



Affine-Invariant Ensemble Transform Methods for Logistic Regression

Jakiw Pidstrigach¹ · Sebastian Reich¹

Received: 16 May 2021 / Accepted: 13 September 2021 / Published online: 21 January 2022
© The Author(s) 2022

Abstract

We investigate the application of ensemble transform approaches to Bayesian inference of logistic regression problems. Our approach relies on appropriate extensions of the popular ensemble Kalman filter and the feedback particle filter to the cross entropy loss function and is based on a well-established homotopy approach to Bayesian inference. The arising finite particle evolution equations as well as their mean-field limits are affine-invariant. Furthermore, the proposed methods can be implemented in a gradient-free manner in case of nonlinear logistic regression and the data can be randomly subsampled similar to mini-batching of stochastic gradient descent. We also propose a closely related SDE-based sampling method which again is affine-invariant and can easily be made gradient-free. Numerical examples demonstrate the appropriateness of the proposed methodologies.

Keywords Logistic regression · Bayesian inference · Interacting particle systems · Affine invariance · Ensemble Kalman filter · Langevin dynamics

Mathematics Subject Classification 62J02 · 65C05 · 62F15

Communicated by Teresa Krick and Hans Munthe-Kaas.

Invited paper based on the FoCM 2021 Online Seminar lecture *Statistical inverse problems and affine-invariant gradient flow structures in the space of probability measures* presented by Sebastian Reich in June 2021.

SR has been partially funded by Deutsche Forschungsgemeinschaft (DFG) - Project-ID 318763901 - SFB1294..

✉ Sebastian Reich
sebastian.reich@uni-potsdam.de

Jakiw Pidstrigach
pidstrigach@uni-potsdam.de

¹ Institut für Mathematik, Universität Potsdam, Karl-Liebknecht-Str. 24/25, 14476 Potsdam, Germany

1 Introduction

Statistical inference for logistic regression and classification problems is a well studied problem from both parameter optimisation and Bayesian inference perspectives [5]. While many classical optimisation and Markov chain Monte Carlo (MCMC) methods are applicable, an efficient inference in the context of semi-parametric models still poses computational challenges. With this paper, we address this challenge by investigating the extension of coupling-of-measures and ensemble transform methodologies [29] from the squared error loss function to the cross entropy loss function typically used for logistic regression and classification.

More specifically, previous work on ensemble methods for Bayesian inference, such as the ensemble Kalman filter (EnKF) [8] and the feedback particle filter (FPF) [35,38,39], have almost exclusively focused on the squared error loss function of the form

$$\Psi_{\text{data}}(\theta) = \frac{1}{2}(g(\theta) - t)^T \Gamma^{-1}(g(\theta) - t), \quad (1)$$

with $g : \mathbb{R}^D \rightarrow \mathbb{R}^N$ a forward map, $t \in \mathbb{R}^N$ the data, $\Gamma \in \mathbb{R}^{N \times N}$ the measurement error covariance matrix, and $\theta \in \mathbb{R}^D$ the parameters to be estimated. Notable exceptions include the application of ensemble Kalman inversion (EKI) [17] and the modified EnKF formulation of [13] to the training of neural networks with a cross entropy loss function. While these methods seek to minimise a regularised cross entropy loss function, the more recent work [14] has also investigated ensemble-based sampling methods for Bayesian inference in the context of logistic regression and classification. This work relies on the time-stepping of appropriate stochastic differential equations (SDEs) and is in line with several other ensemble-based sampling methods such as the ensemble preconditioned MCMC methods of [19], the ensemble Kalman sampler (EKS) [9], and affine invariant Langevin dynamics (ALDI) [10]. Contrary to such invariance-of-measures based SDE methodologies, the approach taken in this paper is instead founded on the homotopy approach to Bayesian inference, as first formulated in a time-continuous framework in [6,26], and which is close in spirit to the iterative application of the EnKF [31] and parameter estimation methods based on sequential Monte Carlo (SMC) methods [4].

In addition to expanding the homotopy-based approach to logistic regression, we address two further important concepts, namely affine-invariance of the proposed methods [10,11,19] and sub-sampling of data points, as widely used in stochastic gradient descent [20]. Both concepts are used to improve the computational efficiency of optimisation and sampling methods. Take, for example, a two dimensional Gaussian random variable with mean zero and covariance matrix

$$\Sigma = \begin{pmatrix} \epsilon^2 & 0 \\ 0 & 1 \end{pmatrix}.$$

A random walk Metropolis–Hastings algorithm will sample inefficiently whenever ϵ^2 is vastly different from one. An affine-invariant modification, on the other hand, will sample this problem as efficiently as if ϵ were set to $\epsilon = 1$ [11]. Sub-sampling

replaces the exact gradient of a cost functional by a cheaper to evaluate stochastic approximation, which agrees with the exact gradient in expectation.

We also discuss the possibility for derivative-free implementations [8–10,17], localisation [8,29] via dropouts [32], and efficient linearly implicit time-stepping methods [2].

Our main contributions with regard to Bayesian homotopy methods for logistic regression are:

- an extension of affine-invariant ensemble transform approaches, such as the EnKF, to logistic regression,
- an affine-invariant generalisation of the FPF and its application to logistic regression,
- an extension of data sub-sampling (mini-batches) to Bayesian homotopy methods.

In a further step, we combine these homotopy approaches with SDE-based sampling methods in order to derive

- a derivative-free and affine-invariant SDE-based sampling methods for logistic regression.

We demonstrate the appropriateness of the proposed methods by means of a set of numerical experiments. Extensions to nonlinear and multi-class logistic regression [5] are also discussed. We also briefly discuss the application to sigmoidal Cox processes [1].

The layout of this paper is as follows. The required mathematical background material on both logistic regression is collected in Sect. 2. The homotopy approach to Bayesian inference is summarised in Sect. 3. There we also present an affine-invariant formulation of the homotopy approach and discuss and analyse data sub-sampling in the spirit of [20]. Section 4 develops three different algorithmic approaches for the implementation of homotopy-based Bayesian inference for logistic regression. More specifically, we propose an affine-invariant modification of the FPF and two extensions of the EnKF to logistic regression. We also discuss robust numerical implementations combining dropouts [32] with localisation [8,29] and linearly implicit time-stepping methods [2]. Section 5 combines SDE-based sampling methods with a homotopy-based drift term in order to derive a gradient-free and affine-invariant algorithm for Bayesian inference. While the proposed method is only exact for Gaussian measures, it can be used for faster equilibration and approximate inference provided the posterior distribution is close to Gaussian. Numerical results are presented and evaluated in Sect. 6 and provide a proof-of-concept while a more detailed exploration is left to follow-up work. Several possible extensions of the proposed methods, namely multi-class classification, nonlinear logistic regression, and sigmoidal Cox processes, are discussed in Sect. 7 followed by a summary statement in Sect. 8.

2 Mathematical Problem Formulation

This paper is motivated by the classical logistic regression problem arising from classification into $L > 1$ classes, denoted by C_l , over an input space $x \in \mathbb{R}^J$ [5]. In

order to simplify the exposition, we start with the binary classification problem, that is, $L = 2$ and focus on the linear case. The extension to nonlinear and multi-class logistic regression is discussed in Sect. 7.

The posterior distribution for class C_1 is defined as a logistic sigmoid

$$\sigma(a) = \frac{1}{1 + \exp(-a)}$$

acting on a linear combination of x -dependent and vector-valued features $\phi_x \in \mathbb{R}^D$ so that

$$\pi(C_1|\phi_x) = \sigma(\theta^T \phi_x). \quad (2)$$

The model has D adjustable parameters $\theta \in \mathbb{R}^D$. The probability of the complimentary class C_2 is given by $\pi(C_2|\phi_x) = 1 - \pi(C_1|\phi_x)$. Given a data set $\{(x_n, t_n)\}_{n=1}^N$ with labels $t_n \in \{0, 1\}$, the negative log-likelihood function is given by the cross-entropy function

$$\Psi_{\text{data}}(\theta) = - \sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\}, \quad (3)$$

where $y_n := y_{x_n}(\theta)$ with

$$y_x(\theta) := \sigma(\theta^T \phi_x). \quad (4)$$

Here $t_n = 1$ for samples x_n which are assigned to class C_1 and $t_n = 0$ for samples from the complementary class C_2 .

We introduce the further shorthand

$$\phi_n = \phi_{x_n}$$

and, using

$$\frac{d}{da} \sigma(a) = \sigma(a)(1 - \sigma(a)),$$

obtain the gradient of the cross-entropy function

$$\nabla_{\theta} \Psi_{\text{data}}(\theta) = \sum_{n=1}^N (y_n - t_n) \phi_n. \quad (5)$$

We rewrite this gradient more compactly as

$$\nabla_{\theta} \Psi_{\text{data}}(\theta) = \Phi(y(\theta) - t)$$

using

$$y(\theta) = (y_1, \dots, y_N)^T \in \mathbb{R}^N, \quad t = (t_1, \dots, t_N)^T \in \mathbb{R}^N,$$

and

$$\Phi = (\phi_1, \dots, \phi_N) \in \mathbb{R}^{D \times N}.$$

The term $y(\theta) - t$ is often called the innovation in the Kalman filter literature. The Hessian matrix of second-order derivative is given by

$$D_{\theta}^2 \Psi_{\text{data}}(\theta) = \sum_{n=1}^N \phi_n y_n (1 - y_n) \phi_n^T = \Phi R(\theta) \Phi^T,$$

where $R(\theta) \in \mathbb{R}^{N \times N}$ denotes the diagonal matrix with entries $r_{nn} = y_n(1 - y_n)$.

Let us denote a minimiser of $\Psi_{\text{data}}(\theta)$ by θ_{MLE} . We note that the maximum likelihood estimators (MLEs) are not uniquely determined in the over-parametrised setting, that is, when $D \gg N$. For that reason and in order to avoid overfitting when $D \approx N$, we introduce a Gaussian prior probability density function (PDF) over the parameters via

$$\pi_{\text{prior}}(\theta) \propto \exp(-\Psi_{\text{prior}}(\theta))$$

with

$$\Psi_{\text{prior}}(\theta) := \frac{1}{2}(\theta - m_{\text{prior}})^T \Sigma_{\text{prior}}^{-1}(\theta - m_{\text{prior}}),$$

that is, the posterior parameter distribution π_{post} is given by

$$\pi_{\text{post}}(\theta) \propto \exp(-\Psi_{\text{post}}(\theta)), \quad \Psi_{\text{post}}(\theta) := \Psi_{\text{data}}(\theta) + \Psi_{\text{prior}}(\theta). \tag{6}$$

Here the prior mean $m_{\text{prior}} \in \mathbb{R}^D$ and covariance matrix $\Sigma_{\text{prior}} \in \mathbb{R}^{D \times D}$ need to be chosen appropriately.

It is known that Ψ_{post} is a strongly convex function since $D_{\theta}^2 \Psi_{\text{data}}(\theta) \geq 0$ and, hence, the MAP estimator, that is,

$$\theta_{\text{MAP}} = \arg \min_{\theta \in \mathbb{R}^D} \Psi_{\text{post}}(\theta),$$

is uniquely defined. Furthermore, assuming the availability of sufficiently many data points, that is $N \gg 1$, the posterior is well approximated by a Gaussian with mean θ_{MAP} and covariance matrix

$$\Sigma_{\text{MAP}} = \left(D_{\theta}^2 \Psi_{\text{post}}(\theta_{\text{MAP}}) \right)^{-1},$$

where

$$D_{\theta}^2 \Psi_{\text{post}}(\theta) = \Phi R(\theta) \Phi^T + \Sigma_{\text{prior}}^{-1} \in \mathbb{R}^{D \times D}$$

denotes the Hessian matrix of second-order derivatives. This is an implication of the Bernstein–von Mises theorem [33].

3 The Homotopy Approach to Bayesian Inference

While there are many Monte Carlo methods available for sampling from the posterior π_{post} , we focus in this paper on ensemble transform methods based on the homotopy approach [6,26]

$$\pi_{\tau}(\theta) \propto \exp(-\tau \Psi_{\text{data}}(\theta)) \pi_{\text{prior}}(\theta) \quad (7)$$

with $\tau \in [0, 1]$, that is, $\pi_0 = \pi_{\text{prior}}$ and $\pi_{\text{post}} = \pi_1$. The PDFs π_{τ} satisfy the evolution equation

$$\partial_{\tau} \pi_{\tau} = -(\Psi_{\text{data}} - [\Psi_{\text{data}}]) \pi_{\tau}, \quad (8)$$

where $\pi_{\tau}[f]$ denotes the expectation value of a function $f: \mathbb{R}^D \rightarrow \mathbb{R}$ under the PDF π_{τ} .

In order to derive the associated evolution equations for random variables $\theta_{\tau} \sim \pi_{\tau}$, we introduce potentials $V_{\tau}: \mathbb{R}^D \rightarrow \mathbb{R}$, which are defined by the elliptic partial differential equation (PDE)

$$\nabla_{\theta} \cdot (\pi_{\tau} \Sigma_{\tau} \nabla_{\theta} V_{\tau}) = -(\Psi_{\text{data}} - \pi_{\tau}[\Psi_{\text{data}}]) \pi_{\tau}, \quad (9)$$

where Σ_{τ} denotes the covariance matrix of π_{τ} , that is,

$$\Sigma_{\tau} = \pi_{\tau} \left[(\theta - \pi_{\tau}[\theta])(\theta - \pi_{\tau}[\theta])^{\text{T}} \right].$$

Hence the evolution equations are given by

$$\frac{d}{d\tau} \theta_{\tau} = -\Sigma_{\tau} \nabla_{\theta} V_{\tau}(\theta_{\tau}), \quad \theta_0 \sim \pi_{\text{prior}}, \quad (10)$$

such that

$$\theta_1 \sim \pi_{\text{post}}. \quad (11)$$

See [26] for further details.

Remark 1 If different models $\mathcal{M}_i, i = 1, \dots, I_M$, need to be compared via their Bayes factors $\mathbb{P}[\mathcal{M}_i]$ [16], then the homotopy approach (8) can also be used and gives rise to the evolution equation

$$\frac{d}{d\tau} \mathbb{P}_{\tau} = -\pi_{\tau}[\Psi_{\text{data}}],$$

$\tau \in [0, 1]$, which is to be solved for each model \mathcal{M}_i with \mathbb{P}_0 its prior probability, π_{τ} its transformed parameter PDFs, as defined by (8), and Ψ_{data} the model's negative log-likelihood function. It then holds that $\mathbb{P}[\mathcal{M}_i] = \mathbb{P}_1$.

Remark 2 It is, of course, feasible to use a stopping time different from $\tau = 1$ in the homotopy approach (10). This has already been hinted at in in the context of optimisation and Tikhonov regularisations in [26]; while subsequently being fully explored by the EKI methodology. See [15,17] for a detailed description in the context of machine learning as well as for an extensive literature survey.

3.1 Affine-Invariance

Note that any vector field $F_\tau(\theta)$, which satisfies

$$\nabla_\theta \cdot (\pi_\tau F_\tau) = -(\Psi_{\text{data}} - \pi_\tau [\Psi_{\text{data}}]) \pi_\tau,$$

will lead to an associated evolution equation

$$\frac{d}{d\tau} \theta_\tau = -F_\tau(\theta_\tau), \quad \theta_0 \sim \pi_{\text{prior}} \tag{12}$$

such that (11) holds. The specific choice $F_\tau = \Sigma_\tau \nabla V_\tau$ in (9) is motivated by the concept of affine-invariance [10,11], which we explain next.

Let

$$\theta = A\bar{\theta} + b \tag{13}$$

denote an invertible affine transformation between two variables $\theta \in \mathbb{R}^D$ and $\bar{\theta} \in \mathbb{R}^D$. Furthermore, the Bayesian inference problem for the variable θ is defined by Ψ_{data} and Ψ_{prior} and the corresponding problem for the transformed variable $\bar{\theta}$ by

$$\bar{\Psi}_{\text{data}}(\bar{\theta}) := \Psi_{\text{data}}(A\bar{\theta} + b), \quad \bar{\Psi}_{\text{prior}}(\bar{\theta}) := \Psi_{\text{prior}}(A\bar{\theta} + b).$$

Hence the PDFs $\bar{\pi}_\tau$ for $\bar{\theta}$ satisfy

$$\bar{\pi}_\tau(\bar{\theta}) = \pi_\tau(A\bar{\theta} + b) |A|, \quad \tau \in [0, 1].$$

Definition 1 A homotopy formulation (12) is called affine-invariant if the associated vector fields F_τ and \bar{F}_τ satisfy

$$F_\tau(A\bar{\theta} + b) = A\bar{F}_\tau(\bar{\theta}) \tag{14}$$

for any suitable choice of $A \in \mathbb{R}^{D \times D}$ and $b \in \mathbb{R}^D$, which implies $\theta_\tau = A\bar{\theta}_\tau + b$ for all $\tau > 0$ if it holds at $\tau = 0$.

Lemma 1 *The choice*

$$F_\tau(\theta) = \Sigma_\tau \nabla_\theta V_\tau(\theta) \tag{15}$$

leads to an affine-invariant homotopy formulation and the transformed vector field is given by

$$\bar{F}_\tau(\bar{\theta}) = \bar{\Sigma}_\tau \nabla_{\bar{\theta}} \bar{V}_\tau(\bar{\theta})$$

with the potential \bar{V}_τ satisfying the elliptic PDE

$$\nabla_{\bar{\theta}} \cdot (\bar{\pi}_\tau \bar{\Sigma}_\tau \nabla_{\bar{\theta}} \bar{V}_\tau) = -\bar{\pi}_\tau (\bar{\Psi}_{\text{data}} - \bar{\pi}_\tau [\bar{\Psi}_{\text{data}}]). \tag{16}$$

Proof Since the covariance matrices satisfy

$$\Sigma_\tau = A \bar{\Sigma}_\tau A^T,$$

the PDE (9) transforms into the PDE (16) under the affine transformation (13). Hence we have

$$\bar{V}_\tau(\bar{\theta}) = V_\tau(A\bar{\theta} + b)$$

and obtain

$$\begin{aligned} \Sigma_\tau \nabla_\theta V_\tau(\theta) &= A \bar{\Sigma}_\tau A^T \nabla_\theta V_\tau(A\bar{\theta} + b) \\ &= A \bar{\Sigma}_\tau \nabla_{\bar{\theta}} V_\tau(A\bar{\theta} + b) \\ &= A \bar{\Sigma}_\tau \nabla_{\bar{\theta}} \bar{V}_\tau(\bar{\theta}), \end{aligned}$$

which verifies (14). □

We note that the affine invariance of (10) is maintained under the forward Euler time discretisation

$$\theta_{\tau_{k+1}} = \theta_{\tau_k} - \Delta\tau \Sigma_{\tau_k} \nabla_\theta V_{\tau_k}(\theta_{\tau_k})$$

with step-size $\Delta\tau > 0$ and $\tau_k = k\Delta\tau, k = 0, \dots, K - 1$ such that $K\Delta\tau = 1$. This immediately follows from property (14).

Remark 3 The choice of the vector field $F_\tau(\theta)$ in (15) is not unique. It is, for example, possible to add any drift term $\Sigma_\tau \nabla_\theta U_\tau(\theta)$ with the potential U_τ chosen such that

$$\nabla_\theta \cdot (\pi_\tau \Sigma_\tau \nabla_\theta U_\tau) = 0. \tag{17}$$

A possible choice is

$$U_\tau(\theta) = \log \pi_\tau + \tau \Psi_{\text{data}}(\theta) - \log \pi_{\text{prior}}(\theta)$$

and the associated operator on the left hand side of (17) corresponds to the nonlinear Fokker–Planck equation arising from the EKS/ALDI mean-field evolution equations [9,10] for sampling from the PDF π_τ . Because of (7), $U_\tau(\theta) = \text{const.}$ and (17) follows trivially. Upon formally replacing the drift term $\Sigma_\tau \nabla_\theta U_\tau(\theta)$ by its EKS/ALDI stochastic mean-field representation, (10) turns into the affine-invariant stochastic evolution equation

$$d\theta_\tau = -\Sigma_\tau \nabla_\theta \{V_\tau(\theta) + \tau \nabla_\theta \Psi_{\text{data}}(\theta) - \log \pi_{\text{prior}}(\theta)\} dt + \Sigma_\tau^{1/2} dW_\tau, \tag{18}$$

where W_τ denotes standard D -dimensional Brownian motion. The associated marginal PDF π_τ still satisfies (8) provided $\pi_0 = \pi_{\text{prior}}$ at initial time $\tau = 0$.

3.2 Data Sub-sampling

If the data sets are large, then it makes sense to randomly sub-sample the data and to only use these subsets in the computation of the gradient (5). More precisely, let Δ denote a randomly selected subset of $\{1, \dots, N\}$ without replacement of cardinality $N' \leq N$. Given such a mini-batch Δ , the full gradient (5) is replaced by its approximation

$$\nabla_{\theta} \Psi_{\text{data}}(\theta) \approx \frac{N}{N'} \sum_{n \in \Delta} (y_n - t_n) \phi_n.$$

More abstractly, data sub-sampling in terms of mini-batches gives rise to random negative log-likelihood functions $\Psi_{\text{data}}^{\gamma}$, where γ characterises the chosen mini-batch Δ^{γ} and is the realisation of an appropriate random variable Γ such that

$$\mathbb{E}_{\Gamma}[\Psi_{\text{data}}^{\gamma}(\theta)] = \Psi_{\text{data}}(\theta) \tag{19}$$

for all parameter values $\theta \in \mathbb{R}^D$. See [20] for a detailed investigation of data sub-sampling in the context of stochastic gradient descent and related methods. If $\Psi_{\text{data}}^{\gamma}$ is now being used in (9) instead of Ψ_{data} , then we denote the resulting potential by V_{τ}^{γ} , which is also unbiased, that is,

$$\mathbb{E}_{\Gamma}[\nabla_{\theta} V_{\tau}^{\gamma}(\theta)] = \nabla_{\theta} V_{\tau}(\theta)$$

due to (19). Let us now time-step the implied mean-field ODEs (10) by the forward Euler method, that is,

$$\theta_{\tau_{k+1}} = \theta_{\tau_k} - \Delta\tau \Sigma_{\tau_k} \nabla_{\theta} V_{\tau_k}^{\gamma_k}(\theta_{\tau_k}), \tag{20}$$

where the γ_k 's are independent realisations of the random variable Γ . A formal application of the modified equation analysis of [20] leads to the stochastic differential equation (SDE)

$$d\theta_{\tau} = -\Sigma_{\tau} \left\{ \nabla_{\theta} V_{\tau}(\theta_{\tau}) d\tau - (\Delta\tau \Omega_{\tau}(\theta_{\tau}))^{1/2} dW_{\tau} \right\}, \tag{21}$$

where W_{τ} denotes standard D -dimensional Brownian motion and

$$\Omega_{\tau}(\theta) := \text{cov}_{\Gamma} [\nabla_{\theta} V_{\tau}^{\gamma}(\theta)].$$

Here $\text{cov}_{\Gamma}[f^{\gamma}]$ denotes the covariance matrix of a vector-valued random function $f^{\gamma} \in \mathbb{R}^D$ with respect to the random variable Γ . We conclude from (21) that the approximation error inflicted on the posterior distribution at $\tau = 1$ can be made arbitrarily small by reducing the time-step $\Delta\tau$ of the forward Euler method (20). While a rigorous mathematical analysis is left to future work, we discuss an application to Cox point processes in Sect. 7.3 below and analyse a Gaussian approximation in the Appendix in more detail.

4 Ensemble Transform Algorithms

We now discuss several Monte Carlo implementations of the proposed affine-invariant homotopy approach to Bayesian inference for logistic regression. All methods start with a set of M samples $\theta_0^{(i)}, i = 1, \dots, M$, from the prior distribution $\pi_0 = \pi_{\text{prior}}$. The three methods considered here differ in the evolution equations for the ensemble $\{\theta_\tau^{(i)}\}$. We introduce the empirical mean

$$\hat{m}_\tau = \frac{1}{M} \sum_{i=1}^M \theta_\tau^{(i)}$$

and the empirical covariance matrix

$$\hat{\Sigma}_\tau = \frac{1}{M-1} \sum_{i=1}^M (\theta_\tau^{(i)} - \hat{m}_\tau)(\theta_\tau^{(i)} - \hat{m}_\tau)^T$$

for $\tau \geq 0$. We also introduce the matrix $\Theta_\tau \in \mathbb{R}^{D \times M}$ of ensemble deviations

$$\Theta_\tau = \left(\theta_\tau^{(1)} - \hat{m}_\tau, \theta_\tau^{(2)} - \hat{m}_\tau, \dots, \theta_\tau^{(M)} - \hat{m}_\tau \right), \tag{22}$$

which leads to the compact representation

$$\hat{\Sigma}_\tau = \frac{1}{M-1} \Theta_\tau \Theta_\tau^T$$

of the empirical covariance matrix. The empirical measure associated with the ensemble $\{\theta_\tau^{(i)}\}$ at time τ is denoted by $\hat{\pi}_\tau$ implying, for example,

$$\hat{m}_\tau = \hat{\pi}_\tau[\theta].$$

More generally,

$$\hat{\pi}_\tau[f] = \frac{1}{M} \sum_{i=1}^M f(\theta_\tau^{(i)})$$

for an observable $f(\theta)$. We describe an affine-invariant extension of the FPF next. The key building stone is a diffusion map approximation to the PDE (9).

4.1 Diffusion Maps and the FPF

The right-hand side of the elliptic PDE (9) leads to the elliptic operator

$$\Delta_\pi \phi := \frac{1}{\pi} \nabla_\theta \cdot (\pi \Sigma \nabla_\theta \phi)$$

with covariance matrix

$$\Sigma = \pi \left[(\theta - \pi[\theta])(\theta - \pi[\theta])^T \right]$$

and (9) can be rewritten as

$$-\Delta_{\pi_\tau} V_\tau = \Psi_{\text{data}} - \pi_\tau [\Psi_{\text{data}}].$$

Following the FPF methodology, as carefully described in [34], we regularise the elliptic PDE problem (9) through the fixed-point equation

$$\tilde{V}_\tau = P_\epsilon \tilde{V}_\tau + \epsilon(\Psi_{\text{data}} - \pi_\tau [\Psi_{\text{data}}])$$

with regularisation parameter $\epsilon > 0$ and semigroup $P_\epsilon = e^{\epsilon \Delta_{\pi_\tau}}$. One then replaces V_τ by \tilde{V}_τ in the evolution Eq. (10).

We note that P_ϵ is the semi-group associated with the mean-field SDE

$$dS_\epsilon = \Sigma_\tau \nabla_\theta \log \pi_\tau(S_\epsilon) d\epsilon + \sqrt{2} \Sigma_\tau^{1/2} dW_\epsilon, \tag{23}$$

where $W_\epsilon, \epsilon \geq 0$, is standard Brownian motion in \mathbb{R}^D . The SDE (23) is affine invariant and is closely related to the EKS/ALDI mean-field equations [9,10].

The diffusion map approximation T_ϵ of P_ϵ is defined as follows:

$$T_\epsilon f(\theta) = \frac{1}{n_\epsilon(\theta)} \int_{\mathbb{R}^D} k_\epsilon(\theta, \theta') f(\theta') \pi_\tau(\theta') d\theta',$$

where $n_\epsilon(\theta) := \int k_\epsilon(\theta, \theta') \pi_\tau(\theta') d\theta'$ is the normalisation constant,

$$k_\epsilon(\theta, \theta') := \frac{g_\epsilon(\theta, \theta')}{\sqrt{\int g_\epsilon(\theta, \theta'') \pi_\tau(\theta'') d\theta''} \sqrt{\int g_\epsilon(\theta', \theta'') \pi_\tau(\theta'') d\theta''}}$$

and

$$g_\epsilon(\theta, \theta') := \exp \left(-\frac{1}{4\epsilon} (\theta - \theta')^T \Sigma_\tau^{-1} (\theta - \theta') \right) \tag{24}$$

is the Gaussian kernel associated with the pure diffusion in (23).

The next step is to replace π_τ by the ensemble-induced empirical measure $\hat{\pi}_\tau$, which results in the empirical approximation

$$T_\epsilon^{(M)} f(\theta) := \frac{1}{n_\epsilon^{(M)}(\theta)} \sum_{i=1}^M k_\epsilon^{(M)}(\theta, \theta_\tau^{(i)}) f(\theta_\tau^{(i)}),$$

where $n_\epsilon^{(M)}(\theta) := \sum_i k_\epsilon^{(M)}(\theta, \theta_\tau^{(i)})$ is the normalisation constant, and

$$k_\epsilon^{(M)}(\theta, \theta') := \frac{g_\epsilon^{(M)}(\theta, \theta')}{\sqrt{\sum_i g_\epsilon^{(M)}(\theta, \theta_\tau^{(i)})} \sqrt{\sum_i g_\epsilon^{(M)}(\theta', \theta_\tau^{(i)})}}.$$

Here $g_\epsilon^{(M)}$ denotes the Gaussian kernel (24) with the covariance matrix Σ_τ replaced by its empirical estimate $\hat{\Sigma}_\tau$. One finally introduces the $M \times M$ Markov matrix T with entries

$$T_{ij} = \frac{1}{n_\epsilon^{(M)}(\theta_\tau^{(i)})} k_\epsilon^{(M)}(\theta_\tau^{(i)}, \theta_\tau^{(j)}) \tag{25}$$

and the finite-dimensional fixed-point equation

$$\tilde{V} = T\tilde{V} + \epsilon \Delta\Psi_{\text{data}}$$

with

$$\Delta\Psi_{\text{data}} := \left(\Psi_{\text{data}}(\theta_\tau^{(1)}) - \hat{\pi}_\tau[\Psi_{\text{data}}], \dots, \Psi_{\text{data}}(\theta_\tau^{(M)}) - \hat{\pi}_\tau[\Psi_{\text{data}}] \right)^T.$$

Given the solution vector $\tilde{V} \in \mathbb{R}^M$ with entries $\tilde{V}^{(j)}, j = 1, \dots, M$, the approximation to the drift term in the ODE (10) is now provided as follows [34]:

$$\hat{\Sigma}_\tau \nabla_\theta \tilde{V}_\tau(\theta) := \hat{\Sigma}_\tau \nabla_\theta \left\{ \frac{1}{n_\epsilon^{(M)}(\theta)} \sum_{j=1}^M k_\epsilon^{(M)}(\theta, \theta_\tau^{(j)}) \left(\tilde{V}^{(j)} + \epsilon \Delta\Psi_{\text{data}}^{(j)} \right) \right\}.$$

Lemma 2 *The affine-invariant FPF drift term for the ensemble members $\theta_\tau^{(i)}, i = 1, \dots, M$, is given by*

$$\hat{\Sigma}_\tau \nabla_\theta \tilde{V}_\tau(\theta_\tau^{(i)}) = \sum_{j=1}^M s_{ij} \theta_\tau^{(j)} \tag{26}$$

with coefficients

$$s_{ij} = \frac{1}{2\epsilon} T_{ij} \left(r_j - \sum_{k=1}^M T_{ik} r_k \right), \quad r_j := \tilde{V}^{(j)} + \epsilon \Delta\Psi_{\text{data}}^{(j)}.$$

Proof The calculations leading to (26) are identical to those for the standard FPF [34] taking note that the covariance matrix $\hat{\Sigma}_\tau$ cancels out. □

The resulting interacting particle system can be propagated in time using the forward Euler method with step size $\Delta\tau = 1/K$, that is,

$$\theta_{\tau_{k+1}}^{(i)} = \theta_{\tau_k}^{(i)} - \Delta\tau \hat{\Sigma}_{\tau_k} \nabla_\theta \tilde{V}_{\tau_k}(\theta_{\tau_k}^{(i)}), \quad i = 1, \dots, M, \tag{27}$$

$k = 0, \dots, K - 1$, with the right-hand side defined by (26). Also note that $M \rightarrow \infty$ leads formally to the mean-field forward Euler method (20).

Remark 4 We note that the Markov matrix T , defined by (25), is non-reversible contrary to the underlying stochastic process (23). This problem can be addressed by a bi-stochastic approximation, which also leads to improved convergence rates. See [37] for more details.

While the FPF leads to a consistent approximation in the limit $M \rightarrow \infty$ and $\epsilon \rightarrow 0$, its application is limited to low-dimensional problems. We next describe two approximations that extend to high-dimensional inference problems.

4.2 Second-Order Methods

One can derive evolution equations for the mean m_τ and the covariance matrix Σ_τ of π_τ from Eq. (8). Furthermore, if one assumes that π_τ is well approximated by a Gaussian PDF with mean m_τ and covariance matrix Σ_τ , the following evolution equations for the mean and the covariance matrix arise.

Lemma 3 *If the temporal PDF π_t in (9) is Gaussian with mean m_τ and covariance matrix Σ_τ , then their evolution equations are given by*

$$\frac{d}{d\tau} m_\tau = -\Sigma_\tau \pi_\tau [\nabla_\theta \Psi_{\text{data}}] = -\Sigma_\tau \Phi(\pi_\tau[y] - t),$$

and

$$\frac{d}{d\tau} \Sigma_\tau = -\Sigma_\tau \pi_\tau [D_\theta^2 \Psi_{\text{data}}] \Sigma_\tau = -\Sigma_\tau \Phi \pi_\tau [R] \Phi^T \Sigma_\tau,$$

respectively.

Proof The evolution equations are derived using the following well-known identity for Gaussian PDFs π_τ (see, for example, [23]):

$$\pi_\tau[\theta (f(\theta) - \pi_\tau[f])] = \Sigma_\tau \pi_\tau[\nabla_\theta f] \tag{28}$$

as well as

$$\pi_\tau[(\theta - \pi_\tau[\theta])(\theta - \pi_\tau[\theta])^T (f(\theta) - \pi_\tau[f])] = \Sigma_\tau \pi_\tau [D_\theta^2 f] \Sigma_\tau \tag{29}$$

for any scalar-valued smooth function $f : \mathbb{R}^D \rightarrow \mathbb{R}$. □

In order to derive corresponding equations for the ensemble $\{\theta_\tau^{(i)}\}$, we introduce

$$\hat{y}_\tau := \hat{\pi}_\tau[y] = \frac{1}{M} \sum_{j=1}^M y_\tau^{(j)}, \quad y_\tau^{(j)} := y(\theta_\tau^{(j)}) \in \mathbb{R}^N,$$

as well as

$$\hat{R}_\tau := \hat{\pi}_\tau[R] = \frac{1}{M} \sum_{j=1}^M R_\tau^{(j)},$$

where the $R_\tau^{(j)} \in \mathbb{R}^{N \times N}$ are diagonal matrices with entries $r_{nn}^{(j)} = y_{\tau,n}^{(j)} (1 - y_{\tau,n}^{(j)})$ for $n = 1, \dots, N$ with

$$y_{\tau,n}^{(j)} = \sigma \left((\theta_\tau^{(j)})^\top \phi_{x_n} \right) \in \mathbb{R}.$$

We finally obtain the evolution equations

$$\frac{d}{d\tau} \hat{m}_\tau = -\hat{\Sigma}_\tau \Phi(\hat{y}_\tau - t) \quad (30)$$

for the ensemble mean and

$$\frac{d}{d\tau} \Theta_\tau = -\frac{1}{2} \hat{\Sigma}_\tau \Phi \hat{R}_\tau \Phi^\top \Theta_\tau \quad (31)$$

for the ensemble deviations (22). Note that these equations can be used to propagate the ensemble $\{\theta_\tau^{(i)}\}$ regardless of the Gaussian assumption made for their derivation. In the following section, we will introduce further approximations which allow us to introduce an affine-invariant gradient flow structure.

4.3 Ensemble Kalman–Bucy Filter

We now turn to a formulation with gradient flow structure in the spirit of the ensemble Kalman–Bucy filter (EnKBF) for quadratic loss functions (1) [3,26,28]. Provided that

$$\hat{y}_\tau = \hat{\pi}_\tau[y] \approx y(\hat{m}_\tau) = y(\hat{\pi}_\tau[\theta])$$

and

$$y_\tau^{(i)} = y \left(\theta_\tau^{(i)} \right) \approx \hat{y}_\tau + \hat{R}_\tau \Phi^\top (\theta_\tau^{(i)} - \hat{m}_\tau),$$

one formally obtains the approximation

$$\frac{d}{dt} \theta_\tau^{(i)} = -\frac{1}{2} \hat{\Sigma}_\tau \Phi (y_\tau^{(i)} + y(\hat{m}_\tau) - 2t). \quad (32)$$

to the second-order Eqs. (30)–(31).

Lemma 4 *The EnKBF equations (32) are of affine-invariant gradient structure*

$$\frac{d}{dt} \theta_\tau^{(i)} = -\hat{\Sigma}_\tau \nabla_{\theta^{(i)}} \Psi_{\text{KBF}}(\{\theta_\tau^{(j)}\}), \tag{33}$$

with potential

$$\Psi_{\text{KBF}}(\{\theta^{(j)}\}) = \frac{1}{2} \sum_{j=1}^M \Psi_{\text{data}}(\theta^{(j)}) + \frac{M}{2} \Psi_{\text{data}}(\hat{m}_\tau).$$

Proof The lemma can be verified by direct calculation and taking note of

$$\nabla_{\theta^{(i)}} \Psi_{\text{data}}(\hat{m}_\tau) = \frac{1}{M} \Phi(y(\hat{m}_\tau) - t)$$

in particular. □

The Eq. (33) also possesses an affine-invariant gradient-flow structure in the space of probability measures in the mean-field limit $M \rightarrow \infty$. See [9,10,28] for details.

Remark 5 The authors of [17] proposed a different modification of the EnKF for classification problems. In our notation, it corresponds to

$$\frac{d}{dt} \theta_\tau^{(i)} = - \sum_{n=1}^N \left\{ \frac{1}{M-1} \sum_{j=1}^M \left(y_{\tau,n}^{(j)} - \frac{1}{M} \sum_{l=1}^M y_{\tau,n}^{(l)} \right) \left(\frac{t_n}{y_{\tau,n}^{(i)}} - \frac{1-t_n}{1-y_{\tau,n}^{(i)}} \right) \theta_\tau^{(j)} \right\}.$$

The modified EnKF formulation of [13] uses an implementation that is closer to ours but does not actually propagate ensembles. Also recall that both methods are to be used for reducing the loss function $\Psi_{\text{data}}(\theta)$ via minimisation instead of sampling from the posterior PDF (6), which is the subject of this paper.

Remark 6 In addition to the EnKBF formulation considered so far, there exists another variant which is based on stochastic perturbations in the innovation [8,26]. We briefly explain how to extend this alternative formulation to logistic regression. We recall that the ensemble deviations Θ_τ satisfy the evolution Eq. (31). The same propagation of the first and second-order moments is achieved by the stochastic equations

$$d(\theta_\tau - m_\tau) = -\Sigma_\tau \Psi \left\{ \pi_\tau [R] \Psi^T (\theta_\tau - m_\tau) d\tau + \pi_\tau [R]^{1/2} dW_\tau \right\}$$

in the mean-field limit $M \rightarrow \infty$, where W_τ denotes standard N -dimensional Brownian motion. This suggests to replace the finite ensemble size formulation (32) by the interacting system of SDEs

$$d\theta_\tau^{(i)} = -\hat{\Sigma}_\tau \Psi \left(y_\tau^{(i)} d\tau + \hat{R}_\tau dW_\tau^{(i)} - t d\tau \right).$$

Here $W_\tau^{(i)}, i = 1, \dots, M$, denote independent N -dimensional Brownian motions.

Again, the forward Euler method (27) can be used to solve (32) or (33) in time. However, the inherent stiffness of the equations of motion can lead to very small time-steps [2]. This issue and the rank deficiency of $\hat{\Sigma}_\tau$ for $M \leq D$ are addressed in the following subsection.

4.4 Dropout and Time-Stepping

We now discuss two numerical issue that are common to all ensemble-base methods considered in this section.

First, we observe that whenever $M \leq D$, then the empirical covariance matrix $\hat{\Sigma}_\tau$ is rank-deficient and the ensemble $\{\theta_\tau^{(i)}\}$ is propagated in the subspace defined by the ensemble at initial time. The concept of covariance localisation has been pioneered in the geoscience community in order to lift this restriction [8,29]. In the context of logistic regression we suggest to instead utilise the idea of dropout training as pioneered in the machine learning community [32]. We employ the concept of dropout as follows: In each time step, a randomly chosen number of entries in $\theta_\tau^{(i)}$, $i = 1, \dots, M$, is set to zero. The resulting modified matrix of ensemble deviations (22) is denoted by $\tilde{\Theta}_\tau$. The empirical covariance matrix is now defined by

$$\hat{\Sigma}_\tau = \frac{1}{(1 - \mu)(M - 1)} \tilde{\Theta}_\tau \tilde{\Theta}_\tau^T, \quad (34)$$

where $\mu \in (0, 1)$ denotes the fraction of entries randomly set to zero. All other aspects of the previously considered algorithms remain unaltered. The estimator (34) underestimates the cross-correlations. An unbiased estimator requires an additional rescaling of the off-diagonal entries by $1/(1 - \mu)$.

Remark 7 There are alternative methods for breaking the sub-space property of ensemble transform methods in case of $M \leq D$. For example, the authors of [17] suggest to randomly perturb the ensemble members after each time step of EKI. In our context, this approach can be viewed as combining the EnKBF with diffusion at low temperature. However, it is not clear that diffusion is sufficient to effectively explore the full parameter space in the context of the homotopy approach considered in this paper.

Second, we return to the issue of efficient time-stepping of the interacting particle systems. We had previously considered the forward Euler method (27). However, the method can encounter severe step-size restrictions due to its restricted domain of stability. Following [2], we therefore consider the following tamed version of the forward Euler discretisation (27) for the EnKBF formulation from Sect. 4.3. Starting from (32) we employ

$$\theta_{\tau_{k+1}}^{(i)} = \theta_{\tau_k}^{(i)} - \frac{\Delta\tau}{2} \hat{\Sigma}_{\tau_k} \Phi \left(I + \Delta\tau \hat{R}_{\tau_k} \Phi^T \hat{\Sigma}_{\tau_k} \Phi \right)^{-1} (y_{\tau_k}^{(i)} + y(\hat{m}_{\tau_k}) - 2t) \quad (35)$$

with step size $\Delta\tau > 0$ for $k = 0, \dots, K - 1$. If the necessary matrix inversions become too computationally expensive, one can diagonalise the inverse as proposed in [2]. The same modification can be applied to the second-order formulation (30)–(31).

5 Ensemble Transform Langevin Dynamics

While we have focused on the homotopy approach in the previous section, we now combine the homotopy approach with overdamped Langevin dynamics in order to sample from the desired target distribution (6) as $\tau \rightarrow \infty$. Let us therefore turn the transport-based differential equation (10) into the McKean–Vlasov SDE

$$d\theta_\tau = -\Sigma_\tau \left\{ \nabla_\theta V_\tau(\theta_\tau) + \frac{1}{2} \Sigma_{\text{prior}}^{-1}(\theta_\tau + m_\tau - 2m_{\text{prior}}) \right\} d\tau + \Sigma_\tau^{1/2} dW_\tau \quad (36)$$

by adding diffusion and a prior-related drift term to (10). This SDE is motivated by the affine-invariant EKS [9] and ALDI [10] variants of standard Langevin dynamics. Note that the crucial differences are that (i) EKS/ALDI uses the gradient $\nabla_\theta \Psi_{\text{prior}}$ of the prior in the drift term while (36) contains

$$\frac{1}{2} \Sigma_{\text{prior}}^{-1}(\theta_\tau + m_\tau - 2m_{\text{prior}})$$

instead, and (ii) in a similar vein, the potential $\Psi_{\text{data}}(\theta)$ is replaced in (36) by a time-dependent potential V_τ which is defined as the solution of the elliptic PDE (9). We will find that these modifications are motivated by the two facts that (i) one can find efficient gradient-free time-stepping methods for the SDE (36) and (ii) that the SDE (36) samples from the posterior PDF in the case of Gaussian distributions as we demonstrate next.

Lemma 5 *If the initial π_0 is Gaussian and Ψ_{data} is a potential of the form*

$$\Psi_{\text{data}}(\theta) = \frac{1}{2}(G\theta - t)^T \Gamma^{-1}(G\theta - t),$$

then the marginal PDFs π_τ of (36) remain Gaussian for all $\tau > 0$, and the SDE (36) has a stationary measure which is given by $\pi_\infty(\theta) \propto \exp(-\Psi_{\text{post}}(\theta))$.

Proof Under the stated assumption, we have

$$-\frac{1}{\pi_\tau(\theta)} \nabla_\theta \cdot (\pi_\tau(\theta) \Sigma_\tau \nabla_\theta V_\tau(\theta)) = (\theta - m_\tau)^T \nabla_\theta V_\tau(\theta) - \nabla_\theta \cdot \Sigma_\tau \nabla_\theta V_\tau(\theta)$$

as well as

$$\begin{aligned} \Psi_{\text{data}}(\theta) - \pi_\tau[\Psi_{\text{data}}] &= \frac{1}{2}(G\theta - t)^T \Gamma^{-1}(G\theta - t) \\ &\quad - \frac{1}{2} \pi_\tau[\theta^T G^T \Gamma^{-1} G\theta] + m_\tau^T G^T \Gamma^{-1} t - \frac{1}{2} t^T \Gamma^{-1} t \\ &= \frac{1}{2}(\theta - m_\tau)^T G^T \Gamma^{-1}(G\theta + Gm_\tau - 2t) \\ &\quad - \frac{1}{2} \pi_\tau[\theta^T G^T \Gamma^{-1} G\theta] + \frac{1}{2} m_\tau^T G^T \Gamma^{-1} Gm_\tau. \end{aligned}$$

Equating both expressions leads us to

$$\nabla_{\theta} V_{\tau}(\theta) = \frac{1}{2} G^T \Gamma^{-1} (G\theta + Gm_{\tau} - 2t). \quad (37)$$

The evolution equations arising from (36) for the mean m_{τ} and the covariance matrix Σ_{τ} are therefore given by

$$\begin{aligned} \frac{d}{d\tau} m_{\tau} &= -\Sigma_{\tau} \left(G^T \Gamma^{-1} (Gm_{\tau} - t) + \Sigma_{\text{prior}}^{-1} (m_{\tau} - m_{\text{prior}}) \right) \\ &= -\Sigma_{\tau} \Sigma_{\text{post}}^{-1} (m_{\tau} - m_{\text{post}}) \end{aligned}$$

and

$$\begin{aligned} \frac{d}{d\tau} \Sigma_{\tau} &= -\Sigma_{\tau} \left(G^T \Gamma^{-1} G + \Sigma_{\text{prior}}^{-1} \right) \Sigma_{\tau} + \Sigma_{\tau} \\ &= -\Sigma_{\tau} \Sigma_{\text{post}}^{-1} \Sigma_{\tau} + \Sigma_{\tau}, \end{aligned}$$

respectively. The associated equilibrium solutions are

$$m_{\infty} = m_{\text{post}} := m_{\text{prior}} - \Sigma_{\text{post}} G^T \Gamma^{-1} (Gm_{\text{prior}} - t)$$

and

$$\Sigma_{\infty} = \Sigma_{\text{post}} := \left(G^T \Gamma^{-1} G + \Sigma_{\text{prior}}^{-1} \right)^{-1},$$

as desired. We also note that the linearity of the gradient (37) implies that the PDFs π_{τ} remain Gaussian. Finally, (37) is indeed the gradient of the potential

$$V_{\tau}(\theta) = \frac{1}{4} (G\theta - t)^T \Gamma^{-1} (G\theta - t) + \frac{1}{2} \theta^T G^T \Gamma^{-1} (Gm_{\tau} - t),$$

which concludes the proof. \square

Remark 8 One can replace the McKean–Vlasov SDE (36) by the following formulation

$$d\theta_{\tau} = -\Sigma_{\tau} \nabla_{\theta} U_{\tau}(\theta_{\tau}) d\tau + \Sigma_{\tau}^{1/2} dW_{\tau}, \quad (38)$$

where the potential U_{τ} now satisfies the elliptic PDE

$$\nabla_{\theta} \cdot (\pi_{\tau} \Sigma_{\tau} \nabla_{\theta} U_{\tau}) = -\pi_{\tau} (\Psi_{\text{post}} - \pi_{\tau} [\Psi_{\text{post}}])$$

instead of (9). Both formulations are equivalent for Gaussian PDFs π_{τ} . While (36) is closer to the homotopy approach of Sect. 4, formulation (38) has a closer resemblance to overdamped Langevin dynamics.

5.1 Asymptotic Behaviour

The case of a general negative log-likelihood function Ψ_{data} is more delicate. However, if the associated posterior (6) is non-Gaussian but is well approximated by a Gaussian, and this also applies to all intermediate PDFs π_τ , that is,

$$\pi_\tau(\theta) \approx \tilde{\pi}_\tau(\theta) \propto \exp\left(-\frac{1}{2}(\theta - m_\tau)^T \Sigma_\tau^{-1}(\theta - m_\tau)\right), \tag{39}$$

then following the arguments from Sect. 4.2, the first two moments satisfy

$$\frac{d}{d\tau} m_\tau = -\Sigma_t \tilde{\pi}_\tau [\nabla_\theta \Psi_{\text{post}}] \tag{40}$$

and

$$\frac{d}{dt} \Sigma_\tau = -\Sigma_t \tilde{\pi}_\tau [D_\theta^2 \Psi_{\text{post}}] \Sigma_t + \Sigma_t. \tag{41}$$

The equation for the covariance matrix is stable if $\tilde{\pi}_\tau [D_\theta^2 \Psi_{\text{post}}]$ is positive-definite for all $\tau \geq 0$. This is, for example, the case for logistic regression.

We now discuss the limiting $\tau \rightarrow \infty$ behavior of (36) in some more detail. Its equilibrium distributions are characterised by the following lemma.

Lemma 6 *Stationary measures π_∞ of the mean-field equation (36) satisfy the PDE*

$$\begin{aligned} & \frac{1}{2} \nabla_\theta \cdot \left(\pi_\infty \Sigma_\infty \left\{ \nabla_\theta \log \pi_\infty + \Sigma_{\text{prior}}^{-1}(\theta + m_\infty - 2m_{\text{prior}}) \right\} \right) \\ & = \pi_\infty (\Psi_{\text{data}} - \pi_\infty [\Psi_{\text{data}}]) \end{aligned}$$

or, in terms of expectation values with respect to test functions ϕ ,

$$\begin{aligned} & \frac{1}{2} \pi_\infty \left[\nabla_\theta \phi \cdot \Sigma_\infty \left\{ \nabla_\theta \log \pi_\infty + \Sigma_{\text{prior}}^{-1}(\theta + m_\infty - 2m_{\text{prior}}) \right\} \right] \\ & = \pi_\infty [\phi (\Psi_{\text{data}} - \pi_\infty [\Psi_{\text{data}}])]. \end{aligned} \tag{42}$$

Proof The lemma follows from the Fokker–Planck equation associated to (36) and the PDE (9). □

Again making the Gaussian approximation (39), formally setting $\tau = \infty$ and recalling Remark 8, we first obtain from (42) the identity

$$\frac{1}{2} \tilde{\pi}_\infty [\nabla_\theta \phi \cdot \Sigma_\infty \nabla_\theta \log \tilde{\pi}_\infty] = \tilde{\pi}_\infty [\phi (\Psi_{\text{post}} - \tilde{\pi}_\infty [\Psi_{\text{post}}])].$$

Then,

$$\frac{1}{2} \tilde{\pi}_\infty [\nabla_\theta \log \tilde{\pi}_\infty] = 0 = \tilde{\pi}_\infty [\nabla_\theta \Psi_{\text{post}}] \tag{43}$$

for $\phi(\theta) = \theta - m_\infty$, and

$$\tilde{\pi}_\infty[D_\theta^2 \log \tilde{\pi}_\infty] = \Sigma_\infty^{-1} = \tilde{\pi}_\infty[D_\theta^2 \Psi_{\text{post}}] \tag{44}$$

for $\phi(\theta) = (\theta - m_\infty)(\theta - m_\infty)^T$, which provide a self-consistent system of equations for the mean m_∞ and the covariance matrix Σ_∞ of the Gaussian approximation (39). Here we have again used (28) and (29), respectively.

Note that (43) and (44) are compatible with (40) and (41), respectively. However, (43) and (44) only require Gaussianity (or a near-Gaussianity) for the equilibrium distribution π_∞ .

5.2 Numerical Implementation

The most appealing aspect of our modified overdamped Langevin dynamics is the possibility of implementing the SDEs (36) or (38), respectively, in a derivative-free manner, that is, without explicit knowledge of the potential V_τ and its gradient.

We describe a suitable time-stepping method in detail for the formulation (36) utilising methods from sequential data assimilation, which are well suited to deal with the $-\Sigma_\tau \nabla_\theta V_\tau$ term in the SDE (36). In other words, we numerically solve the SDE (36) with a step size $\Delta\tau$ by alternating in each time step between (i) a data assimilation step (such as SMC or an ensemble transform filter [29]) with negative log-likelihood

$$l(\theta) = \Delta\tau \Psi_{\text{data}}(\theta)$$

and (ii) the SDE

$$d\theta_t = -\frac{1}{2} \Sigma_\tau \Sigma_{\text{prior}}^{-1} (\theta_t + m_\tau - 2m_{\text{prior}}) dt + \Sigma_\tau^{1/2} dW_t \tag{45}$$

over the same time-interval $\Delta\tau$. More precisely, given $\theta_{\tau_k} \sim \pi_{\tau_k}$ we first find

$$\tilde{\theta}_{\tau_{k+1}} \sim \exp(-\Delta\tau \Psi_{\text{data}}) \pi_{\tau_k} \tag{46}$$

followed by solving (45) with initial $\tilde{\theta}_{\tau_{k+1}}$ over the time-interval $\Delta\tau$ in order to obtain $\theta_{\tau_{k+1}} \sim \pi_{\tau_{k+1}}$ for $k = 0, \dots, K - 1$. Note that (45) can be solved robustly by the tamed forward Euler-Maruyama method

$$\begin{aligned} \theta_{\tau_{k+1}} = & \tilde{\theta}_{\tau_{k+1}} - \frac{\Delta\tau}{2} \tilde{\Sigma}_{\tau_{k+1}} \left(\Sigma_{\text{prior}} + \Delta\tau \tilde{\Sigma}_{\tau_{k+1}} \right)^{-1} (\tilde{\theta}_{\tau_{k+1}} + \tilde{m}_{\tau_{k+1}} - 2m_{\text{prior}}) \\ & + \sqrt{\Delta\tau} \tilde{\Sigma}_{\tau_{k+1}}^{1/2} \mathcal{E}_{k+1} \end{aligned}$$

with \mathcal{E}_{k+1} a D -dimensional standard Gaussian random variable. As for ALDI, the $\Sigma_\tau^{1/2}$ multiplying the Brownian motion in (45) requires a correction term in order to sample from the correct distribution for finite ensemble sizes M . See [10] for details.

Our numerical implementations of (46) in terms of ensembles $\{\theta_{\tau_k}^{(i)}\}$ rely on ensemble transform filters of the form [29]

$$\tilde{\theta}_{\tau_{k+1}}^{(j)} = \sum_{i=1}^M \theta_{\tau_k}^{(i)} s_{ij} \tag{47}$$

for $j = 1, \dots, M$ with appropriate coefficients s_{ij} , which satisfy

$$\sum_{i=1}^M s_{ij} = 1. \tag{48}$$

We collect the coefficients s_{ij} in the $M \times M$ matrix S .

Lemma 7 *An ensemble transform filter is affine-invariant (compare Sect. 3.1) if the associated coefficients \bar{s}_{ij} for the filter in the transformed variable $\bar{\theta}$ satisfy $s_{ij} = \bar{s}_{ij}$ for all $i, j = 1, \dots, M$.*

Proof If

$$\theta_{\tau_k}^{(i)} = A\bar{\theta}_{\tau_k}^{(i)} + b$$

at τ_k , then

$$\theta_{\tau_{k+1}}^{(j)} = \sum_{i=1}^M \theta_{\tau_k}^{(i)} s_{ij} = \sum_{i=1}^M (A\bar{\theta}_{\tau_k}^{(i)} + b) s_{ij} = A \sum_{i=1}^M \bar{\theta}_{\tau_k}^{(i)} \bar{s}_{ij} + b = A\bar{\theta}_{\tau_{k+1}}^{(j)} + b,$$

and we conclude that

$$\theta_{\tau_{k+1}}^{(i)} = A\bar{\theta}_{\tau_{k+1}}^{(i)} + b$$

at τ_{k+1} for all $i = 1, \dots, M$. □

The EnKF and the nonlinear ensemble transform filter (NETF) of [36], for example, are affine-invariant. More specifically, the NETF leads to the transformation matrix

$$S = w_{\tau_k} 1_M^T + \sqrt{M} \left(\text{diag}(w_{\tau_k}) - w_{\tau_k} w_{\tau_k}^T \right)^{1/2}$$

where $w_{\tau_k} \in \mathbb{R}^M$ is the vector of normalised importance weights with entries

$$w_{\tau_k}^{(i)} = \frac{\exp(-\Delta\tau \Psi_{\text{data}}(\theta_{\tau_k}^{(i)}))}{\sum_{j=1}^M \exp(-\Delta\tau \Psi_{\text{data}}(\theta_{\tau_k}^{(j)}))},$$

$1_M \in \mathbb{R}^M$ is the vector of ones, and $\text{diag}(w_{\tau_k}) \in \mathbb{R}^{M \times M}$ is the diagonal matrix with diagonal entries w_{τ_k} . The affine invariance follows from the invariance of the

importance weights w_{τ_k} . The affine invariance of the EnKF has been previously demonstrated, for example, in [22].

Remark 9 The ensemble transform particle filter (ETPF) of [27,29] can also be made affine-invariant by using the cost function

$$V(S) = \frac{1}{2} \sum_{i,j=1}^M s_{ij} (\theta_{\tau_k}^{(i)} - \theta_{\tau_k}^{(j)})^T \hat{\Sigma}_{\tau_k}^{-1} (\theta_{\tau_k}^{(i)} - \theta_{\tau_k}^{(j)}) \tag{49}$$

subject to the constraints $s_{ij} \geq 0$, (48), and

$$\frac{1}{M} \sum_{j=1}^M s_{ij} = w_{\tau_k}^{(i)} \tag{50}$$

in the optimal transport definition of the coefficients s_{ij} in (47). Again, the affine invariance follows from the affine invariance of both the cost function $V(S)$ and the importance weights w_{τ_k} .

6 Numerical Example

In this section, we provide a couple of relatively simple numerical illustrations for the methods proposed in this paper. We start with a low-dimensional toy problem for which all proposed methods can easily be implemented and tested.

Example 1 We follow the example from [5] Section 4.2.1. Specifically, consider the Gaussian likelihood functions

$$\pi(x|C_i) \propto \exp\left(-\frac{1}{2}(x - \mu_i)^T B^{-1}(x - \mu_i)\right)$$

for inputs x to belong to classes C_i , $i = 1, 2$. The implied posterior PDF $\pi(C_1|\phi_x)$ is given by (2) with $\phi_x = (x^T, 1)^T$, that is $D = J + 1$, and the true parameter value is

$$\theta = \begin{pmatrix} B^{-1}(\mu_1 - \mu_2) \\ -\frac{1}{2}\mu_1^T B^{-1}\mu_1 + \frac{1}{2}\mu_2^T B^{-1}\mu_2 + \ln \frac{\pi(C_1)}{\pi(C_2)} \end{pmatrix}. \tag{51}$$

Furthermore, given data points (x_n, t_n) , $n = 1, \dots, N$, the maximum likelihood estimators for $\pi(C_1)$, μ_i , and B are given by

$$\hat{\pi}(C_1) = \frac{N_1}{N}, \quad N_1 = \sum_{n=1}^N t_n,$$

$$\hat{\mu}_1 = \frac{1}{N_1} \sum_{n=1}^N t_n x_n, \quad \hat{\mu}_2 = \frac{1}{N - N_1} \sum_{n=1}^N (1 - t_n) x_n,$$

and

$$\hat{B} = \frac{1}{N} \left\{ \sum_{n=1}^N t_n (x_n - \hat{\mu}_1)(x_n - \hat{\mu}_1)^T + \sum_{n=1}^N (1 - t_n)(x_n - \hat{\mu}_2)(x_n - \hat{\mu}_2)^T \right\},$$

respectively. See Section 4.2.2 in [5] for details. In this example, we will however directly estimate the true parameter value (51) using Bayesian inference.

We consider the special case $J = 2$ and use $B = I$, I the identity matrix, $\mu_1 = (-1, -1)^T$, $\mu_2 = (2, 2)^T$, and $\pi(C_1) = \pi(C_2) = 1/2$ to generate $N = 100$ data points. We note that the chosen values for μ_1, μ_2 , and B lead to

$$\theta = \begin{pmatrix} -3 \\ -3 \\ 3 \end{pmatrix}. \tag{52}$$

We now test the proposed algorithms using (i) an informative Gaussian prior with mean (52) and covariance matrix $\Sigma_{\text{prior}} = I$, and (ii) a less informative Gaussian prior with mean $m_{\text{prior}} = 0$ and covariance matrix $\Sigma_{\text{prior}} = 4I$. We time step the affine-invariant FPF, the second-order filter (30)–(31), and the EnKBF (32) with the forward Euler method using a step size $\Delta\tau = 10^{-3}$. A further reduction of the step size did not change the results in a statistically significant manner. The McKean–Vlasov SDE (36) is also implemented with the NETF as the inner data assimilation step. Furthermore, we implemented the ALDI method [10] for comparison since it is known to exactly sample from the posterior PDF (6). Both ALDI and the McKean–Vlasov SDE (38) were run up to time $\tau = 10$ with step-size $\Delta\tau = 10^{-2}$ at which point the interacting particle systems were considered to be in equilibrium.

The implementation of the FPF, as well as the SDE (36), explicitly involve the evaluation of the negative log-likelihood function (3). We found that we needed to replace (4) by

$$y_n = 0.99 \sigma(\theta^T \phi_{x_n}) + 0.005$$

in order to avoid numerical instabilities due to exceedingly small values of $\ln y_n$. The band-width, ϵ , in (24) has been set to $\epsilon = 0.1$.

All experiments were repeated $L = 1000$ times in order to average out random sampling effects. The numerically computed posterior means and spectral norm of the posterior covariance matrices can be found in Tables 1 and 2, respectively, for the informative prior, and in Tables 3 and 4, respectively, for the less informative prior. Reference values from an ALDI simulation are provided in the captions. While all methods reproduce those reference values very well in case of the informative prior, this picture changes for the less informative prior, where one expects a stronger nonlinear and non-Gaussian behavior of the associated Bayesian inference problem. One finds that the performance of the FPF improves as the ensemble size increases while the second-order method and the EnKBF suffer from a systematic bias. We also find that the McKean–Vlasov SDE formulation (36) behaves rather well over the full

Table 1 Informative prior: We display the ensemble mean averaged over $L = 1000$ repeated experiments for the four proposed methods as a function of the ensemble size M

Ensemble size/method	$M = 50$	$M = 100$	$M = 200$	$M = 400$
McKean–Vlasov SDE	$\begin{pmatrix} -3.34 \\ -3.39 \\ 3.22 \end{pmatrix}$	$\begin{pmatrix} -3.35 \\ -3.39 \\ 3.21 \end{pmatrix}$	$\begin{pmatrix} -3.35 \\ -3.39 \\ 3.21 \end{pmatrix}$	$\begin{pmatrix} -3.35 \\ -3.39 \\ 3.22 \end{pmatrix}$
FPF	$\begin{pmatrix} -3.31 \\ -3.35 \\ 3.20 \end{pmatrix}$	$\begin{pmatrix} -3.33 \\ -3.37 \\ 3.21 \end{pmatrix}$	$\begin{pmatrix} -3.33 \\ -3.37 \\ 3.21 \end{pmatrix}$	$\begin{pmatrix} -3.33 \\ -3.37 \\ 3.21 \end{pmatrix}$
Second-order	$\begin{pmatrix} -3.36 \\ -3.40 \\ 3.22 \end{pmatrix}$	$\begin{pmatrix} -3.36 \\ -3.41 \\ 3.22 \end{pmatrix}$	$\begin{pmatrix} -3.36 \\ -3.41 \\ 3.22 \end{pmatrix}$	$\begin{pmatrix} -3.36 \\ -3.41 \\ 3.22 \end{pmatrix}$
EnKBF	$\begin{pmatrix} -3.27 \\ -3.31 \\ 3.20 \end{pmatrix}$	$\begin{pmatrix} -3.27 \\ -3.32 \\ 3.20 \end{pmatrix}$	$\begin{pmatrix} -3.27 \\ -2.31 \\ 3.19 \end{pmatrix}$	$\begin{pmatrix} -3.27 \\ -3.31 \\ 3.19 \end{pmatrix}$

The reference value from an ALDI simulation is $(-3.32, -3.36, 3.20)^T$. We find that all methods reproduce the reference solution provided by ALDI within small variations

Table 2 Informative prior: We display the spectral norm of the final ensemble covariance matrices averaged over $L = 1000$ repeated experiments for the four proposed methods as a function of the ensemble size M

Ensemble size/method	$M = 50$	$M = 100$	$M = 200$	$M = 400$
McKean–Vlasov SDE	0.89	0.83	0.80	0.78
FPF	0.98	0.90	0.86	0.83
Second-order	0.75	0.75	0.74	0.74
EnKBF	0.80	0.79	0.78	0.78

The ALDI reference value is 0.82. We find that the FPF slightly overestimates the variance for smaller ensemble sizes, while both the EnKBF and the second-order formulation slightly underestimate it across all ensemble sizes

range of ensemble sizes. We conclude that all tested methods are able to qualitatively reproduce the exact parameter value (52) even in the case of the less informative prior. The results also indicate that the posterior distribution is close to Gaussian, while the intermittent PDFs π_τ must deviate significantly from a Gaussian distribution.

In terms of computational cost, the EnKBF and the second-order method perform comparable and are the least expensive by far for all ensemble sizes. While the McKean–Vlasov SDE is twice as expensive as the EnKBF at $M = 50$, this ratio goes up to a factor of five at $M = 400$. The FPF, on the other hand, is ten times as expensive as the EnKFB at $M = 50$ which increases steeply to a factor of one hundred at $M = 400$.

We now consider a high-dimensional extension of the previous example, for which we further investigate the performance of the EnKBF (33) in terms of the tamed time-stepping method (35) for ensemble sizes $M \leq D$, localisation via dropouts, and data sub-sampling.

Table 3 Less informative prior: We display the ensemble mean averaged over $L = 1000$ repeated experiments for the four proposed methods as a function of the ensemble size M

Ensemble size/method	$M = 50$	$M = 100$	$M = 200$	$M = 400$
McKean–Vlasov SDE	$\begin{pmatrix} -2.61 \\ -2.65 \\ 2.15 \end{pmatrix}$	$\begin{pmatrix} -2.60 \\ -2.64 \\ 2.15 \end{pmatrix}$	$\begin{pmatrix} -2.59 \\ -2.62 \\ 2.15 \end{pmatrix}$	$\begin{pmatrix} -2.59 \\ -2.62 \\ 2.15 \end{pmatrix}$
FPF	$\begin{pmatrix} -2.16 \\ -2.15 \\ 1.64 \end{pmatrix}$	$\begin{pmatrix} -2.28 \\ -2.30 \\ 1.83 \end{pmatrix}$	$\begin{pmatrix} -2.39 \\ -2.41 \\ 1.95 \end{pmatrix}$	$\begin{pmatrix} -2.44 \\ -2.47 \\ 2.00 \end{pmatrix}$
Second-order	$\begin{pmatrix} -2.29 \\ -2.32 \\ 1.82 \end{pmatrix}$	$\begin{pmatrix} -2.30 \\ -2.32 \\ 1.83 \end{pmatrix}$	$\begin{pmatrix} -2.31 \\ -2.33 \\ 1.83 \end{pmatrix}$	$\begin{pmatrix} -2.31 \\ -3.34 \\ 1.84 \end{pmatrix}$
EnKBF	$\begin{pmatrix} -2.14 \\ -2.16 \\ 1.73 \end{pmatrix}$	$\begin{pmatrix} -2.16 \\ -2.18 \\ 1.74 \end{pmatrix}$	$\begin{pmatrix} -2.17 \\ -2.19 \\ 1.75 \end{pmatrix}$	$\begin{pmatrix} -2.17 \\ -2.19 \\ 1.76 \end{pmatrix}$

We find that all methods qualitatively reproduce the reference value $(-2.56, -2.59, 2.15)^T$ provided by ALDI. While the second-order method and the EnKBF suffer from a systematic bias, the FPF approaches the reference solution as the ensemble size increases. The McKean–Vlasov SDE formulation yields good approximations for all ensemble sizes

Table 4 Less informative prior: We display the spectral norm of the final ensemble covariance matrices averaged over $L = 1000$ repeated experiments for the four proposed methods as a function of the ensemble size M

Ensemble size/method	$M = 50$	$M = 100$	$M = 200$	$M = 400$
McKean–Vlasov SDE	1.36	1.14	1.03	0.97
FPF	3.61	2.58	2.01	1.66
Second-order	0.46	0.45	0.45	0.45
EnKBF	0.60	0.59	0.59	0.59

The ALDI reference value is 1.18. We find that the FPF overestimates the variance, while both the EnKBF and the second-order formulation underestimate it. While the FPF appears to approach the correct reference value for increased ensemble sizes, both the EnKBF and the second-order filter appear to suffer from a systematic bias. The McKean–Vlasov SDE approach leads to a slightly degraded performance in terms of ensemble variations as the ensemble size increases

Example 2 We consider the simple feature map

$$\phi_x = x,$$

that is $J = D$ such that

$$y_x(\theta) = \sigma(\theta^T x)$$

for $D = 50$. The data is generated by first drawing a $\theta = \theta_{\text{ref}}$ from the standard normal distribution in \mathbb{R}^D . Next, $N = 1000$ data points are generated by randomly drawing $x_n, n = 1, \dots, N$, again from a standard normal distribution in \mathbb{R}^D . Those points are assigned the label $t_n = 1$ with probability $y_{x_n}(\theta_{\text{ref}})$, and $t_n = 0$ otherwise.

Table 5 We display the l_2 -difference between the posterior ensemble means and the true parameter value averaged over $L = 1000$ repeated experiments and its standard deviation for three different implementations of the EnKBF as a function of the ensemble size M

Ensemble size/method	$M = 20$	$M = 60$	$M = 100$
EnKBF	6.26 ± 0.75	2.67 ± 0.68	1.69 ± 0.48
EnKBF (with dropout)	1.29 ± 0.29	1.28 ± 0.30	1.39 ± 0.46
EnKBF (with dropout & mini-batch)	2.14 ± 0.46	1.42 ± 0.25	1.35 ± 0.21

The reference value provided by ALDI is 1.10. We find that localisation by dropout dramatically improves the behavior of the EnKBF especially for small ensemble sizes. Using mini-batches reduces the approximation quality slightly while providing significant computational savings

Table 6 We provide the spectral norm of the computed posterior covariance matrices in the setting described in Table 5

Ensemble size/method	$M = 20$	$M = 60$	$M = 100$
EnKBF	0.014 ± 0.004	0.058 ± 0.012	0.097 ± 0.015
EnKBF (with dropout)	0.043 ± 0.012	0.088 ± 0.022	0.109 ± 0.025
EnKBF (with dropout and mini-batch)	0.041 ± 0.012	0.084 ± 0.021	0.105 ± 0.025

The reference value provided by ALDI is 0.22. We find that both all three EnKBF implementations underestimate the variance. However, dropout localisation increases the ensembles spread which in turn improves the estimated ensemble means as displayed in Table 5

We use the EnKBF tamed time-stepping method (35) with step size $\Delta\tau = 1/200$ for parameter inference. The initial ensemble $\theta_0^{(i)}$, $i = 1, \dots, M$ is drawn from the prior Gaussian distribution with mean $m_{\text{prior}} = 0$ and the covariance matrix Σ_{prior} equal to the identity matrix. We vary the ensemble size M between $M = 20$ and $M = 100$ and use dropout localisation for the computation of the empirical covariance matrices $\hat{\Sigma}_{\tau_n}$ following (34). More specifically, an entry of $\theta_{\tau_n}^{(i)}$ is set to zero if $\eta < 0.5$ where the η 's are i.i.d. uniform random variables from the interval $[0, 1]$. We report in Table 5 the l_2 -difference between θ_{ref} and the ensemble mean \hat{m}_τ at final time $\tau = 1$. We also state the spectral norm of $\hat{\Sigma}_\tau$ at the final time in Table 6. In a second set of experiments, we apply mini-batching to the localised EnKBF with random batches of size $N' = 100$. We also compare the results to those from the standard EnKBF formulation (32). Furthermore, ALDI is run for a sample size of $M = 100$ in order to provide the numerical benchmark values listed in the captions of Tables 5 and 6. All results have been averaged over $L = 1000$ independent simulations. We also report the standard deviations in addition to the averaged values. Our numerical results indicate that localisation by dropout is effective for small ensemble sizes, and that mini-batching does not degrade performance while providing significant computational savings. As for the previous example, the EnKBF is able to capture the posterior mean quite well while the posterior variances are systematically underestimated. Ensemble inflation, as widely used for the EnKF [8,29], could help to further improve the variance estimates. We found that our results are insensitive to the choice of the step size $\Delta\tau$ in (35). However changing the threshold value $\mu = 0.5$ in the dropout localisation (34) affects the results significantly. Smaller values of μ lead in particular to reduced

Table 7 Same results as displayed in Table 5 except that the dropout rate has been reduced from $\mu = 0.5$ to $\mu = 0.2$. While the estimation errors increase for the smallest ensemble sizes the reduced rate improves the performance for larger ensemble sizes

ensemble size/method	$M = 20$	$M = 60$	$M = 100$
EnKBF (with dropout)	3.30 ± 1.01	1.26 ± 0.23	1.12 ± 0.36
EnKBF (with dropout and mini-batch)	3.39 ± 0.86	1.38 ± 0.25	1.19 ± 0.27

l_2 -differences for larger ensemble sizes, M , while being less effective for smaller M . See Table 7.

7 Generalisations and Extensions

In this section, we discuss possible extensions of the proposed methods to nonlinear logistic regression including derivative-free implementations, multi-class logistic regression, and sigmoidal Cox processes.

7.1 Nonlinear Logistic Regression

We generalise linear logistic regression from Sect. 2 to the more general logistic regression problem with $y_x(\theta) := \sigma(f_x(\theta))$ for given functions $f_x : \mathbb{R}^D \rightarrow \mathbb{R}$. Hence the gradient $\nabla_{\theta} y_x(\theta)$ becomes

$$\nabla_{\theta} y_x(\theta) = \phi_x(\theta) (1 - y_x(\theta)) y_x(\theta),$$

where the previously considered feature maps ϕ_x now also depend on the parameters θ , that is,

$$\phi_x(\theta) := \nabla_{\theta} f_x(\theta).$$

The gradient of Ψ_{data} is thus explicitly given by

$$\nabla_{\theta} \Psi_{\text{data}}(\theta) = \sum_{n=1}^N \phi_{x_n}(\theta) (y_{x_n}(\theta) - t_n)$$

and the Hessian matrix of second-order derivatives becomes

$$D_{\theta}^2 \Psi_{\text{data}}(\theta) = \sum_{n=1}^N \left\{ \phi_{x_n}(\theta) y_{x_n}(\theta) (1 - y_{x_n}(\theta)) \phi_{x_n}(\theta)^T + D_{\theta} \phi_{x_n}(\theta) (y_{x_n}(\theta) - t_n) \right\}.$$

We now discuss a derivative-free implementation of the EnKBF (32) for nonlinear logistic regression. The key difference to linear logistic regression is that the matrix $\Phi \in \mathbb{R}^{D \times N}$ now depends on θ and is given by

$$\Phi(\theta) = (\nabla_{\theta} f_{x_1}(\theta), \dots, \nabla_{\theta} f_{x_N}(\theta)).$$

In order to avoid the computation of the gradients of f_x , one replaces $\Sigma_{\tau} \Phi(\theta_{\tau})$ in (32) by the empirical correlation matrix

$$C_{\tau} := \frac{1}{M-1} \sum_{i=1}^M (\theta_{\tau}^{(i)} - \hat{m}_{\tau}) \left(f_{x_1}(\theta_{\tau}^{(i)}), \dots, f_{x_N}(\theta_{\tau}^{(i)}) \right) \in \mathbb{R}^{D \times N}$$

between θ_{τ} and $f_{x_n}(\theta_{\tau})$ for $n = 1, \dots, N$. See [8] for the basic idea in the context of the EnKF and also [13,17] for the specific application of this idea in the context of machine learning. See also the recent work [24,30] on tuneable derivative-free approximations. We also note that this substitution becomes exact in the case of linear functions $f_x(\theta)$. Such a linear approximation is, for example, justified for functions $f_x(\theta)$ defined by sufficiently wide neural networks [18].

7.2 Multi-class Logistic Regression

The extension from two-class to multi-call classification is straightforward, see [5]. The posterior probabilities for L classes $C_l, l = 1, \dots, L$, are given by

$$\pi(C_l | \phi_x) = \frac{\exp(a_{x,l})}{\sum_{j=1}^L \exp(a_{x,j})},$$

where the activation is given by

$$a_{x,l} := \theta_l^T \phi_x.$$

The adjustable parameters are now $\theta = (\theta_1^T, \dots, \theta_L^T)^T \in \mathbb{R}^{LD}$. The negative log-likelihood function becomes

$$\Psi_{\text{data}}(\theta) = - \sum_{n=1}^N \sum_{l=1}^L t_{nl} \ln y_{nl}(\theta)$$

with

$$y_{nl}(\theta) := \frac{\exp(a_{x_n,l})}{\sum_{j=1}^L \exp(a_{x_n,j})}$$

and data $(x_n, \{t_{nl}\}_{l=1}^L)$, where $t_{nl} \in \{0, 1\}$ subject to $\sum_{l=1}^L t_{nl} = 1$ for all $n = 1, \dots, N$. The gradients satisfy

$$\nabla_{\theta_l} \Psi_{\text{data}}(\theta) = \sum_{n=1}^N (y_{nl} - t_{nl}) \phi_n,$$

$l = 1, \dots, L$. The Hessian matrix of second-order derivatives got $D \times D$ block entries

$$D_{\theta_l} D_{\theta_j} \Psi_{\text{data}}(\theta) = \sum_{n=1}^N y_{nl} (\delta_{lj} - y_{nj}) \phi_n \phi_n^T$$

with $\delta_{lj} = 1$ if $l = j$ and $\delta_{lj} = 0$ otherwise. We conclude that these formulas closely resemble the corresponding expressions for binary classification and, hence, all previously considered algorithms easily generalise to multi-class regression.

7.3 Sigmoidal Cox Processes

We consider Cox processes (CPs) on a domain $\mathcal{X} \subset \mathbb{R}^J$. Following [1], the CP is parametrised by an intensity function $\lambda_x : \mathcal{X} \rightarrow \mathbb{R}^+$ which we represent using a sigmoidal random feature model, that is,

$$\lambda_x(\theta) = \lambda^* \sigma(\theta^T \phi_x),$$

where $\lambda^* > 0$ provides the upper threshold value for the intensity function. Note that a Gaussian process is used by [1] while we rely here on the closely related random feature maps ϕ_x [21], which easier integrate into the homotopy approach and the EnKF in particular [12].

See [7] for an efficient Bayesian inference for CPs using variable augmentation. Here we suggest an alternative approach using an associated unbiased estimator of the negative log-likelihood in the forward Euler time-stepping (20).

More specifically, given N events $x_n, n = 1, \dots, N$, the negative log-likelihood function is provided by

$$\Psi_{\text{data}}(\theta) = \int_{\mathcal{X}} \lambda_x(\theta) dx - \sum_{n=1}^N \ln \lambda_{x_n}(\theta).$$

Following the notation of Sect. 3.2 and using I uniformly distributed samples $\hat{x}_i, i = 1, \dots, I$, on \mathcal{X} , we obtain the random unbiased estimator

$$\Psi_{\text{data}}^\gamma(\theta) = \frac{\lambda^*}{I} \sum_{i=1}^I y_{\hat{x}_i}(\theta) - \sum_{n=1}^N \ln y_{x_n}(\theta) - N \ln \lambda^* \tag{53}$$

for this intractable likelihood Ψ_{data} . Here we also used the previously introduced notation (4), that is, $\lambda_x(\theta) = \lambda^* y_x(\theta)$.

We now apply the homotopy approach (10) to the associated Bayesian inference problem. Following the discussion in Sect. 3.2, the evolution converges to the correct posterior if in each time step (20) one uses a different realisation γ_{τ_k} of the unbiased estimator (53) instead of Ψ_{data} in (9) and takes the limit $\Delta\tau \rightarrow 0$.

Example 3 Let us briefly discuss a related problem in a strictly Gaussian setting. We start from the intractable negative log-likelihood function

$$\Psi_{\text{data}}(\theta) := \frac{1}{2} \int_0^1 (\theta^T \phi_x - t)^2 dx.$$

Its gradient is provided by

$$\nabla_{\theta} \Psi_{\text{data}}(\theta) = \int_0^1 \phi_x (\theta^T \phi_x - t) dx$$

and the associated mean-field EnKBF becomes

$$\frac{d}{d\tau} \theta_t = -\frac{1}{2} \Sigma_{\tau} \int_0^1 \phi_x \left((\theta_{\tau} + m_{\tau})^T \phi_x - 2t \right) dx.$$

So far, everything is exact and no approximation has been made. We now replace the integral by its Monte Carlo approximation using I randomly sampled $\hat{x}_{i,k} \sim U[0, 1]$, $i = 1, \dots, I$, and introduce a time step $\Delta\tau$ leading to

$$\theta_{\tau_{k+1}} = \theta_{\tau_k} - \frac{\Delta\tau}{2I} \Sigma_{\tau_k} \sum_{i=1}^I \phi_{\hat{x}_{i,k}} \left((\theta_{\tau_k} + m_{\tau_k})^T \phi_{\hat{x}_{i,k}} - 2t \right),$$

where we re-sample the $\hat{x}_{i,k}$'s in each time-step τ_k , $k = 0, \dots, K - 1$. A theoretical investigation now boils down to studying the limit $\Delta\tau \rightarrow 0$, which can be performed along the modified equation analysis of [20]. We note that the approximation error can be reduced by either increasing I or decreasing the step size $\Delta\tau$.

8 Conclusions

We have presented affine-invariant extensions of the popular EnKF and related methods such as the FPF to logistic regression. In addition, we have proposed a novel SDE-based affine-invariant and derivative-free sampling method which utilises the robust ensemble transform approach for Bayesian inference as an inner time-stepping method. We have also introduced a number of algorithmic choices such as localisation via dropout, linearly implicit time-stepping methods, and mini-batching the data, which help to improve the efficiency of the proposed methods. While we have only considered two simple numerical examples, future work will consider more challenging applications such as the training of the *Sentence Gestalt* model of natural language

comprehension [25]. A numerical exploration of the affine-invariant stochastic homotopy formulation (18) would also be of practical and theoretical interest.

A number of theoretical questions also remain to be answered such as the convergence of the ensemble methods to their mean-field limits, the impact of the various approximations on the computed posterior representations, and the ergodicity and convergence to equilibrium properties of (36) and (38), respectively. Finally, an embedding of the ensemble transform Langevin dynamics into a Markov chain Monte Carlo framework would allow for an exact sampling from the posterior.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

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

Appendix

We provide an analysis of the SDE

$$d\theta_\tau = -\Sigma_\tau \left\{ \nabla_\theta V_\tau(\theta_\tau) d\tau - (\Delta\tau \Omega_\tau)^{1/2} dW_\tau \right\}, \tag{54}$$

which is obtained from the modified SDE (21) by setting $\Omega_\tau(\theta) = \Omega_\tau$, that is, by ignoring the state dependence of the diffusion term, and using the EnKBF drift term from linear regression [3], that is,

$$\nabla_\theta V_\tau(\theta) = \frac{1}{2} G^T \Gamma^{-1} (G\theta + Gm_\tau - 2t).$$

Conditioned on W_τ , the evolution equation for the (conditional) mean becomes

$$dm_\tau = -\Sigma_\tau \left\{ G^T \Gamma^{-1} (Gm_\tau - t) d\tau - (\Delta\tau \Omega_\tau)^{1/2} dW_\tau \right\} \tag{55}$$

and that for the (conditional) covariance matrix

$$\frac{d}{d\tau} \Sigma_\tau = -\Sigma_\tau G^T \Gamma^{-1} G \Sigma_\tau. \tag{56}$$

We can further decompose (55) into an equation for the mean, which we denote by μ_τ , and the covariance P_τ . We obtain the two evolution equations

$$\frac{d}{d\tau}\mu_\tau = -\Sigma_\tau G^T \Gamma^{-1} (G\mu_\tau - t) \quad (57)$$

and

$$\frac{d}{d\tau}P_\tau = -\Sigma_\tau G^T \Gamma^{-1} G P_\tau - P_\tau G^T \Gamma^{-1} G \Sigma_\tau + \Delta\tau \Sigma_\tau \Omega_\tau \Sigma_\tau, \quad (58)$$

respectively. The initial conditions are $\mu_0 = m_0$ and $P_0 = 0$.

We note that (57) and (56) are equivalent to the equations arising from the standard Kalman–Bucy filter with solutions

$$\mu_\tau = \mu_0 - K_\tau (G\mu_0 - t)$$

and

$$\Sigma_\tau = \Sigma_0 - K_\tau G \Sigma_0,$$

respectively, where K_τ denotes the Kalman gain matrix

$$K_\tau = \tau \Sigma_0 G^T (\Gamma + \tau G \Sigma_0 G^T)^{-1}.$$

If we now consider the fully observed case, that is $G = I$, under small measurement error covariance $\Gamma = \epsilon I$, $\epsilon \ll 1$, then we find that

$$\mu_1 \approx \theta_{\text{MLE}} = t, \quad \Sigma_1 \approx \Gamma = \epsilon I, \quad P_1 < \frac{\Delta\tau}{2} \Sigma_1 \Omega_1 \Sigma_1 \approx \epsilon^2 \frac{\Delta\tau}{2} \Omega_1.$$

If we define

$$\Omega_\tau = \pi_\tau [\Omega_\tau(\theta)]$$

for simplicity, where π_τ is the exact filtering distribution for the Kalman–Bucy filter, then $\Omega_0 = \mathcal{O}(\epsilon^{-2})$ while $\Omega_1 = \mathcal{O}(\epsilon^{-1})$. Hence, the variance, P_τ , of the posterior conditional mean, m_τ , is much smaller than the posterior (as well as the frequentist) variance, Σ_τ , at $\tau = 1$ provided the step-size $\Delta\tau$ in (54) is chosen small enough.

References

1. Adams, R., Murray, I., MacKay, D.: Tractable nonparametric Bayesian inference in Poisson processes with Gaussian process intensities. In: Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09, pp. 9–16. Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1553374.1553376>
2. Amezcua, J., Kalnay, E., Ide, K., Reich, S.: Ensemble transform Kalman–Bucy filters. Q.J.R. Meteor. Soc. **140**, 995–1004 (2014). <https://doi.org/10.1002/qj.2186>
3. Bergemann, K., Reich, S.: A localization technique for ensemble Kalman filters. Quarterly Journal of the Royal Meteorological Society **136**, 701–707 (2010). <https://doi.org/10.1002/qj.591>

4. Beskos, A., Crisan, D., Jasra, A.: On the stability of sequential Monte Carlo methods in high dimensions. *The Annals of Applied Probability* **24**, 1396–1445 (2014). <https://doi.org/10.1214/13-AAP951>
5. Bishop, C.M.: *Pattern Recognition and Machine Learning*. Springer, New York (2006)
6. Daum, F., Huang, J., Noushin, A.: Exact particle flow for nonlinear filters. In: I. Kadar (ed.) *Signal Processing, Sensor Fusion, and Target Recognition XIX*, vol. 7697, pp. 92–110. International Society for Optics and Photonics, SPIE (2010). <https://doi.org/10.1117/12.839590>
7. Donner, C., Opper, M.: Efficient Bayesian inference of sigmoidal Gaussian Cox processes. *J. Mach. Learn. Research* **19**(67), 1–34 (2018). URL <http://jmlr.org/papers/v19/17-759.html>
8. Evensen, G.: *Data Assimilation: The Ensemble Kalman Filter*, second edn. Springer, Berlin (2009). <https://doi.org/10.1007/978-3-642-03711-5>
9. Garbuno-Inigo, A., Hoffmann, F., Li, W., Stuart, A.: Interacting Langevin diffusions: Gradient structure and ensemble Kalman sampler. *SIAM J. Appl. Dyn. Syst.* **19**(1), 412–441 (2020). <https://doi.org/10.1137/19M1251655>
10. Garbuno-Inigo, A., Nüsken, N., Reich, S.: Affine invariant interacting Langevin dynamics for Bayesian inference. *SIAM J. Appl. Dyn. Syst.* **19**, 1633–1658 (2020). <https://doi.org/10.1137/19M1304891>
11. Goodman, J., Weare, J.: Ensemble samplers with affine invariance. *Comm. Appl. Math. and Comput. Science* **5**, 65–80 (2010). <https://doi.org/10.2140/camcos.2010.5.65>
12. Gottwald, G., Reich, S.: Supervised learning from noisy observations: Combining machine-learning techniques with data assimilation. *Physica D* **423**, 132911 (2021). <https://doi.org/10.1016/j.physd.2021.132911>
13. Haber, E., Lucka, F., Ruthotto, L.: Never look back—A modified EnKF method and its application to the training of neural networks without back propagation (2018). [arXiv:1805.08034](https://arxiv.org/abs/1805.08034)
14. Heber, F., Ž. Trst'anova, Leimkuhler, B.: Posterior sampling strategies based on discretized stochastic differential equations for machine learning applications. *Journal of Machine Learning Research* **21**(228), 1–33 (2020). URL <http://jmlr.org/papers/v21/19-339.html>
15. Huang, D., Schneider, T., Stuart, A.: Uncented Kalman inversion (2021). [arXiv:2102.01580](https://arxiv.org/abs/2102.01580)
16. Kass, R., Raftery, A.: Bayes factors. *Journal of the American Statistical Association* **90**, 773–795 (1995). <https://doi.org/10.1080/01621459.1995.10476572>
17. Kovachki, N., Stuart, A.: Ensemble Kalman inversion: A derivative-free technique for machine learning tasks. *Inverse Probl.* **35**(9), 095005 (2019). <https://doi.org/10.1088/1361-6420/ab1c3a>
18. Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., Pennington, J.: Wide neural networks of any depth evolve as linear models under gradient descent. *Journal of Statistical Mechanics: Theory and Experiment* **2020**(12), 124002 (2020). <https://doi.org/10.1088/1742-5468/abc62b>
19. Leimkuhler, B., Matthews, C., Weare, J.: Ensemble preconditioning for Markov chain Monte Carlo simulation. *Statistics and Computing* **28**, 277–290 (2018). <https://doi.org/10.1007/s11222-017-9730-1>
20. Li, Q., Tai, C., E, W.: Stochastic modified equations and dynamics of stochastic gradient algorithms I: Mathematical foundations. *Journal of Machine Learning Research* **20**, 1–47 (2019). URL <http://jmlr.org/papers/v20/17-526.html>
21. Liu, F., Huang, X., Chen, Y., Suykens, A.: Random features for kernel approximation: A survey on algorithms, theory and beyond (2020). [arXiv:2004.11154](https://arxiv.org/abs/2004.11154)
22. Nadeem, A., Potthast, R.: Transformed and generalized localization for ensemble methods in data assimilation. *Mathematical Methods in the Applied Sciences* **39**, 619–634 (2016). <https://doi.org/10.1002/mma.3496>
23. Opper, M.: A Bayesian approach to online learning. In: D. Saad (ed.) *Online Learning in Neural Networks*, pp. 363–378. Cambridge University Press (1999). <https://doi.org/10.1017/CBO9780511569920.017>
24. Pavliotis, G., Stuart, A., Vaes, U.: Derivative-free Bayesian inversion using multiscale dynamics (2021). [arXiv:2102.00540](https://arxiv.org/abs/2102.00540)
25. Rabovsky, M., McClelland, J.: Quasi-compositional mapping from form to meaning: A neural network based approach to capturing neural responses during human language comprehension. *Phil. Trans. Royal Soc. B: Biological Sciences* **375**, 20190313 (2019). <https://doi.org/10.1098/rstb.2019.0313>
26. Reich, S.: A dynamical systems framework for intermittent data assimilation. *BIT Numer. Math.* **51**, 235–249 (2011). <https://doi.org/10.1007/s10543-010-0302-4>
27. Reich, S.: A nonparametric ensemble transform method for Bayesian inference. *SIAM J. Sci. Comput.* **35**, A2013–A2024 (2013). <https://doi.org/10.1137/130907367>

28. Reich, S., Cotter, C.: Ensemble filter techniques for intermittent data assimilation. In: M. Cullen, F.M. A., S. Kindermann, R. Scheichl (eds.) *Large Scale Inverse Problems. Computational Methods and Applications in the Earth Sciences, Radon Ser. Comput. Appl. Math.*, vol. 13, pp. 91–134. Walter de Gruyter, Berlin (2013). <https://doi.org/10.1515/9783110282269>
29. Reich, S., Cotter, C.: *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press, New York (2015). <https://doi.org/10.1017/CBO9781107706804>
30. Reich, S., Weissmann, S.: Fokker–Planck particle systems for Bayesian inference: Computational approaches. *SIAM/ASA J. Uncertainty Quantification* **9**, 446–482 (2021). <https://doi.org/10.1137/19M1303162>
31. Sakov, P., Oliver, D., Bertino, L.: An iterative EnKF for strongly nonlinear systems. *Monthly Weather Review* **140**, 1988–2004 (2012). <https://doi.org/10.1175/MWR-D-11-00176.1>
32. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* **15**(56), 1929–1958 (2014). URL <http://jmlr.org/papers/v15/srivastava14a.html>
33. Sullivan, T.: *Introduction to Uncertainty Quantification. Texts in Applied Mathematics*, vol. 63. Springer (2015). <https://doi.org/10.1007/978-3-319-23395-6>
34. Taghvaei, A., Mehta, P., Meyn, S.: Diffusion map-based algorithm for gain function approximation in the feedback particle filter. *SIAM/ASA J. Uncertain. Quantif.* **8**(3), 1090–1117 (2020). <https://doi.org/10.1137/19M124513X>
35. Taghvaei, A., de Wiljes, J., Mehta, P., Reich, S.: Kalman filter and its modern extensions for the continuous-time nonlinear filtering problem. *ASME. J. Dyn. Sys., Meas., Control.* **140**, 030904–030904–11 (2017). <https://doi.org/10.1115/1.4037780>
36. Tödter, J., Kirchgessner, P., Nerger, L., Ahrens, B.: Assessment of a nonlinear ensemble transform filter for high-dimensional data assimilation. *Monthly Weather Review* **144**, 409–427 (2016). <https://doi.org/10.1175/MWR-D-15-0073.1>
37. Wormell, C., Reich, S.: Spectral convergence of diffusion maps: Improved error bounds and an alternative normalisation. *SIAM J. Numer. Anal.* **59**, 1687–1734 (2021). <https://doi.org/10.1137/20M1344093>
38. Yang, T., Blom, H.A.P., Mehta, P.G.: The Continuous-Discrete Time Feedback Particle Filter. In: 2014 American Control Conference, pp. 648–653. American Automatic Control Council (2014). <https://doi.org/10.1109/ACC.2014.6859259>
39. Yang, T., Mehta, P.G., Meyn, S.P.: Feedback particle filter. *IEEE Trans. Automat. Control* **58**, 2465–2480 (2013). <https://doi.org/10.1109/TAC.2013.2258825>

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