left|\hat{X}(r)-\hat{Y}(r)\right|^2 + W_1(\mathrm{law}(\hat{X}(r)),\mathrm{law}(\hat{Y}(r))^2 \,\mathrm{d}s \right] \\
&\leq 16L_G^2(T+G+4)\int_0^t \mathbb{E}\left[\sup_{r\in[0,s]} \left|\hat{X}(r\wedge\tau_G) - \hat{Y}(r\wedge\tau_G)\right|^2 \mathrm{d}s \right]
\end{aligned}$$

for $t \in [0,T]$ and $G \in \mathbb{N}$. Applying Grönwall's lemma yields $P(\hat{X}(t) = \hat{Y}(t)) = 1$ for $t \leq \tau_G$. Since $\tau_G \to \infty$ for $G \to \infty$, the assertion of weak uniqueness holds for all finite times.

About Condition (H5)

To show tightness of the path measures of the controlled particle system, we again refer to the path taken in [Fis14]. Under a linear growth condition on b and a boundedness condition on σ, namely

$$\left|b(t,\varphi,\eta)\right| \leq K_G\left(1+\sup_{s\in[0,t]}\left|\varphi(s)\right|\right), \quad \left|\sigma(t,\varphi,\eta)\right| \leq K_G, \tag{4.11}$$

(H5) is shown to hold. The latter holds for every bounded function σ specified within the scope given in Chapter 3. The former is a version of Equation (3.5).

About Condition (H6)

In Section 3.2 about the well-posedness of the mean field limit, strong existence and uniqueness has been shown for the auxiliary SDE (3.8b) given an input with time marginal measures with finite second moment, i.e. $\int |x|^2 \, \mathrm{d}\eta_t < \infty$ for all $t >= 0$.

About Conditions (H7) and (H8)

The additional conditions to lift the proof to the space $\mathscr{P}_2(\mathbb{R}^d)$ are as follows:

(H7) The particle system has continuous paths in the new probability space, i.e. the (empirical measure) solution of the kernel CBO particle system (3.1) $\mu^N \in C([0,T], \mathscr{P}_2)$.

(H8) The time marginals of the controlled particle system (4.3) are in the new probability space, i.e. $\hat{\mu}^N(t) \in \mathscr{P}_2$.

Clearly, (H7) is proven in Section 3.1. To show condition (H8), we make use of "ol' reliable", the Itô isometry after squaring and taking the expectation of the controlled particle system (4.3), obtaining

$$\mathbb{E}\left|\hat{X}_t^{i,N}\right|^2 \leq 2\mathbb{E}\left|\hat{X}_0^{i,N}\right|^2 + 2K_\sigma T\left(1 + \int_0^t \mathbb{E}\left|u_s^{i,N}\right|^2 \mathrm{d}s\right) + 2\lambda^2 \int_0^t \mathbb{E}\left|\hat{X}_s^{i,N} - m_s^{i,N}\right|^2 \mathrm{d}s\,. \quad (4.12)$$

Using the linear growth from 3.5 and the fact that the energy of the control variable is finite, then subsequently summing this over $i = 1, \ldots, N$ and averaging prepares the ground for Grönwall's lemma. This gives the final estimate

$$\sup_{t\in[0,T]} \mathbb{E}\int |x|^2 \, \mathrm{d}\mu_t^N \leq K_q(N)\,, \quad (4.13)$$

and thus the assertion of condition (H8). Naturally, the moment bound again depends on the number of particles N. Under the additional assumptions of a bounded diffusion coefficient, the applicability of Theorem 4.4 to the kernel CBO system (3.1) has been shown. The root of the rate function is exactly described by the postulated mean field limit (3.7). One can thus conclude this to be indeed the limiting process, which in case of the standard CBO method fills a gap in the literature.

Remark (Large Deviation Property of the standard CBO). The same methods that were employed in this section also work for the standard CBO method if the diffusion coefficient is bounded. Such a bound does neither influence critically the practical applications of the standard CBO method nor the analysis carried out in [CCTT18]. The same is true for the CBO method with the dimension-wise diffusion.

Remark (The Severeness of the Boundedness Assumption on the Diffusion). Looking at the literature it is noteworthy, that the introduction of interaction in the diffusion term increases the difficulty of transition to a mean field significantly. With regards to large deviation results for example, the classic paper [DG87] does not allow interactions in this term at all when deriving its large deviation result. In the light of this, bounding the diffusion coefficient σ seems like an expected concession.

From Laplace to Large Deviations

In [Fis14] there is also an additional condition 4.3 that guarantees the goodness of the rate function, and by that the equivalence of the Laplace principle result and the Large Deviation Property. Since that is not of paramount interest in this context, only a short sketch of a possible proof of that condition is given in the following. The condition

(H') Given a sequence of controls $(u_t^k)_{k\in\mathbb{N}}$ with finite energy $\sup_{k\in\mathbb{N}} \mathbb{E}\left[\int_0^t u_t^k \,\mathrm{d}t\right] \leq \infty$, the law of the solutions to the controlled mean field equation under u^k form a tight family, i.e. $(\mathrm{law}(\hat{\overline{X}}_t))_{k\in\mathbb{N}}$ are tight in the path measure space $\mathscr{P}(C([0,T],\mathbb{R}^d))$.

is analogous to condition (H5) and as such, the same argument based on linear growth of the drift and boundedness of the diffusion applies as well. The large deviation property implies convergence in probability, but in this form cannot be used to determine a rate of convergence in terms of the number of particles N.

Chapter 5

Numerical Experiments

In this chapter, numerical experiments are performed regarding the CBO method with a special interest in differences between finite and full range interaction. While this does include the settings covered by the analysis in the Chapters 3 and 4, it is not limited to them.

5.1 General Remarks about the Implementation

We use, if not stated otherwise, the standard Euler-Maruyama scheme for the time discretization of the stochastic processes. While the strong convergence of the scheme is not necessarily guaranteed for every imaginable interaction kernel, it empirically proves to be good enough for the simple toy examples used in this section. Speaking on a very general level, Lipschitz property of the kernel translates to Lipschitz property of the system coefficients and finite range of the kernel implies linear growth, which usually makes the coefficients relatively well-behaved in this context. Nonetheless, this topic has many intricacies and is something that should be considered when simulating SDEs. There are many schemes that provably cover a bigger class of SDEs than the standard Euler-Maruyama scheme such as the tamed Euler-Maruyama scheme ([HJK12]), the truncated Euler-Maruyama scheme ([Mao15]) as simple-to-implement explicit relatives or alternatively one of many implicit schemes. A good introductory overview on the topic can be found in [KP13].

One further thing worth indicating is that for the computation of the center of

mass, the numerator and denominator are both multiplied with a factor $\exp(\alpha f(x_t^*))$ for $x_t^* = \arg\min_{\{X_t^i\}_{i=1}^N} f(X_t^i)$, i.e. the exponential landscape is readjusted to the level of the smallest function evaluation. This is done to avoid division by very small values.

Algorithm 1 Kernel CBO

Input: $\sigma, \lambda, \alpha, T, \Delta t, \varepsilon, @f, @\phi$
$t \leftarrow 0$
Sample starting config $\mathbf{X}_0 = (X_0^1, \dots, X_0^N)$
while $t < T$ **do**
  Generate N i.i.d. random normal vectors $\Delta W^i \sim \mathcal{N}(0, \sqrt{\Delta t}), \quad \Delta W^i \in \mathbb{R}^d$
  $\texttt{fval} \leftarrow f(\mathbf{X}_t), \quad \texttt{fval} \in \mathbb{R}^N$
  $\texttt{fmin} \leftarrow \min\{\texttt{fval}\}$
  $K \leftarrow KernelDistanceMatrix(\mathbf{X}_t, @\phi), \quad K \in \mathbb{R}^{N \times N}$
  $\mathbf{m}_t \leftarrow WeightedMean(\mathbf{X}_t, K, \texttt{fval}, \texttt{fmin}), \quad \mathbf{m}_t = (m_t^1, \dots, m_t^N), \quad m_t^i \in \mathbb{R}^d$
  $X_{t+\Delta t}^i \leftarrow X_t^i - \lambda(X_t^i - m_t^i)\Delta t + \sigma(|X_t^i - m_t^i|)\Delta W^i$
  $t \leftarrow t + \Delta t$
end while

5.1.1 Possible Modifications

In addition to the ideas discussed in Chapter 2, some ideas of a more computational nature have been introduced in the literature. In [CJLZ21] a batch-wise updating is proposed. In each step, the center of mass is calculated based on a random subset of particles. For the subsequent update of the particle positions either is possible, only updating the subset of particles or updating all particles. This mini-batch method is quite common in other optimization methods, especially in field of machine learning, e.g. [LZCS14]. Since the random subsets already introduce a form of randomness into the algorithm, the diffusion term can sometimes be omitted. The numerical studies with this mini-batch rule show it to be a generally sensible idea.

Another idea, proposed in [FHPS21b], is to reduce the number of particles during the runtime. In each step, the number of particles is reduced according to the reduction of the empirical variance. The intuitive justification for this is that if the variance is lower, fewer particles are needed for the same degree of certainty. Due to further practical considerations a minimal number of particles should be guaranteed

and the reduction might only be exercised every few iterations.

As a final mention, the second proposal in [FHPS21b] concerns the adaption of the parameters during the runtime, a well-known topic in the area of optimization. The essence of most aforementioned convergence guarantees is, that the drift (towards the consensus point) dominates the randomness. Most rigorous conditions on the parameters mentioned in Chapter 2 look very much alike, and imply something like $\lambda \gg d\sigma^2$ or respectively $\lambda \gg \sigma^2$ without the dimension component in the anisotropic variants. To ensure that this is met, but at the same time allow sufficient exploration, it is a natural idea to reduce σ during the runtime of the algorithm. There is no clear way to say how to do that, not even whether it should be pre-set or dependant on the runtime behavior of the system, but the authors e.g. present good numerical results using $\sigma_{n+1} = \sigma_n/(\sigma_0 \log(n+1))$, which is inspired by the cool-off in simulated annealing [GKR94]

All these ideas do not appear out of thin air, but are inspired by similar concepts in heuristic optimization algorithms and stochastic simulations. And while the analytic tractability of the modifications remains to be determined, there are certainly much more in store for the practical use of CBO algorithms. For example [AP13] lists many possibly transferable ideas for the simulation of particle swarms.

5.2 Kernel CBO

This section performs test on the kernel CBO method introduced in Chapter 3. This means different types of kernels are compared exemplarily.

5.2.1 The Epsilon-Effect

In contrast to the standard CBO method, we have introduced the parameter ε for the computation of the weighted center of mass

$$m_t^i = \frac{1}{\varepsilon + \sum_{j=1}^N \phi(X_t^i - X_t^j)\omega_f^\alpha(X_t^i)} \sum_{j=1}^N X_t^j \omega_f^\alpha(X_t^i)\phi(X_t^i - X_t^j)\,.$$

The new parameter disturbs the normalization factor, resulting in a scaling of the desired/standard center of mass by a factor

$$\frac{1}{1+\frac{\varepsilon}{\text{Norm}}} < 1 \quad \text{for} \quad \varepsilon > 0\,.$$

This down-scaling and thus shift towards the origin can also be seen in numerical tests. Additionally, this also means that a consensus state is not an equilibrium state anymore. These properties of the algorithm are truly undesirable. The main motivation for introducing the parameter ε in the first place however was to prevent the denominator from becoming zero in the mean-field system, a problem we need not worry about in the particle system. Therefore, in the following $\varepsilon = 0$ is assumed for all practical purposes.

5.2.2 Finite Range vs. Infinite Range

Using as the kernel the indicator function $\phi(x) = I_{[0,4]}$ immediately reveals one new phenomenon. Since the range of the communication is limited, it is now possible for particles to lose contact with all other particles. In that case, the design of the system is such that they remain stationary. Per the random nature of the algorithm, the loss of contact can happen despite the overall contractive behavior. This is visually underlined by Figure 5.1, where the development of the standard and a finite-range CBO are compared at the start time, an intermediate time and the end time. In both cases the overall behavior is contractive and the minimum at $x^* = (2,2)$ is found by most particles, however not by all particles in the finite-range case.

Ascribing this effect to the finite range is supported by the fact that the same behavior is seen for other full-range kernels such as $\phi(x) = \exp(-x)$, or $\phi(x) = \exp(-x^2)$ and finite-range kernels respectively (e.g. $\phi(x) = \exp(-1/1-r\,x)$ for $x < 1/r$, $\phi(x) = 0$ otherwise). While that seems undesirable at first and perhaps is in many cases, it opens the door towards the possibility of finding multiple (global) minima at once.

As a very simple toy model, we look at a one-dimensional cosine function as the target $f(x) = \cos(x)$. It is clear, that one would a-priori choose the range of the interaction such that it ceases once the minima are found. In the given case that means the range should be smaller than 2π. In Figure 5.2 one can see the final state of a run

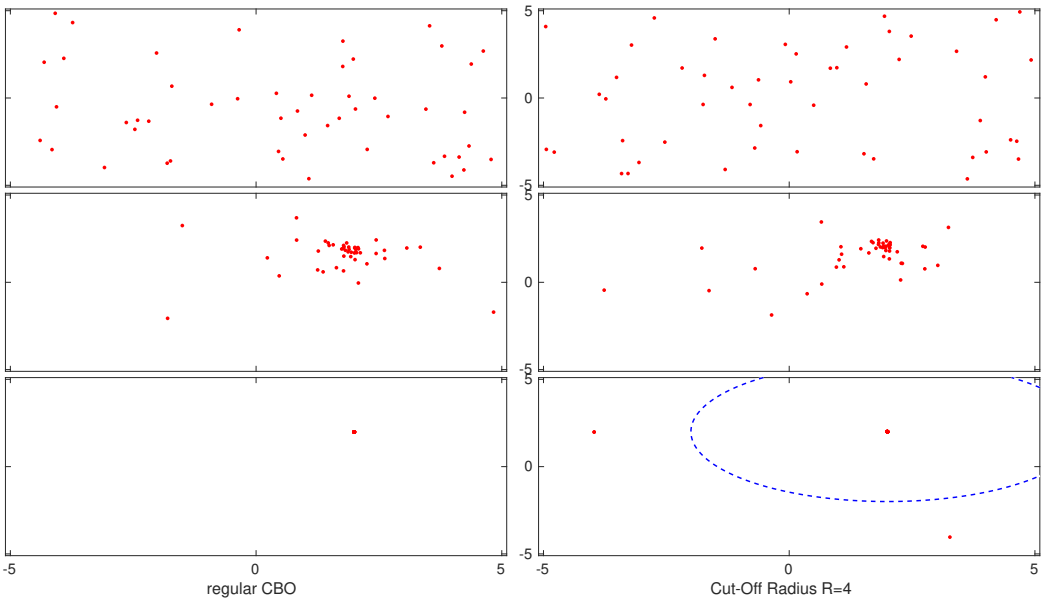

Figure 5.1: Lost particles with finite-range communication

with $N = 20$ particles, represented by the bars, with again $\phi(x) = I_{[0,4]}$ as the communication kernel.

The particles to the left and right interact with the particles in the middle, but not with each other. The influence of the particles in the middle pulls each cluster of particles slightly off the respective minimum. So to be sure to avoid such a situation, one has to choose the interaction range smaller than half the distance between the minima. If that phenomenon is avoided, Figure 5.3 shows what there is to gain by using a finite range interaction. The histograms show the behavior of $N = 500$ particles on the cosine function with minima at $x^*_{1/2} = \pm\pi$. Starting out with a uniform distribution on the interval $[-5,5]$, the standard CBO particles first contract and appear to collect around zero, which marks in fact the maximum of the function. Only after a while do the particles sway to one side. Chance determines which side that is, since the symmetry of the set up is only resolved by the randomness inherent to each single run. Eventually, they go towards the minimum on that side. Unfortunately at $t = 5$ almost all particles are already in the interval $[-\pi, 0]$ and their approach towards the minimum becomes very one-sided. This causes them to find their consensus somewhat off the minimum at around $x = -2.82$.

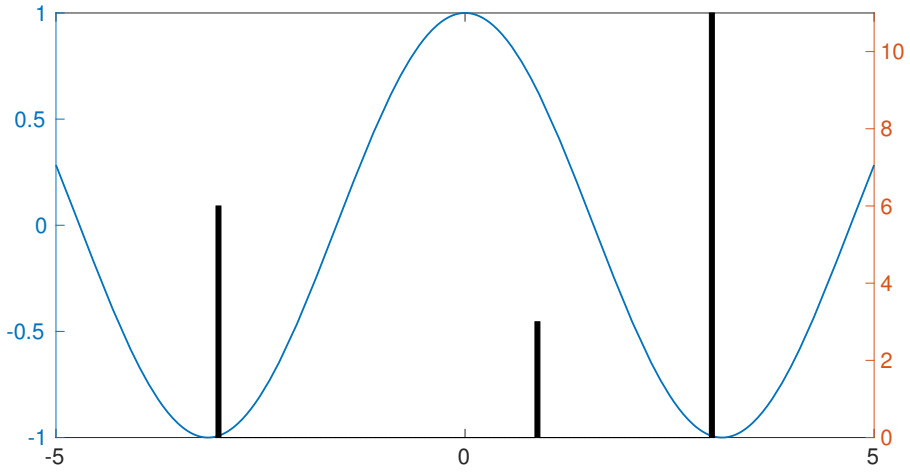

Figure 5.2: An undesirable final state

For the finite range CBO with $\phi(x) = I_{[0,3]}$, the particle behavior is fundamentally different. From the get go, the uniform distribution splits up into two halves, which then each converge towards their respective minimum. This minimum is found earlier compared to the full range CBO and hit more accurately, as neither minimum is approached in the one-sided manner observed in the full range CBO.

This behavior that we described using only one instance of each CBO method happens not only in special cases but rather regularly, underlined by Figure 5.4, where the average mean particle distance from either minimum over several samples of the double well scenario is displayed in time. The finite range is significantly more accurate and fast, even in this very simple example.

The reason for that is, that the difficulty the cosine function poses for the standard CBO method is of a very fundamental nature. The standard CBO, per design, does not allow a consensus state where two minima are marked. But it goes deeper: Even on the mean-field level, since everything is symmetric, the attraction point is always halfway between the two minima and there is no chance either of them is found.

This inherent shortcoming of the standard CBO method manifests numerically and analytically, but can at least for the latter relatively easily be resolved by introducing a finite range interaction. This is by itself interesting enough to at least justify

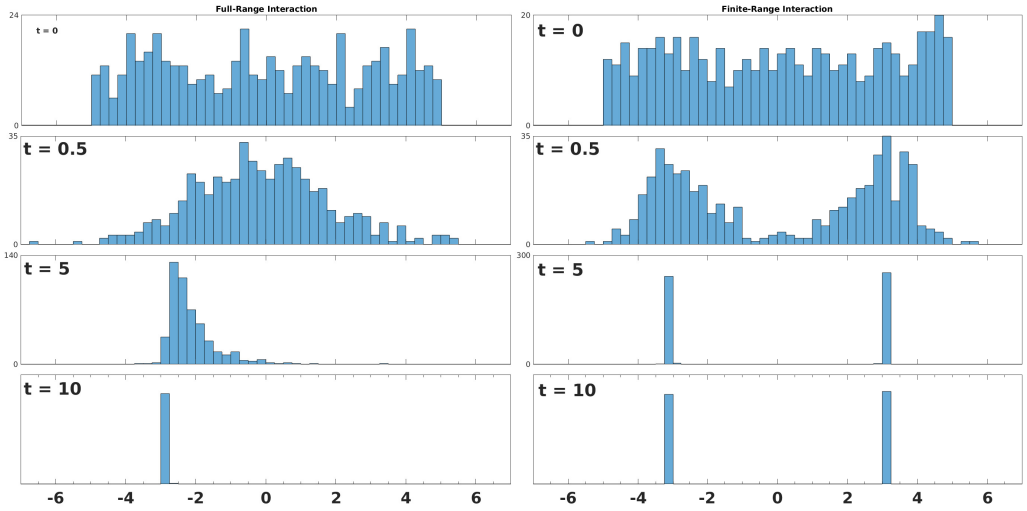

Figure 5.3: Different developments in the double well scenario

opening the door to the analysis of the finite range interaction CBO, as was done with the well-posedness proof in Chapter 3. Obviously, a concentration property cannot expected be proven along the lines of the existing literature, as a singular concentration point is not the expected outcome. The fundamental difference in behavior after a minor change in the formulation might be a good entrance point for gaining future insights about CBO methods as a whole.

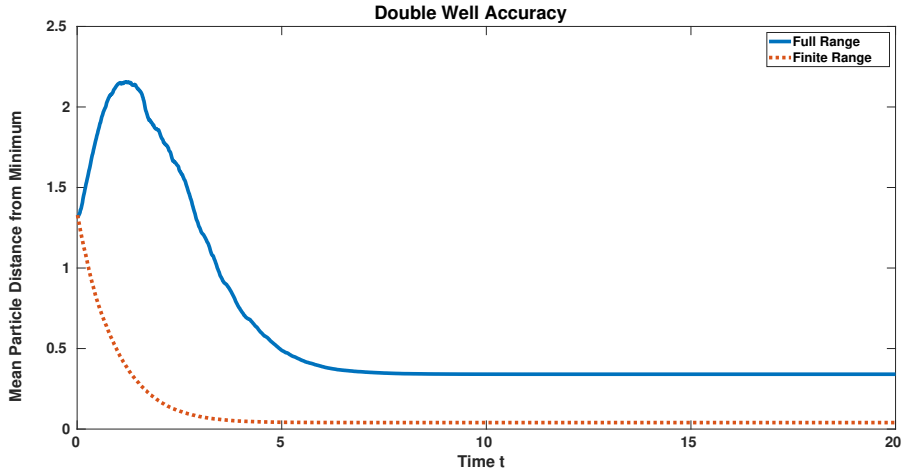

Figure 5.4: Accuracy in the double well scenario

Bibliography

[AGS05] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Birkhäuser Verlag, 2005.

[AP13] G. Albi and L. Pareschi. Binary interaction algorithms for the simulation of flocking and swarming dynamics. *SIAM Journal on Multiscale Modeling and Simulation*, 11:1–29, 03 2013.

[Arn74] L. Arnold. *Stochastic differential equations*. A Wiley-Interscience publication. Wiley, 1974.

[BDF12] A. Budhiraja, P. Dupuis, and M. Fischer. Large deviation properties of weakly interacting processes via weak convergence methods. *The Annals of Probability*, 40(1):74–102, Jan 2012.

[BFOZ19] P. Buttà, F. Flandoli, M. Ottobre, and B. Zegarlinski. A non-linear kinetic model of self-propelled particles with multiple equilibria. *Kinetic & Related Models*, 12(4):791–827, 2019.

[BM17] M. R. Bonyadi and Z. Michalewicz. Particle swarm optimization for single objective continuous space problems: A review. *Evolutionary Computation*, 25(1):1–54, March 2017.

[CCTT18] J. A. Carrillo, Y.-P. Choi, Claudia T., and O. Tse. An analytical framework for consensus-based global optimization method. *Mathematical Models and Methods in Applied Sciences*, 28(06):1037–1066, 2018.

[CHS03] O. Cordon, F. Herrera, and T. Stützle. A review on the ant colony optimization metaheuristic: Basis, models and new trends. *Mathware & soft computing*, 9:141–175, 10 2003.

[CJLZ21] 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.

[COO+18] P. Chaudhari, A. Oberman, S. Osher, S. Soatto, and G. Carlier. Deep relaxation: partial differential equations for optimizing deep neural networks. *Research in the Mathematical Sciences*, 5(30), 2018.

[DG87] D. A. Dawson and J. Gärtner. Large deviations from the mckean-vlasov limit for weakly interacting diffusions. *Stochastics*, 20(4):247–308, 1987.

[Dur96] R. Durrett. *Stochastic Calculus: A Practical Introduction*. Probability and Stochastics Series. Taylor & Francis, 1996.

[FHPS20] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based optimization on hypersurfaces: Well-posedness and mean-field limit, 2020. arXiv:2001.11994.

[FHPS21a] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Anisotropic diffusion in consensus-based optimization on the sphere, 2021. arXiv:2104.00420.

[FHPS21b] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based optimization on the sphere: Convergence to global minimizers and machine learning, 2021. arXiv:2001.11988.

[Fis14] M. Fischer. On the form of the large deviation rate function for the empirical measures of weakly interacting systems. *Bernoulli*, 20(4):1765–1801, Nov 2014.

[FKR21] M. Fornasier, T. Klock, and K. Riedl. Consensus-based optimization methods converge globally in mean-field law, 2021. arXiv:2103.15130.

[GKR94] V. Granville, M. Krivanek, and J. P. Rasson. Simulated annealing: a proof of convergence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):652–656, June 1994.

[GT01] D. Gilbarg and N.S. Trudinger. *Elliptic Partial Differential Equations of Second Order*. Classics in Mathematics. Springer, 2001.

[HJK12] M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong convergence of an explicit numerical method for sdes with nonglobally lipschitz continuous coefficients. *The Annals of Applied Probability*, 22(4):1611–1641, 2012.

[HJK20] 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, 07 2020.

[KP13] P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Stochastic Modelling and Applied Probability. Springer Berlin Heidelberg, 2013.

[KPP19] N. Kantas, P. Parpas, and G. A. Pavliotis. The sharp, the flat and the shallow: Can weakly interacting agents learn to escape bad minima?, 2019. arXiv:1905.04121.

[KS98] I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*. Graduate Texts in Mathematics. Springer-Verlag New York, 1998.

[Lig99] T. M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Grundlehren der mathematischen Wissenschaften. Springer-Verlag Berlin Heidelberg, 1999.

[LZCS14] M. Li, T. Zhang, Y. Chen, and A. Smola. Efficient mini-batch training for stochastic optimization. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 08 2014.

[Mao15] X. Mao. The truncated euler–maruyama method for stochastic differential equations. *Journal of Computational and Applied Mathematics*, 290:370–384, 2015.

[McK66] H. P. McKean. A class of markov processes associated with nonlinear parabolic equations. *Proceedings of the National Academy of Sciences*, 56(6):1907–1911, 1966.

[MN75] S. A. Morris and E. S. Noussair. The Schauder-Tychonoff fixed point theorem and applications. *Matematický časopis*, 25:165–172, 1975.

[MT11] S. Motsch and E. Tadmor. The truncated euler–maruyama method for stochastic differential equations. *Journal of Statistical Physics*, 144(923), 2011.

[Øks03] B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Hochschultext / Universitext. Springer, 2003.

[PTTM17] R. Pinnau, C. Totzeck, O. Tse, and S. Martin. A consensus-based model for global optimization and its mean-field limit. *Mathematical Models & Methods in Applied Sciences*, 27(01):183–204, 2017.

[Sim86] J. Simon. Compact sets in the space lp(o,t; b). *Annali di Matematica Pura ed Applicata*, 146:65–96, 01 1986.

[Tot21] C. Totzeck. Trends in consensus-based optimization, 2021. arXiv:2104.01383.

[Tou11] H. Touchette. A basic introduction to large deviations: Theory, applications, simulations, 2011. arXiv:1106.4146.

[TW20] C. Totzeck and M.-T. Wolfram. Consensus-based global optimization with personal best. *Mathematical Biosciences and Engineering*, 17:6026–6044, 2020.

[YPM+13] B. Yuce, M. Packianather, E. Mastrocinque, D. Pham, and A. Lambiase. Honey bees inspired optimization method: The bees algorithm. *Insects*, 4:646–662, 11 2013.

Appendix A

The Friends we made along the Way

Naturally, not every idea is good enough for a whole thesis to revolve around. Nevertheless, two of the very early ideas deserve to at least be briefly presented. This enterprise happens here in the appendix.

A.0.1 CBO on a Random Graph

If the particles are imagined to be actual physical agents, one of the first questions to arise could be: "How failsafe must their communication be, such that they still exhibit a sensible behavior as a system?" To probe this question, the communication network between the particles is described by a Gilbert random graph, meaning that with a certain probability p there is interaction between each pair of nodes. For $p = 1$ we have a fully connected network and thus revert to our original standard CBO algorithm. For $p = 0$ there is no communication between the particles and therefore each particle suspects the minimum m^i at its own position X^i, resulting in a complete standstill. Anything in between might be of interest. This idea is related to the batch-wise updating as discussed in Section 5.1.1. There, for every update a random subset of nodes is omitted from the communication network. Now here, a random subset of edges is omitted. The fact that the batch-wise updating was already used successfully for a CBO method seems to indicate that there is robustness towards reducing communication. Immediately emerging research questions, perhaps of a mostly computational nature at first, are for instance:

For which values of $p \in (0,1)$ do the particles find a consensus? If they find one,

is it a minimum? If they do not find one, are there maybe still particle clusters indicating local minima? Which properties of the target function appear to be most influential for the answers of the previous questions? How does the updating interval, from "every time step" to "once per run", of the random network influence the system behavior? Does the additionally introduced randomness allow for a reduction of randomness at another point of the algorithm, i.e. under what circumstances might the diffusion term be dropped?

We have no answers for these question, we only provide the lone Figure A.1 to justify posing them with at least some data. The accuracy of the standard CBO method with different degrees of connectedness is shown on a 2-dimensional Rastrigin function. The connections are randomly chosen at every time step.

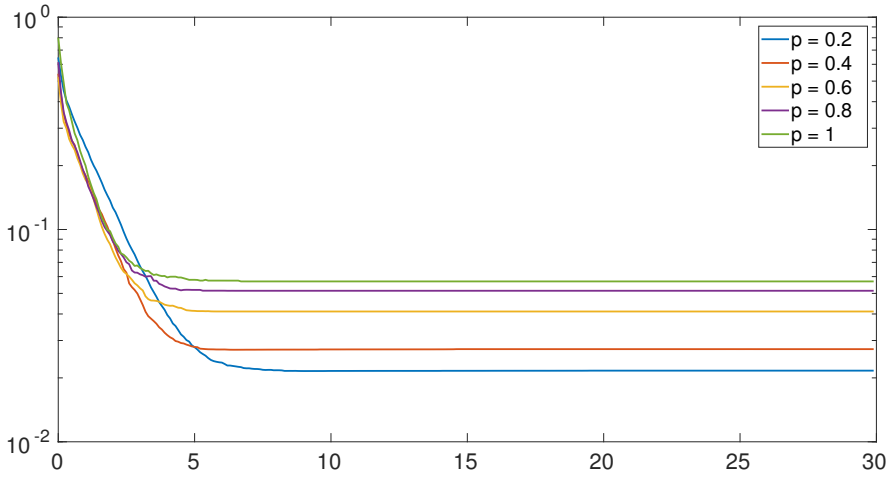

Figure A.1: CBO accuracy on a Rastrigin function with varying degrees of connectedness between the particles.

Down to $p = 0.4$, the CBO on a random Graph seems to slightly outperform the standard method. However, for $p = 0.2$ it already takes a bit longer to approximate the minimum and not included in the figure are the values of $p \lesssim 0.17$ which do not really perform at all anymore. Now, not in every scenario is the picture that drastic, but in all our tests it appeared that some of the connections between particles could be dropped without or with only positive consequences.

A.0.2 CBO with momentum: A Second-Order System

Another natural extension of the CBO is to give the particles momentum. With a twinkle in the eye, one might call this the antithesis of using the particle history. Instead of using past information, the particles now steam forward. One possible formulation of this idea, relying on the notation in the main part of this work, is given by

$$\begin{aligned}
\mathrm{d}X_t^{i,N} &= V_t^{i,N}\,\mathrm{d}t \\
\mathrm{d}V_t^{i,N} &= \lambda\left\{\gamma\left(\frac{\sum_j \phi(\left|X_t^{i,N}-X_t^{j,N}\right|)(X_t^{i,N}-X_t^{j,N})\,w_\alpha^f(X_t^{j,N})}{\sum_j \phi(\left|X_t^{i,N}-X_t^{j,N}\right|)\,w_\alpha^f(X_t^{j,N})}\right)-V_t^{i,N}\right\}\mathrm{d}t+\sigma\left|V_t^{i,N}\right|\mathrm{d}W_t^{i,N},
\end{aligned}$$

where the scaling function γ satisfies $\gamma(0)=0$. The clear interpretation of this is using CBO dynamics to describe the particle acceleration and velocity instead of the velocity and position, and just have the positions result from this. Some minor changes for the CBO dynamics are nonetheless sensible here.

It makes sense to introduce the scaling function γ since otherwise we are comparing differences in position to velocities, as clearly the particles positions cannot be replaced by the velocities for the computation of the attraction point. Possible examples are a cubic attraction potential $\gamma(\xi)=\beta\xi|\xi|^2$, a linear attraction potential $\gamma(\xi)=\beta\xi$ or a root attraction potential $\gamma(\xi)=\beta\xi\,|\xi|^{-\frac{1}{2}}$.

That means, the attraction point determines the direction of the acceleration vector of the particle. The distance from the attraction point is used to compute the magnitude of the acceleration vector via γ. Then the magnitude of the velocity is used as an indicator of the inherent uncertainty the particle has about its position and therefore used for the magnitude of the randomness in the acceleration. It seems also possible to introduce the randomness in the position update, but that did not yield good results in our numerical tests.

Figure A.2 shows a comparison of the standard CBO method and the momentum CBO method in a fairly standard setting on a Rastrigin function. The standard method is faster and more accurate to find the optimum. This is not very surprising, since the momentum method has a phase space with twice the dimension size. A steady state is only reached whenever both, velocity **and acceleration** are zero. The

fact that the momentum method does not approach the optimum as monotonously as the standard method can also be explained easily. The wavy form of the graph is due to overshooting. The particles accelerate towards the optimum and then need to brake and turn around, they have momentum.

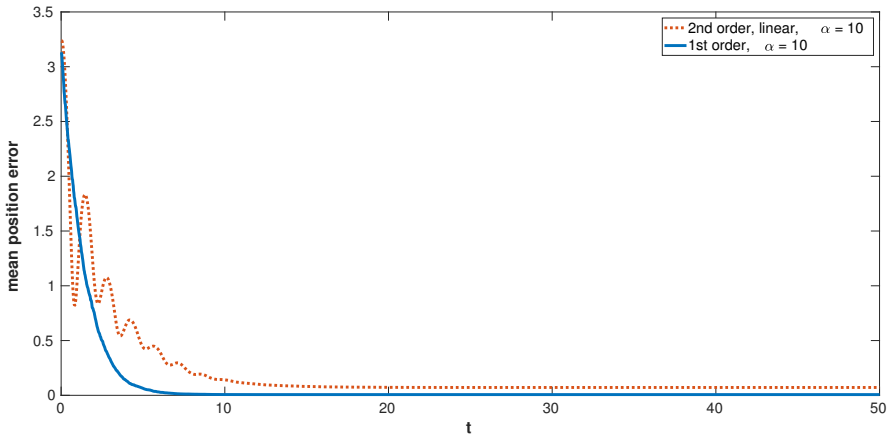

Figure A.2: Accuracy of standard CBO and momentum CBO with linear scaling on a 2-d Rastrigin function

The standard CBO method convinces with its simplicity and does the job in this standard case, but there is a not-so-far-fetched scenario where the momentum method has an advantage. Figure A.3 shows a scenario where the Rastrigin function as the target is shifted such that the optimum is at $x^* = (6,6)$, so that it is outside of the uniform initialisation area $[-5,5]^2$.

The momentum CBO finds the optimum faster and more accurately. The Rastrigin function is overall tilted towards the global optimum but has many local minima which are almost as deep. One can argue that that average tilt is in some sense picked up by the momentum CBO. The attraction point tends towards the lower side of the function and the particles overall not only contract as in the standard case, but also pick up speed towards that direction. They follow the averaged slope of the target function.

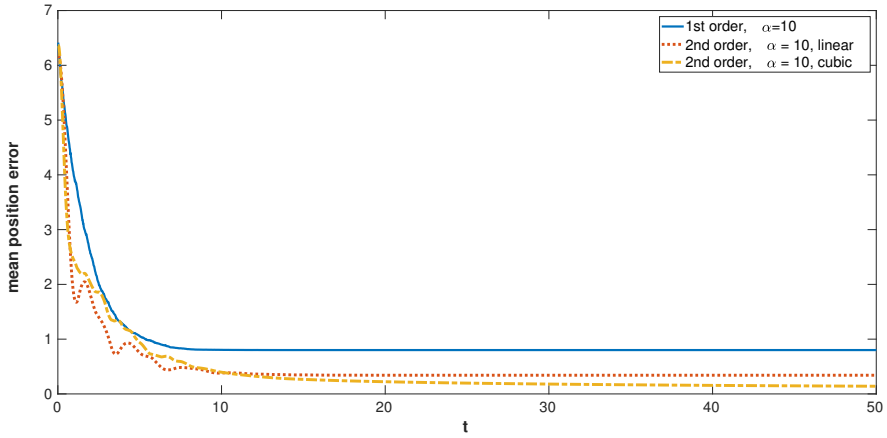

Figure A.3: Standard and momentum CBO with optimum outside initialisation area