

Consensus-based Optimization and Ensemble Kalman Inversion for Global Optimization Problems with Constraints

J. A. Carrillo* C. Totzeck† U. Vaes‡

November 4, 2021

Abstract

We introduce a practical method for incorporating equality and inequality constraints in global optimization methods based on stochastic interacting particle systems, specifically consensus-based optimization (CBO) and ensemble Kalman inversion (EKI). Unlike other approaches in the literature, the method we propose does not constrain the dynamics to the feasible region of the state space at all times; the particles evolve in the full space, but are attracted towards the feasible set by means of a penalization term added to the objective function and, in the case of CBO, an additional relaxation drift. We study the properties of the method through the associated mean-field Fokker–Planck equation and demonstrate its performance in numerical experiments on several test problems.

1 Introduction

Finding efficiently the global minimizer of a function in large dimensions is crucial in today’s data science activity, with applications in various fields of research ranging from modern areas such as machine learning, neural networks, inverse problems, data assimilation and model parameter estimation to more classical areas such as economics, computational biology, mathematical finance and statistical physics. In the most general form, the global optimization task is simply

$$\operatorname{argmin}_{x \in B} f(x) \tag{1.1}$$

for a given objective function f and state space $B \subset \mathbf{R}^d$. This simple-to-describe problem is highly nontrivial for general nonconvex functions f due to the existence of usually many, and in some cases even infinitely many, local minima. Most of the approaches have followed metaheuristic arguments in order to find robust algorithms; see for instance [1, 4, 49, 7, 19, 39] and the references therein.

*Mathematical Institute, University of Oxford, carrillo@maths.ox.ac.uk

†University of Wuppertal, School of Mathematics and Natural Sciences, totzeck@uni-wuppertal.de

‡MATERIALS team, Inria Paris, urbain.vaes@inria.fr

Among global optimization methods, several algorithms were inspired by physical or biological models using the idea of interacting particle/agent systems: ant colony optimization [16], artificial bee colony optimization [38], firefly optimization [55], wind driven optimization [5], and particle swarm optimization [40, 48]. All of these methods have in common the use of interactions of particles with nonlocal terms, the exchange of information among agents and the use of noise terms to explore the landscape of the function f ; see [6] for a review. Another family of global optimization algorithms is known as simulated annealing methods [30, 31]. The simplest examples of simulated annealing methods use the cost function f as the potential in an overdamped Langevin diffusion, which is solved with a time-dependent temperature [33, 43, 46] – a cooling schedule – so that the unique probability measure in the kernel of the Fokker–Planck operator associated with the dynamics at time t converges to a Dirac mass at the global minimizer in the limit as $t \rightarrow \infty$. Despite their simplicity, these methods usually have a slow convergence rate.

Concerning machine learning applications and large scale optimization, there have been a flurry of works in SGD methods, where stochasticity is again introduced as a way of improving exploration of the landscape. An essential ingredient of SGD methods [50, 9, 10] is the mini-batch strategy, also introduced recently in [37, 36] for interacting particle systems, which reduces the computational cost of the iterations. The SGD family of methods relies on the computation of gradients of functions in high-dimensional settings, which may be computationally expensive or even impossible, if the loss function is nondifferentiable. In addition, even when gradients can be calculated at a reasonable computational cost, they may fail to provide useful information on the overall behavior of the loss function if this function exhibits many small-scale oscillations. Gradient-free methods, which rely only on evaluations of the loss function, are therefore an attractive alternative to SGD methods in these settings. In this paper, we consider two particular classes of methods belonging to this category which received a lot of attention lately: the Consensus-Based Optimization (CBO) methods [47, 11, 14, 12], reviewed recently in [53], and methods based on the Ensemble Kalman Filter (EnKF) [17, 21, 34, 51, 25, 26, 13], which are mainly employed in the context of inverse problems but have also proved useful for machine learning tasks [41]. It is shown in [20] that gradient-free ensemble Kalman methods perform better than their gradient-based counterparts in noisy likelihood landscapes. Within the class of ensemble Kalman methods, we focus specifically on the Ensemble Kalman Inversion (EKI) gradient-free method, in the form proposed in [51].

While the EKI method as discussed in Section 3 relies on a nonlocal approximation of gradients, and can be viewed as an exact preconditioned gradient descent in the case of a quadratic objective function [25], the central mechanism of CBO is a relaxation drift towards the weighted average

$$m_f = \frac{\sum_{j=1}^J x^{(j)} e^{-\alpha f(x^{(j)})}}{\sum_{j=1}^J e^{-\alpha f(x^{(j)})}} =: \sum_{j=1}^J w(x^{(j)}) x^{(j)}. \quad (1.2)$$

Here, $x^{(j)}$ are the positions of a number J of explorers evolving in the landscape of the objective function f . As a metaphor for the CBO method, one may think of butterflies in the mountains able to communicate their height and position instantaneously by wireless communication and to

integrate this information in the form of the weighted average m_f , towards which they move with the hope of eventually reaching the global minimum of the landscape given by the function f .

The motivation for the CBO reweighting $e^{-\alpha f(x)}$, which is the Gibbs distribution corresponding to $f(x)$, comes from statistical mechanics: the probability distribution function $\frac{1}{Z_\alpha} e^{-\alpha f(x)}$, where Z_α is the normalization constant, is the unique invariant measure of overdamped Langevin dynamics at temperature α^{-1} in an external potential given by the cost function $f(x)$ [45, Chapter 6]. For sufficiently large α , the weighted average (1.2) provides a good approximation of the particle position $x^{(j^*)}$ minimizing $f(x^{(j)})$, assuming this is unique. Indeed, in this limit all the weights in (1.2) converge to 0, except for the one associated with particle j^* , which converges to 1.

Both the CBO and EKI methods are based on a system of interacting stochastic differential equations, and it is possible to show that their behavior is well described by nonlocal deterministic Fokker–Planck equations when the number of particles is sufficiently large. Studying these continuous equations brings considerable insight, and turns out to be much simpler than working with the particle systems. The mean-field Fokker–Planck equation associated with CBO includes a weighted average which, despite no longer being a finite sum as in (1.2), can still be understood in the limit as $\alpha \rightarrow \infty$. This is achieved through Laplace’s method [44, 18], a classical asymptotic method for integrals.

In its original form [47], the CBO method combines the relaxation drift towards the weighted average m_f with isotropic multiplicative noise proportional in amplitude to the Euclidean distance to the weighted average m_f . The specific form of the multiplicative factor of the noise, which may be viewed as a temperature scheduling, is improved in [12] in order to decrease the computational cost of the method for large dimensional problems; see Section 2 for details. In [12], the authors also demonstrate that approximating m_f based on a randomly selected subset of the particles – a mini-batch – leads to large computational savings. Using mini-batches also introduces extra stochasticity, which is observed to be useful for promoting exploration of the landscape given by f .

In this work, we propose new extensions of the CBO and EKI methods to constrained optimization and constrained inverse problems, respectively. In the case of CBO, we consider problems of the general form

$$\operatorname{argmin}_{x \in B} f(x), \quad \text{with } B = \{x : \mathcal{E}(x) = 0, \mathcal{I}(x) \geq 0\}, \quad (1.3)$$

where $\mathcal{E} : \mathbf{R}^d \rightarrow \mathbf{R}^{N_e}$ and $\mathcal{I} : \mathbf{R}^d \rightarrow \mathbf{R}^{N_i}$ are continuously differentiable functions and the inequality $\mathcal{I}(x) \geq 0$ is understood componentwise. In the case of EKI, we consider subsets of problems of the form (1.3) for which the objective function is of the specific form $f(x) = \frac{1}{2} |\Gamma^{-1}(G(x) - y)|^2$, where $G : \mathbf{R}^d \rightarrow \mathbf{R}^K$ is a map, $\Gamma \in \mathbf{R}^{K \times K}$ is a positive definite matrix, and $y \in \mathbf{R}^K$. In applications, objective functions of this form are used as a simple measure of the misfit between some data y and a model $G(x)$.

The literature on global optimization with constraints is abundant and an extensive review is beyond the scope of this paper, so we summarize hereafter only recent contributions that are specifically related to the CBO and EKI methods. Several recent works propose extensions of CBO

strategies to equality constrained problems, based on imposing that the dynamics is restricted to the feasible manifold at all times [24, 22, 23]. Most of these works consider the specific case of the Euclidean sphere as the feasible region, with applications to robust subspace detection and efficient eigenvalue computations in machine learning as prime examples. The evolution of particles is restricted to the feasible manifold by appropriate projection onto the tangent space. Although this method performs relatively well in very simple settings like the sphere, its implementation is difficult and error-prone for more general constraint manifolds. In addition, since the weighted average is computed in the ambient space, there are cases in which the method does not produce good results, for example when the feasible region is a closed hypersurface enclosing a nonconvex domain. Let us finally emphasize that manifolds other than the sphere are important in optimization problems in machine learning. SGD methods have been studied in Riemannian optimization applications, see for instance [56, 54] and the references therein. Here, the state space is usually a set of matrices with certain constraints, like positive definite matrices in Wasserstein barycenters or Riemannian centroids.

Several methods have also been proposed in the literature for incorporating constraints in EKI. In [2], the authors propose a generalization of EKI such that the iterates produced by the method are guaranteed to lie in the feasible region at each iteration. The method proposed leverages the fact that the update step in the usual ensemble Kalman method can be formulated as an optimization problem, in which linear constraints can be integrated. In [15], a variant of EKI incorporating a projection step is developed for the specific case of box constraints, and the continuous time limit of the method is studied. See also [28], where the continuous-time and mean field limits of the method proposed in [2] are studied.

In this work, we follow an orthogonal approach: we propose extensions of the CBO and EKI methods in which the particles evolve in the full space \mathbf{R}^d , instead of being constrained to the feasible manifold at all times. A requirement for this approach to work, which we take as a standing assumption in the rest of this paper, is that the objective function $f(x)$ can be evaluated anywhere in \mathbf{R}^d , which may or may not hold in applications. The unifying idea behind our extensions of CBO and EKI is the addition of a penalization to the objective function. In the case of only one equality constraint $\mathcal{E}(x) = 0$, for example, we seek a solution to

$$\operatorname{argmin}_{x \in \mathbf{R}^d} \left(f(x) + \frac{1}{\nu} |\mathcal{E}(x)|^2 \right). \quad (1.4)$$

When the penalization parameter ν tends to 0, we expect the solution of this problem to be close to the zero-level set of \mathcal{E} . We emphasize that the use of penalty functions for constrained optimization problems is a standard idea and can be employed, in principle, together with any method for unconstrained global optimization. In this paper, we demonstrate the effectiveness of this simple idea when used in conjunction with CBO and EKI, and we propose specific approaches for its implementation.

In the specific case of CBO, we also propose to use, in addition to a penalization of the form (1.3),

an extra drift term based on the constraint that imposes that particles are asymptotically confined to the manifold. This idea is reminiscent of swarming problems where the Vicsek model with noise on the sphere can be retrieved from the Cucker-Smale model on the whole space in an appropriate limit [8].

The rest of this paper is organized as follows. In [Sections 2](#) and [3](#), we explain our strategy for including constraints in the CBO and the EKI methods, respectively. In [Section 2](#), we also compare the solution to the constrained CBO method and associated the mean-field equations qualitatively for a toy problem in two dimensions, in order to illustrate the method and gain a deeper understanding of its behavior in the many-particle limit. In [Sections 4](#) and [5](#), we investigate the potential of the proposed methods by means of numerical experiments in some reference examples. To this end, we use typical benchmarks [35] in optimization. [Section 6](#) is reserved for conclusions and perspectives for future work.

2 CBO for constrained global optimization problems

In order to illustrate our strategy for incorporating constraints into the CBO methods and its variants, we consider the CBO scheme with component-wise Brownian motion proposed in [14], which reads as

$$dx_t^{(j)} = -(x_t^{(j)} - m_f) dt + \sqrt{2}\sigma(x_t^{(j)} - m_f) \circ dW_j^{(j)}, \quad j = 1, \dots, J, \quad (2.1)$$

where $\{W^{(j)}\}_{1 \leq j \leq J}$ are independent Wiener processes in \mathbf{R}^d and $v_1 \circ v_2$, for vectors v_1 and v_2 in \mathbf{R}^d , is the Hadamard (component-wise) product, that is to say $v_1 \circ v_2 = \text{diag}(v_1)v_2$. The particles are initialized independently according to some given probability density ϱ_0 , and so the law of the initial ensemble is $\varrho_0^{\otimes J}$. The main ingredient of the method, indeed the only mechanism by which the particles interact with the objective function, is the weighted mean m_f given by

$$m_f(\mu_t^J) = \frac{\int x e^{-\alpha f(x)} \mu_t^J(dx)}{\int e^{-\alpha f(x)} \mu_t^J(dx)}, \quad \mu_t^J = \frac{1}{J} \sum_{j=1}^J \delta_{x_t^{(j)}}, \quad (2.2)$$

where $\alpha > 0$ is a parameter. Equivalently,

$$m_f(\mu_t^J) = \frac{\sum_{j=1}^J x_t^{(j)} \exp(-\alpha f(x_t^{(j)}))}{\sum_{j=1}^J \exp(-\alpha f(x_t^{(j)}))}.$$

As mentioned in the introduction, for fixed t the weighted mean $m_f(\mu_t^J)$ in (2.2) converges to the global minimizer of f constrained to the support of the empirical measure μ_t^J , provided this minimizer is unique, in the limit as α tends to infinity.

Assume the problem we aim to solve is (1.1). In order to illustrate the incorporation of equality

and inequality constraints in a unified manner, we introduce $\mathcal{A}: \mathbf{R}^d \rightarrow \mathbf{R}^{N_e+N_i}$ given by

$$(\mathcal{A}(x))_i = \begin{cases} (\mathcal{E}(x))_i & \text{if } i \leq N_e, \\ \min\{(\mathcal{I}(x))_{i-N_e}, 0\} & \text{if } i > N_e. \end{cases}$$

We can then write $B = \{x \in \mathbf{R}^d : \mathcal{A}(x) = 0\}$; that is, we can assume without loss of generality that all the constraints are of equality type. Notice that \mathcal{A} may not be C^1 at the manifold even though \mathcal{E} and \mathcal{I} are, but this is not an issue given that only $\nabla |\mathcal{A}|^2$ appears in the method we propose, and the function $x \mapsto |\mathcal{A}(x)|^2$ is continuously differentiable. To solve the optimization problem (1.1) with constraint $\mathcal{A}(x) = 0$, we modify the scheme (2.1) in the following manner:

1. The first modification is a penalization of the objective function in the weighted average. More precisely, we substitute m_f in (2.1) by

$$m_g(\mu_t^J) = \frac{\int x e^{-\alpha g(x)} \mu_t^J(dx)}{\int e^{-\alpha g(x)} \mu_t^J(dx)}, \quad g(x) := f(x) + \frac{1}{\nu} |\mathcal{A}(x)|^2, \quad (2.3)$$

where ν is a small parameter. By the Laplace principle, for fixed t this new weighted average tends to the minimizer of g (within the support of the empirical measure and assuming this minimizer is unique) in the limit as $\alpha \rightarrow \infty$, which is expected to lie close to the feasible set $\{x : \mathcal{A}(x) = 0\}$ when the parameter ν is sufficiently small.

2. The second modification is the introduction of a drift term that drives particles towards the constraint manifold. For $0 < \varepsilon \ll 1$ we propose the dynamics

$$\begin{aligned} dx_t^{(j)} = & -\frac{1}{\varepsilon} (\nabla |\mathcal{A}|^2)(x_t^{(j)}) dt - \left(x_t^{(j)} - m_g(\mu_t^J) \right) dt \\ & + \sqrt{2} \sigma(x_t^{(j)} - m_g(\mu_t^J)) \circ dW_t^{(j)}. \end{aligned} \quad (2.4)$$

Both mechanisms are useful for driving the particles to the constraint manifold. The first modification is straightforward to implement and does not generate stiffness of the stochastic differential equations driving the particles. Consequently, these can be discretized in time with a step of the same order of magnitude as that used for unconstrained problems. The second modification is helpful for ensuring that the particles move towards the weighted average in a manner independent of the other particles. The additional drift pushes the particles to the manifold in an asymptotically orthogonal manner (indeed $\nabla |\mathcal{A}|^2$ is orthogonal to the contour lines of $|\mathcal{A}|^2$), giving us more function evaluations along the manifold. Despite being detrimental for stability at the discrete level, this additional drift is observed to be useful in our numerical experiments.

Example 2.1. *For global optimization problems subject to just one inequality constraint, i.e. prob-*

lems of the form

$$\operatorname{argmin}_{x \in B} f(x) \quad \text{with} \quad B = \{x : \mathcal{I}(x) \geq 0\}$$

with $\mathcal{I}: \mathbf{R}^d \rightarrow \mathbf{R}$, the above scheme reads

$$\begin{aligned} dx_t^{(j)} = & -\frac{1}{\varepsilon} \nabla(\mathcal{I}^2)(x_t^{(j)}) \chi_{(-\infty, 0)}(\mathcal{I}(x_t^{(j)})) dt - (x_t^{(j)} - m_g(\mu_t^J)) dt \\ & + \sqrt{2} \sigma (x_t^{(j)} - m_g(\mu_t^J)) \circ dW_t^{(j)}, \end{aligned}$$

where $\chi_{(-\infty, 0)}$ denotes the indicator function of the set $(-\infty, 0)$. The weighted average is given by

$$m_g(\mu_t^J) = \frac{\sum_{j=1}^J x_t^{(j)} \exp\left(-\alpha f(x_t^{(j)}) - \alpha \nu^{-1} \mathcal{I}(x_t^{(j)})^2 \chi_{(-\infty, 0)}(\mathcal{I}(x_t^{(j)}))\right)}{\sum_{j=1}^J \exp\left(-\alpha f(x_t^{(j)}) - \alpha \nu^{-1} \mathcal{I}(x_t^{(j)})^2 \chi_{(-\infty, 0)}(\mathcal{I}(x_t^{(j)}))\right)}.$$

2.1 Mean-field CBO

It is possible to prove a propagation of chaos result for (2.1); see [32]. In the many-particle limit $J \gg 1$, the law of $(x_t^{(1)}, \dots, x_t^{(J)})$ approximately tensorises, i.e. it holds approximately that

$$(x_t^{(1)}, \dots, x_t^{(J)}) \sim \mu_t^{\otimes J}, \quad (2.5)$$

where μ_t is a probability measure. In addition, this measure satisfies the partial differential equation

$$\partial_t \mu = \nabla \cdot \left((x - m_f(\mu)) \mu \right) + \sigma^2 \sum_{i=1}^d \partial_{ii} \left((x - m_f(\mu))_i^2 \mu \right), \quad \mu_0 = \varrho_0, \quad (2.6)$$

in the distributional sense. In order to formally derive this equation from (2.5), one can take the limit $J \rightarrow \infty$ in (2.1) and use the law of large numbers to obtain

$$dx_t^{(1)} = -(x_t^{(1)} - m_f(\mu_t)) dt + \sqrt{2} \sigma (x_t^{(1)} - m_f(\mu_t)) \circ dW_t^{(1)}. \quad (2.7)$$

The Fokker–Planck equation corresponding to this equation is then (2.6). A similar formal argument for (2.4) leads to the following nonlocal Fokker–Planck equation

$$\begin{aligned} \partial_t \mu = & \frac{1}{\varepsilon} \nabla \cdot (\nabla |\mathcal{A}|^2(x) \mu) + \operatorname{div}_x \left((x - m_g(\mu)) \mu \right) \\ & + \sigma^2 \sum_{i=1}^d \partial_{ii} \left((x - m_g(\mu))_i^2 \mu \right), \quad \mu_0 = \varrho_0. \end{aligned} \quad (2.8)$$

Remark 2.1. If g in (2.3) satisfies the assumptions in [11], then all the theoretical results proved there for the mean-field Fokker–Planck equation associated with unconstrained CBO apply *mutatis*

mutandis to (2.6) when $\varepsilon = \infty$, i.e. with only the first modification – the penalization.

2.2 Numerical experiments: mean-field versus particles and the penalizations

We illustrate the CBO approach with and without constraints by comparing the solution to the mean-field PDE (2.6) with the solution to the particle system (2.1) in simple two dimensional problems.

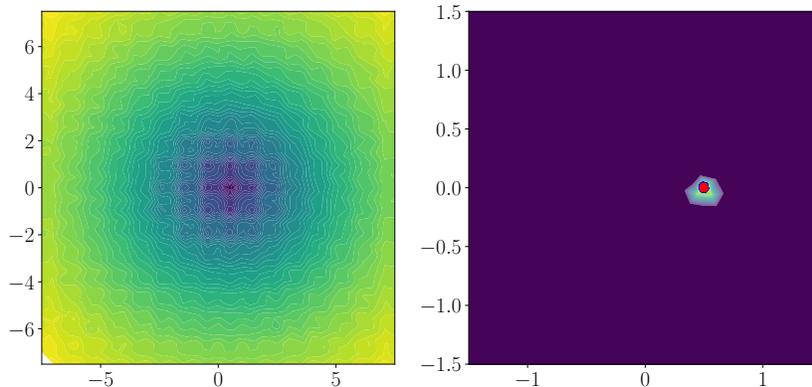


Figure 1: $J = 50$, $M = 100$, $\Delta t_{\text{SDE}} = 0.0005$, $\Delta t_{\text{PDE}} = 0.005$, $h_{\min} = 0.125$, $h_{\max} = 0.5$, $\alpha = 30$, $\nu = 1$, $\sigma = 0.7$, $T = 50$. Here h_{\min} and h_{\max} denote the minimum and maximum characteristic lengths of the mesh employed for the PDE simulation. Left: contour plot of the Ackley function shifted by $(0.5, 0)$. Right: contour plot of the PDE solution which is concentrated at one triangle. The red points show the weighted averages of 100 samples of particle simulations after convergence.

We begin with a comparison in the unconstrained setting. To this end, we use the Ackley benchmark [35] shifted by $(0.5, 0)$; see Figure 1 (left) for a contour plot. The SDE simulation uses 50 particles, which are drawn independently from the normal distribution $\mathcal{N}(0, 3I_2)$, where I_2 is the 2 by 2 identity matrix, at the beginning of the simulation. The PDE solution is initialized according to the same distribution. In view of the shift, the global minimizer of the objective function is not at the center of mass of the initial distribution. The mesh for the PDE solver has a maximal meshfield size of $h_{\max} = 0.5$ and is refined around the origin with a minimal meshfield size of $h_{\min} = 0.125$. In particular, the region around the global minimum is contained in the refined area. Additional details on the numerical implementation and the parameters employed in the simulations are given in Section 4.

Figure 1 (right) illustrates the weighted means after convergence of $M = 100$ independent simulations, as well as the contour plot of the PDE solution after convergence to one triangle. The weighted means lie at the minimizer of the shifted Ackley function and in particular in the support of the PDE solution. This demonstrates that the SDE and PDE models are in good agreement, even with as few as 50 particles. Note that the left panel in Figure 1 covers the domain $[-7.5, 7.5]^2$, while the right panel covers the zoomed-in region $[-1.5, 1.5]^2$.

Next, we study the influence of the penalization (2.3). We consider again the Ackley function shifted by $(0.5, 0)$, with now a constraint given by the circle $B = \{x \in \mathbf{R}^2: x^2 = 9\}$, which is

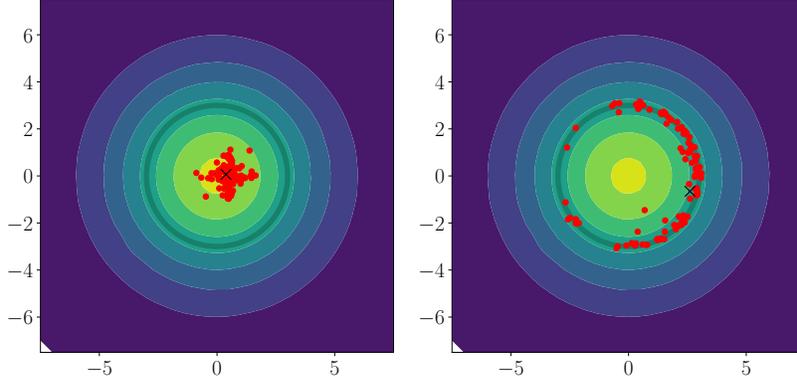


Figure 2: $J = 50, M = 100, T = 0, \alpha = 30, \nu = 0/1$. Left: weighted averages with $\nu = 0$ of 100 samples at the beginning of the simulation. Right: weighted averages with $\nu = 1$ of 100 samples at the beginning of the simulation. The contour plots show the initial distribution of the mean-field solution. The influence of the constraint in the weight is clearly visible.

depicted in gray in [Figure 2](#). In this figure, the filled contours illustrate the initial density of the PDE solution, while the red dots are the weighted averages corresponding to $M = 100$ independent particle simulations at the initial time. The black crosses indicate the positions of the weighted averages of the PDE solution at the initial time, without (left, calculated using [\(1.2\)](#)) or with (right, calculated using [\(2.3\)](#)) penalization. We observe that the weighted average with penalization of the initial PDE solution is already very close to the true global minimizer of the constrained problem. This numerical experiment motivates the strategy of enforcing the constraints through the penalization on the objective functions.

3 EKI with constraints

The EKI was initially proposed as a method for solving inverse problems of the following form: find an unknown parameter $x \in \mathbf{R}^d$ from data $y \in \mathbf{R}^K$ given that

$$y = G(x) + \eta, \quad \eta \sim \mathcal{N}(0, \Gamma), \quad (3.1)$$

where $G : \mathbf{R}^d \rightarrow \mathbf{R}^K$ is a nonlinear map known as the *forward operator*. A point estimator for the solution to [\(3.1\)](#) is usually defined as the minimizer of a least-squares functional of the form

$$\Phi_R(x; y) = \frac{1}{2} |y - G(x)|_\Gamma^2 + \frac{1}{2} |x - a|_\Sigma^2. \quad (3.2)$$

Here $|\cdot|_A$ is a short-hand notation for $|A^{-1/2} \cdot|$ and the second term is a regularization parametrized by a vector $a \in \mathbf{R}^d$ and a matrix $\Sigma \in \mathbf{R}^{d \times d}$. In the Bayesian approach to inverse problems, these parameters can be viewed as the mean and variance of a Gaussian *prior distribution* $\mathcal{N}(a, \Sigma)$ on the unknown x . In this case, the minimizer of $\Phi_R(\cdot; y)$ is the pointwise maximizer of the Bayesian posterior distribution associated with the problem [\[52\]](#). The EKI is an optimization scheme for

the functional $\Phi_R(\cdot; y)$ that can be derived from the ensemble Kalman filter. Its continuous-time version is based on the following interacting particle system:

$$\begin{aligned} dx_t^{(j)} = & -\frac{1}{J} \sum_{k=1}^J \langle G(x_t^{(k)}) - \bar{G}_t, G(x_t^{(j)}) - y \rangle_{\Gamma} (x_t^{(k)} - \bar{x}_t) dt \\ & - \mathcal{C}(\mu_t^J) \Sigma^{-1} (x_t^{(j)} - a) dt, \quad j = 1, \dots, J, \end{aligned} \quad (3.3)$$

where $\mu_t^J = \frac{1}{J} \sum_{j=1}^J \delta_{x_t^{(j)}}$ denotes as before the empirical measure associated with the ensemble, and the quantities \bar{x}_t and \bar{G}_t are defined as

$$\begin{aligned} \bar{x}_t &= \int_{\mathbf{R}^d} x \mu_t^J(dx) = \frac{1}{J} \sum_{j=1}^J x_t^{(j)}, \\ \bar{G}_t &= \int_{\mathbf{R}^d} G(x) \mu_t^J(dx) = \frac{1}{J} \sum_{j=1}^J G(x_t^{(j)}). \end{aligned}$$

The matrix $\mathcal{C}(\mu_t^J)$ is the covariance under the empirical distribution:

$$\mathcal{C}(\mu_t^J) = \int_{\mathbf{R}^d} (x \otimes x) \mu_t^J(dx) = \frac{1}{J} \sum_{j=1}^J (x_t^{(j)} - \bar{x}_t) \otimes (x_t^{(j)} - \bar{x}_t).$$

The first argument in the inner product on the right-hand side of (3.3) is related to consensus, whereas the second argument measures the mismatch with the observed data. In the limit $t \rightarrow \infty$, the particles are expected to concentrate at the global minimizer of Φ_R given in (3.2).

In the case of a linear forward map G , the EKI algorithm takes the form of a gradient descent preconditioned by the ensemble covariance:

$$dx_t^{(j)} = -\mathcal{C}(\mu_t^J) \nabla \Phi_R(x_t^{(j)}) dt, \quad j = 1, \dots, J. \quad (3.4)$$

In the case of a general forward map, the EKI (3.3) can be viewed as a derivative-free approximation of (3.4). This viewpoint, which is adopted in [25, 26], is useful below: we first show how to include constraints for (3.4), and then perform a derivative-free approximation of the resulting system.

In the Bayesian approach to inverse problems, it would be natural to incorporate equality and inequality constraints in the prior distribution on the unknown parameter x . The EKI method, however, applies only to the case of a Gaussian prior distribution, i.e. regularization by a quadratic function in (3.2), and so extensions of this algorithm are required in order to handle nonlinear constraints.

In this section, we demonstrate how a simple penalty-based method similar to that in Section 2 for CBO can be employed for including constraints in EKI. For the sake of simplicity, we assume there is only one equality constraint $\mathcal{A}(x) = 0$. This constraint can be incorporated in the gradient

method (3.4) by adding a penalization term to the regularized least-squares functional:

$$\begin{aligned} dx_t^{(j)} &= -\mathcal{C}(\mu_t^J) \nabla \left(\Phi_R + \frac{1}{\nu} |\mathcal{A}|^2 \right) (x_t^{(j)}) dt, \\ &= -\mathcal{C}(\mu_t^J) \nabla \Phi_R(x_t^{(j)}) dt - \frac{2}{\nu} \mathcal{C}(\mu_t^J) \mathcal{A}(x_t^{(j)}) \nabla \mathcal{A}(x_t^{(j)}) dt, \quad j = 1, \dots, J. \end{aligned}$$

A derivative-free version of this methodology is obtained by approximating gradients in the same manner as in standard EKI:

$$\begin{aligned} dx_t^{(j)} &= -\frac{1}{J} \sum_{k=1}^J \langle G(x_t^{(k)}) - \bar{G}_t, G(x_t^{(j)}) - y \rangle_{\Gamma} (x_t^{(k)} - \bar{x}_t) dt \\ &\quad - \mathcal{C}(\mu_t^J) \Sigma^{-1} (x_t^{(j)} - a) dt - \frac{2}{\nu} \mathcal{C}(\mu_t^J) \mathcal{A}(x_t^{(j)}) \nabla \mathcal{A}(x_t^{(j)}) dt, \end{aligned} \quad (3.5)$$

for $j = 1, \dots, J$. In contrast with the other methods proposed in the literature, the particles forming the ensemble produced by this method are not confined to the feasible region. Therefore, as already emphasized in the introduction, a prerequisite of the method is that the forward map may be evaluated for any $x \in \mathbf{R}^d$, and not only in the feasible region.

The formulation (3.5) also assumes that the gradient of the function \mathcal{A} can be calculated analytically. In the general case, one may employ a gradient-free approximation of $\nabla |\mathcal{A}|^2$,

$$\begin{aligned} dx_t^{(j)} &= -\frac{1}{J} \sum_{k=1}^J \langle G(x_t^{(k)}) - \bar{G}_t, G(x_t^{(j)}) - y \rangle_{\Gamma} x_t^{(k)} dt - \mathcal{C}(\mu_t^J) \Sigma^{-1} (x_t^{(j)} - a) dt, \\ &\quad - \frac{2}{\nu J} \sum_{k=1}^J \langle \mathcal{A}(x_t^{(k)}) - \bar{\mathcal{A}}_t, \mathcal{A}(x_t^{(j)}) \rangle (x_t^{(k)} - \bar{x}_t) dt, \quad j = 1, \dots, J, \end{aligned} \quad (3.6)$$

where $\bar{\mathcal{A}}_t = \frac{1}{J} \sum_{j=1}^J \mathcal{A}(x_t^{(j)})$. The gradient-based and gradient-free constraint terms coincide in the case where \mathcal{A} is linear.

Remark 3.1. *The scheme (3.6) can be rewritten as*

$$dx_t^{(j)} = -\frac{1}{J} \sum_{k=1}^J \langle \mathcal{G}(x_t^{(k)}) - \bar{\mathcal{G}}_t, \mathcal{G}(x_t^{(j)}) - \tilde{y} \rangle_{\tilde{\Gamma}} (x_t^{(k)} - \bar{x}_t) dt, \quad j = 1, \dots, J, \quad (3.7)$$

where $\mathcal{G}(x) = (G(x), x, \mathcal{A}(x))^\top$, $\tilde{y} = (y, a, 0_d)^\top$ and

$$\tilde{\Gamma} = \begin{pmatrix} \Gamma & 0 & 0 \\ 0 & \Sigma & 0 \\ 0 & 0 & \nu \end{pmatrix}.$$

Numerically, implementing (3.6) presents the same level of difficulty as implementing unconstrained EKI. The constraint is interpreted here as an additional observation, with very small associated

noise.

4 Numerical results for CBO

In this section, we discuss numerical results for CBO with constraints. We first present the numerical scheme employed throughout this section in [Section 4.1](#), and then we investigate the qualitative behavior of the method.

In [Section 4.2](#) and [Section 4.3](#), we present numerical results for the case of equality and inequality constraints, respectively. In [Section 4.4](#), we investigate the convergence to the mean-field equation (2.8) in the many-particle limit. In [Section 4.5](#), we compare the proposed method with another approach in the literature [24, 22]. Finally, in [Section 4.6](#), we present a brief parameter study, with the aim of highlighting the role of the parameters entering in the method.

In all the numerical experiments presented in this section, the objective function is defined from the Ackley function in two dimension:

$$f_A(x) = -20 \exp \left(-\frac{1}{5} \sqrt{\frac{1}{2} \sum_{i=1}^2 |x_i|^2} \right) - \exp \left(\frac{1}{d} \sum_{i=1}^2 \cos(2\pi x_i) \right) + e + 20.$$

4.1 Numerical schemes

As mentioned, the introduction of a relaxation drift towards the feasible manifold with $\varepsilon \ll 1$ leads to stiffness of the CBO particle system (2.4). This leads to a stringent limitation on the time step when using an explicit method such as Euler–Maruyama [29]. In the case of a single quadratic equality or inequality constraint, this issue can be mitigated by using a semi-implicit scheme. Let us consider, for example, an equality constraint of the form

$$\mathcal{A}(x) = x^\top A x - c = 0 \tag{4.1}$$

for a positive definite matrix $A \in \mathbf{R}^{d \times d}$ and a positive real number c . This general form of the constraint encompasses elliptic and, in particular, the circular equality constraints considered throughout this section. Moreover, it is simple to extend the method to similar inequality constraints. In order to use CBO with the constraint (4.1), we propose the discretization

$$y_n^{(j)} = x_n^{(j)} - (x_n^{(j)} - m_g(\mu_n^J)) \Delta t + \sqrt{2\Delta t} \sigma(x_n^{(j)} - m_g(\mu_n^J)) \circ \xi_n^{(j)}, \tag{4.2a}$$

$$x_{n+1}^{(j)} = \left(I_d + \frac{4\Delta t}{\varepsilon} \mathcal{A}(x_n^{(j)}) A \right)^{-1} y_n^{(j)}, \tag{4.2b}$$

where $\xi_n^{(j)} \sim \mathcal{N}(0, 1)$, for $n = 1, 2, \dots$ and $j = 1, \dots, J$, are independent. This discretization arises when treating the term $\nabla \mathcal{A}$ in (2.4) implicitly and all the other terms on the right-hand side explicitly. The PDE associated with CBO is solved with the help of the Python free software library **Fenics** [3] on a mesh refined locally around the feasible manifold using the open source software

gmsb [27]. We use the discontinuous Galerkin method with elements of degree zero. The drift term is computed with upwind flux and for the diffusion we use the discretization proposed in [42].

In order to make the SDE and PDE results comparable, we initialize the particle SDE (2.4)-(2.3) with $x_0^{(j)} \sim \mathcal{N}(0, 3I_2)$, and the corresponding PDE (2.8) is initialized with initial condition $\varrho_0(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{x}{3})^2}$, i.e. ϱ_0 is the probability density function of $\mathcal{N}(0, 3I_2)$. Unless otherwise stated, we always assume these initial conditions at the particle and PDE levels. The following parameters are fixed for all simulations: $\sigma = 0.7, \alpha = 30, \Delta t_{\text{PDE}} = 5 \times 10^{-3}, \Delta t_{\text{SDE}} = 5 \times 10^{-4}$. The number of particles J , and the relaxation parameter ε , are specified on a case-by-case basis.

4.2 CBO with equality constraint

We now present the results of particle simulations of the CBO method with constraints (2.4), which incorporates both the relaxation drift towards the constraint and the penalization (2.3), and of PDE simulations of the associated mean field equation (2.8). We begin by considering the following optimization problem with circular equality constraint,

$$\operatorname{argmin}_{x \in B} f_A(x - x_*), \quad x_* = \begin{pmatrix} 3 \\ 0 \end{pmatrix}, \quad B = \{x \in \mathbf{R}^2 : x^2 = 9\}.$$

Here f_A is the Ackley function, which is minimized at $x = 0$, so that the solution of this problem is given by $x_* = (3, 0)$. (Note that this corresponds to the minimizer of the unconstrained global optimization problem.)

The top left panel in Figure 3 illustrates the shifted Ackley function (filled contour) and the constraint set B (gray line). The top right, bottom left and bottom right panels depict the PDE solution (filled contour), as well as the weighted averages corresponding to 100 independent simulations of the particle system (one red dot per independent simulation), at times 0.01, 0.15 and 0.5. We observe that, already for $t = 0.01$, the weighted averages provide a good approximation of the global minimizer.

4.3 CBO with inequality constraint

Now, we investigate the following example with inequality constraint:

$$\operatorname{argmin}_{x \in B} f_A(x - x_*), \quad x_* = \begin{pmatrix} 2 \\ 2 \end{pmatrix}, \quad B = \{x \in \mathbf{R}^2 : x^2 \geq 18\},$$

where f_A is the standard Ackley function in two dimensions. The associated numerical results are illustrated in Figure 4. Again, the red particles indicate the positions of the weighted means of $M = 100$ independent simulations with $J = 50$ particles each. The top left panel of Figure 4 shows the shifted Ackley function (filled contour), and the boundary of the feasible set (gray line). The top right, bottom left and bottom right panels depict the evolution of the PDE solution as well as the weighted means of the SDE simulations. Compared to the previous example, the solution

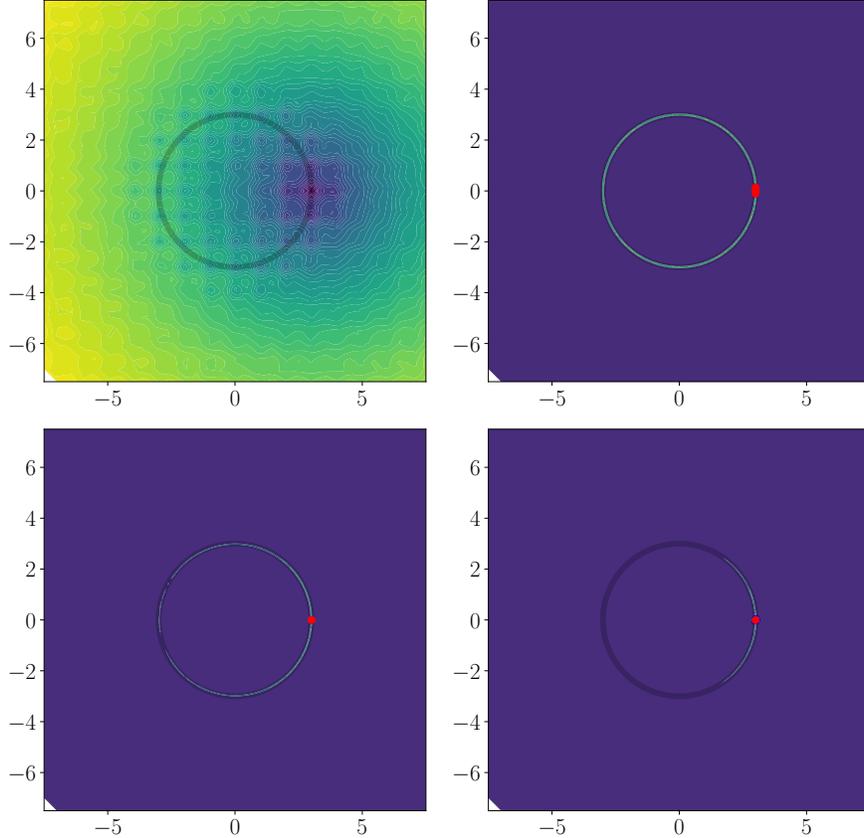


Figure 3: **Top left panel:** Ackley function shifted by $(3, 0)$, and feasible set $B = \{x \in \mathbf{R}^2: x^2 = 9\}$. **Other panels:** PDE solution at $t = 0.01$, $t = 0.15$ and $t = 0.5$ (filled contour), and weighted averages of $M = 100$ independent simulations with $J = 50$ particles each (red dots). We observe that the weighted averages lie in the support of the mean-field solution and provide a good approximation of the global minimizer. The parameters employed for the simulation are the following: $J = 50$, $M = 100$, $\Delta t_{\text{SDE}} = 0.0005$, $\Delta t_{\text{PDE}} = 0.005$, $h_{\min} = 0.125$, $h_{\max} = 0.5$, $\alpha = 30$, $\nu = 1$, $\sigma = 0.7$, $\varepsilon = 0.1$.

to the mean-field PDE is less concentrated, which is expected since the constraint is of inequality type. At $t = 0.01$ (top right plot), not all weighted averages are close to the minimizer. At $t = 0.15$ and $t = 0.5$ (bottom plots), the weighted averages are very close to the minimizer, and the density concentrates around it.

We emphasize that the accuracy of the PDE results is strongly restricted by the mesh size; since this is fixed during a simulation, the PDE solver is unable to capture the evolution of the ensemble after this has concentrated to a very small area. Comparisons between the SDE and PDE simulations are therefore meaningful only in the beginning of the simulation. For this reason we restrict the following numerical studies to particle simulations.

4.4 Many particle limit $J \rightarrow \infty$ for CBO

In the following we investigate the behavior of the proposed CBO scheme as J tends to infinity. For simplicity, we consider toy examples where the true minimizer x_* is known. For the particle scheme

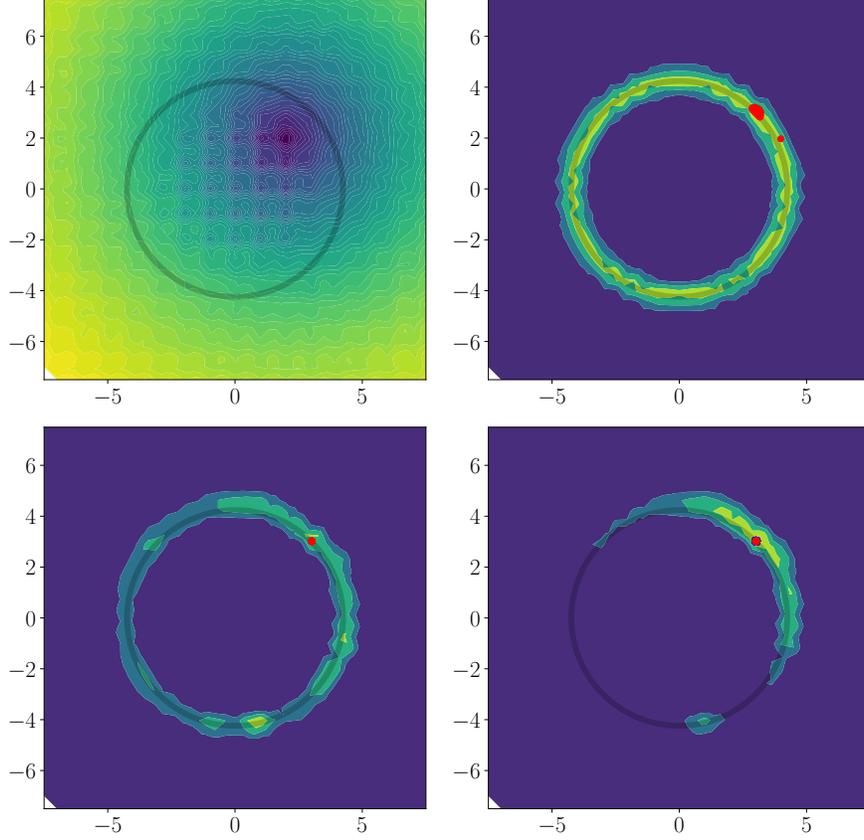


Figure 4: Parameters: $J = 50, M = 100, \Delta t_{\text{SDE}} = 0.0005, \Delta t_{\text{PDE}} = 0.005, h_{\min} = 0.125, h_{\max} = 0.5, \alpha = 30, \nu = 1, \sigma = 0.7, \varepsilon = 0.1$ Top-left: Ackley function shifted by $(2, 2)$ and the constraint $x^2 \geq 18$ (gray). Top-right to bottom-right: mean-field density and weighted means of the particle solution at $t = 0.01, 0.15$ and 0.5 . The weighted averages lie in the support of the mean-field solution and very close to the global minimum of the function.

we perform $M = 100$ independent simulations for every setting considered, each with a number J of particles. For $j = 1, \dots, J$ and $m = 1, \dots, M$, the vector $x_t^{(m,j)} \in \mathbf{R}^d$ denotes the position of the j -th particle in the m -th independent simulation at time t . First, we study in [Figure 5](#) the evolution of the variance (more precisely, the trace of the covariance) of particle positions under an averaged empirical measure,

$$\text{Var}_{X \sim \bar{\mu}_t^J}(X) = \int \left| x - \int y \bar{\mu}_t^J(dy) \right|^2 \bar{\mu}_t^J(dx), \quad \bar{\mu}_t^J = \frac{1}{M} \sum_{m=1}^M \mu_t^{J,m},$$

where $\mu_t^{J,m}$ denotes the empirical measure for the m -th independent simulation at time t , for different ensemble sizes. During the first few time steps, we observe a fast decrease of all variances, as all the particles quickly move toward the feasible manifold. Thereafter, the ensembles continue concentrating but more slowly. As the figure shows, starting from $t \approx 8$ ensembles with a smaller number of particles begin to collapse faster than ensembles with more particles. This plot does not

reveal any information on the proximity of the convergence point to the global minimizer, however.

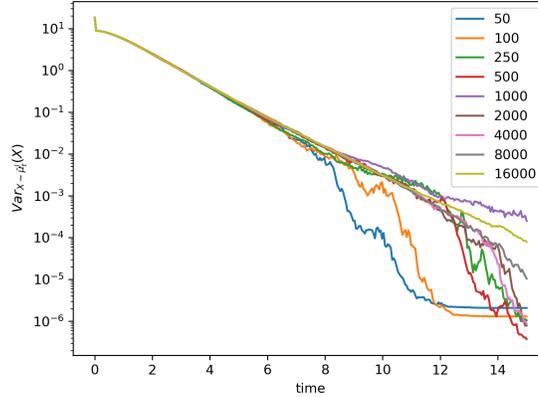


Figure 5: Ensemble variance for different ensemble sizes. For the simulation we used the following parameters: $M = 100$, $\Delta t_{SDE} = 0.0005$, $\alpha = 30$, $\nu = 1$, $\sigma = 0.7$, $\epsilon = 0.1$.

The accuracy of the convergence point is discussed in [Figure 6](#), which depicts the time evolution of the expected value of the Wasserstein distance $W_2(\mu_t^J, \delta_{x_*})$, approximated in practice based on 100 independent simulations. Again, we see a strong decrease at the beginning of the simulation as

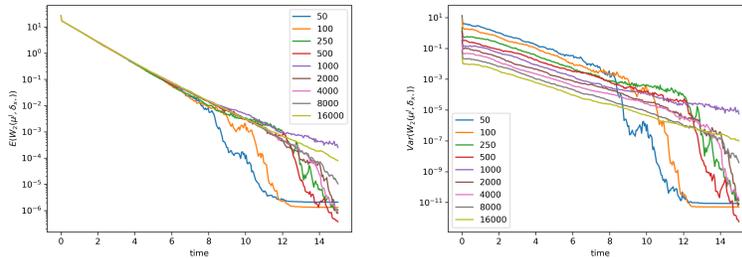


Figure 6: Evolution in time of the expectation (**left**) and variance (**right**) of the Wasserstein distances. The parameters employed here are the same as in [Figure 5](#).

the particles are driven towards the manifold. Then the expected values of the Wasserstein distances decrease exponentially until the ensembles collapse. As we already observed in [Figure 5](#), smaller ensembles collapse earlier than larger ones. The approximation of the true minimizer is very good already for small ensemble sizes.

To illustrate the convergence to the deterministic mean-field dynamics in the limit as $J \rightarrow \infty$, we plot the Wasserstein distances of $W_2(\mu^J, \mu^{16000})$ for different values of J in [Figure 7](#), with this time only one simulation per value of J . At $t = 0$ we see the expected decrease of the Wasserstein distance along the y -axis. The distances vary over time but decrease at the end of the simulation as all ensembles converge towards the global minimizer.

In [Figure 8](#) the average distance of the ensembles to the constraint is illustrated with the help

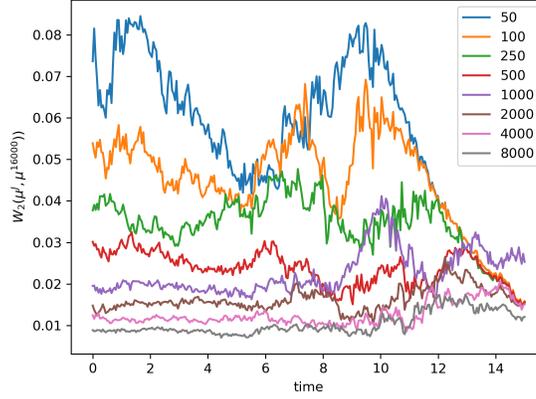


Figure 7: Wasserstein distance to reference solution with $J = 16000$. The parameters employed here are the same as in [Figure 5](#).

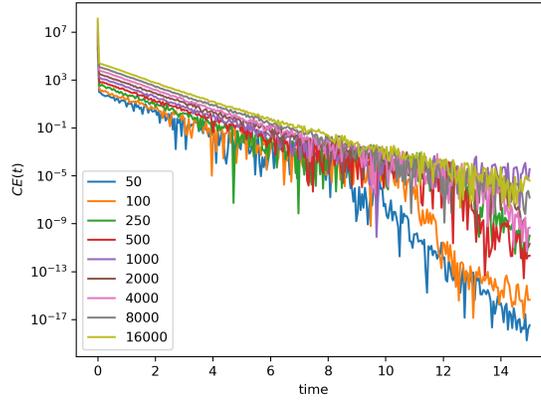


Figure 8: Distance to constraint measured with the quantity defined in (4.3). The parameters employed here are the same as in [Figure 5](#).

of the quantity

$$\text{CE}(t) = \frac{1}{M} \sum_{m=1}^M \left(\frac{1}{J} \sum_{i=1}^J \mathcal{A}(x_t^{(m,i)})^2 \right) \approx \mathbf{E} \left(\int_{\mathbf{R}^d} |\mathcal{A}(x)|^2 \mu_t^J(dx) \right). \quad (4.3)$$

This may be viewed as a “constraint energy”; the smaller $\text{CE}(t)$ is, the closer the ensemble is to the constraint. At the beginning of the simulation, the quantity $\text{CE}(t)$ decreases very fast thanks to the relaxation drift towards the manifold. As time increases, we observe small fluctuations of the constraint energy, but it is likely that these oscillations would disappear in the limit $M \rightarrow \infty$ of many independent simulations. The oscillations are due to the diffusion term acting on the particles, which is stronger for particles far away for the weighted mean. At $t = 15$, all ensembles are close to the constraint.

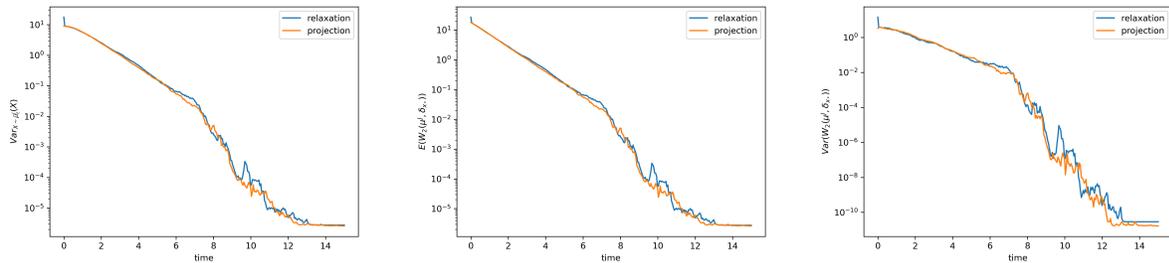


Figure 9: Comparison of our method (2.4) with the projection-based method. The parameters employed for the simulation are the following: $J = 50, M = 100, \alpha = 30, \nu = 1, \Delta t = 0.0005, \sigma = 0.7, \varepsilon = 0.1$.

4.5 Comparison with a projection-based method

In the introduction, we mentioned that, in the case of CBO, the approach discussed here incorporates the constraint via two distinct mechanisms: a penalization added to the objective function, and a relaxation drift towards the manifold. This is in contrast with other methods that ensure that the constraint is satisfied at all times. We therefore make a comparison of the method proposed here and a method using projection to the constraint in every time step, as proposed in [24]. The results of the comparison are presented in Figure 9. In the left panel, which depicts the evolution of the expected value of the variance, we observe that the two graphs differ at the beginning, which is not surprising given that the ensembles are not initially confined to the manifold with our method, and then the two approaches perform comparably. One advantage of our approach is that the initial ensemble can be drawn from any probability distribution over \mathbf{R}^d , whereas for the projection method of [24] the initial ensemble needs to be supported on the constraint. In addition, calculating projections to the feasible manifold at each time step can be computationally costly. The panel in the middle illustrates the evolution of the expected value of the Wasserstein distance to the minimizer, $W_2(\mu^{50}, \delta_{x_*})$, approximated for each method from 100 independent simulations and the evolution of the variance of the Wasserstein distance to the minimizer, again approximated for each method from 100 independent simulations, is shown on the right. As expected, the ensembles corresponding to both methods eventually collapse and provide reasonable approximations of the global minimizer.

We notice that, in all the above plots concerning CBO, the time evolution of the particle ensemble can be divided into roughly three phases. In the first phase, at the beginning of the simulation, the particles are quickly driven to the manifold. In the second phase, the particles move along the feasible manifold towards the solution of the constrained optimization problem, and few fluctuations are observed in the plots. The size of the fluctuations, in logarithmic scale, in the final phase are consistent with the fact that the amplitude of the noise decreases as we get closer to the collapse of the ensemble near the global minimizer at the final stages of the simulation.

To conclude this section, we examine the influence of the parameter ε , which enters in the relaxation drift, on the speed of relaxation towards the constraint. This is illustrated in Figure 10,

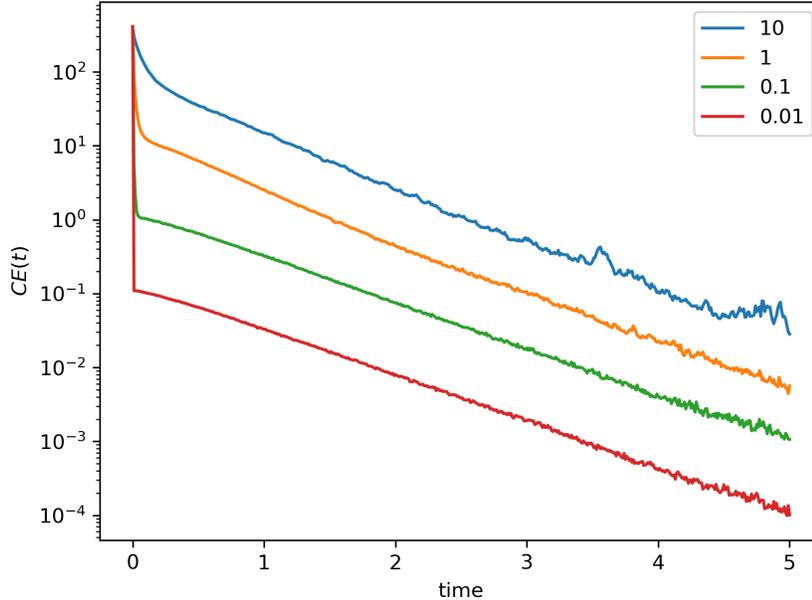


Figure 10: Distance to constraint for different relaxation parameters ϵ . Other parameters of the simulation are: $J = 1000$, $M = 100$, $\alpha = 30$, $\nu = 1$, $\Delta t = 0.0005$, $\sigma = 0.7$.

where the evolution of “constraint energy”, measured using (4.3), is shown for different values of ϵ . For all the values of ϵ considered, the “constraint energy” is approximated based on $M = 100$ independent simulations with $J = 1000$ particles each.

4.6 Numerical experiments for the CBO particle system

In this section, we present additional numerical experiments concerning the particle system. We investigate, in particular, the influence of the parameters ν , ϵ and J on the accuracy of the convergence point of CBO, first for the following optimization problem:

$$\operatorname{argmin}_{x \in B} f_A(x - x_*), \quad x_* = \begin{pmatrix} 2 \\ 2 \end{pmatrix}, \quad B = \{x \in \mathbf{R}^2: x^2 + y^2 = 18\}. \quad (4.4)$$

All the numerical results presented below are generated from CBO with the parameters $\sigma = 0.7$ and $\Delta t = 0.01$. In each simulation, the particles forming the initial ensembles are always drawn from $\mathcal{N}(0, 100I_2)$. The other parameters are specified on a case-by-case basis.

In [Figure 11](#), we illustrate for different values of $\nu \in \{10, 1, .1, .01\}$ and $J \in \{100, 1000\}$, the convergence points of 100 independent simulations (per set of parameters). Here we take $\epsilon = \infty$; that is, we employ the CBO particle system with penalization but without relaxation drift. From the figures corresponding to $J = 100$, we observe that, although smaller values of ν enable to enforce the constraint more effectively, taking ν too small is detrimental for the accuracy of the convergence

point. This effect is observed also for $J = 1000$, but to a much lesser extent, which may indicate a better behavior of the system in the many-particle limit. This numerical experiment highlights the considerable influence of the number of particles on the behavior of the method. In practice, a trade-off needs to be achieved between precision and computational cost.

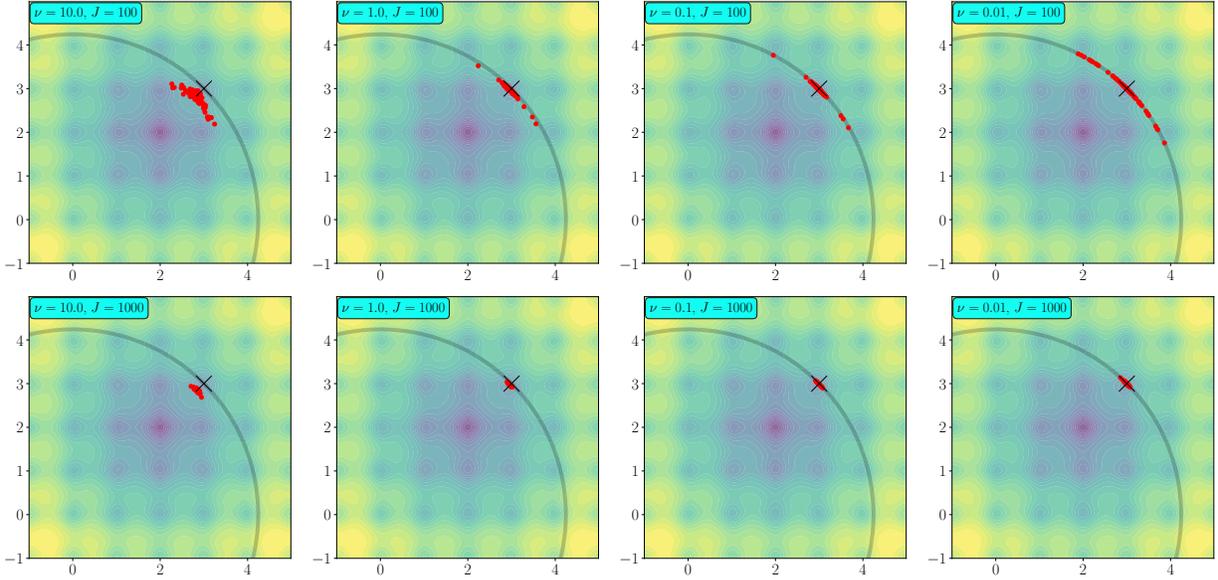


Figure 11: Points of convergence of $M = 100$ simulations of CBO for the constrained optimization problem (4.4), for different values of ν and J and for $\varepsilon = 0$. The black cross indicates the position of the global minimizer under the constraint.

In [Figure 12](#), we illustrate the convergence points of 100 independent simulations with this time $(\nu, \varepsilon) \in \{.1, 10\}^2$, and for fixed $J = 100$.

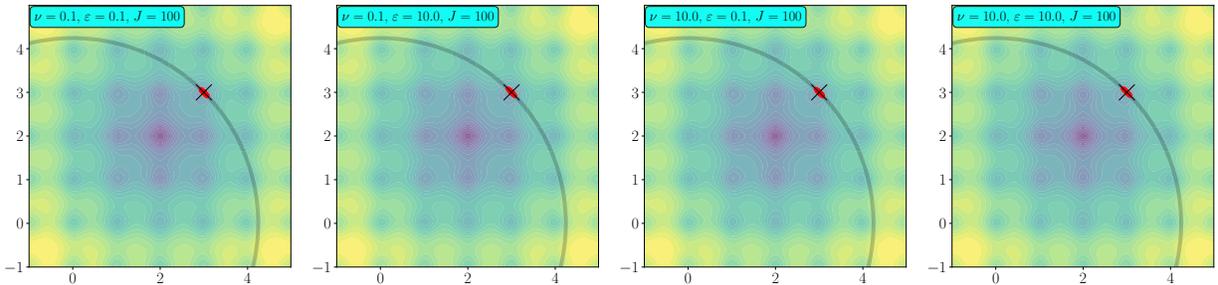


Figure 12: Points of convergence of $M = 100$ simulations of CBO for the constrained optimization problem (4.4), for different values of ν and ε .

For the simulations corresponding to $\varepsilon = .1$, we take a smaller time step $\Delta t = 10^{-3}$ in order to avoid stability issues, and in all the simulations, we employ the semi-implicit scheme (4.2). Interestingly, the presence of the relaxation drift, even with a relatively small amplitude $1/\varepsilon$, considerably improves the performance of the method compared to the case without penalization.

To conclude this section, we present in [Figure 13](#) numerical results for (4.4) with the inequality

constraint $B = \{x \in \mathbf{R}^2: x^2 + y^2 \geq 18\}$ instead of an equality constraint. The problem setting is then the same as in [Figure 3](#). Rather than presenting the points of convergence of the method for several parameter choices as in the previous figures, here we perform only one simulation with a given set of parameters ($\nu = \varepsilon = 1$, $J = 100$), and we illustrate the evolution in time of the ensemble. As the figure shows, the ensemble appears to collapse after roughly 800 iterations, that is at time $t \approx 8$. In addition, the point of convergence is very close to the optimal solution given the constraint.

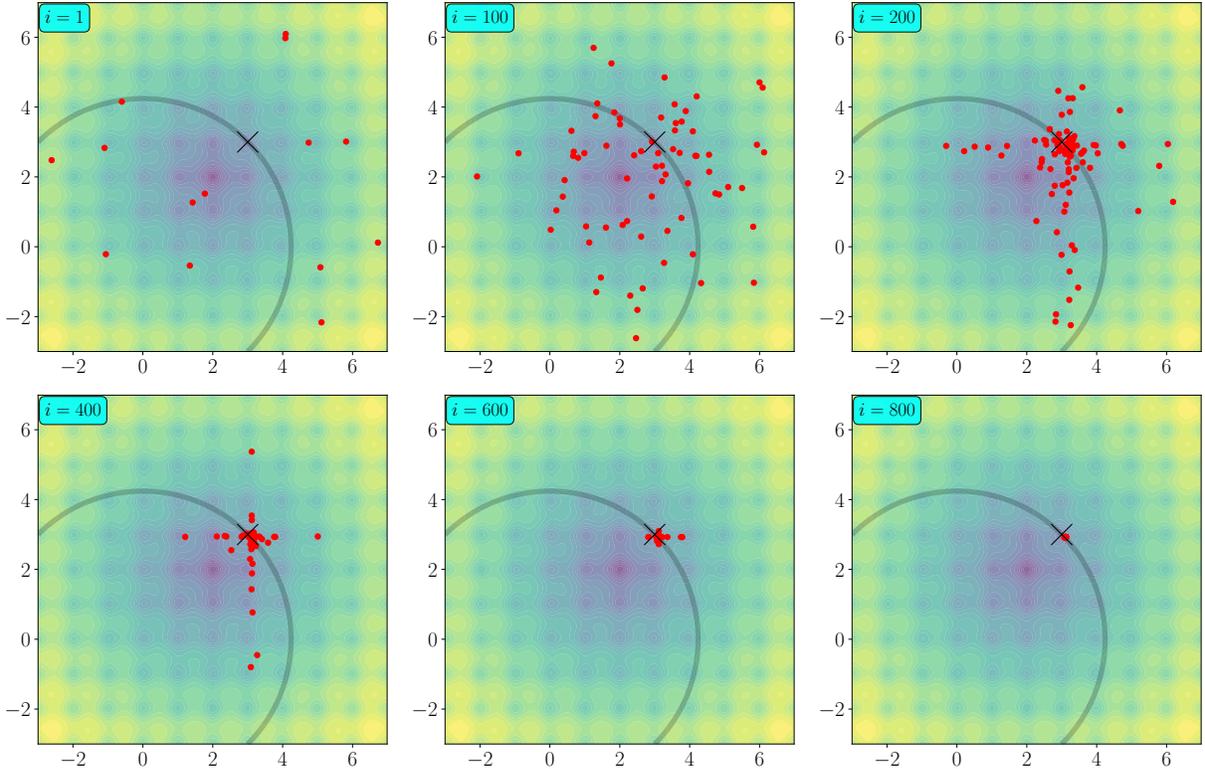


Figure 13: Evolution of the particles for simulations when CBO is employed for minimizing the Ackley function under the constraint that $\{x^2 + y^2 \geq 18\}$. The parameters for this example are $\varepsilon = \nu = 1$ and $J = 100$.

5 Numerical results for EKI with constraints

In this section, we present numerical results for EKI with constraints. We simulate only the particle system and do not present result for the associated mean-field equation. Our numerical experiments aim at illustrating the performance of the method on a toy problem, as a proof of concept.

The numerical schemes employed are described in [Section 5.1](#), and numerical results for a simple inverse problem arising from a PDE application are presented in [Section 5.2](#).

5.1 Numerical schemes

The presence of the relaxation term originating from the constraint poses a challenge also for the discretization of EKI with constraints (3.5), although to a lesser extent than for CBO thanks to the affine-invariance of the method [26]. In this section, we consider specifically the derivative-free dynamics (3.7), and we investigate two different approaches for discretizing this dynamics in time:

- An explicit discretization using the explicit Euler method with adaptive time step. Specifically, the particle positions are evolved according to

$$\begin{aligned} x_{n+1}^{(j)} &= x_n^{(j)} - \Delta t_n \sum_{k=1}^J M_n^{kj} x_n^{(k)}, \quad j = 1, \dots, J, \\ M_n^{kj} &= \frac{1}{J} \langle \mathcal{G}(x_n^{(k)}) - \bar{\mathcal{G}}_n, \mathcal{G}(x_n^{(j)}) - \tilde{y} \rangle_{\tilde{\Gamma}}, \end{aligned} \quad (5.1)$$

and the time step is calculated dynamically according to the method proposed in [41], that is

$$\Delta t_n = \frac{\Delta t_*}{\|M_n\|_2 + \frac{\Delta t_*}{\Delta t_{\max}}}, \quad (5.2)$$

where Δt_* (base time step) and Δt_{\max} (maximum time step) are parameters, and where M_n is the matrix with entries M_n^{kj} . With the notation $X_n = (x_n^{(1)}, \dots, x_n^{(J)})$, equation (5.1) reads in matrix form as

$$X_{n+1} = X_n - \Delta t_n X_n M_n. \quad (5.3)$$

- A semi-implicit discretization given by

$$X_{n+1} = X_n - \Delta t_n X_{n+1} M_n,$$

where the same notation is used as in the previous item. The associated update formula is given by

$$X_{n+1} = X_n (I_d + \Delta t_n M_n)^{-1}. \quad (5.4)$$

Although this approach could be employed with a fixed time step, we obtain faster convergence when the time step is adapted according to (5.2), and so we consider only the latter setting.

Remark 5.1. Note that $\sum_{k=1}^K M_n^{kj} = 0$ for all n and j , by definition of $\bar{\mathcal{G}}_n$. Consequently, the update formulas (5.3) and (5.4) can be rewritten equivalently as

$$X_{n+1} = X_n - \Delta t_n (X_n - \bar{x}_n) M_n, \quad (5.5a)$$

$$X_{n+1} = \bar{x}_n + (X_n - \bar{x}_n) (I_d + \Delta t_n M_n)^{-1}, \quad (5.5b)$$

where $\bar{x}_n = \frac{1}{J} \sum_{j=1}^J x_n^{(j)}$ and $X_n - \bar{x}_n$ denotes, by a slight abuse of notation, the matrix obtained by subtracting \bar{x}_n from each column of X_n . Although mathematically equivalent, these discretizations

are observed in our numerical experiments to be much less sensitive to roundoff errors, and therefore preferable in practice.

5.2 Numerical experiments for EKI with constraints

In this section, we illustrate the performance of the dynamics (3.7), discretized according to (5.5a) or (5.5b), for solving a toy inverse problem with constraints. The problem considered concerns the recovery of the initial condition of a Fokker–Planck equation based on incomplete observation of the solution. More precisely, our aim is to find the parameters of a Gaussian mixture

$$\varrho_0(x) = \sum_{n=1}^N w_n \frac{\exp\left(-\frac{|x - m_n|^2}{2\sigma_n^2}\right)}{\sqrt{2\pi\sigma_n^2}}, \quad (5.6)$$

given noisy observations at time T of the solution $\varrho(x, t)$ to the following initial value problem with initial condition ϱ_0 :

$$\begin{cases} \partial_t \varrho = \partial_x(x\varrho + \partial_x \varrho), & \text{in } \mathbf{R} \times (0, T), \\ \varrho = \varrho_0, & \text{on } \mathbf{R} \times \{t = 0\}. \end{cases} \quad (5.7)$$

We assume that the data is composed of noisy observations

$$y_k = \varrho(x_k, T) + \eta_k, \quad 1 \leq k \leq K,$$

where $x_k \in \mathbf{R}$ are discrete positions and η_k are noise terms drawn independently from $\mathcal{N}(0, \gamma^2)$ for some covariance γ^2 . The observation positions are given by $x_k = -L + (k - 1)\frac{L}{K-1}$, i.e. they are uniformly spread between $-L$ and L , with both ends included.

Recovering the weights. Assuming first that the means and variances of the Gaussian mixture (5.6) are known, we seek to find, as an approximation of the true weights w^\dagger , a vector of weights $w = (w_1, \dots, w_N)^\top$ that minimizes the least-square misfit

$$f(w) = \sum_{k=1}^K \frac{1}{\gamma^2} |y_k - \varrho_e(x_k, T; w)|^2,$$

under the constraints that $\sum_{n=1}^N w_n = 1$ and $w_n \geq 0$ for all $n \in \{1, \dots, N\}$, which guarantee that ϱ_0 is a probability density. Here $\varrho_e(T, x; w)$ is the exact solution to the initial value problem (5.7) when the weights in the initial condition are given by w . This admits the following explicit expression:

$$\varrho_e(x, t; w) = \sum_{n=1}^N w_n \frac{\exp\left(-\frac{|x - m_n(t)|^2}{2\sigma_n(t)^2}\right)}{\sqrt{2\pi\sigma_n(t)^2}}, \quad (5.8)$$

with $m_n(t) = e^{-t}m_n$ and $\sigma_n(t) = \sqrt{1 + (\sigma_n^2 - 1)e^{-2t}}$. Note that, in this problem, the forward model $w \mapsto \{\varrho_\epsilon(x_k, T; w)\}_{k=1}^K$ is linear and so, in the absence of constraints, the ensemble Kalman inversion method performs an exact preconditioned gradient descent.

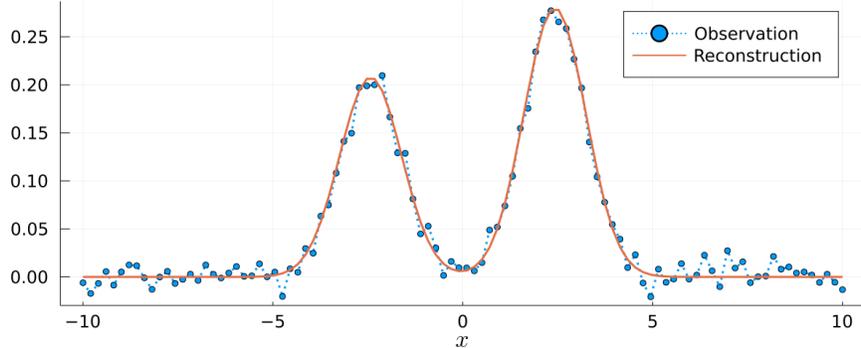


Figure 14: Noisy observations of the solution to the Fokker–Planck equation, in the case of only 2 components in the binary mixture, and solution at time T obtained by using the reconstructed initial condition.

We assume, for simplicity, that the Gaussian mixture is composed of only $N = 2$ components, and that the known means and variances of these components are given by $m_1 = -m_2 = 4$ and $\sigma_1 = \sigma_2 = .1$. For the observations, we take the parameters $L = 10$, $\gamma = .01$ and $K = 100$, and for the adaptation of the time step (5.2) we use the parameters $\Delta t_* = 1$ and $\Delta t_{\max} = \infty$. The ensemble size is taken to be $J = 100$, the final time is $T = .5$, and the noisy observations are illustrated in Figure 14.

We begin by analyzing the influence of the small parameter ν on the solution returned by the method when the particles forming the initial ensemble, each corresponding to the couple of weights, are drawn from $\mathcal{N}(0, I_2)$. For this numerical experiment, we use only the explicit discretization (5.5a). Since the evolution (3.7) is deterministic, and since we observed empirically that the initial ensemble does not have much influence on the point of convergence of the method, we run only one simulation per value of ν . In the table below, we indicate for several values of ν the value of the sum $w_1 + w_2$ at the point of convergence of the method, as well as the distance to the true value of the weights from which the observed data was generated. This was taken to be $w^\dagger = (0.411\dots, 0.588\dots)$.

ν	$w_1 + w_2$	$ w_1 - w_1^\dagger + w_2 - w_2^\dagger $
1	0.9845...	0.0282...
10^{-2}	0.9847...	0.0282...
10^{-4}	0.99289...	0.0282...
10^{-6}	0.999869...	0.0282...
10^{-8}	0.99999868...	0.0282...

As expected, we observe that the smaller ν , the closer the point of convergence is to the feasible manifold. In addition, the proximity to the true weights is hardly affected by changes in the value

of ν . **Figures 15** and **16** depict the evolution of the ensemble for $\nu = 1$ and $\nu = 10^{-8}$. Although the method converges quickly in both cases, the dynamics look qualitatively very different: for $\nu = 1$ the effect of the penalization is not clearly apparent, while for $\nu = 10^{-8}$ the constraint term appears dominant in the penalized objective function (2.3) depicted in the background. In the latter case, the particles first converge to a vicinity of the feasible line $w_1 + w_2 = 1$, and then then move along this line to the optimizer.

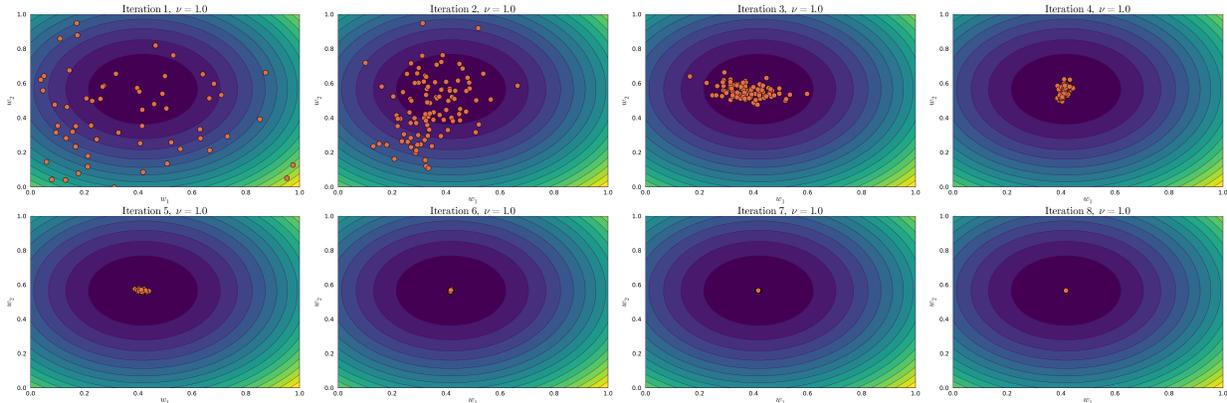


Figure 15: Evolution of the ensemble obtained from the explicit discretization (5.5a) of EKI with constraints, in the case where $\nu = 1$. The penalized objective function (denoted by g in (2.3)) is depicted as a filled contour in the background.

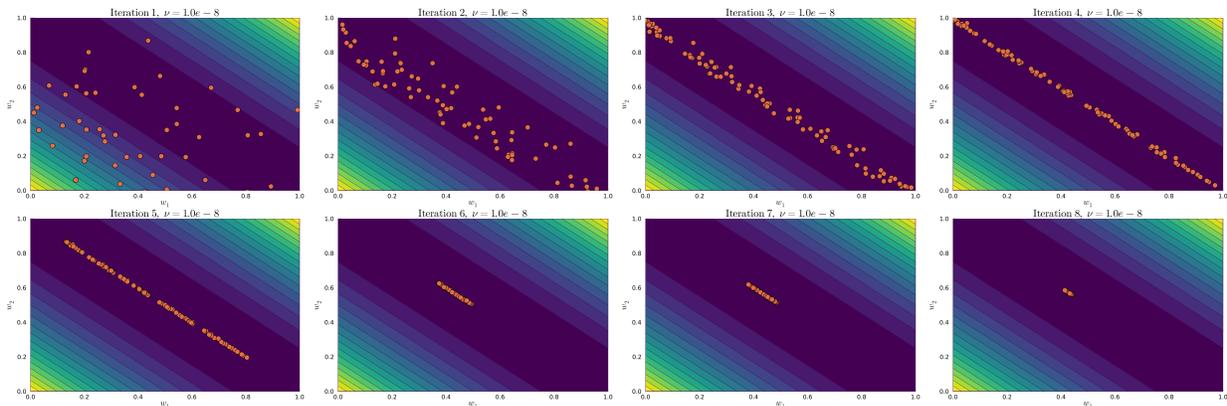


Figure 16: Like **Figure 15**, this plot depicts the evolution the discrete-time dynamics (5.5a), with now $\nu = 10^{-8}$.

To conclude this paragraph, we compare the two discretizations proposed in (5.1), when both methods use the same initial ensemble and with $\nu = 10^{-8}$. In the left panel of **Figure 17**, we illustrate the evolution of the error, measured as $|\bar{w}_1 - w_1^\dagger| + |\bar{w}_2 - w_2^\dagger|$, where \bar{w}_1 and \bar{w}_2 are sample averages of the weights over the particles. In the right panel, we present the evolution of the matrix 2-norm of the sample covariance for each method. Both methods converge quickly, but the figures indicate a slightly faster convergence for the fully explicit discretization.

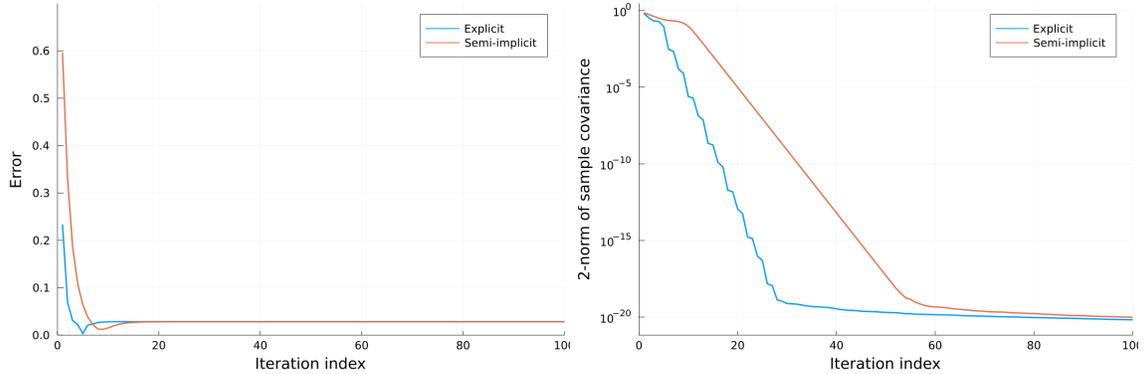


Figure 17: Evolution of the error (**left**) and sample covariance (**right**), for the explicit and semi-implicit discretizations described in Section 5.1. In both cases, the time step is adapted dynamically according to (5.2). We emphasize that the error is computed with respect not to the optimal solution to (5.9) but to the true weights, which were employed to generate the data. It is not surprising, therefore, that the error converges to a strictly positive value in the limit as the iteration index tends to infinity.

Recovering the weights and variances. In this paragraph, we consider the more challenging case where $N = 3$ and both the weights and variances of the initial Gaussian mixture need to be recovered, in which case the forward model is no longer linear. For simplicity, we assume that the means $\{m_n\}_{n=1}^N$ are still known and this time given by $(m_1, m_2, m_3) = (-5, 0, 5)$, and we seek to minimize

$$f(w, v) = \sum_{k=1}^K \frac{1}{\gamma^2} |y_k - \varrho_e(x_k, T; w, \text{abs.}(v))|^2, \quad (5.9)$$

where w is the vector of weights and $v = (\sigma_1^2, \sigma_2^2, \sigma_3^2)$ is the vector of variances, under the constraints

$$\sum_{n=1}^3 w_n = 1, \quad w_n \geq 0 \quad \forall n \in \{1, 2, 3\}, \quad \sigma_n^2 \geq 0 \quad \forall n \in \{1, 2, 3\}.$$

Here the function “abs.” denotes the element-wise absolute value, and the function $(x, t) \mapsto \varrho_e(x, t; w, v)$ denotes the exact solution (5.8) to the Fokker–Planck equation (5.7) when the parameters w and v are employed in the initial condition (5.6). We define the objective function in this manner, with the presence of the “abs.” function, in order to guarantee that this function can be evaluated for any choice of parameters $w \in \mathbf{R}^3$ and $v \in \mathbf{R}^3$, which is a requirement for applying EKI.

Apart from the parameters of the Gaussian mixtures, and unless otherwise specified, all the parameters employed to generate the numerical results presented in the rest of this section are the same as in the previous paragraph. We begin by examining the influence of the small parameter ν on the error and convergence point of the method. We employ to this end the explicit method described in Section 5.1. The table below, in which all figures are truncated after three significant digits, gives the points of convergence of EKI with constraint for different values of ν . We observe again that small values of ν lead to a point of convergence closer to the feasible manifold, as is

expected.

	$w_1 + w_2 + w_3$	w_1	w_2	w_3	σ_1^2	σ_2^2	σ_3^2
Truth	1	0.333	0.476	0.191	0.400	0.100	0.500
$\nu = 1$	0.9844	0.337	0.480	0.167	0.433	0.095	0.303
$\nu = 10^{-2}$	0.9847	0.337	0.480	0.167	0.434	0.095	0.305
$\nu = 10^{-4}$	0.99502	0.342	0.482	0.171	0.467	0.093	0.363
$\nu = 10^{-6}$	0.9999273	0.344	0.483	0.173	0.483	0.092	0.392
$\nu = 10^{-8}$	0.999999270	0.344	0.483	0.173	0.483	0.092	0.393

The solution to the Fokker–Planck equation at time T , when using as parameters in the initial condition the point of convergence of EKI with small parameter $\nu = 10^{-8}$, is illustrated in [Figure 18](#). A good qualitative agreement between the observed and reconstructed solutions is observed.

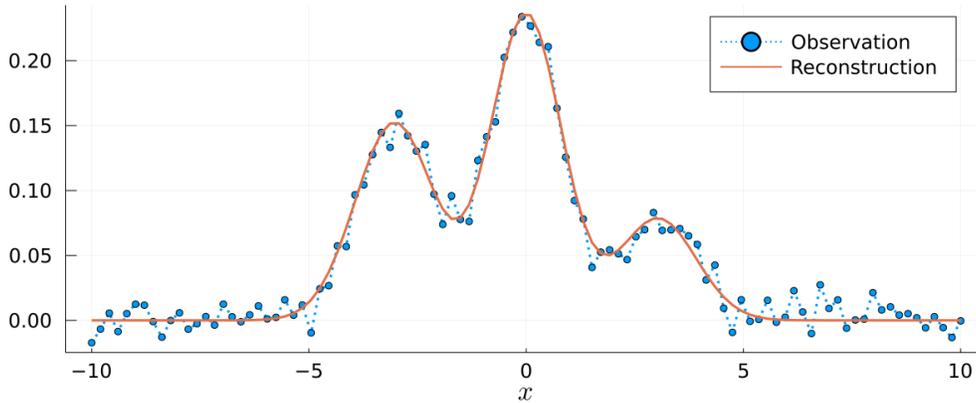


Figure 18: Noisy observations (blue dots) and solution to the Fokker–Planck equation (5.7) at time T when using the reconstructed initial condition.

To conclude this section, we compare in [Figure 19](#) the two discretization methods (5.5a) and (5.5b) in the case where $\nu = 10^{-8}$. All the ensemble members are initialized according to $\mathcal{N}(0, I_6)$. The error depicted in the left panel is computed as

$$\sum_{n=1}^3 |\bar{w}_n - w_n^\dagger| + |\bar{\sigma}_n^2 - (\sigma_n^2)^\dagger|,$$

where w_n^\dagger and $(\sigma_n^2)^\dagger$ are respectively the true weights and variances of the components of the initial Gaussian mixture, whereas \bar{w}_n and $\bar{\sigma}_n^2$ are the average weights and variances over the ensemble. As already remarked in the caption of [Figure 17](#), the error does not decrease to 0 as the number of iterations is increased, because the true value of the mixture parameters do not coincide with the minimizer of $f(w, v)$ in (5.9). As shown in the right panel, the sample variances decrease to a small value below 10^{-15} in less than a hundred iterations.

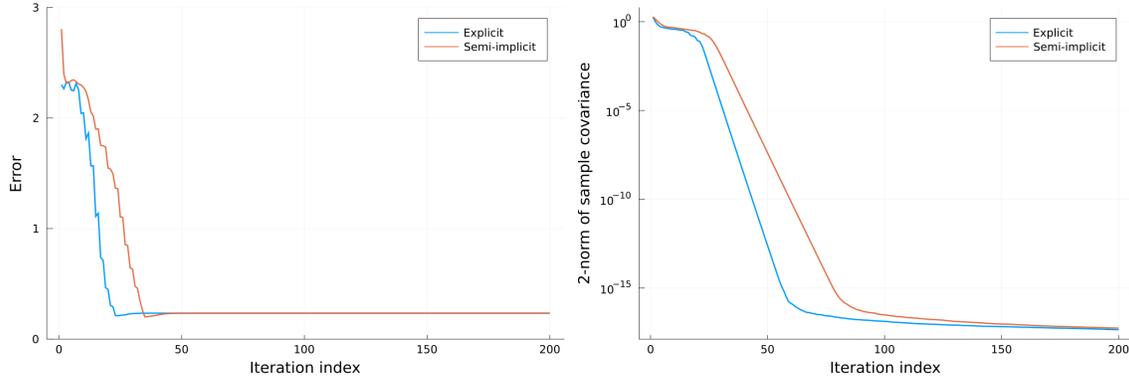


Figure 19: Evolution of the error (**left**) and sample covariance (**right**), for the explicit and semi-implicit discretizations described in Section 5.1. In both cases, the time step is adapted dynamically according to (5.2).

6 Conclusions

In this work, we study a new approach for incorporating constraints in consensus-based optimization and ensemble Kalman methods. We demonstrate that, despite their simplicity and ease of implementation, the proposed methods perform well for a number of toy examples. Our study is mostly qualitative and aims at providing a general idea of the performance of the methods, their behavior in the many-particle limit (in the case of CBO), and the influence of the parameters they contain.

Future investigation would be useful in order to determine the practical value of the proposed approach for realistic high-dimensional constrained optimization problems, and to further assess its efficiency in comparison with other methods in the literature. It would also be worthwhile to obtain quantitative results on the link between the particle systems and their mean-field limits, which could inform the choice of the number of particles employed in practice.

Acknowledgments

JAC was supported by the Advanced Grant Nonlocal-CPD (Nonlocal PDEs for Complex Particle Dynamics: Phase Transitions, Patterns and Synchronization) of the European Research Council Executive Agency (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 883363) and by EPSRC grant number EP/T022132/1. JAC and UV were also supported by EPSRC grant number EP/P031587/1. CT was partly supported by the the European social Fund and by the Ministry Of Science, Research and the Arts Baden-Württemberg (Germany). UV was also supported by the Fondation Sciences Mathématiques de Paris (FSMP), through a postdoctoral fellowship in the “mathematical interactions” program.

References

- [1] E. Aarts and J. Korst. *Simulated annealing and Boltzmann machines*. New York, NY; John Wiley and Sons Inc., 1988.
- [2] D. J. Albers, P.-A. Blancquart, M. E. Levine, E. Esmailzadeh Seylabi, and A. Stuart. Ensemble Kalman methods with constraints. *Inverse Problems*, **35**(9):095007, 28, 2019.
- [3] M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells. The FEniCS project version 1.5. *Archive of Numerical Software*, **3**(100), 2015.
- [4] T. Back, D. B. Fogel, and Z. Michalewicz. *Handbook of evolutionary computation*. IOP Publishing Ltd., 1997.
- [5] Z. Bayraktar, M. Komurcu, J. A. Bossard, and D. H. Werner. The wind driven optimization technique and its application in electromagnetics. *IEEE transactions on antennas and propagation*, **61**(5):2745–2757, 2013.
- [6] L. Bianchi, M. Dorigo, L. M. Gambardella, and W. J. Gutjahr. A survey on metaheuristics for stochastic combinatorial optimization. *Natural Computing: an international journal*, **8**(2):239–287, 2009.
- [7] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: overview and conceptual comparison. *ACM Computing Surveys (CSUR)*, **35**(3):268–308, 2003.
- [8] M. Bostan and J. A. Carrillo. Asymptotic fixed-speed reduced dynamics for kinetic equations in swarming. *Mathematical Models and Methods in Applied Sciences*, **23**(13):2353–2393, 2013.
- [9] L. Bottou et al. Online learning and stochastic approximations. *On-line learning in neural networks*, **17**(9):142, 1998.
- [10] S. Bubeck. Convex optimization: algorithms and complexity. *Foundations and Trends[®] in Machine Learning*, **8**:231–357, 2015.
- [11] J. A. Carrillo, Y.-P. Choi, C. Totzeck, and O. Tse. An analytical framework for consensus-based global optimization method. *Mathematical Models and Methods in Applied Sciences*, **28**(6):1037–1066, 2018.
- [12] J. A. Carrillo, F. Hoffmann, A. M. Stuart, and U. Vaes. Consensus Based Sampling. *arXiv e-prints*, **2106.02519**, 2021.
- [13] J. A. Carrillo and U. Vaes. Wasserstein stability estimates for covariance-preconditioned Fokker–Planck equations. *Nonlinearity*, **34**(4):2275–2295, Feb. 2021.
- [14] J. A. Carrillo, S. Jin, L. Li, and Y. Zhu. A consensus-based global optimization method for high dimensional machine learning problems. *ESAIM. Control, Optimisation and Calculus of Variations*, **27**(suppl.):Paper No. S5, 22, 2021.
- [15] N. K. Chada, C. Schillings, and S. Weissmann. On the Incorporation of Box-Constraints for Ensemble Kalman Inversion. *arXiv e-prints*, **1908.00696**, Aug. 2019.

- [16] B. Chandra Mohan and R. Baskaran. A survey: Ant Colony Optimization based recent research and implementation on several engineering domain. *Expert Systems with Applications*, **39**(4):4618–4627, 2012.
- [17] Y. Chen and D. S. Oliver. Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Mathematical Geosciences*, **44**(1):1–26, Jan. 2012.
- [18] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume **38**. Springer Science & Business Media, 2009.
- [19] M. Dorigo and C. Blum. Ant colony optimization theory: A survey. *Theoretical computer science*, **344**(2):243–278, 2005.
- [20] A. B. Duncan, A. M. Stuart, and M.-T. Wolfram. Ensemble Inference Methods for Models With Noisy and Expensive Likelihoods. *arXiv e-prints*, **2104.03384**, Apr. 2021.
- [21] A. A. Emerick and A. C. Reynolds. Investigation of the sampling performance of ensemble-based methods with a simple reservoir model. *Computational Geosciences*, **17**(2):325–350, 2013.
- [22] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based Optimization on the Sphere II: Convergence to Global Minimizers and Machine Learning. *arXiv e-prints*, **2001.11988**, Jan. 2020.
- [23] M. Fornasier, T. Klock, and K. Riedl. Consensus-based optimization methods converge globally in mean-field law. *arXiv e-prints*, **2103.15130**, 2021.
- [24] M. Fornasier, H. Huang, L. Pareschi, and P. Sünnen. Consensus-based optimization on hypersurfaces: well-posedness and mean-field limit. *Mathematical Models and Methods in Applied Sciences*, **30**(14):2725–2751, 2020.
- [25] A. Garbuno-Inigo, F. Hoffmann, W. Li, and A. M. Stuart. Interacting Langevin diffusions: gradient structure and ensemble Kalman sampler. *SIAM Journal on Applied Dynamical Systems*, **19**(1):412–441, 2020.
- [26] A. Garbuno-Inigo, N. Nüsken, and S. Reich. Affine invariant interacting Langevin dynamics for Bayesian inference. *SIAM Journal on Applied Dynamical Systems*, **19**(3):1633–1658, 2020.
- [27] C. Geuzaine and J.-F. Remacle. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, **79**(11):1309–1331, 2009.
- [28] M. Herty and G. Visconti. Continuous limits for constrained ensemble Kalman filter. *Inverse Problems*, **36**(7):075006, 28, 2020.
- [29] D. Higham. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, **43**(3):525–546.
- [30] R. Holley and D. Stroock. Simulated annealing via Sobolev inequalities. *Communications in Mathematical Physics*, **115**(4):553–569, 1988.

- [31] R. A. Holley, S. Kusuoka, and D. W. Stroock. Asymptotics of the spectral gap with applications to the theory of simulated annealing. *Journal of functional analysis*, **83**(2):333–347, 1989.
- [32] H. Huang and J. Qiu. On the mean-field limit for the consensus-based optimization. *arXiv e-prints*, **2105.12919**, 2021.
- [33] C.-R. Hwang and S.-J. Sheu. Large-time behavior of perturbed diffusion markov processes with applications to the second eigenvalue problem for fokker-planck operators and simulated annealing. *Acta Applicandae Mathematica*, **19**(3):253–295, 1990.
- [34] M. A. Iglesias, K. J. H. Law, and A. M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, **29**(4):045001, 20, 2013.
- [35] M. Jamil and X.-S. Yang. A literature survey of benchmark functions for global optimisation problems. *International Journal of Mathematical Modelling and Numerical Optimisation*, **4**(2):150–194, 2013.
- [36] S. Jin, L. Li, and J.-G. Liu. Convergence of the random batch method for interacting particles with disparate species and weights. *SIAM Journal on Numerical Analysis*, **59**(2):746–768, 2021.
- [37] S. Jin, L. Li, and J.-G. Liu. Random batch methods (RBM) for interacting particle systems. *Journal of Computational Physics*, **400**:108877, 30, 2020.
- [38] D. Karaboga, B. Gorkemli, C. Ozturk, and N. Karaboga. A comprehensive survey: artificial bee colony (ABC) algorithm and applications. *Artificial Intelligence Review*, **42**(1):21–57, 2014.
- [39] J. Kennedy. Particle swarm optimization. In *Encyclopedia of Machine Learning*, pages 760–766. Springer, 2010.
- [40] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN’95-international conference on neural networks*, volume **4**, pages 1942–1948. IEEE, 1995.
- [41] N. B. Kovachki and A. M. Stuart. Ensemble Kalman inversion: a derivative-free technique for machine learning tasks. *Inverse Problems*, **35**(9):095005, 35, 2019.
- [42] D. V. Kulkarni, D. V. Rovas, and D. A. Tortorelli. A discontinuous Galerkin formulation for solution of parabolic equations on nonconforming meshes. In *Domain decomposition methods in science and engineering XVI*. volume **55**, Lect. Notes Comput. Sci. Eng. Pages 651–658. Springer, Berlin, 2007.
- [43] D. Márquez. Convergence rates for annealing diffusion processes. *The Annals of Applied Probability*:1118–1139, 1997.
- [44] P. D. Miller. *Applied asymptotic analysis*, volume **75**. American Mathematical Soc., 2006.
- [45] G. A. Pavliotis. *Stochastic processes and applications*, volume **60** of *Texts in Applied Mathematics*. Springer, New York, 2014, pages xiv+339.

- [46] M. Pelletier. Weak convergence rates for stochastic approximation with application to multiple targets and simulated annealing. *Annals of Applied Probability*:10–44, 1998.
- [47] R. Pinnau, C. Totzeck, O. Tse, and S. Martin. A consensus-based model for global optimization and its mean-field limit. *Mathematical Models and Methods in Applied Sciences*, **27**(1):183–204, 2017.
- [48] R. Poli, J. Kennedy, and T. Blackwell. Particle swarm optimization. *Swarm intelligence*, **1**(1):33–57, 2007.
- [49] C. Reeves. *Genetic algorithms*. Springer, 2003.
- [50] H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*:400–407, 1951.
- [51] C. Schillings and A. M. Stuart. Analysis of the ensemble Kalman filter for inverse problems. *SIAM Journal on Numerical Analysis*, **55**(3):1264–1290, 2017.
- [52] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, **19**:451–559, 2010.
- [53] C. Totzeck. Trends in consensus-based optimization. *arXiv e-prints*, **2104.01383**, 2021.
- [54] M. Weber and S. Sra. Projection-free nonconvex stochastic optimization on Riemannian manifolds, 2021.
- [55] X.-S. Yang. Firefly algorithms for multimodal optimization. In O. Watanabe and T. Zeugmann, editors, *Stochastic Algorithms: Foundations and Applications*, pages 169–178, Berlin, Heidelberg. Springer Berlin Heidelberg, 2009.
- [56] H. Zhang, S. J Reddi, and S. Sra. Riemannian SVRG: fast stochastic optimization on Riemannian manifolds. *Advances in Neural Information Processing Systems*, **29**:4592–4600, 2016.