



# Polarized consensus-based dynamics for optimization and sampling

Leon Bungert<sup>1</sup> · Tim Roith<sup>2</sup> · Philipp Wacker<sup>3</sup>

Received: 25 December 2022 / Accepted: 6 May 2024 / Published online: 31 May 2024  
© The Author(s) 2024

## Abstract

In this paper we propose polarized consensus-based dynamics in order to make consensus-based optimization (CBO) and sampling (CBS) applicable for objective functions with several global minima or distributions with many modes, respectively. For this, we “polarize” the dynamics with a localizing kernel and the resulting model can be viewed as a bounded confidence model for opinion formation in the presence of common objective. Instead of being attracted to a common weighted mean as in the original consensus-based methods, which prevents the detection of more than one minimum or mode, in our method every particle is attracted to a weighted mean which gives more weight to nearby particles. We prove that in the mean-field regime the polarized CBS dynamics are unbiased for Gaussian targets. We also prove that in the zero temperature limit and for sufficiently well-behaved strongly convex objectives the solution of the Fokker–Planck equation converges in the Wasserstein-2 distance to a Dirac measure at the minimizer. Finally, we propose a computationally more efficient generalization which works with a predefined number of clusters and improves upon our polarized baseline method for high-dimensional optimization.

**Keywords** Global optimization · Consensus-based optimization · Polarization · Sampling

**Mathematics Subject Classification** 90C26 · 35Q93 · 35B40 · 65N21

✉ Tim Roith  
tim.roith@desy.de

Leon Bungert  
leon.bungert@uni-wuerzburg.de

Philipp Wacker  
philipp.wacker@canterbury.ac.nz

- <sup>1</sup> Institute of Mathematics, University of Würzburg, Emil-Fischer-Str. 40, 97074 Würzburg, Germany
- <sup>2</sup> Computational Imaging, Helmholtz Imaging, Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany
- <sup>3</sup> School of Mathematics and Statistics, University of Canterbury, Science Road, Christchurch, Ilam 8140, New Zealand

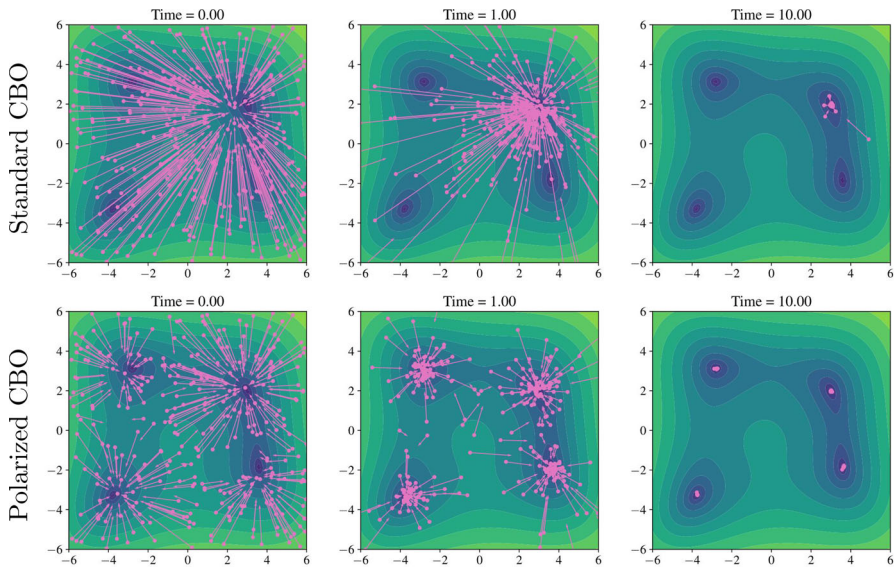
## 1 Introduction

Partially motivated by the success of machine learning methods, which involve the minimization of high-dimensional and strongly non-convex objectives, in recent years the interest in consensus-based methods which do not rely on first-order gradient information has constantly increased. Typically, zero-order optimization methods either construct a surrogate of the gradient and then perform a gradient-descent-type update [1] or use a swarm model [2, 3] where particles are attracted to the particle in the swarm which has the lowest objective value. Notably, the former approach can lead to problems for strongly non-convex objectives with many critical points. This issue is circumvented by swarm models, however, they typically do not allow for a mean-field description which makes their mathematical understanding difficult.

In contrast, consensus-based methods aim to achieve consensus by letting particles  $\{x^{(i)}\}_{i=1}^J$  explore the objective landscape while attracting them to the weighted average of their positions with respect to the Gibbs measure  $\pi \propto \exp(-\beta V)$  of the objective function  $V : \mathbb{R}^d \rightarrow \mathbb{R}$ , where  $\beta > 0$  is an inverse heat parameter. This method was first introduced in [4] as consensus-based optimization (CBO). As opposed to most other swarm methods CBO has a mean-field formulation involving a nonlinear Fokker–Planck equation. In [5] convergence to consensus of this equation was first proved, and [6] showed consensus formation of the discrete particle method. The key property for showing that the consensus is achieved close to the global minimum of the objective  $V$  is that  $\exp(-\beta V)$  concentrates around the global minimizer of  $V$  as  $\beta \rightarrow \infty$ . More recently, [7] presented an improved convergence analysis which weakens some of the assumptions in [5] and directly proves convergence to a point close to the global minimum in a Wasserstein-2 distance.

Following up on the original formulation of CBO, numerous extensions were suggested. In [8] an additional drift term, modelling the time-average of the personal best of every particle, is proposed. In [9] an extension of the CBO method for constrained problems is suggested, and [10] adapted the method for high-dimensional problems from machine learning by introducing random batching and changing the noise model. In a different line of work, CBO methods on hypersurfaces were studied in [11] and applications to machine learning were investigated in [12]. In [13] CBO was enriched by ensemble-based gradient information which comes at low computational cost and can improve upon the baseline method. For an overview of recent developments we also refer to the review [14]. Consensus-based methods have also been transferred to sampling. The work [15] proposes a consensus-based sampling (CBS) method by changing the noise term in CBO to include a weighted sample covariance. This prevents a collapse of the ensemble to full consensus and under suitable assumptions the method was shown to converge to a Gaussian approximation of the distribution  $\exp(-V)$ .

While existing consensus-based methods have proven to work very well for non-convex objectives with many spurious local extreme points, they all suffer from the conceptual drawback that by definition they can at most approximate one minimum, respectively one mode in the context of sampling. Therefore, the goal of this paper is to design consensus-based particle dynamics which support multiple consensus points, or in other words, polarization.



**Fig. 1** Dynamics of standard and the proposed polarized CBO for minimizing the Himmelblau function. The points mark particle locations, the arrows the drift field towards the weighted means

The central idea for the method which we are presenting in this paper is based on the following thought experiment: Assuming that two clusters of particles have formed, each one centered around a global minimum, we do not want to compute a shared weighted mean of their positions. This would pull one of the clusters into the other one. Therefore, we replace the weighted mean—being an integral component of current consensus-based methods—by a collection of localized means which additionally weight particle positions by their proximity to the considered particle. The localization is achieved by a suitable kernel function. This then leads to polarized dynamics where multiple consensus points can be reached which opens up the possibility of finding multiple global minima or multiple modes, respectively. Standard CBO, on the other hand is bound to converge to a single minimum.

As we shall discuss in more detail later, our approach carries strong similarities to bounded confidence models of opinion formation, introduced in [16] but see also [17–19]. There agents only interact with each other if their opinions are sufficiently close, which in the end can lead to formation of multiple consensus points and is referred to as polarization of opinions in [17].

### Main contributions

The focus of this paper is the development of the novel polarized consensus-based dynamics and an extensive numerical evaluation of the method. The mathematical models we propose come with a lot of new theoretical questions, as well, which will be the subject of future investigations. The main contributions and structure of this paper can be summarized as follows:

- Section 2.1: We propose a novel polarized computation of weighted means for CBO methods.

- Section 2.2: We propose an algorithmic variant which uses a predetermined number of cluster points to compute weighted means.
- Section 2.3: We propose a novel polarized computation of weighted covariances for the CBS method and prove that it is unbiased for Gaussian targets.
- Section 2.4: We prove convergence of the mean-field dynamics in the Wasserstein-2 distance for sufficiently well-behaved objective functions (Theorem 4).
- Section 3: We conduct extensive numerical and statistical evaluations of our polarized optimization method, showing that it can find multiple global minima in low and high dimensional optimization problems and can even improve upon standard CBO for the detection of one minimum. We also test our polarized CBS method for sampling from mixtures of Gaussian and a non-Gaussian distribution where it exhibits better performance than standard CBS.

## 2 Models

In this section we describe in detail how standard consensus-based models work and how we generalize them with our polarized approach. For this, we let  $V : \mathbb{R}^d \rightarrow [0, \infty)$  be a (possibly non-convex) objective function and  $\beta > 0$  an inverse heat parameter.

### 2.1 Polarized consensus-based optimization

Given a measure  $\rho \in \mathcal{M}(\mathbb{R}^d)$ , for standard CBO one defines a weighted mean as

$$m_\beta[\rho] := \frac{\int y \exp(-\beta V(y)) d\rho(y)}{\int \exp(-\beta V(y)) d\rho(y)}. \quad (1)$$

Here and in the rest of the paper all integrals will be over  $\mathbb{R}^d$ . The consensus-based optimization method from [4] then amounts to solving the following system of stochastic differential equations (SDEs) for particles  $\{x^{(i)}\}_{i=1,\dots,J}$ :

$$dx^{(i)} = -(x^{(i)} - m_\beta[\rho]) dt + \sigma |x^{(i)} - m_\beta[\rho]| dW^{(i)}, \quad \rho := \frac{1}{J} \sum_{i=1}^J \delta_{x^{(i)}}, \quad (2)$$

where  $\{W^{(i)}\}_{i=1}^J$  denote independent Brownian motions and  $\sigma \geq 0$  determines the strength of randomness in the model. The Fokker–Planck equation associated to (2) is the following PDE

$$\partial_t \rho_t(x) = \operatorname{div} \left( \rho_t(x) (x - m_\beta[\rho_t]) \right) + \frac{\sigma^2}{2} \Delta \left( \rho_t(x) |x - m_\beta[\rho]|^2 \right). \quad (3)$$

As explained above, this dynamical system forces particles to collapse to consensus, meaning that under certain conditions on  $V$  the empirical measures  $\rho(t)$  converge to  $\delta_{\hat{x}}$  as  $t \rightarrow \infty$ , where for  $\beta \rightarrow \infty$  the consensus-point  $\hat{x}$  converges to the global minimizer of  $V$ , see [5, 7].

We now explain our polarized modification of CBO for optimizing objective functions with multiple global minima. Given a measure  $\rho \in \mathcal{M}(\mathbb{R}^d)$  and a kernel function  $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, \infty)$  we define the weighted mean

$$m_{\beta,k}[\rho](x) := \frac{\int k(x, y) y \exp(-\beta V(y)) d\rho(y)}{\int k(x, y) \exp(-\beta V(y)) d\rho(y)}, \quad x \in \mathbb{R}^d. \quad (4)$$

The corresponding polarized optimization dynamics take the form

$$dx^{(i)} = -(x^{(i)} - m_{\beta,k}[\rho](x^{(i)})) dt + \sigma |x^{(i)} - m_{\beta,k}[\rho](x^{(i)})| dW^{(i)}, \quad (5)$$

where  $\rho := \frac{1}{J} \sum_{i=1}^J \delta_{x^{(i)}}$ . The Fokker–Planck equation associated to (5) is the following PDE

$$\partial_t \rho_t(x) = \operatorname{div} \left( \rho_t(x) (x - m_{\beta,k}[\rho_t](x)) \right) + \frac{\sigma^2}{2} \Delta \left( \rho_t(x) |x - m_{\beta,k}[\rho_t](x)|^2 \right). \quad (6)$$

Note that the idea of letting the dynamics of particle  $x^{(i)}$  mainly depend on spatially close particles, as modelled through the kernel  $k$ , has strong similarities to bounded confidence models of opinion dynamics introduced in [16]. In these models, typically there is no objective function to be minimized and the kernel is of the form  $k(x, y) := 1_{|x-y| \leq \kappa}(x, y)$ , where  $\kappa > 0$  is a so-called confidence level. The simplest such dynamics then take the form

$$\frac{dx^{(i)}}{dt} = - \left( x^{(i)} - \frac{1}{N(x^{(i)})} \sum_{j=1}^J 1_{|x^{(i)} - x^{(j)}| \leq \kappa} x^{(j)} \right), \quad (7)$$

where  $N(x^{(i)}) := \#\{1 \leq j \leq J : |x^{(i)} - x^{(j)}| \leq \kappa\}$  denotes the numbers of points which are not farther than  $\kappa$  away from  $x^{(i)}$ . Notably, (7) coincides with (5) for the special case of  $k(x, y) := 1_{|x-y| \leq \kappa}(x, y)$ , a constant objective  $V \equiv \text{const}$ , and  $\sigma = 0$ . More generally, dynamics of the form (5) and (7) can be viewed as processes on co-evolving networks or graphs where the weights between different particles  $x^{(i)}$  and  $x^{(j)}$  depend on the kernel  $k$  and the loss function  $V$ . We refer to [20, 21] for a unified description and the study of mean-field equations for general processes of this form. Furthermore, for  $V = \text{const}$  and  $\sigma = 0$ , a time discretization of Equation (5) via an Euler–Maruyama scheme with stepsize  $dt = 1$ , yields the update

$$x^{(i)} \leftarrow x^{(i)} - (x^{(i)} - m_{\beta,k}[\rho](x^{(i)})) = \frac{\sum_{j=1}^N k(x^{(i)}, x^{(j)}) x^{(j)}}{\sum_{j=1}^N k(x^{(i)}, x^{(j)})},$$

which is known as the *mean-shift* algorithm [22, 23]. In the following, we would like to discuss two important special cases of our model: If one chooses the kernel

$k(x, y) = 1$  for all  $x, y \in \mathbb{R}^d$ , then (4) to (6) simply reduce to the standard CBO setup (1) to (3). Hence, our method is a generalization of CBO. On the other hand, if one chooses the Gaussian kernel  $k(x, y) := \exp\left(-\frac{|x-y|^2}{2\kappa^2}\right)$ , then the weighted mean  $m_{\beta,k}[\rho](x)$  can be rewritten as

$$m_{\beta,k}[\rho](x) := \frac{\int y \exp\left(-\beta V(y) - \frac{|x-y|^2}{2\kappa^2}\right) d\rho(y)}{\int \exp\left(-\beta V(y) - \frac{|x-y|^2}{2\kappa^2}\right) d\rho(y)}. \quad (8)$$

In this case our method can be regarded as standard CBO applied to a spatially varying quadratic regularization of the objective function  $y \mapsto V_x(y)$ , defined as

$$V_x(y) := V(y) + \frac{1}{2\kappa^2\beta} |x - y|^2.$$

Note that the central difference between the standard and our method is that the weighted mean (4) depends on the particle position  $x$  and is not same for all particles. Especially for high-dimensional problems it was shown in [10] that the performance of CBO can be significantly improved when using a coordinate-wise noise model. To be precise they suggest the following replacement in (5)

$$\left| x^{(i)} - m_{\beta,k}[\rho](x^{(i)}) \right| dW^{(i)} \longrightarrow \sum_{n=1}^d \tilde{e}_n \left( x^{(i)} - m_{\beta,k}[\rho](x^{(i)}) \right)_n dW_n^{(i)}, \quad (9)$$

where  $\tilde{e}_n$  denotes the  $n$ th unit vector in  $\mathbb{R}^d$ . This changes the Laplacian term in the corresponding Fokker–Planck equation (6) to

$$\Delta(\rho_t(x) |x - m_{\beta,k}[\rho_t](x)|) \longrightarrow \sum_{n=1}^d \partial_{nn}^2 \left( \rho_t(x) (x - m_{\beta,k}[\rho_t](x))_n \right).$$

For computing an approximate solution of the stochastic CBO dynamics with either of the two discussed noise models, we employ a standard Euler–Maruyama scheme which we sum up in Algorithm 1. There the function **ComputeMean** determines the precise variant of CBO that is being used. In Algorithm 2 we specify this function for the proposed polarized CBO which reduces to standard CBO for a constant kernel  $k \equiv 1$ . In the next section we will discuss an algorithmic variant (see Algorithm 3) of the mean computation (4) which is computationally more efficient and exhibits empirical advantages in high dimensions.

## 2.2 Cluster-based model

In this section we propose an algorithmic alternative to the weighted mean, defined in (4) with the motivation of making the computation of the weighted means more efficient and to also encourage polarizing effects between the different means. Given

**Algorithm 1** General consensus-based optimization

**Input:** Time step size  $dt > 0$ , initial particles  $\{x^{(j)}\}_{j=1}^J$ , diffusion parameter  $\sigma \geq 0$ , noise model  $M$

```

for  $i = 1, \dots, J$  do
   $m^{(i)} \leftarrow \text{ComputeMean}$ 
  if  $M = \text{Standard noise}$  then
     $\xi \sim \mathcal{N}(0, dt \cdot I_{d \times d})$ 
     $n^{(i)} \leftarrow |x^{(i)} - m^{(i)}| \xi$ 
  end if
  if  $M = \text{Coordinate noise}$  then
     $\xi \sim \mathcal{N}(0, dt \cdot I_{d \times d})$ 
     $n^{(i)} \leftarrow \sum_{n=1}^d e_n (x^{(i)} - m^{(i)})_n \xi_n$ 
  end if
   $x^{(i)} \leftarrow x^{(i)} - dt(x^{(i)} - m^{(i)}) + \sigma n^{(i)}$ 
end for

```

**Algorithm 2** ComputeMean for polarized CBO

**Input:** particle  $x^{(i)}$ , particle swarm  $\{x^{(j)}\}_{j=1}^J$ , kernel  $k$ , objective  $V$ , inverse temperature  $\beta > 0$

**Output:**  $m^{(i)} \leftarrow \frac{\sum_{j=1}^J x^{(j)} k(x^{(i)}, x^{(j)}) \exp(-\beta V(x^{(j)}))}{\sum_{j=1}^J k(x^{(i)}, x^{(j)}) \exp(-\beta V(x^{(j)}))}$

particles  $x^{(i)}$  for  $i = 1, \dots, J$ , standard CBO uses one weighted mean whereas our polarized version uses  $J$  many weighted means.

As an alternative model we consider cluster means  $c^{(j)}$  for  $j = 1, \dots, J_c$ , where  $J_c \leq J$ . We encode the probability that the particle  $x^{(i)}$  belongs to the cluster mean  $c^{(j)}$  with  $p_{ij} > 0$ . Given initial probabilities  $p_{ij} > 0$ , we perform the following iterative update for each  $j = 1, \dots, J_c$ ,

$$p_i^{\max} := \max_{j=1, \dots, J_c} p_{ij} \quad \text{for } i = 1, \dots, J, \quad (10a)$$

$$r_{ij} := \left( \frac{p_{ij}}{p_i^{\max}} \right)^\alpha \quad \text{for } i = 1, \dots, J, \quad (10b)$$

$$\tilde{p}_{ij} := r_{ij} k(x^{(i)}, c^{(j)}). \quad (10c)$$

The new values of the probabilities  $p_{ij}$  are then obtained by renormalization over  $j$ , i.e.,

$$p_{ij} \leftarrow \frac{\tilde{p}_{ij}}{\sum_{j=1}^{J_c} \tilde{p}_{ij}} \quad (11)$$

and the clusters means are then updated via

$$c^{(j)} \leftarrow \frac{\sum_{i=1}^J x^{(i)} p_{ij} \exp(-\beta V(x^{(i)}))}{\sum_{i=1}^J p_{ij} \exp(-\beta V(x^{(i)}))}. \quad (12)$$

Here,  $\alpha \geq 0$  is a discounting coefficient. The interpretation of the above scheme is as follows: If the particle  $x^{(i)}$  identifies the cluster center  $c^{(j)}$  as “belonging to it the most”, then  $r_{ij} = 1$  and the probability that  $x^{(i)}$  indeed belongs to  $c^{(j)}$  is only determined by the spatial proximity, encoded through  $k(x^{(i)}, c^{(j)})$ . If, however, it “feels more dedicated” to another cluster mean  $c^{(k)}$  (different from  $c^{(j)}$ ), then  $p_{ij} < p_{ik}$  and thus  $r_{ij} < 1$ . This results in a correction of the particle-cluster correspondence, strengthening the bond to cluster  $j$  of particle  $i$ . The exponent  $\alpha \geq 0$  determines the strength of this additional polarization incentive, with larger  $\alpha \geq 0$  leading to greater polarization. In order to obtain the individual mean for each particle (like in the polarized CBO scheme) we simply compute

$$m^{(i)} := \sum_{j=1}^{J_c} p_{ij} c^{(j)}. \quad (13)$$

We summarize the cluster-based mean computations in Algorithm 3, which can be used in place of Algorithm 2 in the CBO scheme Algorithm 1.

Note that for  $\alpha = \infty$  we have that

$$r_{ij} := \begin{cases} 1, & \text{if } j \in \arg \max_{j=1, \dots, J_c} p_{ij}, \\ 0, & \text{else,} \end{cases}$$

meaning that the only probability which survives is the one for the likeliest cluster centers. Another interesting special case arises when choosing the Gaussian kernel  $k(x, y) := \exp\left(-\frac{|x-y|^2}{2\kappa^2}\right)$ . In this case one gets that

---

### Algorithm 3 ComputeMean for cluster-based method

---

**Input:** particle  $x^{(i)}$ , particle swarm  $\{x^{(j)}\}_{j=1}^J$ , cluster centers  $\{c^{(j)}\}_{j=1}^{J_c}$ , probabilities  $\{p_{ij}\}_{j=1}^{J_c}$ , discounting coefficient  $\alpha \geq 0$ , kernel  $k$ , objective  $V$ , inverse temperature  $\beta > 0$

```

 $p_i^{\max} \leftarrow \max_{j=1, \dots, J_c} p_{ij}$  ▷ Compute probability of likeliest cluster
for  $j = 1, \dots, J_c$  do
   $r_{ij} \leftarrow \left(\frac{p_{ij}}{p_i^{\max}}\right)^\alpha$ 
   $\tilde{p}_{ij} \leftarrow r_{ij} k(x^{(i)}, c^{(j)})$  ▷ Discount probabilities for other clusters
end for
for  $j = 1, \dots, J_c$  do
   $p_{ij} \leftarrow \frac{\tilde{p}_{ij}}{\sum_{j=1}^{J_c} \tilde{p}_{ij}}$  ▷ Normalize probabilities
end for
for  $j = 1, \dots, J_c$  do
   $c^{(j)} \leftarrow \frac{\sum_{i=1}^J x^{(i)} p_{ij} \exp(-\beta V(x^{(i)}))}{\sum_{i=1}^J p_{ij} \exp(-\beta V(x^{(i)}))}$  ▷ Update cluster centers
end for
Output:  $m_i \leftarrow \sum_{j=1}^{J_c} p_{ij} c^{(j)}$ 

```

---



$$\begin{aligned}
p_{ij} &:= \frac{\exp\left(-\frac{1}{2\kappa^2} |x^{(i)} - c^{(j)}|^2 + \log r_{ij}\right)}{\sum_{j=1}^{J_c} \exp\left(-\frac{1}{2\kappa^2} |x^{(i)} - c^{(j)}|^2 + \log r_{ij}\right)} \\
&\rightarrow \begin{cases} 1 & \text{if } |x^{(i)} - c^{(j)}| \leq |x^{(i')} - c^{(j')}| \quad \forall i' = 1, \dots, J, \quad j' = 1, \dots, J_c, \\ 0 & \text{else,} \end{cases}
\end{aligned}$$

as  $\kappa \rightarrow 0$ . So for very small values of  $\kappa$  a hard assignment of the points  $x^{(i)}$  to the clusters  $c^{(j)}$  based on spatial proximity is performed which is reminiscent of the k-means algorithm.

It is important to initialize all quantities correctly in order to obtain a meaningful algorithm. If one naively initializes  $p_{ij} := \frac{1}{J_c}$  for all  $i = 1, \dots, J$  and computes initial cluster centers via (12), then  $c^{(j)} = \frac{\sum_{i=1}^J x^{(i)} \exp(-\beta V(x^{(i)}))}{\sum_{i=1}^J \exp(-\beta V(x^{(i)}))}$  equals the standard CBO weighted mean for all  $j$ . Correspondingly, the probability updates (10) and (11) will leave  $p_{ij}$  untouched. In this case the method reduces precisely to standard CBO. Therefore, we initialize the probabilities randomly by drawing  $\tilde{p}_{ij} \sim \text{Unif}(0, 1)$  and normalizing with (11). Then we compute initial cluster means using (12).

Let us also mention that the complexity of CBO where means are computed with Algorithm 3 is order  $O(J \cdot J_c)$  which is significantly smaller than  $O(J^2)$  which is the complexity of CBO with the baseline polarized mean computation from Algorithm 2. This is due to the fact that the cluster-based method models consensus and polarization of individuals with respect to existing opinionated groups, with the cluster center as a surrogate for these groups, whereas the polarized CBO approach tracks all particles' interactions with each other.

Last, we demonstrate how to obtain a SDE and mean-field interpretation of Algorithm 3, where for conciseness we restrict ourselves to the case  $\alpha = 0$ . In this case (11) and (12) reduce to

$$p_{ij} \leftarrow \frac{k(x^{(i)}, c^{(j)})}{\sum_{j=1}^{J_c} k(x^{(i)}, c^{(j)})}, \quad c^{(j)} \leftarrow \frac{\sum_{i=1}^J x^{(i)} k(x^{(i)}, c^{(j)}) \exp(-\beta V(x^{(i)}))}{\sum_{i=1}^J k(x^{(i)}, c^{(j)}) \exp(-\beta V(x^{(i)}))}.$$

This allows us to express the cluster-based mean computation with  $\alpha = 0$  as the coupled system:

$$dx^{(i)} = -(x^{(i)} - m^{(i)}) dt + \sigma |x^{(i)} - m^{(i)}| dW^{(i)}, \quad (14a)$$

$$m^{(i)} = \frac{\sum_{j=1}^{J_c} k(x^{(i)}, c^{(j)}) c^{(j)}}{\sum_{j=1}^{J_c} k(x^{(i)}, c^{(j)})}, \quad (14b)$$

$$c^{(j)} = \frac{\int x k(x, c^{(j)}) \exp(-\beta V(x)) d\rho_t(x)}{\int k(x, c^{(j)}) \exp(-\beta V(x)) d\rho_t(x)}, \quad (14c)$$

$$\rho_t := \frac{1}{N} \sum_{i=1}^J \delta_{x^{(i)}}. \quad (14d)$$

Note that Algorithm 3 approximates the fixed point equation for  $c^{(j)}$  with one iteration of the fixed point map. The corresponding mean-field system is readily obtained as:

$$\partial_t \rho_t(x) = \operatorname{div}(\rho_t(x)(x - m[\rho_t](x))) + \frac{\sigma^2}{2} \Delta \left( \rho_t(x) |x - m[\rho_t](x)|^2 \right), \quad (15a)$$

$$m[\rho_t](x) = \frac{\sum_{j=1}^{J_c} k(x, c^{(j)}) c^{(j)}}{\sum_{j=1}^{J_c} k(x, c^{(j)})}, \quad (15b)$$

$$c^{(j)} = \frac{\int x k(x, c^{(j)}) \exp(-\beta V(x)) d\rho_t(x)}{\int k(x, c^{(j)}) \exp(-\beta V(x)) d\rho_t(x)}. \quad (15c)$$

We expect that similarly one can derive a SDE and mean-field interpretation of the model with  $\alpha > 0$  but we do not endeavour this here.

### 2.3 Polarized consensus-based sampling

The last model variant that we consider here is an application to sampling. In [15] a sampling version of CBO was proposed and termed consensus-based sampling (CBS). Defining a weighted covariance matrix as

$$C_\beta[\rho] := \frac{\int (y - m_\beta[\rho]) \otimes (y - m_\beta[\rho]) \exp(-\beta V(y)) d\rho(y)}{\int \exp(-\beta V(y)) d\rho(y)}. \quad (16)$$

CBS aims to sample from the measure  $\exp(-V)$  by solving the following system of SDEs:

$$dx^{(i)} = -(x^{(i)} - m_\beta[\rho]) dt + \sqrt{2\lambda^{-1} C_\beta[\rho]} dW^{(i)}, \quad \rho := \frac{1}{J} \sum_{i=1}^J \delta_{x^{(i)}}. \quad (17)$$

Here the parameter  $\lambda$  interpolates between an optimization method ( $\lambda = 1$ ) and a sampling method ( $\lambda = (1 + \beta)^{-1}$ ). For the latter scaling of  $\lambda$  a collapse of  $\rho$  is avoided and  $\rho$  samples from  $\exp(-V)$  if this measure is Gaussian. The Fokker–Planck equation associated with (17) is given by

$$\partial_t \rho_t(x) = \operatorname{div} \left( \rho_t(x)(x - m_\beta[\rho_t]) \right) + \lambda^{-1} \operatorname{div} \left( C_\beta[\rho_t] \nabla \rho_t(x) \right). \quad (18)$$

We can polarize CBS by using the mean from (4) to define a weighted variance, as follows:

$$C_{\beta,k}[\rho](x) := \frac{\int k(x, y)(y - m_{\beta,k}[\rho](x)) \otimes (y - m_{\beta,k}[\rho](x)) \exp(-\beta V(y)) d\rho(y)}{\int k(x, y) \exp(-\beta V(y)) d\rho(y)}, \quad (19)$$

for  $x \in \mathbb{R}^d$ . The corresponding CBS dynamics are then given by

$$\boxed{dx^{(i)} = -(x^{(i)} - m_{\beta,k}[\rho](x^{(i)})) dt + \sqrt{2\lambda^{-1}C_{\beta,k}[\rho](x^{(i)})} dW^{(i)},} \quad (20)$$

where  $\rho := \frac{1}{J} \sum_{i=1}^J \delta_{x^{(i)}}$ . The associated Fokker–Planck equation is

$$\partial_t \rho_t(x) = \operatorname{div}(\rho_t(x)(x - m_{\beta,k}[\rho_t](x))) + \lambda^{-1} \operatorname{div}(C_{\beta,k}[\rho_t](x) \nabla \rho_t(x)). \quad (21)$$

Perhaps a little unexpectedly we can prove that, just as the Fokker–Planck equation of CBS (18), for Gaussian kernels of arbitrary width our version (21) leaves Gaussians measures invariant, which is an important consistency property.

**Proposition 1** *The polarized CBS dynamics in sampling mode with a Gaussian kernel leaves a Gaussian target measure invariant. More precisely, let*

$k(x, y) := \exp\left(-\frac{1}{2}(x - y)^T \Sigma_1^{-1}(x - y)\right)$  *for a symmetric and positive definite covariance matrix*  $\Sigma_1 \in \mathbb{R}^{d \times d}$ , *and let*

$$V(y) := \frac{1}{2}(x - m)^T \Sigma_2^{-1}(x - m)$$

*for some*  $m \in \mathbb{R}^d$  *and a symmetric and positive definite covariance matrix*  $\Sigma \in \mathbb{R}^{d \times d}$ . *Then*  $\rho^*$  *defined as*

$$\rho^*(x) := \exp(-V(x))$$

*is a stationary solution of the Fokker–Planck equation (21) for any*  $\beta > 0$  *and for*  $\lambda = (1 + \beta)^{-1}$ .

**Proof** We use the following formula for the product of two Gaussians, which can be found, for instance, in [24], to obtain

$$\begin{aligned} k(x, y) \exp(-(1 + \beta)V(y)) &= \exp\left(-\frac{1}{2}(x - y)^T \Sigma_1^{-1}(x - y)\right) \\ &\quad \exp\left(-\frac{1 + \beta}{2}(x - m)^T \Sigma_2^{-1}(x - m)\right) \\ &= c_x \exp\left(-\frac{1}{2}(y - m_x)^T \Sigma_3(y - m_x)\right), \end{aligned}$$

where  $c_x > 0$  is a normalization constant and

$$\begin{aligned} \Sigma_3 &:= \left(\Sigma_1^{-1} + (1 + \beta)\Sigma_2^{-1}\right)^{-1}, \\ m_x &:= \Sigma_3 \left(\Sigma_1^{-1}x + (1 + \beta)\Sigma_2^{-1}m\right). \end{aligned}$$

Using this we obtain

$$\begin{aligned}\int k(x, y) \exp(-\beta V(y)) d\rho^*(y) &= \int k(x, y) \exp(-(1 + \beta)V(y)) dy \\ &= c_x \int \exp\left(-\frac{1}{2}(y - m_x)^T \Sigma_3(y - m_x)\right) dy \\ &= c_x (2\pi)^{\frac{d}{2}} \det(\Sigma_3)^{\frac{1}{2}}.\end{aligned}$$

Similarly, we obtain

$$\begin{aligned}\int y k(x, y) \exp(-V(y)) d\rho^*(y) &= c_x \int y \exp\left(-\frac{1}{2}(y - m_x)^T \Sigma_3(y - m_x)\right) dy \\ &= c_x (2\pi)^{\frac{d}{2}} \det(\Sigma_3)^{\frac{1}{2}} m_x.\end{aligned}$$

Combining the two we obtain

$$\mathfrak{m}_{\beta, k}[\rho^*](x) = m_x.$$

Similarly, we can compute the weighted covariance as

$$\begin{aligned}C_{\beta, k}[\rho^*](x) &= \frac{\int (x - \mathfrak{m}_{\beta, k}[\rho^*](x)) \otimes (x - \mathfrak{m}_{\beta, k}[\rho^*](x)) k(x, y) \exp(-\beta V(y)) d\rho^*(y)}{\int k(x, y) \exp(-V(y)) d\rho^*(y)} \\ &= \frac{c_x \int (x - m_x) \otimes (x - m_x) c_x \exp\left(-\frac{1}{2}(y - m_x)^T \Sigma_3(y - m_x)\right) dy}{c_x (2\pi)^{\frac{d}{2}} \det(\Sigma_3)^{\frac{1}{2}}} = \Sigma_3.\end{aligned}$$

Hence, we get for  $\lambda = (1 + \beta)^{-1}$ :

$$\begin{aligned}&\rho^*(x)(x - \mathfrak{m}_{\beta, k}[\rho^*](x)) + \lambda^{-1} C_{\beta, k}[\rho^*](x) \nabla \rho^*(x) \\ &= \rho^*(x)(x - m_x) + (1 + \beta) \Sigma_3 \nabla \rho^*(x) \\ &= \rho^*(x)(x - m_x) - (1 + \beta) \Sigma_3 \nabla V(x) \rho^*(x) \\ &= \rho^*(x) \left( (x - m_x) - (1 + \beta) \Sigma_3 \Sigma_2^{-1} (x - m) \right).\end{aligned}$$

Now we use the definition of  $m_x = \Sigma_3(\Sigma_1^{-1}x + (1 + \beta)\Sigma_2^{-1}m)$  to obtain

$$\begin{aligned}&x - m_x - (1 + \beta) \Sigma_3 \Sigma_2^{-1} (x - m) \\ &= x - \Sigma_3(\Sigma_1^{-1}x + (1 + \beta)\Sigma_2^{-1}m) - (1 + \beta) \Sigma_3 \Sigma_2^{-1} (x - m) \\ &= x - \Sigma_3 \left( \Sigma_1^{-1} + (1 + \beta) \Sigma_2^{-1} \right) x = 0,\end{aligned}$$

using that  $\Sigma_3 = \left( \Sigma_1^{-1} + (1 + \beta) \Sigma_2^{-1} \right)^{-1}$ . This proves that  $\rho^*$  is a stationary solution of the Fokker–Planck equation (21).  $\square$

## 2.4 Towards mean-field analysis

In this section we will make some remarks on the analysis of the Fokker–Planck equation (6) for polarized CBO which we repeat here for convenience:

$$\partial_t \rho_t(x) = \operatorname{div} \left( \rho_t(x) (x - m_{\beta,k}[\rho_t](x)) \right) + \frac{\sigma^2}{2} \Delta \left( \rho_t(x) |x - m_{\beta,k}[\rho_t](x)|^2 \right).$$

We shall first explain why current analytical frameworks for the mean-field analysis of consensus-based optimization do not carry over to this equation. Then, we shall prove convergence as  $t \rightarrow \infty$  of solutions to this equation to a singular measure located at a minimizer, restricting ourselves to the zero temperature limit ( $\beta = 0$ ) and to sufficiently nice objective functions  $V$ .

Weak solutions of this equation are continuous curves of probability measures  $t \mapsto \rho_t$  such that

$$\begin{aligned} \frac{d}{dt} \int \phi(x) d\rho_t(x) &= - \int \nabla \phi(x) \cdot (x - m_{\beta,k}[\rho_t](x)) d\rho_t(x) \\ &\quad + \frac{\sigma^2}{2} \int \Delta \phi(x) |x - m_{\beta,k}[\rho_t](x)|^2 d\rho_t(x) \end{aligned}$$

holds true for all smooth and compactly supported test functions  $\phi \in C_c^\infty(\mathbb{R}^d)$ .

Using the Leray–Schauder fixed point theorem, existence proofs for this equation without a kernel, i.e.,  $k(x, y) = 1$ , were given in [5, 7] under mild assumptions on the objective  $V$  and for initial distributions with finite fourth-order moment. Taking into account that

$$k(x, y) \exp(-\beta V(y)) = \exp \left( -\beta \left( V(y) - \frac{1}{\beta} \log k(x, y) \right) \right)$$

we expect that under reasonable Lipschitz-like assumptions on the logarithm of the kernel these arguments translate to our case, but we leave this for future work. The biggest challenge is that for standard CBO  $t \mapsto m_\beta[\rho_t]$  is a continuous curve in  $\mathbb{R}^d$ , whereas  $t \mapsto m_{\beta,k}[\rho_t](\cdot)$  is not. Rather, it is a curve in a space of vector fields. This requires more sophisticated compactness arguments than the one made in [5, 7].

Besides existence, in the literature there exist two different approaches to proving formation of consensus around the global minimizer of  $V$  for the Fokker–Planck equation (3) associated to standard CBO. The first one was presented in [5] and constitutes a two-step approach. First, they prove that the non-weighted standard variance

$$V(\rho_t) := \int |x - E(\rho_t)|^2 d\rho_t(x), \quad \text{where } E(\rho_t) := \int y d\rho_t(y),$$

decreases to zero along solutions  $\rho_t$  of (3). This obviously implies that  $\rho_t$  converges to a Dirac measure  $\delta_{\tilde{x}}$  concentrated on some point  $\tilde{x} \in \mathbb{R}^d$ . Second, the Laplace principle is invoked in order to conclude that  $\tilde{x}$  lies close to the global minimizer of  $V$  if  $\beta$

is chosen sufficiently large. This analytical approach uses heavily that the weighted mean  $m_\beta[\rho_t]$  for CBO *does not* depend on the spatial variable  $x$  which is a linearity property that our weighted mean  $m_{\beta,k}[\rho_t](x)$  does not enjoy. Furthermore, by design our method does in general not converge to a single Dirac mass and hence its classical variance does not converge to zero.

A different approach is presented in [7] where the authors propose a more unified strategy. For this they fix a point  $\hat{x} \in \mathbb{R}^d$  (which later will be the global minimizer of  $V$ ) and define the variance-type function

$$V[\rho_t] := \frac{1}{2} \int |x - \hat{x}|^2 d\rho_t(x). \quad (22)$$

The authors note that  $V[\rho_t] = \frac{1}{2} \mathcal{W}_2^2(\rho_t, \delta_{\hat{x}})$  and so convergence of the variance implies convergence of  $\rho_t$  to  $\delta_{\hat{x}}$  in the Wasserstein-2 distance. They derive a differential inequality for  $V[\rho_t]$  which in combination with the Laplace principle allows them to show some kind of semi-convergence behavior, i.e., for every  $\varepsilon > 0$  there exists  $\beta > 0$  such that  $V[\rho_t]$  decreases exponentially fast until it hits the threshold  $V[\rho_t] \leq \varepsilon$ .

When it comes to generalizing this approach to our setting there are two main obstacles. First, in the case of several global minimizers  $\{\hat{x}_i : 1 \leq i \leq N\}$  the Wasserstein-2 distance between  $\rho_t$  and the empirical measure  $\frac{1}{N} \sum_{i=1}^N \delta_{\hat{x}_i}$  does not equal  $\frac{1}{N} \sum_{i=1}^N \int |x - \hat{x}_i|^2 d\rho_t(x)$ . Indeed, the latter quantity is a very bad upper bound for the desired Wasserstein-2 distance since it is bounded from below by a positive number.

In the following, we therefore present an alternative approach for analyzing convergence of the Fokker–Planck equation (6) which is based on choosing the Lyapunov function

$$L_\beta[\rho_t] := \frac{1}{2} \int |x - m_{\beta,k}[\rho_t](x)|^2 d\rho_t(x)$$

and the reasons for this choice will become clear in a moment. We consider the setting of a Gaussian kernel  $k(x, y) := \exp\left(-\beta \frac{|x-y|^2}{2\kappa^2}\right)$  of variance  $\kappa^2/\beta$  for  $\kappa > 0$  and consider the limit as  $\beta \rightarrow \infty$ . The case of  $\beta < \infty$  is then treated using the quantitative Laplace principle as, for instance, in [7]. Let us assume that the support of  $\rho$  equals  $\mathbb{R}^d$ . In this setting the Laplace principle implies that for  $\beta \rightarrow \infty$  one has

$$m_{\beta,k}[\rho](x) = \frac{\int y \exp\left(-\beta \left(\frac{|x-y|^2}{2\kappa^2} + V(y)\right)\right) d\rho(y)}{\int \exp\left(-\beta \left(\frac{|x-y|^2}{2\kappa^2} + V(y)\right)\right) d\rho(y)} \rightarrow \arg \min_{y \in \mathbb{R}^d} \frac{|x-y|^2}{2\kappa^2} + V(y),$$

where we assume that  $\kappa$  is sufficiently small such that  $y \mapsto \frac{|x-y|^2}{2\kappa^2} + V(y)$  has a unique minimizer for every  $x \in \mathbb{R}^d$ . Note that this is possible, for instance, if  $V$  is  $C^2$  and the smallest eigenvalue of its Hessian matrix is bounded from below. Therefore, we consider the limiting dynamics as  $\beta \rightarrow \infty$ , governed by the drift field  $x - p(x)$ ,

where

$$p(x) := \arg \min_{y \in \mathbb{R}^d} \frac{|x - y|^2}{2\kappa^2} + V(y) \quad (23)$$

is known as the *proximal operator* of  $\kappa^2 V$ . Correspondingly, the Lyapunov function becomes

$$\mathbb{L}[\rho_t] := L_\infty[\rho_t] := \frac{1}{2} \int |x - p(x)|^2 \, d\rho_t(x)$$

and we would like to emphasize that, using properties of the proximal operator, one has

$$\mathbb{L}[\rho] = 0 \iff x \in \arg \min V \quad \text{for } \rho\text{-almost every } x \in \mathbb{R}^d.$$

The following analysis treats the model case where the loss function  $V$  is strongly convex, sufficiently smooth, and additionally satisfies a certain derivative bound in case we consider a Fokker–Planck equation with non-vanishing diffusion term. While this is just the first step toward a comprehensive analysis of polarized CBO method for much larger classes functions, the following results introduce a set of important techniques which—we believe—will be a cornerstone for future analysis. The core ideas of the following analysis are based on discussions of the first author with Massimo Fornasier and Oliver Tse and an extension to non-convex loss functions and  $\beta < \infty$  is ongoing work.

We start with the following decay property of  $\mathbb{L}[\rho_t]$  for solutions of the associated Fokker–Planck equation:

**Proposition 2** (Exponential decay of the Lyapunov function) *Let  $t \mapsto \rho_t$  be a weak solution of the Fokker–Planck equation*

$$\partial_t \rho_t(x) = \operatorname{div} \left( \rho_t(x)(x - p(x)) \right) + \frac{\sigma^2}{2} \Delta \left( \rho_t(x) |x - p(x)|^2 \right). \quad (24)$$

*We pose the following assumptions on the loss function  $V$ :*

- *If  $\sigma = 0$  assume that  $V \in C^2(\mathbb{R}^d)$  and there exists  $\mu > 0$  such that*

$$D^2 V \succcurlyeq \mu \mathbb{I}$$

- *If  $\sigma > 0$  assume additionally that  $V \in C^3(\mathbb{R}^d)$  and satisfies*

$$\sup_{x \in \mathbb{R}^d} |Dv(x)| \left| D^3 V(x) \right| < \infty.$$

Then there exists constants  $C_1, C_2 > 0$  such that, if  $\sigma < C_1$ , it holds for all  $t > 0$  that

$$\frac{d}{dt} \mathbb{L}[\rho_t] \leq -C_2 \mathbb{L}[\rho_t] \quad (25)$$

and consequently

$$\mathbb{L}[\rho_t] \leq \mathbb{L}[\rho_0] \exp(-C_2 t). \quad (26)$$

**Proof** Differentiating  $\mathbb{L}[\rho_t]$  in time and using the weak form of the PDE implies

$$\begin{aligned} \frac{d}{dt} \mathbb{L}[\rho_t] &= - \int \nabla \frac{1}{2} |x - p(x)|^2 \cdot (x - p(x)) \, d\rho_t(x) \\ &\quad + \frac{\sigma^2}{4} \int \Delta |x - p(x)|^2 |x - p(x)|^2 \, d\rho_t(x). \end{aligned} \quad (27)$$

We continue by computing the derivatives that appear in this expression. First, it holds

$$\partial_i \frac{1}{2} |x - p(x)|^2 = \partial_i \sum_j \frac{1}{2} |x_j - p_j(x)|^2 = \sum_j (\delta_{ij} - \partial_i p_j(x)) (x_j - p_j(x)). \quad (28)$$

Next we observe that by definition of the proximal operator  $p(x)$  it holds

$$0 = p(x) - x + \kappa^2 \nabla V(p(x)). \quad (29)$$

Differentiating this equation yields

$$0 = Dp(x) - \mathbb{I} + \kappa^2 D^2 V(p(x)) Dp(x) \quad (30)$$

and therefore

$$Dp(x) = \left( \mathbb{I} + \kappa^2 D^2 V(p(x)) \right)^{-1}, \quad (31)$$

which is a symmetric matrix. Using (28) we get

$$-\nabla \frac{1}{2} |x - p(x)|^2 \cdot (x - p(x)) = -(\mathbb{I} - Dp(x))(x - p(x)) \cdot (x - p(x))$$

For estimating this expression from above it suffices to bound the eigenvalues of  $M := \mathbb{I} - Dp(x)$  from below. By assumption we have  $D^2 V \succcurlyeq \mu \mathbb{I}$ , which implies that  $M \succcurlyeq \left(1 - \frac{1}{1+\kappa^2\mu}\right) \mathbb{I} = \frac{\kappa^2\mu}{1+\kappa^2\mu} \mathbb{I}$  and therefore we can bound the first term in (27)

$$-\nabla \frac{1}{2} |x - p(x)|^2 \cdot (x - p(x)) \leq -\frac{\kappa^2\mu}{1+\kappa^2\mu} |x - p(x)|^2. \quad (32)$$



If  $\sigma = 0$  we can already conclude the proof, using Gronwall's inequality. For  $\sigma > 0$  we have to bound the second term in (27), coming from the diffusion. Using (28) and the product rule it follows

$$\begin{aligned} \partial_i^2 \frac{1}{2} |x - p(x)|^2 &= - \sum_j \partial_i^2 p_j(x) (x_j - p_j(x)) \\ &\quad + \sum_j (\delta_{ij} - \partial_i p_j(x)) (\delta_{ij} - \partial_i p_j(x)). \end{aligned} \quad (33)$$

Consequently, we obtain

$$\begin{aligned} \Delta \frac{1}{2} |x - p(x)|^2 &= \sum_i \partial_i^2 \frac{1}{2} |x - p(x)|^2 \\ &= - \sum_{i,j} \partial_i^2 p_j(x) (x_j - p_j(x)) + \sum_{i,j} (\delta_{ij} - \partial_i p_j(x)) (\delta_{ij} - \partial_i p_j(x)) \\ &= - \sum_{i,j} \partial_i^2 p_j(x) (x_j - p_j(x)) + d - 2 \sum_i \partial_i p_i(x) + \sum_{i,j} \partial_i p_j(x) \partial_i p_j(x) \\ &= - \sum_{i,j} \partial_i^2 p_j(x) (x_j - p_j(x)) + d - 2 \operatorname{trace}(Dp(x)) + \operatorname{trace}(Dp(x) Dp(x)^T) \\ &= - \sum_{i,j} \partial_i^2 p_j(x) (x_j - p_j(x)) + d + \operatorname{trace}(Dp(x) (Dp(x)^T - 2\mathbb{I})). \end{aligned}$$

It also holds  $0 \preceq Dp(x) \preceq \mathbb{I}$  and therefore  $2\mathbb{I} - Dp(x) \succeq \mathbb{I}$ . This allows us to estimate  $\operatorname{trace}(Dp(x) (Dp(x)^T - 2\mathbb{I})) = -\operatorname{trace}(Dp(x) (2\mathbb{I} - Dp(x))) \leq -\operatorname{trace}(Dp(x)) \leq 0$ .

Going back to the previous formula for the Laplacian, we obtain the estimate

$$\Delta \frac{1}{2} |x - p(x)|^2 \leq - \sum_{i,j} \partial_i^2 p_j(x) (x_j - p_j(x)) + d \quad (34)$$

and it remains to bound the first term. For this we need second derivatives of  $p(x)$ . Writing (30) in coordinates gives

$$\partial_i p_j(x) = \delta_{ij} - \kappa^2 \sum_r \partial_r \partial_j V(p(x)) \partial_i p_r(x).$$

Taking another derivative with respect to the  $i$ th variable and using the product rule yields

$$\partial_i^2 p_j(x) = -\kappa^2 \sum_r \partial_r \partial_j V(p(x)) \partial_i^2 p_r(x) - \kappa^2 \sum_{r,s} \partial_s \partial_r \partial_j V(p(x)) \partial_i p_s(x) \partial_j p_r(x).$$

We have to solve this equation for the second derivatives of  $p$  for which we define the matrix  $A = A_{ij} := \partial_i^2 p_j(x)$  and the matrix  $B = B_{ij} :=$

$\sum_{r,s} \partial_s \partial_r \partial_j V(p(x)) \partial_i p_s(x) \partial_j p_r(x)$ . Then the previous equation is equivalent to the linear system

$$A = -\kappa^2 A D^2 V(p(x)) - \kappa^2 B$$

which is solved by

$$A = -\kappa^2 B (\mathbb{I} + \kappa^2 D^2 V(p(x)))^{-1}. \quad (35)$$

Using the definition of  $A$  and (34) we get

$$\Delta \frac{1}{2} |x - p(x)|^2 \leq -\kappa^2 \sum_{i,j} A_{ij} \partial_j V(p(x)) + d$$

and it remains to uniformly bound the first term.

Using (35), the definition of  $B$  as well as (31), we get the following estimate in terms of the matrix/tensor norms

$$\Delta \frac{1}{2} |x - p(x)|^2 \leq C_d \kappa^4 \left| D^3 V(p(x)) \right| \left| \nabla V(p(x)) \right| \left| \left( \mathbb{I} + \kappa^2 D^2 V(p(x)) \right)^{-1} \right|^3 + d,$$

where  $C_d$  is a dimensional constant. By assumption the right hand side is uniformly bounded by some constant  $C > 0$  and going back to (27) we obtain

$$\begin{aligned} \frac{d}{dt} \mathbb{L}[\rho_t] &\leq -\frac{\kappa^2 \mu}{1 + \kappa^2 \mu} \int |x - p(x)|^2 d\rho_t(x) + \frac{\sigma^2}{2} C \int |x - p(x)|^2 d\rho_t(x) \\ &= -\left( \frac{2\kappa^2 \mu}{1 + \kappa^2 \mu} - \sigma^2 C \right) \mathbb{L}[\rho_t]. \end{aligned}$$

Since  $\mu > 0$ , we can choose  $\sigma > 0$  sufficiently small—for instance,  $\sigma^2 < \frac{1}{C} \frac{\kappa^2 \mu}{1 + \kappa^2 \mu}$ —such that the brackets are strictly positive. Then we can conclude the proof using Gronwall's inequality.  $\square$

**Example 1** (The one-dimensional case) In the case of one spatial dimension  $d = 1$ , we can bound the Laplacian  $\Delta \frac{1}{2} |x - p(x)|^2$ —which in this case is just the second derivative—more accurately. In this case  $B = V'''(p(x))(p'(x))^2$  and hence  $A = -\kappa^2 \frac{V'''(p(x))(p'(x))^2}{1 + \kappa^2 V''(p(x))}$ . Plugging in  $p'(x) = \frac{1}{1 + \kappa^2 V''(p(x))}$  we obtain

$$A = -\kappa^2 \frac{V'''(p(x))}{(1 + \kappa^2 V''(p(x)))^3}.$$

Hence, we obtain the estimate

$$\Delta \frac{1}{2} |x - p(x)|^2 \leq \kappa^4 \frac{V'''(p(x))V'(p(x))}{(1 + \kappa^2 V''(p(x)))^3} + 1.$$

We give three examples of loss functions  $V$  which satisfy all assumptions of Proposition 2. The first one is the quadratic loss  $V_1(x) = ax^2 + bx + c$  with  $a > 0, b, c \in \mathbb{R}$  for which it holds  $V_1''' = 0$  and hence the Laplacian is bounded by 1. The second one is  $V_2(x) = x^2 + \ln(x + \sqrt{1+x^2})$  for which we have (the loose bound)  $|V'''V'| \leq 1$  and so a valid bound for the Laplacian is  $\kappa^4 + 1$ . The third loss is  $V_3(x) = x^2 + x + 1 - (x+1)\ln(x+1)$  for  $x \geq 0$ , extended to an even function on  $\mathbb{R}$ . Also here one has (the loose) bound  $|V'''V'| \leq 1$  and again a valid bound for the Laplacian is  $\kappa^4 + 1$ .

Next we prove a compactness property for measures for which  $L[\rho_t]$  is uniformly bounded in time. For this we also require that the loss  $V$  has a Lipschitz continuous gradient, which is a standard assumption in nonlinear optimization.

**Proposition 3** *Let  $V \in C^2(\mathbb{R}^d)$  satisfy  $\mu\mathbb{I} \preceq D^2V \preceq L\mathbb{I}$  for some  $0 < \mu \leq L < \infty$  and assume that*

$$\sup_{t>0} L[\rho_t] < \infty.$$

*Then there exists a subsequence  $(\rho_{t_n})_{n \in \mathbb{N}}$  for  $t_n \rightarrow \infty$  as  $n \rightarrow \infty$  which converges to a probability measure  $\rho_\infty$  in the Wasserstein-2 distance.*

**Proof** Let  $x^* := \arg \min V$ . By assumption it holds

$$\frac{\mu}{2} |p(x) - x^*|^2 + \langle \nabla V(p(x)), p(x) - x^* \rangle + V(p(x)) \leq V(x^*) \leq V(p(x))$$

which implies

$$\frac{\mu\kappa^2}{2} |p(x) - x^*|^2 \leq \kappa^2 \langle \nabla V(p(x)), x^* - p(x) \rangle.$$

By definition of  $p(x)$  it holds  $\kappa^2 \nabla V(p(x)) = x - p(x)$  and therefore

$$\frac{\mu\kappa^2}{2} |p(x) - x^*|^2 \leq \langle x - p(x), x^* - p(x) \rangle \leq |x - p(x)| |x^* - p(x)|.$$

Hence, we have showed that

$$|p(x) - x^*| \leq \frac{2}{\mu\kappa^2} |x - p(x)|.$$

Using this as well as the inequality  $|a + b|^2 \leq 2|a|^2 + 2|b|^2$  we obtain

$$\begin{aligned} \int |p(x)|^2 d\rho_t(x) &\leq 2 \int |x^* - p(x)|^2 d\rho_t(x) + 2 |x^*|^2 \\ &\leq \frac{8}{\mu^2\kappa^4} \int |x - p(x)|^2 d\rho_t(x) + 2 |x^*|^2. \end{aligned} \quad (36)$$

By assumption one has

$$\sup_{t>0} \int |x - p(x)|^2 d\rho_t(x) < \infty$$

which, together with change of variables and (36), implies

$$\sup_{t>0} \int |y|^2 d(p_{\#}\rho_t)(y) = \sup_{t>0} \int |p(x)|^2 d\rho_t(x) < \infty.$$

Therefore, by compactness a subsequence  $(p_{\#}\rho_{t_n})_{n \in \mathbb{N}}$  of the push-forward measures converges to some probability measure  $\sigma_{\infty}$  in the Wasserstein-2 distance as  $n \rightarrow \infty$ .

First, we claim that under our assumptions the map  $x \mapsto p(x)$  is a biLipschitz homeomorphism with Lipschitz continuous inverse  $y \mapsto p^{-1}(y) := y + \kappa^2 \nabla V(y)$ . The fact that these two maps are Lipschitz continuous is obvious, noting that the proximal operator of a convex function is 1-Lipschitz and the map  $p^{-1}$  is Lipschitz because  $\nabla V$  is Lipschitz. Furthermore, we observe that

$$p^{-1}(p(x)) = p(x) + \kappa^2 \nabla V(p(x)) = x$$

by definition of the proximal operator so  $p^{-1}$  is a left inverse. To see that also  $p(p^{-1}(x)) = x$  we note that

$$p(p^{-1}(x)) = \arg \min_{y \in \mathbb{R}^d} \frac{|y - x - \kappa^2 \nabla V(x)|^2}{2\kappa^2} + V(y).$$

This is a strongly convex optimization problem and hence the optimality conditions

$$0 = y - x - \kappa^2 \nabla V(x) + \kappa^2 \nabla V(y)$$

are necessary and sufficient. They have the (unique) solution  $y = x$  and therefore  $p(p^{-1}(x)) = x$  which shows that  $p^{-1}$  is truly the inverse of  $p$ .

Next, we claim that the measures  $\rho_{t_n}$  converge to  $\rho_{\infty} := p_{\#}^{-1}\sigma_{\infty}$  in the Wasserstein-2 distance as  $n \rightarrow \infty$ . To see this, let  $\pi_n$  denote an optimal coupling of  $p_{\#}\rho_{t_n}$  and  $\sigma_{\infty}$  such that

$$W_2(p_{\#}\rho_{t_n}, \sigma_{\infty}) = \iint |x - y|^2 d\pi_n(x, y).$$

We define  $\tilde{\pi}_n := (p^{-1} \times p^{-1})_{\#}\pi_n$  and first argue that this a coupling of  $\rho_{t_n}$  and  $\rho_{\infty}$ . To see this, we compute

$$\tilde{\pi}_n(A \times \mathbb{R}^d) = \pi_n(p(A) \times p(\mathbb{R}^d)) = \pi_n(p(A) \times \mathbb{R}^d) = p_{\#}\rho_{t_n}(p(A)) = \rho_{t_n}(A)$$

for every Borel set  $A \subset \mathbb{R}^d$ , where we used that  $p$  is invertible and that the first marginal of  $\pi_n$  is  $p_{\#}\rho_{t_n}$ . Hence, the first marginal of  $\tilde{\pi}_n$  is  $\rho_{t_n}$ . Similarly, the second marginal can be shown to be  $\rho_{\infty}$  and hence  $\pi_n$  is a coupling. We therefore obtain

$$\begin{aligned} W_2(\rho_{t_n}, \rho_{\infty}) &\leq \iint |x - y|^2 \, d\tilde{\pi}_n(x, y) \\ &= \iint \left| p^{-1}(x) - p^{-1}(y) \right|^2 \, d\pi_n(x, y) \\ &\leq (1 + \kappa^2 L)^2 \iint |x - y|^2 \, d\pi_n(x, y) \\ &= W_2(p_{\#}\rho_{t_n}, \sigma_{\infty}) \rightarrow 0, \quad \text{as } n \rightarrow \infty, \end{aligned}$$

where we used the Lipschitzness of  $\nabla V$  which implies Lipschitzness of  $p^{-1}$ .  $\square$

**Theorem 4** *Under the conditions of Propositions 2 and 3 and assuming  $L[\rho_0] = 0$  it holds that*

$$W_2(\rho_t, \delta_{x^*}) \rightarrow 0 \quad \text{as } t \rightarrow \infty,$$

where  $x^* := \arg \min V$ .

**Proof** By Proposition 2 it holds  $L[\rho_t] \rightarrow 0$  and, in particular,  $t \mapsto L[\rho_t]$  is uniformly bounded. Hence, one can apply Proposition 3 to obtain that a subsequence of  $\rho_t$  converges to some probability measure  $\rho_{\infty}$  in the Wasserstein-2 distance and hence also weakly. Since  $\rho \mapsto L[\rho]$  is an integral of the continuous and lower-bounded function  $x \mapsto \frac{1}{2} |x - p(x)|^2$  against  $\rho$ , it is lower semicontinuous with respect to weak convergence of measures. It follows  $L[\rho_{\infty}] = 0$  and hence  $\rho_{\infty} = \delta_{x^*}$ , where  $x^*$  is the global minimizer of  $V$ . The uniqueness of the minimizer implies that the whole sequence  $\rho_t$  converges to  $\delta_{x^*}$  and this concludes the proof.  $\square$


### 3 Numerical examples

In this section we evaluate the numerical performance of the proposed algorithms. In all our experiments we chose a time step parameter of  $dt = 0.01$ . The code to reproduce all numerical experiments is available on [GitHub](https://github.com/TimRoith/PolarCBO).<sup>1</sup>


#### 3.1 Unimodal Ackley function

In the first example we perform a consistency check for our method for finding the unique global minimum of the Ackley function [26], defined as

<sup>1</sup> <https://github.com/TimRoith/PolarCBO>, note that this repository is now deprecated in favor of the package CBXPy [27].



**Fig. 2** Dynamics of standard CBO for minimizing the Ackley function. The points mark particle locations, the arrows the drift field towards the shared weighted mean



**Fig. 3** Dynamics of the proposed polarized CBO for minimizing the Ackley function. The points mark particle locations, the arrows the drift field towards the individual weighted means


$$A(x) := -20 \exp\left(-\frac{0.2}{\sqrt{d}} |x|\right) - \exp\left(\frac{1}{d} \sum_{n=1}^d \cos(2\pi x_n)\right) + e + 20. \quad (37)$$

This function has a global minimum at  $0 \in \mathbb{R}^d$  with  $A(0) = 0$  and in this experiment we choose  $d = 2$  and minimize the shifted version  $V(x) := A(x - (3, 2))$  which has its global minimum at  $(3, 2) \in \mathbb{R}^2$ . We compare the dynamics of standard CBO with our proposed polarized variant at three different time points in Figs. 2 and 3. We use the Gaussian kernel  $k(x, y) := \exp\left(-\frac{|x-y|^2}{2\kappa^2}\right)$  with standard deviation  $\kappa = \infty$  for standard CBO and  $\kappa = 1$  for polarized CBO. Furthermore, we choose  $\beta = 1$ . In this simple situation with a unique global minimum we observe that both standard and polarized CBO find the global minimum and do not get stuck in local minima. Notably, the polarized variant converges slightly slower than standard CBO which is due to the localization effect of the kernel with a relatively small standard deviation.

### 3.2 Multimodal Rastrigin function

In this example we evaluate different choices of kernel functions for minimizing a Rastrigin-type function with three global minima. The original Rastrigin function

**Fig. 4** A variant of the Rastrigin function with three global minima



[27] on  $\mathbb{R}^d$  is defined as

$$R(x) := 10d + \sum_{n=1}^d (x_n^2 - 10 \cos(2\pi x_n)) \quad (38)$$

and has a global minimum at  $x = 0$  with  $U(0) = 0$ . In this experiment we choose  $d = 2$  and minimize the product

$$V(x) = \frac{1}{8} R(x) R(x - (3, 2)) R(x + (1, 3.5))$$

which has three global minima. Note that this function is very non-convex and at the same time extremely flat around its minima, see Fig. 4 for a surface plot.

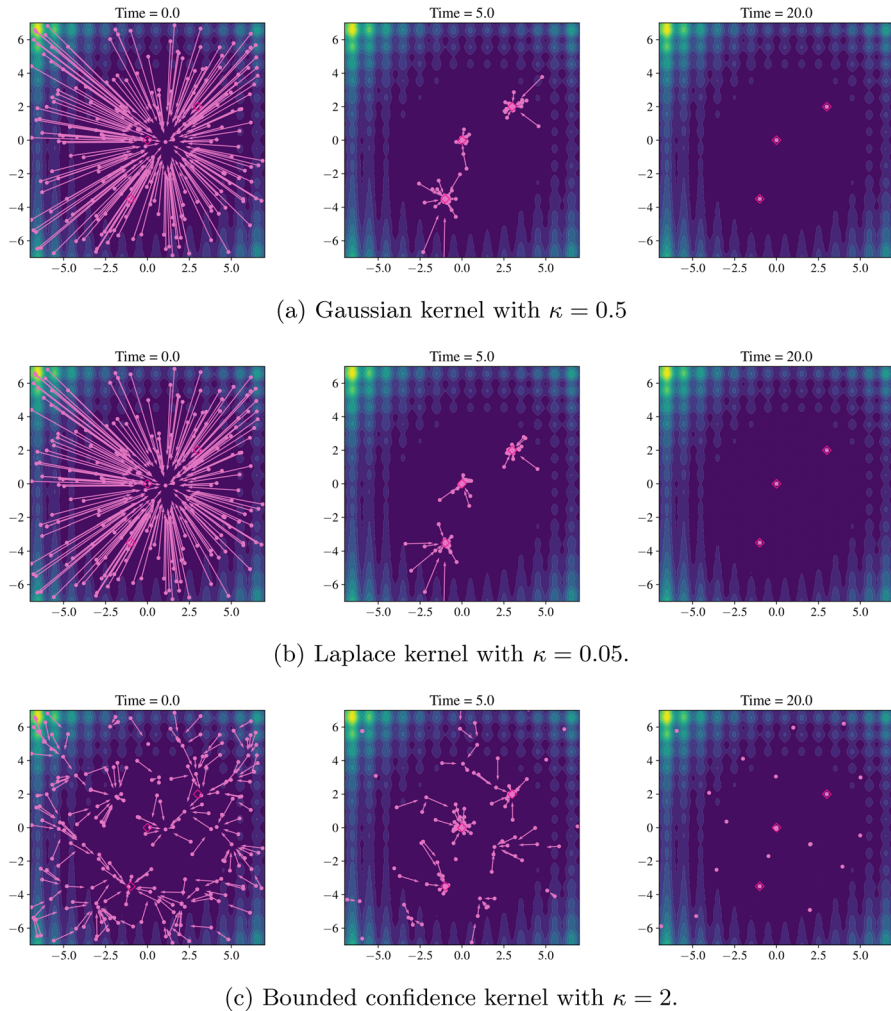
We consider the three different kernels

$$k(x, y) = \exp\left(-\frac{|x - y|^2}{2\kappa^2}\right) \quad \text{Gaussian kernel,} \quad (39)$$

$$k(x, y) = \exp\left(-\frac{|x - y|}{\kappa}\right) \quad \text{Laplace kernel,} \quad (40)$$

$$k(x, y) = 1_{|x - y| \leq \kappa}(x, y) \quad \text{bounded confidence kernel,} \quad (41)$$

and the corresponding results of our method are depicted in Fig. 5. Again, we choose  $\beta = 1$ . The kernel parameters  $\kappa$  were chosen sufficiently small for the methods to detect all three minima, which are marked with blue diamonds. While the Gaussian and the Laplace kernel work similarly well—note however that the Laplace kernel needs a much smaller value of  $\kappa$  than the Gaussian—the bounded confidence kernel (cf. the discussion in Sect. 2.1) works suboptimally for the task of minimization. While it manages to detect all three minima a lot of particles get stuck in suboptimal consensus points which can be explained by the fact that the kernel has compact support.



**Fig. 5** Dynamics of polarized CBO with different kernels for minimizing a Rastrigin-type function with three global minima, marked by red diamonds. For all kernels all minima were detected. Gaussian and Laplace kernel work especially well, whereas the bounded confidence kernel generates too many consensus points due to its compact support

### 3.3 Multimodal Ackley function

We proceed with quantitative evaluations of our method for a multimodal version of the Ackley function (37), defined as

$$V(x) := \prod_{i=1}^N A(x - z_i),$$



where  $\{z_i \in \mathbb{R}^d : i = 1, \dots, N\}$  are points which constitute the global minimizers of  $V$ , and  $U$  is the standard Ackley function defined in (37). In Table 1 we plot how many of the three minima in dimension  $d = 2$  were detected by the proposed polarized CBO method with a Gaussian kernel and different values of the standard deviation  $\kappa$ . That means, we show the percentage of runs that detected at least 1, 2, or 3 minima. Here, we employed the standard noise model as specified in (5) with  $\sigma = 1.0$  and  $\beta = 1.0$ . For completeness we also include results for standard CBO (i.e.,  $\kappa = \infty$ ) which by definition can detect at most one minimum.

We say that a method detects a minimum if at convergence there exists a weighted mean  $m_{\beta, \kappa}[\rho](x)$  which is closer than 0.25 in the infinity norm to the minimum. For standard CBO where  $m_{\beta, \kappa}[\rho](x) = m_\beta[\rho]$  for all  $x$  this coincides with the definition of success in [4, 15]. For our experiments we employ  $N = 3$  different minima  $z_1, z_2, z_3 \in \mathbb{R}^d$  with

$$\begin{aligned} (z_1)_i &:= \begin{cases} -2 & \text{if } i \bmod 2 = 0, \\ 1 & \text{else,} \end{cases} & (z_2)_i &:= \begin{cases} 2 & \text{if } i \bmod 2 = 0, \\ -1 & \text{else,} \end{cases} \\ (z_3)_i &:= \begin{cases} -1 & \text{if } i \bmod 2 = 0, \\ -3 & \text{else} \end{cases} \end{aligned}$$

for  $i = 1, \dots, d$ . While polarized CBO works very well in the two-dimensional example, Table 2 shows that in dimension  $d = 10$ , it fails to detect more than one minimum. However, the cluster method from 3 exhibits significantly improved behavior and manages to detect all three minima frequently. Here, we employed the coordinate-wise noise model from (9) with  $\sigma = 7.5$ . Additionally we employed a simple scheduling for the parameter  $\beta$ , where we start with  $\beta = 30$  and increase it in each step via

$$\beta \leftarrow 1.01 \cdot \beta$$

up to a limit of  $\beta_{\max} = 10^7$ . Here, one could potentially employ more sophisticated approaches as proposed in [15]. For the cluster-based methods we chose  $\alpha = 5.0$  in Algorithm 3.

Even for  $\kappa = \infty$ , where the kernel has no influence anymore, the method works very well. Note that this case does not correspond to standard CBO. Furthermore, Table 2 shows that our polarized method can outperform standard CBO at finding at least one minimum, albeit at the cost of higher complexity.

We also test the cluster method in  $d = 30$  dimensions and the results can be found in Table 3. We observe that it is harder to find multiple minima, however for  $J = 1600$  the method finds at least two minima in over 50% of the runs for  $\kappa \geq 1$ . For smaller particle numbers the algorithm performs better for  $\kappa \leq 1$ , although the percentage of runs, where multiple minima are found is very low.

**Table 1** Performance of polarized CBO for minimizing the multimodal Ackley function with  $d = 2$  and three global minima: Averaging over 100 independent runs of 1000 iterations we plot how many percent of the succeeded in detecting at least 1, 2, or 3 of the minima

# minima	$J = 25$			$J = 50$			$J = 100$			$J = 200$		
	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$
Polarized												
$\kappa = 0.1$	33%	7%	0%	86%	59%	24%	100%	96%	67%	100%	100%	97%
$\kappa = 0.5$	100%	62%	5%	100%	78%	18%	100%	93%	41%	100%	100%	84%
$\kappa = 1$	100%	5%	0%	100%	12%	0%	100%	14%	0%	100%	24%	0%
CBO	100%	0%	0%	100%	0%	0%	100%	0%	0%	100%	0%	0%

Note that, by definition, standard CBO ( $\kappa = \infty$ ) can detect at most one minimum

**Table 2** Performance of polarized and cluster CBO for minimizing the multimodal Ackley function with  $d = 10$  and ceteris paribus

# minima	$J = 50$			$J = 100$			$J = 200$			$J = 400$		
	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$
Polarized												
$\kappa = 0.001$	5%	0%	0%	18%	0%	0%	26%	0%	0%	63%	1%	0%
$\kappa = 0.01$	26%	0%	0%	56%	0%	0%	80%	1%	0%	79%	3%	0%
$\kappa = 0.1$	36%	0%	0%	68%	0%	0%	73%	0%	0%	75%	0%	0%
CBO	32%	0%	0%	55%	0%	0%	74%	0%	0%	72%	0%	0%
Cluster												
$\kappa = 10^{-7}$	26%	2%	0%	75%	26%	2%	98%	59%	13%	100%	89%	28%
$\kappa = 0.1$	11%	1%	0%	68%	13%	0%	98%	77%	19%	100%	96%	39%
$\kappa = \infty$	6%	0%	0%	65%	11%	0%	98%	73%	15%	100%	92%	41%

The first method is slightly better than CBO but fails to detect more than one minimum in most cases. The cluster variant works well for a wide range of parameters  $\kappa$ . It manages to detect all minima in some runs, already for  $J = 200$  particles. Note that  $\kappa = \infty$  *does not* correspond to CBO for this method

**Table 3** Performance of cluster CBO for minimizing the multimodal Ackley function with  $d = 30$  and ceteris paribus


# minima	$J = 200$			$J = 400$			$J = 800$			$J = 1600$		
	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$	$\geq 1$	$\geq 2$	$\geq 3$
Cluster												
$\kappa = 0.01$	15%	0%	0%	61%	5%	0%	92%	28%	1%	99%	45%	1%
$\kappa = 0.1$	9%	0%	0%	59%	6%	0%	94%	34%	2%	98%	49%	2%
$\kappa = 1$	4%	0%	0%	57%	5%	0%	88%	40%	0%	97%	57%	4%
$\kappa = 10$	5%	0%	0%	58%	7%	0%	89%	35%	1%	98%	58%	1%
$\kappa = 100$	4%	0%	0%	59%	7%	0%	89%	32%	0%	98%	57%	2%
$\kappa = \infty$	5%	0%	0%	58%	6%	0%	88%	33%	0%	98%	54%	2%

### 3.4 Multimodal sampling

In this section we consider the task of sampling from a bimodal mixture of Gaussians.

$$\exp(-V(x)) := \exp\left(-\frac{(x_1 - a_1)^2}{0.2} - \frac{(x_2 - a_2)^2}{0.2}\right) + \frac{1}{2} \exp\left(-\frac{(x_1 - b_1)^2}{8} - \frac{(x_2 - b_2)^2}{0.5}\right),$$

where the tuples  $a = (a_1, a_2)$  and  $b = (b_1, b_2)$  determine the centers of the clusters. In Fig. 6 we plot the result of standard CBS and our polarized variant using Gaussian kernels with three different standard deviations  $\kappa \in \{0.4, 0.6, 0.8\}$ . For both methods we choose  $\beta = 1$ . Note that, at least for convex potentials  $V$  with bounded Hessian,



**Fig. 6** Dynamics of standard CBS and our polarized version, sampling from a mixture of Gaussians (top row: far apart, bottom row: closer together). The points mark particle locations, the arrows the drift field towards the weighted mean(s)


standard CBS is known to exhibit a Gaussian steady state. Being designed as a method for unimodal sampling, there is not much hope CBS can work in this multimodal situation. Still we include CBS results for comparison.

In contrast, our polarized modification of CBS manages to isolate the two modes. Note that in Proposition 1 we proved that polarized CBS with a Gaussian kernel is unbiased when the target measure is a Gaussian. We do not expect this to be true for target measures which are a mixture of Gaussians. Still, our results in the first row of Fig. 6 show that, if the two Gaussians are sufficiently far apart, our method (second to fourth column) seems close to being unbiased. Standard CBS (left column) can, by design, find at most one mode and here it successfully detects the lower Gaussian. Note that we use the same number of particles for both algorithms, namely  $J = 400$ , which is why the bottom mode for CBS looks more densely sampled.

The situation is different when the two modes are closer together, as shown in the bottom row of Fig. 6. Here, standard CBS fails to detect even one mode whereas our polarized version does achieve that. However, as expected in this multimodal situation the result is not a perfect sample but appears to be biased. Furthermore, if the clusters are close, the sensitivity with respect to the choice of the kernel width  $\kappa$  is larger and  $\kappa$  has to be chosen sufficiently small in order to generate enough samples for the lower cluster.

### 3.5 Non-Gaussian sampling

Our final experiment deals with sampling from a non-Gaussian distribution. Again, we use  $\beta = 1$ . According to the theoretical analysis in [15] standard CBS is not expected



**Fig. 7** Standard and polarized CBS for sampling from a non-Gaussian distribution. The orange crosses mark the position of the weighted means

to correctly sample from such a distribution but rather from a Gaussian approximation. While we do not claim that our polarized method generates an exact sample, our result in Fig. 7 show that polarized CBS approximates the non-Gaussian distribution much better than standard CBS. This is due to the fact that the weighted means, indicated as orange crosses, do not collapse to a single point but concentrate in the region of high probability mass.

## 4 Conclusion and outlook

In this article we presented a polarized version of consensus-based optimization and sampling dynamics for objectives with multiple global minima or modes. For this we localized the dynamics such that every particle is primarily influenced by close-by particles. We proved that in the case of sampling from a Gaussian distribution this does not introduce a bias. We also suggested a cluster-based version of our polarized dynamics which is computationally more efficient. Our extensive numerical experiments suggested a large potential of our method for detecting multiple global minima or modes, improving over standard consensus-based methods.

There is a lot of room for future work regarding well-definedness of the Fokker–Planck equation derived above, and stability and convergence of both the mean-field and particle system to consensus using less restrictive assumptions than the ones used in this paper. We are convinced that the Lyapunov function  $L[\rho_t]$  which we studied in Sect. 2.4 will be very helpful in this endeavour. Future work will also focus on further numerical improvements of our method, in particular, incorporating a batching strategy similar to the one from [10] to further improve performance in high-dimensional optimization. Finally, a long-term goal will be to find multimodal sampling methods which are provably consistent for multimodal and non-Gaussian distributions. Given that even gradient-free sampling from unimodal non-Gaussian distributions is still relatively poorly understood, with [13] being a promising approach, this will be a challenging task for future work.

**Acknowledgements** Part of this work was also done while LB and TR were in residence at Institut Mittag-Leffler in Djursholm, Sweden during the semester on *Geometric Aspects of Nonlinear Partial Differential Equations* in 2022, supported by the Swedish Research Council under grant no. 2016-06596. LB is very thankful to Massimo Fornasier and Oliver Tse for insightful discussions on the analytical aspects of this paper which took place at the Lorentz Center in Leiden during the workshop “Purpose-driven particle systems” in Spring 2023. Part of this work was done while LB was affiliated with the University of Bonn and the Technical University of Berlin. PW acknowledges support from MATH+ project EF1-19: Machine Learning Enhanced Filtering Methods for Inverse Problems, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689). TR acknowledges support by the German Ministry of Science and Technology (BMBF) under grant agreement No. 05M2020 (DELETO) and was supported by the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 777826 (NoMADS). Part of this work was done while TR was affiliated with the Friedrich-Alexander Universität Erlangen-Nürnberg. TR acknowledges support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF. This research was supported in part through the Maxwell computational resources operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

## Declarations

**Conflict of interest** On behalf of all authors, the corresponding author declares that there are no conflict of interest.

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

## References

- Conn, A.R., Scheinberg, K., Vicente, L.N.: Introduction to Derivative-free Optimization. SIAM, Philadelphia, PA (2009)
- Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of ICNN’95 - International Conference on Neural Networks, vol. 4, pp. 1942–1948 (1995)
- Reynolds, C.W.: Flocks, herds and schools: A distributed behavioral model. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, pp. 25–34 (1987)
- Pinnau, R., Totzeck, C., Tse, O., Martin, S.: A consensus-based model for global optimization and its mean-field limit. *Math. Models Methods Appl. Sci.* **27**(01), 183–204 (2017)
- Carrillo, J.A., Choi, Y.-P., Totzeck, C., Tse, O.: An analytical framework for consensus-based global optimization method. *Math. Models Methods Appl. Sci.* **28**(06), 1037–1066 (2018)
- Ha, S.-Y., Jin, S., Kim, D.: Convergence of a first-order consensus-based global optimization algorithm. *Math. Models Methods Appl. Sci.* **30**(12), 2417–2444 (2020)
- Fornasier, M., Klock, T., Riedl, K.: Consensus-based optimization methods converge globally in mean-field law (2021). [arXiv:2103.15130](https://arxiv.org/abs/2103.15130)
- Totzeck, C., Wolfram, M.-T.: Consensus-based global optimization with personal best. *Math. Biosci. Eng.* **17**(5), 6026–6044 (2020). <https://doi.org/10.3934/mbe.2020320>
- Carrillo, J.A., Totzeck, C., Vaes, U.: Consensus-based optimization and ensemble kalman inversion for global optimization problems with constraints (2021). [arXiv:2111.02970](https://arxiv.org/abs/2111.02970)

10. Carrillo, J.A., Jin, S., Li, L., Zhu, Y.: A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM: Control Optim. Calculus of Varia.* **27**, 5 (2021)
11. Fornasier, M., Huang, H., Pareschi, L., Sünnen, P.: Consensus-based optimization on hypersurfaces: well-posedness and mean-field limit. *Math. Models Methods Appl. Sci.* **30**(14), 2725–2751 (2020)
12. Fornasier, M., Pareschi, L., Huang, H., Sünnen, P.: Consensus-based optimization on the sphere: convergence to global minimizers and machine learning. *J. Mach. Learn. Res.* **22**(237), 1–55 (2021)
13. Schillings, C., Totzeck, C., Wacker, P.: Ensemble-based gradient inference for particle methods in optimization and sampling (2022). [arXiv:2209.15420](https://arxiv.org/abs/2209.15420)
14. Totzeck, C.: Trends in consensus-based optimization. In: *Active Particles*, volume 3, pp. 201–226. Springer, Cham (2022)
15. Carrillo, J.A., Hoffmann, F., Stuart, A.M., Vaes, U.: Consensus-based sampling. *Stud. Appl. Math.* **148**(3), 1069–1140 (2022)
16. Deffuant, G., Neau, D., Amblard, F., Weisbuch, G.: Mixing beliefs among interacting agents. *Adv. Complex Syst.* **3**(01–04), 87–98 (2000)
17. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* **5**(3) (2002)
18. Gómez-Serrano, J., Graham, C., Le Boudec, J.-Y.: The bounded confidence model of opinion dynamics. *Math. Models Methods Appl. Sci.* **22**(02), 1150007 (2012)
19. Fortunato, S., Latora, V., Pluchino, A., Rapisarda, A.: Vector opinion dynamics in a bounded confidence consensus model. *Int. J. Mod. Phys. C* **16**(10), 1535–1551 (2005)
20. Burger, M.: Kinetic equations for processes on co-evolving networks. *Kinetic Relat. Models* **15**(2), 187–212 (2022). <https://doi.org/10.3934/krm.2021051>
21. Burger, M.: Network structured kinetic models of social interactions. *Vietnam J. Math.* **49**(3), 937–956 (2021)
22. Schnell, P.: Eine methode zur auffindung von gruppen. *Biom. Z.* **6**(1), 47–48 (1964)
23. Fukunaga, K., Hostetler, L.: The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans. Inf. Theory* **21**(1), 32–40 (1975)
24. Petersen, K.B., Pedersen, M.S., et al.: The matrix cookbook. *Tech. Univ. Denmark* **7**(15), 510 (2008)
25. Bailo, R., Barbaro, A., Gomes, S. N., Riedl, K., Roith, T., Totzeck, C., Vaes, U. (2024). CBX: Python and Julia packages for consensus-based interacting particle methods. [arXiv preprint arXiv:2403.14470](https://arxiv.org/abs/2403.14470)
26. Ackley, D.: A Connectionist Machine for Genetic Hillclimbing. The Springer International Series in Engineering and Computer Science. Springer, New York, NY (2012)
27. Rastrigin, L.A.: Systems of extremal control. *Nauka* (1974)

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