

# Subsampling in ensemble Kalman inversion

Matei Hanu<sup>1,\*</sup> , Jonas Latz<sup>2</sup>  and Claudia Schillings<sup>1</sup>

<sup>1</sup> Fachbereich Mathematik und Informatik, Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany

<sup>2</sup> Maxwell Institute for Mathematical Sciences and School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom

E-mail: [matei.hanu@fu-berlin.de](mailto:matei.hanu@fu-berlin.de)

Received 22 February 2023; revised 30 June 2023

Accepted for publication 11 July 2023

Published 28 July 2023



CrossMark

## Abstract

We consider the ensemble Kalman inversion (EKI) which has been recently introduced as an efficient, gradient-free optimisation method to estimate unknown parameters in an inverse setting. In the case of large data sets, the EKI becomes computationally infeasible as the data misfit needs to be evaluated for each particle in each iteration. Here, randomised algorithms like stochastic gradient descent have been demonstrated to successfully overcome this issue by using only a random subset of the data in each iteration, so-called subsampling techniques. Based on a recent analysis of a continuous-time representation of stochastic gradient methods, we propose, analyse, and apply subsampling-techniques within EKI. Indeed, we propose two different subsampling techniques: either every particle observes the same data subset (single subsampling) or every particle observes a different data subset (batch subsampling).

Keywords: inverse problems, stochastic optimisation, data subsampling, particle system, piecewise-deterministic Markov process

(Some figures may appear in colour only in the online journal)

\* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

## 1. Introduction

A large variety of physical, biological and social systems and processes have been described by mathematical models. Those models can be used to analyse and predict the behaviour of the associated processes and systems. In case, a model shall be employed to describe a particular system, the model needs to be calibrated with respect to observation of that particular system. This calibration process that fits the model to data is often called *inversion*. Inversion forms the basis for, e.g. numerical weather prediction, medical image processing, and many machine learning methods. Several inversion techniques have been proposed and studied: we often distinguish variational/optimisation-based approaches and Bayesian/statistical approaches. Throughout this work, we consider a class of methods that can be seen as being in between those two approaches: the ensemble Kalman inversion (EKI) framework going back to [20, 31]. EKI is based on an Ensemble Kalman–Bucy Filter that is iteratively applied to solve an inverse problem. In the linear setting the resulting algorithm is usually given in the form of a preconditioned gradient flow, that is an ordinary differential equation describing the dynamics of an ensemble of particles.

EKI becomes computationally infeasible, if the considered amount of data is too large: the data cannot be stored in the memory at once and, thus, EKI is not applicable. The same problem arises also in other traditional optimisation algorithms, like gradient descent or the Gauss–Newton method. In the past decades, randomised algorithms that optimise in each time step only with respect to a subsample of the data set have become popular. A subsample is a (often randomly chosen) subset of the considered data set. The foundation for all of these stochastic optimisation algorithms is the stochastic gradient descent (SGD) algorithm going back to [28]. Stochastic gradient descent and its variants have become especially popular in the machine learning community.

The idea of randomised subsampling in the EKI framework has been proposed in [22], where it is, indeed, applied to train a neural network. There, the subsampling has been introduced after discretising the preconditioned gradient flow, but no analysis has been presented. A recent work by [23] explains how subsampling can be represented in continuous-time settings and how these can be analysed. In the present work, we aim at using this theory in the context of EKI to analyse subsampling methodology at the ODE level.

### 1.1. Literature overview

Since its introduction in [15] the ensemble Kalman filter (EnKF) has been widely used in both inverse problems as well as data assimilation problems. The EnKF is very appealing for many applications due to its straightforward implementation and robustness w.r. to small ensemble sizes [3, 4, 18–20, 24]. Stability results can be found in [32, 33]. Convergence analysis based on the continuous time limit of the EKI has been developed in [5, 6, 9, 30, 31]. However, to obtain convergence results in the parameter space some form of regularisation is usually needed. We mainly consider Tikhonov regularisation which was analysed for the EKI in [12] for example. Recently there has been further analysis on Tikhonov regularisation for the stochastic EKI as well as adaptive Tikhonov strategies to improve the original variant [35]. Considering large ensemble sizes, an analysis of the mean-field limit is presented in [10, 14].

A historical overview of the Kalman filter and some of its extensions can be found [10].

After their introduction by Robbins and Monro [28], the stochastic gradient descent method has in the recent past been further analysed by, e.g. [8]. Stochastic gradient descent is often

computationally advantageous compared to normal gradient descent [26] due to computational efficiency as well as being able to escape local minimisers in non-convex optimisation problems [13, 34]. As mentioned earlier, the theory employed in this work is based on the continuous-time analysis of stochastic gradient descent by Latz [23] that was further generalised in [21], but is somewhat orthogonal to the diffusion-based continuous-time analysis of SGD of, e.g. [25].

## 1.2. Contributions and outline

In the following, we focus on the case, where the data misfit is computationally infeasible due to large data. Inspired by the success story of randomised gradient descent methods, we introduce subsampling strategies to EKI to ensure feasibility of the method also in the large data regime. We summarise our contribution below:

1. We introduce two subsampling techniques for EKI: single subsampling and batch subsampling.
2. We present an analysis of the subsampling techniques for linear forward operators, in particular we analyse stability of the subsampling schemes and give conditions under which we obtain asymptotic stability. Indeed, the resulting dynamical system approximates the EKI solution.
3. We illustrate our results with two examples: estimation of the source term for an elliptic partial differential equation and estimation of the diffusion coefficient for a parabolic partial differential equation.

Whilst analysing the subsampling EKI, we generalise some results from [23] to more general flows. These generalisations may be of independent interest.

This work is structured as follows. We introduce problem setting and EKI methods in section 2. We discuss stochastic approximations of certain flows in general and the subsampling in EKI in particular in section 3; before analysing them in section 4. We show numerical examples in section 5 and conclude the work in section 6.

## 2. Problem setting and mathematical background

Let  $(\Omega, \mathcal{A}, \mathbb{P})$  be a probability space,  $X$  be a separable Hilbert space and  $Y := \mathbb{R}^{N_{\text{obs}}}$ , with  $N_{\text{obs}} \in \mathbb{N} = \{1, 2, \dots\}$ . We will refer to  $X$  as *parameter space* and to  $Y$  as *data space*. Let now  $n, m \in \mathbb{N}$ . We sometimes associate finite-dimensional spaces  $\mathbb{R}^n$  with the basic inner product  $\langle \cdot, \cdot \rangle$  and the associated Euclidean norm  $\| \cdot \|$  or the weighted inner product  $\langle \cdot, \cdot \rangle_{\Gamma} := \langle \cdot, \Gamma^{-1} \cdot \rangle$  and its associated weighted norm  $\| \cdot \|_{\Gamma}$ , where  $\Gamma \in \mathbb{R}^{n \times n}$  is symmetric positive definite. Further we denote by  $\mathcal{B}X := \mathcal{B}(X, \| \cdot \|)$  (or respectively  $\mathcal{B}X := \mathcal{B}(X, \| \cdot \|_{\Gamma})$ ), the Borel- $\sigma$ -algebra on  $X$ . Given an additional space  $\mathbb{R}^m$  we define the tensor product of vectors in  $x \in \mathbb{R}^n$  and  $y \in \mathbb{R}^m$  by  $x \otimes y := xy^{\top}$ .

In the following, we focus on linear inverse problems of the form

$$A\theta^{\dagger} + \eta^{\dagger} = y^{\dagger}, \quad (2.1)$$

where  $\theta^{\dagger} \in X$  is the *true parameter*,  $y^{\dagger} \in Y$  is the *observed data set*,  $\eta^{\dagger} \in Y$  is *observational noise*, and  $A : X \rightarrow Y$  is a compact operator. In a Bayesian setting, we model  $\theta^{\dagger}, \eta^{\dagger}$  as random

variables  $\theta : \Omega \rightarrow X$  and  $\eta : \Omega \rightarrow Y$ , where  $\theta \perp \eta$ . Assuming that the noise is normally distributed, i.e.  $\eta \sim \mathcal{N}(0, \Gamma)$  and non-degenerate, the posterior  $\mu^y$  is characterised through Bayes' formula:

$$d\mu^y(\theta) = \frac{1}{Z} \exp\left(-\frac{1}{2}\|y - A\theta\|_{\Gamma}^2\right) d\mu_0(\theta),$$

where  $\mu_0$  denotes the prior distribution on the unknown parameters and  $Z = \mathbb{E}_{\mu_0} \exp(-\frac{1}{2}\|y - A\theta\|_{\Gamma}^2)$  is the normalisation constant. We will focus in the following on the computation of a point estimate of the unknown parameters, the maximum *a posteriori* (MAP) estimate, which is a minimiser of a regularised version of the potential

$$\Phi(\theta) = \Phi(\theta; y) = \frac{1}{2}\|y - A\theta\|_{\Gamma}^2. \quad (2.2)$$

In order to handle large data settings, i.e.  $N_{\text{obs}}$  large, we introduce a subsampling strategy, i.e. we partition the data  $y^\dagger$  into multiple subsets  $y_1^\dagger, \dots, y_{N_{\text{sub}}}^\dagger$ , such that  $(y_1^\dagger, \dots, y_{N_{\text{sub}}}^\dagger) = y^\dagger$ ,  $N_{\text{sub}} \in \mathbb{N}$ ,  $N_{\text{sub}} \geq 2$ , and  $I := \{1, \dots, N_{\text{sub}}\}$ . To this end, we define *data subspaces*  $Y_1, \dots, Y_{N_{\text{sub}}}$ , such that  $Y := \prod_{i \in I} Y_i$ . Moreover, we assume that the noise  $\eta$  has independent entries with respect to this splitting of the data space  $Y$ . In particular, we assume that there are covariance matrices  $\Gamma_i : Y_i \rightarrow Y_i$ ,  $i \in I$ , such that  $\Gamma$  has the following block diagonal structure:

$$\Gamma = \begin{pmatrix} \Gamma_1 & & & \\ & \Gamma_2 & & \\ & & \ddots & \\ & & & \Gamma_{N_{\text{sub}}} \end{pmatrix}.$$

Finally, we split the operator  $A$  into a family of  $(A_i)_{i \in I}$ , where

$$A = \begin{pmatrix} A_1 \\ \vdots \\ A_{N_{\text{sub}}} \end{pmatrix}.$$

Then, we can equivalently represent the inverse problem (2.1) by the family of inverse problems

$$\begin{aligned} A_1 \theta^\dagger + \eta_1^\dagger &= y_1^\dagger \\ &\vdots \\ A_{N_{\text{sub}}} \theta^\dagger + \eta_{N_{\text{sub}}}^\dagger &= y_{N_{\text{sub}}}^\dagger, \end{aligned}$$

where  $\eta_i^\dagger$  is a realisation of  $\eta_i \sim \mathcal{N}(0, \Gamma_i)$  for  $i \in I$ .

**Remark 2.1.** Note that the diagonal structure of the noise can always be guaranteed by multiplying (2.1) with the inverse square root of  $\Gamma$ . To simplify notation we will, w.l.o.g. assume for the remaining discussion that  $\Gamma = \text{Id}_k$  and correspondingly  $\Gamma_i = \text{Id}_{k_i}$ .

We will then consider the potentials of the respective data subset

$$\Phi_i(\theta) := \frac{1}{2}\|A_i \theta - y_i^\dagger\|^2 \quad (i \in I).$$

To overcome the ill-posedness of the inverse problem, we often consider a regularised version of the potential in the form of

$$\Phi^{\text{reg}}(\theta) := \Phi^{\text{reg}}(\theta; y) := \frac{1}{2} \|y - A\theta\|^2 + \frac{\alpha}{2} \|\theta\|_{C_0}^2,$$

where  $C_0$  is a self-adjoint, trace-class operator and  $\alpha > 0$ . This corresponds to the MAP estimate in case of a Gaussian prior with covariance  $\alpha C_0$ , cp. [12, 29]. Assuming a Gaussian prior distribution with mean equal to zero, we can incorporate the regularisation via

$$\tilde{A} = \begin{pmatrix} A \\ (\alpha C_0)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y} = \begin{pmatrix} y^\dagger \\ 0 \end{pmatrix}$$

allowing us to write

$$\Phi^{\text{reg}}(\theta) = \frac{1}{2} \|\tilde{y} - \tilde{A}\theta\|^2. \tag{2.3}$$

Note that the forward operator  $A$  is usually not injective, whereas the regularisation results in an injective operator  $\tilde{A}$ .

Similarly, when we split the forward operator we consider

$$\tilde{A}_i = \begin{pmatrix} A_i \\ \left(\frac{\alpha}{N_{\text{sub}}} C_0\right)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y}_i = \begin{pmatrix} y_i^\dagger \\ 0 \end{pmatrix}$$

leading to the family of potentials

$$\Phi_i^{\text{reg}}(\theta) = \frac{1}{2} \|\tilde{y}_i - \tilde{A}_i\theta\|^2, \quad (i \in I).$$

that satisfies

$$\Phi^{\text{reg}}(\theta) = \sum_{i=1}^{N_{\text{sub}}} \Phi_i^{\text{reg}}(\theta).$$

### 2.1. EKI and its variants

We aim to solve the inverse problem (2.1) using the EKI framework. We will focus in the following on the continuous-time limit of the Kalman inversion, cp. [30].

We define the initial ensemble to be  $\theta_0 = (\theta_0^{(j)})_{j \in J} \in X^{N_{\text{ens}}}$  assuming w.l.o.g. that  $(\theta_0^{(j)} - \bar{\theta}_0)_{j \in J}$  is a linearly independent family, with  $N_{\text{ens}} \in \mathbb{N}, N_{\text{ens}} \geq 2$ , and  $J := \{1, \dots, N_{\text{ens}}\}$ . The *basic EKI* proceeds by moving the particles in the parameter space according to the following dynamical system

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\widehat{C}^{\theta y}_t (A\theta^{(j)}(t) - y^\dagger) \quad (j \in J) \\ \theta(0) &= \theta_0, \end{aligned} \tag{2.4}$$

where

$$\widehat{C}^{\theta y}_t := \frac{1}{N_{\text{ens}} - 1} \sum_{j=1}^{N_{\text{ens}}} (\theta^{(j)}(t) - \bar{\theta}(t)) \otimes (A\theta^{(j)}(t) - A\bar{\theta}(t)), \quad \bar{\theta}(t) = \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \theta^{(j)} \quad (t \geq 0).$$

The linearity of the forward model leads to the following equivalent reformulation of the dynamical system

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\widehat{C}_t D_\theta \Phi(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0, \end{aligned} \quad (2.5)$$

where

$$\widehat{C}_t := \frac{1}{N_{\text{ens}} - 1} \sum_{j=1}^{N_{\text{ens}}} (\theta^{(j)}(t) - \bar{\theta}(t)) \otimes (\theta^{(j)}(t) - \bar{\theta}(t)).$$

Intuitively, the dynamic represents parallel gradient flows minimizing (2.2) which are coupled through the empirical covariance  $\widehat{C}_t$ . This empirical covariance can be viewed as a preconditioner. The following reformulation of the right hand side

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= - \sum_{k=1}^{N_{\text{ens}}} \langle A\theta^{(k)}(t) - A\bar{\theta}(t), A\theta^{(j)}(t) - y^\dagger \rangle (\theta^{(k)}(t) - \bar{\theta}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0 \end{aligned}$$

reveals the so-called *subspace property*, i.e. the particles  $(\theta_t^{(j)})_{j \in J}$  lie in the span of the initial ensemble  $(\theta_0^{(j)})_{j \in J}$  for any time  $t \geq 0$ , cp. [20]. Hence, we can assume that the parameter space is finite dimensional with  $X := \mathbb{R}^{N_{\text{ens}}}$  and will, w.l.o.g. do so in the following.

We define the Tikhonov-regularised ensemble Kalman inversion (TEKI) by the solution of the following ODE

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\widehat{C}_t D_\theta \Phi^{\text{reg}}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0. \end{aligned} \quad (2.6)$$

The ensemble of particles converges to the empirical mean following the dynamics given by (2.6). This results in the so-called ensemble collapse, and thus the degeneracy of the ensemble covariance operator, which results in an algebraic convergence rate rather than an exponential rate (compared to gradient flows). *Variance inflation* is a technique mitigating this effect by adding an operator to the degenerate  $\widehat{C}_t$ . Let  $C_{\text{vi}} : X \rightarrow X$  be a covariance operator on  $X$ . We define the *variance-inflated EKI* as the solution of the ODE

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -(\widehat{C}_t + \alpha_{\text{vi}} C_{\text{vi}}) D_\theta \Phi(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0, \end{aligned} \quad (2.7)$$

for  $\alpha_{\text{vi}} > 0$ .

By replacing  $\Phi$  by  $\Phi^{\text{reg}}$  in (2.7), one obtains the *variance-inflated TEKI*

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -(\widehat{C}_t + \alpha_{\text{vi}} C_{\text{vi}}) D_\theta \Phi^{\text{reg}}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0, \end{aligned}$$

for  $\alpha_{\text{vi}} > 0$ .

**2.1.1. Well-posedness and convergence analysis.** We summarise in this section the main results on the well-posedness and convergence properties of EKI with regularisation and variance inflation.

**Theorem 2.2** ([30, theorem 3.1],[12, theorem 3.2]). *Let  $u^1(0), \dots, u^J(0)$  be a given initial ensemble, we denote by  $S = \text{span}\{u^{(j)}, j \in \{1, \dots, J\}\}$  the span of the initial ensemble. Then the ODE systems (2.5)–(2.7) have unique global solutions  $u^{(j)}(t) \in C^1([0, \infty); S)$  for all  $j \in \{1, \dots, J\}$ .*

The convergence of the EKI estimate to the true parameter is restricted to the span of the initial ensemble  $S$ . More precisely, the particles stay in the affine space  $\theta_0^\perp + \mathcal{E}$  for all  $t \geq 0$ , cp. [12, corollary 3.8], where

$$\mathcal{E} = \text{span}\{e^{(1)}(0), \dots, e^{(N_{\text{ens}})}(0)\}$$

with

$$e^{(j)} = \theta^{(j)} - \bar{\theta}, \quad (j \in \{1, \dots, N_{\text{ens}}\})$$

$\theta_0^\perp = \bar{\theta}(0) - P_{\mathcal{E}}\bar{\theta}(0)$ , and  $P_{\mathcal{E}}$  being the projection onto the subspace  $\mathcal{E}$ .

This implies that the accuracy of EKI is bounded below by the accuracy of the best approximation in  $\theta_0^\perp + \mathcal{E}$ . We summarise in the following the main convergence results for the various variants of EKI.

**Theorem 2.3** ([30, theorem 3.3],[12, theorem 3.13]). *Assume that*

$$\text{span}\{Ae^{(1)}(0), \dots, Ae^{(N_{\text{ens}})}(0)\} = \mathbb{R}^{N_{\text{obs}}}.$$

*Then, the residuals of EKI mapped under the forward operator converge to 0. It holds that*

- *the rate of convergence for EKI without variance inflation is*

$$\|A\theta^{(j)} - y\|_{\Gamma}^2 \in \mathcal{O}(t^{-1}) \quad \forall j \in \{1, \dots, J\},$$

- *the rate of convergence for EKI with variance inflation is*

$$\|A\theta^{(j)} - y\|_{\Gamma}^2 \in \mathcal{O}(e^{-ct}) \quad \forall j \in \{1, \dots, J\},$$

*for a constant  $c > 0$ .*

*Assume that  $\mathcal{E} = X$ . Then, the particles of TEKI converge to the minimiser of the regularised least-squares problem  $\theta_{\text{reg}}^\dagger$ , which is given by  $\theta_{\text{reg}}^\dagger := (A^T A)^{-1} A^T y$  and respectively  $\theta_{\text{reg}}^\dagger := (\tilde{A}^T \tilde{A})^{-1} \tilde{A}^T \tilde{y}$  when considering regularisation. It holds that*

- *the rate of convergence for TEKI without variance inflation is*

$$\|\theta^{(j)} - \theta_{\text{reg}}^\dagger\|_X^2 \in \mathcal{O}(t^{-1}) \quad \forall j \in \{1, \dots, J\},$$

- *the rate of convergence for TEKI with variance inflation is*

$$\|\theta^{(j)} - \theta_{\text{reg}}^\dagger\|_X^2 \in \mathcal{O}(e^{-ct}) \quad \forall j \in \{1, \dots, J\}$$

*for a constant  $c > 0$ .*

The assumption on the affine space  $\mathcal{E} = X$  is rather restrictive and usually not satisfied in practice. We discuss in the following the generalisation of the convergence result to the more

general setting  $\theta_0^\perp + \mathcal{E} \subset X$ . The best approximation in this space is given by the solution  $\theta_{\mathcal{E}}^\dagger$  of the following constrained optimisation problem

$$\min_{\theta \in \mathcal{E}} \frac{1}{2} \|\tilde{A}(\theta + \theta_0^\perp) - \tilde{y}\|^2,$$

which can be equivalently formulated as the unconstrained optimisation problem

$$\min_{c \in \mathbb{R}^{N_{\text{ens}} - 1}} \frac{1}{2} \|\tilde{A}Ec - (\tilde{y} - \tilde{A}\theta_0^\perp)\|^2, \quad (2.8)$$

with  $E$  denoting a basis of  $\mathcal{E}$  (w.l.o.g.  $\dim(\mathcal{E}) = N_{\text{ens}} - 1$ ). Then, TEKI can be formulated in the coordinate system of  $E$  and convergence results can be straightforwardly generalised. Note that convergence then follows to the minimiser of 2.8, i.e. the best approximation of 2.3 in the affine space  $\theta_0^\perp + \mathcal{E}$ . We refer to [12] for more details on the derivation of the convergence result. We will denote in the following the optimiser of the constrained optimisation problem by  $\theta^*$ , i.e.

$$\theta^* \in \operatorname{argmin}_{\theta \in \mathcal{E}} \frac{1}{2} \|\tilde{A}(\theta + \theta_0^\perp) - \tilde{y}\|^2,$$

### 3. Subsampling in continuous time

In this work, we are interested in certain stochastic approximations of ODEs, indeed, we are studying EKIs in which we randomly replace the potential  $\Phi^{\text{reg}}$  by one of the  $(\Phi_i^{\text{reg}})_{i \in I}$ . We now introduce and study a framework in which we are able to consider the subsampling of (actually, more general) flows, before then discussing the subsampling in EKI.

#### 3.1. A general framework and result

Let  $\mathbf{F}_i : X \times [0, \infty) \rightarrow X$  be Lipschitz continuous for  $i \in I$ . Moreover, we define  $\bar{\mathbf{F}} = \sum_{i \in I} \mathbf{F}_i / N_{\text{sub}}$ . We study the full dynamical system  $(\theta(t))_{t \geq 0}$ , given by

$$\begin{aligned} \dot{\theta}(t) &= -\bar{\mathbf{F}}(\theta(t), t) & (t > 0) \\ \theta(0) &= \theta_0 \in X \end{aligned} \quad (3.1)$$

and also the subsampled dynamical system

$$\begin{aligned} \dot{\theta}(t) &= -\mathbf{F}_i(\theta(t), t) & (t > 0) \\ \theta(0) &= \theta_0 \in X. \end{aligned} \quad (3.2)$$

We denote the flows with respect to  $(\mathbf{F}_i)_{i \in I}$  by  $(\varphi_t^{(i)})_{i \in I, t \geq 0}$ : Let  $t > 0, i \in I$ , and  $\theta_0 \in X$ , then

$$\dot{\varphi}_t^{(i)}(\theta_0) = -\mathbf{F}_i(\varphi_t^{(i)}(\theta_0)), \quad \varphi_0^{(i)}(\theta_0) = \theta_0.$$

In the same way, we denote the flow with respect to  $\bar{\mathbf{F}}$  by  $\bar{\varphi}_t$ .

Solving or approximating the full dynamical system (3.1) can be computationally expensive, especially if  $N_{\text{obs}}$  is large. We will now discuss an approximation strategy for this full dynamical system that replaces the full system at any point in time by a randomly selected subsampled system. Hence, in any small time interval, we only need to evaluate (3.2) for some  $i \in I$ .

We define the subsampled system through a continuous-time Markov process (CTMP) on  $I$ , which we call  $(i(t))_{t \geq 0}$ . Let  $\eta : [0, \infty) \rightarrow (0, \infty)$  be continuously differentiable and bounded



from above. We refer to  $\eta(t)$  as *learning rate* at time  $t \geq 0$ . Let  $\mathbf{i} : [0, \infty) \times \Omega \rightarrow I$  be the CTMP with transition rate matrix

$$A(t) := \frac{1}{(N_{\text{sub}} - 1)\eta(t)} \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \cdots & 1 \end{pmatrix} - \frac{N_{\text{sub}}}{(N_{\text{sub}} - 1)\eta(t)} \cdot \text{id}_I \quad (t \geq 0) \quad (3.3)$$

and initial distribution  $\mathbf{i}(0) \sim \text{Unif}(I)$ .  $(\mathbf{i}(t))_{t \geq 0}$  is the stochastic process characterised by Algorithm 1.

---

**Algorithm 1.** Sampling  $(\mathbf{i}(t))_{t \geq 0}$ .

---

1: initialise  $\mathbf{i}(0) \sim \text{Unif}(I)$  and  $t_0 \leftarrow 0$

2: sample  $\Delta$  with survival function

$$\mathbb{P}(\Delta \geq t | t_0) := \mathbf{1}[t < 0] + \exp\left(-\int_0^t \eta(u + t_0)^{-1} du\right) \quad (t \in [-\infty, \infty])$$

3: set  $\mathbf{i}|_{(t_0, t_0 + \Delta)} \leftarrow \mathbf{i}(t_0)$

4: sample  $\mathbf{i}(t_0 + \Delta) \sim \text{Unif}(I \setminus \{\mathbf{i}(t_0)\})$

5: increment  $t_0 \leftarrow t_0 + \Delta$  and go to 2

---

Hence, the process  $(\mathbf{i}(t))_{t \geq 0}$  is a piecewise constant process that jumps from one state to another after random waiting times. There are several other characterisations of the process  $(\mathbf{i}(t))_{t \geq 0}$ , we refer the reader to [1] for general CTMPs on discrete spaces. The algorithmic procedure above goes back to Gillespie [17]. Properties of this particular CTMP have been studied in [23]. We can now define the stochastic approximation process for  $(\mathbf{F}_i)_{i \in I}$  and  $(\mathbf{i}(t))_{t \geq 0}$ .

**Definition 3.1.** We define the *stochastic approximation process* given by the family of flows  $(\mathbf{F}_i)_{i \in I}$  and the index process  $(\mathbf{i}(t))_{t \geq 0}$  by the tuple  $(\mathbf{i}(t), \theta(t))_{t \geq 0}$ , with

$$\begin{aligned} \dot{\theta}(t) &= -\mathbf{F}_{\mathbf{i}(t)}(\theta(t), t) \quad (t > 0) \\ \theta(0) &= \theta_0 \in X. \end{aligned}$$

In the following, we are interested in the long time behaviour of the stochastic approximation process.

**Assumption 3.2** Let  $K \in \mathbb{N}$  and  $X := \mathbb{R}^K$ . Let (i) and (ii) hold for any  $i \in I$ :

(i)  $\mathbf{F}_i \in C^1(X \times [0, \infty), X)$ ,

(ii) the flow  $\varphi_t^{(i)}$  contracts quickly—in particular, we have a measurable function  $h : [0, \infty) \rightarrow \mathbb{R}$ , with  $\int_0^\infty h(t) dt = \infty$  such that

$$\langle \mathbf{F}_i(\varphi_t^{(i)}(\theta_0), t) - \mathbf{F}_i(\varphi_t^{(i)}(\theta_1), t), \varphi_t^{(i)}(\theta_0) - \varphi_t^{(i)}(\theta_1) \rangle_X \leq -h(t) \|\varphi_t^{(i)}(\theta_0) - \varphi_t^{(i)}(\theta_1)\|^2$$

for any two initial values  $\theta_0, \theta_1 \in X$ .

Note that assumption 3.2(ii) already implies that the flow of  $-\bar{\mathbf{F}}$  is exponentially contracting. The Banach fixed-point theorem implies that the flow has a unique stationary point, which we denote by  $\theta^* \in X$ . We now generalise one of the main results of [23] by showing that the stochastic process converges to the unique stationary point  $\theta^*$  of the flow  $(\bar{\varphi}_t)_{t \geq 0}$  if the learning rate goes to zero. Convergence is measured in terms of the Wasserstein distance

$$d_W(\pi, \pi') = \inf_{H \in C(\pi, \pi')} \int_{X \times X} \min\{1, \|\theta - \theta'\|^q\} dH(\theta, \theta'),$$

where  $q \in (0, 1]$  and  $C(\pi, \pi')$  is the set of couplings of the probability measures  $\pi, \pi'$  on  $(X, \mathcal{B}X)$ .

**Theorem 3.3.** *Let assumption 3.2(i) and (ii) hold for a constant  $h$  and a stochastic approximation process  $(\mathbf{i}(t), \theta(t))_{t \geq 0}$  with initial values  $(i_0, \theta_0) \in I \times X$ . Moreover, assume that  $\lim_{t \rightarrow \infty} \eta(t) = 0$ . Then,*

$$\lim_{t \rightarrow \infty} d_W(\delta(\cdot - \theta^*), \mathbb{P}(\theta(t) \in \cdot | \theta_0, i_0)) = 0.$$

**Proof.** Please see theorem A.1. □

Hence, the process converges to the Dirac measure  $\delta(\cdot - \theta^*)$ .

### 3.2. EKI with subsampling

We introduce two possibilities of subsampling: EKI with single subsampling and EKI with batch-subsampling. In the first case, we choose a single subsample  $y_i^\dagger$  from  $y^\dagger$  essentially replace the potential in (2.4) by the subsampled potential  $\Phi_i$ . In the mini-batching case, we pick a total of  $N_{\text{ens}}$  subsamples  $(y_{i(1)}^\dagger, \dots, y_{i(N_{\text{ens}})}^\dagger)$ , i.e. one subsample for each particle in the ensemble. Then, we evolve each of the ensemble members with respect to their data subsample, i.e.  $\theta^{(j)}(t)$  is evolved with respect to  $\Phi_{i(j)}$ , for  $j = 1, \dots, N_{\text{ens}}$ .

After one remark, we continue by defining the single subsampling.

**Remark 3.4.** When defining the subsampling algorithms, we will only refer to the basic EKI. Of course, it is possible to combine subsampling with Tikhonov regularisation. In this case, we replace the potential  $\Phi_i$  in (3.4) with

$$\Phi'_i(\theta) := \Phi_i(\theta) + \frac{\alpha}{2N_{\text{sub}}} \|\theta\|_X^2$$

for  $i \in I$ . In the same way, one can combine subsampling with variance inflation, by replacing the empirical covariance  $\widehat{C}_t$  with the inflated covariance  $(\widehat{C}_t + \alpha' C_{\text{vi}})$ .


**3.2.1. Single subsampling.** The essential idea is now the following: at every time step, we follow the EKI flow of the potential  $\Phi_i$ , for one random  $i \in I$ , as determined by  $(\mathbf{i}(t))_{t \geq 0}$ . Indeed, the *EKI with single subsampling* is defined via the dynamical system

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\widehat{C}_t D_\theta \Phi_{i(t)}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0. \end{aligned} \tag{3.4}$$


We illustrate this single subsampling strategy in figure 1.

**3.2.2. Batch-subsampling.** In batch-subsampling, we define a set  $\widehat{I} \subseteq I^{N_{\text{ens}}}$  such that for all  $i \in \{1, \dots, N_{\text{sub}}\}, j \in J$ , the number of elements in  $\{\widehat{i} \in \widehat{I} : \widehat{i}_j = i\}$  is identical. Moreover, we define a stochastic process  $\mathbf{i} : [0, \infty) \times \Omega \rightarrow \widehat{I}$ . Here, the coordinate processes  $(\mathbf{i}(t; j))_{t \geq 0}$ , for  $j = 1, \dots, N_{\text{ens}}$ , are stochastically independent CTMPs with transition rate matrix  $(A(t))_{t \geq 0}$ , as given in (3.3). The process  $(\mathbf{i}(t; j))_{t \geq 0}$  now represents the data subset with which the particle  $\theta^{(j)}$  is evolved, for  $j = 1, \dots, N_{\text{ens}}$ . Hence, the *EKI with batch-subsampling* is given by

$$\frac{d\theta^{(j)}(t)}{dt} = -\widehat{C}_t D_\theta \Phi_{i(t; j)}(\theta^{(j)}(t)) \quad (j \in J)$$



**Figure 1.** Cartoon of EKI with single subsampling: same data subset  $y_1, y_2, \dots$  for each ensemble member.



**Figure 2.** Cartoon of EKI with batch subsampling: different data subsets  $y_1, y_2, \dots$  for each ensemble member.

$$\theta(0) = \theta_0.$$

We illustrate the batch subsampling strategy in figure 2.

#### 4. Analysis of the EKI with subsampling in the linear setting

We present in the following a convergence theory for the two subsampling strategies in the linear setting. Our goal is to verify the assumptions 3.2, in particular the condition

$$-\langle \theta_1 - \theta_2, \mathbf{F}_i(\theta_1, t) - \mathbf{F}_i(\theta_2, t) \rangle \leq -h(t) \|\theta_1 - \theta_2\|^2,$$

for  $t$  large enough, with  $-\mathbf{F}_i(\theta(t), t)$  denoting the right hand side of the dynamical system and  $h : [0, \infty) \rightarrow \mathbb{R}$  being a measurable function. We will see that the analysis of the ensemble collapse will play a central role for the construction of the function  $h$ . In order to derive convergence results in the parameter space, we will focus in the following on the regularised setting, i.e. we consider the potential  $\Phi^{\text{reg}}$ .

##### 4.1. TEKI with variance inflation

The variance inflation allows to explicitly control the ensemble collapse by controlling the preconditioner in the following way:

**Theorem 4.1 (Single Subsampling TEKI with variance inflation).** Let  $(\theta^{(j)}(t))_{t \geq 0, j \in J}$  satisfy

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -(\widehat{C}_t + \alpha_{vi}C_{vi})D_{\theta}\Phi_{i(t)}^{reg}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0. \end{aligned}$$

with index process  $(\mathbf{i}(t))_{t \geq 0}$  and

$$\Phi_{i(t)}^{reg}(\theta^{(j)}(t)) = \frac{1}{2} \|\tilde{y}_{i(t)} - \tilde{A}_{i(t)}\theta\|^2, \quad \tilde{A}_{i(t)} = \begin{pmatrix} A_{i(t)} \\ \left(\frac{\alpha}{N_{sub}}C_0\right)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y} = \begin{pmatrix} y_{i(t)}^{\dagger} \\ 0 \end{pmatrix},$$

with  $\alpha, \alpha_{vi} > 0$ ,  $\widehat{C}_t$  denoting the empirical covariance matrix of the particles  $\theta^{(j)}(t)$ , and  $C_0, C_{vi}$  being symmetric positive definite matrices. The tuple  $(\mathbf{i}(t), \theta(t))_{t \geq 0}$  denotes the single subsampling TEKI solution with variance inflation. Then,

$$\lim_{t \rightarrow \infty} d_W \left( \delta(\cdot - \theta^*), \mathbb{P}(\theta^{(j)}(t) \in \cdot | \theta_0, i_0) \right) = 0 \quad (j \in J).$$

**Proof.** By theorem 3.3, we know that convergence of  $(\theta^{(j)}(t))$ , for  $j \in J$ , follows if assumptions 3.2 are satisfied. The continuous differentiability of the right hand side follows straightforwardly from the definition of the TEKI. The inequality holds with

$$h(t) = \alpha_{vi} \lambda_{\min}(C_{vi}) \left( \min_{1, \dots, N_{sub}} \lambda_{\min}(\tilde{A}_i^T \tilde{A}_i) \right),$$

where obviously  $\int_0^{\infty} h(t) dt = \infty$  is fulfilled. The details on the construction of  $h$  are given in lemma A.5.  $\square$

**Theorem 4.2 (Batch subsampling TEKI with variance inflation).** Let  $(\theta^{(j)}(t))_{t \geq 0, j \in J}$  satisfy

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -(\widehat{C}_t + \alpha_{vi}C_{vi})D_{\theta}\Phi_{i(t;j)}^{reg}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0. \end{aligned}$$

with index process  $(\mathbf{i}(t;j))_{t \geq 0, j \in 1, \dots, N_{ens}}$  and

$$\Phi_{i(t;j)}^{reg}(\theta^{(j)}(t)) = \frac{1}{2} \|\tilde{y}_{i(t;j)} - \tilde{A}_{i(t;j)}\theta\|^2, \quad \tilde{A}_{i(t;j)} = \begin{pmatrix} A_{i(t;j)} \\ \left(\frac{\alpha}{N_{sub}}C_0\right)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y}_{i(t;j)} = \begin{pmatrix} y_{i(t;j)}^{\dagger} \\ 0 \end{pmatrix},$$

with  $\alpha, \alpha_{vi} > 0$ ,  $\widehat{C}_t$  denoting the empirical covariance matrix of the particles  $\theta^{(j)}(t)$ , and  $C_0, C_{vi}$  being symmetric positive definite matrices. The tuple  $(\mathbf{i}(t;j), \theta(t))_{t \geq 0}$  denotes the batch subsampling TEKI solution with variance inflation. Then,

$$\lim_{t \rightarrow \infty} d_W \left( \delta(\cdot - \theta^*), \mathbb{P}(\theta^{(j)}(t) \in \cdot | \theta_0, i_0) \right) = 0,$$

for  $j \in J$ .

**Proof.** The proof follows the same lines as the proof of theorem 4.1.  $\square$

#### 4.2. TEKI without variance inflation

We have seen that the control on the smallest eigenvalue of the preconditioners, i.e. the empirical covariances, is crucial in order to prove convergence. We will consider in the following

the more general case, where the smallest eigenvalue converges to 0 with a rate such that  $\int_0^\infty h(t)dt = \infty$  still holds true and convergence follows by theorem 3.3.

**Theorem 4.3 (Single Subsampling TEKI without variance inflation).** *Let  $(\theta^{(j)}(t))_{t \geq 0, j \in J}$  satisfy*

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\widehat{C}_t D_\theta \Phi_{i(t)}^{reg}(\theta^{(j)}(t)) \quad (j \in J) \\ \theta(0) &= \theta_0. \end{aligned} \tag{4.1}$$

with index process  $(i(t))_{t \geq 0}$  and

$$\Phi_{i(t)}^{reg}(\theta^{(j)}(t)) = \frac{1}{2} \|\tilde{y}_{i(t)} - \tilde{A}_{i(t)}\theta\|^2, \quad \tilde{A}_{i(t)} = \begin{pmatrix} A_{i(t)} \\ \left(\frac{\alpha}{N_{sub}} C_0\right)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y} = \begin{pmatrix} y_{i(t)}^\dagger \\ 0 \end{pmatrix},$$

with  $\alpha > 0$ ,  $\widehat{C}_t$  denoting the empirical covariance matrix of the particles  $\theta^{(j)}(t)$ , and  $C_0$  being a symmetric positive definite matrix. The tuple  $(i(t), \theta(t))_{t \geq 0}$  denotes the single subsampling TEKI solution without variance inflation. Then,

$$\lim_{t \rightarrow \infty} d_W \left( \delta(\cdot - \theta^*), \mathbb{P}(\theta^{(j)}(t) \in \cdot | \theta_0, i_0) \right) = 0,$$

for  $j \in J$ .

**Proof.** The continuous differentiability of the right hand side follows with the same argument as in theorem 4.3. The inequality holds with  $h(t) = \lambda_{\min}(\widehat{C}_t^1) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(\tilde{A}_i^T \tilde{A}_i)$ , where  $\int_0^\infty h(t)dt = \infty$  is fulfilled due to  $\lambda_{\min}(\widehat{C}_t^1) \in \mathcal{O}(t^{-1})$ . The details on the minimal eigenvalue are given in lemma A.6. □

In the case of batch subsampling, the rate of convergence of each particle is exponential for a fixed data set, as the ensemble collapse is prevented (under suitable assumptions of the data). This is contrast to the single subsampling case, where the rate of convergence is algebraic due to the ensemble collapse. This can be shown as follows:

**Lemma 4.4.** *Given the  $N_{sub}$  subsamples  $(\tilde{y}_{(1)}^\dagger, \dots, \tilde{y}_{(N_{sub})}^\dagger)$ , assume that the centred initial ensemble is a generator of the full space  $X$ , i.e.  $\text{span}\{e_0^{(j)}, j \in J\} = X$ . We further assume that  $\sum_{j=1}^{N_{ens}} (\theta_j^\dagger - \bar{\theta}^\dagger)(\theta_j^\dagger - \bar{\theta}^\dagger)^\top$  has full rank  $d$ . Then the particles converge to the true solution  $\theta^\dagger$  exponentially fast, i.e.  $\theta^{(j)} \rightarrow \theta_i^\dagger$  with  $\theta_i^\dagger$  denoting the minimiser of  $\Phi_i(\theta) = \frac{1}{2} \|\tilde{A}_i\theta - \tilde{y}_i\|^2$ .*

**Proof.** Please see lemma A.4. □

The assumption on the ensemble spread being a generating set of the parameter space  $X$  is rather restrictive and usually not satisfied in practice. Generalisation of the result is straightforward when working in the coordinate system of the linear subspace spanned by the initial ensemble using the projection  $P_{\mathcal{E}}$ .

The batch subsampling case leads to exponential convergence rates for fixed data, however, as the minimum eigenvalue depends on the initial data, the convergence result cannot be readily applied. We modify the process such that we can control the minimum eigenvalue similar to the variance inflation technique above. However, instead of considering a static lower bound, we now allow the minimum eigenvalue to decrease at a certain rate.

**Theorem 4.5 (Batch subsampling TEKI with diminishing variance inflation).** *Let  $(\theta^{(j)}(t))_{t \geq 0, j \in J}$  satisfy*

$$\frac{d\theta^{(j)}(t)}{dt} = -(\widehat{C}_t + \frac{\alpha_{vi}}{1+t} C_{vi}) D_\theta \Phi_{i(t;j)}^{reg}(\theta^{(j)}(t)) \quad (j \in J) \tag{4.2}$$

$$\theta(0) = \theta_0.$$

with index process  $(i(t;j))_{t \geq 0, j \in 1, \dots, N_{ens}}$  and

$$\Phi_{i(t;j)}^{reg}(\theta^{(j)}(t)) = \frac{1}{2} \|\tilde{y}_{i(t;j)} - \tilde{A}_{i(t;j)} \theta\|^2, \quad \tilde{A}_{i(t;j)} = \begin{pmatrix} A_{i(t;j)} \\ (\frac{\alpha}{N_{sub}} C_0)^{\frac{1}{2}} \end{pmatrix}, \quad \tilde{y} = \begin{pmatrix} y_{i(t;j)}^\dagger \\ 0 \end{pmatrix},$$

with  $\alpha, \alpha_{vi} > 0, \widehat{C}_t$  denoting the empirical covariance matrix of the particles  $\theta^{(j)}(t)$ , and  $C_0, C_{vi}$  being symmetric positive definite matrices. The tuple  $(i(t;j), \theta(t))_{t \geq 0}$  denotes the batch subsampling TEKI solution with variance inflation. Then,

$$\lim_{t \rightarrow \infty} d_W(\delta(\cdot - \theta^*), \mathbb{P}(\theta^{(j)}(t) \in \cdot | \theta_0, i_0)) = 0 \quad (j \in J).$$

**Proof.** The proof follows the same lines as the proof of theorem 4.1. The only difference being that we have here  $h(t) = \frac{\alpha_{vi}}{t} \lambda_{\min}(C_{vi}) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(A_i^T A_i)$ , The diminishing rate however is slow enough to obtain  $\int_0^\infty h(t) dt = \infty$ .  $\square$

We now move on to proving the convergence of the subsampled EKI. We start by defining an auxiliary EKI subsampling process. Let  $\varepsilon \in (0, 1)$  and

$$\frac{d\theta^{(j,\varepsilon)}(t)}{dt} = -\widehat{C}_t D_\theta \Phi_{i'(t,\varepsilon;j)}(\theta^{(j,\varepsilon)}(t)) \quad (j \in J)$$

$$\theta^{(j,\varepsilon)}(0) = \theta_0^{(j)} \quad (j \in J),$$

where  $(i'(t,\varepsilon))_{t \geq 0}$  is the CTMP on  $\widehat{I}$  with transition rate matrix  $B(t) := A(t1[t \in [0, 1/\varepsilon]] + 1/\varepsilon 1[t \in (1/\varepsilon, \infty)])$ .

**Proposition 4.6.** *Let  $\varepsilon > 0$ . Then, the process  $(\theta^{(\cdot,\varepsilon)}(t), \xi^\varepsilon(t), i'(t,\xi))_{t \geq 0}$  has a unique stationary measure  $\mu_\varepsilon$ . Moreover, for every  $\theta_0 \in X^{N_{ens}}$  and  $i_0 \in \widehat{I}$  there are  $c, c' > 0$ , with*

$$d_W(\mu_\varepsilon, \mathbb{P}((\theta^{(\cdot,\varepsilon)}(t), \xi^\varepsilon(t), i'(t,\xi)) \in \cdot | \theta(0) = \theta_0, i'(t) = i_0)) \rightarrow 0 \quad (t \rightarrow \infty).$$

**Proof.** 1. We note that for any initial value  $\theta_0 \in X^{N_{ens}}$  there is a compact set  $M \subseteq X^{N_{ens}}, M \ni \theta_0$  from which the process  $(\theta^{(j,\varepsilon)}(t), i'(t,\varepsilon))_{t \geq 0}$  cannot escape. See [2] for details. Moreover, we note that under the hypothesis of theorem 4.4, we can assure that  $\theta^{(\cdot,\varepsilon)}(t) \notin X_{diag} := \{(\theta_1, \dots, \theta_{N_{ens}}) \in X^{N_{ens}} : \exists j, j' \in J, \theta_j = \theta_{j'}\}$

2. We now show that two coupled processes  $(\theta^{(j,\varepsilon)}(t), i'(t,\varepsilon))_{t \geq 0}$  and  $(\theta_*^{(j,\varepsilon)}(t), i'(t,\varepsilon))_{t \geq 0}$  starting at different initial points  $\theta_0, \theta_{0,*}$  contract in the Wasserstein-2 distance. Let  $(i'(t,\varepsilon; \cdot))_{t \geq 0}$  be a realisation of  $(i'(t,\xi; \cdot))_{t \geq 0}$ . Then, again we try to find a continuous function  $h_{i'} : [0, \infty) \rightarrow (0, \infty)$  with

$$-(\theta - \theta_*, \widehat{C}_t D_\theta \Phi_{i'(t,\varepsilon; \cdot)}(\theta) - \widehat{C}_t D_\theta \Phi_{i'(t,\varepsilon; \cdot)}(\theta_*)) \leq -h_{i'}(t) \|\theta - \theta_*\|^2.$$

Lemma A.5 gives us  $h_{i'}(t) = \frac{\alpha_{vi}}{1+t} \lambda_{\min}(C_{vi}) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(A_i^T A_i)$ .

Note that we are using variance inflation, however one that is diminishing at rate  $t^{-1}$ . This is slow enough so that we obtain  $\int_0^\infty h_{i'}(t) dt = \infty$ .  $\square$

In proposition 4.6, we showed that the auxiliary process  $(\theta^\varepsilon(t))_{t \geq 0}$  is ergodic ( $\varepsilon > 0$ ) and converges to a stationary measure. In the following, we will show that  $(\theta^\varepsilon(t))_{t \geq 0} \rightarrow (\theta(t))_{t \geq 0}$  as  $\varepsilon \rightarrow 0$  and then also that  $\theta_t \rightarrow P_Y \theta^\dagger$  as  $t \rightarrow \infty$ .

**Theorem 4.7.** *Under the assumptions of proposition 4.6, we have*

$$d_W(\delta(\cdot - P_Y \theta^\dagger), \mathbb{P}(\theta(t) \in \cdot)) \rightarrow 0 \quad (t \rightarrow \infty).$$

**Proof.** The result follows from [23].  $\square$

## 5. Numerical experiments

We now test our methodology in two numerical experiments: first, we aim to estimate the source term in a 1D parabolic PDE using measurements from its solution. Then, we estimate the log-diffusion coefficient in a 2D elliptic PDE, again using measurements of the solution.

### 5.1. 1D-Heat equation

In this experiment we consider a one-dimensional Heat equation, given by the following differential equation

$$\begin{aligned} \frac{\partial u(x,t)}{\partial t} - \frac{\partial^2 u(t,x)}{\partial x^2} &= f(x) & (t > 0, x \in (0, 1)) \\ u(0,x) &= 0 & (x \in (0, 1)) \\ u(t,0), u(t,1) &= 0 & (t \geq 0). \end{aligned}$$


Our goal in this inverse problem is to estimate the unknown forcing  $f$  from perturbed measurement data that we have obtained from the solution. Indeed, we define  $A = \mathcal{O} \circ L^{-1}$ , where  $L = \frac{d}{dt} - \frac{d^2}{dx^2}$ , and  $\mathcal{O} : H_0^1([0, 1], \mathbb{R}) \rightarrow \mathbb{R}^K$ ,  $(p(\cdot)) \mapsto \mathcal{O}(p(\cdot)) = (p(x_1), \dots, p(x_K))^T$  is an equidistant observation operator on  $[0, 1] \times \mathbb{R}_{\geq 0}$  and  $p \in H_0^1([0, 1], \mathbb{R})$  is a solution operator of the PDE. The inverse problem is given by:

$$u = Af + \eta,$$

where  $\eta \sim \mathcal{N}(0, \Gamma)$ , with  $\Gamma = 0.1^2 \text{Id}_K$ .

The forcing is assumed to be a Gaussian random field with zero mean and covariance  $C(s, t) = \sigma^2 \exp(-\frac{|s-t|^2}{L_{sc}})$ , where  $(s, t) \in [0, 1] \times [0, 1]$ ,  $\sigma^2 = 10$  and  $L_{sc} = 0.1$ . We simulate the random field using a KL-expansion, which is truncated after 8 terms, i.e.  $f(x, \omega) = \sum_{i=1}^8 \lambda_i^{1/2} e_i(x) \xi_i(\omega)$ , where  $\lambda_i$  are the largest eigenvalues of  $C(s, t)$ ,  $e_i(x)$  the corresponding eigenfunctions and  $\xi_i$  standard normal distributed random variables. Furthermore, we use a spatial step size of  $h = 0.01$ , a time step size of  $\Delta = 0.05$  and a time horizon of  $T = 0.3$ . Hence, we have 6 time steps. The PDE-solution is then computed through the Crank-Nicolson method.

To solve the inverse problem we will consider the numerical computed solutions of  $u(t, x)$  at each time step as one subsample. Therefore we have  $N_{\text{sub}} = 6$  many subsamples. Additionally the choice of the step size  $h$ , yields  $N_{\text{dim}(X)} = 99$  interior points that we seek to estimate. Furthermore, we have  $N_{\text{obs}} = 6 \cdot 99 = 594$  many observations and use  $N_{\text{ens}} = 5$  particles. Our initial ensemble is assumed to have the same distribution of the Gaussian random field that we choose for the forcing  $f$ . We sample from the random field by making  $N_{\text{ens}}$  independent draws from  $f(x, \omega)$ . Therefore, the generated solutions will lie in the subspace  $f_0^\perp + \mathcal{E}$ , where  $\mathcal{E}$  is the



**Figure 3.** Mean absolute errors of computed solutions in the parameter (left) and observation space (right). The red line illustrates the EKI, the blue line single subsampling and the green line batch subsampling.

linear span of the centred initial ensemble and  $f_0^\perp = \bar{f}(0) - P_{\mathcal{E}}\bar{f}(0)$ . The best approximation in this space is given by the solution  $f_{\mathcal{E}}^\dagger$  of the constrained optimisation problem (2.8), i.e.

$$\min_{c \in \mathbb{R}^{N_{\text{ens}}-1}} \frac{1}{2} \|\tilde{A}Ec - (\tilde{y} - \tilde{A}f_0^\perp)\|^2,$$

with  $E$  denoting a basis of  $\mathcal{E}$  and  $\tilde{A}$  and  $\tilde{y}$  correspond to the regularised versions of the respective variable. This solution can be computed analytically and is given by

$$c_{\mathcal{E}}^\dagger = ((\tilde{A}E)^T(\tilde{A}E))^{-1} (\tilde{A}E)^T(\tilde{y} - \tilde{A}\theta_0^\perp).$$

Thus the reference solution in the parameter space is given by

$$f_{\mathcal{E}}^\dagger = Ec_{\mathcal{E}}^\dagger + \theta_0^\perp.$$


We simulate  $N = 32$  many runs and illustrate the mean absolute error of the runs in the parameter space as well as in the observation space for single-subsampling, batch-subsampling and compare it to the EKI. Lastly, we also illustrate the mean ensemble collapse for each particle. The ODE solutions of the EKI as well as our subsampling methods are computed using MATLABs ode45 ODE solver.

**5.1.1. TEKI with variance inflation.** The first conducted experiment uses constant variance inflation of the magnitude  $\alpha_{vi} = 0.01$  and regularisation of  $\beta = 10$ . Furthermore, the learning rate decays at exponential speed, i.e.  $\eta(t) = a \exp(-bt)$ , with  $a = 0.01$  and  $b = 10$ . We compute the solution up until  $T = 1$ . With those parameter choices we obtain approx.  $2 \cdot 1e5$  many data changes. We illustrate our results in semi-log plots, i.e. we only apply the log function on to the y-axis.

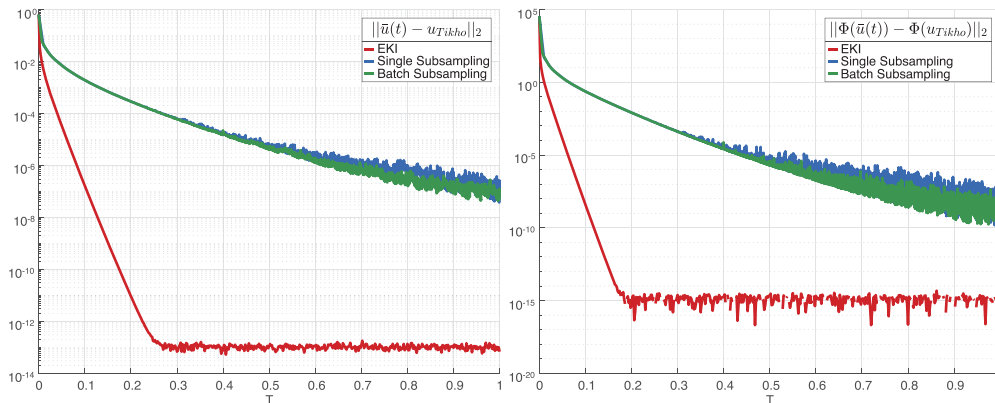
Figure 3 shows the mean relative error over all  $N = 32$  runs, w.r.t the Tikhonov solution. The left figure shows the mean error in the parameter space and the right figure in the observation space. We can see that both methods converge towards the true solution at exponential rate, due to the linear decay in the semi-log plots. They only differentiate themselves in the constants.

In figure 4 the mean ensemble collapse of the particles is illustrated. Again we can see that the collapse happens at an exponential rate. Interesting to note is that there seem to be more fluctuations in batch-subsampling than in single-subsampling. And that batch-subsampling is also a bit slower than single-subsampling(see different scaling on y-axis).





**Figure 4.** Mean ensemble collapse of all  $N_{\text{ens}} = 5$  particles for the three methods. The different colours represent the ensemble collapse of one ensemble member respectively. Left figure: EKI; middle figure: single subsampling; right: batch subsampling.




**Figure 5.** Mean absolute errors of computed solutions in the parameter (left) and observation space (right). The red line illustrates the EKI, the blue line single-subsampling and the green line batch-subsampling.

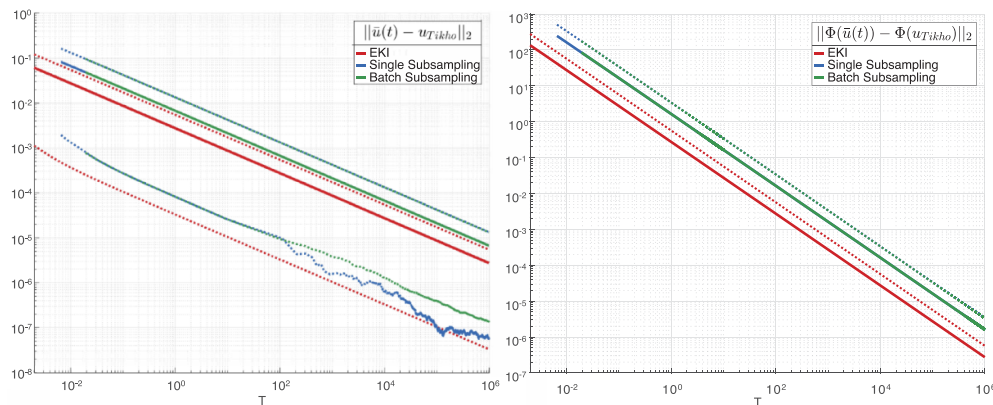
**5.1.2. TEKI with diminishing variance inflation.** In this simulation we consider diminishing variance inflation. We illustrate the results using the same parameters as chosen in the experiment with constant variance inflation. We let variance inflation vanish at a linear rate, i.e. the gradient flow is given by (4.2) with  $\alpha_{vi} = 0.01$ . We use again an exponential decaying learning rate and we illustrate the results for EKI, single-subsampling and batch-subsampling.

We can see in figure 5 similar results as in 3. Both subsampling methods converge at a similar rate towards the solution. However, they are both converging at a slower rate than the EKI. Furthermore, this experiment shows less noise in the subsampling approaches as opposed to a constant variance inflation. Figure 6 shows the ensemble collapse of all methods. We get similar results to the experiment conducted before.

**5.1.3. TEKI without variance inflation.** In this subsection we consider experiments without variance inflation. In theorem 4.1 we proved convergence for single-subsampling. However, we were not able to obtain a result for batch-subsampling. We illustrate numerical results, indicating that we can also expect convergence for batch-subsampling without variance inflation.




**Figure 6.** Mean ensemble collapse of all  $N_{\text{ens}} = 5$  particles for the three methods. The different colours represent the ensemble collapse of one ensemble member respectively. Left figure: EKI; middle figure: single-subsampling; right figure: batch-subsampling.




**Figure 7.** Mean absolute errors  $\pm$  standard deviation of computed solutions in the parameter (left) and observation space (right). The red line illustrates the EKI, the blue line single-subsampling and the green line batch-subsampling.

We consider again  $\beta = 10$  as regularisation parameter and compute the solution up until  $T = 1e6$ . This time we use a linear decaying learning rate  $\eta(t) = (at + b)^{-1}$ , with  $a, b = 100$ . However, due to the decreasing switching times of the subsets, the algorithm becomes computationally slow. We therefore only use the linear decaying learning rate up until  $T = 10^1$ . Afterwards we consider  $1e5$  equidistant switching times. Up until  $T = 1$  we obtain approx 6000 data switches. Finally, we illustrate our results in log-log plots, since we expect a similar convergence rate as the EKI, which is algebraic. Through Log-log plots, it is easier to observe this convergence rate as opposed to semi-log y plots. Again we conduct  $N = 32$  experiments and illustrate the mean absolute error.

Figure 7 depicts the mean absolute error in the parameter and observation space to the Tikhonov solution. In those experiments we also included the mean errors  $\pm$  one standard deviation. Both subsampling methods show similar results to the EKI. Our methods are therefore suitable alternatives to the EKI, since we obtain the same convergence rate, however we have less computational costs. Note that the mean minus standard deviation is negative in the right picture and therefore not shown.



**Figure 8.** Mean ensemble collapse of all  $N_{\text{ens}} = 5$  particles for the three methods. The different colours represent the ensemble collapse of one ensemble member respectively. Left figure: depicts EKI; middle figure: single-subsampling; right figure: batch-subsampling.

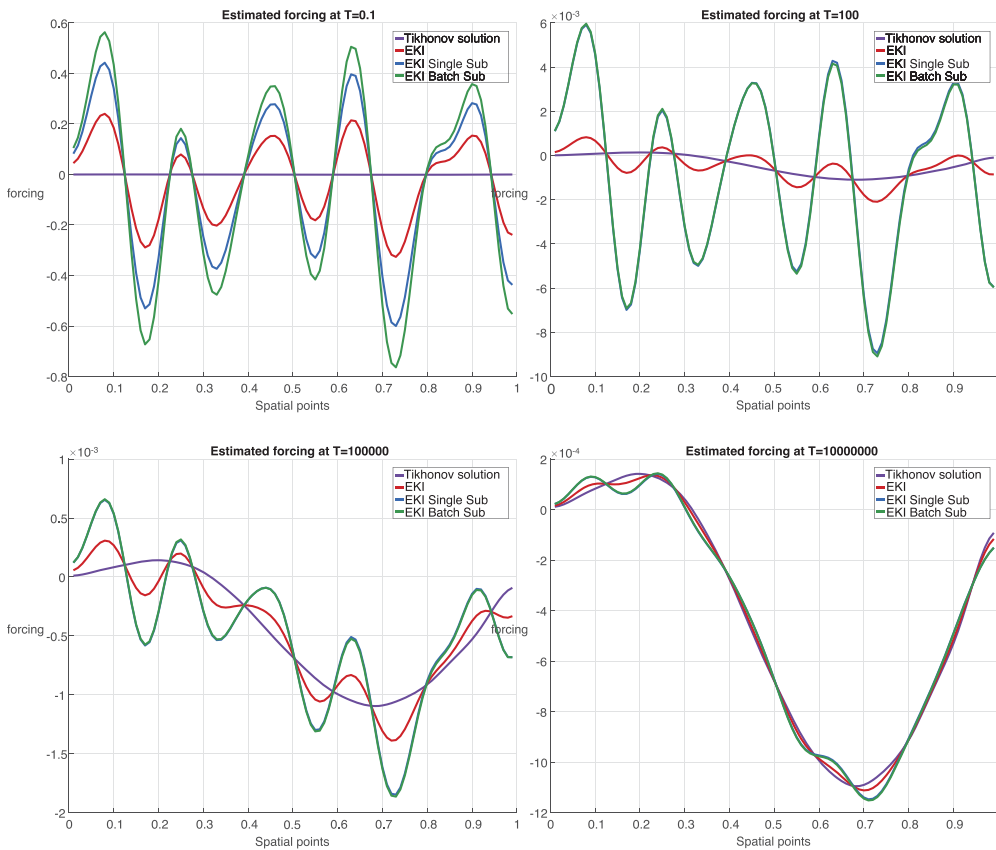


**Figure 9.** Absolute error of computed solutions in the parameter (left) and observation space (right). The red line illustrates the EKI, the blue line single-subsampling and the green line batch-subsampling.

We can see in figure 8 that the ensemble collapse also happens at an algebraic rate for EKI as well as both of our subsampling methods.

Furthermore, we illustrate the results of one single run: We simulate our prior distribution  $X$  also by the same KL-expansion, using 18 terms. We then make  $N_{\text{ens}} = 20$  independent draws to simulate our initial ensemble. Moreover, we do not use variance inflation and use a regularisation of  $\beta = 10$ , we compute the solution until time  $T = 1e7$  and use until time  $T = 10$  the same linear decaying learning rate as in the previous experiment with  $a, b = 10$ . Afterwards we consider again a constant learning rate. We obtain around 600 switching times until  $T = 10$  using those parameters.

In figure 9, we can see again that the error of all three solutions behaves similarly. Single and batch subsampling compute nearly identical solutions, therefore the error of those two methods are almost identical. The only difference to EKI is the starting time at which the solution is evaluated. The randomly computed starting time of our subsampling approaches is a little higher than the starting point of the EKI. Therefore, there is a slight shift in the error curves.



**Figure 10.** Development of computed solution over time in comparison with the Tikhonov solution at times  $T = 10^{-1}$  (upper left),  $T = 10^2$  (upper right),  $T = 10^5$  (lower left) and  $T = 10^7$  (lower right). The red line illustrates the EKI, the blue line single-subsampling, the green line batch-subsampling and purple is the reference solution  $\theta^*$ .

Furthermore, we illustrate how the computed solution behaves over time. We compare the solutions to the Tikhonov solution.

Figure 10 shows the development of the computed solution over time. Here one can see that there is a slight difference between single and batch-subsampling at  $T = 0.1$ . However, both methods quickly converge to one another as one can see in the computed results at  $T = 100$ . As mentioned above the EKI is slightly faster due to beginning a bit earlier. However at  $T = 10000000$  all methods approximate the Tikhonov solution quite similarly.

## 5.2. Nonlinear 2D Darcy flow

We introduce in this section one experiment with a non-linear forward operator  $G$ . Even though our theory only covers the linear case, we will illustrate that subsampling also leads to good results in the nonlinear setting. As subsampling strategy we only consider single-subsampling. The example is motivated by [11, 16].

Consider the following elliptic PDE.

$$\begin{cases} -\nabla \cdot (\exp(u)\nabla p) = f, & x \in D \\ p = 0, & x \in \partial D \end{cases}, \quad (5.1)$$

where  $D = (0, 1)^2$ . We seek to recover the unknown diffusion coefficient  $u^\dagger \in C^1(D) = X$ , given observation of the solution  $p \in H_0^1(D) \cap H^2(D) := \mathcal{V}$ . Furthermore, we assume that the scalar field  $f \in \mathbb{R}$  is known.

The observations are given by

$$y = \mathcal{O}(p) + \eta,$$

where  $\mathcal{O}(p) : \mathcal{V} \rightarrow \mathbb{R}^K$  is the observation Operator, that considers  $K$  randomly chosen points in  $X$ , i.e.  $\mathcal{O}(p) = (p(x_1), \dots, p(x_K))$ . Finally,  $\eta$  denotes the noise on our data and is assumed to be Gaussian, i.e. a realisation of  $\mathcal{N}(0, \Gamma)$ , where  $\Gamma = 0.1^2 \text{Id}_K$

Then our inverse problem is given by

$$y = \mathcal{G}(u) + \eta,$$

where  $\mathcal{G} = \mathcal{O} \circ G$  and  $G : X \rightarrow \mathbb{R}^K$  denotes the solution operator of the PDE (5.1). We solve the PDE on a uniform mesh with a grid size of  $h = 2^{-8}$  using an FEM method with continuous, piecewise linear finite element basis functions. We model our prior distribution as the random field

$$u(x, \omega) = \sum_{i=1}^s \lambda_i^{1/2} e_i(x) \xi_i(\omega),$$

where we have  $\lambda_i = (\pi^2(k_j^2 + l_j^2) + \tau^2)^{-\alpha}$  and  $e_i(x) = \cos(\pi x_1 k_j) \cos(\pi x_2 l_j)$  with  $\tau = 0.01, \alpha = 2, s = 25, (k_j, l_j)_{j \in \{1, \dots, s\}} \in \{1, \dots, s\}^2$ . The variables  $\xi_i$  are i.i.d standard normal variables. Afterwards we make  $N_{\text{ens}} = 10$  independent draws for our initial ensemble.

The dimension of the parameter space is  $d = 2^8$  due to the grid size. We take  $K = 30$  observations and divide them into  $N_{\text{sub}} = 5$  many subsets.

We use a linear decaying learning  $\gamma(t) = (a + bt)^{-1}$ , where  $a, b = 10$ . As regularisation factor we consider  $\beta = 10$  and compute the solution up until time  $T = 10^5$ . Again we note that due to the decrease of the switching times of the data sets the algorithm becomes computationally very slow. Therefore, we only use a linear decaying switching rate up until time  $T = 10^1$  from there on we consider  $10^5$  equidistant switching times.

Note that we again work in a subspace that is smaller than  $X$ . Therefore, we need to compare the computed solution with the respective one given the subspace.

The particles stay in the affine space  $u_0^\perp + \mathcal{E}$  for all  $t \geq 0$ . Therefore, the reference solution is given by  $u_0^\perp + u_{\mathcal{E}, j}^\dagger$ . We formulate it as a constrained optimisation problem

$$\min_{u \in \mathcal{E} + \theta_0^\perp} \frac{1}{2} \|G(u) - y\|_\Gamma^2 + \frac{\beta}{2} \|u\|^2,$$

and use MATLABs `fmincon` solver to compute it.

Figure 11 shows the computed solutions given by the different algorithms. The left picture is the solution given by MATLABs `fmincon` solver. In the middle, the solution computed by the normal EKI is depicted and on the right the result of the single-subsampling algorithm shown. One can see that both algorithms compute visually very similar solutions as the optimiser.




Figure 11. Comparison of computed diffusion coefficient  $u$ . Left figure: optimiser; middle figure: EKI; right figure: single-subsampling.





Figure 12. Mean absolute errors  $\pm$  standard deviation of computed solutions in the parameter (left) and observation space (right). The red line illustrates the EKI and the blue line single-subsampling.

Figure 12 depicts the mean errors of  $N = 32$  runs  $\pm$  one standard deviation in the parameter space (left subplot) as well as of the functionals (right subplot) evaluated in the corresponding solutions. We can see that both errors behave similarly and are converging towards zero at an algebraic rate. As in the linear example we note that the mean minus standard deviation is negative in the right picture and therefore not shown.

Figure 13 shows the mean ensemble collapse of the  $N = 32$  runs. Again the left figure shows the results of the EKI, whereas the right for single subsampling. We can see that in both methods the collapse occurs at a similar rate. One should note that single subsampling has a smoother result, this however is only due to the amount of observations where we evaluate our solutions. Due to the required frequent changes in subsampling, we obtain more observations.



**Figure 13.** Mean ensemble collapse of all  $N_{\text{ens}} = 10$  particles for the two methods. The different colours represent the ensemble collapse of one ensemble member respectively. Left figure: EKI; right figure: single-subsampling.

## 6. Conclusions

We have introduced subsampling schemes for EKI to allow the application of the method also in the large data regime. Based on recent results on continuous stochastic gradient processes [23], two subsampling approaches, (i) single subsampling, where each particle obtains the same data set when switching the data and (ii) batch-subsampling where data sets may differ for each particle, have been considered in the continuous-time setting. By applying Tikhonov regularisation and variance inflation on both methods (i) and (ii) we were able to show convergence of the schemes to the solution of the original EKI version. For the non-variance inflated variant of (i) convergence results with an algebraic rate could be proven. For batch-subsampling we were only able to show convergence when using a vanishing variance inflation over time. However, our numerical experiments in section 5.1 also showed similar convergence results for non-variance inflated batch-subsampling. The analysis requires the control of the eigenvalues of the empirical covariance w.r. to the initial ensemble. This will be subject to future work. Further, we also considered in section 5.2 a numerical experiment for a non-linear forward operator. Single-subsampling without variance inflation shows similar convergence results as the original EKI. Analysis of subsampling techniques for non-linear forward operators will be also subject for future work.

### Data availability statement

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: <https://github.com/matei1996> [36].

### Acknowledgments

M H is grateful to the DFG RTG1953 ‘Statistical Modeling of Complex Systems and Processes’ for funding of this research. J L thanks the Engineering and Physical Sciences



Research Council for their support through Grant EP/S026045/1. C S acknowledges support from MATH+ Project EF1-19: Machine Learning Enhanced Filtering Methods for Inverse Problems and EF1-20: Uncertainty Quantification and Design of Experiment for Data-Driven Control, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—The Berlin Mathematics Research Center MATH+ (EXC-2046/1, Project ID: 390685689). We are also grateful to the state of Baden-Württemberg through bwHPC.

### Appendix. Proof of theorem 3.3

We now prove theorem 3.3 which we first recall below.

**Theorem A.1 (theorem 3.3).** *Let assumption 3.2(i) and (ii) hold for a function  $h$  and a stochastic approximation process  $(\mathbf{i}(t), \theta(t))_{t \geq 0}$  with initial values  $(i_0, \theta_0) \in I \times X$ . Then,*

$$\lim_{t \rightarrow \infty} d_W(\delta(\cdot - \theta^*), \mathbb{P}(\theta(t) \in \cdot | \theta_0, i_0)) = 0.$$

To prove this theorem, we first define an auxiliary process  $(\mathbf{i}^{(\varepsilon)}(t), \theta^{(\varepsilon)}(t))_{t \geq 0}$  that converges to  $(\mathbf{i}(t), \theta(t))_{t \geq 0}$  as  $\varepsilon \rightarrow 0$ , but has a bounded transition rate matrix. For this process, we show Wasserstein ergodicity. Let  $\varepsilon > 0$  and  $B(t, \varepsilon) := A(-\log(\varepsilon + \exp(-t)))$ ,  $t \geq 0$  where  $(A(t))_{t \geq 0}$  is the transition rate matrix given in (3.3). We then define  $(\mathbf{i}^{(\varepsilon)}(t))_{t \geq 0}$  to be the stochastic process with transition rate matrix  $B(t, \varepsilon)$  and  $(\theta^{(\varepsilon)}(t))_{t \geq 0}$  to be the related stochastic approximation process. Moreover, we denote by  $K_{t|t_0}^{(\varepsilon)} := \mathbb{P}(\theta^{(\varepsilon)}(t) \in \cdot | \theta^{(\varepsilon)}(t_0) = \cdot)$  the Markov kernel associated with  $(\theta^{(\varepsilon)}(t))_{t \geq 0}$ , we ignore the underlying dependency on  $(\mathbf{i}^{(\varepsilon)}(t))_{t \geq 0}$ . Similarly, we write  $K_{t|t_0} := \mathbb{P}(\theta(t) \in \cdot | \theta(t_0) = \cdot)$ .

In a first auxiliary result, we show that  $(\theta^{(\varepsilon)}(t))_{t \geq 0}$  is ergodic and converges to a unique stationary measure.

**Lemma A.2.** *Let assumption 3.2(i) and (ii) hold and let  $\varepsilon > 0$ . Then, there is a unique probability measure  $\mu_\varepsilon$  such that for any initial distribution  $\mu_0 := \mathbb{P}(\theta^{(\varepsilon)}(0) \in \cdot)$  with finite second moment, we have*

$$\lim_{t \rightarrow \infty} d_W(\mu_\varepsilon, \mathbb{P}(\theta^{(\varepsilon)}(t) \in \cdot)) = 0.$$

**Proof.** 1. Let  $\mu_0, \mu_0^\dagger$  be to different initial distributions. Moreover, let  $(\theta^{(\varepsilon)}(t))_{t \geq 0}, (\theta_{\dagger}^{(\varepsilon)}(t))_{t \geq 0}$  be two realisations of the stochastic approximation process with  $\mu_0 = \mathbb{P}(\theta^{(\varepsilon)}(0) \in \cdot)$  and  $\mu_0^\dagger = \mathbb{P}(\theta_{\dagger}^{(\varepsilon)}(0) \in \cdot)$ . Furthermore, we assume that the processes are coupled through the associated index processes. Indeed, we assume that  $(\mathbf{i}^{(\varepsilon)}(t))_{t \geq 0}$  and  $(\mathbf{i}_{\dagger}^{(\varepsilon)}(t))_{t \geq 0}$  are almost surely identical. Then,

$$d_W(\mu_0 K_{t|0}^{(\varepsilon)}, \mu_0^\dagger K_{t|0}^{(\varepsilon)})^2 \leq \mathbb{E}[\|\theta^{(\varepsilon)}(t) - \theta_{\dagger}^{(\varepsilon)}(t)\|^2] \quad (t \geq 0).$$

Assumption 3.2(ii) implies that

$$\begin{aligned} \frac{d\|\theta^{(\varepsilon)}(t) - \theta_{\dagger}^{(\varepsilon)}(t)\|^2}{dt} &= 2\langle \mathbf{F}_i(\theta^{(\varepsilon)}(t)) - \mathbf{F}_i(\theta_{\dagger}^{(\varepsilon)}(t)), \theta^{(\varepsilon)}(t) \\ &\quad - \theta_{\dagger}^{(\varepsilon)}(t) \rangle_X \leq -2h(t)\|\theta^{(\varepsilon)}(t) - \theta_{\dagger}^{(\varepsilon)}(t)\|^2. \end{aligned}$$



Now, the Grönwall inequality implies that

$$\|\theta^{(\varepsilon)}(t) - \theta_{\dagger}^{(\varepsilon)}(t)\|^2 \leq \exp\left(-2 \int_0^t h(t) dt\right) \|\theta^{(\varepsilon)}(0) - \theta_{\dagger}^{(\varepsilon)}(0)\|^2.$$

Taking expectations on both sides, we obtain:

$$d_W(\mu_0 K_{t|0}^{(\varepsilon)}, \mu_0^\dagger K_{t|0}^{(\varepsilon)})^2 \leq \exp\left(-2 \int_0^t h(t) dt\right) \mathbb{E} \left[ \|\theta^{(\varepsilon)}(0) - \theta_{\dagger}^{(\varepsilon)}(0)\|^2 \right] \quad (\text{A.1})$$

which converges to 0 as  $t \rightarrow \infty$  by assumption 3.2(ii).

2. To show the assertion of the theorem, we now choose some  $\delta > 0$  and consider the process  $\theta_{\dagger}^{(\varepsilon)}(t) := \theta^{(\varepsilon)}(t + \delta)$ . The contraction property (A.1) implies that

$$d_W(\mu_0 K_{t|0}^{(\varepsilon)}, \mu_0 K_{t+\delta|0}^{(\varepsilon)}) \rightarrow 0, \quad (t \rightarrow \infty).$$

Thus, the sequence  $(\mu_0 K_{\delta n|0}^{(\varepsilon)})_{n=0}^\infty$  is a Cauchy sequence on the Wasserstein space associated to  $d_W$ . Thus, due to the completeness of Wasserstein spaces, see [7], we have that  $\mu_0 K_{t|0}$  converges to some probability distribution  $\mu_\varepsilon$  as  $t \rightarrow \infty$  in  $d_W$ . Again due to the contraction given in (A.1),  $\mu_\varepsilon$  does not depend on the initial distribution  $\mu_0$  and is, thus, unique.  $\square$

In the second auxiliary result, we show that  $(\theta^{(\varepsilon)}(t))_{t \geq 0}$  converges to  $(\theta(t))_{t \geq 0}$ , as  $\varepsilon \downarrow 0$ , and that  $\pi^\varepsilon$  converges to  $\delta(\cdot - \theta^*)$ . This result is a small extension of proposition 4 in [23] and the proof proceeds identically.

**Proposition A.3.** *Let assumption 3.2(i) and (ii) hold. Then, we have*

1.  $d_W(\mu_0 K_{t|t_0}^{(\varepsilon)}, \mu_0 K_{t|t_0}) \leq \alpha'(\varepsilon)$  ( $\varepsilon > 0$ ) for any initial distribution  $\mu_0$ ,
2.  $d_W(\mu_\varepsilon, \delta(\cdot - \theta^\varepsilon)) \leq \alpha''(\varepsilon)$  ( $\varepsilon > 0$ ),

where  $\alpha', \alpha'' : [0, \infty) \rightarrow [0, \infty)$  are continuous and equal to 0 at 0.

The proof of theorem 3.3 now consists in a simple rearrangement of the auxiliary results above.

**Proof of theorem 3.3.** By the triangular inequality, we have

$$d_W(\delta(\cdot - \theta^*), \mathbb{P}(\theta(t) \in \cdot | \theta_0, i_0)) \leq d_W(\delta(\cdot - \theta^*), \mu_\varepsilon) + d_W(\mu_\varepsilon, \mathbb{P}(\theta^{(\varepsilon)}(t) \in \cdot)) + d_W(\mathbb{P}(\theta^{(\varepsilon)}(t) \in \cdot), \mathbb{P}(\theta(t) \in \cdot))$$

where the last term in the sum is identical to  $d_W(\mu_0 K_{t|t_0}^{(\varepsilon)}, \mu_0 K_{t|t_0}^{(\varepsilon)})$ . By lemma A.2 and proposition A.3, we have  $d_W(\delta(\cdot - \theta^*), \mathbb{P}(\theta(t) \in \cdot | \theta_0, i_0)) \rightarrow 0$ , as  $t \rightarrow \infty$ .  $\square$

**Lemma A.4 (lemma 4.4).** *Given the  $N_{sub}$  subsamples  $(\tilde{y}_{(1)}^\dagger, \dots, \tilde{y}_{(N_{sub})}^\dagger)$ , assume that the centred initial ensemble is a generator of the full space  $X$ , i.e.  $\text{span}\{e_0^{(j)}, j \in J\} = X$ . We further assume that  $\sum_{j=1}^{N_{ens}} (\theta_j^\dagger - \bar{\theta}^\dagger)(\theta_j^\dagger - \bar{\theta}^\dagger)^\top$  has full rank  $d$ . Then the particles converge to the true solution  $\theta^\dagger$  exponentially fast, i.e.  $\theta^{(j)} \rightarrow \theta_i^\dagger$  with  $\theta_i^\dagger$  denoting the minimiser of  $\Phi_i(\theta) = \frac{1}{2} \|\tilde{A}_i \theta - \tilde{y}_i\|^2$ .*

**Proof.** Note that we do not switch the data. Therefore, the subset that each particle obtains does not depend on time.

The gradients  $g^{(j)} = \nabla \Phi_{(i;j)}(\theta) = \tilde{A}_{(i;j)}^\top (\tilde{A}_{(i;j)} \theta^{(j)} - \tilde{y}_{(i;j)})$  satisfy

$$\frac{d\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)} u^{(j)}(t)}{dt} = -\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)} \widehat{C}_t \tilde{A}_{(i;j)}^\top (\tilde{A}_{(i;j)} \theta^{(j)}(t) - \tilde{y}_{(i;j)}).$$

We will prove in the following that  $g^{(j)} \rightarrow 0$  exponentially fast as  $t \rightarrow \infty$ . This is a sufficient and necessary optimality condition as  $A_{i^{(j)}}^\top \tilde{A}_{(i;j)}$  is positive definite due to regularisation. We obtain

$$\frac{1}{2} \frac{d\|g^{(j)}\|_{\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)}}^2}{dt} = -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} \langle e^{(k)}, g^{(j)} \rangle^2 \leq 0, \quad (\text{A.2})$$

i.e. the gradients are monotonically decreasing. To prove convergence, we will now show that  $\frac{1}{2} \frac{d}{dt} \|g^{(j)}\|_{\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)}}^2 < 0$ . Note that, if  $g^{(j)} \neq 0$  and  $\{e^{(k)}\}_{k=1}^{N_{\text{ens}}}$  is still a generating set of  $X$  at time  $t$ , then there exists at least one  $k \in \{1, \dots, N_{\text{ens}}\}$  such that  $\langle e^{(k)}, g^{(j)} \rangle \neq 0$ . The quantity  $e^{(j)}$  satisfies

$$\frac{de^{(j)}(t)}{dt} = -\widehat{C}_t v^{(j)} = -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} \langle v^{(j)}, e^{(k)} \rangle e^{(k)},$$

with  $v^{(j)} = g^{(j)} - \bar{g}$ . Thus the dynamical behaviour of empirical covariance is given by

$$\frac{d}{dt} \widehat{C} = D\widehat{C} + \widehat{C}D^\top$$

with  $D = -\frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} v^{(j)} \otimes e^{(j)}$ .

Therefore, the rank of the empirical covariance stays constant over time (see [27]) and the members  $\{e^{(j)}\}_{j=1}^{N_{\text{ens}}}$  still form a generating set of  $X$  at time  $t$ .

Then, there exists at least one  $k$  in (A.2) such that  $\langle e^{(k)}, g^{(j)} \rangle^2 \neq 0$ , i.e. the gradients converge to 0.

This implies the convergence  $\theta^{(j)} \rightarrow \theta_j^\dagger$  due to the strong convexity. By assumption, the limit of the empirical covariance has full rank  $d$ , since  $\theta_j \rightarrow \theta_j^\dagger$ , i.e. the minimal eigenvalue  $\lambda_{\min}(\widehat{C}_t)$  of the empirical covariance is bounded from below uniformly in time. Thus, we have

$$\frac{1}{2} \frac{d\|g^{(j)}\|_{\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)}}^2}{dt} = -\langle g^{(j)}, C(t)g^{(j)} \rangle \leq -\frac{\lambda_{\min}}{\lambda_{\max}((\tilde{A}_i^\top \tilde{A}_i)^{-1})} \|g^{(j)}\|_{\tilde{A}_{(i;j)}^\top \tilde{A}_{(i;j)}}^2,$$

where  $\lambda_{\min} > 0$  denotes the lower bound on the minimal eigenvalue of the empirical covariance.  $\square$

**Lemma A.5.** *The particles  $\theta^{(j)}$  converge at exponential speed to the unique solution of the regularised data misfit, i.e.  $\rho^{(j)}(t) = A_{i^{(j)}} \theta^{(j)} - y_{i^{(j)}} \rightarrow 0$ . Hence, there exists a (unique)  $\theta_j^\dagger \in S = \text{span}\{\theta_0^{(1)}, \dots, \theta_0^{(j)}\}$  such that  $\theta^{(j)} \rightarrow \theta_j^\dagger$ .*

Furthermore, let  $\theta_1$  and  $\theta_2$  be two coupled process with different initial values  $\theta_1(0), \theta_2(0)$ . Then there exists a measurable function  $h: [0, \infty) \rightarrow \mathbb{R}$  such that the following holds

$$-\langle \theta_1 - \theta_2, (\widehat{C}_t^1 + \alpha_{vi} C_{vi}) D_\theta \Phi_{i'(t, \varepsilon; \cdot)}(\theta_1) - (\widehat{C}_t^1 + \alpha_{vi} C_{vi}) D_\theta \Phi_{i'(t, \varepsilon; \cdot)}(\theta_2) \rangle \leq -h(t) \|\theta_1 - \theta_2\|^2,$$

for  $t$  large enough. We have

1. *Single-Subsampling with variance inflation:*

$$h(t) = \alpha_{vi} \lambda_{\min}(C_{vi}) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(\tilde{A}_i^T \tilde{A}_i).$$

2. *Batch-Subsampling with variance inflation:*

$$h(t) = \alpha_{vi} \lambda_{\min}(C_{vi}) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(A_i^T A_i)$$

3. *Single-Subsampling without variance inflation:*

$$h(t) = \lambda_{\min}(\hat{C}_t) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(\tilde{A}_i^T \tilde{A}_i).$$

4. *Batch-Subsampling with diminishing variance inflation:*

$$h(t) = \frac{\alpha_{vi}}{t} \lambda_{\min}(C_{vi}) \min_{i \in \{1, \dots, N_{sub}\}} \lambda_{\min}(A_i^T A_i).$$

**Proof.** 1. We first consider single subsampling with variance inflation

The gradients  $\tilde{A}_{(i,j)}^T \rho^{(j)}$  for a constant (w.r. to time and particle) data stream satisfy the following differential equation for all subsets  $i$ .

$$\frac{d\tilde{A}_i^T \rho^{(j)}(t)}{dt} = -\tilde{A}_i^T \tilde{A}_i [\hat{C}_t + \alpha_{vi} C_{vi}] \tilde{A}_i^T \rho^{(j)}(t).$$

The norm of the gradients thus satisfies

$$\frac{1}{2} \frac{d\|\tilde{A}_i^T \rho^{(j)}(t)\|_{\tilde{A}_i^T \tilde{A}_i}^2}{dt} \leq -\alpha_{vi} \frac{\lambda_{\min}(C_{vi})}{\lambda_{\max}(\tilde{A}_i^T \tilde{A}_i)} \|\rho^{(j)}(t)\|_{\tilde{A}_i^T \tilde{A}_i}^2,$$

which implies the exponential convergence of the mapped residuals and with the injectivity of the modified forward operator the exponential convergence in the parameter space to the (unique) solution of the regularised data misfit.

Therefore, we have  $\theta^{(j)} \rightarrow \theta_j^\dagger$ .

Then  $\theta_j^\dagger$  is an equilibrium point of

$$\mathbf{F}_i(\theta^{(j)}, t) = (\hat{C}_t + \alpha_{vi} C_{vi}) D_\theta \Phi_i^{\text{reg}}(\theta^{(j)}(t)) \quad (j \in J).$$

Hence, there exists a function  $\kappa(t) \geq 0$  for all  $t > 0$  that converges exponentially fast to 0 for  $t \rightarrow \infty$  such that:

$$\|\theta_t^{(j)} - P_Y \theta_j^\dagger\| \leq \kappa(t) \quad \forall t \geq 0.$$

Due to linearity w.r. to the initial values, we obtain

$$\|\theta(t, \theta_0)^{(j)} - \theta(t, \theta_1)^{(j)}\| = \|\theta(t, \theta_0)^{(j)} - P_Y \theta_j^\dagger - (\theta(t, \theta_1)^{(j)} - P_Y \theta_j^\dagger)\| \leq 2\kappa(t).$$

The next step is to consider equation (ii) from assumption 3.2. Note that  $\theta_1$  and  $\theta_2$  are column vectors consisting of the stacked particle vectors.

Therefore, we define the following matrices to represent the gradient flow for the stacked vector. We set:  $\tilde{A} = \text{diag}\{\tilde{A}_1, \tilde{A}_2, \dots, \tilde{A}_{N_{sub}}\}$ ,  $\tilde{A}^T = \text{diag}\{\tilde{A}_1^T, \tilde{A}_2^T, \dots, \tilde{A}_{N_{sub}}^T\}$ ,  $\mathbf{C}_t = \text{diag}\{\hat{C}_t + \alpha_{v1} C_{v1}, \hat{C}_t + \alpha_{v2} C_{v2}, \dots, \hat{C}_t + \alpha_{vN_{sub}} C_{vN_{sub}}\}$ .

Then the dynamics are given by:

$$\frac{d\theta}{dt} = -\hat{C}_t \tilde{A}^T (\tilde{A} \theta - y).$$

We want to show:

$$-\langle \theta_1 - \theta_2, \mathbf{C}_t^1 D_\theta \Phi_{i'(t, \varepsilon; \cdot)}(\theta_1) - \mathbf{C}_t^2 D_\theta \Phi_{i'(t, \varepsilon; \cdot)}(\theta_2) \rangle \leq -h_{i'}(t) \|\theta_1 - \theta_2\|^2,$$

We can split the left hand side into two parts:

$$\begin{aligned} & - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_1) - \mathbf{C}_t^2 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) \rangle \\ & = - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_1) - \mathbf{C}_t^1 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) \\ & \quad + \mathbf{C}_t^1 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) - \mathbf{C}_t^2 D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) \rangle \\ & = - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 [D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_1) - D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2)] \rangle \\ & \quad - \langle \theta_1 - \theta_2, [\mathbf{C}_t^1 - \mathbf{C}_t^2] D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) \rangle. \end{aligned}$$

Substituting the corresponding gradient flows into the equations, we obtain

$$\begin{aligned} & - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 [D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_1) - D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2)] \rangle \\ & \quad - \langle \theta_1 - \theta_2, [\mathbf{C}_t^1 - \mathbf{C}_t^2] D_\theta \Phi_{i'(t,\varepsilon;\cdot)}(\theta_2) \rangle \\ & = - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 \tilde{A}^T \tilde{A} (\theta_1 - \theta_2) \rangle \\ & \quad - \langle \theta_1 - \theta_2, [\mathbf{C}_t^1 - \mathbf{C}_t^2] \tilde{A}^T (\tilde{A} \theta_2 - \tilde{y}) \rangle. \end{aligned}$$

We consider both terms separately. For the first part we obtain

$$\begin{aligned} - \langle \theta_1 - \theta_2, \mathbf{C}_t^1 \tilde{A}^T \tilde{A} (\theta_1 - \theta_2) \rangle & \leq - \lambda_{\min}(\mathbf{C}_t^1) \lambda_{\min}(\tilde{A}^T \tilde{A}) \|\theta_1 - \theta_2\| \\ & = - \lambda_{\min}(\hat{\mathbf{C}}_t^1 + \alpha_{vi} \mathbf{C}_{vi}) \lambda_{\min}(\tilde{A}^T \tilde{A}) \|\theta_1 - \theta_2\| \\ & \leq - \alpha_{vi} \lambda_{\min}(\mathbf{C}_{vi}) \left( \min_{i \in \{1, \dots, N_{\text{sub}}\}} \lambda_{\min}(\tilde{A}_i^T \tilde{A}_i) \right) \|\theta_1 - \theta_2\|^2, \end{aligned}$$

where we used the positive definiteness of  $\hat{\mathbf{C}}_t^1$  for every  $t \geq 0$  in the third step. For the second term we obtain

$$\begin{aligned} - \langle \theta_1 - \theta_2, [\mathbf{C}_t^1 - \mathbf{C}_t^2] \tilde{A}^T (\tilde{A} \theta_2 - \tilde{y}) \rangle & \leq |\langle \theta_1 - \theta_2, [\mathbf{C}_t^1 - \mathbf{C}_t^2] \tilde{A}^T (\tilde{A} \theta_2 - \tilde{y}) \rangle| \\ & \leq \|\theta_1 - \theta_2\| \|\mathbf{C}_t^1 - \mathbf{C}_t^2\| \|\tilde{A}^T (\tilde{A} \theta_2 - \tilde{y})\|. \end{aligned}$$

We can compare the rates of convergence. Considering the results from above we have

$$\|\theta_1 - \theta_2\|^2 \in \mathcal{O}(\kappa(t)^2),$$

and also

$$\|\tilde{A}^T (\tilde{A} \theta_2 - \tilde{y})\| \in \mathcal{O}(\kappa(t)).$$

Finally, we have for the covariance matrices

$$\begin{aligned} \|\mathbf{C}_t^1 - \mathbf{C}_t^2\| & = \left\| \frac{1}{J} \sum_{j=1}^J (u_1^{(j)} - \bar{u}_1)(u_1^{(j)} - \bar{u}_1)^T - (u_2^{(j)} - \bar{u}_2)(u_2^{(j)} - \bar{u}_2)^T \right\| \\ & \leq \frac{1}{J} \sum_{j=1}^J \|u_1^{(j)}(u_1^{(j)})^T - u_2^{(j)}(u_2^{(j)})^T\| + \|u_1^{(j)}(\bar{u}_1)^T - u_2^{(j)}(\bar{u}_2)^T\| \\ & \quad + \|\bar{u}_1(u_1^{(j)})^T - \bar{u}_2(u_2^{(j)})^T\| + \|\bar{u}_1(\bar{u}_1)^T - \bar{u}_2(\bar{u}_2)^T\|. \end{aligned}$$

Since we know that all particles converge with rate  $\kappa(t)$ , the mean values also converge with the same rate. By using triangle inequality we obtain the following

$$\|\mathbf{C}_t^1 - \mathbf{C}_t^2\| \in \mathcal{O}(\kappa(t)).$$

Hence

$$\|\theta_1 - \theta_2\| \|\mathbf{C}_t^1 - \mathbf{C}_t^2\| \|\tilde{\mathbf{A}}^T(\tilde{\mathbf{A}}\theta_2 - \tilde{\mathbf{y}})\| \in \mathcal{O}(\kappa(t)^3),$$

showing that the second term converges faster and we can therefore neglect it.

2. In batch subsampling with variance inflation the only difference is that the forward operator and data both depend on the particle, i.e. the mapped residuals satisfy

$$\frac{d\rho^{(j)}}{dt} = -A_{i(t;j)} \left[ \hat{\mathbf{C}}_t + \alpha_{vi} C_{vi} \right] A_{i(t;j)}^T \rho^{(j)}(t)$$

with  $\rho^{(j)} = A_{i(t;j)} \theta^{(j)} - \tilde{y}_{i(t;j)}$ . The exponential convergence of each particle to the minimiser of the functional  $\frac{1}{2} \|\tilde{\mathbf{A}}_{i(t;j)} \theta^{(j)} - \tilde{y}_{i(t;j)}\|^2 + \frac{\alpha}{2} \|\theta\|_{C_0}$  follows again from standard arguments with the Lyapunov function  $\|\tilde{\mathbf{A}}_{i(t;j)}^T \rho^{(j)}\|^2$ . Hence, the convexity analysis does not change.

3. If we do not use variance inflation the gradients  $\tilde{\mathbf{A}}_{i(t;j)}^T \rho^{(j)}$  satisfy the following differential equation

$$\frac{d\tilde{\mathbf{A}}_{i(t;j)}^T \rho^{(j)}(t)}{dt} = -\tilde{\mathbf{A}}_{i(t;j)}^T \tilde{\mathbf{A}}_{i(t;j)} \hat{\mathbf{C}}_t \tilde{\mathbf{A}}_{i(t;j)}^T \rho^{(j)}(t).$$

Again, by basic Lyapunov theory we obtain convergence at an algebraic speed, which however is enough to do the same analysis as above. Similar to above we obtain  $h(t) = \lambda_{\min}(\hat{\mathbf{C}}_t) (\min_{i \in \{1, \dots, N_{\text{sub}}\}} \lambda_{\min}(A_i^T A_i))$

4. Batch subsampling with diminishing variance inflation: theorem 4.4 gives us the exponential convergence to the respective solution. Then the convexity analysis is similar to the above calculations and we obtain  $h(t) = \frac{\alpha_{vi}}{1+t} \lambda_{\min}(C_{vi}) (\min_{i \in \{1, \dots, N_{\text{sub}}\}} \lambda_{\min}(A_i^T A_i))$ .  $\square$

#### Subsampling without variance inflation

**Lemma A.6.** For the regularised single-subsampling EKI flow, given by the solution of (4.1), the following lower bound for the smallest eigenvalue  $\lambda_{\min}(t)$  of the empirical covariance  $\hat{\mathbf{C}}(t)$  holds

$$\lambda_{\min}(t) \geq \left( 2ct + \frac{1}{\lambda_{\min}(0)} \right)^{-1}.$$

**Proof.** The proof follows the ideas used in [12, theorem 3.5].

W.l.o.g we assume that our initial ensemble is a generator of  $\mathbb{R}^d$ . Otherwise we consider the dynamics of the corresponding coordinates, which correspond to the minimisation problem (2.8). The particles satisfy (3.4). Substituting the covariance matrix

$$\widehat{C}_t = \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \left( \theta^{(j)}(t) - \bar{\theta}_t \right) \left( \theta^{(j)}(t) - \bar{\theta}_t \right)^T = \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} e_t^{(j)} \left( e_t^{(j)} \right)^T$$

into the dynamics of the particles gives

$$\begin{aligned} \frac{d\theta^{(j)}(t)}{dt} &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} e_t^{(k)} \left( e_t^{(k)} \right)^T \tilde{A}_{i(t)}^T \left( \tilde{A}_{i(t)} \theta^{(j)}(t) - \tilde{y}_{i(t)} \right) \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} e_t^{(k)} \left( \tilde{A}_{i(t)} e_t^{(k)} \right)^T \left( \tilde{A}_{i(t)} \theta^{(j)}(t) - \tilde{y}_{i(t)} \right) \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} D_{kj} e_t^{(k)}, \end{aligned}$$

where we set  $D_{kj} := \langle \tilde{A}_{i(t)} e_t^{(k)}, (\tilde{A}_{i(t)} \theta^{(j)}(t) - \tilde{y}_{i(t)}) \rangle$ .

Next we consider the dynamics of the weighted particles, i.e.

$$\begin{aligned} \frac{d\bar{\theta}(t)}{dt} &= \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \frac{d\theta^{(j)}(t)}{dt} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} D_{kj} e_t^{(k)} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} \langle \tilde{A}_{i(t)} e_t^{(k)}, (\tilde{A}_{i(t)} \theta^{(j)}(t) - \tilde{y}_{i(t)}) \rangle e_t^{(k)} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} F_k e_t^{(k)}, \end{aligned}$$

where we set  $F_k := \langle \tilde{A}_{i(t)} e_t^{(k)}, (\tilde{A}_{i(t)} \bar{\theta} - \tilde{y}_{i(t)}) \rangle$ .

The difference of the scalars  $D_{jk}$  and  $F_k$  is given by

$$\begin{aligned} &\langle \tilde{A}_{i(t)} e_t^{(k)}, (\tilde{A}_{i(t)} \theta^{(j)}(t) - \tilde{y}_{i(t)}) \rangle - \langle \tilde{A}_{i(t)} e_t^{(k)}, (\tilde{A}_{i(t)} \bar{\theta} - \tilde{y}_{i(t)}) \rangle \\ &= \langle \tilde{A}_{i(t)} e_t^{(k)}, \tilde{A}_{i(t)} e_t^{(j)} \rangle := E_{kj}. \end{aligned}$$

Obviously we have  $E_{kj} = E_{jk}$ . With this we can quantify the dynamics of the centred particles, i.e.

$$\begin{aligned} \frac{de^{(j)}(t)}{dt} &= \frac{d\theta^{(j)}(t) - \bar{\theta}(t)}{dt} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} D_{kj} e_t^{(k)} + \frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} F_k e_t^{(k)} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} (D_{kj} - F_k) e_t^{(k)} \\ &= -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} E_{kj} e_t^{(k)}. \end{aligned}$$

Finally, we obtain for the dynamics of the empirical covariance  $\widehat{C}_t$

$$\begin{aligned} \frac{d\widehat{C}(t)}{dt} &= \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \frac{d}{dt} \left( e_t^{(j)} \left( e_t^{(j)} \right)^T \right) \\ &= \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} \left( -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} E_{kj} e_t^{(k)} \right) \left( e_t^{(j)} \right)^T + \frac{1}{N_{\text{ens}}} \sum_{j=1}^{N_{\text{ens}}} e_t^{(j)} \left( -\frac{1}{N_{\text{ens}}} \sum_{k=1}^{N_{\text{ens}}} E_{kj} e_t^{(k)} \right)^T \\ &= -\frac{2}{N_{\text{ens}}^2} \sum_{j,k=1}^{N_{\text{ens}}} E_{kj} e_t^{(k)} \left( e_t^{(j)} \right)^T. \end{aligned}$$

Now let  $\lambda_{\min}(t)$  be the smallest eigenvalue of  $\widehat{C}(t)$  with unit-norm eigenvector  $v(t)$ . Then we have

$$0 = \frac{d}{dt} \|v(t)\|_X^2 = 2 \langle v(t), \frac{d}{dt} v(t) \rangle.$$

The dynamics of the smallest eigenvalue are then given by:

$$\begin{aligned} \frac{d\lambda_{\min}(t)}{dt} &= \frac{d}{dt} \langle v(t), \widehat{C}(t)v(t) \rangle \\ &= \langle v(t), \frac{d}{dt} (\widehat{C}(t)v(t)) \rangle + \langle \frac{d}{dt} v(t), \widehat{C}(t)v(t) \rangle \\ &= \langle v(t), \frac{d}{dt} (\widehat{C}(t)v(t)) \rangle + \lambda_{\min} \langle \frac{d}{dt} v(t), v(t) \rangle \\ &= \langle v(t), \frac{d}{dt} (\widehat{C}(t)v(t)) \rangle \\ &= -\frac{2}{N_{\text{ens}}^2} \sum_{j,k=1}^{N_{\text{ens}}} E_{kj} \langle v(t), e_t^{(k)} \left( e_t^{(j)} \right)^T v \rangle. \end{aligned}$$

Note that the following holds

$$\langle v(t), e_t^{(k)} \left( e_t^{(j)} \right)^T v \rangle = \langle \left( e_t^{(k)} \right)^T v(t), \left( e_t^{(j)} \right)^T v \rangle = \langle e_t^{(k)}, v \rangle \langle e_t^{(j)}, v \rangle.$$

Substituting this and  $E_{kj}$  into the latter equation, gives us

$$\begin{aligned} \frac{d\lambda_{\min}(t)}{dt} &= -\frac{2}{N_{\text{ens}}^2} \sum_{j,k=1}^{N_{\text{ens}}} \langle \tilde{A}_{i(t)} e_t^{(k)}, \tilde{A}_{i(t)} e_t^{(j)} \rangle \langle e_t^{(k)}, v \rangle \langle e_t^{(j)}, v \rangle \\ &= -\frac{2}{N_{\text{ens}}^2} \sum_{j,k=1}^{N_{\text{ens}}} \langle \tilde{A}_{i(t)} e_t^{(k)} \langle e_t^{(k)}, v \rangle, \tilde{A}_{i(t)} e_t^{(j)} \langle e_t^{(j)}, v \rangle \rangle \\ &= -2 \langle \tilde{A}_{i(t)} \widehat{C}(t)v(t), \tilde{A}_{i(t)} \widehat{C}(t)v(t) \rangle \\ &= -2 \| \tilde{A}_{i(t)} \widehat{C}(t)v(t) \|^2 \\ &\geq -2 \| \tilde{A}_{i(t)} \|^2 \lambda_{\min}^2(t). \end{aligned}$$

Setting  $c = \max_{i \in \{1, \dots, N_{\text{sub}}\}} \|A_{i(t)}\|^2$ , we obtain for this ODE the following lower bound for the solution

$$\lambda_{\min}(t) \geq \left( 2ct + \frac{1}{\lambda_{\min}(0)} \right)^{-1}.$$

## ORCID iDs

Matei Hanu  <https://orcid.org/0000-0002-6494-8971>

Jonas Latz  <https://orcid.org/0000-0002-4600-0247>

## References

- [1] Anderson W 1991 *Continuous-Time Markov Chains: An Applications-Oriented Approach (Applied Probability)* (Springer)
- [2] Benaïm M, Borgne S L, Malrieu F and Zitt P-A 2012 Quantitative ergodicity for some switched dynamical systems *Electron. Commun. Probab.* **17** 1–14
- [3] Bergemann K and Reich S 2009 A localization technique for ensemble Kalman filters *Q. J. E. Meteorol. Soc.* **136** 701–7
- [4] Bergemann K and Reich S 2010 A mollified ensemble Kalman filter *Q. J. R. Meteorol. Soc.* **136** 1636–43
- [5] Blömker D, Schillings C, Wacker P and Weissmann S 2019 Well posedness and convergence analysis of the ensemble Kalman inversion *Inverse Problems* **35** 085007
- [6] Blömker D, Schillings C, Wacker P and Weissmann S 2021 Continuous time limit of the stochastic ensemble Kalman inversion: strong convergence analysis *SIAM J. Numer. Anal.* **60** 3181–215
- [7] Bolley F 2008 *Separability and Completeness for the Wasserstein Distance* (Springer) pp 371–7
- [8] Bottou L, Curtis F E and Nocedal J 2016 Optimization methods for large-scale machine learning *SIAM Rev.* **60** 223–311
- [9] Bungert L and Wacker P 2021 Complete deterministic dynamics and spectral decomposition of the linear ensemble Kalman inversion *SIAM/ASA J. Uncertain. Quantif.* **11** 320–57
- [10] Calvello E, Reich S and Stuart A M 2022 Ensemble Kalman methods: a mean field perspective (arXiv:2209.11371)
- [11] Chada N K, Schillings C, Tong X T and Weissmann S 2020 Consistency analysis of bilevel data-driven learning in inverse problems *Commun. Math. Sci.* **20** 123–64
- [12] Chada N K, Stuart A M and Tong X T 2020 Tikhonov regularization within ensemble Kalman inversion *SIAM J. Numer. Anal.* **58** 1263–94
- [13] Choromanska A, Henaff M, Mathieu M, Arous G B and LeCun Y 2014 The loss surfaces of multilayer networks *18th Int. Conf. on Artificial Intelligence and Statistics (AISTATS) (San Diego, CA, USA, 2015)* PLMR 38 pp 192–204
- [14] Ding Z and Li Q 2020 Ensemble Kalman inversion: mean-field limit and convergence analysis *Stat. Comput.* **31** 9
- [15] Evensen G 2003 The Ensemble Kalman filter: theoretical formulation and practical implementation *Ocean Dyn.* **53** 343–67
- [16] Garbuno-Inigo A, Hoffmann F, Li W and Stuart A M 2020 Interacting langevin diffusions: Gradient structure and ensemble kalman sampler *SIAM J. Appl. Dyn. Syst.* **19** 412–41
- [17] Gillespie D T 1977 Exact stochastic simulation of coupled chemical reactions *J. Phys. Chem.* **81** 2340–61
- [18] Iglesias M A 2014 Iterative regularization for ensemble data assimilation in reservoir models *Comput. Geosci.* **19** 177–212
- [19] Iglesias M A 2016 A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems *Inverse Problems* **32** 025002
- [20] Iglesias M A, Law K J H and Stuart A M 2013 Ensemble Kalman methods for inverse problems *Inverse Problems* **29** 045001
- [21] Jin K, Latz J, Liu C and Schönlieb C 2021 A continuous-time stochastic gradient descent method for continuous data (arXiv:2112.03754)
- [22] Kovachki N B and Stuart A M 2019 Ensemble Kalman inversion: a derivative-free technique for machine learning tasks *Inverse Problems* **35** 095005



- [23] Latz J 2021 Analysis of stochastic gradient descent in continuous time *Stat. Comput.* **31** 39
- [24] Li G and Reynolds A 2009 Iterative ensemble Kalman filters for data assimilation *SPE J.* **14** 496–505
- [25] Li Q, Tai C and E W 2019 Stochastic modified equations and dynamics of stochastic gradient algorithms i: mathematical foundations *J. Mach. Learn. Res.* **20** 40
- [26] Nocedal J and Wright S 2006 *Numerical Optimization (Operations Research and Financial Engineering)* (Springer)
- [27] Reid W 1946 A matrix differential equation of Riccati type *Am. J. Math.* **68** 237–46
- [28] Robbins H and Monro S 1951 A stochastic approximation method *Ann. Math. Stat.* **22** 400–7
- [29] Sanz-Alonso D, Stuart A M and Taeb A 2018 Inverse problems and data assimilation with connections to machine learning (arXiv:1810.06191)
- [30] Schillings C and Stuart A 2017 Convergence analysis of ensemble Kalman inversion: the linear, noisy case *Appl. Anal.* **97** 107–23
- [31] Schillings C and Stuart A M 2016 Analysis of the ensemble Kalman filter for inverse problems *SIAM J. Numer. Anal.* **55** 1264–90
- [32] Tong X T, Majda A J and Kelly D 2015 Nonlinear stability of the ensemble Kalman filter with adaptive covariance inflation *Commun. Math. Sci.* **14** 1283–313
- [33] Tong X T, Majda A J and Kelly D 2016 Nonlinear stability and ergodicity of ensemble based Kalman filters *Nonlinearity* **29** 657–91
- [34] Vidal R, Bruna J, Giryes R and Soatto S 2017 Mathematics of deep learning (arXiv:1712.04741)
- [35] Weissmann S, Chada N K, Schillings C and Tong X T 2022 Adaptive Tikhonov strategies for stochastic ensemble Kalman inversion *Inverse Problems* **38** 045009
- [36] Hanu M 2023 *EKI-subsampling* (available at: <https://github.com/matei1996/EKI-Subsampling>)