

PAPER • OPEN ACCESS

Adaptive regularisation for ensemble Kalman inversion

To cite this article: Marco Iglesias and Yuchen Yang 2021 *Inverse Problems* **37** 025008

View the [article online](#) for updates and enhancements.



IOP | ebooksTM

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Adaptive regularisation for ensemble Kalman inversion

Marco Iglesias*  and Yuchen Yang

School of Mathematical Sciences, University of Nottingham, Nottingham, NG72RD, United Kingdom

E-mail: marco.iglesias@nottingham.ac.uk

Received 29 June 2020, revised 22 November 2020

Accepted for publication 10 December 2020

Published 22 January 2021



CrossMark

Abstract

We propose a new regularisation strategy for the classical ensemble Kalman inversion (EKI) framework. The strategy consists of: (i) an adaptive choice for the regularisation parameter in the update formula in EKI, and (ii) criteria for the early stopping of the scheme. In contrast to existing approaches, our parameter choice does not rely on additional tuning parameters which often have severe effects on the efficiency of EKI. We motivate our approach using the interpretation of EKI as a Gaussian approximation in the Bayesian tempering setting for inverse problems. We show that our parameter choice controls the symmetrised Kullback–Leibler divergence between consecutive tempering measures. We further motivate our choice using a heuristic statistical discrepancy principle. We test our framework using electrical impedance tomography with the complete electrode model. Parameterisations of the unknown conductivity are employed which enable us to characterise both smooth or a discontinuous (piecewise-constant) fields. We show numerically that the proposed regularisation of EKI can produce efficient, robust and accurate estimates, even for the discontinuous case which tends to require larger ensembles and more iterations to converge. We compare the proposed technique with a standard method of choice and demonstrate that the proposed method is a viable choice to address computational efficiency of EKI in practical/operational settings.

Keywords: ensemble Kalman inversion, parameter identification, derivative-free inversions

(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

Ensemble Kalman inversion (EKI) [1–5] is a collection of algorithms that import ideas from the ensemble Kalman filter (EnKF) developed in [6]. While EnKF was devised for assimilating data into models for numerical weather prediction and oceanography, EKI has been applied to solve parameter identification problems arising from multiple disciplines. The initial versions of EKI were proposed for the calibration of oil-reservoir models [7, 8], and then transferred in [1, 2] to more generic PDE-constrained inverse problems settings. The development of EKI as an iterative solver for parameter identification problems has led to numerous applications including the calibration of climate [9], turbulent flow [10], finance [11] and biomedical [12] models. EKI has been used for imaging and non-destructive testing including electrical impedance tomography (EIT) [13], seismic inversion [14], characterisation of thermophysical properties of walls [15] and composite materials [16]. More recently, EKI has also been applied for the solution of machine learning tasks [17].

There is clear promise for the potential use of EKI as a practical and operational tool for PDE-constrained calibration and imaging arising from multiple applications in science and engineering. However, most EKI algorithms still rely on the appropriate selection of user-defined parameters that control the stability, accuracy and computational efficiency of EKI. Unfortunately, the lack of a general theory for the convergence of EKI means that there is no principled approach to select these parameter in an optimal fashion. The seamless version of EKI proposed in [18] and further developed in [19–22] has provided enormous theoretical advances for understanding EKI within the more general frameworks of stochastic differential equations (SDEs). However, to the best of our knowledge, the selection of those crucial EKI parameters are still chosen empirically. Among those parameters, the choice of the regularisation (inflation) parameter in EKI, or alternatively, the step size from the discretisation of the SDE formulation of EKI often depends on additional tuning parameters which can only be informed after problem-specific numerical testing which maybe computationally intensive.

The aim of this work is to introduce a simple, yet computationally efficient, regularisation strategy for EKI that does not rely on further tuning parameters. Our main focus is large-scale/high resolution imaging/identification settings in which there are two fundamental challenges: (i) the forward model is computationally costly and (ii) the unknown must be parameterised in a highly-complex manner so that key features from the truth can be extracted via EKI. Both challenges are intertwined since the latter means that EKI requires large ensembles and more iterations to achieve accurate identifications; thus the total computational cost of EKI can become unfeasible. An efficient and robust regularisation strategy within EKI is a key requirement in these practical settings.

1.1. The inverse problem framework with ensemble Kalman inversion

We work on a generic setting where the properties that we wish to identify are functions which, in turn, play the role of inputs (e.g. coefficients) of partial differential equations (PDEs) describing the underlying experiment/process. The observable variables in these experiment/processes are functionals of the solution of these PDEs. Under this inverse problem setting, the forward problem can be written in terms of a nonlinear operator $\mathcal{F} : \mathcal{X} \rightarrow \mathbb{R}^M$ that maps physical properties from an admissible space \mathcal{X} to the space of observable quantities \mathbb{R}^M . Our work is focussed on using EKI for solving the *classical (deterministic) inverse problem* that consists of

estimating the underlying physical property of the medium, denoted by κ^\dagger , given noisy measurements of $\mathcal{F}(\kappa^\dagger)$, which we assume are given by

$$y = \mathcal{F}(\kappa^\dagger) + \eta, \quad (1)$$

where η is the unknown measurement error. We further assume η is drawn from a centred Gaussian measure with known covariance Γ . In order to address this inverse problem, we define a class of suitable parameterisations $\mathcal{P} : \mathcal{H} \rightarrow \mathcal{X}$ that enable us to characterise our estimate of the unknown as $\kappa = \mathcal{P}(u)$, i.e. in terms of an input (function) $u \in \mathcal{H}$ which we calibrate within the EKI framework. We pose the parameterised (in terms of u) inversion via

$$u^* = \arg \min_{u \in \mathcal{S}_0} \Phi(u; y), \quad (2)$$

where $\mathcal{S}_0 \subset \mathcal{H}$ is a user-defined subspace of admissible solutions and $\Phi(u; y)$ is the functional defined by

$$\Phi(u; y) \equiv \frac{1}{2} \|\Gamma^{-1/2}(y - \mathcal{G}(u))\|^2, \quad (3)$$

where $\mathcal{G} = \mathcal{F} \circ \mathcal{P}$ is the forward (parameter-to-output) map and $\|\cdot\|$ denotes the M -dimensional Euclidean norm. We use EKI to approximate (2) and to provide an estimate of κ^\dagger via $\kappa^* = \mathcal{P}(u^*)$.

For the PDE-constrained identification problems that we wish to address, \mathcal{G} is often a compact operator which leads to the ill-posedness of (2) in the sense of stability [23, 24]. Although numerous regularisation techniques [23, 25] can be used to address ill-posedness, most of them require the computation of the Fréchet derivative of the forward map, as well as the corresponding adjoint operator. This constitutes a substantial limitation in many practical applications where the map \mathcal{F} is only accessible via commercial software simulations with no adjoint built-in functionalities. Such a practical limitation has given rise to a large body of work on EKI techniques that stem from EnKF [6] and which, in turn, do not require derivatives of the forward map.

While there are several versions of EKI algorithms to approximate (2), here we focus on the classical, perturbed-observation EKI displayed in algorithm 1. This generic version of EKI involves selecting an initial ensemble of J particles $\{u_0^{(j)}\}_{j=1}^J \subset \mathcal{H}$. Then, each particle is iteratively updated according to the following formula

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{u\mathcal{G}}(C_n^{\mathcal{G}\mathcal{G}} + \alpha_n \Gamma)^{-1}(y + \sqrt{\alpha_n} \xi_n - \mathcal{G}(u_n^{(j)})), \quad (4)$$

where α_n is a tuning (regularisation) parameter, $\xi_n \sim N(0, \Gamma)$ is perturbation of the data, and $C_n^{u\mathcal{G}}$ and $C_n^{\mathcal{G}\mathcal{G}}$ are empirical covariances defined in (16) and (17). The running estimate of the unknown is obtained via the ensemble mean

$$\bar{u}_{n+1} \equiv \frac{1}{J} \sum_{j=1}^J u_{n+1}^{(j)}. \quad (5)$$

Using informal arguments, it can be shown (see appendix A and the work in [2]) that the ensemble mean \bar{u}_{n+1} can be seen as an ensemble approximation of

$$m_{n+1} = \arg \min_{u \in \mathcal{S}_0} \left\{ \frac{1}{2} \left\| \Gamma^{-1/2} (y - \mathcal{G}(m_n) - D\mathcal{G}(m_n)(u - m_n)) \right\|^2 + \frac{\alpha_n}{2} \left\| C_n^{-1/2} (u - m_n) \right\|_{\mathcal{H}}^2 \right\} \quad (6)$$

where $D\mathcal{G}$ denotes the Fréchet derivative of \mathcal{G} , C_n is a covariance operator that we define in appendix A, and $\mathcal{S}_0 \equiv \text{span}\{u_0^{(j)}\}_{j=1}^J$ is the subspace generated by the initial ensemble. If C_n is the identity operator, (6) is the standard Levenberg–Marquardt (LM) scheme [26] applied for the computation of (2). Note that (6) can also be interpreted as iterative Tikhonov regularisation applied to the linearisation of \mathcal{G} .

The link between EKI and (6) is very useful because (i) it motivates EKI as a derivative-free solver for (2) and (ii) it reveals the role of α_n in (4) as a Tikhonov regularisation parameter. According to the theory in [27], α_n must be carefully selected, together with the stopping criteria, in order to ensure the stability of the LM scheme. The approach for selecting α_n in the LM proposed in [27] has been adapted to the EKI framework in [2, 16], and subsequently used in [3, 12, 14, 28–30]. As we discuss in the next section, this approach relies on tuning parameters that, unless carefully chosen, can lead to unnecessary large number of iterations n^* . Since the main computational cost of algorithm 1 is $n^* \times J$, it is clear that a large n^* is detrimental to the computational efficiency of EKI.

1.2. EKI as a Gaussian approximation in the Bayesian tempering setting

Although the goal of most of the existing applications of EKI is to solve the deterministic problem in (2), the role of EKI within the Bayesian setting for parameter identification can be useful for identifying suitable choices of the regularisation parameter α_n . In the Bayesian setting we put a prior measure, $\mu_0(u) = \mathbb{P}(u)$, on the unknown u that we wish to infer. Given measurements, y , the Bayesian inverse problem consists of approximating the posterior $\mu(u) \equiv \mathbb{P}(u|y)$ which, from Bayes' rule [31] is given by

$$\mu(\mathrm{d}u) \propto \mu_0(\mathrm{d}u) \exp \left[-\frac{1}{2} \left\| \Gamma^{-1/2} (y - \mathcal{G}(u)) \right\|^2 \right], \quad (7)$$

where we have made the standard assumption that $y = \mathcal{G}(u) + \eta$ with $\eta \sim N(0, \Gamma)$. Modern computational approaches [28, 32–35] for high-dimensional Bayesian inverse problems use the tempering approach that consists of introducing N intermediate measures $\{\mu_n\}_{n=1}^N$ between the prior and the posterior. These measures are defined by

$$\mu_{t_n}(\mathrm{d}u) \propto \mu_{t_0}(\mathrm{d}u) \exp \left[-\frac{t_n}{2} \left\| \Gamma^{-1/2} (y - \mathcal{G}(u)) \right\|^2 \right], \quad (8)$$

where $\{t_n\}_{n=1}^N$ are tempering parameters that satisfy:

$$t_0 \equiv 0 < t_1 < t_2 < \dots < t_N < t_{N+1} \equiv 1. \quad (9)$$

Note that $n = 0$ and $n = N + 1$ in (8) yields the prior ($\mu_{t_0} = \mu_0$) and posterior ($\mu_{t_{N+1}} = \mu$), respectively. From expression (8) we obtain the following recursive formula for the intermediate measures:

$$\frac{\mu_{t_{n+1}}(\mathrm{d}u)}{\mu_{t_n}(\mathrm{d}u)} \propto \exp \left[-\frac{1}{2} \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}(u)) \right\|^2 \right], \quad (10)$$

where

$$\alpha_n^{-1} = t_{n+1} - t_n. \quad (11)$$

From (9) it follows that

$$\sum_{n=0}^N \alpha_n^{-1} = 1. \quad (12)$$

In (11) we employ the same notation that we use for the regularisation parameter in EKI (see equation (4)), because the ensemble computed at the n th iteration of EKI, is an ensemble approximation of a Gaussian measure that, in turn, approximates¹ the intermediate distribution μ_{t_n} in (8) (see [16] and appendix A.2). Therefore, in the Bayesian setting, controlling the regularisation parameter α_n in EKI means to gradually transition between prior and posterior in order to facilitate more accurate sampling of the intermediate measures.

The link between EKI and the Bayesian tempering setting has been recently explored in [16], where the selection of α_n is borrowed from the adaptive-tempering sequential Monte Carlo (SMC) method of [33]. In this approach, α_n is selected based on user-defined threshold to the effective sample size (ESS) which, in SMC methods, is used to determine the quality in the population of particles that approximate each intermediate measure μ_{t_n} . However, this selection of α_n involves the computation of the likelihood between consecutive measures (10). When the number of observations is large and/or a the measurement noise is small, this likelihood can take very small values unless large α_n 's are chosen. Consequently, many iterations may be needed to satisfy condition (12). Furthermore, the approach of [16] requires the aforementioned user-defined threshold on the ESS which may substantially affect the efficiency of EKI.

1.3. Our contribution

In this work we propose a novel adaptive choice of α_n in EKI that does not require any tuning parameters. This new selection of α_n depends on the number of observations, M , as well as the values of the least-squares functional (3) for the ensemble of particles, i.e.

$$\Phi_n \equiv \left\{ \Phi(u_n^{(j)}; y) \right\}_{j=1}^J = \left\{ \frac{1}{2} \left\| \Gamma^{-1/2} (y - \mathcal{G}(u_n^{(j)})) \right\|^2 \right\}_{j=1}^J. \quad (13)$$

More specifically, we select α_n^{-1} via

$$\alpha_n^{-1} = \min \left\{ \max \left\{ \frac{M}{2\langle \Phi \rangle_n}, \sqrt{\frac{M}{2\langle \Phi, \Phi \rangle_n}} \right\}, 1 - t_n \right\}, \quad (14)$$

where $\langle \Phi \rangle_n$ and $\langle \Phi, \Phi \rangle_n$ denote the empirical mean and variance of Φ_n , and

$$t_n = \begin{cases} \sum_{j=0}^{n-1} \alpha_j^{-1} & \text{if } n \geq 1, \\ 0 & \text{if } n = 0. \end{cases} \quad (15)$$

Note that in (14) we require that $\alpha_n^{-1} \leq 1 - t_n$ in order to enforce constraint (12) from the tempering/annealing setting using $N = n^*$ intermediate measures. More specifically, we propose to stop the EKI algorithm at the iteration, n^* , such that

$$\alpha_{n^*}^{-1} = 1 - t_{n^*}$$

¹ Even when J is large, these approximations are only exact in the linear-Gaussian case (i.e. \mathcal{G} linear and μ_0 Gaussian).

which, from (15), yields constraint (12). Noting that the selection of α_n^{-1} in (14) is determined by the data misfit of the particles, we refer to our regularisation strategy as the *data misfit controller* (DMC).

We motivate our DMC from the Bayesian perspective of EKI in which α_n^{-1} is the difference between consecutive tempering/annealing parameters (see equation (11)). We show that the selection of α_n^{-1} in (14) can be obtained by controlling, via imposing a threshold, the symmetrised Kullback–Leibler divergence (or Jeffreys’ divergence) between two consecutive tempering measures. We then use a statistical discrepancy principle to select this threshold.

Jeffreys’ divergence has been employed to design efficient selection of the step size between tempering measures when using MCMC to approximate the target posterior. For example, in the context of the *tempered transitions* method, the work of [36] uses Jeffreys’ divergence to select step sizes that increase acceptance rates. Jeffreys’ divergence has also been used for path sampling: a method commonly used to estimate normalisation constants (e.g. for Bayesian model selection). The work of [37], for example, shows that by controlling the Jeffreys’ divergence between tempering measures, the error in the path sampling estimator can be minimised. Both the works in [36, 37], as well as various references therein, employ continuous tempering in the finite-dimensional setting in order to characterise Jeffreys’ divergence. Here we extend some of those results to our infinite-dimensional framework for EKI.

We encode our regularisation strategy for EKI in an algorithm (**EKI-DMC**) that we test via EIT with the complete electrode model (CEM). We show that the choice of α_n via the DMC that we import from the Bayesian formulation, leads to a stable, robust and computationally efficient EKI algorithm for solving classical (deterministic) inverse problems posed via (2). We demonstrate the computational and practical advantages of **EKI-DMC** over existing approaches in which α_n is borrowed from the theory for the LM scheme. We conduct test with two parameterisations of the unknown conductivity that allows to consider both continuous and piece-wise constant conductivities. In particular, we investigate the performance **EKI-DMC** to infer regions characterised via a level-set function which is, in turn, parameterised with an anisotropic Whittle–Matern (WM) field.

In section 2 we further review the literature on existing approaches for the selection of α_n for the classical EKI framework. The proposed DMC (14) is motivated in section 3, where we use continuous tempering to derive approximations to Jeffreys’ divergence between tempering measures (subsection 3.2), as well as a statistical discrepancy principle to (subsection 3.3) for selecting the threshold on Jeffreys’ divergence. In section 4 we provide a numerical investigation of the performance of the proposed **EKI-DMC** algorithm for EIT. In section 5 we discuss final remarks and conclusions. For completeness in appendix A we motivate the classical EKI from the Bayesian tempering scheme, and from which we also draw links between the LM algorithm and EKI. Some technical proofs are included in appendix B.

2. Literature review

We discuss some existing regularisation strategies for EKI within the inversion setting posed in terms of the unregularised least-squares formulation of (2), and which leads to the classical EKI formulation in (4). We highlight that there is a new alternative EKI methodology and algorithms proposed in [4, 5] that arise from regularising (2) with a Tikhonov-like term. In addition, the review is focussed on PDE-constrained inversion; for a review of modern Kalman methods in the context of the data assimilation framework we refer the reader to the recent work on [38] and references therein.

Algorithm 1. Generic EKI (with perturbed observations).**Input:**

- 1) $\{u_0^{(j)}\}_{j=1}^J$: Initial ensemble of inputs.
- 2) Measurements y and covariance of measurement errors Γ .

Set $\{u_n^{(j)}\}_{j=1}^J = \{u_0^{(j)}\}_{j=1}^J$ and $\theta = 0$ **while** $\theta < 1$ **do**

- (1) Compute $\mathcal{G}_n^{(j)} = \mathcal{G}(u_n^{(j)})$, $j \in \{1, \dots, J\}$
- (2) Compute regularisation parameter α_n and check for convergence criteria
if converged **then**
set $\theta = 1$ and $n^* = n$
- (3) Update each ensemble member

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{u\mathcal{G}}(C_n^{\mathcal{G}\mathcal{G}} + \alpha_n\Gamma)^{-1}(y + \sqrt{\alpha_n}\xi_n - \mathcal{G}_n^{(j)}), \quad j \in \{1, \dots, J\}$$

where

$$C_n^{\mathcal{G}\mathcal{G}} \equiv \frac{1}{J-1} \sum_{j=1}^J (\mathcal{G}_n^{(j)} - \bar{\mathcal{G}}_n) \otimes (\mathcal{G}_n^{(j)} - \bar{\mathcal{G}}_n) \quad (16)$$

$$C_n^{u\mathcal{G}} \equiv \frac{1}{J-1} \sum_{j=1}^J (u_n^{(j)} - \bar{u}_n) \otimes (\mathcal{G}_n^{(j)} - \bar{\mathcal{G}}_n) \quad (17)$$

with $\bar{u}_n \equiv \frac{1}{J} \sum_{j=1}^J u_n^{(j)}$ and $\bar{\mathcal{G}}_n \equiv \frac{1}{J} \sum_{j=1}^J \mathcal{G}_n^{(j)}$. $n \leftarrow n + 1$ **end****output:** $\{u_{n^*}^{(j)}\}_{j=1}^J$ converged ensemble**2.1. EKI as an iterative solver for identification problems**

The initial versions of EKI [1] for generic identification problems proposed to use the classical EnKF [6] update formula² as an iterative solver for (2) by introducing an artificial dynamical system. For various PDE-constrained identification problems, the work of [1] numerically showed that this early version of EKI approximated well the solutions of (2) (with $\mathcal{S}_0 = \text{span}\{u_0^{(j)}\}_{j=1}^J$) within the first few iterations. However, they noted the algorithm became unstable if it was allowed to iterate after the data misfit (2) had reached the noise level δ defined by

$$\delta = \left\| \Gamma^{-1/2} (y - \mathcal{G}(u^\dagger)) \right\|, \quad (18)$$

where u^\dagger is the truth. This lack of stability led to the work of [2] where links between EKI and the regularising LM scheme of [27] were first established and used to develop a regularising version of EKI. For the LM scheme, the work of [27] ensures that, under certain assumptions of the forward map \mathcal{G} , the scheme in (6) converges to the solution of (2) (as $\delta \rightarrow 0$) provided that (i) the regularisation parameter α_n satisfies

²The classical EnKF consists of equation (4) with $\alpha_n = 1$ fixed throughout the iterations.

$$\rho \left\| \Gamma^{-1/2}(y - \mathcal{G}(m_n)) \right\| \leq \alpha_n \left\| \Gamma^{1/2}(D\mathcal{G}(m_n)\mathcal{C}_n D\mathcal{G}^*(m_n) + \alpha_n \Gamma)(y - \mathcal{G}(m_n)) \right\|, \quad (19)$$

where $\rho < 1$ is a tuning parameter, and (ii) that the algorithm is terminated at an iteration level n^* determined by the following *discrepancy principle*

$$\left\| \Gamma^{-1/2}(y - \mathcal{G}(m_{n^*})) \right\| \leq \tau \delta < \left\| \Gamma^{-1/2}(y - \mathcal{G}(m_n)) \right\| \quad 0 \leq n < n^*, \quad (20)$$

where τ is another tuning parameter that must satisfy $\tau > 1/\rho$. In [2], these regularisation strategies from LM were adapted for the selection of α_n in EKI via using derivative-free Gaussian approximations in (19) and (20). We refer to the approach from [2] as **EKI-LM** (see algorithm 2).

The numerical results of [2] showed that **EKI-LM** enabled stability and accuracy for sufficiently large ensembles. Further work that has explored **EKI-LM** can be found in [3, 28, 29] as well as some practical applications including seismic tomography [14], modelling of intracranial pressure [12] and time fractional diffusion inverse problems [30].

Despite of addressing stability in EKI, the approach **EKI-LM** suffers from a potential practical limitation that arises from the fact that it relies on the tuning parameters ρ and τ in (21). Larger ρ 's yield larger α_n and in turn more iterations to converge. Smaller ρ 's, while desirable for computational efficiency, result in larger τ 's which can lead to larger data misfit and possible loss of accuracy from stopping too early (via (22)). Selecting these parameters in a computationally optimal fashion becomes more crucial when EKI is combined with complex parameterisations of the unknown. As we discuss in subsection 4.2, there are cases for which, instead of using EKI directly on the physical property that we wish to infer, we need to parameterise the unknown to be able to capture properties that are not necessarily encoded in the initial ensemble. For example, in [2] the LM approach for EKI was applied with a level-set parameterisation of the unknown in order to characterise discontinuous properties (i.e. discontinuous conductivity in the context of EIT). In comparison to the simpler case in which EKI directly estimates a physical property of interest, [2] found that not only they needed a larger ensemble to achieve converge, but also more EKI iterations. The application of **EKI-LM** with level-set parameterisations for seismic imaging in [14] reported up to 40 iterations to achieve convergence. The numerical results of [3] also show that when EKI is combined with various others parameterisations of the unknown, the number of iterations of EKI can become large even for simple 1D and 2D forward modelling settings. We aim at addressing these very same issues with the selection of α_n that we propose in (14) and that we motivate in section 3 using the Bayesian perspective of EKI discussed earlier.

2.2. EKI as Gaussian approximation in linearised Bayesian tempering

Although the Bayesian perspective of EKI as a Gaussian approximation within the tempering setting was initially mentioned in [18] and further developed in [16], the early work of [7, 39] established a strong link between EKI and the Bayesian setting which led exactly to the same EKI scheme from algorithm 1. However, instead of using tempering, they used a heuristic approach in which the data was inverted multiple times with noise inflated by $\sqrt{\alpha_n}$. From algebraic manipulations they figured out that, in order to accurately sample from the posterior in the linear-Gaussian case, their inflation parameter α_n must satisfy (12) with $N = n^*$, where n^* is the total number of EKI iterations. Note that, in the tempering setting, (12) is simply a consequence of the definition of α_n in (11) as well as the definitions of t_0 and t_{N+1} in (9). Moreover, in the general Bayesian tempering settings, expression (12) holds for the general

Algorithm 2. (EKI-LM) EKI with LM selection of α_n and stopping criteria.

Input: Same from Algorithm 1, $\rho < 1$, and $\tau > \frac{1}{\rho}$ as well as the noise level δ .

Set $\{u^{(j)}\}_{j=1}^J = \{u_0^{(j)}\}_{j=1}^J$ and $\theta = 0$

while $\theta < 1$ **do**

(1) Compute $\mathcal{G}^{(j)} \equiv \mathcal{G}(u^{(j)})$, $j \in \{1, \dots, J\}$

(2) Compute α_n such that

$$\rho \|\Gamma^{-1/2}(y - \bar{\mathcal{G}}_n)\| \leq \alpha_n \|\Gamma^{1/2}(\mathcal{C}_n^{\mathcal{G}} + \alpha_n \Gamma)^{-1}(y - \bar{\mathcal{G}}_n)\| \quad (21)$$

if

$$\|\Gamma^{-1/2}(y - \bar{\mathcal{G}}_n)\| \leq \tau \delta. \quad (22)$$

then

set $n^* = n$.

break

(3) Update each ensemble member using (4) (see also Step 3 from Algorithm 1)

end

output: $\{u_{n^*}^{(j)}\}_{j=1}^J$ converged ensemble

nonlinear case and without any assumptions on the prior. Nonetheless, those considerations for the linear-Gaussian case from [39] enabled them to (i) justify the use of (12) in the nonlinear case and (ii) to propose a simple selection of α_n given by $\alpha_n = n^*$ with n^* selected *a priori* (which trivially satisfies (12)). This approach, referred to as ensemble smoother with multiple data assimilation, has been popularised in petroleum engineering applications (see [40] and references therein). However, from the insight gained from the link between EKI and the LM scheme, this selection of α_n was discouraged in [41] since the stability of EKI, as shown in [2, 41], requires α_n to be large at the beginning of the iterations, and gradually decrease as the data misfit approaches the noise level. Further versions of ES-MDA [42] adopted selections of α_n similar to those initially proposed in [2] based on the LM scheme.

As discussed in the previous section, a selection of α_n based on the adaptive-tempering SMC method of [33] is proposed in [16]. More specifically, α_n is selected so that

$$\left[\sum_{j=1}^J (\mathcal{W}_n^{(j)}[\alpha_n])^2 \right]^{-1} = J_*, \quad (23)$$

where J_* is a tuning user-defined parameter and

$$\mathcal{W}_n^{(j)}[\alpha_n] = \frac{\exp \left[-\frac{1}{2} \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}(u_n^{(j)})) \right\|^2 \right]}{\sum_{s=1}^J \exp \left[-\frac{1}{2} \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}(u_n^{(s)})) \right\|^2 \right]}. \quad (24)$$

The left-hand side of (23) is the ESS of the ensemble approximation of μ_{i_n} . For further details we refer the reader to [16], where the selection of α_n according to (23) was implemented in a batch-sequential EKI framework to sequentially solve an inverse problem that arises in resin transfer moulding. The same EKI methodology was applied in [15] for parameter identification of the heat equation, including identification of thermal conductivity and heat capacitance given boundary measurement of heat flux. Both time-dependent applications tackled in [15, 16]

involved the inversion of small number of measurements (e.g. < 30) at each observation time, and with relatively large noise informed by measurement protocols. However, as discussed in the previous section, the selection of α_n via solving (23) becomes problematic for large number of observations and/or small observational noise, since (24) requires the computation of the likelihood (10). Another limitation is that the efficiency of the approach used in [15, 16] relies on the tuning parameter J_* in (23).

2.2.1. EKI as a discretisation of stochastic differential equations. The pioneering work of [18] has shown that EKI can be derived as a discretisation of an SDE system for the ensemble of particles in EKI. This so-called continuous-time limit or seamless formulation of EKI has led to further theoretical understanding of the EKI framework [19, 20, 22]. In addition, using alternative discretisation schemes of the seamless formulation of EKI, leads to new EKI algorithms in both the context of optimisation [17, 19] and sampling within the Bayesian approach [21].

In the seamless formulation of EKI, the regularisation parameter α_n from the classical EKI becomes the inverse of the mesh-size/discretisation step. The work of [17] proposes to choose this parameter according to

$$\alpha_n^{-1} = \frac{\alpha_0^{-1}}{\|U\|_F + \epsilon}, \quad (25)$$

where α_0 and ϵ are user defined parameters, $\|\cdot\|_F$ denotes the Frobenius norm and U is a matrix with entries $U_{j,k} = (\mathcal{G}(u_n^{(k)}) - \bar{\mathcal{G}}_n)^T \Gamma^{-1} (\mathcal{G}(u_n^{(j)}) - y)$. To the best of our knowledge, this selection of α_n has only been used for new different EKI algorithms including the ones arising from forward-Euler [17] and the implicit split-step method [21]. While (25) is also a perfectly valid choice of α_n for classical EKI, we emphasise that it relies on the choice of the tuning parameters α_0 and ϵ .

3. The proposed regularisation framework for EKI

The aim of this section is to motivate the new approach (DMC) that we propose via (14) to select α_n^{-1} for the classical EKI setting given in (4). We motivate our approach using tempering within the Bayesian setting for inverse problems in which α_n^{-1} is the step size between consecutive tempering measures. After introducing some notation and assumptions in subsection 3.1, in subsection 3.2 we define tempering in the continuous setting which we require in order to compute an approximation of Jeffreys' divergence between two consecutive measures. Under this approximation, we further show that our selection of α_n^{-1} (14) controls Jeffreys' divergence. The control parameter is then selected following a heuristic approach based on a discrepancy principle that we introduced in subsection 3.3.

3.1. Notation, definitions and assumptions

We assume that the space of inputs, \mathcal{H} , is a real-valued separable Hilbert space. The Borel-sigma algebra over \mathcal{H} is denoted by $\mathcal{B}(\mathcal{H})$. In this section we assume that the prior is a Gaussian measure $\mu_0 = \mathcal{N}(m, \mathcal{C}) : \mathcal{B}(\mathcal{H}) \rightarrow [0, 1]$ on $(\mathcal{H}, \mathcal{B}(\mathcal{H}))$, satisfying $\mu_0(\mathcal{H}) = 1$. Given two probability measures $\hat{\mu}$ and μ on a measurable space (\mathcal{Y}, Σ) , if $\hat{\mu}$ is absolutely continuous with respect to μ , we define the Kullback–Leibler divergence of $\hat{\mu}$ with respect to μ by [43],

$$D_{\text{KL}}(\hat{\mu} \parallel \mu) := \int_{\mathcal{Y}} \log \left(\frac{d\hat{\mu}}{d\mu} \right) d\hat{\mu},$$

where $\frac{d\hat{\mu}}{d\mu}$ denotes the Radon–Nikodym derivative of $\hat{\mu}$ with respect to μ . For two equivalent probability measures $\hat{\mu}$ and μ (i.e. $\hat{\mu} \ll \mu$ and $\mu \ll \hat{\mu}$) we define the symmetrized KL divergence or Jeffrey’s divergence:

$$D_{\text{KL},2}(\mu, \hat{\mu}) := D_{\text{KL}}(\hat{\mu} \parallel \mu) + D_{\text{KL}}(\mu \parallel \hat{\mu}). \quad (26)$$

We consider the following assumption on the forward operator:

Assumption 1. Assume that the forward map $\mathcal{G} : \mathcal{H} \rightarrow \mathbb{R}^M$ satisfies the following two conditions:

- (a) For every $\epsilon > 0$ there is an $\beta = \beta(\epsilon) \in \mathbb{R}$ such that, for all $u \in \mathcal{H}$,

$$\left\| \Gamma^{-1/2} \mathcal{G}(u) \right\| \leq \exp(\epsilon \|u\|_{\mathcal{H}}^2) + \beta$$

- (b) For every $r > 0$ there is an $K = K(r) > 0$ such that,

$$\left\| \Gamma^{-1/2} (\mathcal{G}(u_1) - \mathcal{G}(u_2)) \right\| \leq K \|u_1 - u_2\|_{\mathcal{H}}$$

for all $u_1, u_2 \in \mathcal{H}$ with $\max \{ \|u_1\|_{\mathcal{H}}, \|u_2\|_{\mathcal{H}} \} < r$.

This assumption was used in [[31], theorem 4.1] to establish well-posedness of the posterior in the infinite-dimensional framework. Let us recall the least-square functional $\Phi(u; y)$ defined in (3). For any $y \in \mathbb{R}^M$, Fernique’s theorem and condition (a) in assumption 1 implies that $\Phi(\cdot; y)^m$ is square integrable, for any $m \geq 1$, with respect to the prior $\mu_0 = \mathcal{N}(m, \mathcal{C})$. Integrability of $\Phi(\cdot; y)^m$ is used through the rest of this section and for the proofs in appendix B.

3.2. Controlling Jeffreys’ divergence between successive tempering measures

In this section we show that the data-misfit controller (14) ensures that the level of information between any two consecutive tempering measures is approximately bounded by a user-defined threshold $\theta > 0$. We measure the level of information between these measures via Jeffreys’ divergence defined in (26). Our aim is to show that the selection of $\alpha_n^{-1} = t_{n+1} - t_n$ via (14) yields

$$D_{\text{KL},2}(\mu_{t_{n+1}}, \mu_{t_n}) \leq \theta, \quad n \in \{1, \dots, N\} \quad (27)$$

In subsection 3.3 we give heuristic motivations for the choice of θ based on a statistical discrepancy principle.

Our first step is the explicit characterisation of $D_{\text{KL},2}(\mu_{t_{n+1}}, \mu_{t_n})$ which we latter approximate in order to propose our adaptive selection of α_n^{-1} . This approximation, require us to work on the continuous tempering setting which we now introduce.

3.2.1. Continuous tempering. For $t \in [0, 1]$, we define measure μ_t such that the Radon–Nikodym derivative of μ_t with respect to the prior $\mu_0 = \mathcal{N}(m, \mathcal{C})$ satisfies, for almost every $u \in \mathcal{H}$,

$$\frac{d\mu_t}{d\mu_0}(u) = N_t^{-1} \exp(-t\Phi(u; y)), \quad (28)$$

where N_t is the normalising constant defined by

$$N_t = \int_{\mathcal{H}} \exp(-t\Phi(u; y)) \mu_0(du). \quad (29)$$

The following lemma establishes the well-posedness of $\{\mu_t | t \in [0, 1]\}$. The proof can be found in appendix B.1.

Lemma 1. *The probability measure μ_t from (28) exists on the probability space $(\mathcal{H}, \mathcal{B}(\mathcal{H}), \mu_0)$ and is equivalent to μ_0 , i.e. $\mu_t \ll \mu_0$ and $\mu_0 \ll \mu_t$.*

Using integrability of the data misfit with respect to the prior, and the equivalence proven in the previous lemma, let us define the mean and variance of the data misfit:

$$\langle \Phi \rangle_t := \int_{\mathcal{H}} \Phi(w; y) \mu_t(dw) \quad \langle \Phi, \Phi \rangle_t := \int_{\mathcal{H}} (\Phi(w; y) - \langle \Phi \rangle_t)^2 \mu_t(dw) \quad (30)$$

for all $t \in [0, 1]$.

It is not difficult to see that, given $0 = t_0 < t_1 < \dots < t_{N+1} = 1$, lemma 1 also establishes the equivalence of the probability measures $\{\mu_{t_n} : n = 0, 1, \dots, N + 1\}$. Furthermore, from the integrability of the data misfit, it follows that the Jeffreys' divergence between consecutive measures is well-defined and, using its definition in (26), it can be written as

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) = \alpha_n^{-1} (\langle \Phi \rangle_{t_n} - \langle \Phi \rangle_{t_{n+1}}). \quad (31)$$

The finite-dimensional version of (31) was proven in [[37], proposition 3.2] in the context of path sampling.

Let us note that, since $\Phi \geq 0$,

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) \leq \alpha_n^{-1} \langle \Phi \rangle_{t_n}. \quad (32)$$

3.2.2. Approximating $D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}})$. In practice, the expression for $D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}})$ provided in (31) cannot be used in (27) to find α_n^{-1} since, at time t_n , measure $\mu_{t_{n+1}}$ is unknown. Moreover, since $t_{n+1} = t_n + \alpha_n^{-1}$, measure $\mu_{t_{n+1}}$ depends on the choice of α_n . In order to propose a practical choice of α_n^{-1} via (31), let us now discuss some approximations. Note that if $\langle \Phi \rangle_{t_{n+1}} \ll \langle \Phi \rangle_{t_n}$, then the term $\langle \Phi \rangle_{t_n} \geq 0$ in the right-hand side of (31) can be neglected and hence

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) \approx \alpha_n^{-1} \langle \Phi \rangle_{t_n} \leq \theta \Rightarrow \alpha_n^{-1} \leq \frac{\theta}{\langle \Phi \rangle_{t_n}}. \quad (33)$$

However, approximating $D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}})$ with its upper bound in (32) may not be accurate when $\langle \Phi \rangle_{t_{n+1}} \approx \langle \Phi \rangle_{t_n}$ which is likely to happen when α_n^{-1} is small. We address this case via a first order approximation of $\langle \Phi \rangle_t$ (as a function of t), around $t = t_n$. We first need the following proposition that we prove in appendix B.2 and which extends, to the present infinite-dimensional setting, the results from [[36], section 2] for tuning tempered transitions.

Proposition 1. *For any bounded $t \geq 0$, let $u_t \sim \mu_t$ be an \mathcal{H} -valued random variable, where μ_t is the probability measure determined via formula (28). For any square integrable functional $g : \mathcal{H} \rightarrow \mathbb{R}$, the expected value $\mathbb{E}\{g(u_t)\}$, as a function of t , is differentiable and its derivative $\frac{d}{dt} \mathbb{E}\{g(u_t)\}$ satisfies:*

$$\frac{d}{dt} \mathbb{E}\{g(u_t)\} = -\text{Cov}\{g(u_t), \Phi(u_t; y)\}, \quad (34)$$

where Cov denotes covariance between two random variables.

The following corollary is a simple consequence of theorem 1 with $g(\cdot) = \Phi(\cdot, y)$.

Corollary 1. For any $t \in [0, 1]$, let μ_t denote the probability measure determined via the tempering setting (28). The mean $\langle \Phi \rangle_t$ defined in (30), as a function of $t \in [0, 1]$, is differentiable. Moreover, its derivative is given by

$$\frac{d}{dt} \langle \Phi \rangle_t = -\langle \Phi, \Phi \rangle_t.$$

We now come back to the case in which $\alpha_n^{-1} = t_{n+1} - t_n$ needs to be small so that (27) can be satisfied for a given threshold $\theta > 0$. Using corollary 1 we have

$$\langle \Phi \rangle_{t_{n+1}} \approx \langle \Phi \rangle_{t_n} + \alpha_n^{-1} \frac{d}{dt} \langle \Phi \rangle_{t_n} = \langle \Phi \rangle_{t_n} - \alpha_n^{-1} \langle \Phi, \Phi \rangle_{t_n}.$$

Thus,

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) \approx \alpha_n^{-2} \langle \Phi, \Phi \rangle_{t_n} \leq \theta \Rightarrow \alpha_n^{-1} \leq \sqrt{\frac{\theta}{\langle \Phi, \Phi \rangle_{t_n}}}. \quad (35)$$

We postulate that the approximation used in (35) is reasonable provided that it honours inequality (32), i.e.

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) \approx \alpha_n^{-2} \langle \Phi, \Phi \rangle_{t_n} \leq \alpha_n^{-1} \langle \Phi \rangle_{t_n}.$$

Whenever this condition is not satisfied, we approximate $D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}})$ via its upper bound as in (33). In summary, we propose the following approximation

$$D_{\text{KL},2}(\mu_{t_n}, \mu_{t_{n+1}}) \approx \min \{ \alpha_n^{-1} \langle \Phi \rangle_{t_n}, \alpha_n^{-2} \langle \Phi, \Phi \rangle_{t_n} \}. \quad (36)$$

We define α_n^{-1} :

$$\alpha_n^{-1} = \min \left\{ \max \left\{ \frac{\theta}{\langle \Phi \rangle_{t_n}}, \sqrt{\frac{\theta}{\langle \Phi, \Phi \rangle_{t_n}}} \right\}, 1 - t_n \right\}, \quad (37)$$

which satisfies (27) using the approximation (36), while we ensure that $\alpha_n^{-1} \leq 1 - t_n$ in order to satisfy (12) with $n^* = N$; hence n^* defines the natural stopping iteration level.

Expression (37) has the same form as the DMC in (14) with (i) $\langle \Phi \rangle_{t_n}$ and $\langle \Phi, \Phi \rangle_{t_n}$ approximated with empirical mean and variance from the ensemble of particles (recall $u_n^{(j)}$ are approximate samples of μ_{t_n}), and (ii) the choice $\theta = M/2$ which we motivate from heuristic arguments in the next subsection.

3.3. A statistical discrepancy principle

According to Morozov's discrepancy principle [44], if $u \in \mathcal{H}$ satisfies

$$\left\| \Gamma^{-1/2}(y - \mathcal{G}(u)) \right\| \leq \delta, \quad (38)$$

where δ is the noise level defined in (18), then u is an acceptable estimate of the classical (deterministic) inverse problem (1). In the context of iterative regularisation [25], a form of the discrepancy principle is often used as an early stopping rule (e.g. see expression (20)) to avoid

instabilities. In the Bayesian setting, however, the noise level is a random variable. In fact, under the assumption of additive centred Gaussian noise (1), the squared noise level δ^2 (see (18)) is a realisation of a chi-square random variable with M degrees of freedom. Therefore, its mean and variance are:

$$\mathbb{E} \{ \delta^2 \} = M, \quad \text{Var} \{ \delta^2 \} = 2M.$$

Inspired by the classical discrepancy principle, we then postulate that an acceptable estimate, u , must satisfy either one of the following criteria

C1 (accuracy)

$$\mathbb{E} \left\{ \left\| \Gamma^{-1/2} (y - \mathcal{G}(u)) \right\|^2 \right\} \leq \mathbb{E} \{ \delta^2 \} \equiv M. \quad (39)$$

C2 (uncertainty)

$$\text{Var} \left\{ \left\| \Gamma^{-1/2} (y - \mathcal{G}(u)) \right\|^2 \right\} \leq \text{Var} \{ \delta^2 \} \equiv 2M. \quad (40)$$

We invoke this statistical discrepancy principle to determine the step size α_n^{-1} of the n th Bayesian sub-problem (41) that arises from transition between tempering measures μ_{t_n} and $\mu_{t_{n+1}}$. More specifically, let us consider the following remark.

Remark 1. We can view (10) as an iterative application of Bayes rule, where at the n th iteration level we have a prior $\mu_{t_n}(u)$ and a likelihood defined by the observational model

$$y = \mathcal{G}(u) + \sqrt{\alpha_n} \eta, \quad \eta \sim N(0, \Gamma). \quad (41)$$

In other words, (10) defines a sequence of Bayesian inverse problems similar to the original one³ but with a Gaussian error $\sqrt{\alpha_n} \eta$ that has a covariance matrix Γ inflated by α_n . We can then think of $\mu_{t_{n+1}}$ given by (10) as the distribution of $u|y$ under the observational model (41) and prior μ_{t_n} .

We apply (39) and (40) to the sub-problem in remark 1. To this end, we use $\alpha_n \Gamma$ and $u_n \sim \mu_{t_n}$ instead of Γ and u in formulas (39) and (40) to obtain

$$2\alpha_n^{-1} \langle \Phi \rangle_{t_n} = \mathbb{E} \left\{ \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}(u_n)) \right\|^2 \right\} \leq M \quad (42)$$

$$4\alpha_n^{-2} \langle \Phi, \Phi \rangle_{t_n} = \text{Var} \left\{ \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}(u_n)) \right\|^2 \right\} \leq 2M, \quad (43)$$

where we have used the definition of Φ in (3) as well as (30). By choosing α_n^{-1} via (37) with $\theta = M/2$, we enforce that one of the above criteria (42) or (43) are satisfied, while ensuring that $t_{n+1} = t_n + \alpha_n^{-1} \leq 1$ as required. In other words, the proposed statistical discrepancy principle applied to each intermediate Bayesian sub-problem arising from tempering, yields the DMC (14) which, from the previous subsection, controls Jeffreys' divergence with threshold $\theta = M/2$. We incorporate the data-misfit controller for the selection of α_n^{-1} in EKI; we summarise the scheme in algorithm 3 (EKI-DMC).

³ Recall the observational model for original problem is $y = \mathcal{G}(u) + \eta$ with $\eta \sim N(0, \Gamma)$.

Algorithm 3. EKI with DMC (EKI-DMC).**Input:** Same from Algorithm 1Set $n = 0$ and $t_0 = 0$ **while** $t_n < 1$ **do**(1) Compute $\mathcal{G}_n^{(j)} = \mathcal{G}(u_n^{(j)})$, $j \in \{1, \dots, J\}$ (2) Compute α_n via (14)

(3) Update each ensemble member using (4) (see also Step 3 from Algorithm 1)

 $n \leftarrow n + 1$ $t_n \leftarrow t_n + \alpha_n^{-1}$ **end****output:** $\{u_n^{(j)}\}_{j=1}^J$ converged ensemble

Remark 2. For the applications that we discuss in section 4, we have find that the mean of the data misfit $\langle \Phi \rangle_{t_n}$ is considerably smaller than the standard deviation $\langle \Phi, \Phi \rangle_{t_n}$. Therefore,

$$\frac{\langle \Phi \rangle_{t_n}^2}{\langle \Phi, \Phi \rangle_{t_n}} \leq 1 < \frac{M}{2}, \Rightarrow \frac{M}{2\langle \Phi, \Phi \rangle_{t_n}} > \sqrt{\frac{M}{2\langle \Phi, \Phi \rangle_{t_n}}}. \quad (44)$$

Hence, in this case, α_n^{-1} is given by condition (42) rather than condition (43). Intuitively, (44) is likely to occur in problems where the prior is sufficiently wide so that it comprises the truth. However, for a prior that is centred too far from the truth and with small variance, we would then expect for $\langle \Phi \rangle_{t_n}$ to be larger than $\langle \Phi, \Phi \rangle_{t_n}$. In this case, DMC selects α_n^{-1} according to (43) in order to allow for larger steps in the tempering settings.

Remark 3. It is not difficult to see that, except for the case in which $\langle \Phi, \Phi \rangle_{t_n} \theta = \langle \Phi \rangle_{t_n}^2$, the selection of α_n^{-1} via (37) satisfies only one of the criteria (42) and (43) while violating the other one. As pointed by one of the anonymous referees, if we instead select α_n^{-1} via

$$\alpha_n^{-1} = \min \left\{ \frac{\theta}{\langle \Phi \rangle_{t_n}}, \sqrt{\frac{\theta}{\langle \Phi, \Phi \rangle_{t_n}}}, 1 - t_n \right\}, \quad (45)$$

then both criteria (42) and (43) will be satisfied. While (45) is indeed a valid choice of α_n^{-1} , this is a suboptimal choice for the numerical experiments that we discuss in section 4. From the observations we made in remark 2, using (45) will lead to smaller α_n^{-1} 's than those obtained by (37) and, thus, to many more iterations which makes the EKI algorithm more costly.

4. Numerical testing

In this section we test the performance of **EKI-DMC** in the context of EIT with the CEM that we introduce in subsection 4.1. We use two parameterisations of the unknown that we introduce in subsection 4.2. Implementation aspects and measures of performance are discussed in subsections 4.3 and 4.4. In subsection 4.5 and 4.6 we discuss numerical results which includes comparing the performance of **EKI-DMC** and **EKI-LM**.

4.1. Complete electrode model

Given electric currents $\{I_k\}_{k=1}^{m_e}$ injected through a set of surface electrodes $\{e_k\}_{k=1}^{m_e}$ placed on the boundary, ∂D , of D , the CEM consist of finding $[v, \{V_k\}_{k=1}^{m_e}]$ where v is the voltage in D and $\{V_k\}_{k=1}^{m_e}$ are the voltages on the electrodes. The dependent variables $[v, \{V_k\}_{k=1}^{m_e}]$ are given by the solution to [45]:

$$\nabla \cdot \kappa \nabla v = 0 \quad \text{in } D, \quad (46)$$

$$v + z_k \kappa \nabla v \cdot \mathbf{n} = V_k \quad \text{on } e_k, k = 1, \dots, m_e, \quad (47)$$

$$\nabla v \cdot \mathbf{n} = 0 \quad \text{on } \partial D \setminus \cup_{k=1}^{m_e} e_k, \quad (48)$$

$$\int_{e_k} \kappa \nabla v \cdot \mathbf{n} \, ds = I_k \quad k = 1, \dots, m_e, \quad (49)$$

where κ is the electric conductivity of D and $\{z_k\}_{k=1}^{m_e}$ are the electrodes' contact impedances. We consider an experimental setting consisting of n_p current patterns $I_1 = \{I_{1,k}\}_{k=1}^{m_e}, \dots, I_{n_p} = \{I_{n_p,k}\}_{k=1}^{m_e}$. For each of these patters $\{I_{j,k}\}_{k=1}^{m_e}$, we denote by $\{V_{j,k}\}_{k=1}^{m_e}$ the prediction of voltages at the electrodes defined by the CEM (46)–(49). For simplicity we assume that the contact impedances of the electrodes are known. We define the map

$$\mathcal{F}(\kappa) = V \equiv [\{V_{1,k}\}_{k=1}^{m_e}, \dots, \{V_{n_p,k}\}_{k=1}^{m_e}],$$

that for every conductivity, produces voltage measurements. EIT consist of finding the conductivity κ^\dagger of D given measurements of $V^\dagger = \mathcal{F}(\kappa^\dagger)$. For a review of the EIT problem we refer the reader to [46].

4.2. Parameterisations of EKI

In this subsection we introduce two maps, \mathcal{P}_1 and \mathcal{P}_2 , that we use to parameterise the unknown physical property κ^\dagger that we wish to estimate using the EKI framework introduced earlier.

4.2.1. Parameterisation \mathcal{P}_1 . Smooth properties. Let us first discuss the case in which prior knowledge suggest that the unknown physical property κ^\dagger is a smooth function. For simplicity we only consider the 2D case that we use in the numerical experiments of subsection 4.5 and 4.6 but we emphasise that the approach can be used for 1D and 3D settings. Let us introduce a *WM parameterisation* of the unknown $\kappa(x)$ that we wish to infer via EKI. The WM parameterisation involves a positive smoothness parameter denoted by ν , an amplitude scale σ , and two intrinsic length scales, L_1 and L_2 , along the horizontal and vertical direction, respectively. Given $\Theta = (\nu, \sigma, L_1, L_2) \in \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R}_+^2$, we define an operator \mathcal{W}_Θ that maps every $\omega \in H^{-1-\epsilon}(D)$ (with $\epsilon > 0$ arbitrary) to $\Psi = \mathcal{W}_\Theta \omega$, that satisfies the following fractional PDE in the domain D

$$(\mathbb{I} - \nabla \cdot \text{diag}(L_1^2, L_2^2) \nabla)^{(\nu+1)/2} \Psi = 4\sigma^2 \pi \frac{\Gamma(\nu+1)}{\Gamma(\nu)} \sqrt{L_1 L_2} \omega, \quad (50)$$

with Robin boundary conditions (BCs) on ∂D :

$$\Psi - \zeta_R \text{diag}(L_1^2, L_2^2) \nabla \Psi \cdot \mathbf{n} = 0. \quad (51)$$

where ζ_R is a tuning parameter. In (50) Γ denotes the gamma function, \mathbb{I} is the identity operator, and

$$\text{diag}(L_1^2, L_2^2) \equiv \begin{pmatrix} L_1^2 & 0 \\ 0 & L_2^2 \end{pmatrix}.$$

The proposed WM parameterisation of κ is defined by

$$\kappa = \mathcal{P}_1(\lambda, \Theta, \omega) \equiv \lambda \exp(\mathcal{W}_\Theta \omega), \quad (52)$$

where λ a positive scaling factor for κ . Note that (52) enforces that the electrical conductivity is positive. We can succinctly write (52) as,

$$\kappa = \mathcal{P}_1(u), \quad \text{where } u = (\lambda, \Theta, \omega) = (\lambda, \nu, \sigma, L_1, L_2, \omega). \quad (53)$$

The motivation behind this parameterisation comes from the theory of Gaussian random fields (GRFs). In particular, the work of [47] that shows that if $\omega \in H^{-1+\epsilon}$ is Gaussian white noise (i.e. $\omega \sim N(0, \mathbb{I})$) then

$$\log \kappa = \log(\lambda) + \mathcal{W}_\Theta \omega \sim N(\log \lambda, C_\Theta),$$

where C_Θ is the covariance operator induced by the Matern autocorrelation function defined by

$$\text{ACF}_\Theta(x) = \sigma^2 \frac{1}{2^{\nu-1} \Gamma(\nu)} \|x\|_{L_1, L_2}^\nu K_\nu \left(\|x\|_{L_1, L_2} \right),$$

where K_ν is the modified Bessel function of the second kind of order ν , and

$$\|x\|_{L_1, L_2} \equiv \sqrt{\frac{x_1^2}{L_1^2} + \frac{x_2^2}{L_2^2}}$$

It can also be shown that if $\log \kappa \sim N(\log \lambda, C_\Theta)$, then almost surely $\log \kappa \in H^{\nu-\epsilon}(D)$ [48] which further shows the role of the smoothness parameter ν .

Our choice for the BCs in (51) follows from the work of [49] that shows that Robin BCs are better suited, compared to Neumann and Dirichlet, to alleviate (via the appropriate choice of ζ_R) undesirable boundary effects which arise from the discretisation of GRFs. Let us reiterate that our goal here is to introduce parameterisations for EKI and then suitable initial ensembles on the parameters. Whether the initial ensemble yields Gaussian (or log-Gaussian) properties is not our main concern. However, it would not be advisable to select parameterisations that, for example, restrict the values of the physical property near the boundary which is something that we would expect if $\zeta_R = 0$ in (51).

The WM parameterisation in (52) allows us to incorporate the smoothness and lengthscales of the underlying field as part of the unknown that we can estimate with EKI. In contrast to the work of [3] where the lengthscales are estimated under isotropic assumptions, here we consider the anisotropic case. We focus on vertical/horizontal anisotropy but a rotation matrix can be further introduced in (50) to characterise properties with an arbitrary (and unknown) direction of anisotropy [49]

A straightforward approach to generate the initial ensemble for EKI with the parameterisation from (52) is to specify (hyper prior) densities, π_λ , π_{L_1} , π_{L_2} , π_σ and π_ν , and produce samples:

$$u_0^{(j)} \equiv (\lambda^{(j)}, \nu^{(j)}, \sigma^{(j)}, L_1^{(j)}, L_2^{(j)}, \omega^{(j)}), \sim \pi_\lambda \otimes \pi_\nu \otimes \pi_\sigma \otimes \pi_{L_1} \otimes \pi_{L_2} \otimes N(0, \mathbb{I}) \quad (54)$$

Remark 4 (KL expansion). For some geometries of D , the eigenfunctions and eigenvalues of a Matern covariance C_Θ are available in closed form [48]. In this case, an equivalent formulation of the WM parameterisations can be defined in terms of the spectral decomposition of C_Θ (i.e. KL expansion under the prior). Of course, for computational purposes WM parameterisations can be defined on a larger (simple) domain that encloses D and simply restrict the physical property to the domain of interest. While the spectral/KL approach is more standard, the approach we use here based on the work of [47] allows to naturally define the parameterisation in (53) via (50).

Remark 5 (smoothness parameter and amplitude scales). For simplicity, in this work the smoothness parameter ν and the amplitude scales σ in the WM are fixed, i.e. we leave these parameters out of the inversion. For the experiments that we discuss in the subsequent sections, the spatial variability in the unknown can be captured quite effectively only via $\omega(x)$. However, we recognise that including a variable σ can be beneficial in some cases as reported in [49]. Similarly, in the context of the experiments reported later, where the aim is mainly to recover medium anomalies, we find that sensible (fixed) choices of ν are sufficient, provided that the lengthscales are properly estimated via EKI. Nevertheless, as shown in [3], EKI can be used to estimate the smoothness parameter ν .

4.2.2. Parameterisation \mathcal{P}_2 . Piecewise-constant functions. In order to characterise piecewise-smooth functions we use the level-set approach initially proposed in [50] for deterministic inverse problems and, more recently, for fully Bayesian [48, 51] and EKI [2] settings. Our main modelling assumption is that the unknown property has a background value potentially heterogeneous, and that possible anomalies/defects consist of regions with (also possibly heterogeneous) higher/lower values than those in the background field. More specifically, let us first define the parameterisation

$$\kappa(x) = \mathcal{H}(\{\kappa_\iota\}_{\iota \in \{l,b,h\}}, f) \equiv \begin{cases} \kappa_l & f(x) \leq \zeta_1 \\ \kappa_b & \zeta_1 < f(x) \leq \zeta_2 \\ \kappa_h & f(x) > \zeta_2, \end{cases} \quad (55)$$

where κ_b , κ_l and κ_h denote the background, low-value and high-value fields that characterise the unknown physical property. For simplicity we assume these variables take only constant values. The level-set function denoted by $f(x)$ determines the background $\Omega_b \equiv \{x : \zeta_1 < f(x) \leq \zeta_2\}$ as well as the region of low $\Omega_l \equiv \{x : f(x) \leq \zeta_1\}$ and high $\Omega_h \equiv \{x \in f(x) > \zeta_2\}$ values. We assume that f is a smooth field with some variability that we enforce via a second level of parameterisation in terms of \mathcal{P}_1 introduced earlier. More specifically, we consider

$$f = \log \lambda_f + \mathcal{W}_{\Theta_f} \omega_f = \log(\mathcal{P}_1(u_f)), \quad (56)$$

where \mathcal{P}_1 is defined according to (52) and

$$u_f = (\lambda_f, \nu_f, \sigma_f, L_{1,f}, L_{2,f}, \omega_f).$$

Combining (55) with (56) we can write

$$\kappa = \mathcal{P}_2(u), \quad \text{with } u \equiv \{\{\kappa_\iota\}_{\iota \in \{l,b,h\}}, u_f\}, \quad (57)$$

where

$$\begin{aligned} \mathcal{P}_2(u) &= \mathcal{H}(\{\kappa_\iota\}_{\iota \in \{l,b,h\}}, \log(\mathcal{P}_1(u_f))) \\ &= \mathcal{H}(\{\kappa_\iota\}_{\iota \in \{l,b,h\}}, \log(\mathcal{P}_1(\lambda_f, \nu_f, \sigma_f, L_{1,f}, L_{2,f}, \omega_f))). \end{aligned} \quad (58)$$

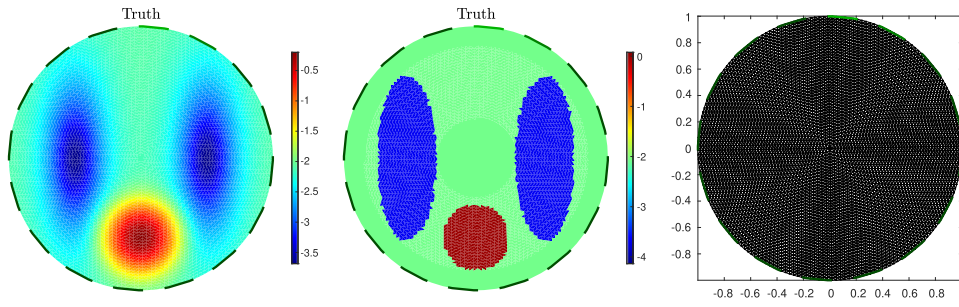


Figure 1. True log conductivity $\log(\kappa^\dagger)$ for **Exp_EIT₁** (left) and **Exp_EIT₂** (middle). Right: mesh for the generation with synthetic data for all EIT experiments.

The selection of the initial ensemble for u_f can be done similarly to the one in (54). This parameterisation can be extended for the case in which $\{\kappa_\nu\}_{\nu \in \{1, b, h\}}$ are unknown functions via using \mathcal{P}_1 to parameterise each of these functions.

4.3. Implementation

For the numerical experiments discussed in the following subsections, algorithms 2 and 3 are implemented in MATLAB. Let us recall that for these algorithms we need to construct the forward map $\mathcal{G} = \mathcal{F} \circ \mathcal{P}$ where \mathcal{F} is the operator induced by the CEM and \mathcal{P} is any of the parameterisations defined earlier. The numerical implementation of the CEM from subsection 4.1 is conducted using MATLAB software EIDORS [52]. The experimental setting consists of (i) a circular domain of unit radius centred at the origin, (ii) 16 surface electrodes with contact impedances of values 0.01 Ohms, (iii) an adjacent injection pattern with an electric current of 0.1 Amps, and (iv) measurements at each electrode. All these parameters are assumed known and fixed for the inversions of this section. Synthetic data are generated using the mesh from figure 1 (right) with 9216 elements, while a coarser mesh of 7744 elements is used for inversions. The total number of measurements is $M = 16^2 = 256$.

For the evaluation of the WM parameterisation $\kappa = \mathcal{P}_1(u)$, we solve (50) and (51) using the techniques from [47] and which restrict us to the cases in which $\nu \in \mathbb{N}$. The discretisation of the operator $\mathbb{I} - \nabla \cdot \text{diag}(L_1, L_2) \nabla$ with BCs from (51) is performed via cell-centred finite differences. The PDE in (50) and (51) is solve in a square domain, equal to, or enclosing D (the domain of definition for the PDE encoded in \mathcal{F}). The actual field $\kappa(x)$ that we pass into the CEM is an interpolation of $\mathcal{P}(u)$ on D . Parameterisations \mathcal{P}_2 is based on the truncation of the level-set function, $f(x)$ so the implementation is straightforward once all the fields κ_ν ($\nu \in 1, b, h$) and $f(x)$ are computed. Given these construction of \mathcal{G} , the rest of the steps in algorithms 2 and 3 are computed in a straightforward manner.

4.4. Measures of performance for EKI

Given the ensemble $\{u_n^{(j)}\}$ computed via EKI at the n iteration of the scheme ($n = 0$ corresponds to the prior ensemble), our estimate of the unknown property κ^\dagger is given by

$$\kappa_n \equiv \mathcal{P}(\bar{u}_n) = \mathcal{P} \left(\frac{1}{J} \sum_{j=1}^J u_n^{(j)} \right). \quad (59)$$

where \mathcal{P} is either \mathcal{P}_1 or \mathcal{P}_2 defined above. We measure the accuracy in terms of the relative error with respect to (w.r.t.) the truth defined by

$$\mathcal{E}_n = \frac{\|\kappa_n - \kappa^\dagger\|_{L^2(D)}}{\|\kappa^\dagger\|_{L^2(D)}}. \quad (60)$$

We often visualise some transformed ensemble members (mainly for the initial ensemble $n = 0$), i.e.

$$\kappa_n^{(j)} \equiv \mathcal{P}(u_n^{(j)}), \quad j \in \{1, \dots, J\}, \quad (61)$$

but note that our estimate κ_n in (59) does not involve taking the average of the particles in (61); this would be particularly detrimental for \mathcal{P}_2 since averaging (61) will not preserve discontinuities. For these two parameterisations we also visualise an estimate of the level-set function given by

$$f_n \equiv \log \mathcal{P}_1(\bar{u}_{f,n}) \quad (62)$$

We additionally monitor the following data-misfit quantities

$$\mathcal{DM}_{1,n} = \left\| \Gamma^{-1/2} \left[y - \frac{1}{J} \sum_{j=1}^J \mathcal{G}(u_n^{(j)}) \right] \right\| \quad (63)$$

$$\mathcal{DM}_{2,n} = \left\| \Gamma^{-1/2} (y - \mathcal{G}(\bar{u}_n)) \right\| \quad (64)$$

$$\mathcal{DM}_{3,n} = \left[\frac{1}{J} \sum_{j=1}^J \left\| \Gamma^{-1/2} (y - \mathcal{G}(u_n^{(j)})) \right\|^2 \right]^{1/2}. \quad (65)$$

4.5. Numerical experiment **Exp_EIT₁**. Continuous conductivity

For the first series of experiments the true conductivity, κ^\dagger , is a $C^\infty(D)$ function that we specify analytically; the plot of κ^\dagger is displayed in figure 1 (left). Synthetic data are constructed via $y = V^\dagger + \eta$ where $V^\dagger = \mathcal{F}(\kappa^\dagger)$ is computed with the CEM and η is a realisation from $N(0, \Gamma)$. We chose $\Gamma = \text{diag}(\gamma_1, \dots, \gamma_M)$ where

$$\begin{aligned} \gamma_m &= (10^{-2} |V_m^\dagger|)^2 + (10^{-3} |\max\{V_m^\dagger\}_{m=1}^M - \min\{V_m^\dagger\}_{m=1}^M|)^2 \\ m &= 1, \dots, M. \end{aligned} \quad (66)$$

The first term in the right-hand side of (66) corresponds to adding 1% Gaussian noise. The second term is added simply to avoid small variances from very small voltage (noise-free) measurements.

For this experiment we use the parameterisation of smooth functions, $\kappa = \mathcal{P}_1(u)$, from (52) with $u = (\lambda, L_1, L_2, \omega)$. Note that we have removed ν and σ from the inversion as we discussed in remark 5; we use fixed values $\nu = 3$ and $\sigma = 1.5$. The unknown consist of 3 scalars and 1 function, $\omega(x)$, that we discretise on a 100×100 grid. Upon discretisation, the dimension of the unknown is $\dim(\mathcal{U}) = 10\,003$. We follow the discussion of subsection 4.2.1 (see equation (54)) for the selection of the initial ensemble. More specifically, we select J particles $u_0^{(j)} = (\lambda^{(j)}, L_1^{(j)}, L_2^{(j)}, \omega^{(j)}) \sim \mu_0$ where we define the following prior

$$\mu_0 \equiv U[5 \times 10^{-3}, 1] \otimes U[0.15, 0.6] \otimes U[0.15, 0.6] \otimes N(0, \mathbb{I}) \quad (67)$$

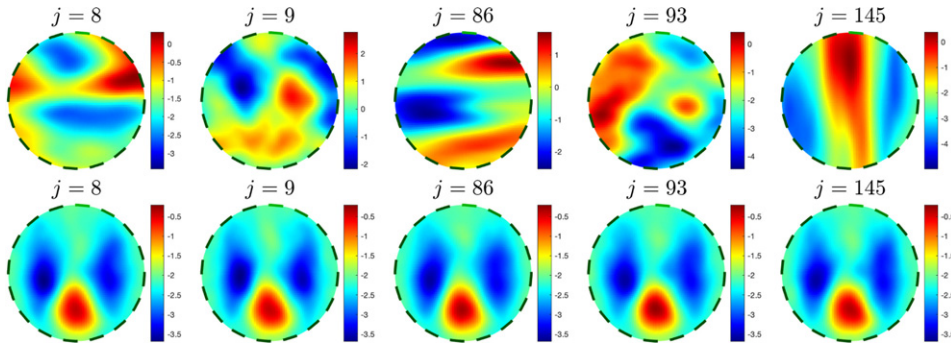


Figure 2. **Exp_EIT₁**. Top row: logarithm of five members from the prior ensemble $\{\kappa_0^{(j)}\}_{j=1}^J$. Bottom row: logarithm of five realisations from the final (converged) ensemble $\{\kappa_{n^*}^{(j)}\}_{j=1}^J$.

Table 1. Number of iterations n^* for EKI using the DMC (**EKI-DMC**) with various choices of J .

	Exp_EIT₁	Exp_EIT₂
DM $J = 100$	10.00 ± 0.53	13.20 ± 2.50
DM $J = 200$	10.00 ± 0.00	12.83 ± 0.59
DM $J = 400$	10.30 ± 0.47	14.53 ± 0.63
DM $J = 800$	11.03 ± 0.18	16.87 ± 0.51

where $U[a, b]$ denotes the uniform distribution on the interval $[a, b]$.

In figure 2 (top) we show plots for the logarithm of $\kappa_0^{(j)} = \mathcal{P}_1(u_0^{(j)})$ for five of members of the initial ensemble. Note that our choice of smoothness parameter $\nu = 3$ produces an initial ensemble that is quite smooth (recall from subsection 4.2.1 that the draws belong to $H^{3-\epsilon}(D)$). We also observe substantial differences in the degree of anisotropy which arises from our selection of a reasonably wide distribution of intrinsic lengthscales that we use to produce the initial ensemble.

4.5.1. Results from the inversion using EKI-DMC with various choices of J . Synthetic data and initial ensembles produced as described earlier are used as inputs for **EKI-DMC** (algorithm 3) with different choices of ensemble size J : 100, 200, 400, 800. For each choice of J , we conduct 30 experiments with different random selections of the initial ensemble. The plots of $(\log) \kappa_{n^*} = \mathcal{P}_1(\bar{u}_{n^*})$ (i.e. upon convergence) from one of these experiments, computed for each J , are displayed in figure 3 together with the truth (right panel). In figure 4 (left) we show boxplots of the relative error w.r.t. the truth \mathcal{E}_{n^*} (60), computed at the final iteration, from the set of 30 experiments conducted for each J . Boxplots of the data misfits defined in (63)–(65) are shown in the right panel of figure 4 where the mean noise level, approximated via $\delta = \sqrt{M}$, is indicated via the horizontal dotted red line. The average (over the 30 experiments) number of iterations to converge, n^* , is displayed in table 1. These experiments suggest that the choice of $J = 200$ provides a reasonable value of the data misfit (i.e. around the noise level). Furthermore, the average relative error w.r.t. the truth for $J = 200$ (approximately 30%) does not improve substantially as we increase J . Given these considerations, we select $J = 200$ for subsequent runs of **Exp_EIT₁**.

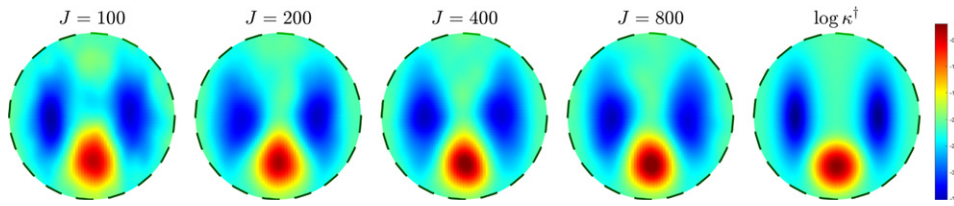


Figure 3. Exp_EIT₁. Logarithm of $\kappa_{n^*} \equiv \mathcal{P}_1(\bar{u}_{n^*})$ computed via the **EKI-DMC** with ensemble size (from left to right) $J = 100, 200, 400, 800$. Right panel shows the log of the truth.

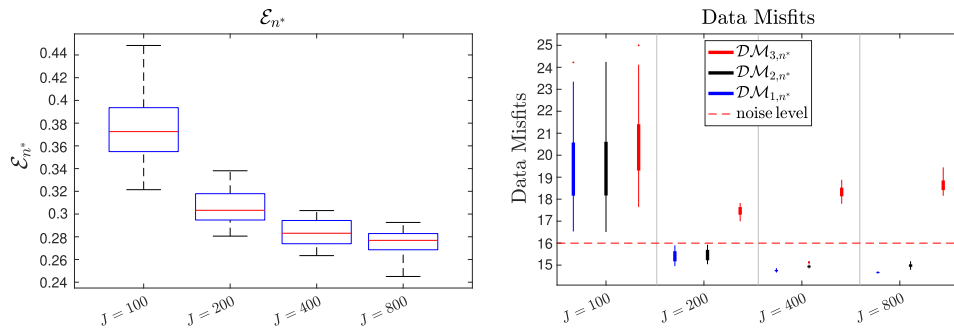


Figure 4. Exp_EIT₁. Error with respect to the truth (left), \mathcal{E}_{n^*} (see (60)), and data misfits from (63)–(65) (right) computed at the final iteration n^* via **EKI-DMC**. The noise level estimated by $\delta = \sqrt{M}$ is indicated with the dotted red-line in the right panel.

4.5.2. Further results from one run of EKI-DMC with $J = 200$. In figure 2 (bottom) we displayed 5 members of the final (converged) ensemble of $(\log \kappa_n^{(j)})$ corresponding to the initial ensemble from figure 2 (top). Plots of the log of $\kappa_n = \mathcal{P}_1(\bar{u}_n)$ (see equation (59)), at some of the intermediate iterations $1 \leq n < n^* = 10$ can be found in figure 5. In figure 6 (right) we plot, as a function of n , the values of the means $\bar{\lambda}$, \bar{L}_1 , and \bar{L}_2 , of the ensembles $\{\lambda^{(j)}\}_{j=1}^J$, $\{L_1^{(j)}\}_{j=1}^J$ and $\{L_2^{(j)}\}_{j=1}^J$, respectively⁴. Note that EKI produces a larger lengthscale in the vertical direction. This comes as no surprise since the truth κ^\dagger (see figure 1 (left)) has two inclusions of lower conductivity with larger correlation along the vertical direction. Finally, in figure 6 (right and middle) we display, for each of the 30 runs, the relative error w.r.t. the truth as well as the (log) data misfit (63), as a function of the iteration number n . We note that **EKI-DMC** is very robust and accurate across ensembles.

4.5.3. Comparison between EKI-DMC and EKI-LM. We compare the performance of algorithms 2 and 3 using the same set of 30 initial ensembles for each algorithm using $J = 200$. In [2, 3], this selection of J was sufficient to provide stable and accurate estimates for EIT. We consider different choices of the input ρ in **EKI-LM** and, for simplicity, we set $\tau = 1/\rho + 10^{-6}$. In figure 7 we show boxplots of the error w.r.t. the truth (right) and the data misfit \mathcal{DM}_{1,n^*}

⁴ We call from (52) and (53) that λ , L_1 and L_2 are scalar components of the unknown parameter u that we estimate via EKI. For ease in the notation we do not use the subscript n on these variables but we emphasise these are updated at each iteration of EKI.

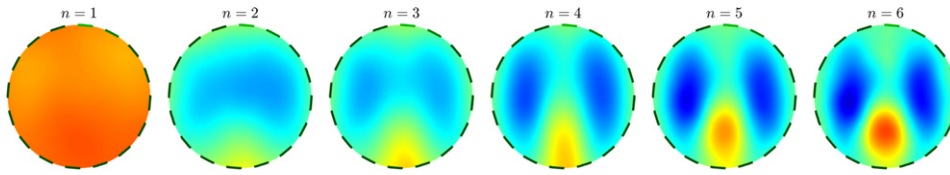


Figure 5. Exp_EIT₁. Logarithm of $\kappa_n = \mathcal{P}_1(\bar{u}_n)$ computed via the **EKI-DMC** at various intermediate iterations n ($1 \leq n \leq n^*$) computed using one ensemble of size $J = 200$.

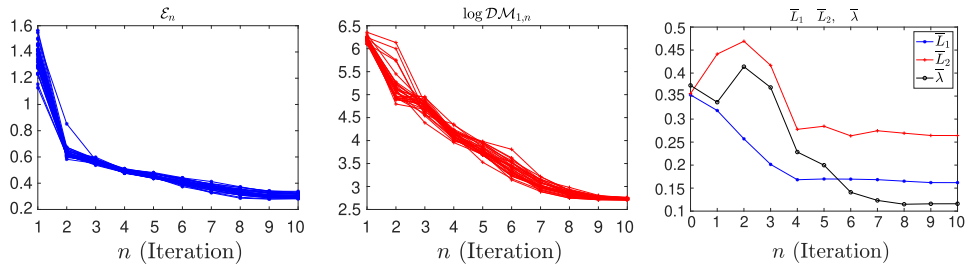


Figure 6. Exp_EIT₁. Plots of the relative error w.r.t. the truth (left), data misfit $\mathcal{DM}_{1,n}$ (middle) and means \bar{L}_1 , \bar{L}_2 and $\bar{\lambda}$ (right) as a function of n . The left and middle panels show results from 30 runs.

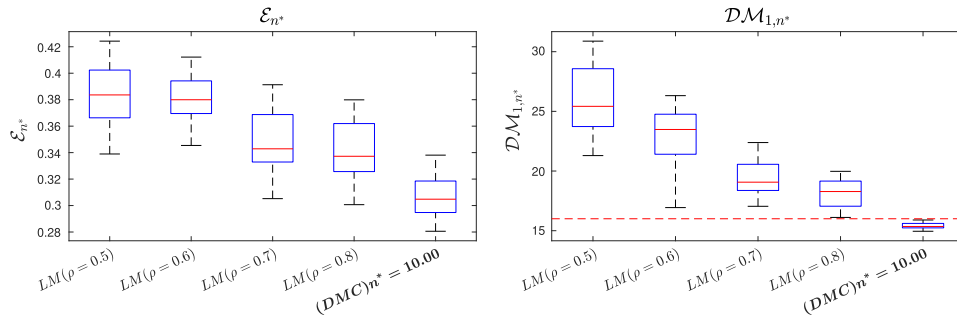


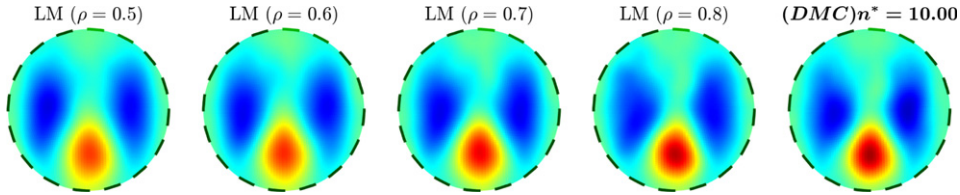
Figure 7. Exp_EIT₁. Error with respect to the truth (left), \mathcal{E}_{n^*} (see (60)), and data misfit \mathcal{DM}_{1,n^*} (63) (right) computed at the final iteration n^* using **EKI-DMC** and **EKI-LM** with various choices of ρ . The noise level estimated by $\delta = \sqrt{M}$ is indicated with the dotted red-line in the right panel.

(left) obtained with several choices of ρ for **EKI-LM**; the results from **EKI-DMC** for $J = 200$ are also included in these plots. Similar behaviour is observed for \mathcal{DM}_{2,n^*} and \mathcal{DM}_{3,n^*} and so these plots are omitted. The number of iterations for **EKI-LM** to achieve convergence, n^* , is displayed in table 2. For one of the 30 runs, in figure 8 we show the plots of the log of $\kappa_{n^*} = \mathcal{P}_1(\bar{u}_{n^*})$ computed with both algorithms using the selections of parameters described above.

While **EKI-DMC** does not depend on any tuning parameter, the results above show that **EKI-LM** is highly dependent on the selection of ρ (which is specified *a priori*). From table 2

Table 2. Number of iterations n^* for EKI using the LM approach (**EKI-LM**) with various choices of ρ .

	Exp_EIT ₁	Exp_EIT ₂
LM $\rho = 0.5$	8.57 ± 0.50	10.30 ± 0.79
LM $\rho = 0.6$	10.53 ± 0.51	14.10 ± 1.79
LM $\rho = 0.7$	14.40 ± 0.56	20.03 ± 1.45
LM $\rho = 0.8$	21.80 ± 0.76	32.53 ± 2.65

**Figure 8.** Exp_EIT₁. Logarithm of $\kappa_{n^*} = \mathcal{P}_1(\bar{u}_{n^*})$ computed using the same initial ensemble ($J = 200$) with **EKI-LM** for several choices of ρ . In the top-right panel with display the log of κ_{n^*} that we obtain with **EKI-DMC**.

we see that, for $\rho = 0.6$, the computational cost of **EKI-LM** is similar to the cost of **EKI-DMC**; i.e. convergence in approximately 10 iterations (in average). For this ρ , the average error computed with **EKI-LM** ($\approx 38\%$) is, however, larger than the error obtained via **EKI-DMC** ($\approx 30.5\%$). Although the improvement in accuracy of **EKI-DMC** over **EKI-LM** may not be overly impressive, from figure 8 (computed from one run) we note that the area of high conductivity is not accurately captured by **EKI-LM** with $\rho < 0.8$. For these experiments, if instead of $\rho = 0.6$ we choose, say $\rho = 0.8$ (see table 2), the computational cost of **EKI-LM** doubles without improving its accuracy with respect to **EKI-DMC**.

Furthermore, from figure 7 we also observe that both the data-misfit \mathcal{DM}_{1,n^*} and the error w.r.t. the truth obtained via **EKI-LM** decreases as we increase ρ . This behaviour is expected since large ρ implies smaller τ and, hence, smaller data misfits upon convergence (see stopping criteria for **EKI-LM** in equation (22)). Even though the measures of performance of **EKI-LM** (see figure 7) seem to approach those of **EKI-DMC** for increasing ρ , experiments (not shown) conducted with even larger ρ 's, (i.e. $\rho \geq 0.9$) do not yield the same level of accuracy that we achieve with **EKI-LM**. This is due to the fact that the stopping criteria in **EKI-LM** does not allow the algorithm to iterate once \mathcal{DM}_{1,n^*} is smaller than $\tau\delta = \tau\sqrt{M}$ (recall $\tau > 1/\rho$). However, for this particular problem, it seems that iterating below this threshold yields more accurate estimates. Indeed, we note that for the estimates obtained by **EKI-DMC**, the value of \mathcal{DM}_{1,n^*} is smaller than the $\tau\delta$ for all the 30 runs. We conclude that, for this particular problem setting, the stopping criteria in **EKI-LM** is not optimal.

In order to further understand the potential limitation of the stopping criteria in **EKI-LM** as proposed in [2], in figure 9 we display the behaviour of $\log \alpha_n$ as a function of n , computed from one run (with same initial ensemble) of **EKI-DMC** and **EKI-LM** with $\rho = 0.6$. As mentioned above, for this particular selection of ρ , both **EKI-LM** and **EKI-DMC** converged in 10 iterations so the cost of running both algorithms is the same. We find that the selection of α_n via **EKI-LM** yields quite large values of α_n compared to those obtained via **EKI-DMC**. Larger α_n 's, as discussed in subsection 2.1, means slower convergence. Although we can obtain smaller α_n 's using **EKI-LM** with smaller ρ 's, this will result

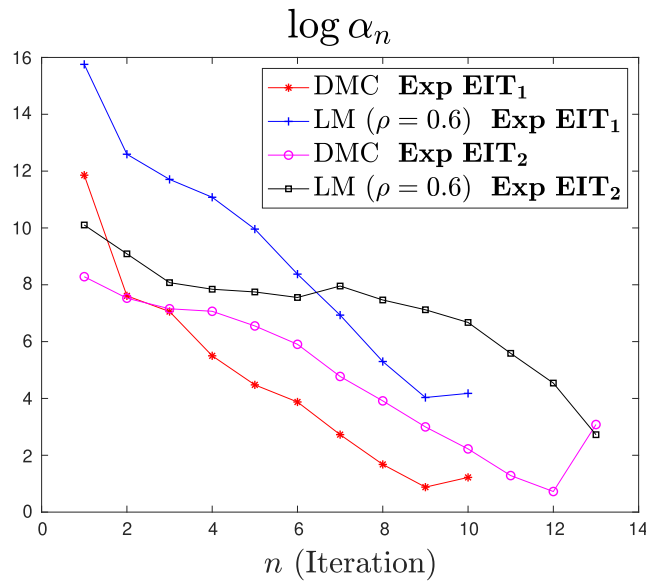


Figure 9. Logarithm of the regularisation parameter α_n as a function of n computed with **EKI-DMC** and **EKI-LM** for EIT experiments.

in larger τ 's and, consequently, earlier termination with is detrimental to the accuracy of the algorithm.

The aforementioned limitation of the stopping criteria in **EKI-LM** can be addressed, for example, by adopting the same stopping rule that we use in **EKI-DMC**. In other words, we may replace (22) by imposing that the α_n 's in **EKI-LM** satisfy (12). Such an approach has been used, for example, in [42] in the context of history matching of petroleum reservoirs. We have repeated our experiments (not shown) using (12) as stopping rule for **EKI-LM** which show that, indeed, using (12) enable us to run the algorithm beyond the noise level and, hence, obtain more accurate estimates at a lower computational cost. However, the role of ρ still determines the performance of **EKI-LM** and, while (12) allows us to use smaller ρ 's (thus smaller α_n 's and faster convergence), for some of these small values the algorithm loses stability.

In summary, we may find problem-specific tuning parameters ρ and τ in **EKI-LM**, or modifications of **EKI-LM** by adopting (12) as stopping rule, that will yield comparable level of accuracy and even similar computational cost to those achieved by **EKI-DMC**. However, finding optimal stopping criteria and/or tuning parameters may require thorough numerical testing which is often computationally intensive and unsuitable for practical large-scale applications.

4.6. Numerical experiment **Exp_EIT₂**. Discontinuous (piecewise constant) conductivity

The true conductivity for the experiments of this section is the piecewise constant function with plot displayed in figure 1 (middle). The background, low and high conductivity regions have constant values:

$$\kappa_b^\dagger = 0.125, \quad \kappa_l^\dagger = 0.025, \quad \kappa_h^\dagger = 1.0, \quad (68)$$

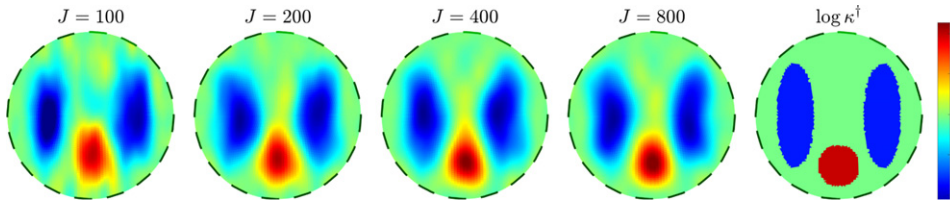


Figure 10. Logarithm of κ_{n^*} computed via **EKI-DMC** ensemble size (from left to right) $J = 100, 200, 400, 800$. Right panel shows $\log(\kappa^\dagger)$. This results use the parameterisation from **Exp_EIT₁**.

respectively. Synthetic data are generated in the same way as described in the previous subsection, and with noise covariance from (66).

In figure 10 we show the plots of $(\log) \kappa_{n^*}$ computed from one run with **EKI-DMC** using the same parameterisation (for continuous functions) and initial ensemble from **Exp_EIT₁**. While these results show that we can identify the three different regions of different conductivity, the reconstruction of the interface between these regions is quite inaccurate because of smoothness enforced by the parameterisation used within EKI. We overcome this limitation by means of a level-set parameterisations within the EKI framework.

4.6.1. Level-set parameterisation and prior ensemble. For the series of experiment in this subsection we test EKI with the level-set parameterisation $\kappa = \mathcal{P}_2(u)$ from (57). As discussed in remark 5 we remove σ_f and ν_f from the inversion. Here we use the values $\nu_f = 2$ and $\sigma_f = 0.5$. In addition, we also leave λ_f out of the inversion; we take $\lambda_f = 1$. This selection is made for simplicity since we expect to capture all the variability of the level-set function via the term $\mathcal{W}_{\Theta_f} \omega_f$ in (56). The unknown reduces to

$$u = (\kappa_1, \kappa_b, \kappa_h, L_{1,f}, L_{2,f}, \omega_f),$$

which consists of 5 scalars and one function, ω_f , which we discretise also on a 100×100 grid. Hence, the dimension of the unknown u is $\dim(\mathcal{U}) = 10\,005$. We select the initial ensemble according to

$$\begin{aligned} u_0^{(j)} &\sim U[0.015, 0.075] \otimes U[0.1, 0.4] \otimes U[0.65, 1.1] \\ &\otimes U[0.15, 0.6] \otimes U[0.15, 0.6] \otimes N(0, \mathbb{I}), \end{aligned} \quad (69)$$

From (56) we know that this selection produces a centred ensemble of initial level-set functions. It is worth noticing from (69) that there is no overlapping among the support of the distributions for the conductivity values (i.e. κ_1 , κ_b and κ_h) on each region. Furthermore, each of these intervals contain the true values from (68). Therefore, we work under the assumption that (i) there is clear contrast between the unknown values on each region and (ii) we have good knowledge of the range of possible values for the unknown conductivity in each of those regions.

In figure 11 we show the plots of some $\log \kappa_0^{(j)} = \log \mathcal{P}_2(u_0^{(j)})$ (top panels) and the corresponding level-set functions $f_0^{(j)} = \log(\mathcal{P}_1(u_{f,0}^{(j)}))$ (bottom panels) from five ensemble members. Our selection $\nu_f = 2$ imposes smoothness in the level-set function and, in turn, in the interface between the three regions of low, high and background conductivity. From figure 11 (bottom) we notice that the ensemble of level-sets displays anisotropy induced by randomising

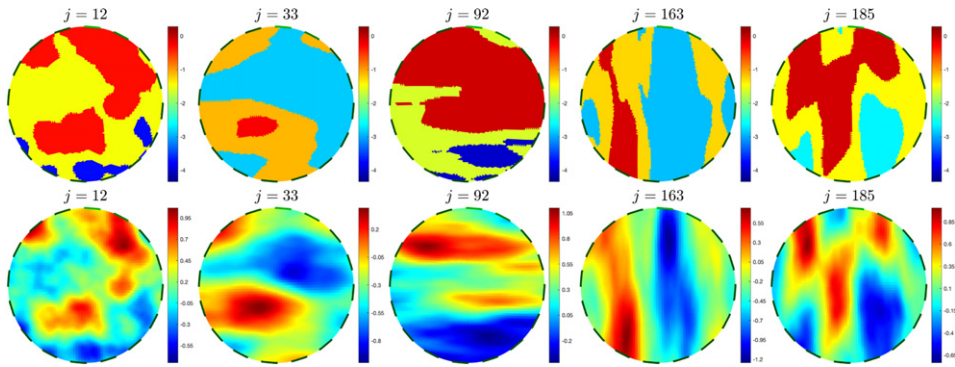


Figure 11. Exp_EIT₂. Five members of the initial ensemble of $(\log) \{\kappa_0^{(j)}\}_{j=1}^J$ (top row) and their corresponding level-set set function $\{f_0^{(j)}\}_{j=1}^J$ (bottom row).

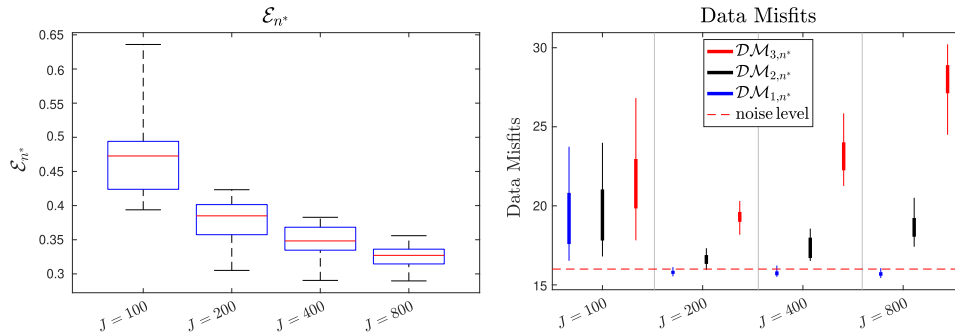


Figure 12. Exp_EIT₂. Error with respect to the truth (left), \mathcal{E}_{n^*} (see (60)), and data misfits from (63)–(65) (right) computed at the final iteration n^* with **EKI-DMC** with different choices of J . The noise level estimated by $\delta = \sqrt{M}$ is indicated with the dotted red-line in the right panel.

the intrinsic lengthscales in the level-set function. This variability can be seen in the corresponding interface between regions of different conductivities (top row figure 11). Note that the values of conductivity within each region are variable across particles.

4.6.2. Results from several choices of J in **EKI-DMC.** In figure 12 we show boxplots of the relative error w.r.t. the truth as well as data misfits (63)–(65). As before, these are results from 30 **EKI-DMC** runs using different initial ensembles for each choice of J . We can clearly see a decrease in the error w.r.t. the truth as we increase J , while the data misfits $\mathcal{DM}_{1,n}$ and $\mathcal{DM}_{2,n}$ achieve values close to the noise level (indicated via dotted red line) for $J \geq 200$. We see that, again, $J = 200$ is a good compromise between accuracy and cost. Using $J = 400$ will double the cost with a marginal improvement in accuracy. The choice of $J = 200$ also yields reasonable visual results as we can verify from figure 13 where, for one run of **EKI-DMC**, we display the log of $\kappa_{n^*} = \mathcal{P}_2(\bar{u}_{n^*})$ (top panels) and the level-set function $f_{n^*} = \log(\mathcal{P}_1(\bar{u}_{f,n^*}))$ (bottom panels).

Although we note that the average data misfit \mathcal{DM}_{3,n^*} in figure 12 seems to increase with J , experiments (not shown) with larger J suggest that \mathcal{DM}_{3,n^*} eventually stabilises. The increase

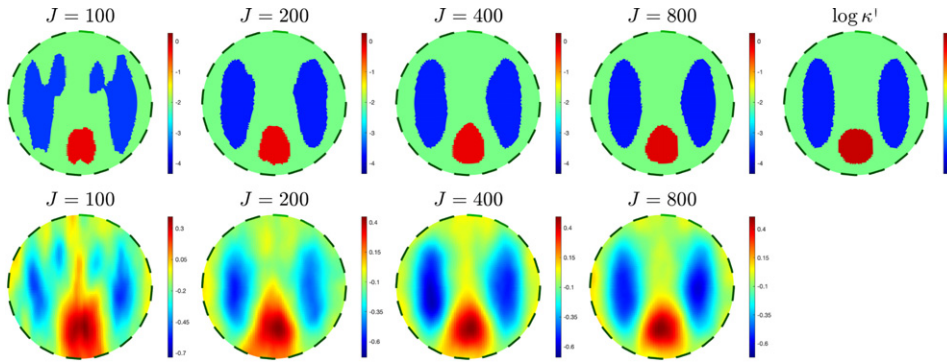


Figure 13. Exp_EIT₂. Plots of $(\log) \kappa_{n^*}$ (top row) and the corresponding level-set function f_{n^*} (bottom row) computed with EKI-DMC with different choice of ensemble size J . Top-right panel shows $\log(\kappa^\dagger)$.

in \mathcal{DM}_{3,n^*} can be attributed to the fact that larger J 's produces a better spread/coverage of the distribution of particles; some of these particles yield a larger data misfit within EKI⁵. The ensemble mean \bar{u}_n , however, is quite accurate (see figure 13, for $J = 200$) and so the corresponding $\kappa_n = \mathcal{P}_2(\bar{u}_n)$ yields reasonable values of the \mathcal{DM}_{2,n^*} (i.e. close to the noise level).

4.6.3. Results from one run of EKI-DMC with $J = 200$. Figure 14 shows some members from the final (converged) ensemble corresponding to the initial ensemble from figure 11; the top panels are $\log \kappa_{n^*}^{(j)} = \mathcal{P}_2(u_{n^*}^{(j)})$ while the bottom panels show the level-set functions $f_{n^*}^{(j)} = \log(\mathcal{P}_1(u_{f,n^*}^{(j)}))$. This figure shows that there is, indeed, significant variability across particles of the converged ensemble $\kappa_{n^*}^{(j)}$ and, in turn, possible large spread in the values of the data misfit obtained using each of these conductivities (hence potentially large values of $\mathcal{DM}_{3,n}$). The average error w.r.t. the truth and (log) data misfit $\mathcal{DM}_{1,n}$ are shown in figure 15, as a function of the iteration number n (these are results from our 30 runs). Note that, in contrast to the previous experiment, the error displays more variability across ensembles.

Plots of $\log \kappa_n$ and the level-set function f_n , at some of the intermediate iterations $1 \leq n < n^*$, are displayed in the top and bottom panels of figure 16, respectively. We can see that EKI not only estimates the shape (via the level-set function) of the regions with different conductivity but also the conductivity values on each region. In figure 17 (top) we plot, as a function of n , the values of $\bar{L}_{1,f}$, $\bar{L}_{2,f}$, $\bar{\lambda}_1$, $\bar{\lambda}_b$ and $\bar{\lambda}_h$, i.e. the means of the ensembles $\{L_{1,f}^{(j)}\}_{j=1}^J$, $\{L_{2,f}^{(j)}\}_{j=1}^J$, $\{\lambda_1^{(j)}\}_{j=1}^J$, $\{\lambda_b^{(j)}\}_{j=1}^J$, and $\{\lambda_h^{(j)}\}_{j=1}^J$ (we reiterate that these variables are updated at each iteration of EKI). The ensemble mean for the intrinsic lengthscale of the level-set function in the vertical direction is larger than the horizontal one. Again, this is due to the presence of the regions of low conductivity which are longer in the vertical direction. From the middle-right panel we can see that the true conductivity in the background region $\bar{\lambda}_b^\dagger$ is recovered quite accurately. In contrast, the mean values $\bar{\lambda}_1$ and $\bar{\lambda}_h$ do not seem to vary much with respect to the mean.

⁵ From the Bayesian perspective some of these particles have small (approximate) posterior probability.

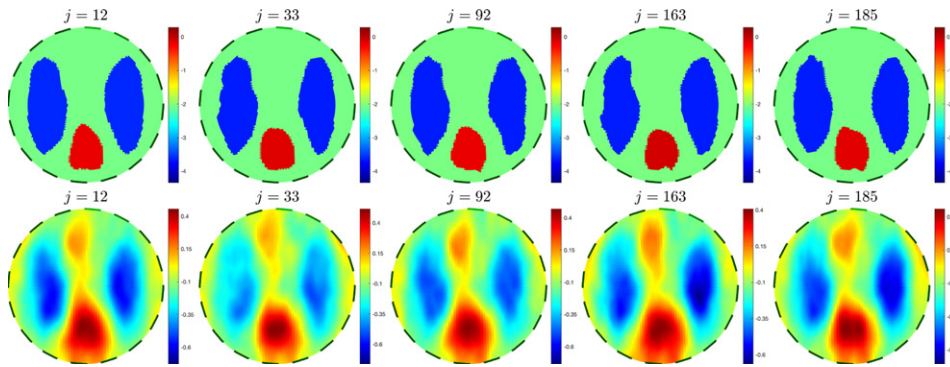


Figure 14. Exp_EIT₂. Five members of the converged ensemble of $(\log) \{\kappa_n^{(j)}\}_{j=1}^J$ (top row) and their corresponding level-set set function $\{f_n^{(j)}\}_{j=1}^J$ (bottom row).

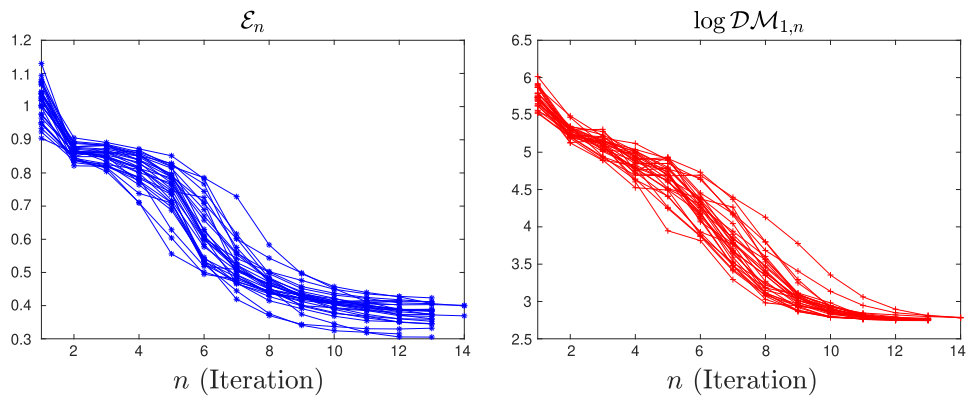


Figure 15. Exp_EIT₂. Plots of the relative error w.r.t. the truth (left) and data misfit $\mathcal{DM}_{1,n}$ (right) computed from 30 runs with EKI-DMC.

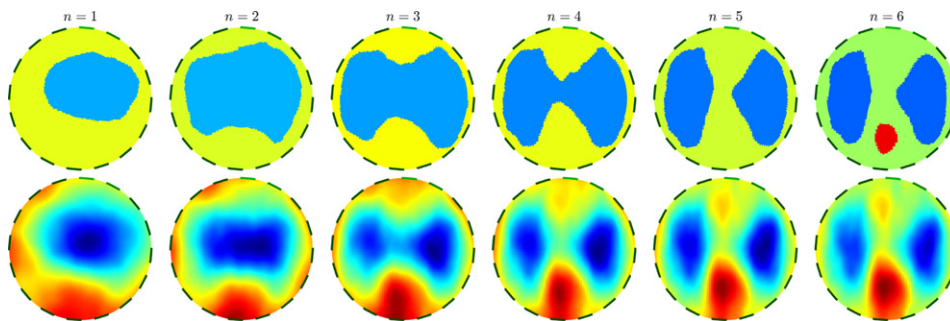


Figure 16. Exp_EIT₂. Logarithm of κ_n (top) and the corresponding level-set f_n (bottom) computed with one run of EKI-DMC (with $J = 200$) at various intermediate iterations n ($1 \leq n \leq n^*$).

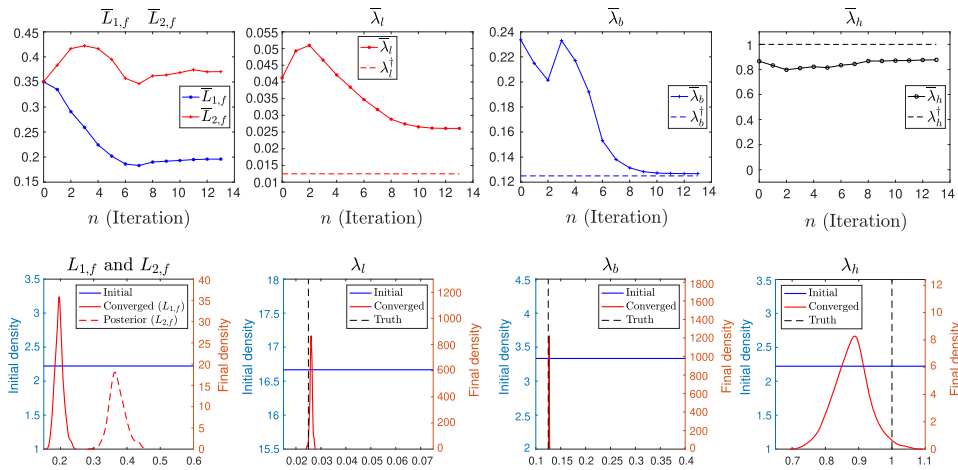


Figure 17. Exp_EIT₂. Top: plots of the means $\bar{L}_{1,f}$ and $\bar{L}_{2,f}$ (left) as well as the mean conductivity values $\bar{\lambda}_l$ (left-middle), $\bar{\lambda}_b$ (middle-right) and $\bar{\lambda}_h$ (right) corresponding to the low, background and high conductivity regions. Bottom: densities from the initial and final (converged) ensemble of $L_{1,f}$, $L_{2,f}$, λ_l , λ_b , λ_h .

Although we are mainly focus on the deterministic case here, to further appreciate the accuracy of the inversion for these variables, in figure 17 (bottom) we show their probabilistic densities approximated from the initial and final (converged) ensembles. For most of these variables we can see that the converged ensemble has a much smaller variance compared to the initial one. For the conductivity values we note that the lower and background values are identified accurately with the ensemble mean; the higher value is captured in the tail of the final ensemble.

4.6.4. Comparison between EKI-DMC and EKI-LM. In figure 18 we compare the performance of **EKI-DMC** and **EKI-LM** using 30 different initial ensembles. For **EKI-LM** we explore different choices of ρ . We note that **EKI-DMC** outperforms **EKI-LM** for all our choices of ρ . From table 2 note that for $\rho > 0.7$ the computational cost of **EKI-LM** is approximately twice the cost of **EKI-DMC** and the cost even triple if we choose $\rho = 0.8$. The plots of $\log \kappa_n$ from one run with the three algorithms (same initial ensemble) are shown in figure 19, where we see that all these runs perform well. Similarly conclusions to those in **Exp_EIT₁** are also drawn for this case. Namely, **EKI-DMC** is more accurate than **EKI-LM** in the chosen metrics although visually we achieve good performance from both algorithms and inputs. Experiments (not shown) indicate that substantial improvement in the performance of **EKI-LM** can be achieved by using the same stopping rule given by (12) instead of (22). However, some choices of ρ can result in inaccurate estimates in terms of error w.r.t. the truth. The main advantage of **EKI-DMC** over **EKI-LM** is that the former does not rely on tuning parameters whose optimal selection is crucial for the stability and accuracy of the scheme.

4.7. Further discussions

Let us recall that the proposed stopping criteria for EKI given by (12) comes naturally from the Bayesian tempering scheme. The tempering parameter $t_{N+1} = 1$ in (8) corresponds to the posterior $\mu_{N+1} = \mu$. From the deterministic perspective in which we aim approximating (2), this stopping criteria might seem rather arbitrary, and one may argue that further iterations of

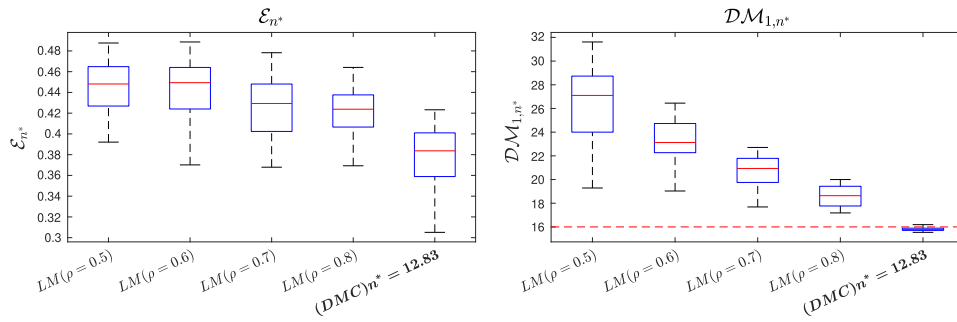


Figure 18. Exp_EIT₂. Error with respect to the truth (left), \mathcal{E}_{n^*} (see (60)), and data misfit \mathcal{DM}_{1,n^*} (63) (right) computed at the final iteration n^* using **EKI-DMC** and **EKI-LM** with various choices of ρ . The noise level estimated by $\delta = \sqrt{M}$ is indicated with the dotted red-line in the right panel.

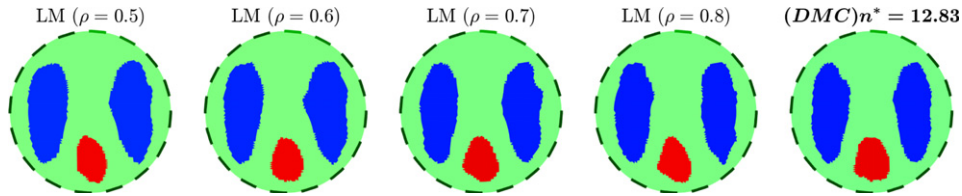


Figure 19. Exp_EIT₂. Logarithm of κ_{n^*} computed using the same initial ensemble ($J = 200$) with **EKI-LM** using different selections of the parameter ρ . In the right panel with display the log of κ_{n^*} that we obtain using **EKI-DMC**.

the EKI algorithm (i.e. for $t_{N+1} > 1$) may produce more accurate approximations of (2). Our numerical examples show that the stopping criteria from (12) yields values of the data misfit which are consistent with the discrepancy principle (38) where, as stated earlier, the noise level δ is approximated by $\sqrt{M} = 16$. In particular, from the experiments with multiple runs (figures 6 and 15) we see that the (log) data misfit approaches the (log) noise level $\log \delta \approx 2.77$ as we reach the stopping criteria given by (12). In fact, for these experiments we note that the error w.r.t. the truth, \mathcal{E}_n , stabilises before we reach the stopping criteria. Hence, earlier stopping (i.e. $t_{N+1} < 1$) could have avoided one or two unnecessary iterations. In contrast, experiments (not shown) confirm that additional iterations (i.e. for $t_{N+1} > 1$) not only do not decrease \mathcal{E}_n , but they will eventually lead to an increase in \mathcal{E}_n and hence loss of accuracy. These observations align with existing work suggesting that the appropriate early stopping is crucial for the stability of EKI. As discussed in subsection 2.1, the work in [1, 2] has shown that the error with respect to the truth increase with the number of iterations once the data misfit takes values below the noise level (i.e. the discrepancy principle is not satisfied).

Similarly, even though we motivate the selection of α_n in the DMC by controlling (via Jeffrey's divergence) the transitions between consecutive tempering measures μ_{n+1} and μ_n , the aforementioned results from figures 6 and 15 constitute numerical evidence that DMC leads to a stable/regularised EKI algorithm that addresses the ill-posedness of the classical inverse problem (2). Intuitively, controlling the transitions between consecutive measures yields stable transitions between the statistics of these measures which are, in turn, approximated with Gaussians using EKI (see appendix A). Therefore, one may expect that, for sufficiently large

ensemble size, EKI also leads to stable computation of the ensemble mean. Nonetheless, theoretical work beyond the scope of this manuscript is needed in order to deepen our understanding of the link between the transition of tempering measures and the regularisation properties of **EKI-DMC** that arises from the proposed selection of α_n and the stopping criteria induced by constraint (12).

5. Conclusions

We introduced the DMC: a new adaptive regularisation strategy within the classical EKI setting. This led to an algorithm, **EKI-DMC**, that in contrast to existing EKI approaches, does not require any tuning parameters. Although we focus on the solution of deterministic identification problem, the proposed DMC is motivated from the Bayesian perspective of EKI within the tempering setting, where the inverse of the regularisation parameter, α_n^{-1} , controls the transition between two consecutive intermediate measures. The Bayesian tempering setting provides us with a condition that these parameters must satisfy ($\sum_{n=1}^{n^*} \alpha_n^{-1} = 1$) to bridge the prior and posterior. We encode this condition for the termination of **EKI-DMC** together with our new method for choosing α_n . We show that the selection satisfies a heuristic statistical discrepancy principle which also controls Jeffreys' divergence between consecutive measures.

We applied **EKI-DMC** for the solution of EIT with the CEM in which the conductivity was parameterised via different maps that enabled us, via the EKI framework applied to the appropriate parameters, to characterise both smooth and piece-wise constants conductivities. WM fields were at the core of these parameterisations which included the intrinsic lengthscales as inputs that we estimate within EKI. For the piece-wise constant case, we use a truncated WF field (the level-set function) to characterise discontinuities between different regions. As with any other EKI algorithms, our results show that the performance of **EKI-DMC** relies on reasonable choices of ensemble size, J . For sufficiently large choices of J , our experiments show that **EKI-DMC** is quite robust and capable of producing accurate identification of physical properties suitably parameterised.

We conducted a performance comparison between **EKI-DMC** and the **EKI-LM** approach of [2]. In most cases **EKI-DMC** outperforms the accuracy of **EKI-LM** in terms of error w.r.t. the truth and data misfit, but we noted that suitable choices of tuning parameters and possible modifications to the stopping rule can produce comparable performance to **EKI-DMC**. We recognise that similar performance can also be achieved by using different choices of α_n in EKI including those discussed in section 2. We reiterate that, as with **EKI-LM** and its variants, most regularisation approaches for EKI rely on the choice of additional tuning parameters which, of course, can be tuned to display comparable or even better performance to **EKI-DMC**. However, there is no principled approach for the optimal selection of those parameters. Optimal tuning parameters can only be informed via careful numerical investigations on the given problem-specific setting. Our results show that **EKI-DMC** is a robust self-tuning regularisation strategy of EKI, ideally suited for practical and operational settings for which finding optimal choices of tuning parameters in existing EKI approaches may not be computationally feasible.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council [Grant No. EP/P006701/1]; through the EPSRC Future Composites Manufacturing Research Hub.

We would like to thank the two anonymous reviewers for their insightful feedback and suggestions.

Appendix A. Motivation of EKI from the Bayesian tempering approach

We introduce a series of approximations and Gaussian assumptions that lead to the EKI algorithm (algorithm 1).

A.1. Linearisation and Gaussian approximations

Suppose that the collection of tempering parameters $\{t_n\}_{n=1}^N$ in (9) have been specified. Our objective now is to construct a sequence of Gaussian approximations of each measure μ_n in (8). To this end, let $\nu_0 = N(m_0, \mathcal{C}_0)$ be a Gaussian approximation of the prior μ_0 , and let us denote by $D\mathcal{G}$ the Frechet derivative of \mathcal{G} . We recursively construct a sequence of Gaussian approximations $\{\nu_n = N(m_n, \mathcal{C}_n)\}_{n=1}^{N+1}$ of $\{\mu_n\}_{n=1}^{N+1}$ via the following expression

$$\nu_{n+1}(du) \propto \nu_n(du) \exp \left[-\frac{1}{2} \left\| (\alpha_n \Gamma)^{-1/2} (y - \mathcal{G}_n - D\mathcal{G}_n(u - m_n)) \right\|^2 \right], \quad (\text{A.1})$$

where for ease in the notation we have defined $\mathcal{G}_n \equiv \mathcal{G}(m_n)$ and $D\mathcal{G}_n \equiv D\mathcal{G}(m_n)$. We note the right-hand side of (A.1) involves the linearisation of the forward map around the mean of $\nu_n = N(m_n, \mathcal{C}_n)$.

Recursive formulas for the mean and covariance of $\nu_n = N(m_n, \mathcal{C}_n)$ can be obtained by completing the square in (A.1) (see theorem 6.20 in [31]). Indeed, since $\nu_n = N(m_n, \mathcal{C}_n)$ and the model $L(u) \equiv D\mathcal{G}_n(u - m_n)$ is linear, then $\nu_{n+1} = N(m_{n+1}, \mathcal{C}_{n+1})$ with

$$m_{n+1} = m_n + \mathcal{C}_n D\mathcal{G}_n^* (D\mathcal{G}_n \mathcal{C}_n D\mathcal{G}_n^* + \alpha_n \Gamma)^{-1} (y - \mathcal{G}_n) \quad (\text{A.2})$$

$$\mathcal{C}_{n+1} = \mathcal{C}_n - \mathcal{C}_n D\mathcal{G}_n^* (D\mathcal{G}_n \mathcal{C}_n D\mathcal{G}_n^* + \alpha_n \Gamma)^{-1} D\mathcal{G}_n \mathcal{C}_n, \quad (\text{A.3})$$

where $D\mathcal{G}_n^*$ denotes the adjoint of $D\mathcal{G}$ evaluated at m_n .

Remark 6 (LM from linearised Bayesian tempering). Using standard arguments (see for example lemma 3.1 in [2]) it can be shown that m_n in (A.2) satisfies (6). In the case where \mathcal{C}_n is the identity operator, (6) yields the standard LM iterative scheme [26]. For the modified version in (6), we note that the recursive formula for the mean involves introducing the precision operator for u in the regularisation term in the right-hand side of (6).

Remark 7 (the linear-Gaussian case). Note that, if $\mu_0 = \nu_0$ (i.e. the prior is Gaussian) and the forward map \mathcal{G} is linear, then we have that $D\mathcal{G}_n(u - m_n) = \mathcal{G}(u - m_n) = \mathcal{G}(u) - \mathcal{G}_n$. Hence, (A.1) and (10) coincide and so $\mu_n = \nu_n$ for all $n = 0, \dots, N + 1$. In particular, in the linear-Gaussian case the final measure ν_n^* coincides with the posterior (i.e. $\nu_{N+1} = \mu_{N+1} = \mu$).

A.2. Derivative-free ensemble approximation

We now introduce further approximations that will lead to the EKI algorithm under consideration. Let us denote by u_n a random variable such that $u_n \sim \nu_n = N(m_n, \mathcal{C}_n)$. Denote by \mathbb{E}_n expectation with respect to ν_n . Let us consider the first order approximation

$$\mathcal{G}(u_n) \simeq \mathcal{G}_n + D\mathcal{G}_n(u_n - m_n) \quad (\text{A.4})$$

that we used in (A.1) to define our Gaussian approximations for the tempering scheme introduced in appendix A.1. From (A.4) it follows that

$$\mathbb{E}_n[\mathcal{G}(u_n)] \simeq \mathcal{G}_n, \quad \mathcal{G}(u_n) - \mathbb{E}_n[\mathcal{G}(u_n)] \simeq D\mathcal{G}_n(u_n - m_n). \quad (\text{A.5})$$

Hence,

$$\begin{aligned} \text{Cov}_n(u_n, \mathcal{G}(u_n)) &\equiv \mathbb{E}_n[(u_n - m_n) \otimes (\mathcal{G}(u_n) - \mathbb{E}_n[\mathcal{G}(u_n)])] \\ &\simeq \mathcal{C}_n D\mathcal{G}_n^*, \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} \text{Cov}_n(\mathcal{G}(u_n)) &\equiv \mathbb{E}_n[(\mathcal{G}(u_n) - \mathbb{E}_n[\mathcal{G}(u_n)]) \otimes (\mathcal{G}(u_n) - \mathbb{E}_n[\mathcal{G}(u_n)])] \\ &\simeq D\mathcal{G}_n \mathcal{C}_n D\mathcal{G}_n^*, \end{aligned} \quad (\text{A.7})$$

where we have used the fact that $\mathbb{E}_n[(u_n - m_n) \otimes (u_n - m_n)] = \mathcal{C}_n$. If we use approximations (A.5)–(A.7) in (A.2) and (A.3) we note that

$$m_{n+1} \simeq \tilde{m}_{n+1} \equiv m_n + \text{Cov}_n(u_n, \mathcal{G}(u_n))(\text{Cov}_n(\mathcal{G}(u_n)) + \alpha_n \Gamma)^{-1}(y - \mathbb{E}_n[\mathcal{G}(u_n)]) \quad (\text{A.8})$$

$$\mathcal{C}_{n+1} \simeq \tilde{\mathcal{C}}_{n+1} \equiv \mathcal{C}_n - \text{Cov}_n(u_n, \mathcal{G}(u_n))(\text{Cov}_n(\mathcal{G}(u_n)) + \alpha_n \Gamma)^{-1} \text{Cov}_n(\mathcal{G}(u_n), u_n). \quad (\text{A.9})$$

These approximations to the mean and covariance of the sequence of approximate measures $\{\nu_n\}_{n=1}^{N+1}$ do not involve derivatives of the forward map. However, the covariance and cross-covariance that appear in (A.8) and (A.9) cannot be computed in closed form. This issue is overcome by using particle approximations of each approximate measure $\nu_n = N(m_n, \mathcal{C}_n)$. In other words, we consider approximations

$$\nu_n^J(u_n) = \frac{1}{J} \sum_{j=1}^J \delta(u_n - u_n^{(j)}), \quad (\text{A.10})$$

where we assume that $u_n^{(j)} \sim N(m_n, \mathcal{C}_n)$. The classical EKI update formula (4) for the ensemble of particles $\{u_n^{(j)}\}_{j=1}^J$ is defined in such a way, that the corresponding ensemble mean and covariance approximate those in (A.8) and (A.9) as $J \rightarrow \infty$. To see this more clearly, let us note that the ensemble approximation of $\mathbb{E}_n[u_n]$ and $\mathbb{E}_n[\mathcal{G}(u_n)]$ are \bar{u}_n and $\bar{\mathcal{G}}_n$ defined in (5). The ensemble approximations of $\text{Cov}_n(\mathcal{G}(u_n))$, $\text{Cov}_n(u_n, \mathcal{G}(u_n))$ and \mathcal{C}_n , denoted by $\mathcal{C}_n^{\mathcal{G}\mathcal{G}}$, $\mathcal{C}_n^{u\mathcal{G}}$ and \mathcal{C}_n^{uu} , are defined by (16), (17) and

$$\mathcal{C}_n^{uu} \equiv \frac{1}{J-1} \sum_{j=1}^J (u_n^{(j)} - \bar{u}_n) \otimes (u_n^{(j)} - \bar{u}_n), \quad (\text{A.11})$$

respectively. From (4) we note that

$$\bar{u}_{n+1} = \bar{u}_n + \mathcal{C}_n^{u\mathcal{G}}(\mathcal{C}_n^{\mathcal{G}\mathcal{G}} + \alpha_n \Gamma)^{-1}(y + \sqrt{\alpha_n} \bar{\xi}_n - \bar{\mathcal{G}}_n). \quad (\text{A.12})$$

It can be shown (see for example [53] for a rigorous proof in finite dimensions) that

$$\begin{aligned} \bar{u}_{n+1} &\rightarrow \tilde{m}_{n+1} \equiv m_n + \text{Cov}_n(u_n, \mathcal{G}(u_n))(\text{Cov}_n(\mathcal{G}(u_n)) + \alpha_n \Gamma)^{-1} \\ &\quad \times (y - \mathbb{E}_n[\mathcal{G}(u_n)]) \simeq m_{n+1} \end{aligned} \quad (\text{A.13})$$

$$\begin{aligned} \mathcal{C}_{n+1}^{uu} &\rightarrow \tilde{\mathcal{C}}_{n+1} \equiv \mathcal{C}_n - \text{Cov}_n(u_n, \mathcal{G}(u_n))(\text{Cov}_n(\mathcal{G}(u_n)) + \alpha_n \Gamma)^{-1} \text{Cov}_n \\ &\quad \times (\mathcal{G}(u_n), u_n) \simeq \mathcal{C}_{n+1}. \end{aligned} \quad (\text{A.14})$$

as $J \rightarrow \infty$. Moreover, for the particles in (4) we have

$$\nu_{n+1}^J(u_{n+1}) = \frac{1}{J} \sum_{j=1}^J \delta(u_{n+1} - u_{n+1}^{(j)}) \rightarrow \tilde{\nu}_{n+1} \equiv N(\tilde{m}_{n+1}, \tilde{\mathcal{C}}_{n+1}).$$

The development above (informally) shows that if $\nu_n^J \simeq \nu_n = N(m_n, \mathcal{C}_n)$, then $\tilde{\nu}_{n+1}^J \simeq \tilde{\nu}_{n+1} = N(\tilde{m}_{n+1}, \tilde{\mathcal{C}}_{n+1})$. Furthermore, if (A.5)–(A.7) are accurate enough, then $N(\tilde{m}_{n+1}, \tilde{\mathcal{C}}_{n+1}) \simeq \nu_{n+1} = N(m_{n+1}, \mathcal{C}_{n+1})$ and so the EKI ensemble $\tilde{\nu}_{n+1}^J$ approximates ν_{n+1} . Recall that the ν_n 's are Gaussian approximations of the tempered distributions μ_n . Hence the regularisation parameter α_n in EKI is the inverse of the difference between consecutive tempering parameters (see equation (11)) which is, in turn, a Gaussian approximation of μ_{n+1} .

Remark 8 (squared-root EKI). It is worth noticing that other approaches can be used to approximate $N(\tilde{m}_{n+1}, \tilde{\mathcal{C}}_{n+1})$ above. This includes the so-called ensemble square-root formulations [54] in which the particles are cleverly updated so that their sample mean and covariance coincide (exactly) with \tilde{m}_{n+1} and $\tilde{\mathcal{C}}_{n+1}$. While these has been shown to be beneficial for very small samples (i.e. < 50), we note that $N(\tilde{m}_{n+1}, \tilde{\mathcal{C}}_{n+1})$ does not, in general, coincides with $\nu_{n+1} = N(m_{n+1}, \mathcal{C}_{n+1})$ (unless \mathcal{G} is linear).

Remark 9 (EKI as a derivative-free approximation of LM). For sufficiently large J , the mean of the ensemble \bar{u}_{n+1} approximates \tilde{m}_{n+1} and so m_{n+1} which is, in turn, the iteration of the LM scheme in (6) (see remark 9). Therefore, we can interpret EKI as a derivative-free approximation of the LM scheme constrained to the subspace generated by the initial ensemble $\{u_0^{(j)}\}_{j=1}^J$. More specifically, the ensemble mean \bar{u}_n define by the recursive formula (5) satisfies the following subspace invariance property (see theorem 2.1 in [1])

$$\bar{u}_n \in \mathcal{S}_0 \equiv \text{span}\{u_0^{(j)}\}_{j=1}^J$$

for all $n \in \mathbb{N}$. We expect EKI to produce approximate solutions to (4) within the subspace defined above. While the numerical experiments from [2] provides evidence of such a claim, to the best of our knowledge, the convergence of EKI in this context is still an open problem.

Appendix B. Lemmas, theorems, and proofs

Let $y \in \mathbb{R}^M$ fixed. We recall that from assumption 1 on the forward map \mathcal{G} , the data misfit $\Phi(\cdot, y) : \mathcal{H} \rightarrow [0, +\infty)$ is a square integrable functional on the probability space $(\mathcal{H}, \mathcal{B}(\mathcal{H}), \mu_0)$, i.e.

$$\int_{\mathcal{H}} |\Phi(u; y)|^2 \mu_0(du) < \infty. \quad (\text{B.1})$$

This implies that $\Phi(\cdot, y) : \mathcal{H} \rightarrow [0, +\infty)$ is absolutely integrable functional on the probability space $(\mathcal{H}, \mathcal{B}(\mathcal{H}), \mu_0)$, i.e.

$$\int_{\mathcal{H}} |\Phi(u; y)| \mu_0(du) \leq M_0 < \infty. \quad (\text{B.2})$$

B.1. Proof of lemma 1

Proof. Assumption 1 allows us to use the techniques from [[31], theorems 4.1 and 4.2] and [[55], theorem 3.4] to show that $\mu_t \ll \mu_0$. We now prove $\mu_0 \ll \mu_t$, i.e. $\forall \mathcal{X} \in \mathcal{B}(\mathcal{H})$, $\mu_t(\mathcal{X}) = 0 \Rightarrow \mu_0(\mathcal{X}) = 0$. We prove this statement by contradiction. Assume that there exists $\mathcal{X} \in \mathcal{B}(\mathcal{H})$ such that $\mu_t(\mathcal{X}) = 0$ and $\mu_0(\mathcal{X}) > 0$. Since $\mu_0(\mathcal{X}) > 0$, we can define an another probability measure, $\nu_{\mathcal{X}}$, on the measurable space $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, such that for all $\mathcal{S} \in \mathcal{B}(\mathcal{X})$,

$$\nu_{\mathcal{X}}(\mathcal{S}) := \frac{\mu_0(\mathcal{S})}{\mu_0(\mathcal{X})}. \quad (\text{B.3})$$

Using $\mu_t \ll \mu_0$ and expressions (28) and (B.3) we have

$$\begin{aligned} \mu_t(\mathcal{X}) &:= \int_{\mathcal{X}} \mu_t(\mathrm{d}u) = \int_{\mathcal{X}} \frac{\mathrm{d}\mu_t}{\mathrm{d}\mu_0}(u) \mu_0(\mathrm{d}u) \\ &= \frac{1}{N_t} \int_{\mathcal{X}} \exp(-t\Phi(u; y)) \mu_0(\mathrm{d}u) \\ &= \frac{\mu_0(\mathcal{X})}{N_t} \left(\int_{\mathcal{X}} \exp(-t\Phi(u; y)) \nu_{\mathcal{X}}(\mathrm{d}u) \right). \end{aligned}$$

We use Jensen's inequality, formula (B.3) as well as (B.2) to obtain

$$\begin{aligned} \mu_t(\mathcal{X}) &\geq \frac{\mu_0(\mathcal{X})}{N_t} \exp\left(-t \int_{\mathcal{X}} \Phi(u; y) \nu_{\mathcal{X}}(\mathrm{d}u)\right) \\ &= \frac{\mu_0(\mathcal{X})}{N_t} \exp\left(-\frac{t}{\mu_0(\mathcal{X})} \int_{\mathcal{X}} \Phi(u; y) \mu_0(\mathrm{d}u)\right) \\ &\geq \frac{\mu_0(\mathcal{X})}{N_t} \exp\left(-\frac{tM_0}{\mu_0(\mathcal{X})}\right) > 0 \end{aligned}$$

which contradicts our assumption that $\mu_t(\mathcal{X}) = 0$. □

B.2. Proof of theorem 1

Let us first prove:

Theorem 1. For any bounded $t \geq 0$, the normalisation constant N_t , defined in formula (29), as a function of t , is differentiable and its derivative $\frac{\mathrm{d}N_t}{\mathrm{d}t}$ is given by

$$\frac{1}{N_t} \frac{\mathrm{d}N_t}{\mathrm{d}t} = - \int_{\mathcal{H}} \Phi(u; y) \mu_t(\mathrm{d}u), \quad (\text{B.4})$$

where μ_t is the probability measure determined via formula (28).

Proof. For any $u \in \mathcal{H}$ and for any bounded $t \geq 0$, let us define

$$f_u(t) := \exp(-t\Phi(u; y)).$$

Then, N_t defined in formula (29) can be rewritten as

$$N_t = \int_{\mathcal{H}} f_u(t) \mu_0(\mathrm{d}u).$$

Notice that the derivative of f_u is bounded by $\Phi(u; y)$ for any $t \geq 0$, since

$$\left| \frac{df_u}{dt}(t) \right| = \left| -\Phi(u; y) \exp(-t\Phi(u; y)) \right| \leq \Phi(u; y).$$

Recall that $\Phi(\cdot; y)$ is absolutely integrable. Thus, according to the dominated convergence theorem, the derivative of N_t can be calculated by

$$\begin{aligned} \frac{dN_t}{dt} &= \frac{d}{dt} \left(\int_{\mathcal{H}} f_u(t) \mu_0(du) \right) = \int_{\mathcal{H}} \frac{df_u}{dt} \mu_0(du) \\ &= - \int_{\mathcal{H}} \Phi(u; y) \exp(-t\Phi(u; y)) \mu_0(du). \end{aligned}$$

From the proof of lemma 1 we know N_t is strictly greater than 0 for any bounded $t \geq 0$. Therefore, N_t' can be divided by N_t ,

$$\frac{1}{N_t} \frac{dN_t}{dt} = - \frac{1}{N_t} \int_{\mathcal{H}} \Phi(u; y) \exp(-t\Phi(u; y)) \mu_0(du).$$

Expression (B.4) follows from lemma 1 which establishes the equivalence between μ_t and μ_0 . \square

Expression (B.4) can be used to characterise N_t by integrating on both sides, arriving at

$$\log N_t - \log N_0 = - \int_0^t \langle \Phi \rangle_t dt,$$

where we have used the notation introduced in (30). This extends, to the infinite-dimensional setting, the characterisation of the normalising constant obtained, for example, in [36].

We now prove theorem 1:

Proof. $\mathbb{E} \{g(u_t)\}$ can be rewritten by changing the measure from μ_t to μ_0 (recall $\mu_t \ll \mu_0$),

$$\begin{aligned} \mathbb{E} \{g(u_t)\} &:= \int_{\mathcal{H}} g(u) \mu_t(du) = \int_{\mathcal{H}} g(u) \frac{d\mu_t}{d\mu_0}(u) \mu_0(du) \\ &= \frac{1}{N_t} \int_{\mathcal{H}} g(u) \exp(-t\Phi(u; y)) \mu_0(du) = \frac{1}{N_t} \int_{\mathcal{H}} f_u(t) \mu_0(du), \end{aligned}$$

where $f_u(t)$ is defined by

$$f_u(t) := g(u) \exp(-t\Phi(u; y))$$

for any $u \in \mathcal{H}$ and any bounded $t \geq 0$. By applying the chain rule, the derivative of $\mathbb{E} \{g(u_t)\}$ is given by

$$\begin{aligned} \frac{d}{dt} \mathbb{E} \{g(u_t)\} &= \frac{d}{dt} \left(\frac{1}{N_t} \int_{\mathcal{H}} f_u(t) \mu_0(du) \right) \\ &= - \frac{\frac{dN_t}{dt} \int_{\mathcal{H}} f_u(t) \mu_0(du)}{N_t^2} + \frac{\frac{d}{dt} \left(\int_{\mathcal{H}} f_u(t) \mu_0(du) \right)}{N_t}. \end{aligned} \quad (\text{B.5})$$

Notice that the derivative of f_u is bounded by $|g(u)\Phi(u; y)|$ for any $t \geq 0$, since

$$\left| \frac{df_u}{dt}(t) \right| = \left| -g(u)\Phi(u; y) \exp(-t\Phi(u; y)) \right| \leq |g(u)\Phi(u; y)|.$$

Since both g and Φ are square integrable, it follows from Cauchy–Schwarz inequality that $g \cdot \Phi$ is absolutely integrable. Thus, according to the dominated convergence theorem, we have

$$\begin{aligned} \frac{d}{dt} \left(\int_{\mathcal{H}} f_u(t) \mu_0(\mathbf{d}u) \right) &= \int_{\mathcal{H}} \frac{df_u}{dt}(t) \mu_0(\mathbf{d}u) \\ &= - \int_{\mathcal{H}} \Phi(u; y) f_u(t) \mu_0(\mathbf{d}u). \end{aligned} \quad (\text{B.6})$$

On the other hand, we note that the conditions of theorem 1 apply. Then, we substitute (B.6) and (B.4) in (B.5) to arrive at

$$\frac{d}{dt} \mathbb{E} \{g(u_t)\} = \frac{\int_{\mathcal{H}} \Phi(u; y) \mu_t(\mathbf{d}u) \int_{\mathcal{H}} f_u(t) \mu_0(\mathbf{d}u)}{N_t} - \frac{\int_{\mathcal{H}} \Phi(u; y) f_u(t) \mu_0(\mathbf{d}u)}{N_t}.$$

According to lemma 1, μ_t and μ_0 are equivalent for any bounded $t \geq 0$. Hence, we change measure in the right-hand side of the equation above using formula (28). Thus,

$$\frac{d}{dt} \mathbb{E} \{g(u_t)\} = \int_{\mathcal{H}} \Phi(u; y) \mu_t(\mathbf{d}u) \int_{\mathcal{H}} g(u) \mu_t(\mathbf{d}u) - \int_{\mathcal{H}} \Phi(u; y) g(u) \mu_t(\mathbf{d}u).$$

Since both Φ and g are square integrable, it is not difficult to see that the right-hand side in the previous expression gives the desired result in (34). \square

ORCID iDs

Marco Iglesias  <https://orcid.org/0000-0002-8952-717X>

References

- [1] Iglesias M A, Law K J H and Stuart A M 2013 Ensemble Kalman methods for inverse problems *Inverse Problems* **29** 045001
- [2] Iglesias M A 2016 A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems *Inverse Problems* **32** 025002
- [3] Chada N K, Iglesias M A, Roininen L and Stuart A M 2018 Parameterizations for ensemble Kalman inversion *Inverse Problems* **34** 055009
- [4] Chada N K, Stuart A M and Tong X T 2019 Tikhonov regularization within ensemble Kalman inversion (arXiv:1901.10382)
- [5] Chada N K and Tong X T 2019 Convergence acceleration of ensemble Kalman inversion in nonlinear settings (arXiv:1911.02424)
- [6] Evensen G 1994 Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics *J. Geophys. Res.* **99** 10143–62
- [7] Emerick A A and Reynolds A C 2012 History matching time-lapse seismic data using the ensemble Kalman filter with multiple data assimilations *Comput. Geosci.* **16** 639–59
- [8] Chen Y and Oliver D S 2012 Ensemble randomized maximum likelihood method as an iterative ensemble smoother *Math. Geosci.* **44** 1–26
- [9] Schneider T, Lan S, Stuart A and Teixeira J 2017 Earth system modeling 2.0: a blueprint for models that learn from observations and targeted high-resolution simulations *Geophys. Res. Lett.* **44** 12396–417
- [10] Xiao H, Wu J-L, Wang J-X, Sun R and Roy C J 2016 Quantifying and reducing model-form uncertainties in Reynolds-averaged Navier-Stokes simulations: a data-driven, physics-informed Bayesian approach *J. Comput. Phys.* **324** 115–36

- [11] Albani V, Ascher U M, Ascher U M, Yang X and Zubelli J P 2017 Data driven recovery of local volatility surfaces *Inverse Problems Imaging* **11** 799–823
- [12] Wang J-X, Hu X and Shadden S C 2019 Data-augmented modeling of intracranial pressure *Ann. Biomed. Eng.* **47** 714–30
- [13] Calvetti D, Nakkireddy S and Somersalo E 2020 A Bayesian filtering approach to layer stripping for electrical impedance tomography *Inverse Problems* **36** 055014
- [14] Muir J B and Tsai V C 2019 Geometric and level set tomography using ensemble Kalman inversion *Geophys. J. Int.* **220** 967–80
- [15] De Simon L, Iglesias M, Jones B and Wood C 2018 Quantifying uncertainty in thermophysical properties of walls by means of Bayesian inversion *Energy Build.* **177** 220–45
- [16] Iglesias M, Park M and Tretyakov M V 2018 Bayesian inversion in resin transfer molding *Inverse Problems* **34** 105002
- [17] Kovachki N B and Stuart A M 2019 Ensemble Kalman inversion: a derivative-free technique for machine learning tasks *Inverse Problems* **35** 095005
- [18] Schillings C and Stuart A M 2017 Analysis of the ensemble Kalman filter for inverse problems *SIAM J. Numer. Anal.* **55** 1264–90
- [19] Schillings C and Stuart A M 2018 Convergence analysis of ensemble Kalman inversion: the linear, noisy case *Appl. Anal.* **97** 107–23
- [20] Blömker D, Schillings C, Wacker P and Weissmann S 2019 Well posedness and convergence analysis of the ensemble Kalman inversion *Inverse Problems* **35** 085007
- [21] 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
- [22] Blömker D, Schillings C and Wacker P 2018 A strongly convergent numerical scheme from ensemble Kalman inversion *SIAM J. Numer. Anal.* **56** 2537–62
- [23] Engl H, Hanke M and Neubauer A 1996 *Regularization of Inverse Problems* vol 375 (Berlin: Springer)
- [24] Kaipio J and Somersalo E 2005 *Statistical and Computational Inverse Problems* (New York: Springer)
- [25] Kalltenbacher B, Neubauer A and Scherzer O 2008 *Iterative Regularization Methods for Nonlinear Ill-Posed Problems* 1st edn (*Radon Series on Computational and Applied Mathematics*) (Berlin: de Gruyter & Co)
- [26] Marquardt D W 1963 An algorithm for least-squares estimation of nonlinear parameters *J. Soc. Ind. Appl. Math.* **11** 431–41
- [27] Hanke M 1997 A regularizing Levenberg - marquardt scheme, with applications to inverse groundwater filtration problems *Inverse Problems* **13** 79–95
- [28] Ruchi S, Dubinkina S and Iglesias M 2019 Transform-based particle filtering for elliptic bayesian inverse problems *Inverse Problems* **35** 115005
- [29] Dubinkina S and Ruchi S 2020 Comparison of regularized ensemble Kalman filter and tempered ensemble transform particle filter for an elliptic inverse problem with uncertain boundary conditions *Comput. Geosci.* **24** 149–60
- [30] Yang F-l and Yan L 2020 A non-intrusive reduced basis EKI for time fractional diffusion inverse problems *Acta Math. Appl. Sin. Engl. Ser.* **36** 183–202
- [31] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559
- [32] Beskos A, Jasra A, Muzaffer E A and Stuart A M 2015 Sequential Monte Carlo methods for Bayesian elliptic inverse problems *Stat. Comput.* **25** 727–37
- [33] Kantas N, Beskos A and Jasra A 2014 Sequential Monte Carlo methods for high-dimensional inverse problems: a case study for the Navier–Stokes equations *SIAM/ASA J. Uncertain. Quantification* **2** 464–89
- [34] Latz J, Papaioannou I and Ullmann E 2018 Multilevel sequential2 Monte Carlo for bayesian inverse problems *J. Comput. Phys.* **368** 154–78
- [35] Tawn N G and Roberts G O 2019 Accelerating parallel tempering: quantile tempering algorithm (quanta) *Adv. Appl. Probab.* **51** 802–34
- [36] Behrens G, Friel N and Hurn M 2012 Tuning tempered transitions *Stat. Comput.* **22** 65–78
- [37] Lefebvre G, Steele R and Vandal A C 2010 A path sampling identity for computing the Kullback–Leibler and J divergences *Comput. Stat. Data Anal.* **54** 1719–31
- [38] de Wiljes G, Reich S and Stannat W 2018 Long-time stability and accuracy of the ensemble Kalman–Bucy filter for fully observed processes and small measurement noise *SIAM J. Appl. Dyn. Syst.* **17** 1152–81

- [39] Emerick A A and Reynolds A C 2013 Ensemble smoother with multiple data assimilation *Comput. Geosci.* **55** 3–15
- [40] Evensen G 2018 Analysis of iterative ensemble smoothers for solving inverse problems *Comput. Geosci.* **22** 885–908
- [41] Iglesias M A 2014 Iterative regularization for ensemble data assimilation in reservoir models *Comput. Geosci.* **19** 177–212
- [42] Le D H, Emerick A A and Reynolds A C 2016 An adaptive ensemble smoother with multiple data assimilation for assisted history matching *SPE J.* **21** 2195–207
- [43] Kullback S and Leibler R A 1951 On information and sufficiency *Ann. Math. Stat.* **22** 79–86
- [44] Mühlig H and Morozov V A 1984 *Methods for Solving Incorrectly Posed Problems* (Berlin: Springer)
- [45] Somersalo E, Cheney M and Isaacson D 1992 Existence and uniqueness for electrode models for electric current computed tomography *SIAM J. Appl. Math.* **52** 1023–40
- [46] Borcea L 2002 Electrical impedance tomography *Inverse Problems* **18** R99
- [47] Lindgren F, Rue H and Lindström J 2011 An explicit link between Gaussian fields and Gaussian markov random fields: the stochastic partial differential equation approach *J. R. Stat. Soc. B* **73** 423–98
- [48] Dunlop M M, Iglesias M A and Stuart A M 2017 Hierarchical Bayesian level set inversion *Stat. Comput.* **27** 1555–84
- [49] Roininen L, Huttunen J M J and Lasanen S 2014 Whittle–Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography *Inverse Problems Imaging* **8** 561–86
- [50] Santosa F 1996 A level-set approach for inverse problems involving obstacles fadil santosa *ESAIM: Control, Optim. Calc. Var.* **1** 17–33
- [51] Iglesias M, Lu Y and Stuart A 2016 A Bayesian level set method for geometric inverse problems *Interfaces Free Bound.* **18** 181
- [52] Adler A and Lionheart W R B 2006 Uses and abuses of eiders: an extensible software base for eit *Physiol. Meas.* **27** S25
- [53] Mandel J, Cobb L and Beezley J D 2011 On the convergence of the ensemble Kalman filter *Appl. Math.* **56** 533–41
- [54] Kwiatkowski E and Mandel J 2015 Convergence of the square root ensemble kalman filter in the large ensemble limit *SIAM/ASA J. Uncertain. Quantification* **3** 1–17
- [55] Dashti M and Stuart A M 2017 *The Bayesian Approach to Inverse Problems* (Berlin: Springer) pp 311–428