

Less interaction with forward models in Langevin dynamics

Martin Eigel, Robert Gruhlke, David Sommer

submitted: December 21, 2022

Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: martin.eigel@wias-berlin.de
robert.gruhlke@wias-berlin.de
david.sommer@wias-berlin.de

No. 2987
Berlin 2022



2020 Mathematics Subject Classification. 65N21, 62F15, 65N75, 65C30, 35Q84, 65K10.

Key words and phrases. Langevin dynamics, interacting particle systems, Bayesian inference, computational optimal transport, Wasserstein distance, mean-field Fokker-Planck equation, homotopy.

We thank Christian Bayer, Paul Hagemann, Matthias Liero and Claudia Schillings for fruitful discussions on the topics of this paper. ME acknowledges the partial support by the DFG SPP 2998 “Theoretical Foundations of Deep Learning”. RG acknowledges the support by the DFG SPP 1886 “Polymorphic Uncertainty Modelling for the Numerical Design of Structures”. DS acknowledges support by the ProFIT project “ReLkat – Reinforcement Learning for complex automation engineering”. This study does not have any conflicts to disclose.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: <http://www.wias-berlin.de/>

Less interaction with forward models in Langevin dynamics

Martin Eigel, Robert Gruhlke, David Sommer

Abstract

Ensemble methods have become ubiquitous for the solution of Bayesian inference problems. State-of-the-art Langevin samplers such as the Ensemble Kalman Sampler (EKS), Affine Invariant Langevin Dynamics (ALDI) or its extension using weighted covariance estimates rely on successive evaluations of the forward model or its gradient. A main drawback of these methods hence is their vast number of required forward calls as well as their possible lack of convergence in the case of more involved posterior measures such as multimodal distributions. The goal of this paper is to address these challenges to some extent. First, several possible adaptive ensemble enrichment strategies that successively enlarge the number of particles in the underlying Langevin dynamics are discussed that in turn lead to a significant reduction of the total number of forward calls. Second, analytical consistency guarantees of the ensemble enrichment method are provided for linear forward models. Third, to address more involved target distributions, the method is extended by applying adapted Langevin dynamics based on a homotopy formalism for which convergence is proved. Finally, numerical investigations of several benchmark problems illustrate the possible gain of the proposed method, comparing it to state-of-the-art Langevin samplers.

1 Introduction

Consider the inverse problem of finding an unknown $y \in \mathbb{R}^D$ from an observation $\Delta \in \mathbb{R}^K$ for $D, K \in \mathbb{N}$, where

$$\Delta = \mathcal{G}(y) + \eta, \quad (1.1)$$

with a deterministic forward operator $\mathcal{G}: \mathbb{R}^D \rightarrow \mathbb{R}^K$ and centered Gaussian observational noise $\eta \sim \mathcal{N}(0, \Gamma)$ with positive definite covariance matrix $\Gamma \in \mathbb{R}^{K, K}$. In the Bayesian framework, a prior distribution μ_{prior} is associated with the unknown y . Then, for a given measurement $\Delta = \tilde{\Delta}$, the prior is updated via Bayes' rule to yield a posterior distribution μ_* [24]. In case that μ_{prior} has a Lebesgue density π_{prior} , then under mild assumptions [41] there exists a posterior density π_* , which is the Lebesgue density of μ_* given by

$$\pi_*(y) \propto \exp(-L(y))\pi_{\text{prior}}(y), \quad (1.2)$$

with the log-likelihood potential $L(y) := 1/2|\tilde{\Delta} - \mathcal{G}(y)|_{\Gamma}^2$. Here, $|\cdot|$ denotes the standard Euclidean norm and $|\cdot|_M = |M^{-1/2} \cdot|$ for any symmetric positive definite M . If the prior density is positive on \mathbb{R}^D , we define the potential

$$\Phi(y) := L(y) - \log \pi_{\text{prior}}(y), \quad (1.3)$$

such that the posterior density becomes

$$\pi_*(y) \propto \exp(-\Phi(y)). \quad (1.4)$$

In this work we are concerned with the problem of sampling from the posterior distribution given by (1.4) using interacting particle methods based on Langevin dynamics. While non-interacting particle systems exhibit slow convergence in time, their extensions to interacting particle systems such as

ALDI and the EKS have demonstrated a superior convergence speed. However, such state-of-the-art methods still require a vast number of forward calls to solve the underlying model. Furthermore, convergence to more involved posteriors such as multimodal distributions may become arbitrarily slow or cannot be guaranteed at all. To address these drawbacks, we propose to extend existing methods by the following two strategies.

The first concept is *ensemble enrichment*. It allows to work with ensembles of small batch sizes for large parts of the process, utilizing the contained information to build larger ensembles with the desired distribution at a later time. This substantially reduces the number of necessary calls of the forward model $\mathcal{G}(\cdot)$.

The second concept is based on the notion of *homotopy*. Here, instead of directly working with a particle system based on the posterior, we utilize intermediate measures obtained from interpolation between a simple auxiliary measure and the posterior. Such a preconditioning of the particle ensemble potentially increases the convergence speed over time significantly, especially for multimodal distributions. We coin this new methodology LIDL¹.

1.1 Related work

There is a vast amount of literature on different sampling methods such as Markov Chain Monte Carlo methods (MCMC) [36, 3, 35] and more recently methods based on Langevin dynamics [37] or Stein variational gradient descent [26], to name just a few. While the idea of ensemble enrichment can in principle be deployed for any ensemble sampler, in this work we focus on the class of Langevin based samplers. This is due to the fact that Langevin samplers lend themselves quite naturally to ensemble enrichment schemes like *slicing* (adding together ensembles defined by the Langevin process at different time points) and appropriately scaled *random kicks* (adding noisy duplicates of the existing batch, such that the covariance structure is preserved).

In analytical chemistry, the term *sample enrichment* is common to denote certain ways of manipulating samples to give them desirable properties [40, 7]. In order to differentiate from this term and to emphasize the connection to particle ensemble methods, we coin our form of enrichment *ensemble enrichment*. In this work, we understand ensemble enrichment as the addition of new samples $\{y^i\}_{i=n+1}^N$ to an existing sample batch $\{y^i\}_{i=1}^n$. In a very general setting, an enrichment can be considered admissible if both sample batches are drawn independently from the same distribution. Depending on the context however, weaker conditions may be sufficient. In the setting of Langevin sampling considered here, an enrichment is suitable if the distance from the posterior to the enriched sample batch can be bounded by the respective distance of the batch before enrichment.

The use of Langevin dynamics in Bayesian inference requires the posterior to be an invariant measure of the chosen dynamics. Historically, the most basic dynamics with this property is given by the first order overdamped Langevin equation [30]. As a state-of-the-art Langevin method, we use the Affine Invariant Langevin Dynamics (ALDI) [16] in this work. ALDI is a modification of the Ensemble Kalman Sampler (EKS) [33], which ensures affine invariance [18] as well as convergence in total variation to the posterior, even with a finite number of particles. Ensemble Langevin methods such as EKS and ALDI have strong links to ensemble Kalman filters (EKF) [11, 25, 32] and ensemble Kalman inversion (EKI) [23], which was pointed out in [33]. A recent survey of ensemble Kalman methods and their application to Bayesian inverse problems can be found in [4]. The analysis of EKF and EKI

¹An acronym meaning *Less Interaction with forward models in Langevin Dynamics*. It should be noted that the first author has objections against this term but was overruled by majority vote.

for Bayesian inversion has been extensively studied [38, 39, 1, 8]. In [9] convergence of the EKS to the limiting Fokker-Planck equation in expected 2-Wasserstein distance is shown. To the best of our knowledge, a comparable result for ALDI has not been shown yet.

1.2 Contribution

- We propose a modification of existing Langevin samplers, coined as LIDL, utilizing successive ensemble enrichment to substantially reduce the number of required forward calls in the Langevin dynamics. The effect of ensemble enrichment is explained in Section 2, ensemble enrichment strategies are introduced in Section 3 and we present the analysis of the method in Section 4. For the family of ensemble distributions $(\hat{\mu}_t)_{t \geq 0}$ generated by the Langevin process, we show analytical consistency of the scheme in the following sense: under suitable consistency assumptions on the initial density π_0 (which need not coincide with the prior), the underlying potential Φ and the ensemble enrichment scheme, we find for $\delta > 0$ a configuration of the enrichment and time $T_\delta > 0$ such that the expected 2-Wasserstein distance of the solution to the posterior is smaller than δ , i.e.,

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_{T_\delta}, \mu_*)] < \delta.$$

This result is a combination of the corresponding consistency result for ALDI, shown in Theorem 4.1, and the definition of consistent enrichment schemes given in Definition 4.2, yielding the main result in Corollary 4.2.

- We introduce the concept of *ensemble enrichment*, and provide a rigorous mathematical description of ensemble enrichment schemes \mathcal{E} as stochastic mappings operating on families of empirical measures $(\hat{\mu}_t)_{t \geq 0}$ in Section 3. For carrying out the enrichment in Langevin processes, we propose several informed choices of enrichments, namely *slicing* (adding up time slices $\hat{\mu}_{t \pm \Delta t}$ of the process at different times, Section 3.1.1), *diffusion propagation* (propagating the current ensemble using only the computationally cheap diffusion part of the Langevin process, Section 3.1.2) and *generalized transport* (sampling i.i.d. from a learned random variable $Y \sim \hat{\mu}_t$ fitted to the current ensemble, Section 3.1.4).
- In order to accelerate the convergence of Langevin dynamics in the case of multimodal invariant measures, we introduce the terminology of *2-Wasserstein stable homotopy maps* \mathcal{H} in Section 4.2. Such a function maps between an auxiliary initial potential $\mathcal{H}(0) = \Psi$ and the posterior potential $\mathcal{H}(1) = \Phi$. This framework enables to construct a time inhomogeneous drift term $b_{\mathcal{H}}(t, Y_t) = b(\mathcal{H}(s(t)), Y_t)$ in the underlying Langevin dynamics that reads

$$dy_t^{(i)} = b_{\mathcal{H}}(t, Y_t)dt + \Gamma(Y_t)dW_t^{(i)}.$$

Based on Assumption 4.1 of local convergence in expected Wasserstein-2 distance of the associated particle system, we provide a convergence analysis in the case of piecewise constant inhomogeneity in Theorem 4.2. We show by numerical evidence that different homotopy designs affect the overall convergence speed over time, especially for multimodal distributions. The suggested homotopy approach can be conceived as a preconditioning tool.

- The final contribution concerns the numerical investigation of our method. Based on the consistency results of Corollary 4.2 and Theorem 4.2, we propose in Section 5 a principled way of

tracking convergence of numerical realizations of the method in debiased Sinkhorn divergence (3.14). More precisely, we consider the convergence in distribution

$$\mathcal{S}_\epsilon(\hat{\mu}_t, \mu_*^{(\bar{b})}) \xrightarrow{d} \mathcal{S}_\epsilon(\tilde{\mu}_*^{(\bar{b})}, \mu_*^{(\bar{b})}), \quad t \rightarrow \infty,$$

where $\hat{\mu}_t$ is the ensemble distribution generated by the Langevin process (producing a total of \bar{b} approximate posterior samples), $\mu_*^{(\bar{b})}$ and $\tilde{\mu}_*^{(\bar{b})}$ are empirical measures of $\bar{b} \in \mathbb{N}$ particles drawn i.i.d. from the posterior. Moreover, initial ideas for adaptive ensemble enrichment schemes and homotopy designs are presented in Sections 5.1.2 and 5.1.3. We investigate the performance of the proposed method on several numerical test problems in Sections 5.1 and 5.2. This includes a challenging multimodal posterior and a high-dimensional Darcy equation often used as benchmark in Uncertainty Quantification.

2 Methodology

Suppose we are interested to obtain a batch of $\bar{b} \in \mathbb{N}$ samples drawn from the posterior distribution μ_* . Suppose further that we have a method which, relying on the evaluation of the potential Φ or its gradient $\nabla\Phi$, maps \bar{b} particles/samples drawn from some initial measure μ_0 to particles approximately being distributed according to $\otimes_{i=1}^{\bar{b}} \mu_*$ in N_{iter} iterations. We will refer to this method as the *particle propagator*. One such particle propagator is given by Langevin dynamics based particle systems [37, 15, 16, 33]. Here, N_{iter} depends on the distance $d(\mu_0, \mu_*)$ between μ_0 and μ_* , where d is a suitable metric on the set of probability measures. Note that we do not require the initial measure μ_0 to be identical to the prior μ_{prior} although this is a natural choice in the setup of Bayesian inference. We assume one iteration step of \bar{b} particles to require \bar{b} forward calls (as is the case e.g. for ALDI [16, 33]), leading to a total workload of

$$\#\text{forward calls} = \bar{b}N_{\text{iter}}.$$

Our goal throughout this manuscript is to reduce the number of forward calls significantly. For this we propose to construct a sequence of auxiliary measures from which an increasing number of particles are drawn over time and eventually propagated to the sought posterior distribution. We may formally divide this approach into two stages. The first stage deals with the design choice of the underlying particle system in terms of a modification of the potential Φ through the introduction of surrogates or homotopy maps. The second stage concerns the propagation of particles to the true posterior using successive ensemble enrichment. In practical application, both stages are intertwined in that ensemble enrichment may be performed while using variations of the potential.

Stage I: Auxiliary measures Initially, one chooses an accessible distribution μ_0 , which not necessarily coincides with the prior distribution. Next, some auxiliary potential Ψ is introduced and instead of directly computing with Φ , we use the intermediate potentials

$$\mathcal{H}(s) := (1 - s)\Psi + s\Phi, \quad s \in [0, 1], \quad (2.1)$$

replacing the potential Φ in the particle propagator. This homotopy approach carries out a linear interpolation from the potential $\Psi = \mathcal{H}(0)$ to $\Phi = \mathcal{H}(1)$ in this example. By using this potential, we introduce intermediate target measures with density proportional to $\exp(-\mathcal{H}(s))$. From the perspective of the particle dynamics, the particles are moved towards the final target measure, i.e. the

posterior distribution defined by $\mathcal{H}(1) = \Phi$, by passing through the intermediate measures associated with $\mathcal{H}(s)$. From the classical optimization point of view, this type of homotopy can be seen as a sequence of Tikhonov regularizations with parameters $(1 - s)/s \in [0, \infty]$. The effect of different designs of interpolations between Ψ and Φ is discussed in the numerical examples in Section 5.1.3. It turns out that this approach is crucial to enable and accelerate the propagation of particles to multimodal distributions with our approach. We refer to Figure 2.2 for an illustration of the technique and to Section 4.2 for a more in depth discussion.

We now discuss potential choices of the start distribution μ_0 . A canonical first choice would be $\mu_0 = \mu_{\text{prior}}$. However, this might not be ideal if the prior distribution is not a good approximation of the posterior. A reasonable alternative is to choose μ_0 as a Gaussian approximation of the posterior distribution.

As a third option, assume that a surrogate model of the forward operator \mathcal{G} or of the potential Φ is available. We further assume that it takes b_{surr} evaluations of the exact forward map to construct this surrogate and that the cost of evaluating the surrogate potential $\hat{\Phi}$ is negligible compared to the cost of evaluating Φ . Then, we choose μ_0 with density proportional to $\exp(-\hat{\Phi})$. Consequently, samples can be drawn with negligible workload through a particle propagator using the surrogate potential $\hat{\Phi}$. This idea is illustrated in Figure 2.1. Depending on the quality of the surrogate, such a μ_0 is a potentially better proxy to μ_* than μ_{prior} in the sense that

$$d(\mu_0, \mu_*) < d(\mu_{\text{prior}}, \mu_*). \quad (2.2)$$

We refer to evaluations of both the surrogate potential and $\mathcal{H}(0)$ as *free calls*.

In applications the exact solution of the forward model \mathcal{G} is often unknown and only approximations $\mathcal{G}_h \approx \mathcal{G}$ are available. The influence of this approximation on the convergence to the true posterior has e.g. been considered in [5]. Note that successive improvement of the approximation during particle propagation potentially leads to significant total workload reduction. However, this idea and its interaction with the concepts presented in this paper is beyond the scope of this work and we only comment on it in Section 6.

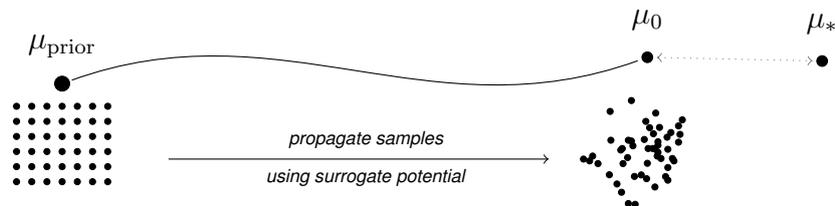


Figure 2.1: Stage I: Samples drawn from the prior measure μ_{prior} are propagated to μ_0 , the approximate posterior measure introduced through the substitution by the chosen surrogate.

Stage II: Intermediate ensemble enrichment In the second stage of our method, the exact forward map is used. We start by propagating b_0 samples drawn from μ_0 obtained by the Stage I procedure and iterate them $n_0 \in \mathbb{N}$ times with the exact forward model. This yields b_0 samples distributed according to some intermediate measure $\mu_1^{(b_0)}$ defined on $\mathbb{R}^{b_0 D}$. Since we have now used the true potential, it is reasonable to assume that

$$d\left(\mu_1^{(b_0)}, \bigotimes_{i=1}^{b_0} \mu_*\right) < d\left(\mu_0^{(b_0)}, \bigotimes_{i=1}^{b_0} \mu_*\right), \quad (2.3)$$

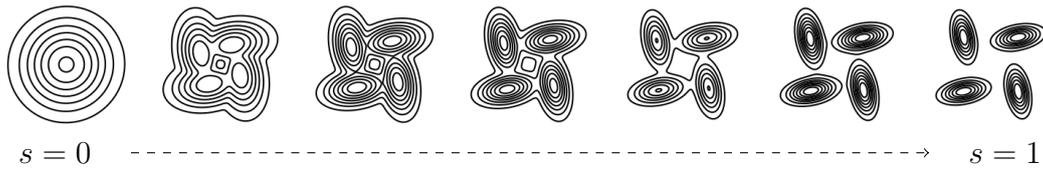


Figure 2.2: Illustration of the homotopy approach towards a multimodal distribution with Ψ chosen as the Gaussian approximation for $h = 0$ on the left-hand side to the target distribution on the right-hand side for $h = 1$.

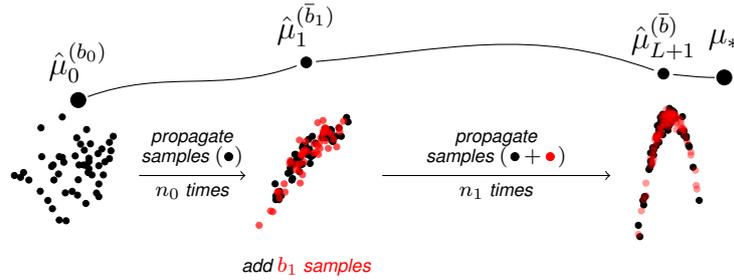


Figure 2.3: Stage II: Samples are propagated and enriched a single time, i.e. $L = 1$, eventually leading to $\bar{b} = \bar{b}_1 = b_0 + b_1$ samples defining the final empirical measure $\hat{\mu}_{L+1}^{(\bar{b})}$.

where $\mu_0^{(b_0)} = \bigotimes_{i=1}^{b_0} \mu_0$ defines a product measure on $\mathbb{R}^{b_0 D}$. Within the ensemble enrichment, we change the viewpoint from distributions on the product space to ensemble distributions on \mathbb{R}^D . Throughout this manuscript, we shall use the wedge symbol $\hat{\cdot}$ to flag a distribution as an empirical measure or an ensemble distribution. For instance, the intermediate distribution $\mu_1^{(b_0)}$ on the product space $\mathbb{R}^{b_0 D}$ gives rise to an empirical (random) measure $\hat{\mu}_1^{(b_0)}$ on \mathbb{R}^D .

Next, we define a *number of enrichments* $L \in \mathbb{N}_0$ and *batch sizes* $b_0, \dots, b_L \in \mathbb{N}$ with $\bar{b} = \sum_{\ell=0}^L b_\ell$. Furthermore, we denote by $\bar{b}_k = \sum_{\ell=0}^k b_\ell$ the *partial summed batch sizes*. Formally, an *ensemble enrichment scheme* \mathcal{E} produces an empirical measure $\hat{\mu}_{\text{enriched}}^{(b+a)}$ from an empirical measure $\hat{\mu}_{\text{original}}^{(b)}$ for some arbitrary $b \in \mathbb{N}$ and a dependent number of additional samples $a \in \mathbb{N}$. For now, we tacitly assume that both distributions are close enough in the sense that

$$\mathbb{E} \left[d \left(\hat{\mu}_{\text{original}}^{(b)}, \hat{\mu}_{\text{enriched}}^{(b+a)} \right) \right] < \epsilon \quad (2.4)$$

for some suitable threshold $\epsilon > 0$. We then apply \mathcal{E} to the current batch of b_0 samples at level $\ell = 1$ to draw b_1 additional samples, yielding a total of \bar{b}_1 samples distributed according to some measure $\hat{\mu}_1^{(\bar{b}_1)}$, which is close to $\hat{\mu}_1^{(b_0)}$ in the sense of (2.4). This enriched batch is propagated through the dynamics $n_1 \in \mathbb{N}$ times, yielding \bar{b}_1 samples distributed according to an intermediate measure $\hat{\mu}_2^{(\bar{b}_1)}$. The process continues iteratively for $\ell = 2, \dots, L$: at level ℓ the samples distributed according to the intermediate measure $\mu_\ell^{(\bar{b}_{\ell-1})}$ are enriched with b_ℓ samples by application of \mathcal{E} . The new batch of b_ℓ samples is distributed according to $\hat{\mu}_\ell^{(\bar{b}_\ell)}$. This batch is then iterated $n_\ell \in \mathbb{N}$ times through the dynamics, leading to \bar{b}_ℓ samples drawn from a measure $\hat{\mu}_{\ell+1}^{(\bar{b}_\ell)}$. Finally, we get an ensemble distributed according to $\hat{\mu}_{L+1}^{(\bar{b})}$. The described process is sketched for the case of a single enrichment (i.e. $L = 1$)

in Figure 2.3. With $n_{\text{iter}} := \sum_{\ell=0}^L n_{\ell}$, the number of forward calls in this scheme is given as

$$\#\text{forward calls} = b_0 n_{\text{iter}} + b_1 (n_{\text{iter}} - n_0) + \dots + b_L (n_{\text{iter}} - \sum_{\ell=0}^{L-1} n_{\ell}) = \bar{b} n_{\text{iter}} - \sum_{\ell=1}^L b_{\ell} \sum_{k=0}^{\ell-1} n_k. \quad (2.5)$$

To quantify the computational cost reduction, let $b_{\ell} = p_{\ell} \bar{b}$ for a percentile $p_{\ell} \in (0, 1)$ such that $\sum_{\ell=0}^L p_{\ell} = 1$ and let $n_{\ell} = c_{\ell} n_{\text{iter}}$ for some $c_{\ell} \in (0, 1)$ with $\sum_{i=0}^L c_{\ell} = 1$. Then, (2.5) becomes

$$\#\text{forward calls} = \bar{b} n_{\text{iter}} \underbrace{\left(1 - \sum_{\ell=1}^L p_{\ell} \sum_{k=0}^{\ell-1} c_k \right)}_{=:c}. \quad (2.6)$$

Hence, a reduction factor of $c \in [1/n_{\text{iter}}, 1)$ is obtained that only depends on the parameters L , p_{ℓ} and c_{ℓ} for $\ell = 0, \dots, L$. Eventually, the number of forward map evaluations is effectively reduced to

$$\#\text{forward calls} = b_{\text{surr}} + c \bar{b} n_{\text{iter}}. \quad (2.7)$$

There are two main design aspects of the method that require a further analysis: first, the choice of the particle propagator and second, how to perform the ensemble enrichment. The remainder of this section is concerned with answering the first question. For this, we provide an overview of Langevin sampling methods, which are at the center of our method. The question of how to generate new samples is discussed subsequently in Section 3.

2.1 Langevin sampling methods

The class of particle propagators considered in this work is based on interacting particle systems. The starting point is a *single-particle* first order overdamped Langevin process of the form

$$dy_t = -\nabla \Phi(y_t) dt + \sqrt{2} dW_t, \quad (2.8)$$

where W_t is D -dimensional Brownian motion and the potential Φ is defined by the Bayesian inference problem via (1.3). The probability density function π_t of the particle y_t at time $t \geq 0$ satisfies the Fokker-Planck equation [29]

$$\partial_t \pi_t = \nabla \cdot (\pi_t \nabla (\log \pi_t - \log \pi_*)). \quad (2.9)$$

It can be seen immediately that π_* is a stationary solution, in particular the right-hand side of (2.9) equals zero for $\pi_t = \pi_*$. Hence, the posterior μ_* is an invariant measure of the process (2.8). Furthermore, a desirable property of the process is ergodicity. This means that μ_* is the only invariant measure and all initial measures converge to μ_* in a suitable sense as $t \rightarrow \infty$. Consequently, y_t defined by (2.8) is distributed according to the posterior in the limit $t \rightarrow \infty$. Ergodicity is ensured, for example, under certain growth conditions on the potential Φ [16].

There exists a vast literature on extensions to dynamics of the form (2.8) usually with the goal of accelerating convergence to the posterior by introducing interaction between multiple particles [15, 16, 33]. Collecting $B \in \mathbb{N}$ particles $\{y_t^{(i)}\}_{i=1}^B$ at time t into a vector

$$Y_t = \text{vec}(y_t^{(1)}, y_t^{(2)}, \dots, y_t^{(B)}) \in \mathbb{R}^{DB}, \quad t \geq 0,$$

many interacting particle approaches admit the general form

$$dy_t^{(i)} = -A(Y_t)\nabla_{y_t^{(i)}}\mathcal{V}(Y_t)dt + \Gamma(Y_t)dW_t^{(i)} \quad \text{for } i = 1, \dots, B. \quad (2.10)$$

Here, $A(Y_t) \in \mathbb{R}^{D \times D}$, $\Gamma(Y_t) \in \mathbb{R}^{D \times J}$ for some $J \in \mathbb{N}$, $W_t^{(i)}$ are independent J -dimensional Brownian motions for $i = 1, \dots, J$ and $\mathcal{V}: \mathbb{R}^{DB} \rightarrow \mathbb{R}$ usually depends on the potential Φ . For convenience, we introduce a block notation for (2.10), which becomes

$$dY_t = -\mathbf{A}(Y_t)\nabla_{Y_t}\mathcal{V}(Y_t)dt + \mathbf{\Gamma}(Y_t)dW_t, \quad (2.11)$$

where $\mathbf{A}(Y_t) \in \mathbb{R}^{DB \times DB}$ and $\mathbf{\Gamma}(Y_t) \in \mathbb{R}^{DB \times JB}$ are block matrices with all B blocks equal to $A(Y_t)$ and $\Gamma(Y_t)$ respectively. Moreover,

$$\nabla_{Y_t}\mathcal{V}(Y_t) = \text{vec}\left(\nabla_{y_t^{(i)}}\mathcal{V}(Y_t), i = 1, \dots, B\right) \quad (2.12)$$

and W_t is JB -dimensional Brownian motion.

Choosing $A \equiv I_d$, $\Gamma \equiv \sqrt{2}I_d$ and $\mathcal{V}(Y_t) = \sum_i^B \Phi(y_t^{(i)})$ in (2.10) leads to a particle system where each particle follows the process (2.8) independently. In this case, there is no interaction between the particles. Moreover, the system lacks *affine invariance* [18], meaning that it does not retain its convergence properties under linear affine transformations of the state variables. Both properties can be important to enable accelerated convergence [16], e.g. in the case of multimodal or highly skewed posteriors. We hence review some more involved methods of the form (2.10) that address these issues. For this, we define the time-dependent ensemble mean

$$\bar{y}_t = \frac{1}{B} \sum_{i=1}^B y_t^{(i)} \in \mathbb{R}^D, \quad (2.13)$$

and

$$\bar{Y}_t = \text{vec}(\bar{y}_t, \dots, \bar{y}_t) \in \mathbb{R}^{DB}, \quad (2.14)$$

as well as the empirical covariance

$$C(Y_t) = \frac{1}{B} \sum_{i=1}^B (y_t^{(i)} - \bar{y}_t)(y_t^{(i)} - \bar{y}_t)^T \in \mathbb{R}^{D \times D}. \quad (2.15)$$

Furthermore, we define the product posterior density

$$\tilde{\pi}_*(Y) = \prod_{i=1}^B \pi_*(y^{(i)}), \quad \text{for } Y = \text{vec}(y^{(1)}, \dots, y^{(B)}) \in \mathbb{R}^{DB}, \quad (2.16)$$

with $y^{(i)} \in \mathbb{R}^D$ for $i = 1, \dots, B$. In what follows, we briefly review some extensions of the classic Langevin dynamics from (2.8) in historical order.

Scaled first order overdamped Langevin dynamics As a first step towards affine invariance, (2.8) is extended by the introduction of a positive definite scaling matrix $C \in \mathbb{R}^{D \times D}$. The choice $A \equiv C$, $\Gamma \equiv \sqrt{2}C^{1/2}$ and $\mathcal{V}(Y) = \sum_i^B \Phi(y^{(i)})$ in (2.10) leads to the dynamics

$$dy_t^{(i)} = -C\nabla\Phi(y_t^{(i)})dt + \sqrt{2}C^{1/2}dW_t \quad \text{for } i = 1, \dots, B. \quad (2.17)$$

Ideally, C should be close to the (unknown) posterior covariance matrix. The Fokker-Planck equation for the PDF $\pi_t^{(i)}$ of the i -th particle at time t now reads

$$\partial_t \pi_t = \nabla \cdot (\pi_t C \nabla (\log \pi_t - \log \pi_*)), \quad (2.18)$$

retaining the posterior μ_* as invariant measure. This system has no interaction between particles and is not affine invariant for time-homogeneous C . However, affine invariance can be achieved for $C = C(t)$ equal to the process covariance matrix [16]. This is a key observation, laying the groundwork for the following sections.

Ensemble Kalman Sampler (EKS, cf. [15]) The EKS employs a time dependent scaling $C = C(Y_t)$ given by the empirical covariance. In particular, this scaling depends nonlinearly on the current ensemble. Setting $A(Y_t) = C(Y_t)$, $\Gamma(Y_t) = \sqrt{2}C(Y_t)^{1/2}$ and $\mathcal{V}(Y_t) = \sum_i^B \Phi(y_t^{(i)})$ in (2.10), the process equations become

$$dy_t^{(i)} = -C(Y_t) \nabla \Phi(y_t^{(i)}) dt + \sqrt{2}C(Y_t)^{1/2} dW_t^{(i)}, \quad i = 1, \dots, B. \quad (2.19)$$

This choice of scaling is motivated by the large particle limit $B \rightarrow \infty$, which formally leads to the mean field equation

$$dy_t = -C(\pi_t) \nabla \Phi(y_t) dt + \sqrt{2}C(\pi_t)^{1/2} dW_t \quad (2.20)$$

with the true process covariance $C(\pi_t) = \mathbb{E}_{y \sim \pi_t} [(y - \mathbb{E}_{\pi_t})(y - \mathbb{E}_{\pi_t})^\top]$. Here, $\mathbb{E}_\rho = \int_{\mathbb{R}^D} y \rho(y) dy$ denotes the first moment of a probability density ρ on \mathbb{R}^D . The corresponding Fokker-Planck equation

$$\partial_t \pi_t = \nabla \cdot (\pi_t C(\pi_t) \nabla (\log \pi_t - \log \pi_*)) \quad (2.21)$$

and its associated process (2.20) are shown to be affine invariant in [16], while retaining the invariant measure μ_* .

Affine Invariant Langevin Dynamics (ALDI, cf. [16]) It is shown in [28] and further elaborated on in [16] that the posterior μ_* is actually not invariant under the particle system (2.19), i.e. in the case of finitely many particles. Instead, the finite ensemble version (2.19) requires an additional correction term in the potential. To see this, note that in block notation (2.19) becomes

$$dY_t = C(Y_t) \nabla \ln \tilde{\pi}_*(Y_t) dt + \sqrt{2}C(Y_t)^{1/2} dW, \quad (2.22)$$

where $C(Y) \in \mathbb{R}^{DB \times DB}$ is a block diagonal matrix with B block entries $C(Y)$. The associated Fokker-Planck equation for the product density $\tilde{\pi}_t$ on \mathbb{R}^{DB} now reads

$$\partial_t \tilde{\pi}_t = \nabla \cdot (\tilde{\pi}_t C \nabla (\log \tilde{\pi}_t - \log \tilde{\pi}_*) + \tilde{\pi}_t \nabla C), \quad (2.23)$$

with the vector valued divergence $(\nabla C)_i = \sum_{j=1}^{NB} \partial_j C_{ij}$. Due to this divergence term, $\tilde{\pi}_*$ is not invariant under (2.23) for a finite number of particles. However, a straightforward calculation (see [28] for details) yields

$$\nabla C(Y) = \frac{D+1}{B} (Y - \bar{Y}). \quad (2.24)$$

Hence, the undesired term in the Fokker-Planck equation can be eliminated by replacing the drift in (2.22) accordingly with

$$C(Y_t) \nabla \ln \tilde{\pi}_*(Y_t) \longrightarrow C(Y_t) \nabla \ln \tilde{\pi}_*(Y_t) + \frac{D+1}{B} (Y_t - \bar{Y}_t).$$

ALDI adds the correction term to the drift term of the EKS. Hence, \mathcal{V} from (2.10) takes the form $\mathcal{V}(Y) = \sum_{i=1}^B \Phi(y^{(i)}) - \frac{D+1}{2} \log |C(Y)|$. Using the identity $C(Y) \nabla_{y^{(i)}} \log |C(Y)| = \frac{2}{B} (y^{(i)} - \bar{y})$, this leads to

$$dy_t^{(i)} = -C(Y_t) \nabla \Phi(y_t^{(i)}) dt + \frac{D+1}{B} (y_t^{(i)} - \bar{y}_t) dt + \sqrt{2} C(Y_t)^{1/2} dW_t^{(i)}. \quad (2.25)$$

Under strong growth bound conditions on Φ , $\nabla \Phi$ and $\text{Hess } \Phi$, and given $B > D+1$, ALDI is ergodic, i.e. $\tilde{\pi}_t$ converges to $\otimes_{i=1}^B \mu_*$ as $t \rightarrow \infty$ in total variation distance. The correction term not only makes $\tilde{\pi}_*$ invariant under the process, it also retains affine invariance for the finite particle system and its gradient-free version [16].

In practice, ALDI is used with a non-symmetric generalization of the square root

$$C(Y_t)^{1/2} = \frac{1}{\sqrt{B}} \left(y_t^{(1)} - \bar{y}_t, \dots, y_t^{(B)} - \bar{y}_t \right) \in \mathbb{R}^{D \times B} \quad (2.26)$$

such that $C(Y) = C(Y)^{1/2} (C(Y)^{1/2})^\top$. In contrast to the proper symmetric square root, (2.26) can be obtained without additional computational cost. Recently, the formalism of ALDI was extended to other forms of time-dependent scaling matrices such as *localized empirical covariances* [33] of the form

$$C^{(i)}(Y_t) = \sum_{j=1}^B \omega_t^{ij} (y_t^{(j)} - \bar{y}_t^{(i)}) (y_t^{(j)} - \bar{y}_t^{(i)})^\top, \quad (2.27)$$

for weights $\omega_t^{ij} \in \mathbb{R}$. This approach has been demonstrated to be effective for multimodal distributions [33], but it requires the actual computation of the square root $C(Y_t^{(i)})^{1/2}$ at each time step to build the localized empirical covariances. The respective cubic scaling in the number of samples might become computationally costly compared to the generalized non-symmetric square root (2.26) used in ALDI. Hence, to handle multimodalities in this work, we instead apply homotopy techniques to alter the potential, as in 2.1, retaining standard empirical covariances.

3 Ensemble enrichment

The goal of this section is to introduce various enrichment strategies that lead to a reduction of the interaction with the forward model within the Langevin dynamics. To that end, we will distinguish between three types of discrete measures. First, we define a general *discrete measure* $\hat{\mu}^{(B)}$ for $B \in \mathbb{N}$ through a set of points $\{y^{(i)}\}_{i=1}^B$ in \mathbb{R}^D by

$$\hat{\mu}^{(B)} = \frac{1}{B} \sum_{i=1}^B \delta_{y^{(i)}}, \quad (3.1)$$

where δ_y denotes the Dirac measure concentrated in $y \in \mathbb{R}^D$. If the points $\{y^{(i)}\}_{i=1}^B$ in (3.1) are i.i.d. samples from some distribution μ on \mathbb{R}^D , we will call $\mu^{(B)} = \hat{\mu}^{(B)}$ (a realization of) an *empirical measure* associated with μ . Third, if $\{y^{(i)}\}_{i=1}^B = \{y_t^{(i)}\}_{i=1}^B$ corresponds to the solution of a B -particle Langevin process at time t , we will call $\hat{\mu}_t = \hat{\mu}^{(B)}$ an *ensemble distribution*. In all of these cases, we suppress the dependence of the measure on the (random) set of points. In the last case, we additionally suppress the dependence on the batch size B , as it will be clear by context. Note that the $\{y_t^{(i)}\}$ are not necessarily i.i.d. in this case (in fact they are not independent for the solution of (2.25)).

We will sometimes use the particle ensemble $\{y_t^{(i)}\}_{i=1}^B$ and the corresponding ensemble distribution $\hat{\mu}_t$ interchangeably. The set of discrete measures on \mathbb{R}^D of the form (3.1) for arbitrary $B \in \mathbb{N}$ and an arbitrary set of points $\{y^{(i)}\}_{i=1}^B$ in \mathbb{R}^D will be denoted by $\mathcal{M}(\mathbb{R}^D)$. The set of all families $(\hat{\mu}_t)_{t \geq 0}$ with $\hat{\mu}_t \in \mathcal{M}(\mathbb{R}^D)$ we denote by $\mathcal{M}_{t \geq 0}(\mathbb{R}^D)$. Subsequently, the notion of ensemble enrichment on such families of discrete measures in combination with an enrichment time t_a and a number $a \in \mathbb{N}$ of requested additional particles is defined.

Definition 3.1. *Let $(\Omega, \sigma, \mathbb{P})$ be an abstract probability space. Then, a map $\mathcal{E}: \mathcal{M}_{t \geq 0}(\mathbb{R}^D) \times \Omega \times \mathbb{R}_+ \times \mathbb{N} \rightarrow \mathcal{M}(\mathbb{R}^D)$ is called an ensemble enrichment scheme if the following condition (C) is satisfied.*

(C) *If $\hat{\mu}_t$ is determined by an ensemble of b particles at $t = t_a$, then for $a \in \mathbb{N}$, $\mathcal{E}[(\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a]$ is a discrete measure determined by $b + a$ particles.*

The dependence on $\omega \in \Omega$ models random effects, e.g. introduced through random perturbations or random selection of particles. In practice, such dependence is given as

$$\mathcal{E}[(\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a] = \mathcal{E}[(\hat{\mu}_t)_{t \geq 0}, \zeta(\omega), t_a, a] \quad (3.2)$$

with some random variable ζ independent of the $(\hat{\mu}_t)_{t \geq 0}$. If (C) is satisfied, we call $\hat{\mu}_{t_a}$ and $\{y_{t_a}^{(i)}\}$ the *original measure* and *original batch*, respectively. Moreover, $\mathcal{E}[(\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a]$ is the *enriched measure* and its corresponding sample batch the *enriched batch*. When the associated family of measures is clear from the context, we call this the application of an enrichment scheme to a batch $\{y_{t_a}^{(i)}\}$. This can be considered the more natural viewpoint and appeals to our intuitive understanding of enrichment as adding samples to an existing batch. The reason we formally define an enrichment scheme as a map acting on a family of measures is two-fold. First, we want an enrichment to be able to rely on past and future values of the process $\{y_t^{(i)}\}$. Second, the operation on measures instead of particle ensembles allows to view the solution provided by our LIDL method as a time-continuous process. To make this point clear, note that we cannot define the result of the process described in Section 2 as a time-dependent ensemble $\{y_t^{(i)}\}_{i=1}^B$ since the batch size B changes over time. This however can be achieved in terms of a family of ensemble distributions $(\hat{\mu}_t)_{t \geq 0}$. This formulation is agnostic with regard to the current batch size, which is made precise in Definition 4.1.

Having formally defined what an ensemble enrichment scheme is, the question what constitutes a *good* enrichment strategy arises immediately. Since in the end we are concerned with convergence of the sampling scheme, the distance of the enriched measure to the posterior has to be controlled. In the remainder of this paper we use the *Kantorovich–Rubinstein metric* (or *p-Wasserstein distance*) of measures for $p = 2$. Let $\mathcal{D}_p(\mathcal{V})$ be the space of measures on a metric space \mathcal{V} with finite second moments. The 2-Wasserstein distance of measures $\mu, \nu \in \mathcal{D}_2(\mathbb{R}^D)$ is then defined by

$$\mathcal{W}_2(\mu, \nu) = \left[\min_{\pi \in \mathcal{D}(\mathbb{R}^D \times \mathbb{R}^D)} \{ \langle \pi, c_2 \rangle : \pi_1 = \mu, \pi_2 = \nu \} \right]^{1/2}, \quad (3.3)$$

where $\pi_1 = \int_{\mathbb{R}^D} d\pi(\cdot, y)$ and $\pi_2 = \int_{\mathbb{R}^D} d\pi(x, \cdot)$ are the marginals of the transport plan and $c_2(x, y) = \frac{1}{2}|x - y|^2$. The 2-Wasserstein metric space is denoted by $\mathcal{W}_2(\mathbb{R}^D) := (\mathcal{D}_2(\mathbb{R}^D), \mathcal{W}_2)$. The following enrichment strategies aim to approximately follow the flow towards the posterior measure, as initiated by the original measure.

3.1 Enrichment schemes

Let $a, b \in \mathbb{N}$ with $a \leq b$ and let the measure $\hat{\mu} \in \mathcal{M}(\mathbb{R}^D)$ be determined by $b \in \mathbb{N}$ particles $y^{(i)}$. Furthermore, let the random variable ζ_a encode the uniformly at random selection of a particles out of $\{y^{(i)}\}_{i=1}^B$ with realizations denoted by $\zeta_a(\omega)[\hat{\mu}]$ being a discrete measure determined by a particles. For instance, if the first $i = 1, \dots, a$ particles have been selected by ζ_a , the resulting measure is $\frac{1}{a} \sum_{i=1}^a \delta_{y^{(i)}}$.

3.1.1 Slicing

The idea of (time)-*slicing* relies on the idea to enrich the batch of size b at time t_a by adding particles from other batches associated to a selection of finitely many ensemble distributions $\hat{\mu}_t$ from the family $(\hat{\mu}_t)_{t \geq 0}$. A canonical slicing is defined at time points around t_a . For this let, $0 < \Delta t < t_a$ and assume the batch size of the family of measures to be equal to $b \geq a$ in the neighborhood $[t_a - \Delta t, t_a + \Delta t]$.

We define the *forward slicing* enrichment $\mathcal{E}_{+\Delta t}$ via

$$\mathcal{E}_{+\Delta t}((\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a) = \frac{1}{b+a} (b\hat{\mu}_{t_a} + a\zeta_a(\omega) [\hat{\mu}_{t_a+\Delta t}]). \quad (3.4)$$

In the idealized setting where $b = a$ and $\mathcal{W}_2(\hat{\mu}_t, \mu_*)$ is monotonically decreasing in t , we get that the 2-Wasserstein distance of (3.4) to the posterior is bounded by $\mathcal{W}_2(\hat{\mu}_{t_a}, \mu_*)$ and hence, (3.4) should at least preserve the distance to the posterior at time t_a . This property is obviously desirable but it comes at the cost of carrying out the forward model evaluations necessary to compute the measure $\hat{\mu}_{t_a+\Delta t}$ from $\hat{\mu}_{t_a}$. Note that in practice the measure family is only available for $t \leq t_a$ without additional computational burden.

To avoid the additional forward model calls, one may instead use a *backward slicing* scheme $\mathcal{E}_{-\Delta t}$ defined by

$$\mathcal{E}_{-\Delta t}((\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a) = \frac{1}{b+a} (b\hat{\mu}_{t_a} + a\zeta_a(\omega) [\hat{\mu}_{t_a-\Delta t}]). \quad (3.5)$$

Here, no extra forward model evaluations are needed since $\hat{\mu}_{t_a-\Delta t}$ is readily available at time t_a . However, compared to the forward slicing scheme, we only get a 2-Wasserstein distance bound proportional to $\mathcal{W}_2(\hat{\mu}_{t_a-\Delta t}, \mu_*)$ when assuming monotonic decrease of the Wasserstein distance in t .

3.1.2 Diffusion propagation

Recall that computing the measure $\hat{\mu}_{t_a+\Delta t}$ in the forward slicing scheme (3.4) involves propagating $\hat{\mu}_{t_a}$ through the underlying Langevin process (2.10) from time t_a to $t_a + \Delta t$. For small Δt , we can define an approximation to $\hat{\mu}_{t_a+\Delta t}$ by propagating the particle ensemble $\{y_{t_a}^{(i)}\}$ associated with $\hat{\mu}_{t_a}$ using only the diffusion part of (2.10),

$$dy_t^{(i)} = \Gamma(Y_t) dW_t^{(i)}. \quad (3.6)$$

We encode the effect of this *diffusion propagation* on $\hat{\mu}_{t_a}$ by a random variable $\zeta_W^{\Delta t}$ with realizations $\zeta_W^{\Delta t}(\omega)[\hat{\mu}_{t_a}] \approx \hat{\mu}_{t_a+\Delta t}$. By $\zeta_a \circ \zeta_W^{\Delta t}$ with realizations $(\zeta_a \circ \zeta_W^{\Delta t})(\omega)[\hat{\mu}_{t_a}] = (\zeta_a(\omega) \circ \zeta_W^{\Delta t}(\omega))[\hat{\mu}_{t_a}]$, we denote the concatenation with the previously defined random selection. Hereby, we choose ζ_a and

$\zeta_W^{\Delta t}$ to be independent.

The diffusion propagation scheme is then defined by

$$\mathcal{E}_{+\Delta t}^{\text{diff}}((\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a) = \frac{1}{b+a} (b\hat{\mu}_{t_a} + a(\zeta_a \circ \zeta_W^{\Delta t})(\omega) [\hat{\mu}_{t_a}]). \quad (3.7)$$

This scheme entirely avoids any forward model calls, while also not relying on the history of the process as in the backward slicing scheme. This approximation is also motivated by the numerical time discretization of (2.10). To that end, consider an Euler-Maruyama discretization $\tilde{y}_t^{(i)}$ of the process $y_t^{(i)}$ with time step Δt defined through

$$\tilde{y}_{t+\Delta t}^{(i)} = \tilde{y}_t^{(i)} - \Delta t A(\tilde{Y}_t) \nabla_{\tilde{y}_t^{(i)}} \mathcal{V}(\tilde{Y}_t) + \sqrt{\Delta t} \Gamma(\tilde{Y}_t) \xi_t^{(i)}. \quad (3.8)$$

Here, $\xi_t \sim \mathcal{N}(0, I_d)$ are i.i.d. increments and $\tilde{Y}_t = \text{vec}(\tilde{y}_t^{(1)}, \dots, \tilde{y}_t^{(b)})$. Vice-versa a discretization of (3.6) leads to the relation

$$\tilde{y}_{t+\Delta t}^{(i)} = \tilde{y}_t^{(i)} + \sqrt{\Delta t} \Gamma(\tilde{Y}_t) \xi_t^{(i)}, \quad (3.9)$$

where we used the same notation to underline similarities to (3.8). While in (3.8) the drift part scales with Δt , the diffusion part scales with $\sqrt{\Delta t}$, and hence an approximation via (3.9) is justified provided $\Delta t \ll 1$ and $A(Y_t) \nabla_{y_t^{(i)}} \mathcal{V}(Y_t)$ being bounded. Consequently, the diffusion propagation approximates a forward slice, without the computational burden of evaluating \mathcal{V} . Note that, in this time-discrete setting, the randomness of the propagation, formerly represented by $\zeta_W^{\Delta t}$ is now encoded in the increments $\xi_t^{(i)}$ for $t = t_a$.

In this time-discrete setting, the randomness of the propagation, which was formerly encoded in $\zeta_W^{\Delta t}$ is now encoded in the increments $\xi_k^{(i)}$.

3.1.3 Random kicks

The discrete viewpoint taken in (3.8) and (3.9) allows for the interpretation of the diffusion step as a special case of a *random kicks* method

$$\mathcal{E}_\eta((\hat{\mu}_t)_{t \geq 0}, \omega, t_a, a) = \frac{1}{b+a} (b\hat{\mu}_{t_a} + a(\zeta_a \circ \zeta_\eta)(\omega) [\hat{\mu}_{t_a}]), \quad (3.10)$$

where the random variable ζ_η encodes adding perturbation noise $\eta \sim \bigotimes_{i=1}^b \mathcal{D}(\mathbb{R}^D)$ to the particles $\{y_{t_a}^{(i)}\}$ determining $\hat{\mu}_{t_a}$. Note that (3.9) can be seen as a highly informed choice of noise, utilizing the underlying process. In the case of the ALDI method (2.25), it preserves the covariance structure of the ensemble.

The question arises why one would consider other less informed random kicks at all. We note that in the low batch size regime one important strength of diffusion propagation, namely preserving the covariance rank as in the case of ALDI, turns out to be detrimental. To see this, recall (2.25) with covariance matrix $C(Y_t)$ given by (2.15) and generalized non-symmetric square root $C(Y_t)^{1/2}$ given by (2.26). Consider now a batch size $b < D$ such that the covariance matrix has at most rank b and is hence not positive definite. While any proper perturbation noise η independent of the particle ensemble (with sufficiently large enrichment size a) leads to a positive definite covariance matrix with probability 1, (3.8) and (3.9) produce additional samples in the range of the covariance matrix, not increasing its rank at all. This means that the covariance matrix is not positive definite at any point in

the future (up to numerical instability) and particles remain in the linear subspace associated to the range of the covariance. Hence, preserving the rank of the covariance via methods like slicing and diffusion propagation only makes sense for sufficiently large batch sizes $b \geq D$, leading to full rank with probability 1. As a simple and purely heuristic choice of a single random kick for the low batch size regime, we propose using scaled Gaussian noise

$$\tilde{y}_{t+\Delta t}^{(i)} = \tilde{y}_t^{(i)} + \sqrt{\Delta t} \xi_k^{(i)} \quad (3.11)$$

for i.i.d. increments $\xi_k^{(i)} \sim \mathcal{N}(0, I_d)$ and a step size Δt depending on $C(\tilde{Y}_t)$. We leave the discussion of other informed choices of noise as a topic for future work.

Finally, we note that all of these methods can be extended to the case $a > b$, e.g., by adding together multiple forward/backward slices at different times $t_a \pm \Delta t_i$. The corresponding modifications of (3.4), (3.5), (3.7) and (3.10) are straightforward and left for the reader.

3.1.4 Generalized transport approach

A conceptually different approach relies on the intermediate learning of random variables, that approximately follow the distribution of $y_t \sim \pi_t$ given by (2.20) and (2.21) at time point $t = t_a$. This allows for a fast generation of samples due to a functional representation, which takes the form

$$\mathcal{M}_t(x_t) \stackrel{d}{\approx} y_t \quad (3.12)$$

with a suitable model class \mathcal{M}_t and some auxiliary random variable x_t . Here, the approximation quality should be controlled in the same metric used to analyze the convergence of the particle propagator. Once such representation is found for a fixed $t = t_a$, samples can be drawn by sampling from x_t and propagating through \mathcal{M}_t .

A setup of particular importance arises for $t = T \gg 1$. If μ_T is close to μ_* , then the left-hand side of (3.12) provides an approximate functional access to the posterior distribution. This concept is similar to the case of transport maps, where x_T is distributed with respect to the prior distribution and \mathcal{M}_T is a diffeomorphism [42, 34, 27, 2]. However, here we relax the assumption of bijectivity or continuity properties of \mathcal{M} in order to enable accurate approximations of more involved distributions such as multimodal ones even when x_T is a unimodal distribution. Moreover, x_T is not necessarily distributed with respect to the prior distribution but rather defined on some latent space.

Another approach based on (3.12) is denoted as *sequential learning of generalized transport* and described in the following. Let $0 = t_0 < t_1 < \dots < t_L < \infty$ and

$$x_{t_\ell} = y_{t_{\ell-1}}, \quad \ell = 1, \dots, L.$$

This design leads to the form

$$y_{t_L} \stackrel{d}{\approx} M_{t_L} \circ \dots \circ M_{t_1}(y_0)$$

for a suitable model class \mathcal{M}_{t_ℓ} . Hence, the compositional structure imitates the associated particle propagator flow. This concept appears to be closely related to *stochastic normalizing flows* [44, 21]. The discussion of such a compositional approach is subject to future research.

Another special case of (3.12) is of the form

$$y_t \stackrel{d}{\approx} \mathcal{M}_t(x). \quad (3.13)$$

In particular, the auxiliary random variable remains the same over the time horizon and only the model class is updated. A particular design for (3.13) is realized by *generative adversarial networks* (GANs), where typically x corresponds to a standard normal multivariate Gaussian distribution. As an alternative, the use of a *compressed Wasserstein polynomial chaos expansion* (WPCE) was proposed in [20]. Here, x determines the family of orthonormal stochastic polynomials used for the approximation.

In order to realize the representation (3.13), one can utilize techniques from computational optimal transport in the framework of unsupervised learning. We define the debiased Sinkhorn divergence [12] based on (3.3) for $\epsilon > 0$ by

$$\mathcal{S}_\epsilon(\mu, \nu) = \mathcal{W}_{c,\epsilon}(\mu, \nu) - \frac{1}{2} (\mathcal{W}_{c,\epsilon}(\mu, \mu) + \mathcal{W}_{c,\epsilon}(\nu, \nu)), \quad (3.14)$$

for measures $\mu, \nu \in \mathcal{D}_2(\mathbb{R}^D)$, where

$$\begin{aligned} \mathcal{W}_{c,\epsilon}(\mu, \nu) &= \min_{\pi \in \mathcal{D}(\mathbb{R}^D \times \mathbb{R}^D)} \langle \pi, c \rangle + \epsilon \text{KL}(\pi, \mu \otimes \nu), \\ &\text{subject to } \pi \geq 0, \quad \pi_1 = \mu, \quad \pi_2 = \nu. \end{aligned} \quad (3.15)$$

Here, KL denotes the Kullback-Leibler divergence. Then, samples y_t define an discrete measure $\hat{\mu}_t$ and samples from x propagated through $\mathcal{M}_t = \mathcal{M}[\theta_t]$ define a discrete measure $\hat{\nu}[\theta_t]$ depending on the parameter θ_t . The desired coefficient θ_t is defined as the minimizer of

$$\min_{\theta} \mathcal{S}_\epsilon(\hat{\mu}_t, \hat{\nu}[\theta]). \quad (3.16)$$

4 Theoretical Foundations

We henceforth use ALDI (2.25) as our particle propagator of choice. Together with a chosen enrichment scheme from Section 3, this allows to define our method in a rigorous way.

Definition 4.1 (LIDL). *Set $L \in \mathbb{N}_0$ and define the index set $I_L = \{0, 1, 2, \dots, L\}$. Moreover, let*

- (i) $(t_\ell)_{\ell \in I_L}$, $t_\ell \geq t_0 = 0$ for all ℓ be a strictly monotonically increasing sequence of time points with $t_{L+1} = \infty$,
- (ii) $(b_\ell)_{\ell \in I_L}$, $b_\ell \in \mathbb{N}$ be a sequence of batch sizes,
- (iii) $\{y_0^i\}_{i=1}^{b_0}$ be an initial sample batch,
- (iv) \mathcal{E} be an enrichment strategy.

Then, an instance of a LIDL run with parameters $((t_\ell)_\ell, (b_\ell)_\ell, \{y_0^i\}, \mathcal{E})$ applied to the inverse problem defined by (1.4) produces a family of ensemble distributions $(\hat{\mu}_t)_{t \geq 0}$ on \mathbb{R}^D by repeating the following two steps for $\ell = 0, \dots, L$:

(Step 1) Solve (2.25) with initial conditions $\{y_\ell^i\}$ for $0 \leq t \leq t_{\ell+1} - t_\ell$. Denote the resulting solution by $\{y_{\ell,t}^i\}$ and the corresponding ensemble distribution by $\hat{\mu}_{\{y_{\ell,t}^i\}}$. Set

$$\hat{\mu}_t = \hat{\mu}_{\{y_{\ell,t-t_\ell}^i\}} \quad \text{for } t \in [t_\ell, t_{\ell+1}). \quad (4.1)$$

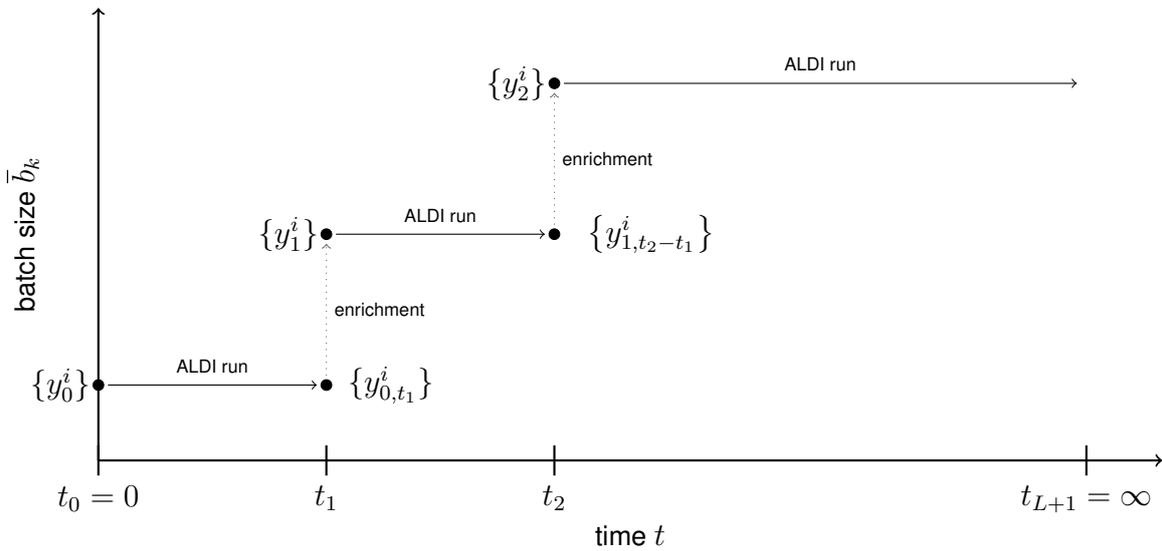


Figure 4.1: Schematic of a run of LIDL with $L = 2$ enrichment steps. An initial batch of samples $\{y_0^i\}$ is propagated through the ALDI dynamics (2.25) for $t \in [0, t_1]$. The resulting samples $\{y_{0,t_1}^i\}$ are enriched via some enrichment strategy \mathcal{E} to receive a larger sample batch $\{y_1^i\}$, which is again propagated and then enriched at $t = t_2$. Finally, a full run of ALDI is performed on the final batch $\{y_2^i\}$.

(Step 2) If $\ell \leq L - 1$, get the next initial sample batch $\{y_{\ell+1}^i\}$ by applying the ensemble enrichment \mathcal{E} with $b_{\ell+1}$ new samples to $\{y_{\ell,t_{\ell+1}}^i\}$.

The random process $(\hat{\mu}_t)_{t \geq 0}$ is called the solution of LIDL. A visualization of steps 1 and 2 are depicted in Figure 4.1.

Remark 4.1. Some comments on the definition are in order.

- 1 Setting $t_{L+1} = \infty$ is a formality, guaranteeing that LIDL returns a measure $\hat{\mu}_t$ for all $t \geq 0$. The last stage is equivalent to solving (2.25) with initial conditions $\{y_L^i\}$.
- 2 When $L = 0$, the LIDL solution becomes identical to the solution of (2.25) for all t . Hence, ALDI can be seen as a special case of LIDL.
- 3 We could define the method more generally by admitting a sample propagator like EKS, ALDI, etc. as an additional parameter. For the sake of simplicity, we only work with ALDI and neglect this dependence in the rest of the paper.

4.1 Convergence analysis in the linear case

Throughout this section, we consider the Bayesian inverse problem (1.1) in the special case of a Gaussian prior and a linear forward map. Hence, let $\mathcal{G}(\cdot) = A \cdot$ for some $A \in \mathbb{R}^{K \times D}$ and

$$\pi_{\text{prior}}(y) \propto \exp\left(-\frac{1}{2}(y - y_0)^\top \Gamma_0^{-1}(y - y_0)\right) = \exp\left(-\frac{1}{2}|y - y_0|_{\Gamma_0}^2\right), \quad (4.2)$$

with prior mean $y_0 \in \mathbb{R}^D$ and prior covariance matrix $\Gamma_0 \in \mathbb{R}^{D \times D}$. In this special case, the posterior is again Gaussian with density

$$\pi_*(y) \propto \exp\left(-\frac{1}{2}|y - y^*|_{P^{-1}}^2\right), \quad (4.3)$$

with the posterior precision matrix

$$P = C(\mu_*)^{-1} = A^\top \Gamma^{-1} A + \Gamma_0^{-1} \quad (4.4)$$

and the posterior mean

$$y^* = \mathbb{E}_{\mu_*} = P^{-1}(A^\top \Gamma^{-1} \delta + \Gamma_0^{-1} y_0). \quad (4.5)$$

In the following, we denote by $\lambda_{\min}(M)$ the smallest eigenvalue of a symmetric positive definite matrix $M \in \mathbb{R}^{D,D}$.

As a first step in the convergence analysis, we aim for a consistency result in the expected 2-Wasserstein distance \mathcal{W}_2 between the measures $\hat{\mu}_t$ generated by LIDL and the posterior measure μ_* . By consistency we mean that for any $\delta > 0$ there is a configuration $((t_\ell)_\ell, (b_\ell)_\ell, \{y_0^i\}, \mathcal{E})$ of LIDL and a time $T_\delta > 0$ such that

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_{T_\delta}, \mu_*)] \leq \delta. \quad (4.6)$$

We have established in Remark 4.1 that ALDI can be seen as a special case of LIDL with $L = 0$ enrichment steps. Hence, consistency as defined above can be achieved in a trivial manner if ALDI can be shown to be B-T-consistent, meaning that for $\delta > 0$ there are T_δ, B_{T_δ} such that the ensemble distribution $\hat{\mu}_t$ generated by ALDI satisfies (4.6). The following theorem provides sufficient conditions for consistency in this sense.

Theorem 4.1 (B-T-consistency of ALDI). *Suppose $\mathcal{G}(\cdot) = A \cdot$ is linear and π_{prior} is given by (4.2). Furthermore, let $\pi_0 \in \mathcal{C}^2$ be a density with bounded higher moments and for $B \in \mathbb{N}$ let $\{y_t^i\}_{i=1}^B$ be the solution of (2.25) with initial condition $\{y_0^i\}_{i=1}^B$ drawn i.i.d from π_0 . Let $\delta > 0$ and $T > 0$ be such that $\mathcal{W}_2(\pi(T), \pi_*) < \delta$ where $\pi(T)$ is the solution of (2.21) with initial condition π_0 . Furthermore assume that*

$$\frac{\lambda_{\min}(P)\lambda_0(t)}{2} \geq 1, \quad \text{for all } 0 \leq t \leq T, \quad (4.7)$$

where $\lambda_0(t) = (\lambda_{\min}(C(Y_t))^{1/2} + \lambda_{\min}(C(\pi(t)))^{1/2})^2$. Then, there exists $T_\delta \leq T$ and $B_{T_\delta} > 0$ such that the solution $\{y_t^i\}_{i=1}^B$ with $B = B_{T_\delta}$ and its corresponding ensemble distribution $\hat{\mu}_{T_\delta}^{B_{T_\delta}}$ satisfy

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_{T_\delta}^{B_{T_\delta}}, \mu_*)] \leq \delta. \quad (4.8)$$

Proof. The proof can be found in Appendix A. □

Remark 4.2. *Since $\pi(t)$ converges to π_* exponentially fast in 2-Wasserstein distance (see Theorem A.1 in the appendix), we can always find a $T > 0$ such that $\mathcal{W}_2(\pi(T), \pi_*) < \delta$. The second assumption (4.7) on T is a technical one and comes from an application of the Ando-Hemmen inequality (see the proof of Lemma 5.4 in [9]).*

Remark 4.3. *The proof uses a triangle argument. For $0 \leq t \leq T$ this leads to a bound of the form*

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_t^B, \mu_*)] \leq ce^{-t}\mathcal{W}_2(\pi_0, \pi_*) + c(t, D, \epsilon) \begin{cases} B^{-1/2+\epsilon}, & D \leq 4, \\ B^{-2/D}, & D > 4, \end{cases} \quad (4.9)$$

where $0 < \epsilon < 1/2$, $c > 0$ is a constant depending only on the initial density π_0 and the posterior density π_* and $c(t, D, \epsilon) > 0$ is a constant depending on t, D, ϵ . The first term comes from the Fokker-Planck solution and decays exponentially fast for $t \rightarrow T$. The constant $c(t, D, \epsilon)$ for the remaining term however grows exponentially with T and we hence obtain no monotonicity for $t \rightarrow T$. Nevertheless, the bound (4.9) yields monotonicity in the following sense: if (4.8) is satisfied for some T_δ, B_{T_δ} , then for any $t \in [T_\delta, T]$ we find a $B_t \geq B_{T_\delta}$ such that (4.8) is also satisfied for t, B_t . Note that this already provides a theoretical motivation for an ensemble enrichment: to achieve the same error bound δ for times $t > T_\delta$, we may have to increase the batch size.

Corollary 4.1 (B-T-consistency of LIDL). *Under the same conditions as in Theorem 4.1, for every $\delta > 0$ there exists a configuration $((t_\ell)_\ell, (b_\ell)_\ell, \{y_0^i\}, \mathcal{E})$ of LIDL and a time $T_\delta > 0$ such that the solution $\hat{\mu}_t$ defined by (4.1) satisfies*

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_{T_\delta}, \mu_*)] \leq \delta. \quad (4.10)$$

Proof. Choosing $L = 0$ and $T_\delta, b_0 = B_{T_\delta}$ as in Theorem 4.1 yields one such configuration. \square

A full ALDI run is not the only configuration with B-T-consistency. In particular, we require the last sample batch $\{y_L^i\}$ to satisfy the conditions of Theorem 4.1. One of the conditions is that $\{y_L^i\}$ is drawn i.i.d. from some \mathcal{C}^2 -density with bounded higher moments. This condition comes from the smoothness required of a strong solution of the Fokker-Planck equation. This can be formulated as an assumption on the enrichment scheme.

Definition 4.2 (Consistent Enrichment). *We call an enrichment strategy \mathcal{E} consistent if it creates enriched ensembles drawn i.i.d. from some \mathcal{C}^2 density with finite higher moments.*

As the following remark shows, this is a purely theoretical condition. In practice, any enrichment scheme can be seen as an arbitrarily close approximation of a consistent one.

Remark 4.4. *The subset of measures in $\mathcal{D}_2(\mathbb{R}^D)$ with $\mathcal{C}^2(\mathbb{R}^D)$ -Lebesgue density is dense in $\mathcal{W}_2(\mathbb{R}^D)$, see Lemma D.1 for $p = 2$ and $k = 2$. Now let \mathcal{E} be an arbitrary ensemble enrichment scheme taking the original batch $\{y_1^i\}_{i=1}^{b_1}$ and generating an enriched batch $\{y_2^i\}_{i=1}^{b_2}$ for $b_2 > b_1$ with ensemble distribution $\hat{\mu}$. Then, for any $\epsilon > 0$ there exists a regular measure $\mu_\epsilon \in \mathcal{W}_2(\mathbb{R}^D)$ with Lebesgue density in $\mathcal{C}^2(\mathbb{R}^D)$ such that $\mathcal{W}_2(\hat{\mu}, \mu_\epsilon) < \epsilon$. Hence, the enriched ensemble approximately can be seen as being drawn i.i.d. from a distribution with \mathcal{C}^2 density.*

With this notion of consistency in place, the following result immediately follows.

Corollary 4.2 (B-T-consistency of LIDL). *Assume the conditions of Theorem 4.1 to be satisfied. Let \mathcal{E} be a consistent ensemble enrichment scheme in the sense of Definition 4.2. Then, for every configuration $((t_\ell), (b_\ell)_\ell, \{y_0^i\}, \mathcal{E})$ of LIDL with b_0, \dots, b_L sufficiently large and $\delta > 0$, there exists a time $T_\delta > 0$ such that the solution $\hat{\mu}_t$ defined by (4.1) satisfies*

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_{T_\delta}, \mu_*)] \leq \delta. \quad (4.11)$$

Proof. Since \mathcal{E} is consistent, the last batch $\{y_L^i\}$ is drawn i.i.d. from a \mathcal{C}^2 -density with bounded higher moments. Moreover, the conditions of Theorem 4.1 apply to the starting ensemble $\{y_L^i\}$ by assumption. Hence, there exists a T_δ and B_{T_δ} such that (4.11) holds true, provided $\bar{b}_L = b_0 + \dots + b_L \geq B_\delta$. \square

4.2 Convergence of homotopy approach

Motivated by the particular homotopy approach from Section 2, define *switch points* $0 = s_0 < s_1 < \dots < s_K = 1$ for some $K \in \mathbb{N}$. Let $L_{\log}^1(\mathbb{R}^D) := \{\phi: \mathbb{R}^D \rightarrow \mathbb{R} : \exp(-\phi) \in L^1(\mathbb{R}^D)\}$. In what follows, the constant Z_s^{-1} is a generic L^1 normalization constant enumerated by $s \in [0, 1]$ and associated with a measure μ_s . First, we define a class of feasible homotopy functions, which we call \mathcal{W}_2 -stable homotopies.

Definition 4.3 (\mathcal{W}_2 -stable homotopy). *Let $\Phi_i: \mathbb{R}^D \rightarrow \mathbb{R}$ such that $\mu_i := Z_i^{-1} \exp(-\Phi_i) \in \mathcal{W}_2(\mathbb{R}^D)$ for $i = 0, 1$. A mapping $\mathcal{H}: [0, 1] \rightarrow L_{\log}^1(\mathbb{R}^D)$ is denoted a \mathcal{W}_2 -stable homotopy between Φ_0 and Φ_1 if*

$$(i) \quad \mathcal{H}(0) = \Phi_0, \quad \mathcal{H}(1) = \Phi_1, \quad (\text{interpolation})$$

$$(ii) \quad \mu_s := Z_s^{-1} \exp(-\mathcal{H}(s)) \in \mathcal{W}_2(\mathbb{R}^D), \quad s \in (0, 1), \quad (\text{consistency})$$

$$(iii) \quad \mathcal{W}_2(\mu_{s_1}, \mu_{s_2}) \leq \phi(|s_1 - s_2|), \quad s_1, s_2 \in [0, 1], \quad (\text{stability})$$

for continuous $\phi: [0, 1] \rightarrow \mathbb{R}$ with $\phi(s) \rightarrow 0$ as $s \searrow 0$.

With this preparation, let \mathcal{H} be a \mathcal{W}_2 -stable homotopy between some auxillary potential $\mathcal{H}(0) = \Psi$ and the posterior potential $\mathcal{H}(1) = \Phi$. In order to still be able to work with Itô diffusion processes as in (2.10), we aim to define an inhomogeneous drift term $f_{\mathcal{H}} = f_{\mathcal{H}}(t, \cdot)$ that is piecewise constant in t . For this, we define a time horizon partition $\mathcal{T}: \{0, \dots, K\} \rightarrow \mathbb{R}_+$ with

$$\mathcal{T}(0) = 0, \quad \mathcal{T}(k) < \mathcal{T}(k+1), \quad k = 0, \dots, K-1. \quad (4.12)$$

Then, for $Y = \text{vec}(y^{(i)})_i \in \mathbb{R}^{DB}$ with $B \in \mathbb{N}$ and $\mathcal{T}(K+1) = \infty$, let

$$f_{\mathcal{H}}(t, Y) := f(\mathcal{H}(s_k), Y), \quad t \in [\mathcal{T}(k), \mathcal{T}(k+1)), \quad k = 0, \dots, K. \quad (4.13)$$

Then, for $k = 0, \dots, K$ consider the time-partitioned dynamics

$$dy_t^{(i)} = f(\mathcal{H}(s_k), Y_t)dt + \Gamma(Y_t)dW_t^{(i)}, \quad t \in [\mathcal{T}(k), \mathcal{T}(k+1)), \quad (4.14)$$

yielding the full process

$$dy_t^{(i)} = b_{\mathcal{H}}(t, Y_t)dt + \Gamma(Y_t)dW_t^{(i)}, \quad t \geq 0. \quad (4.15)$$

For the intended application, the drift term $f_{\mathcal{H}}$ corresponds to the drift term in (2.25), replacing Φ with $\mathcal{H}(s_k)$, i.e.

$$f(\mathcal{H}(s_k), Y_t) = -C(Y_t)\nabla_{y_t^{(i)}}\mathcal{H}(s_k)\left(y_t^{(i)}\right) + \frac{D+1}{B}\left(y_t^{(i)} - \bar{y}_t\right).$$

The partitioning of the dynamics up to time $\mathcal{T}(K)$ via (4.14) can be seen as a preconditioner with auxiliary potentials prior to starting the ALDI run with the posterior potential Φ at $t = \mathcal{T}(K)$.

Assumption 4.1 (Local convergence). *Let $\delta: [0, 1] \rightarrow [\underline{\delta}, \infty]$ for $\underline{\delta} > 0$ be a parameter dependent convergence radius. For $s \in [0, 1]$, let $\hat{\mu}_{0,s}$ be an arbitrary random probability measure with*

$$\mathbb{E}[\mathcal{W}_2(\mu_s, \hat{\mu}_{0,s})] < \delta(s). \quad (4.16)$$

Then, the propagator SDE (4.14) implies expected convergence to μ_s in the following sense:

For $0 < \epsilon_s < \delta(s)$ there exists $T_s := T_s(\epsilon_s) > 0$ and $B_s = B_s(\epsilon_s) > 0$ such that for all $B \geq B_s$,

$$\mathbb{E} [\mathcal{W}_2(\hat{\mu}_{T_s, s}^B, \mu_s)] \leq \epsilon_s. \quad (4.17)$$

Here, $\hat{\mu}_{t, s}^B$ denotes the random ensemble distribution associated to $(y_t^{(i)})_{i=1}^B$ as the solution of (4.14) with $y_0^{(i)} \sim \hat{\mu}_{0, s}$.

Note that Assumption 4.1 refers to local convergence in terms of the starting distribution only.

Theorem 4.2 (Convergence of homotopy approach). *Let \mathcal{H} be a \mathcal{W}_2 stable homotopy between $\mathcal{H}(0) = \Psi \in L_{\log}^1(\mathbb{R}^D)$ and the posterior potential $\mathcal{H}(1) = \Phi$ and let Assumption 4.1 be satisfied. Let $\epsilon > 0$. Then there exists $K \in \mathbb{N}$, switch points s_0, \dots, s_K and a corresponding horizon partition \mathcal{T} , $T_\epsilon > 0$ and $B_\epsilon > 0$ such that*

$$\mathbb{E} [\mathcal{W}_2(\hat{\mu}_{T_\epsilon}^{B_\epsilon}, \mu_*)] < \epsilon,$$

where $\hat{\mu}_t^B$ denotes the ensemble distribution of $(y_t^{(i)})_{i=1}^B$.

Proof. For $B \in \mathbb{N}$ let $\hat{\mu}_{t, s}^B$ denote the B -particle ensemble distribution at time t using the propagation (2.10) with

$$-A(Y_t) \nabla_{y_t^{(i)}} \mathcal{V}(Y_t) = -C(Y_t) \nabla_{y_t^{(i)}} \mathcal{H}(s) \left(y_t^{(i)} \right) + \frac{D+1}{B} \left(y_t^{(i)} - \bar{y}_t \right)$$

with initial condition $\hat{\mu}_{0, s}^B$. Since \mathcal{H} is \mathcal{W}_2 -stable, it follows from the stability assumption in Definition 4.3 that for all $s \in [0, 1]$

$$\mathcal{W}_2(\mu_0, \mu_s) \leq \phi(s).$$

Since ϕ is continuous and $\delta(s)$ is uniformly bounded from below by $\underline{\delta}$, for arbitrary but fixed $0 < \epsilon_0 < \underline{\delta}$ there is a maximal $s_1 \in (0, 1]$ such that

$$\phi(s_1) + \epsilon_0 < \delta(s_1).$$

Now let $\hat{\mu}_{0, 0}$ be an initial ensemble distribution with $\mathbb{E}[\mathcal{W}_2(\mu_0, \hat{\mu}_{0, 0})] < \delta(0)$. Then by the triangle inequality and using Assumption 4.1 we find $T_0 = T_0(\epsilon_0)$, $B_0 = B_0(\epsilon_0)$ such that

$$\mathbb{E} [\mathcal{W}_2(\mu_{s_1}, \hat{\mu}_{T_0, 0}^{B_0})] \leq \mathcal{W}_2(\mu_{s_1}, \mu_{s_0}) + \mathbb{E} [\mathcal{W}_2(\mu_{s_0}, \hat{\mu}_{T_0, s_0}^{B_0})] \leq \phi(s_1) + \epsilon_0 \leq \delta(s_1)$$

with initial ensemble distribution $\hat{\mu}_{0, 0}^{B_0}$ with particles drawn from $\hat{\mu}_{0, 0}$. Hence, $\hat{\mu}_{T_0, 0}^{B_0}$ is a random measure that is within the convergence radius of the expected 2-Wasserstein distance to μ_{s_1} and the setting of Assumption 4.1 holds again. We then define the first part of the horizon partition as

$$\mathcal{T}(1) = T_0.$$

With similar arguments for arbitrary but fixed $\epsilon_k < \underline{\delta}$, we define maximum s_{k+1} with $s_k \leq s_{k+1} \leq 1$ inductively for $k = 1, \dots$, satisfying

$$\phi(s_{k+1} - s_k) + \epsilon_k < \delta(s_{k+1}).$$

Then, for the chosen ϵ_k we again find a time horizon T_k and sample size B_k such that

$$\mathbb{E} \left[\mathcal{W}_2(\mu_{s_{k+1}}, \hat{\mu}_{T_k, s_k}^{B_k}) \right] \leq \delta(s_{k+1}).$$

This yields the time horizon update

$$\mathcal{T}(k+1) := \mathcal{T}(k) + T_k.$$

Since $\underline{\delta} > 0$ and ϕ is continuous, this procedure stops in finite time. Concretely, there exists $K \in \mathbb{N}$ with $s_K = 1$. \square

Remark 4.5. *Provided that Theorem 4.1 holds true for $\mathcal{H}(s)$ instead of ϕ and $\exp(-\mathcal{H}(s))$ is proportional to a Gaussian for some $s \in [0, 1]$, we can then set $\delta(s) = \infty$. This becomes especially relevant when choosing $\Psi = \mathcal{H}(0)$ such that $\exp(-\Psi)$ is proportional to a Gaussian approximation of μ_* in the non-linear setup.*

Remark 4.6. *The design of the piecewise constant drift term is motivated by the numerical realisation of the homotopy approach. In a more general framework, one may want to define an inhomogenous drift term $f(t, Y)$ using some proper time scaling between the SDE time t and the homotopy switch design s of the form $s = s(t)$. This would lead to a particle system of the form*

$$dy_t^{(i)} = f(\mathcal{H}(s(t)), Y_t) dt + \Gamma(Y_t) dW_t^{(i)}.$$

The analysis of the requirements on general $s(t)$ as in remark 4.6 are out of the scope of this work, but we discuss several designs in the following numerics section.

5 Numerical examples

This section is devoted to the numerical investigation of the ideas presented in this manuscript. Since our error analysis is carried out in the 2-Wasserstein metric, numerical errors are discussed to some extent in this error discrepancy. In particular, since we are concerned with the setup of ensembles, only discrete measures are examined. It is well-known that the computation of the 2-Wasserstein distance of two discrete measures is equivalent to a constrained assignment problem. Its computation can easily become cumbersome with a growing number of samples. Consequently, we perform the numerical error analysis in terms of an approximation of the Wasserstein error, namely the debiased Sinkhorn metric \mathcal{S}_ϵ introduced in Section 3.1.4. For the approximation error introduced by the Sinkhorn metric in terms of the regularizing parameter $\epsilon > 0$ we refer to [17]. Throughout this section, we choose $\epsilon = 0.1$. The computation of the debiased Sinkhorn metric is realised with the python package `GeomLoss` by Jean Feydy [12].

Random error variables Since we want to track the convergence in the case of finite batch size \bar{b} as reliably as possible, we define the following random error variables

$$\text{EP}_t: \omega \mapsto \mathcal{S}_\epsilon(\hat{\mu}_t, \mu_*^{(\bar{b})}), \quad (t\text{-ensemble-posterior}) \quad (5.1)$$

$$\text{PP}: \omega \mapsto \mathcal{S}_\epsilon(\mu_*^{(\bar{b})}, \tilde{\mu}_*^{(\bar{b})}). \quad (\text{posterior-posterior}) \quad (5.2)$$

Here, the ensemble distribution $\hat{\mu}_t$ is defined by (4.1) and $\mu_*^{(\bar{b})}, \tilde{\mu}_*^{(\bar{b})}$ are independent empirical measures of \bar{b} samples drawn from the true posterior distribution, respectively. As in the previous sections,

\bar{b} corresponds to the total number of posterior samples that are produced with our method. For the sake of readability, we suppress the dependence of $\mathbb{E}P_t$ on the particular used batch sizes for $\hat{\mu}_t$ and the total batch size \bar{b} as this becomes clear from the context. Furthermore, since we are dealing with particular instances of ensemble trajectories, the corresponding ensemble distributions are random measures and thus it is natural to investigate the expected Wasserstein (Sinkhorn) error

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_t, \mu_*^{(\bar{b})})] \approx \mathbb{E}[\mathcal{S}_\epsilon(\hat{\mu}_t, \mu_*^{(\bar{b})})] = \mathbb{E}[\mathbb{E}P_t]. \quad (5.3)$$

Under the assumption that (4.7) is satisfied for sufficiently large T , this expectation can be controlled for suitable configurations of ALDI/LIDL since by Theorem D.1 and Corollary 4.2 for any $\delta > 0$ there exist a time T_δ and a batch size \bar{b}_{T_δ} such that

$$\mathbb{E}[\mathcal{W}_2(\hat{\mu}_t, \mu_*^{(\bar{b}_{T_\delta})})] \leq \mathbb{E}[\mathcal{W}_2(\hat{\mu}_t, \mu_*)] + \mathbb{E}[\mathcal{W}_2(\mu_*, \mu_*^{(\bar{b}_{T_\delta})})] \leq \delta. \quad (5.4)$$

However, for fixed batch sizes \bar{b} , (5.3) cannot be expected to be close to 0 for any $t > 0$ unless \bar{b} tends to infinity. In order to still track some type of convergence based on (5.3) for a finite number of particles, we consider the so-called expected *posterior-posterior* error for finite batch size given by

$$\mathbb{E}[\mathcal{W}_2(\mu_*^{(\bar{b})}, \tilde{\mu}^{(\bar{b})})] \approx \mathbb{E}[\mathcal{S}_\epsilon(\mu_*^{(\bar{b})}, \tilde{\mu}^{(\bar{b})})] = \mathbb{E}[\text{PP}]. \quad (5.5)$$

The random variable PP and its expectation $\mathbb{E}[\text{PP}]$ are constant in t . Moreover, since PP is defined upon different realisations of finite posterior samples it is a non-negative random variable. However, its expectation $\mathbb{E}[\text{PP}]$ goes to 0 in the limit $\bar{b} \rightarrow \infty$. In the numerical error analysis we then examine the convergence

$$\mathbb{E}P_t \xrightarrow{d} \text{PP}, \quad (5.6)$$

$$\mathbb{E}[\mathbb{E}P_t] \longrightarrow \mathbb{E}[\text{PP}], \quad (5.7)$$

as $t \rightarrow \infty$. We remark that the investigation of the expected 2-Wasserstein error is also motivated by the so-called *concentration around expectation* property. In particular, let $\mu \in \mathcal{W}_2(\mathbb{R}^D)$ and $\mu^{(B)}$ be an empirical measure determined by B i.i.d. samples drawn from μ . Then, McDiarmid's inequality yields [43]

$$\mathbb{P} \left(\mathcal{W}_2^2(\mu, \mu^{(B)}) \geq \mathbb{E}[\mathcal{W}_2^2(\mu, \mu^{(B)})] + t \right) \leq \exp(-2Bt^2).$$

Double Sinkhorn error As discussed above, the aforementioned quantities do not converge to 0 for a finite batch size. The random variables $\mathbb{E}P_t$ and PP are non-negative scalar-valued. Due to the ergodic property of the ALDI propagation [16], we expect that $\mathbb{E}P_t$ converges to PP in distribution (5.6). Let $\mu_{\mathbb{E}P_t}$ and μ_{PP} be the distribution of $\mathbb{E}P_t$ and PP, respectively. Then, we define the so-called *double Sinkhorn error* as the mapping

$$t \rightarrow \mathcal{S}_\epsilon(\mu_{\mathbb{E}P_t}, \mu_{\text{PP}}). \quad (5.8)$$

Note that $\mathcal{S}_\epsilon(\mu_{\mathbb{E}P_t}, \mu_{\text{PP}}) \rightarrow 0$ implies (5.6). The numerical realisation of (5.8) is examined as follows. For a number of runs $R \in \mathbb{N}$, let $(\hat{\mu}_t^r)_t$, $r = 1, \dots, R$, be the family of measures generated by the r -th ALDI/LIDL run. Moreover, let

$$\mu_*^{r,\bar{b}}, \tilde{\mu}_*^{r,\bar{b}}, \tilde{\tilde{\mu}}_*^{r,\bar{b}}$$

be independent random empirical measures of \bar{b} particles sampled from the posterior for every r . Then, the sets

$$\{\mathcal{S}_\epsilon(\hat{\mu}_t^r, \mu_*^{r,\bar{b}})\}_{r=1}^R, \quad \{\mathcal{S}_\epsilon(\tilde{\mu}_*^{r,\bar{b}}, \tilde{\tilde{\mu}}_*^{r,\bar{b}})\}_{r=1}^R \quad (5.9)$$

are interpreted in this instance as an empirical measure approximating μ_{EP_t} and μ_{PP} , respectively. Finally, these empirical measures are used to approximate (5.8) with error converging to zero as $t, r \rightarrow \infty$.

Expected trajectory difference and slope To adaptively choose the enrichment times t_ℓ , we are interested in tracking the convergence speed of $\mathbb{E}[\text{EP}_t]$ to $\mathbb{E}[\text{PP}]$. In the case of convergence, for any $k \in \mathbb{N}$, we have

$$\text{DIFF}(t) := \mathbb{E}[\text{EP}_t] - \mathbb{E}[\text{EP}_{t-k\Delta t}] \longrightarrow 0, \quad (5.10)$$

which we will use for a first heuristic for adaptivity. Towards a second heuristic, suppose that we have an idealized setting with a smooth and strictly monotonically decreasing $\mathbb{E}[\text{EP}_t]$. Then, a suitable measure for the convergence speed would be the *expected trajectory slope*

$$\text{SLOPE}(t) := \frac{\mathbb{E}[\text{EP}_t] - \mathbb{E}[\text{EP}_{t-\Delta t}]}{\Delta t} = \frac{d}{dt}\mathbb{E}[\text{EP}_t] + \mathcal{O}(\Delta t), \quad t \geq \Delta t, \quad (5.11)$$

with time step $\Delta t > 0$. However, $\mathbb{E}[\text{EP}_t]$ and $\mathbb{E}[\text{EP}_{t-k\Delta t}]$ depend on the unknown posterior, which we cannot access during computation. Furthermore, a solution of ALDI/LIDL produces only a *single* realization of a family of ensemble distributions $(\hat{\mu})_t$ and hence does not provide access to the expectation. Therefore, we aim to approximate the above quantities based on information available from such a single realization $(\hat{\mu}_t)_t$. For this, we devise two simple heuristic approaches. Let $t_a \geq 0$ be a possible time point for ensemble enrichment. Moreover, let $N \in \mathbb{N}$ be a history depth with $t_a - N\Delta t \geq 0$ and $0 \leq t_1, \dots, t_N = t_a$ with $t_i = t_{i+1} - \Delta t$ for $i = 1, \dots, N-1$. The heuristics based on the difference (5.10) and the slope (5.11) respectively are given as follows.

- Let $N_1, N_2 \in \mathbb{N}$ be local averaging parameters with $N = N_1 + N_2$. We then define the *difference heuristic* by

$$\widehat{\text{DIFF}}(t_a) = \left| \frac{1}{N_1} \sum_{i=1}^{N_1} \mathcal{S}_\epsilon(\hat{\mu}_{t_i}, \hat{\mu}_{t_a}) - \frac{1}{N_2} \sum_{i=N_1+1}^N \mathcal{S}_\epsilon(\hat{\mu}_{t_i}, \hat{\mu}_{t_a}) \right|. \quad (5.12)$$

Given some threshold $0 < \text{tol} < 1$, we encourage an ensemble enrichment at time t_a if $\widehat{\text{DIFF}}(t_a) < \text{tol} * \widehat{\text{DIFF}}(t_{\text{last}})$, where t_{last} is the time point of the last enrichment, if at least one enrichment has already taken place. If no enrichment has occurred yet, we choose some small fixed positive value $\text{ref} > 0$ and set $\widehat{\text{DIFF}}(t_{\text{last}}) \equiv \text{ref}$.

- For the *slope heuristic*, let

$$\begin{aligned} \widehat{\text{SLOPE}}(t_a) &:= \frac{1}{N-2} \sum_{i=2}^{N-1} \frac{\mathcal{S}_\epsilon(\hat{\mu}_{t_i}, \hat{\mu}_{t_a}) - \mathcal{S}_\epsilon(\hat{\mu}_{t_{i-1}}, \hat{\mu}_{t_a})}{\Delta t} \\ &= \frac{1}{(N-2)\Delta t} [\mathcal{S}_\epsilon(\hat{\mu}_{t_{N-1}}, \hat{\mu}_{t_a}) - \mathcal{S}_\epsilon(\hat{\mu}_{t_1}, \hat{\mu}_{t_a})]. \end{aligned} \quad (5.13)$$

Given some threshold $0 < \text{tol} < 1$, we encourage ensemble enrichment at time t_a if $\widehat{\text{SLOPE}}(t_a) > -\text{tol} * \left| \widehat{\text{SLOPE}}(t_{\text{last}}) \right|$, where t_{last} is defined analogously to the difference heuristic.

This choice of heuristics may additionally be motivated as follows. First, instead of the expectations in (5.10) and (5.11) we use history based averaging to alleviate stochastic fluctuation. Second, we substitute $\hat{\mu}_{t_a}$ for $\mu_*^{(\bar{b})}$, since, in the ideal setup of ergodicity of ALDI with b particles, and for $t_a \rightarrow \infty$, the measure $\hat{\mu}_{t_a}$ is an instance of an empirical measure $\mu_*^{(b)}$ defined through b i.i.d. samples from the posterior μ_* .

Error with respect to forward calls To carry out the numerical investigation and numerical error analysis, we make use of the error development with respect to both the time t and the number of *forward calls*. Let Δt denote the uniform time discretisation step and let $T = n_{\text{iter}}\Delta t < \infty$ for some $n_{\text{iter}} \in \mathbb{N}$ denote the maximum time point for which we compute the solution of the underlying particle system. Moreover, let an instance of ALDI be realised with $\bar{b} \in \mathbb{N}$ particles and of LIDL with enrichment sample sizes $(b_0, \dots, b_L) \in \mathbb{N}_0^{L+1}$, where $L \in \mathbb{N}_0$ denotes the number of enrichment stages, such that $\bar{b} = \sum_{\ell=0}^L b_\ell$. Furthermore, for $L > 0$ let

$$t_\ell = k_\ell \Delta t,$$

with $0 = t_0 < t_1 < \dots < t_L < T$ for a suitable number of local iterations $k_\ell \in \mathbb{N}$ denote the time points at which the ensemble enrichment with b_ℓ additional particles takes place. Then, for $k_0 = k_{-1} := 0$ we define the number of forward calls for an instance of ALDI or LIDL for $t \leq T$ as follows:

$$\text{FC}_{\text{ALDI}}(t) = \bar{b} \cdot \arg \max_{k \in \mathbb{N}} \{k \Delta t \leq t\}, \quad (5.14)$$

$$\text{FC}_{\text{LIDL}}(t) = \bar{b}_\ell \cdot \arg \max_{k \in \mathbb{N}} \{k \Delta t \leq t - t_\ell : t_\ell \leq t \leq t_{\ell+1}\} + \sum_{\ell: t_\ell < t} \bar{b}_{\ell-1} \cdot (k_\ell - k_{\ell-1}), \quad (5.15)$$

where $\bar{b}_\ell := \sum_{m=0}^\ell b_m$ for $\ell = 0, \dots, L$ denotes the partial sum of batch sizes. When using the homotopy techniques, we implicitly make use of some function $s(t)$ relating the homotopy time scale s to the time scale t of the particle system. As discussed in Section 4.2, this $s(t)$ is chosen to be piecewise constant in t , corresponding to the piecewise constant definition of the potentials from (4.14). When required, we may interpret $t \mapsto s(t)$ in terms of the number of function calls, formally defining

$$s_{\text{ALDI}}(\text{fc}) = s(\text{FC}_{\text{ALDI}}^{-1}(\text{fc})), \quad (5.16)$$

$$s_{\text{LIDL}}(\text{fc}) = s(\text{FC}_{\text{LIDL}}^{-1}(\text{fc})), \quad (5.17)$$

for a number of function calls $\text{fc} \in \mathbb{N}$.

Outline In the following, several problems are investigated and discussed. First, some validation experiments for Gaussian posteriors are performed in Sections 5.1.1 and 5.1.2. Second, the application and efficacy of the homotopy approach in the case of multimodal posteriors is examined in Section 5.1.3. Third, we investigate a high-dimensional linear second order PDE problem in Section 5.2. In most examples, we use diffusion propagation (3.6) as our enrichment scheme of choice, as we did not discover substantial differences in performance to other methods.

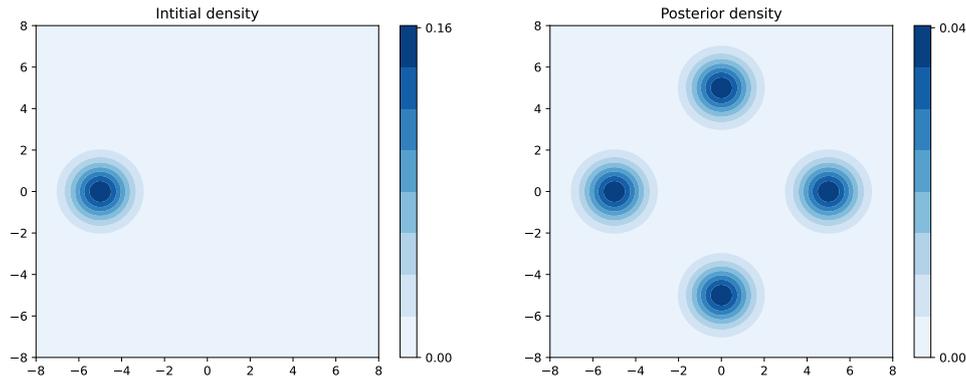


Figure 5.1: *Initial distribution (5.19) and posterior density $\exp(-\Phi)$ with Φ from (5.18) for $K = 4$.*

5.1 Gaussian mixtures: Unimodal and multimodal posterior

For a number of mixtures $K \in \{1, 4\}$, define the potential

$$\Phi : \mathbb{R}^2 \longrightarrow \mathbb{R}, \quad \Phi(x) = -\ln \left(\frac{1}{2\pi K \sqrt{|\Sigma|}} \sum_{i=1}^K \exp \left(-\frac{1}{2} \|x - x_i\|_{\Sigma}^2 \right) \right), \quad (5.18)$$

with $\Sigma = I_2$ being the identity matrix in \mathbb{R}^2 and $x_i = (\cos(i\frac{\pi}{2}), \sin(i\frac{\pi}{2}))$ for $i = 1, \dots, 4$. Then, the posterior density π_* is given by $\pi_* = \exp(-\Phi)$. As initial distribution for $\{y_0^i\}$ in the particle system, we choose a Gaussian with density

$$\pi_0(x) = \frac{1}{2\pi \sqrt{|\Sigma|}} \exp \left(-\frac{1}{2} \|x - x_3\|_{\Sigma}^2 \right). \quad (5.19)$$

This initial density π_0 and the posterior π_* corresponding to $K = 4$ are depicted in Figure 5.1.

Hence, for $K = 1$ the task of a given particle propagator consists in translating the initial density from $(-5, 0)$ to $(5, 0)$. Moreover, for $K = 4$ the initial condition is close to a local minimum of the potential Φ .

5.1.1 The unimodal case: proof of concept

First, we compare the performance of EKS [15], ALDI [16] and LIDL on the translation problem, i.e. $K = 1$. The dynamics (2.25) are discretized using the Euler-Maruyama method with step size $\Delta t = 0.05$ for a finite time horizon with $T = 10$ leading to 200 time steps. We choose a total batch size of $\bar{b} = 400$ for both methods and $L = 3$ enrichment stages with $b_0 = \dots = b_3 = 100$ for LIDL. Ensemble enrichment is done via diffusion steps (3.9) at fixed time steps $(t_1, t_2, t_3) = (1, 2, 3)$. In between ensemble enrichment stages, LIDL uses ALDI as a particle propagator.

As shown in Figure 5.2, all methods achieve similar convergence speed with regards to required time steps, but LIDL converges faster in terms of forward calls using this simple enrichment strategy. Note that ALDI and EKS perform similarly on this problem due to the large batch size $\bar{b} = 400$. The smaller \bar{b} gets the more we expect ALDI to outperform the EKS since the correction term entering (2.25) is inversely proportional to the batch size (this is demonstrated for the Darcy problem in Section 5.2). ALDI's ergodicity even for small batch sizes $b > D + 1$ is crucial for LIDL's ensemble enrichment strategy, which is why we use ALDI instead of the EKS as a particle mover. Due to its superior performance, we mainly use ALDI as a benchmark henceforth.

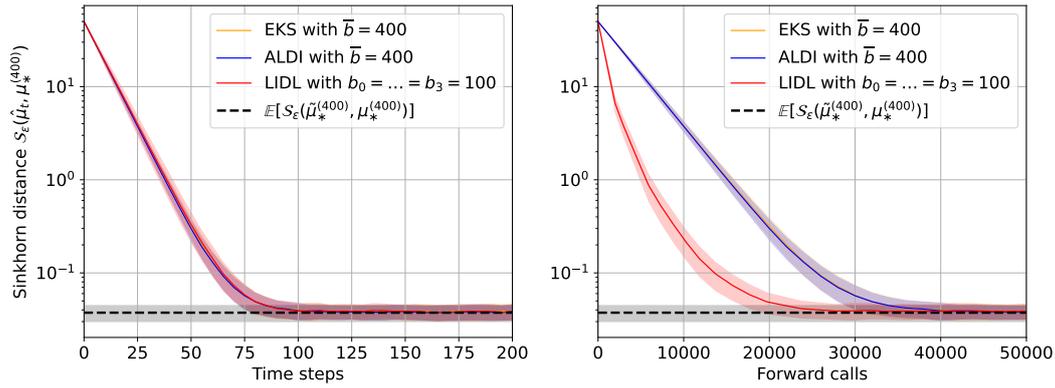


Figure 5.2: *Convergence history of the EKS, ALDI and LIDL for the case of translation problem with potential Φ from (5.18) with $K = 1$. The shaded areas and the bold interior lines represent standard deviation and mean of the error random variables PP (black) from (5.2) and EP_t from (5.1) generated by the EKS (orange), ALDI (blue) and LIDL (red). While LIDL leads to similar convergence speed with respect to the time scale (left) it outperforms the classical particle systems in terms of required forward calls (right).*

5.1.2 Motivating adaptive Enrichment

A simple equidistant enrichment scheme as in Section 5.1.1 is by no means optimal. In particular, the LIDL method used in Figure 5.2 is configured to trace ALDI in terms of convergence over the time steps (left plot in Figure 5.2). However, this is not necessarily desirable as Figure 5.3 shows. Here, two ALDI runs with a total particle number of $\bar{b} = 50$ and $\bar{b} = 400$ each are compared with a LIDL run with $b_0 = 50, b_1 = 350$ and the equidistant enrichment $b_i = 100$ for $i = 0, \dots, 3$ from Section 5.1.1. The ensemble enrichment to obtain 350 samples from 50 samples is realised with diffusion propagation (3.9) using $\delta = k\Delta t$ with $k = 1, \dots, 7$. The non-equidistant enrichment approach clearly lacks behind the ALDI scheme with $\bar{b} = 400$ particles with respect to the time steps (left plot). But examining the convergence with respect to forward calls (right plot), we see that this behaviour is actually beneficial as long the solver is in the region of decay for that particular particle number $\bar{b} = 50$. Roughly speaking, LIDL follows the trajectory of a 50-particle ALDI run as long as it has not converged. Once convergence with that particle number is reached, LIDL jumps (by means of the ensemble enrichment) onto the trajectory corresponding to a 400-particle ALDI run. As can be seen in the right plot in Figure 5.3, hardly any convergence speed is lost due to the enrichment. Here, the gray shaded areas can be seen as the maximum possible accuracy for the specific choice of finite batch size $B = 50, 400$ up to the desired amount of particles \bar{b} .

In Figure 5.4, all four methods are tracked with respect to the double Sinkhorn error from (5.8) approximated empirically using (5.9) with $R = 210$ independent runs for each of the four instances. We observe that at a double Sinkhorn error of around 10^{-7} the curves start to fluctuate, which is explained by the empirical approximation with $R < \infty$. When comparing both enrichment strategies for LIDL, we see that the $(50, 350)$ approach outperforms the $(100, 100, 100, 100)$ setup in terms of forward calls. This motivates to find adaptive schemes in terms of time points t_ℓ and batch sizes b_ℓ of ensemble enrichment.

Fixed batch size enrichment at heuristically determined time points While adaptivity in the batch size is a topic for future work, we already present preliminary results for our time point heuristics.

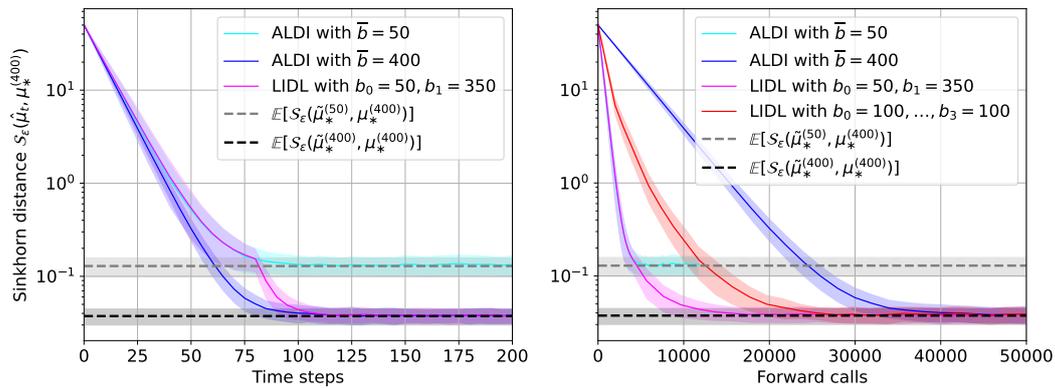


Figure 5.3: Performance of LIDL and ALDI with different choice of batch sizes and enrichment times. Instead of tracing the convergence of the 400-particle ALDI run (blue) over time steps, LIDL (magenta) follows the 50-particle ALDI run (cyan) until it is close to its accuracy limit (light gray plateau). LIDL then jumps to the 400-particle trajectory while saving a substantial number of forward calls compared to the full 400-particle ALDI run. For comparison, the suboptimal LIDL configuration from Figure 5.2 is added (red).

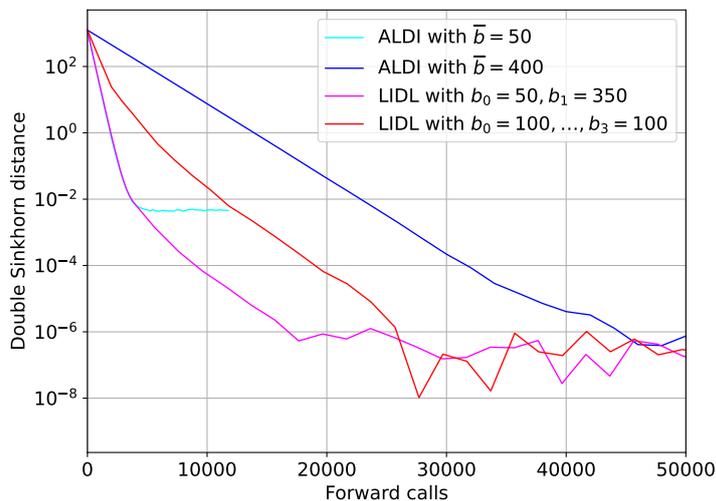


Figure 5.4: Double-Sinkhorn distance (5.8) for the four methods from Figure 5.3. While the full ALDI run with 400 particles (blue) has still not converged for 50000 forward calls, LIDL (magenta) seems to converge after ~ 30000 forward calls with a distance fluctuating stochastically around 10^{-7} .

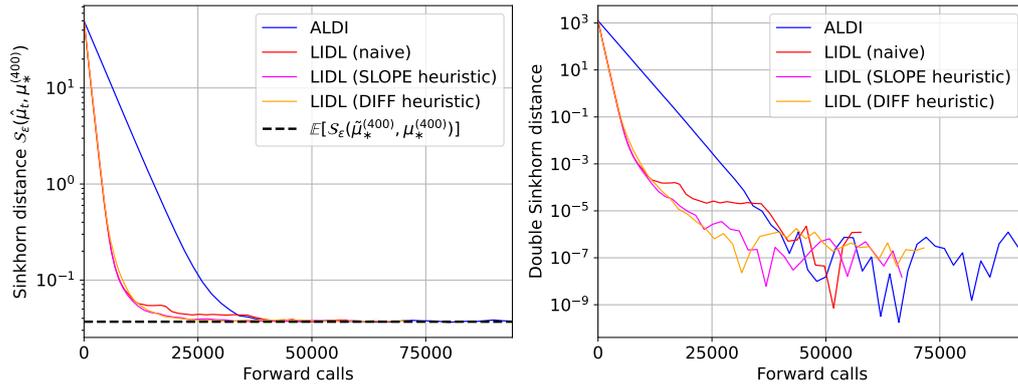


Figure 5.5: *Difference and slope heuristics applied to the unimodal Gaussian posterior from Section 5.1.1. Both heuristics clearly outperform ALDI as well as the naive LIDL implementation.*

The difference and slope heuristics (5.12) and (5.13) are designed with the goal to recognize the occurrence of plateaus as seen in the dashed grey line in Figure 5.2. In the emergence of such a plateau, (5.12) should be positive and close to 0 while (5.13) should be negative and close to 0. We deploy both heuristics under the exact same conditions as in Section 5.1.1 for a LIDL run with 3 enrichments and $b_0 = \dots = b_3 = 100$. In practice, we check the heuristics only every 5 time steps to save computation time. We choose $\text{tol} = 0.5$ and $\text{ref} = 1$ for both heuristics. In our case this leads to average enrichment times $(\bar{t}_1, \bar{t}_2, \bar{t}_3) = (2.6, 3.65, 4.75)$ for the difference heuristic and $(\bar{t}_1, \bar{t}_2, \bar{t}_3) = (3.2, 4.3, 5.95)$ for the slope heuristic. Note that the required forward calls (5.15) at time t now vary from run to run. For a fixed t , we denote the average of (5.15) over all LIDL runs by $\overline{\text{FC}}_{\text{LIDL}}(t)$. For an infinite number of runs $R \rightarrow \infty$ this becomes the expected number of forward calls at time t with a given heuristic. In Figure 5.5, we consider this average when plotting both EP_t and the double Sinkhorn distance with respect to the number of forward calls. It can be seen that for this problem both heuristics perform better than the naive LIDL configuration with fixed equidistant enrichment stages. In particular, both heuristics deploy ensemble enrichment before a suboptimal convergence plateau is reached. At the same time, the methods utilize smaller batch sizes long enough to benefit from the fast convergence speed. While the unimodal Gaussian is only a very simple test case, these promising results encourage future work on adaptive heuristics.

5.1.3 The multimodal case: homotopy approach

For $K = 4$ the task is more challenging since the center of the initial density (5.19) is now close to a local minimum of the potential (5.18) (see Figure 5.1). Figure 5.6 showcases the particle movement for $K = 4$ with ALDI-based LIDL. In Stage 1, an auxiliary potential Ψ – in this case a zero mean Gaussian with covariance matrix 5Σ – is used to precondition the data. In this particular case, the auxiliary potential is used to get the particles out of the vicinity of the local minimum near $(-5, 0)$. In stage 2, LIDL is deployed with successive ensemble enrichments in $L = 3$ stages and a homotopy based potential \mathcal{H} interpolating between $\mathcal{H}(0) = \Psi$ and $\mathcal{H}(1) = \Phi$.

Comparison of homotopy switch designs $s(t)$ We now shift the focus to homotopy based approaches and compare these with a standard ALDI, which we denote as *plain* ALDI. In the experiment, $\bar{b} = 200$ particles are propagated for a time horizon $T = 40$ using step size $\Delta t = 0.01$ leading to a total of 800 iterations. We deploy 3 different interpolation variations s_{ALDI} for the homo-

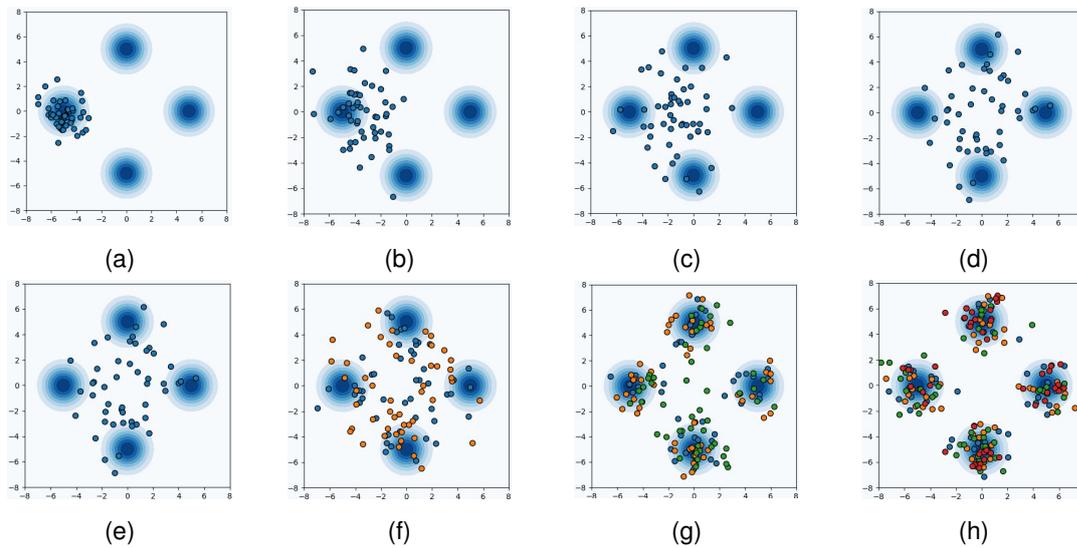


Figure 5.6: **Stage I (top row):** An auxiliary potential corresponding to a centered Gaussian distribution with covariance $5I_2$ is used to precondition a batch of $b_0 = 50$ initial samples. The particles are propagated with time moving forward from Figure (a) to (d) using ALDI. In (d), the samples follow the distribution associated to the auxiliary potential. In particular, they have been moved out of the vicinity of the local minimum near $(-5, 0)$.

Stage II (bottom row): The potential used in stage I is successively replaced via linear homotopy by the true multimodal potential (depicted as a contour plot behind the particles). The original batch (blue) of 50 particles is enriched 3 times by $b_1 = b_2 = b_3 = 50$ particles each (orange, green, red) during ALDI propagation, yielding a total batch of $\bar{b} = 200$ approximate posterior samples in (h).

topology based approach, starting with an auxiliary zero-mean Gaussian potential $\mathcal{H}(0)$ with covariance $8I_2$ for $t \in [0, 2]$ (we have chosen $5I_2$ in Figure 5.6 for the purpose of presentation. In fact, a more spread out auxiliary distribution is more advantageous). Then, for $t \in [a, b] := [2, 18]$ we use a linear $(\alpha_1 t + \alpha_2)$, convex (βt^4) and concave $(1 - \gamma(t - b)^4)$ speed of change for $\alpha_1, \alpha_2, \beta, \gamma \in \mathbb{R}$ such that the corresponding interpolation maps are continuous with values $s(t) = 0$ at $t \leq 2$ and $s(t) = 1$ at $t \geq 18$. The corresponding s_{ALDI} are illustrated in dashed lines in Figure 5.7. Note since we are using ALDI without ensemble enrichment, the representation of s with respect to time t and forward calls fc is a linear rescaling only. Since the posterior distribution is more spread out due the multimodality, we choose $R = 2000$ runs to approximate the expectations. The results of this experiment are illustrated in Figure 5.7. As a first observation, the strategy of plain ALDI leads to very slow convergence. The three homotopy based approaches achieve fast initial convergence due to the auxiliary potential $\mathcal{H}(0)$. The subsequent performance depends on the corresponding homotopy switches $s_{\text{ALDI}}(t)$. We observe that the concave version performs the best, almost carrying on with the initial convergence speed for times $t > 2$. This numerical result clearly motivates the further investigation of potential optimal choices of homotopy switch designs.

From Figure 5.7 we conclude that the design of the homotopy has a significant impact on the convergence speed. In particular, the concave case is a faster switch design compared to classical linear interpolation within the desired time horizon. We note that in the numerical investigation this switch design can be critical if the switches towards the potential Φ are too fast. Then, no convergence or at least a very slow convergence is observed asymptotically. The slow convergence then matches the convergence speed of the plain non-homotopy based SDE scheme.

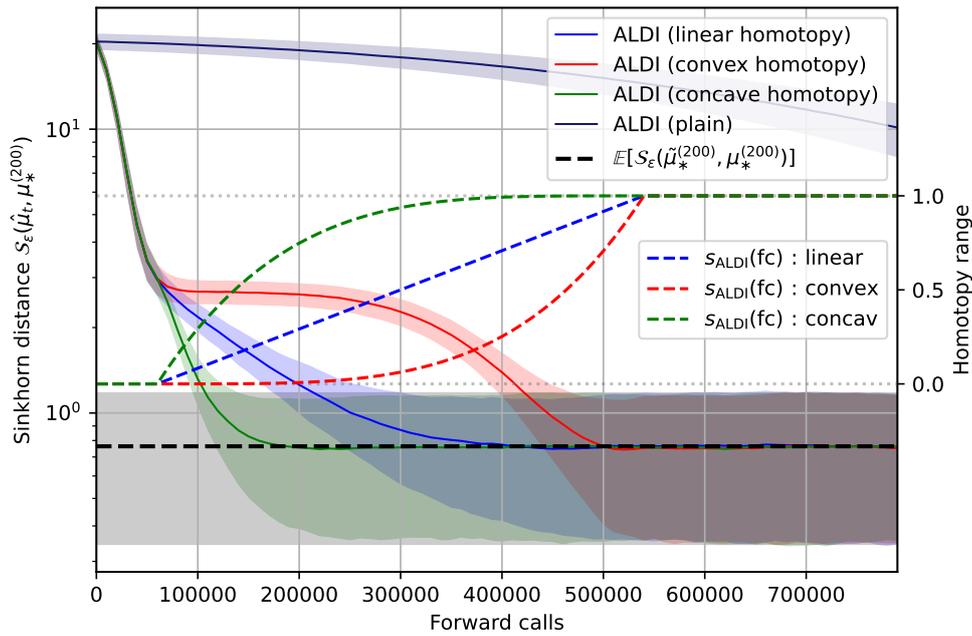


Figure 5.7: Comparison of different homotopy switch designs with the plain ALDI strategy based on $\bar{b} = 200$ particles. While the slow convergence of plain ALDI underlines the potential of the homotopy approach, possible improvements of the actual switch point design becomes clear when comparing a linear, convex and concave homotopy variation.

Combination of homotopy and ensemble enrichment As a next experiment, we investigate the interaction and the possible advantages of both, homotopy and ensemble enrichment. Here, we consider the case of linear homotopy and concave switch design combined with $L = 3$ enrichment stages. In particular for the linear homotopy we use

$$s(t) = \begin{cases} 0, & t < 0.2T, \\ \frac{10}{7T}t - \frac{2}{7}, & t \in [0.2T, 0.9T], \\ 1, & t > 0.9T, \end{cases} \quad (5.20)$$

with a time horizon $T = 40$. For the concave homotopy, we use the switch design

$$s(t) = \begin{cases} 0, & t < 0.1T, \\ 1 - \frac{1}{(0.8T)^4}(0.9T - t)^4, & t \in [0.1T, 0.9T], \\ 1, & t > 0.9T, \end{cases} \quad (5.21)$$

with $T = 40$ for ALDI and $T = 60$ for LIDL. The results of the experiments are depicted in Figures 5.8 and 5.9. Here, the time interval on which $s(t) \equiv 0$ is considered as free forward calls since no evaluation of the posterior potential Φ is performed. Consequently, in both experiments the respective ALDI and LIDL schemes reach plateaus, indicating convergence to the auxiliary measure associated to $\exp(-\Psi)$. As an observation we note that the ensemble enrichment has an impact on the actual switch design s_{LIDL} as a function of function calls as defined in (5.17). While $s_{\text{ALDI}}(\text{fc})$ is a linear rescaling of its t -dependent version $s_{\text{ALDI}}(t)$, the effect is non-linear in the case of LIDL. An interesting observation can be made in the linear homotopy case. In Figure 5.8, the enrichment scheme yields

a concave (piecewise linear) switch design (red line) when interpreted as a function of function calls. The fact that LIDL with linear homotopy turns out to require less forward calls may be explained by the faster switch speed. In particular, the use of smaller batch sizes for a fixed time horizon requires less forward calls spent on the auxiliary intermediate measures. Once the maximum batch size \bar{b} is reached, LIDL recovers the convergence speed of ALDI. We again use time step $\Delta t = 0.01$ and time horizon $T = 40$ with $R = 2000$ runs with a maximum number of $\bar{b} = 200$ particles. A total of $L = 3$ ensemble enrichment stages $(b_0, b_1, b_2, b_3) = (20, 40, 60, 80)$ are realized at $(t_1, t_2, t_3) = (12, 15, 18)$, based on forward slicing.

The ensemble enrichment strategy is more involved in the concave homotopy setup shown in Figure 5.9. Here, the construction of suitable time steps for the enrichment is not straightforward. In particular, in the numerical investigation not every enrichment concept allowed for increased convergence speed. For example, using the equidistant enrichment scheme from the linear homotopy case leads to a significant slow-down of convergence. This phenomena may be explained by observing that such equidistant enrichment leads to a switch design $s_{\text{LIDL}}(\text{fc})$ that changes too quickly towards the posterior potential.

Consequently, the homotopy and the sample enrichment design in general need to interact. In our experiment we construct an a priori enrichment scheme as follows. Let $0 = s_0 < s_1 < s_2 < s_3 < s_4 = 1$ be an equidistant partition of the switch range $[0, 1]$. We then define the time point for an enrichment and some $\gamma > 0$ by

$$t_i := \arg \min |s(t) - s_i^\gamma|, \quad i = 1, 2, 3. \quad (5.22)$$

In the experiment we use $\gamma = 1.0$ and the batch size enrichment $b_\ell = 100$ for $\ell = 0, \dots, L$.

The chosen concave switch design of order 4 seems to be a reasonable choice. It can be observed that significantly increasing the interpolation speed leads to a severe slow-down of the convergence speed. This effect may be controlled by choosing smaller time steps. While an increased switch speed may enable a good approximation of the posterior for a smaller time horizon T , the introduced stiffness in the SDE system may be computationally prohibitive due to the necessity of a smaller time step Δt . Consequently, a proper automatic design of the homotopy speed and the enrichment stages should also involve an adaptive time stepping scheme. In Figure 5.9, we additionally plot the current partial batch sizes. The initial lower batch sizes lead to a faster switch (red dashed line) and increase the convergence speed locally. However, due to the choice of $T = 60$ for LIDL the switch $s_{\text{LIDL}}(\text{fc})$ becomes slower than $s_{\text{ALDI}}(\text{fc})$, leading to an expected slow-down of convergence. In this example, a total of $L = 3$ ensemble enrichment stages $(b_0, b_1, b_2, b_3) = (50, 50, 50, 50)$ are realized at $t_i > 0$ as defined in (5.22) based on diffusion propagation. While Figure 5.9 displays the result of one hand crafted design, it may still be far away from optimality. This motivates the development of fully adaptive schemes that are capable to automatically construct promising switch and enrichment stage designs including time step adaptivity to adjust the switch speed.

For the enrichment schemes, we use forward slicing in the linear homotopy case and diffusion propagation for the concave setup. In the numerical investigation, no significant difference between the schemes can be observed.

5.2 High dimensional example: parametric Darcy equation

In this section the one-dimensional parametric Darcy equation as discussed in [16] is investigated. Consider the inverse problem of determining the permeability field $a(x) > 0$ based on the solution p

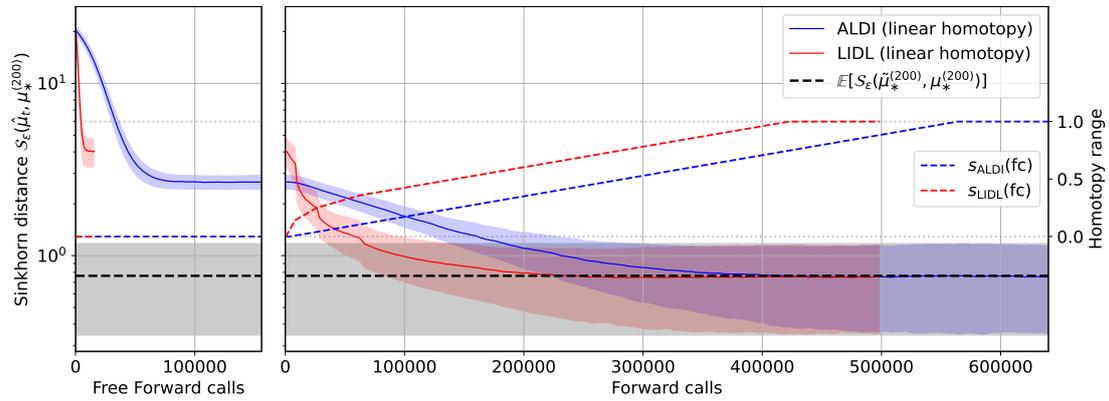


Figure 5.8: LIDL with ensemble enrichment and homotopy approach applied to the potential Φ from (5.18) for $K = 4$. Both ALDI and LIDL use auxiliary potential Ψ corresponding to a centered Gaussian distribution with covariance $8I$. Here, we consider the linear homotopy defined in (5.20). The left subplot shows the convergence when using the potential $\mathcal{H}(s)$ for $s = 0$ referred to as free forward calls since no evaluation of Φ is involved. The combination of linear homotopy and sample enrichment leads to faster convergence.

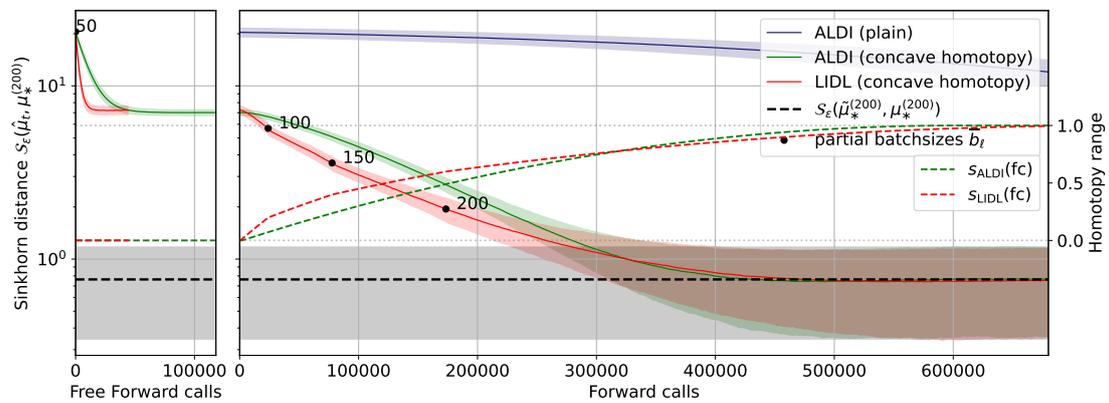


Figure 5.9: LIDL with ensemble enrichment and homotopy approach applied to the potential Φ from (5.18) for $K = 4$ with concave switch design $s(t)$ defined in (5.21). Both ALDI and LIDL use auxiliary potential Ψ corresponding to a centered Gaussian distribution with covariance $8I$. The black dots mark the partial batch sizes used in the subsequent propagation within the LIDL scheme up to $\bar{b} = 200$ particles. The left subplot displays the convergence when using the potential $\mathcal{H}(s)$ for $s = 0$ referred as free forward calls since no evaluation of Φ is involved. This particular design leads to fast initial convergence but is still suboptimal.

of an elliptic PDE

$$-\partial_x(a(x)\partial_x p(x)) = f(x), \quad \text{for } x \in I = [0, 2\pi).$$

For $K = 10$ we define observation points \hat{x}_k in I by

$$\hat{x}_k = \frac{2\pi(k-1)}{K}, \quad k = 1, \dots, K,$$

and consider noisy measurements δ_k of the solution p at these discrete points defined by

$$\delta_k := p(\hat{x}_k) + \eta_k,$$

where the measurement errors η_j are i.i.d. Gaussians with zero mean and variance $\sigma_R = 10^{-4}$. As in [16], this infinite-dimensional problem is made finite-dimensional by introducing a computational grid

$$x_i := \frac{2\pi i}{D}, \quad i = 0, \dots, D-1.$$

We choose $D = 50$ for the computational grid $\mathcal{M} = \{x_1, \dots, x_D\}$ and consider the finite-difference formulation

$$\frac{a_{i+1/2}(p_{i+1} - p_i) - a_{i-1/2}(p_i - p_{i-1})}{h^2} = -f_i,$$

with mesh size $h = \frac{2\pi}{D}$ and $p_i \approx p(x_i)$, $f_i = f(x_i)$ and $a_{i-1/2} = \exp(u_i)$. Since $\{f_i\}_{i=0}^{D-1}$ is known and fixed, the forward operator \mathcal{G} maps $\{u_i\}_{i=1}^D \in \mathbb{R}^D$ to the restriction of the discrete solution $\{p_i\}_{i=0}^{D-1} \in \mathbb{R}^D$ to the observation grid $\{\hat{x}_j\}_{j=1}^K \in \mathbb{R}^K$. Note that the observation grid has to be a subset of the computational grid.

In this example, we use the forcing

$$f_i = \exp\left(\frac{-(2x_i - 2\pi)^2}{40}\right) - \frac{3}{5}.$$

The prior on u is defined to be centered Gaussian and covariance matrix P_0 defined by

$$P_0^{-1} = 4h \left(\frac{\mu}{D} 1_D 1_D^\top - \Delta_h \right)^2,$$

where Δ_h is the standard finite-difference discretization of the Laplace operator with periodic boundary on I with uniform mesh width h , $1_D = (1, \dots, 1)^\top \in \mathbb{R}^D$ and $\mu = 10^2$ leads to a penalization of deviations of the spatial mean of u away from 0.

Having all these quantities in place, we generate observations according to

$$\delta_j = p_\ell + \eta_j, \quad \ell = \frac{D}{K}j = 5j, \quad \eta_j \sim \mathcal{N}(0, \sigma_R), \quad j = 1, \dots, K,$$

where the indices ℓ are defined such that all observations lie on the observation grid. The discrete observations p_i on the computational grid are generated with

$$a_{i-1/2}^\dagger = \exp(u_i^\dagger), \quad u_i^\dagger = \frac{1}{2} \sin(x_i - h/2),$$

for $i = 1, \dots, D$.

In Figure 5.10, different sampling methods are applied to the Darcy problem. We consider EKS/ALDI setups with a total number of samples $\bar{b} = 240$ as well as EKS/ALDI-based LIDL setups with

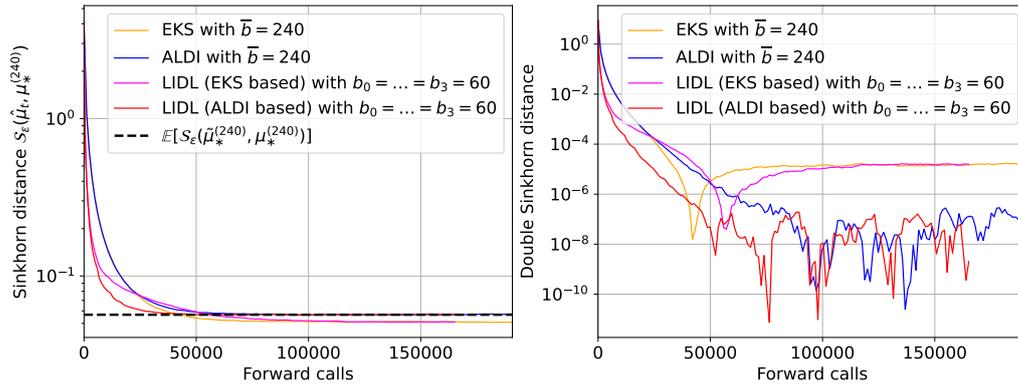


Figure 5.10: Performance of different samplers with $\bar{b} = 240$ for the Darcy problem measured in expected Sinkhorn divergence $\mathbb{E}P_t$ (left) and double Sinkhorn distance (5.8) (right). While ALDI-based LIDL seems to converge the fastest with a double Sinkhorn distance of 10^{-8} after ~ 50000 function calls, EKS and EKS-based LIDL do not converge to the correct posterior distribution. Posterior samples are generated with an affine invariant Markov Chain Monte Carlo (MCMC) sampler using the Python package `emcee` [13].

$(t_1, t_2, t_3) = (1, 1.5, 1.75)$ and $b_0 = \dots = b_3 = 60$, i.e. batch sizes only slightly above the minimum amount $b = 52$ of samples necessary to ensure ALDI's ergodicity in the linear case. The time step is $\Delta t = 0.01$ and all methods are run up to $T = 8$. Expectations are approximated by empirical averages over $R = 70$ runs each. As is expected, due to the rather small batch size EKS and EKS-based LIDL do not converge to the true posterior distribution. This can be observed in the Sinkhorn distance plot on the left, where both methods produce an expectation slightly lower than $\mathbb{E}[S_\epsilon(\tilde{\mu}_*^{(240)}, \mu_*^{(240)})]$. In the right plot, we can observe a stagnation in the double Sinkhorn distance (5.8) at around 10^{-5} for the EKS-based methods, while ALDI and ALDI-based LIDL seem to fluctuate around 10^{-8} . Comparing all methods, ALDI-based LIDL seems to perform best, achieving a double Sinkhorn distance of 10^{-8} in roughly half the amount of forward calls ($\sim 50,000$) as ALDI ($\sim 90,000$).

6 Conclusion and Outlook

This work devises a strategy which can significantly reduce the computational effort in state-of-the-art ensemble samplers based on Langevin dynamics to solve Bayesian inference problems. A significant improvement when compared to previous techniques is achievement by the notions of intermediate ensemble enrichment and \mathcal{W}_2 -stable homotopy maps.

To increase the number of particles during propagation, several ensemble enrichment strategies were introduced. These are designed with the underlying goal to preserve the distance of the ensemble distribution to the target measure. The dynamical system is then restarted with the new set of particles as initial condition.

The concept of homotopy was introduced to enhance (or sometimes to enable at all) the convergence to more involved target distributions. Moreover, this framework allows to use standard estimators for covariance matrices within the dynamical system. This is advantageous in comparison to weighted estimators as discussed in [33], allowing for very cheap computations of a (generalized) square root of the covariance that enters the diffusion part of the underlying SDE.

A convergence analysis was carried out for both concepts. In the numerical examples, the efficacy of the method was demonstrated for unimodal and multimodal posterior distributions with Gaussian tails. Additionally, a high-dimensional and non-linear PDE problem given by the random Darcy equation was investigated.

As an outlook, we suggest two promising directions for future work:

- **$(t_\ell, b_\ell, \Psi, s(t))$ -adaptivity of the Sampler.** We have only presented simple heuristics for adaptivity with respect to the enrichment times t_ℓ in (5.12) and (5.13). More sophisticated strategies, involving the enrichment batch sizes b_ℓ as well as adaptive adjustments of the surrogate potential Ψ and homotopy switch $s(t)$ are desirable. While this paper proves the general feasibility of the ensemble enrichment and homotopy approaches, the question of *when* to add *how many* samples to facilitate fast convergence is an open research problem.
- **Hierarchy of approximate potentials.** The homotopy approach is based on the idea to subsequently increase the complexity of the potential, here interpreted as a preconditioning tool. In reality, the posterior potential Φ is not available but instead only some approximation Φ_h with $\Phi_h \rightarrow \Phi$ as $h \rightarrow 0$ can be used. By introducing a monotone decreasing function $h = h(t)$ representing the approximation error over time, we may write the inhomogeneous drift term as

$$b(t, Y_t) = b(\mathcal{H}(s(t), h(t)), Y_t),$$

where now $\mathcal{H}(s, h)$ is the interpolation at switch point s between an auxiliary potential and Φ_h . A suitable design of $t \rightarrow (s(t), h(t))$ then potentially leads to a significantly decreased workload for reaching a desired accuracy threshold.

Acknowledgements

We thank Christian Bayer, Paul Hagemann, Matthias Liero and Claudia Schillings for fruitful discussions on the topics of this paper. ME acknowledges the partial support by the DFG SPP 2998 “Theoretical Foundations of Deep Learning”. RG acknowledges the support by the DFG SPP 1886 “Polymorphic Uncertainty Modelling for the Numerical Design of Structures”. DS acknowledges support by the ProFIT project “ReLkat – Reinforcement Learning for complex automation engineering”. This study does not have any conflicts to disclose.

References

- [1] D. Blömker, C. Schillings, P. Wacker, and S. Weissmann. Well posedness and convergence analysis of the ensemble Kalman inversion. *Inverse Problems*, 35(8):085007, 2019.
- [2] M. Brennan, D. Bigoni, O. Zahm, A. Spantini, and Y. Marzouk. Greedy inference with structure-exploiting lazy maps. *Advances in Neural Information Processing Systems*, 33:8330–8342, 2020.
- [3] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- [4] E. Calvello, S. Reich, and A. M. Stuart. Ensemble Kalman methods: A mean field perspective, 2022.

- [5] D. Calvetti, M. Dunlop, E. Somersalo, and A. Stuart. Iterative updating of model error for Bayesian inversion. *Inverse Problems*, 34(2):025008, 2018.
- [6] J. A. Carrillo and U. Vaes. Wasserstein stability estimates for covariance-preconditioned Fokker–Planck equations. *Nonlinearity*, 34(4):2275–2295, feb 2021.
- [7] V. David, T. Galaon, and E. Bacalum. Sample enrichment by solid-phase extraction for reaching parts per quadrillion levels in environmental analysis. *Chromatographia*, 82, 08 2019.
- [8] Z. Ding and Q. Li. Ensemble Kalman inversion: Mean-field limit and convergence analysis. *Statistics and Computing*, 31(1):1–21, 2021.
- [9] Z. Ding and Q. Li. Ensemble Kalman sampler: Mean-field limit and convergence analysis. *SIAM Journal on Mathematical Analysis*, 53(2):1546–1578, 2021.
- [10] S. S. Dragomir and M. City. Some Gronwall type inequalities and applications. URL: <http://rgmia.vu.edu.au/SSDragomirWeb.html>, 2002.
- [11] G. Evensen. *Data Assimilation: The Ensemble Kalman Filter*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [12] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trouvé, and G. Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2681–2690. PMLR, 2019.
- [13] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *PASP*, 125:306–312, 2013.
- [14] N. Fournier and A. Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3):707–738, 2015.
- [15] 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.
- [16] A. Garbuno-Inigo, N. Nusken, and S. Reich. Affine invariant interacting Langevin dynamics for Bayesian inference. *SIAM Journal on Applied Dynamical Systems*, 19(3):1633–1658, 2020.
- [17] A. Genevay, L. Chizat, F. Bach, M. Cuturi, and G. Peyré. Sample complexity of Sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583. PMLR, 2019.
- [18] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80, 2010.
- [19] T. H. Gronwall. Note on the derivatives with respect to a parameter of the solutions of a system of differential equations. *Annals of Mathematics*, pages 292–296, 1919.
- [20] R. Gruhlke and M. Eigel. Low-rank Wasserstein polynomial chaos expansions in the framework of optimal transport. *arXiv preprint arXiv:2203.09358*, 2022.
- [21] P. Hagemann, J. Hertrich, and G. Steidl. Stochastic normalizing flows for inverse problems: a markov chains viewpoint. *SIAM/ASA Journal on Uncertainty Quantification*, 10(3):1162–1190, 2022.

- [22] N. J. Higham. *Functions of Matrices*. Society for Industrial and Applied Mathematics, 2008.
- [23] M. A. Iglesias, K. J. Law, and A. M. Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.
- [24] J. Kaipio and E. Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.
- [25] K. Law, A. Stuart, and K. Zygalakis. Data assimilation. *Cham, Switzerland: Springer*, 214, 2015.
- [26] Q. Liu and D. Wang. Stein variational gradient descent: A general purpose Bayesian inference algorithm. *Advances in neural information processing systems*, 29, 2016.
- [27] Y. Marzouk, T. Moselhy, M. Parno, and A. Spantini. An introduction to sampling via measure transport. *arXiv preprint arXiv:1602.05023*, 2016.
- [28] N. Nüsken and S. Reich. Note on interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. *arXiv preprint arXiv:1908.10890*, 2019.
- [29] B. Øksendal. Stochastic differential equations. In *Stochastic differential equations*, pages 65–84. Springer, 2003.
- [30] G. Pavliotis and A. Stuart. *Multiscale methods: averaging and homogenization*. Springer Science & Business Media, 2008.
- [31] A. Perov. K voprosu o strukture integral'noi voronki. *Nauc. Dokl. Vysšei Školy. Ser FMN*, 2, 1959.
- [32] S. Reich and C. Cotter. *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press, 2015.
- [33] S. Reich and S. Weissmann. Fokker–Planck particle systems for Bayesian inference: Computational approaches. *SIAM/ASA Journal on Uncertainty Quantification*, 9(2):446–482, 2021.
- [34] D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- [35] C. Robert and G. Casella. A Short History of Markov Chain Monte Carlo: Subjective Recollections from Incomplete Data. *Statistical Science*, 26(1):102 – 115, 2011.
- [36] G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probability surveys*, 1:20–71, 2004.
- [37] G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.
- [38] 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.
- [39] C. Schillings and A. M. Stuart. Convergence analysis of ensemble Kalman inversion: the linear, noisy case. *Applicable Analysis*, 97(1):107–123, 2018.
- [40] D. Stalter, L. I. Peters, E. O'Malley, J. Y.-M. Tang, M. Revalor, M. J. Farré, K. Watson, U. von Gunten, and B. I. Escher. Sample enrichment for bioanalytical assessment of disinfected drinking water: Concentrating the polar, the volatiles, and the unknowns. *Environmental Science & Technology*, 50(12):6495–6505, 2016. PMID: 27153244.

- [41] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta numerica*, 19:451–559, 2010.
- [42] C. Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [43] J. Weed and F. Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- [44] H. Wu, J. Köhler, and F. Noé. Stochastic normalizing flows. *Advances in Neural Information Processing Systems*, 33:5933–5944, 2020.

A Proof of theorem 4.1

The proof follows closely the proof of a comparable result for the Ensemble Kalman Sampler shown in [9, Theorem 3.1]. In the linear case, the dynamics of ALDI (2.25) become

$$\begin{aligned} dy_t^i &= -\text{Cov}_{y_t, \mathcal{G}_t} \Gamma^{-1} (\mathcal{G}(y_t^i) - \tilde{\Delta}) dt - \text{Cov}_{y_t} \Gamma_0^{-1} (y_t^i - y_0) dt + \frac{D+1}{B} (y_t^i - \bar{y}_t) dt + \sqrt{2\text{Cov}_{y_t}} dW_t^i \\ &= -\text{Cov}_{y_t} P (y_t^i - y^*) dt + \frac{D+1}{B} (y_t^i - \bar{y}_t) dt + \sqrt{2\text{Cov}_{y_t}} dW_t^i, \quad i = 1, \dots, B. \end{aligned} \quad (\text{A.1})$$

Here, we adopt the notation from [9] for covariance matrices, i.e.

$$\begin{aligned} \text{Cov}_{x_t, y_t} &= \frac{1}{B} \sum_{i=1}^B (x_t^i - \bar{x}_t) (y_t^i - \bar{y}_t)^\top, \quad \text{Cov}_{y_t} = \text{Cov}_{y_t, y_t} \\ \text{Cov}_\rho &= \mathbb{E}_{x \sim \rho} [(x - \mathbb{E}_\rho) (x - \mathbb{E}_\rho)^\top], \\ \text{Cov}_{y_t, \mathcal{G}_t} &= \text{Cov}_{y_t} A^\top, \end{aligned}$$

for particle ensembles $\{x_t^i\}_{i=1}^B$, $\{y_t^i\}_{i=1}^B$ and densities ρ . For better readability, we use this form for the rest of the Appendix. We will frequently make use of Itô's lemma, stating that for any twice differentiable function f of an Itô diffusion process X_t , we have

$$dX_t = \mu_t dt + \sigma_t dW_t \implies df(X_t) = \left[(\nabla f)^\top \mu_t + \frac{1}{2} \text{Tr}[\sigma_t^\top \text{Hess}_f \sigma_t] \right] dt + (\nabla f)^\top \sigma_t dW_t \quad (\text{A.2})$$

When considering Wasserstein distances between measures, we will often identify a measure with its density, writing, for example, $\mathcal{W}_2(\hat{\mu}_t^B, \pi_*)$ instead of $\mathcal{W}_2(\hat{\mu}_t^B, \mu_*)$.

Before getting into details we will briefly sketch the main ingredients of the proof.

A.1 Roadmap of the Proof

The central idea is a triangle argument

$$\mathbb{E}(\mathcal{W}_2(\pi_*, \hat{\mu}_t^B)) \leq \underbrace{\mathcal{W}_2(\pi_*, \pi(t))}_{\text{(I)}} + \underbrace{\mathbb{E}(\mathcal{W}_2(\pi(t), \hat{\nu}_t^B))}_{\text{(II)}} + \underbrace{\mathbb{E}(\mathcal{W}_2(\hat{\nu}_t^B, \hat{\mu}_t^B))}_{\text{(III)}}, \quad (\text{A.3})$$

where $\pi(t) \equiv \pi(\cdot, t)$ is the solution of the limiting Fokker-Planck equation

$$\begin{aligned} \partial_t \pi &= \nabla \cdot (\pi \text{Cov}_{\pi(t)} \nabla \Phi(y)) + \text{Tr}(\text{Cov}_{\pi(t)} D^2 \pi), \\ \pi(\cdot, 0) &= \pi_0(\cdot), \end{aligned} \quad (\text{A.4})$$

and $\hat{\nu}_t^B$ is the ensemble distribution induced by the solution $\{z_t^i\}_{i=1}^B$ of the process

$$\begin{aligned} dz_t^i &= -\text{Cov}_{\pi(t)} \nabla \Phi(z_t^i) dt + \sqrt{2\text{Cov}_{\pi(t)}} dW_t^i \\ &= -\text{Cov}_{\pi(t)} P (z_t^i - y^*) dt + \sqrt{2\text{Cov}_{\pi(t)}} dW_t^i, \quad i = 1, \dots, B. \end{aligned} \quad (\text{A.5})$$

An intuition regarding the three terms can be given as follows. (I) concerns the distance of the posterior to the solution $\pi(t)$ of the limiting Fokker Planck equation. Since (A.5) produces i.i.d. samples of $\pi(t)$ at time t , (II) concerns the distance between π and empirical measures of B samples each, drawn from π . Finally, (III) measures the distance of the ensemble distributions defined by the ALDI right hand side to the ideal (but unavailable) ensembles defined by (A.5). It remains to be shown that, given $\delta > 0$, we find suitable T_δ and B_{T_δ} such that the right hand side in (A.3) is bounded by δ . The arguments bounding the first two terms can already be found in [9] and are repeated here for the sake of completeness.

A.2 Bounding (I)

This term is independent of the particle ensemble $\hat{\mu}_t^B$ and hence of the number of particles B . Its behaviour is governed by the solution $\pi(t)$ of the underlying Fokker-Planck equation (A.4). As an immediate consequence of the following result, the \mathcal{W}_2 -distance between $\pi(t)$ and the posterior decays to 0 exponentially fast.

Theorem A.1 (Proposition 3.8 in [6]). *Let ρ_1 and ρ_2 be two solutions of the nonlinear nonlocal mean field equation*

$$\partial_t \rho = \nabla (\text{Cov}_{\rho(t)}(\nabla \Phi(y)\rho + \sigma \nabla \rho)) \quad (\text{A.6})$$

with linear forward model \mathcal{G} . Let the covariances and expectations of the initial conditions $\rho_1(0)$ and $\rho_2(0)$ satisfy

$$\begin{aligned} |\text{Cov}_{\rho_1(0)}|_2, |\text{Cov}_{\rho_2(0)}|_2, |B^{-1}|_2 &\leq M, \\ |\text{Cov}_{\rho_1(0)}^{-1}|_2, |\text{Cov}_{\rho_2(0)}^{-1}|_2, |B|_2 &\leq m, \\ |\mathbb{E}_{\rho_1(0)}|_2, |\mathbb{E}_{\rho_2(0)}|_2 &\leq R \end{aligned} \quad (\text{A.7})$$

for some constants $M, m, R > 0$. Then it holds that

$$W_2(\rho_1(t), \rho_2(t)) \leq C(1 + m^4 M^4 + m^4 M^{7/2} R) \frac{e^{-\sigma t}}{\sqrt{\alpha(t)^{1+[1 \wedge \sigma]}}} W_2(\rho_1(0), \rho_2(0)), \quad (\text{A.8})$$

where C is a constant only dependent on the dimension d and

$$\alpha(t) = \begin{cases} 2t + 2, & \sigma = 0, \\ \sigma^{-1}(1 - e^{-2\sigma t}) + e^{-2\sigma t}, & \sigma > 0. \end{cases}$$

The Fokker Planck equation (A.6) becomes (A.4) for $\sigma = 1$. Now, setting $\rho_1(0) = \pi_0$ and $\rho_2(0) = \pi_*$, (A.8) yields (I) $\rightarrow 0$ exponentially fast with $t \rightarrow \infty$.

A.3 Bounding (II)

This term is again independent of the dynamics used to sample $\hat{\mu}_t^B$ and depends only on the rate of convergence of the ensembles $\hat{\nu}_t^B$ defined by (A.5) to their continuous limit (A.4). We cite the corresponding result from [9].

Theorem A.2 (Proposition 5.1 in [9]). *Let π solve the Fokker-Planck equation (A.4) and let $\{z_t^i\}_{i=1}^B$ solve (A.5) with initial data $\{z_{t=0}^i\}_{i=1}^B$ drawn i.i.d. from $\pi_0 \in \mathcal{C}^2$ with finite higher moments. Let $\hat{\nu}_t^B$ be the ensemble distribution defined by $\{z_t^i\}$, then for any $t > 0$ and $0 < \epsilon < 1/2$, there exists a constant C , depending on t , the dimension D and ϵ but not on B such that*

$$\mathbb{E}(\mathcal{W}_2(\hat{\nu}_t^B, \pi(t))) \leq C \begin{cases} B^{-1/2+\epsilon}, & D \leq 4, \\ B^{-2/D}, & D > 4. \end{cases} \quad (\text{A.9})$$

Proof. The proof uses Theorem D.1 and the boundedness of the higher moments of $\hat{\nu}_t^B$ (see Proposition 5.3 in [9]). \square

A.4 Bounding (III)

This is the main part of the proof, where we show that the results from [9] for the EKS carry over for the modified dynamics defined by ALDI. For each result, we will cite the corresponding result for the EKS, so that the reader may compare. We start with the main result, yielding the required boundedness.

Theorem A.3 (Compare Proposition 5.2 in [9]). *Let $\{y_t^i\}_{i=1}^B$ be the solution of (A.1) and $\{z_t^i\}_{i=1}^B$ solve (A.5), where $\{z_0^i\}_{i=1}^B = \{y_0^i\}_{i=1}^B$ are drawn i.i.d. from the distribution induced by $\pi_0 \in C^2$ with finite high moments. Let $\hat{\nu}_t^B$ and $\hat{\mu}_t^B$ denote the ensemble distributions defined by $\{z_t^i\}$ and $\{y_t^i\}$ respectively. Suppose (4.7) holds true. Then for any $0 < \epsilon < 1/2$, there exists a constant C depending only on D, T and ϵ such that*

$$\mathbb{E}(\mathcal{W}_2(\hat{\nu}_T^B, \hat{\mu}_T^B)) \leq \left(\frac{1}{B} \sum_{i=1}^B \mathbb{E}|z_T^i - y_T^i|^2 \right)^{1/2} \leq CB^{-1/2+\epsilon}. \quad (\text{A.10})$$

The assertion results from the combination of several lemmas. First, we define

$$x_t^i = y_t^i - z_t^i \quad \text{and} \quad p_t^i = x_t^i - \bar{x}_t, \quad q_t^i = z_t^i - \bar{z}_t$$

for convenience. We need that the higher moments of x_t^i and p_t^i are bounded for all time.

Lemma A.1 (Compare Lemma 5.3 in [9]). *Under the same conditions as in Theorem A.3, for all $2 \leq p < \infty$ and $T > 0$, there is a constant C_p independent of B, t such that*

$$\mathbb{E}|x_t^i|^p = \mathbb{E}|x_t^1|^p \leq C_p, \quad \mathbb{E}|p_t^i|^p = \mathbb{E}|p_t^1|^p \leq C_p \quad (\text{A.11})$$

for all $1 \leq i \leq B$ and $0 \leq t \leq T$.

Lemma A.2 (Compare Lemma 5.4 in [9]). *Under the same conditions as in Theorem 5, for any $0 \leq \alpha < 1$ and $T > 0$, if there is a constant \hat{C} independent of B, t such that*

$$\mathbb{E}|x_t^i|^2 \leq \hat{C}B^{-\alpha} \quad (\text{A.12})$$

for all $1 \leq i \leq B$ and $0 \leq t \leq T$, then for any $0 < \epsilon < 1/2$ and $1 \leq i \leq B$, there exists a constant C independent of B, t such that

$$\mathbb{E}|p_t^i|^2 = \mathbb{E} \left| x_t^i - \frac{1}{B} \sum_{k=1}^B x_t^k \right|^2 \leq \tilde{C}B^{-1/2-\alpha/2+\epsilon} \quad (\text{A.13})$$

for all $1 \leq i \leq B$ and $0 \leq t \leq T$.

Lemma A.3 (Compare Lemma 5.5 in [9]). *Under the same conditions as lemma A.2, we have for any $0 < \epsilon < 1/2$ and $T > 0$ a constant \tilde{C} independent of B, t such that*

$$\mathbb{E}|x_t^i|^2 \leq \tilde{C}B^{-1/2-\alpha/2+\epsilon} \quad (\text{A.14})$$

for all $1 \leq i \leq B$ and $0 \leq t \leq T$.

Proof of Theorem A.3. By lemma A.1, (A.12) is satisfied for $\alpha_0 = 0$, yielding

$$\mathbb{E}|x_t^i|^2 \leq \tilde{C}B^{-1/2+\epsilon}$$

by lemma A.3. Hence, (A.12) is satisfied for $\alpha_1 = 1/2 - \epsilon$ and recursively for $\alpha_{n+1} = 1/2 + \alpha_n/2 - \epsilon$. The limit of this sequence is $\lim_{n \rightarrow \infty} \alpha_n = 1 - 2\epsilon$, yielding

$$\mathbb{E}|z_t^i - y_t^i|^2 = \frac{1}{B} \sum_{i=1}^B \mathbb{E}|z_t^i - y_t^i|^2 \leq \tilde{C} B^{-1+2\epsilon}.$$

□

It remains to show the three lemmas A.1-A.3, for which we need some more intermediate results.

Lemma A.4 (Compare Lemma 4.1 from [9]). *Let $\{y_t^i\}_{i=1}^B$ be the solution of (A.1) with initial conditions $\{y_0^i\}_{i=1}^B$ sampled i.i.d from $\pi_0 \in \mathcal{C}^2$. If the initial p -th moment is finite, i.e.*

$$(\mathbb{E}|y_0^i|^p)^{1/p} < M, \quad 1 \leq i \leq B$$

for $p \geq 2$ and some $M > 0$ independent of B , then the boundedness also holds true for $\mathbb{E}|e_t^i|$, where $e_t^i = y_t^i - \bar{y}_t$, namely there is a $C > 0$ depending only on p , so that

$$(\mathbb{E}|e_t^i|^p)^{1/p} < 2(\kappa(P))^{1/2} M \exp(Ct)$$

for all $t \geq 0$ and $1 \leq i \leq B$, where $\kappa(P) = \|P\|_2 \|P^{-1}\|_2$ is the condition number of P .

Proof of Lemma A.4. Following the proof of lemma 4.1 in [9], we will show the claim for $2p$ with $p \geq 1$ and define

$$\mathbf{e}_t^i = \sqrt{P} e_t^i, \quad V_p(\mathbf{e}) = \frac{1}{B} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^p = \frac{1}{B} \sum_{i=1}^B |e_t^i|^{2p}, \quad h_p(t) = \mathbb{E}V_p = \mathbb{E}|\mathbf{e}_t^1|^{2p}.$$

Since $\lambda_{\min}(P) > 0$, boundedness of $\mathbb{E}|e_t^i|^p$ follows immediately from boundedness of $h_p(t)$, which we now show. First, we get

$$\begin{aligned} (h_p(0))^{1/2p} &= (\mathbb{E}|\mathbf{e}_0^1|^{2p})^{1/2p} \leq \|P\|_2^{1/2} (\mathbb{E}|e_0^1|^{2p})^{1/2p} \\ &\leq \|P\|_2^{1/2} \left((\mathbb{E}|y_0^1|^{2p})^{1/2p} + (\mathbb{E}|\bar{y}_0|^{2p})^{1/2p} \right) \\ &\leq 2\|P\|_2^{1/2} M, \end{aligned}$$

where we have used the boundedness of the initial moments. Our next intermediate goal is to show

$$h_p(t) \leq h_p(0) + C(p, D) \int_0^t h_p(s) ds$$

with a constant $C(p, D)$ depending on p and D . To that end, note that by (A.1) we have

$$d\bar{y}_t = -\text{Cov}_{y_t} P(\bar{y}_t - y^*) dt + \frac{D+1}{B} \sum_{i=1}^B (y_t^i - \bar{y}_t) + \sqrt{2\text{Cov}_{y_t}} d\bar{W}_t \quad (\text{A.15})$$

and with $\text{Cov}_{y_t} = \text{Cov}_{e_t} = \sqrt{P}^{-1} \text{Cov}_{e_t} \sqrt{P}^{-1}$, we find

$$\begin{aligned} de_t^i &= d\sqrt{P}(y_t^i - \bar{y}_t) = -\sqrt{P} \text{Cov}_{e_t} P e_t^i dt + \frac{D+1}{B} (\mathbf{e}_t^i - \bar{\mathbf{e}}_t) dt + \sqrt{P} \sqrt{2\text{Cov}_{e_t}} d(W_t^i - \bar{W}_t) \\ &= -\text{Cov}_{e_t} \mathbf{e}_t^i dt + \frac{D+1}{B} (\mathbf{e}_t^i - \bar{\mathbf{e}}_t) dt + \sqrt{P} \sqrt{2\text{Cov}_{e_t}} d(W_t^i - \bar{W}_t). \end{aligned} \quad (\text{A.16})$$

Noting that

$$V_p(\mathbf{e}_t) = \sum_{i=1}^B f(\mathbf{e}_t^i), \quad (\text{A.17})$$

with

$$f(x) = \frac{1}{B}|x|^{2p}, \quad \nabla f(x) = \frac{2p}{B}|x|^{2p-2}x, \quad \text{Hess}_f(x) = \frac{2p}{B}|x|^{2p-2}I_d + \frac{4p(p-1)}{B}|x|^{2p-4}xx^\top,$$

we apply Itô's lemma (A.2) to $V_p(\mathbf{e}_t)$ to receive

$$\begin{aligned} dV_p(\mathbf{e}_t) &= -\frac{2p}{B} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^{p-1} \langle \mathbf{e}_t^i, \text{Cov}_{\mathbf{e}_t} \mathbf{e}_t^i \rangle dt \\ &\quad + \frac{2p}{B} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^{p-1} \langle \mathbf{e}_t^i, \sqrt{P} \sqrt{2\text{Cov}_{\mathbf{e}_t}} d(W_t^i - \bar{W}_t) \rangle \\ &\quad + \frac{4(B-1)p(p-1)}{B^2} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^{p-2} \text{Tr} \{ (\mathbf{e}_t^i \otimes \mathbf{e}_t^i) \text{Cov}_{\mathbf{e}_t} \} dt \\ &\quad + \frac{2(B-1)p}{B^2} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^{p-1} \text{Tr} \{ \text{Cov}_{\mathbf{e}_t} \} dt \\ &\quad + \frac{2p}{B} \sum_{i=1}^B \langle \mathbf{e}_t^i, \mathbf{e}_t^i \rangle^{p-1} \left\langle \mathbf{e}_t^i, \frac{D+1}{B} \left[\mathbf{e}_t^i - \sum_{k=1}^B \mathbf{e}_t^k \right] \right\rangle dt, \end{aligned} \quad (\text{A.18})$$

for $p \geq 2$, with the third term vanishing in the case $p = 1$. In order to bound these terms in expectation, note that for $l \in \mathbb{N}$ with $l \leq p$ we have by Hölder's inequality

$$\begin{aligned} \mathbb{E} \left[\langle \mathbf{e}_t^j, \mathbf{e}_t^j \rangle^{p-l} \langle \mathbf{e}_t^k, \mathbf{e}_t^k \rangle^l \right] &\leq \mathbb{E} \left[\langle \mathbf{e}_t^j, \mathbf{e}_t^j \rangle^p \right]^{\frac{p-l}{p}} \mathbb{E} \left[\langle \mathbf{e}_t^k, \mathbf{e}_t^k \rangle^p \right]^{\frac{l}{p}} \\ &= \mathbb{E} \left[\langle \mathbf{e}_t^j, \mathbf{e}_t^j \rangle^p \right]. \end{aligned} \quad (\text{A.19})$$

Furthermore, we will make use of the inequalities

$$\langle x, y \rangle \leq \frac{\langle x, x \rangle + \langle y, y \rangle}{2} \quad \text{and} \quad \langle x, y \rangle^2 \leq \langle x, x \rangle \langle y, y \rangle \leq \frac{\langle x, x \rangle^2 + \langle y, y \rangle^2}{2}, \quad (\text{A.20})$$

holding for arbitrary elements x, y of a linear space with inner product $\langle \cdot, \cdot \rangle$. Taking the expectation in

(A.18) and noting that the first term is always nonnegative, we find

$$\begin{aligned}
h_p(t) - h_p(0) &\leq \frac{4(B-1)p(p-1)}{B^3} \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-2} \langle \mathbf{e}_s^i, \mathbf{e}_s^k \rangle^2 \right] ds \\
&\quad + \frac{2(B-1)p}{B^3} \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-1} \langle \mathbf{e}_s^k, \mathbf{e}_s^k \rangle \right] ds \\
&\quad + \frac{2p(D+1)}{B^2} \left\{ \int_0^t \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^p \right] ds - \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-1} \langle \mathbf{e}_s^i, \mathbf{e}_s^k \rangle \right] ds \right\} \\
&\leq \frac{4(B-1)p(p-1)}{B^3} \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-2} \frac{\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^2 + \langle \mathbf{e}_s^k, \mathbf{e}_s^k \rangle^2}{2} \right] ds \\
&\quad + \frac{2(B-1)p}{B^3} \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-1} \langle \mathbf{e}_s^k, \mathbf{e}_s^k \rangle \right] ds \\
&\quad + \frac{2p(D+1)}{B^2} \left\{ \int_0^t \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^p \right] ds + \int_0^t \sum_{i,k=1}^B \mathbb{E} \left[\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle^{p-1} \frac{\langle \mathbf{e}_s^i, \mathbf{e}_s^i \rangle + \langle \mathbf{e}_s^k, \mathbf{e}_s^k \rangle}{2} \right] ds \right\} \\
&\leq \frac{4(B-1)p(p-1)}{B} \int_0^t h_p(s) ds + \frac{2(B-1)p}{B} \int_0^t h_p(s) ds \\
&\quad + \frac{2p(D+1)}{B^2} \left\{ B \int_0^t h_p(s) ds + B^2 \int_0^t h_p(s) ds \right\} \\
&\leq \frac{4(B-1)p(p-1) + 2(B-1)p + 2p(D+1) + 2Bp(D+1)}{B} \int_0^t h_p(s) ds \\
&\leq (4p(p-1) + 2p + 2p(D+1) + 2p(D+1)) \int_0^t h_p(s) ds := C(p, D) \int_0^t h_p(s) ds.
\end{aligned}$$

In total we get by Gronwall's inequality

$$h_p(t) \leq (2\|B\|_2^{1/2} M)^{2p} e^{C(p,D)t}$$

and hence

$$\begin{aligned}
(\mathbb{E}|e_t^j|^{2p})^{1/2p} &= (\mathbb{E}|e_t^1|^{2p})^{1/2p} = \left(\mathbb{E}|\sqrt{P}^{-1} \mathbf{e}_t^1|^{2p} \right)^{1/2p} \\
&\leq \|P^{-1}\|_2^{1/2} (\mathbb{E}|\mathbf{e}_t^1|^{2p})^{1/2p} \\
&\leq \|P^{-1}\|_2^{1/2} \left((2\|P\|_2^{1/2} M)^{2p} e^{C(p,D)t} \right)^{1/2p} \\
&= 2M\kappa(P)^{1/2} e^{C(p,D)t}.
\end{aligned}$$

□

Lemma A.5 (Proposition 4.1 from [9]). *Let the conditions of lemma A.4 be satisfied for $2p$ (i.e. the $2p$ -th initial moments are finite). Then we have*

$$(\mathbb{E}|y_t^i - \bar{y}_t|^p)^{1/p} \leq C e^{Ct} \quad \text{and} \quad (\mathbb{E}\|\text{Cov}_{y_t}\|_2^p)^{1/p} \leq C e^{Ct}$$

as well as

$$(\mathbb{E}|y_t^i|^p)^{1/p} \leq C e^{C e^{Ct}}$$

with $C > 0$ being independent of B and t .

Proof of Lemma A.5. The first inequality follows directly from lemma A.4 and Hölder's inequality. The

second inequality follows from

$$\begin{aligned}
(\mathbb{E}\|\text{Cov}_{y_t}\|_2^p)^{1/p} &\leq \frac{1}{B} \sum_{i=1}^B \mathbb{E}(\|(y_t^i - \bar{y}_t)(y_t^i - \bar{y}_t)^\top\|_2^p)^{1/p} \\
&= \frac{1}{B} \sum_{i=1}^B (\mathbb{E}|y_t^i - \bar{y}_t|^{2p})^{1/p} \\
&= \frac{1}{B} \sum_{i=1}^B ((\mathbb{E}|y_t^i - \bar{y}_t|^{2p})^{1/2p})^2 \\
&\leq Ce^{Ct},
\end{aligned}$$

where we have again used lemma A.4 in the last inequality. Similar to the proof of lemma A.4, we now set

$$\mathbf{y}_t^i = \sqrt{P}y_t^i, \quad K_p(\mathbf{y}) = \frac{1}{B} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^p = \frac{1}{B} \sum_{i=1}^B |\mathbf{y}_t^i|^{2p}$$

and

$$g_p(t) = \mathbb{E}K_p = \mathbb{E}|\mathbf{y}_t^1|^{2p}.$$

We will show the remaining inequality for $2p$ instead of p , by which Hölder's inequality yields the assertion. The dynamics of the scaled particles \mathbf{y}_t^i is given by

$$d\mathbf{y}_t^i = -\text{Cov}_{y_t}(\mathbf{y}_t^i - \mathbf{y}^*)dt + \frac{D+1}{B}(\mathbf{y}_t^i - \bar{\mathbf{y}}_t)dt + \sqrt{B}\sqrt{2\text{Cov}_{y_t}}dW_t^i.$$

The proof now follows similar steps to that of lemma A.4. By Itô's lemma we get the dynamics of $K_p(\mathbf{y}_t)$ and, by taking the expectation, that of $g_p(t)$, which we then bound by Gronwall's lemma. More precisely, Itô's lemma (A.2) yields

$$\begin{aligned}
dK_p(\mathbf{y}_t) &= -\frac{2p}{B} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^i, \text{Cov}_{y_t}(\mathbf{y}_t^i - \mathbf{y}^*) \rangle dt \\
&\quad + \frac{2p}{B} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^i, \sqrt{B}\sqrt{2\text{Cov}_{y_t}}dW_t^i \rangle \\
&\quad + \frac{4p(p-1)}{B} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-2} \text{Tr}\{(\mathbf{y}_t^i \otimes \mathbf{y}_t^i)\text{Cov}_{y_t}\} dt \tag{A.21} \\
&\quad + \frac{2p}{B} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \text{Tr}\{\text{Cov}_{y_t}\}dt \\
&\quad + \frac{2p(D+1)}{B^2} \sum_{i=1}^B \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^i, \mathbf{y}_t^i - \bar{\mathbf{y}}_t \rangle dt,
\end{aligned}$$

with the last term corresponding to the ALDI correction term. Using (A.19) and (A.20), the expectation

of this term term can be bounded according to

$$\begin{aligned}
& \frac{2p(D+1)}{B^2} \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^i, \mathbf{y}_t^i - \bar{\mathbf{y}}_t \rangle \right] \\
& \leq \frac{2p(D+1)}{B} g_p(t) + \frac{2p(D+1)}{B^3} \sum_{i,k=1}^B \mathbb{E} \left[\frac{\langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^p + \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^k, \mathbf{y}_t^k \rangle}{2} \right] \\
& \leq \frac{2p(D+1)}{B} g_p(t) + \frac{p(D+1)}{B} g_p(t) + \frac{p(D+1)}{B} g_p(t) \\
& = C(p, D) g_p(t).
\end{aligned}$$

The second term in (A.21) vanishes in expectation. The first, third and fourth terms can be bounded in the same way as in the proof of Proposition 4.1 in [9]. The computations are technical but not very insightful and since the interested reader may follow the arguments in that work step by step, we will not repeat them here in detail, but simply state the result. For the first term, we get with Hölder's inequality and lemma A.4

$$\begin{aligned}
& - \frac{2p}{B} \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \langle \mathbf{y}_t^i, \text{Cov}_{\mathbf{y}_t}(\mathbf{y}_t^i - \mathbf{y}^*) \rangle \right] \\
& \leq 2p |\mathbf{y}^*| \left(\frac{1}{B} \sum_{i=1}^B \mathbb{E} \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^p \right)^{(p-1/2)/p} \left(\frac{1}{B} \sum_{i=1}^B \mathbb{E} \langle \mathbf{e}_t^k, \mathbf{e}_t^k \rangle^{2p} \right)^{1/(2p)} \quad (\text{A.22}) \\
& \leq 2p |\mathbf{y}^*| g_p^{(p-1/2)/p}(t) C e^{C(p,D)t}.
\end{aligned}$$

For the third and fourth terms we receive

$$\begin{aligned}
& \frac{4p(p-1)}{B} \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-2} \text{Tr} \{ (\mathbf{y}_t^i \otimes \mathbf{y}_t^i) \text{Cov}_{\mathbf{y}_t} \} \right] \\
& \leq 4p(p-1) \left(\frac{1}{B} \sum_{i=1}^B \mathbb{E} \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^p \right)^{(p-1)/p} \left(\frac{1}{B} \sum_{k=1}^B \mathbb{E} \langle \mathbf{e}_t^k, \mathbf{e}_t^k \rangle^p \right)^{1/p} \quad (\text{A.23}) \\
& \leq 4p(p-1) g_p^{(p-1)/p}(t) C e^{C(p,D)t}
\end{aligned}$$

and

$$\begin{aligned}
& \frac{2p}{B} \sum_{i=1}^B \mathbb{E} \left[\langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^{p-1} \text{Tr} \{ \text{Cov}_{\mathbf{y}_t} \} \right] \\
& \leq 2p \left(\frac{1}{B} \sum_{i=1}^B \mathbb{E} \langle \mathbf{y}_t^i, \mathbf{y}_t^i \rangle^p \right)^{(p-1)/p} \left(\frac{1}{B} \sum_{k=1}^B \mathbb{E} \langle \mathbf{e}_t^k, \mathbf{e}_t^k \rangle^p \right)^{1/p} \quad (\text{A.24}) \\
& \leq 2p g_p^{(p-1)/p}(t) C e^{C(p,D)t}.
\end{aligned}$$

In total, we arrive at

$$g_p'(t) \leq C g_p(t) + C e^{Ct} \left(g_p^{(p-1)/p}(t) + g_p^{(p-1/2)/p}(t) \right).$$

The result now follows from lemma B.3. □

Proof of Lemma A.1. The lemma results directly from boundedness of the higher moments of \mathbf{y}_t^i (Lemma A.5) and z_t^i (see Proposition 5.3 in [9]). □

Proof of Lemma A.2. First, note that

$$\mathbb{E}|p_t^i|^2 = \mathbb{E}|x_t^i - \bar{x}_t|^2 = \mathbb{E}(|x_t^i|^2 + |\bar{x}_t|^2 - 2\langle x_t^i, \bar{x}_t \rangle) = \mathbb{E}(|x_t^i|^2 - |\bar{x}_t|^2) = \mathbb{E}\left(\frac{1}{B} \sum_{i=1}^B |x_t^i|^2 - |\bar{x}_t|^2\right) \quad (\text{A.25})$$

due to the symmetry between the particles. Using (A.1) and (A.5), we have

$$\begin{aligned} dx_t^i &= \left(-\text{Cov}_{x_t+z_t} P(x_t^i + z_t^i) + \text{Cov}_{\pi(t)} Pz_t^i\right) dt + (\text{Cov}_{x_t+z_t} - \text{Cov}_{\pi(t)}) Py^* dt \\ &\quad + \frac{D+1}{B} (y_t^i - \bar{y}_t) dt + \left(\sqrt{2\text{Cov}_{x_t+z_t}} - \sqrt{2\text{Cov}_{\pi(t)}}\right) dW_t^i. \end{aligned} \quad (\text{A.26})$$

Now, applying Itô's formula (A.2) and replacing $\text{Cov}_{\pi(t)}$ in the second and third terms with Cov_{z_t} , we find

$$\begin{aligned} d|x_t^i|^2 &= -2\langle x_t^i, \text{Cov}_{x_t+z_t} P x_t^i \rangle dt - 2\langle x_t^i, (\text{Cov}_{x_t+z_t} - \text{Cov}_{z_t}) P z_t^i \rangle dt \\ &\quad + 2\langle x_t^i, (\text{Cov}_{x_t+z_t} - \text{Cov}_{z_t} P y^*) \rangle dt + \text{Tr} \left[\left(\sqrt{\text{Cov}_{x_t+z_t}} - \sqrt{\text{Cov}_{\pi(t)}} \right)^2 \right] dt \\ &\quad + 2\left\langle x_t^i, \frac{D+1}{B} (y_t^i - \bar{y}_t) \right\rangle dt + 2\left\langle x_t^i, \left(\sqrt{2\text{Cov}_{\pi(t)}} - \sqrt{2\text{Cov}_{\pi(t)}} \right) dW_t^i \right\rangle + R_t^i dt \end{aligned} \quad (\text{A.27})$$

and

$$\begin{aligned} d|\bar{x}_t|^2 &= -2\langle \bar{x}_t, \text{Cov}_{x_t+z_t} P \bar{x}_t \rangle dt - 2\langle \bar{x}_t, (\text{Cov}_{x_t+z_t} - \text{Cov}_{z_t}) P \bar{z}_t \rangle dt \\ &\quad + 2\langle \bar{x}_t, (\text{Cov}_{x_t+z_t} - \text{Cov}_{z_t} P y^*) \rangle dt + \text{Tr} \left[\left(\sqrt{\text{Cov}_{x_t+z_t}} - \sqrt{\text{Cov}_{\pi(t)}} \right)^2 \right] dt \\ &\quad + 2\sum_{i=1}^B \left\langle \bar{x}_t, \frac{D+1}{B^2} (y_t^i - \bar{y}_t) \right\rangle dt + 2\langle \bar{x}_t, \left(\sqrt{2\text{Cov}_{\pi(t)}} - \sqrt{2\text{Cov}_{\pi(t)}} \right) d\bar{W}_t \rangle + \bar{R}_t dt, \end{aligned} \quad (\text{A.28})$$

with

$$\begin{aligned} R_t^i &= 2\langle x_t^i, (\text{Cov}_{\pi(t)} - \text{Cov}_{z_t}) P z_t^i \rangle - 2\langle x_t^i, (\text{Cov}_{\pi(t)} - \text{Cov}_{z_t}) P y^* \rangle, \\ \bar{R}_t &= 2\langle \bar{x}_t, (\text{Cov}_{\pi(t)} - \text{Cov}_{z_t}) P \bar{z}_t \rangle - 2\langle \bar{x}_t, (\text{Cov}_{\pi(t)} - \text{Cov}_{z_t}) P y^* \rangle. \end{aligned} \quad (\text{A.29})$$

Combining (A.28) and (A.27), we finally arrive at

$$\begin{aligned} d\left(\frac{1}{B} \sum_{i=1}^B |x_t^i|^2 - |\bar{x}_t|^2\right) &= -\frac{2}{B} \sum_{i=1}^B \langle p_t^i, \text{Cov}_{p_t+q_t} P p_t^i \rangle dt - \frac{2}{B} \sum_{i=1}^B \langle p_t^i, (\text{Cov}_{p_t+q_t} - \text{Cov}_{q_t}) P q_t^i \rangle dt \\ &\quad + 2\left(1 - \frac{1}{B}\right) \text{Tr} \left[\left(\sqrt{\text{Cov}_{x_t+z_t}} - \sqrt{\text{Cov}_{\pi(t)}} \right)^2 \right] dt + \left(\frac{1}{B} \sum_{i=1}^B R_t^i - \bar{R}_t \right) dt \\ &\quad + \frac{D+1}{B^2} \sum_{i=1}^B \left[2\langle x_t^i, (y_t^i - \bar{y}_t) \rangle - 2\left\langle \bar{x}_t, \frac{1}{B} \sum_{i=1}^B (y_t^i - \bar{y}_t) \right\rangle \right] dt \\ &\quad + \frac{2}{B} \sum_{i=1}^B \left\langle (x_t^i - \bar{x}_t), \left(\sqrt{2\text{Cov}_{x_t+z_t}} - \sqrt{2\text{Cov}_{\pi(t)}} \right) d(W_t^i - \bar{W}_t) \right\rangle. \end{aligned} \quad (\text{A.30})$$

The expectation of the last term vanishes by the properties of Brownian motion. The first four terms have the same form as in equation (5.22) in [9] and can be treated in a similar way. Note that the terms are not identical to the terms in that work, since the underlying process y_t and hence x_t, p_t, q_t are different. They can however be treated in a completely analogous way, tracing the arguments one by one. Hence, we will only state the resulting bounds here, and refer to equations (5.22)-(5.27) in [9] for details. The first two terms can be bounded in expectation according to

$$\begin{aligned} &-\frac{2}{B} \mathbb{E} \left[\sum_{i=1}^B \langle p_t^i, \text{Cov}_{p_t+q_t} P p_t^i \rangle dt + \sum_{i=1}^B \langle p_t^i, (\text{Cov}_{p_t+q_t} - \text{Cov}_{q_t}) P q_t^i \rangle dt \right] \\ &\leq -\lambda_{\min}(P) \mathbb{E} \left(\|\text{Cov}_{x_t+z_t} - \text{Cov}_{z_t}\|_F^2 \right) + 3\|P\|_F \text{Var}(\pi(t)) \mathbb{E}|p_t^1|^2 + CB^{-1/2-\alpha(1-\epsilon)}. \end{aligned} \quad (\text{A.31})$$

For the third term, we get

$$\mathbb{E} \left[\text{Tr} \left[\left(\sqrt{\text{Cov}_{x_t+z_t}} - \sqrt{\text{Cov}_{\pi(t)}} \right)^2 \right] \right] \leq \lambda_0(t)^{-1} \mathbb{E} \|\text{Cov}_{p_t+q_t} - \text{Cov}_{q_t}\|_F^2 + C_\epsilon B^{-1/2-\alpha/2+\epsilon\alpha/4} \quad (\text{A.32})$$

with a constant C_ϵ depending only on ϵ and $\lambda_0(t) = (\lambda_{\min}(\text{Cov}_{x_t+z_t})^{1/2} + \lambda_{\min}(\text{Cov}_{\pi(t)})^{1/2})^2$ coming from the Ando-Hemmen inequality (see Theorem 6.2 in [22]).

Finally, for the fourth term, we get

$$\mathbb{E} \left[\frac{1}{B} \sum_{i=1}^B R_t^i - \bar{R}_t \right] \leq C_\epsilon B^{-1/2-\alpha/2+\epsilon\alpha/4}, \quad (\text{A.33})$$

again with a constant C_ϵ depending on ϵ . The fifth term comes from the ALDI correction term. To bound it, we note that

$$\begin{aligned} \langle x_t^i, (y_t^i - \bar{y}_t) \rangle &\leq \mathbb{E}|x_t^i| (\mathbb{E}|y_t^i| + \mathbb{E}|\bar{y}_t|) \\ &\leq 2(\mathbb{E}|x_t^i|^2)^{1/2} \mathbb{E}(|y_t^i|^2)^{1/2} \\ &\leq C_1 \end{aligned}$$

with a constant C_1 independent of B, t , where we have used Hölder's inequality in the second inequality and Lemmas A.5 and A.1 in the last inequality. Similarly, we get

$$\left\langle \bar{x}_t, \frac{1}{B} \sum_{i=1}^B (y_t^j - \bar{y}_t) \right\rangle \leq \frac{2}{B} \mathbb{E}|x_t^i| \sum_{i=1}^B \mathbb{E}|y_t^i| = 2\mathbb{E}|x_t^i| \mathbb{E}|y_t^i| \leq C_1,$$

to arrive at

$$\mathbb{E} \left[\frac{D+1}{B^2} \sum_{i=1}^B \left[2 \langle x_t^i, (y_t^i - \bar{y}_t) \rangle - 2 \left\langle \bar{x}_t, \frac{1}{B} \sum_{i=1}^B (y_t^i - \bar{y}_t) \right\rangle \right] \right] \leq \frac{2(D+1)}{B^2} B C_1 = 2(D+1)C_1 B^{-1}.$$

Hence, taking the expectation in (A.30), we find

$$\begin{aligned} \frac{d\mathbb{E}|p_t^1|^2}{dt} &\leq 3\|P\|_2 \text{Var}(\pi(t)) \mathbb{E}|p_t^1|^2 - (\lambda_{\min}(B) - 2\lambda_0^{-1}) \mathbb{E} \|\text{Cov}_{p_t+q_t} - \text{Cov}_{q_t}\|_F^2 \\ &\quad + C B^{-1/2-\alpha(1-\epsilon)} + C_\epsilon B^{-1/2-\alpha/2+\epsilon\alpha/4} + 2(D+1)C_1 B^{-1}. \end{aligned} \quad (\text{A.34})$$

The second term is always negative by assumption (4.7). Furthermore we note that since $0 \leq \alpha < 1$ and $0 < \epsilon < 1/2$ all exponents of B appearing in (A.34) are smaller than $-1/2 - \alpha/2 + \epsilon$, leading to

$$\frac{d\mathbb{E}|p_t^1|^2}{dt} \leq 3\|P\|_2 \text{Var}(\pi(t)) \mathbb{E}|p_t^1|^2 + C B^{-1/2-\alpha/2+\epsilon}. \quad (\text{A.35})$$

Now, by equation (2.2) in [6] we have

$$\text{Cov}_{\pi(t)} = (1 - e^{-2t})P^{-1} + e^{-2t}\text{Cov}_{\pi_0}^{-1} \quad (\text{A.36})$$

and hence $\text{Var}(\pi(t)) \leq M$ for some $M > 0$. Integrating (A.35) leads to

$$\mathbb{E}|p_t^1|^2 \leq C \int_0^t \mathbb{E}|p_s^1|^2 ds + Ct B^{-1/2-\alpha/2+\epsilon} \quad (\text{A.37})$$

and applying lemma B.1 yields

$$\mathbb{E}|p_t^1|^2 \leq Ct B^{-1/2-\alpha/2+\epsilon} + \int_0^t C_s B^{-1/2-\alpha/2+\epsilon} e^{C(t-s)} ds = C(t) B^{-1/2-\alpha/2+\epsilon}, \quad (\text{A.38})$$

finishing the proof. □

Proof of Lemma A.3. The proof is the same as that of Lemma 5.5 in [9]. □

B Gronwall inequalities

Lemma B.1 (Gronwall [19]). *Let $u(t)$ be a nonnegative function satisfying*

$$u(t) \leq \beta(t) + \int_0^t c(s)u(s)ds \quad \text{for all } t \geq 0,$$

where β and c are continuous nonnegative functions for $t \geq 0$. Then

$$u(t) \leq \beta(t) + \int_0^t \beta(s)c(s)e^{\int_s^t c(r)dr}ds \quad \text{for all } t \geq 0.$$

Proof. See e.g. Theorem 1 in [10]. □

Lemma B.2 (Perov, [31] or Theorem 21 in [10]). *Let $u(t)$ be a nonnegative function satisfying*

$$u'(t) \leq c(t)u(t) + a(t)u(t)^\alpha,$$

where $0 \leq \alpha < 1$ and c and a are continuous nonnegative functions for $t > 0$. Then we have

$$u(t) \leq \left[u^{1-\alpha}(0)e^{(1-\alpha)\int_0^t c(s)ds} + e^{(1-\alpha)\int_0^t c(s)ds} \int_0^t (1-\alpha)a(s)e^{(\alpha-1)\int_0^s a(r)dr}ds \right]^{\frac{1}{1-\alpha}}.$$

Proof. Let $v(t)$ solve

$$v'(t) = c(t)v(t) + a(t)v(t)^\alpha, \quad v(0) = u(0).$$

This is a classic Bernoulli-type ODE, which can be solved by noting that the ODE becomes linear in $v^{1-\alpha}$:

$$(v^{1-\alpha})'(t) = (1-\alpha)c(t)v^{1-\alpha}(t) + (1-\alpha)a(t).$$

Applying the variation of parameters to $v^{1-\alpha}$ we get

$$v(t) = \left[u^{1-\alpha}(0)e^{(1-\alpha)\int_0^t c(s)ds} + \int_0^t (1-\alpha)a(s)e^{(1-\alpha)\int_0^s c(r)dr}ds \right]^{\frac{1}{1-\alpha}}.$$

The assertion now immediately follows, since $u(t) \leq v(t)$ by construction. □

Lemma B.3. *Let $u(t)$ be a nonnegative function satisfying*

$$u'(t) \leq c(t)u(t) + a(t)u^\alpha(t) + b(t)u^\beta(t),$$

where $0 < \alpha < \beta \leq 1$ and c, a and b are continuous and nonnegative functions for $t \geq 0$. Then $u(t)$ is bounded for all $t \geq 0$ by

$$u(t) \leq \left[u^{1-\alpha}(0)e^{(1-\alpha)\int_0^t c(s)+b(s)ds} + \int_0^t (1-\alpha)(a(s) + b(s))e^{(1-\alpha)\int_0^s c(r)+b(r)dr}ds \right]^{\frac{1}{1-\alpha}}. \quad (\text{B.1})$$

Proof. For $u(t) < 1$ we have $u^\beta(t) \leq u^\alpha(t)$ and hence

$$\begin{aligned} u'(t) &\leq c(t)u(t) + (a(t) + b(t))u^\alpha(t) \\ &\leq (c(t) + b(t))u(t) + (a(t) + b(t))u^\alpha(t). \end{aligned}$$

For $u(t) \geq 1$, we have $u^\beta(t) \leq u(t)$ and hence

$$\begin{aligned} u'(t) &\leq (c(t) + b(t))u(t) + a(t)u^\alpha(t) \\ &\leq (c(t) + b(t))u(t) + (a(t) + b(t))u^\alpha(t). \end{aligned}$$

In total, we arrive at $u'(t) \leq (c(t) + b(t))u(t) + (a(t) + b(t))u^\alpha(t)$ for all $t \geq 0$ and can use lemma B.2. □

C Existence and Uniqueness of a solution of ALDI

Existence and uniqueness of a strong solution by means of construction of a Lyapunov function of the system has already been proven in the original work on ALDI [16]. We provide an alternative proof, using the same Lyapunov function that is used in [9] for existence and uniqueness of the EKS solution, merely for the interest of the reader.

Theorem C.1 (Compare Thm 3.2 in [9]). *Suppose G is linear and $(y_0^i)_{i=1}^B$ are i.i.d. Then for all $t \geq 0$ there exists a unique strong solution $(y_t^i)_{i=1}^B$ (up to \mathbb{P} -indistinguishability) of the set of coupled SDEs defined by (2.25).*

Proof. Proof of Theorem C.1 We consider the stacked SDE

$$dY_t = F(Y_t)dt + G(Y_t)dW_t,$$

where $Y_t = (y_t^i)_{i=1}^B \in \mathbb{R}^{DB}$, $W_t = (W_t^i)_{i=1}^B$ and

$$F(Y_t) = \left(-\text{Cov}_{y_t} P(y_t^i - u^*) + \frac{D+1}{B} (y_t^j - \bar{y}_t) \right)_{i=1}^B \in \mathbb{R}^{DB},$$

$$G(Y_t) = \text{diag} \left(\sqrt{2\text{Cov}_{y_t}} \right)_{i=1}^B \in \mathbb{R}^{DB \times DB},$$

where $\text{diag}(D_i)_{i=1}^B$ is a block diagonal matrix with entries $(D_i)_i^B$ on the diagonal.

We prove the assertion by showing existence of a Lyapunov function of the system, i.e. a function $V \in \mathcal{C}^2(\mathbb{R}^{DB}, \mathbb{R})$ such that

(1) there exists a $c > 0$ so that

$$\mathcal{L}V(Y) := \nabla V(Y) \cdot F(Y) + \frac{1}{2} \text{tr} [G(Y)^\top \text{Hess}_V(Y) G(Y)] \leq cV(Y)$$

for all $Y \in \mathbb{R}^{DB}$.

(2) we have

$$\inf_{|Y| > R} V(Y) \longrightarrow \infty$$

as $R \longrightarrow \infty$.

The function we choose is the same that is used in [9] for the EKS:

$$V(Y) = V_1(Y) + V_2(Y) = \frac{1}{B} \sum_{i=1}^B |y^j - \bar{y}|^2 + |\bar{y} - y^*|_P^2.$$

Towards (1), note that

$$\nabla V_1(Y) = \left(\frac{2}{B} (y^j - \bar{y}) \right)_{i=1}^B,$$

$$\nabla V_2(Y) = (2P(\bar{y} - y^*))_{i=1}^B$$

and that the diagonal blocks of $\text{Hess}_{V_2}(Y)$ and $\text{Hess}_{V_1}(Y)$ are given by

$$\frac{2}{B} \text{diag}(P)_{i=1}^B$$

and

$$\frac{2}{B} \left(1 - \frac{1}{B}\right) \text{diag}(I_d)_{i=1}^B$$

respectively.

This yields

$$\begin{aligned} \nabla V_1(Y) \cdot F(Y) &= -\frac{2}{B} \sum_{i=1}^B \langle y^i - \bar{y}, \text{Cov}_y B(y^i - y^*) \rangle + \frac{2(D+1)}{B^2} \sum_{i=1}^B \langle y^i - \bar{y}, y^i - \bar{y} \rangle \\ &= -\frac{2}{B} \sum_{i=1}^B \langle y^i - \bar{y}, \text{Cov}_y B(y^i - \bar{y}) \rangle + \frac{2(D+1)}{B} V_1(Y) \leq \frac{2(D+1)}{B} V_1(Y), \end{aligned}$$

where we have used that both Cov_y and B are positive semi-definite, as well as

$$\begin{aligned} \nabla V_2(Y) \cdot F(Y) &= -2 \sum_{i=1}^B \langle P(\bar{y} - y^*), \text{Cov}_y P(y^i - y^*) \rangle + \frac{2(D+1)}{B} \sum_{i=1}^B \langle P(\bar{y} - y^*), y^i - \bar{y} \rangle, \\ &= -2B \langle P(\bar{y} - y^*), \text{Cov}_y P(\bar{y} - y^*) \rangle + 2(D+1) \langle P(\bar{y} - y^*), \bar{y} - \bar{y} \rangle \leq 0. \end{aligned}$$

Furthermore, we have

$$\begin{aligned} \frac{1}{2} \text{Tr}[G(Y)^\top \text{Hess}_{V_1}(Y) G(Y)] &= \frac{1}{2} \text{Tr}[\text{Hess}_{V_1}(Y) G(Y) G(Y)^\top] = \text{Tr}[\text{Hess}_{V_1}(Y) \text{diag}(\text{Cov}_y)_{i=1}^B] \\ &= \frac{2}{B} \sum_{i=1}^B \left(1 - \frac{1}{B}\right) \text{Tr}[\text{Cov}_y] = 2 \left(1 - \frac{1}{B}\right) \text{Tr}[\text{Cov}_y] \\ &= 2 \left(1 - \frac{1}{B}\right) \frac{1}{B} \sum_{i=1}^B (y^i - \bar{y})^\top (y^i - \bar{y}) = 2 \left(1 - \frac{1}{B}\right) V_1(Y) \end{aligned}$$

and

$$\begin{aligned} \frac{1}{2} \text{Tr}[G(Y)^\top \text{Hess}_{V_2}(Y) G(Y)] &= \text{Tr}[\text{Hess}_{V_2}(Y) \text{diag}(\text{Cov}_y)_{i=1}^B] = 2 \text{Tr}[P \text{Cov}_y] \\ &= \frac{2}{B} \sum_{i=1}^B (y^i - \bar{y})^\top P (y^i - \bar{y}) \leq 2 \|P\| V_1(Y). \end{aligned}$$

In total, we get

$$\mathcal{L}V(Y) \leq \left(\frac{2(D+1)}{B} + 2 \left(1 - \frac{1}{B}\right) + 2\|P\| \right) V(Y)$$

for all Y , proving (1). Now, towards (2) assume there is a sequence $(Y_n)_{n \in \mathbb{N}}$ in \mathbb{R}^{DB} with $V(Y_n) < M$ for some $M > 0$ and all n . By the construction of V , we have

$$|y_n^i - \bar{y}_n| < \sqrt{MB}, \quad |\bar{y}_n - y^*| < \sqrt{M}$$

and hence

$$|Y_n|^2 = \sum_{i=1}^B |y_n^i|^2 = \sum_{i=1}^B |y_n^i - \bar{y}_n + \bar{y}_n - y^* + y^*|^2 < \sum_{i=1}^B \left(\sqrt{M}(\sqrt{B} + 1) + |y^*| \right)^2$$

for all n , meaning the sequence $(|Y_n|)_{n \in \mathbb{N}}$ is also bounded. \square

D Wasserstein spaces and convergence

For $p \geq 1$ let $\mathcal{D}_p(X)$ denote the set of probability measures over a Banach space X with norm $\|\cdot\|$ that have finite p -th moment. On $\mathcal{D}_p(X)$ we define the p -Wasserstein metric as

$$\mathcal{W}_p(\mu, \nu)^p = \inf_{\pi \in \Pi_p(\mu, \nu)} \int_{X \times X} \|x - y\|^p d\pi(x, y),$$

with $\Pi_p(\mu, \nu) = \{\pi \in \mathcal{D}_p(X \times X) \mid \pi(\cdot \times X) = \mu, \pi(X \times \cdot) = \nu\}$. Furthermore for $X = \mathbb{R}^D$, we denote with $F_p^k(X)$ with $k \in \mathbb{N} \cup \{\infty\}$ the set of regular measures with p -th moment and a Lebesgue density in $\mathcal{C}^k(X)$, in particular

$$F_p^k = \{\mu \in \mathcal{D}_p(\mathbb{R}^D) \mid \exists f := d\mu/dx \text{ and } f \in \mathcal{C}^k(\mathbb{R}^D)\}.$$

Lemma D.1. *Let $p \geq 1$. Then, for any $k \in \mathbb{N} \cup \{\infty\}$ the subset $F_p^k \subset \mathcal{D}_p$ is dense in \mathcal{D}_p with respect to the \mathcal{W}_p metric.*

Proof. Consider a non-negative convolution kernel $\rho \in \mathcal{C}_c^\infty(\mathbb{R}^D)$ with $\|\rho\|_1 = 1$, where $\|\cdot\|_1$ denotes the L^1 -norm over \mathbb{R}^D . Then define $\rho_\epsilon(x) = \epsilon^{-D} \rho(\epsilon^{-1}x)$ and

$$f_\epsilon(x) = (\rho_\epsilon * \mu)(x) = \int_{\mathbb{R}^D} \rho_\epsilon(x - y) d\mu(y).$$

Then $f_\epsilon \in \mathcal{C}^\infty(\mathbb{R}^d)$ is a Lebesgue density defining a measure μ_ϵ on \mathbb{R}^D . Now consider the transport plan or coupling

$$\pi_\epsilon(dx, dy) = \rho_\epsilon(x - y) d\mu(y) dx.$$

Then $\pi_\epsilon \in \mathcal{W}_2(\mathbb{R}^D \times \mathbb{R}^D)$ with marginal projections $\mu_\epsilon(\cdot \times \mathbb{R}^D) = \mu_\epsilon$ and $\pi_\epsilon(\mathbb{R}^D \times \cdot) = \mu$. Then it holds

$$\begin{aligned} \mathcal{W}_p(\mu, \mu_\epsilon)^p &\leq \int_{\mathbb{R}^D \times \mathbb{R}^D} |x - y|^p d\pi_\epsilon \\ &= \int_{\mathbb{R}^D} \int_{\mathbb{R}^D} |x - y|^p \rho_\epsilon(x - y) d\mu(y) dx \\ &= \int_{\mathbb{R}^D} \int_{\mathbb{R}^D} \epsilon^p |z|^p \rho(z) d\mu(y) dz \\ &= \epsilon^p \int_{\mathbb{R}^D} |z|^p \rho(z) dz, \end{aligned}$$

where we used the substitution $z = \epsilon^{-1}(x - y)$. Now, clearly, $\mathcal{W}_p(\mu, \mu_\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. \square

Theorem D.1 (Theorem 1 in [14]). *Let μ be a probability measure on \mathbb{R}^D and let $\mu^{(B)} = \frac{1}{B} \sum_{i=1}^B \delta_{y^i}$ be an empirical measure of i.i.d. samples $\{y^i\}_{i=1}^B$ drawn from μ . Let $p > 0$ and assume that*

$$M_q(\mu) := \int_{\mathbb{R}^D} |y|^q d\mu(y) < \infty \quad \text{for some } q > p. \quad (\text{D.1})$$

Then, there exists a constant C depending only on p, d, q such that, for all $N \geq 1$,

$$\mathbb{E}[(\mathcal{W}_p(\hat{\mu}^{(B)}, \mu))] \leq CM_q^{p/q}(\mu) \cdot \begin{cases} B^{-1/2} + B^{-(q-p)/q} & \text{if } p > d/2 \text{ and } q = 2p, \\ B^{-1/2} \log(1 + B) + B^{-(q-p)/q} & \text{if } p = d/2 \text{ and } q = 2p, \\ B^{-p/d} + B^{-(q-p)/q} & \text{if } p \in (0, d/2) \text{ and } q = d/(d-p). \end{cases} \quad (\text{D.2})$$