

Iterative regularization for ensemble data assimilation in reservoir models

Marco A. Iglesias

Received: date / Accepted: date

Abstract We propose the application of iterative regularization for the development of ensemble methods for solving Bayesian inverse problems. In concrete, we construct (i) a variational iterative regularizing ensemble Levenberg-Marquardt method (IR-enLM) and (ii) a derivative-free iterative ensemble Kalman smoother (IR-ES). The aim of these methods is to provide a robust ensemble approximation of the Bayesian posterior. The proposed methods are based on fundamental ideas from iterative regularization methods that have been widely used for the solution of deterministic inverse problems [25]. In this work we are interested in the application of the proposed ensemble methods for the solution of Bayesian inverse problems that arise in reservoir modeling applications. The proposed ensemble methods use key aspects of the regularizing Levenberg-Marquardt scheme developed by Hanke [20] and that we recently applied for history matching in [22]. Unlike most existing methods where the stopping criteria and regularization parameters are typically selected heuristically, in the proposed ensemble methods the discrepancy principle is applied for (i) the selection of the regularization parameters and (ii) the early termination of the scheme. The discrepancy principle is key for the theory of iterative regularization and the purpose of the present work is to apply this principle for the development of ensemble methods defined as iterative updates of solutions to linear ill-posed inverse problems.

The regularizing and convergence properties of iterative regularization methods for deterministic inverse problems have long been established. However, the approximation properties of the proposed ensemble methods in the context of Bayesian inverse problems is an open problem. In the case where the forward operator is linear and the prior is Gaussian, we show that the tunable parameters of the proposed IR-enLM and IR-ES can be chosen so that the resulting schemes coincide with the standard randomized maximum likelihood (RML) and the ensemble smoother (ES), respectively. Therefore, the proposed methods sample from the posterior in the linear-

M. A. Iglesias
University of Nottingham
E-mail: marco.iglesias@nottingham.ac.uk

Gaussian case. Similar to RML and ES methods, in the nonlinear case, one may not conclude that the proposed methods produce samples from the posterior. The present work provides a numerical investigation of the performance of the proposed ensemble methods at capturing the posterior. In particular, we aim at understanding the role of the tunable parameters that arise from the application of iterative regularization techniques. The numerical framework for our investigations consists of using a state-of-the-art MCMC method for resolving the Bayesian posterior from synthetic experiments. The resolved posterior via MCMC then provides a gold standard against to which compare the proposed IR-enLM and IR-ES. Our numerical experiments show clear indication that the regularizing properties of the regularization methods applied for the computation of each ensemble have significant impact of the approximation properties of the proposed ensemble methods at capturing the Bayesian posterior. Furthermore, we provide a comparison of the proposed regularizing methods with respect to some unregularized methods that have been typically used in the literature. Our numerical experiments showcase the advantage of using iterative regularization for obtaining more robust and stable approximation of the posterior than unregularized methods.

Keywords First keyword · Second keyword · More

1 Introduction

We use key elements of *iterative regularization* techniques to develop robust ensemble methods for solving Bayesian inverse problems that arise in subsurface flow. A novel aspect of these ensemble methods is that the *discrepancy principle* is used for (i) the selection of the regularization parameter that stabilizes the update of each ensemble member and (ii) the stopping criteria that avoids data overfitting. These strategies have been used for theoretically establishing the convergence and regularizing properties of some well-known iterative regularization methods aimed at solving deterministic nonlinear ill-posed inverse problems. In this work we extend these strategies for the solution of Bayesian inverse problems where the aim is to capture the Bayesian posterior. In particular, we use the aforementioned elements from iterative regularization to improve the robustness of RML-based and Kalman-based methodologies. The ensemble methods proposed in the present manuscript constitute a proof-of-concept which demonstrates the potential advantages of incorporating fundamental components of iterative regularization such as the aforementioned discrepancy principle. This principle, and more generally, iterative regularization, can be potentially combined with existing frameworks for ensemble data assimilation. It is our conjecture that when standard methods are combined with systematic approaches that take into account the mathematical structure of the underlying inverse problem, further increase in robustness, stability and accuracy for capturing the Bayesian posterior may be obtained. While the present work is focused in large-scale inverse problems that arise in reservoir flow modeling, the proposed methods can be applied for the solution of generic PDE-constrained inverse problems.

1.1 Background of the proposed methods

In a recent publication [22] we studied the application of iterative regularization methods for computing *inverse estimates* of unknown geologic properties by means of history matching. More concretely, we developed an application of the regularizing Levenberg-Marquardt (LM) scheme [20] for solving history matching problems posed as the minimization of

$$\Phi(u) \equiv \|\Gamma^{-1/2}(y - G(u))\|_Y^2 \quad (1)$$

where u is the unknown geologic property, y denotes production data, Γ is the measurements error covariance. $G : X \rightarrow Y$ is the forward operator that arises from the reservoir model; G maps the parameter space X (of admissible geologic properties) to the observation space Y . In typical reservoir modeling applications, geologic properties need to be defined on thousands or even millions of gridblocks. Due to the resulting large size of the space X , the computation of a minimizer of (1) is unstable (ill-posed) in the sense that an arbitrarily small data misfit (1) may not necessarily correspond to an estimates u that is close to the optimal. This ill-posedness, that arises from the mathematical structure of G , requires regularization. In the approach proposed in [22], stable approximations to the minimizer of (1) were generated with an application of the regularizing LM scheme. In that application, the regularizing LM scheme was initialized with the prior mean \bar{u} . Additionally, it used the prior error covariance C for the regularization that was built into the iterative scheme.

The regularizing LM scheme uses the *discrepancy principle* [17] to impose a limit on how close the model predictions must fit the data. More precisely, for some $\tau > 1$, the regularizing LM scheme terminates whenever at the m iteration we find that

$$\|\Gamma^{-1/2}(y - G(u_{m+1}))\|_Y \leq \tau\eta \leq \|\Gamma^{-1/2}(y - G(u_m))\|_Y \quad (2)$$

where u_m is the LM estimate at the m -th iteration and the *noise level* η is defined by

$$\eta \equiv \|\Gamma^{-1/2}(y - G(u^\dagger))\|_Y \quad (3)$$

where u^\dagger denotes the truth properties. In other words, the noise level η is the weighted data misfit for the truth u^\dagger . The discrepancy principle is the key to ensure the regularizing properties of iterative regularization methods [25] and in particular of the regularizing LM scheme that we applied in [22] for computing inverse estimates of geologic properties by means of history matching. Although the truth u^\dagger is unknown, it is reasonable to assume that knowledge of the noise level η is available (see discussion in [22, Section 3]).

While [22] provides numerical evidence of the robustness and accuracy of the regularizing LM scheme for recovering u^\dagger , a minimizer of (4) may not be unique. In other words, different approximations to a minimizer of (4) may be found. For example, initializing the regularizing LM scheme with a different (from the prior mean \bar{u}) initial guess may lead to different approximation to the minimizer of (1). Similarly, a small perturbation of the production data y may result in a significantly different, yet stable, estimate of the geologic properties. Both uncertainties in prior knowledge and observational noise give rise to uncertainty in the computation of inverse estimates

of the geologic properties that we often perform with least-squares. The quantification of this uncertainty is essential for the optimal management and decision-making involved in reservoir applications. In this paper we extend the approach of [22] and apply fundamental ideas from iterative regularization techniques to develop ensemble methods that, within the Bayesian framework [32], quantify uncertainty of these inverse estimates. More precisely, we provide numerical evidence that this methods produce useful information of the posterior that arises from the Bayesian inverse problem of conditioning the geologic properties to data from the reservoir dynamics.

1.2 The Bayesian posterior and its ensemble approximation

The most standard approach to quantify inverse estimates of the geologic properties is the Bayesian framework. The aim is to merge uncertainties, both in prior knowledge and observational noise, with the mathematical model that describes the reservoir flow. The prior uncertainty in geologic properties is incorporated in terms of a *prior* probability distribution $\mathbb{P}(u)$. The data y and the geologic properties u are related via the forward operator G by

$$y = G(u) + \xi \quad (4)$$

where ξ is a vector of random noise. The *likelihood* of the measured data y given a particular instance of u is denoted by $\mathbb{P}(y|u)$. In the Bayesian framework, the uncertainty of the inverse estimates of u given y is quantified with the conditional *posterior* probability $\mathbb{P}(u|y)$, which from Bayes' rule, is given by

$$\frac{\mathbb{P}(u|y)}{\mathbb{P}(u)} \propto \mathbb{P}(y|u). \quad (5)$$

Since the forward operator G that arises from typical reservoir models is highly non-linear and/or the prior distribution of geologic properties $\mathbb{P}(u)$ is not Gaussian, the posterior $\mathbb{P}(u|y)$ cannot be described with a few parameters. In general, the characterization of the posterior may be conducted by means of sampling. Unfortunately, this approach often requires millions of forward model evaluations which is only feasible for small problems where a relatively coarse grid is used to discretize the geologic properties [13, 29]. The full characterization of the posterior by means of sampling is therefore impractical. However, sampling can be used for the purpose of benchmarking and assessing methods that can be actually used in practice [13, 29, 23]. That is the case of *ensemble methods* which have become the paradigm for capturing features of the posterior $\mathbb{P}(u|y)$ within a reasonable amount of computing resources and time. The general objective of these methods is to provide fast, accurate and robust estimates of the posterior (and/or its mapping under the forward operator) by means of an ensemble $\{u^{(j)}\}_{j=1}^{N_e}$ of N_e inverse estimates of the geologic properties of the reservoir. From this ensemble we may compute statistical information of the unknown u as well as quantities of interest related to the model predictions $G(u)$. Furthermore, an ensemble of reservoir properties that accurately captures the posterior can be used to quantify the uncertainty in future production scenarios (e.g. drilling new wells) and/or additional subsurface applications (e.g. injection of CO₂ in

a depleted reservoir). Moreover, efficient ensemble methods may be combined with optimal control modules that closes the loop of reservoir management and monitoring [28]. It is clear that ensemble methods that provide a robust and accurate approximation of the Bayesian posterior are fundamental for geologic uncertainty quantification in subsurface modeling applications.

Fully rigorous Bayesian sampling is very expensive and often impractical; to overcome this difficulty there exist two main frameworks to generate an ensemble approximation of the posterior which, although not rigorously justifiable in general, are widely used within the reservoir modeling community: (i) the randomized maximum likelihood (RML) method and (ii) Kalman-based techniques. The RML method consist of an ensemble of realizations computed by minimizing a randomized version of the cost functional that defines the maximum a posteriori (MAP) estimate [30]. In other words, each ensemble member of RML involves an optimization problem that may be addressed with similar methods to the ones used for computing the MAP [33, 13, 29, 23]. When a relatively small number of measurements is available, the computation of the derivative of the forward operator (or sensitivity) may be feasible provided the adjoint of the derivative can be computed. In this case, standard Gauss-Newton or Levenberg-Marquardt methods can be potentially applied for the minimization of each ensemble member of RML [23, 27]. However, in some cases, large amount of data need to be assimilated and the computation of the sensitivity matrix becomes computationally prohibitive. This limitation of Newton-type methods can be addressed by the application of quasi-Newton methods where only the action of the sensitivity is needed. In [16, 34] and [13], for example, the LBFGS was applied for the computation of the MAP and the ensemble of RML, respectively. In some cases, a SVD decomposition of the sensitivity matrix of the MAP may be used as a reparameterization of the search space of the iterative scheme (e.g. LM or LBFGS) used for the optimization. The effect of this reparameterization is twofold. On the one hand, a truncated SVD parameterization reduces the parameter space and could potentially alleviate the ill-posedness of the inverse problem. On the other hand, the aforementioned SVD decomposition may be computed by means of a Lanczos algorithm with the aid of an adjoint method, thus avoiding the explicit computation of the sensitivity matrix [33, 34].

Ensemble Kalman-based methods use the Kalman formula for generating an ensemble of inverse estimates of the geologic properties [21]. Each ensemble member may be updated in iterative or noniterative and in sequential (filtering) or all-at-once (smoothing) basis [1]. As stated before, variational methods often require an adjoint code for the efficient computation of the sensitivity matrix (or its multiplication by a vector). In contrast, ensemble Kalman-based methods are derivative-free approaches where the ensemble updates are based on simple computations involving covariances and crosscovariances computed directly from ensemble usually initialized with samples from the prior. Thus, ensemble Kalman methods can be used in a black-box fashion. However, it is widely known the lack of stability (or roughness) of ensemble Kalman methods when the ensemble size is small with respect to the number of parameters or measurements. Therefore, in the context of reservoir applications, attention has been recently given to the regularization of ensemble Kalman based methods by means of localization [10, 4], inflation [15]. For a recent review and comparison

of ensemble methods we refer the reader to the work of [13]. It is also important to mention the work of [21] that uses Kalman updates within an iterative ensemble scheme to solve PDE-constrained parameter identification problems. By a subspace property proven in [21], the space of admissible solutions is defined as the subspace generated by the prior ensemble. Thus, the regularization results from defining the search space on a compact subspace. However, the numerical results of [21] suggest that an early termination of the scheme is additionally required to further stabilize the iterative scheme.

Methods that combine ideas from both RML and Kalman-based techniques have been developed in [5, 6, 26, 18]. The essence of these methods is to preserve the properties of RML to capture the posterior, while using a derivative-free ensemble-based optimization method. For instance, in the LM iteration step for the minimization of each ensemble member in RML, the LM-enRML approach of [6] replaces the sensitivity matrix with an adjoint-free approximation computed from the ensemble. Similarly, LM-enRML uses the prior ensemble to approximate the prior error covariance that appears in the LM scheme. The results reported in [6] indicate that LM-enRML provides a good approximation of the posterior for a small problem where the posterior can be computed analytically. However, the approximation properties of LM-enRML for capturing the posterior for large problems is still an open problem.

1.3 Contribution of this work

In this paper we contribute to both mainstreams of ensemble methods previously described. On the one hand, we extend our implementation of the regularizing LM of [22] and develop a variational Iterative Regularizing ensemble LM method (IR-enLM). The aim of this method is to generate a stabilized ensemble of inverse estimates computed from a randomized least-squares minimization. On the other hand, we apply fundamental ideas from iterative regularization and construct an Iterative Regularizing ensemble (Kalman) Smoother (IR-ES). For these proposed ensemble methods the computation of each ensemble member can be posed as the solution of a deterministic nonlinear ill-posed inverse problem that we address by means of iterative regularization methods. The resulting ensemble, and the main goal of these methods is to provide a robust approximation of the Bayesian posterior. However, the regularization of each ensemble member and the resulting approximation of the Bayesian posterior are two independent aspects that combined yield the proposed techniques. The main objective of the present work is to explore the connection between these two aspects and determine the capabilities of these ensemble methods for capturing the posterior. This objective requires understanding the role of the tunable parameters that arise from the application of iterative regularization methods. While such role has been established both theoretically and numerically in the context of deterministic inverse problems, here we provide an extensive numerical study to understand these tunable parameters in the context of Bayesian inverse problems that arise in reservoir modeling applications. Our aim is to obtain a range of tunable parameters that corresponds to accurate and computationally feasible approximations of the posterior.

The main focus of the proposed methods is to address two sources of ill-posedness that could potentially affect the accuracy and robustness of ensemble methods. The first source is due to the high dimension that results from the discretization of the parameter space X as well as the intrinsic properties of G . As previously discussed, these properties give rise to the lack of stability in the computation of the inverse problem that we aim at solving, for each ensemble member, in standard ensemble methods. The second source of ill-posedness arises in ensemble Kalman-type methods with a small ensemble size relative to the number of observations/parameters. In ensemble methods formulated in the observation space, for example, the covariance of the model predictions is rank-deficient and, even though the error covariance matrix may add stabilization, the resulting matrix inversion in the Kalman update may be ill-conditioned hence detrimental to the stable computations of the ensemble. As we stated above, heuristic fixes such as covariance inflation and localization have been used to address such ill-posedness. Here we propose an alternative approach based on ideas from iterative regularization methods.

While both aforementioned sources of ill-posedness are different in essence, computationally, they both result in a large condition number of the matrix that is inverted when the Kalman formula is applied at each iteration of the ensemble method. Therefore, even though iterative regularization is designed to alleviate the first source of ill-posedness, in the present work we show that iterative regularization may also have a significant effect in the regularization/stabilization of ensemble Kalman methods with small ensemble size. Addressing ill-posedness due to small ensemble sizes is a relevant contribution of the present work; the high computational cost associated with the evaluation of $G(u)$ imposes severe limitation in the computational cost, which is, in turn, proportional to the ensemble size. Thus, having a small ensemble relative to the observations/parameters is inherent to ensemble-based inverse modeling in subsurface applications.

Although the proposed methods are based on iterative regularization for which a convergence theory is available in the context of deterministic inverse problems [25] the corresponding numerical analysis of the approximation for the solution of Bayesian inverse problems is still an open problem beyond the scope of the present manuscript. In the trivial case where the forward operator is linear and the prior is Gaussian, we provide specific conditions for which the proposed methods provide the proper sampling of the posterior. For the nonlinear case, we use the framework of [23] to assess the performance of the proposed method for recovering the mean and variance of the posterior. In particular, we use a state-of-the-art MCMC method for functions [8] that enable us to sample posteriors that result from synthetic data from moderate size experiments.

We reiterate that a key objective of the present work is to develop and promote the use of ideas from iterative regularization for the design of ensemble approximation of the posterior that arises from reservoir modeling applications. Thus, the proposed IR-enLM and IR-ES are a proof-of-concept of iterative regularization ideas for ensemble approximations to the Bayesian posterior. We do not suggest to *replace* existing methods but rather to demonstrate a methodology to *enhance* those methods by the careful choice of regularization parameters and stopping criteria; we show the potential advantages of taking into account methodologies that have been rigor-

ously constructed for the solution of deterministic inverse problems. In particular, the proposed IR-enLM and IR-ES are based on the discrepancy principle for the early termination and selection of parameters. Further applications of iterative regularization can be developed in combination with state-of-the-art frameworks such as the SVD parameterization of [33] and the LM approaches of [6]. Those applications, however, are beyond the scope of the present work.

The paper is organized as follows. Relevant aspects of variational iterative regularization are discussed in Section 2. In particular, a brief description of RML is presented in subsection 2.1. The proposed IR-enLM algorithm is introduced in subsection 2.2. The linear-Gaussian case is addressed in subsection 2.5. Differences between IR-enLM and a standard implementation of RML is discussed in 2.6. In Section 3 we study Kalman-based ensemble methods. More concretely, 3.1 we briefly introduce the standard ES. A regularized version of this ES algorithm is introduced in 3.2. The proposed IR-ES is then presented in subsection 3.3. Numerical examples of the proposed methods are presented in Section 4. In subsection 4.1 we describe the generation of synthetic data for the experiments. The resulting posteriors are resolved with the methodology described in subsection 4.2. The numerical experiments of the implementation of IR-enLM and IR-ES are displayed in subsections 4.3 and 4.4, respectively. A comparison between IR-ES and ES-MDA is presented in subsection 4.6 [12]. Comparison with unregularized methods is presented in 4.6. Final remarks are provided in Section 5. The forward models under consideration are described in the Appendix.

2 Variational iterative regularization

In this section we define a variational ensemble method that quantifies uncertainty of inverse estimates of geologic properties. The proposed method is analogous to the randomized maximum likelihood (RML) method of [30] which we review in the following subsection. For simplicity, throughout the rest of this document we consider that the prior distribution of the geologic properties is the Gaussian measure $\mathbb{P}(u) = N(\bar{u}, C)$ where as before, C is the prior covariance and \bar{u} is the prior mean. In addition, we consider that the observational noise ξ in (4) is centered Gaussian with covariance Γ . Under these assumptions, (5) becomes

$$\mathbb{P}(u|y) \propto \exp\{-J(u, y)\} \quad (6)$$

where

$$J(u, y) \equiv \frac{1}{2} \|\Gamma^{-1/2}(y - G(u))\|_Y^2 + \frac{1}{2} \|C^{-1/2}(u - \bar{u})\|_X^2 \quad (7)$$

2.1 The randomized maximum likelihood (RML) method

From (6)-(7) is easy to see that the posterior distribution $\mathbb{P}(u|y)$ is maximized for $u \in X$ that minimizes

$$J(u, y) \equiv \frac{1}{2} \|\Gamma^{-1/2}(y - G(u))\|_Y^2 + \frac{1}{2} \|C^{-1/2}(u - \bar{u})\|_X^2 \quad (8)$$

Under certain conditions, such a u exists [32] and is often referred as the maximum a posteriori (MAP) estimate. Since the operator G is nonlinear for the case of interest for the present application, the MAP estimator may not be unique and so multiple modes of the posterior may exist. Except in the linear-Gaussian case, the computation of one minimizer of (7) does not provide sufficient information to fully characterize the posterior. In the context of Bayesian subsurface inverse problems [30], RML has been proposed to approximate the posterior with an ensemble $\{u_{RML}^{(j)}\}_{j=1}^{N_e}$ of N_e realizations defined by

$$u_{RML}^{(j)} = \operatorname{argmin}_u J_{RML}^{(j)}(u) \quad (9)$$

where

$$J_{RML}^{(j)} = \frac{1}{2} \|\Gamma^{-1/2}(y^{(j)} - G(u))\|_Y^2 + \frac{1}{2} \|C^{-1/2}(u - u^{(j)})\|_X^2 \quad (10)$$

and

$$u^{(j)} \sim N(\bar{u}, C), \quad y^{(j)} = y + \xi^{(j)}, \quad \xi^{(j)} \sim N(0, \Gamma). \quad (11)$$

In the case where $G(u)$ is a linear operator (and since we assume that $\mathbb{P}(u)$ is Gaussian) the ensemble obtained with (9) are samples of the posterior distribution [30]. In the nonlinear case, however, the analysis of the approximation properties of RML is still an open problem. Nevertheless, multiple methods for computing or approximating the minimizer of (9) have been proposed in recent years [33, 6, 26]. As stated in the preceding section, the main focus of these RML-based ensemble methods has been the numerical efficiency and computational feasibility that are essential in practice. While some of these methods may be more efficient and suitable than others, they are all expected to provide an approximation of the posterior similar to the one provided by the ensemble defined in (9). Even though most RML-based methods have been assessed in terms of their capabilities for data-fitting, some recent publications have established the numerical efficacy of some RML implementations for capturing the Bayesian posterior. In concrete, [13, 23, 29] have used MCMC to fully characterize the posterior which was, in turn, used to assess the performance of ensemble methods including some implementations of RML. For problems where the dimension of the parameter space is very small, it has been found that some of those variational implementations exhibits the best performance at approximating some moments of the posterior [13, 29]. However, for larger problems, the work in [23] reported an implementation of RML that overestimated the variance of the posteriors from some test problems which were, in turn, fully resolved with a state-of-the-art MCMC method for functions. This suboptimal behavior of RML can be attributed to (i) the larger size of experiments considered in [23] and (ii) the optimization method used for the minimization of (7). More specifically, [23] uses an ‘‘unregularized’’ LM proposed in [2] and used in [27, 30] for minimizing (7).

In the general nonlinear case, it is quite clear that the RML ensemble approximation of the Bayesian posterior will strongly depend on the numerical technique that employed for the minimization of (9). On the other hand, this minimization that we perform to compute each ensemble member is nothing but a regularized version of a

deterministic nonlinear ill-posed inverse problems. Therefore, it should come as no surprise that (9) may be subject to the numerical instabilities discussed in [22] for the computation of the MAP estimate. More concretely, for each ensemble member, the stable computation of (9) requires that the “prior term” $\frac{1}{2} \|C^{-1/2}(u - u^{(j)})\|_X^2$ provides enough regularization to the ill-posed problem of minimizing $\frac{1}{2} \|\Gamma^{-1/2}(y^{(j)} - G(u))\|_Y^2$. The stabilization of (9) thus relies on the proper selection of the operators Γ and C . These, however, are typically chosen according to measurement and prior geological information, respectively. Therefore, for some choices of C and Γ , numerical instabilities in the computation of (9) could potentially arise unless some additional form of regularization is applied. In [22] we showed synthetic experiments where the computation of the MAP estimator (i.e. minimizer of (7)) with an unregularized LM method did not ensure the proper regularization of the history matching problem. That is, even though data were successfully fitted, the corresponding estimates of the geologic properties were significantly different from the truth. While recovering the truth was the main focus of [22], in the present work we are interested in the solution to the Bayesian inverse problem, i.e. the posterior distribution. In Section 4 of the present document we therefore show that, when the unregularized LM method of [2, 30] is implemented to minimize (9) in RML, suboptimal approximations of the Bayesian posterior may be obtained. This motivates our application of iterative regularization for the development of a robust and stable variational ensemble approximations of the posterior. It is worth mentioning that the aforementioned instabilities that may arise from the minimization of the MAP (or alternatively RML) in subsurface flow modeling application [30, 27] have been also observed and addressed in some implementations of incremental variational data assimilation methods for numerical weather prediction applications [19].

It is important to remark that some recent implementations of RML use a truncated parameterization of the SVD of the sensitivity matrix [33] within a LM algorithm. This non-standard LM approach has been developed for computational efficiency because it avoids the explicit computations of the sensitivity that is needed in the LM method of [2]. However, the truncation of the aforementioned SVD parameterization is an additional form of regularization [24]. This truncation has the potential disadvantage of removing eigenvalues of the sensitivity matrix that may be relevant for the proper characterization of petrophysical properties in complex geologies. It is also worth mentioning the recent work of [6] where an RML-based approach was used in conjunction with an LM method that, as indicated in Section 1 uses ensemble approximations of the sensitivity and prior error covariance. While the approach of [6] uses the standard stopping criteria and regularizing parameters, it also incorporates truncated SVD of the variables of interest, thereby inducing a regularization. Moreover, [6] reported the use localization as an additional form of regularization. In the following section we present an approach where neither truncation of the spectrum nor localization is used. Instead, basic ideas from iterative regularization are applied to stabilize the computations of the ensemble. Nonetheless, as indicated earlier, the application from iterative regularization techniques could be potentially merged with existing methods that have been successfully used in practice.

2.2 An Iterative Regularizing ensemble LM (IR-enLM) method

We propose a variational ensemble method that aims at providing stable and robust ensemble approximations of the Bayesian posterior. Each member $u_{IR}^{(j)}$ of the ensemble is a stable approximation of a minimizer of

$$\Phi^{(j)}(u) \equiv \|\Gamma^{-1/2}(y^{(j)} - G(u^{(j)}))\|_Y, \quad j \in \{1, \dots, N_e\} \quad (12)$$

computed with the regularizing LM scheme, initialized with a random sample from the prior, i.e. $u_0^{(j)} \equiv u^{(j)} \sim N(\bar{u}, C)$. The data $y^{(j)}$ in (12) is a perturbation of the original data y as defined in (11). Note that (12) excludes the prior term of (9). For the sake of clarity, in the subsequent lines we briefly outline key aspects of the regularizing LM scheme applied to the computation of stable minimizers of (12). For full details on the implementation of this method for history matching, the reader is referred to [22]. The theory of the regularizing properties and the convergence of the regularizing LM scheme are found in [20]. This includes the conditions on the forward model G for which the theory ensure convergence to stable solutions (see also [22]).

The regularizing LM scheme applied to each ensemble member $j \in \{1, \dots, N_e\}$ at each iteration level m involves the computation of $u_{m+1}^{(j)} = u_m^{(j)} + \Delta u_m^{(j)}$, where the increment $\Delta u_m^{(j)}$ is defined as the minimizer of

$$J_{LM}^{(j)}(w) = \frac{1}{2} \|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}) - DG(u_m^{(j)})w)\|_Y^2 + \frac{1}{2} \alpha_m^{(j)} \|C^{-1/2}w\|_X^2 \quad (13)$$

which from standard arguments can be expressed as

$$\Delta u_m^{(j)}(\alpha_m^{(j)}) = CDG^*(u_m^{(j)})[DG(u_m^{(j)})CDG^*(u_m^{(j)}) + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})] \quad (14)$$

Let us apply $DG(u_m^{(j)})$ from both sides of the previous expression to find

$$DG(u_m^{(j)})\Delta u_m^{(j)}(\alpha_m^{(j)}) = H[H + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})]$$

where we have defined the operator $H = DG(u_m^{(j)})CDG^*(u_m^{(j)})$. From the previous expression it follows that

$$\begin{aligned} & y^{(j)} - G(u_m^{(j)}) - DG(u_m^{(j)})\Delta u_m^{(j)}(\alpha_m^{(j)}) \\ &= y^{(j)} - G(u_m^{(j)}) - H[H + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})] \\ &= (I - H[H + \alpha_m^{(j)}\Gamma]^{-1})[y^{(j)} - G(u_m^{(j)})] \\ &= (H + \alpha_m^{(j)}\Gamma - H)[H + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})] \\ &= \alpha_m^{(j)}\Gamma[H + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})] \end{aligned} \quad (15)$$

where I is the identity matrix in Y . Thus,

$$\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}) - DG(u_m^{(j)})\Delta u_m^{(j)}(\alpha_m^{(j)})) = \alpha_m^{(j)}\Gamma^{1/2}[H + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})] \quad (16)$$

The regularizing LM scheme of Hanke has two main components that ensure the regularization of the minimization of (12): (i) the selection of the regularization parameter $\alpha_m^{(j)}$ and (ii) the stopping criteria. These are both based on the discrepancy principle. In concrete, according to the theory of Hanke [20], $\alpha_m^{(j)}$ must satisfy

$$\|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}) - DG(u_m^{(j)})\Delta u_m^{(j)}(\alpha_m^{(j)}))\|_Y^2 \geq \rho^2 \|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}))\|_Y^2 \quad (17)$$

for some $\rho \in (0, 1)$. From (16) and recalling that $H = DG(u_m^{(j)})CDG^*(u_m^{(j)})$, the previous expression can be written as

$$\begin{aligned} \alpha_m^{(j)} \|\Gamma^{1/2}[DG(u_m^{(j)})CDG^*(u_m^{(j)}) + \alpha_m^{(j)}\Gamma]^{-1}[y^{(j)} - G(u_m^{(j)})]\|_Y \\ \geq \rho \|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}))\|_Y \end{aligned} \quad (18)$$

The existence of such $\alpha_m^{(j)}$ has been shown in [25] and its actual computation can be carried out with a simple iterative scheme [22] (see also Algorithm 1 below). The selection of $\alpha_m^{(j)}$ based on (18) ensures the regularization of the minimizer of (13). However, to fully stabilize the computation of a minimizer of (12), the iterative regularizing LM scheme is terminated according to the following discrepancy principle

$$\|\Gamma^{-1/2}(y^{(j)} - G(u_{k+1}^{(j)}))\|_Y \leq \tau \eta^{(j)} \leq \|\Gamma^{-1/2}(y^{(j)} - G(u_k^{(j)}))\|_Y \quad (19)$$

which is the application of (2) for minimizing (12). In the previous expression $\eta^{(j)}$ is the noise level corresponding to the data set $y^{(j)}$. In other words,

$$\eta^{(j)} \equiv \|\Gamma^{-1/2}(y^{(j)} - G(u^\dagger))\|_Y \quad (20)$$

Since we do not have access to the truth, (and therefore to $G(u^\dagger)$) and estimate of (20) is required. Let us recall from our definition (11) that $y^{(j)} = y + \xi^{(j)}$ and so (20) becomes

$$\begin{aligned} \eta^{(j)} \equiv \|\Gamma^{-1/2}(y^{(j)} - G(u^\dagger))\|_Y &\leq \|\Gamma^{-1/2}(y - G(u^\dagger))\|_Y + \|\Gamma^{-1/2}\xi^{(j)}\|_Y \\ &= \eta + \|\Gamma^{-1/2}\xi^{(j)}\|_Y \end{aligned} \quad (21)$$

where we have used the definition for the noise level (3). On the other hand, from the definition of the perturbed observations (11) we have

$$\frac{1}{N_e} \sum_{j=1}^{N_e} \Gamma^{-1/2}(y^{(j)} - G(u^\dagger)) = \frac{1}{N_e} \sum_{j=1}^{N_e} \Gamma^{-1/2}(y + \xi^{(j)} - G(u^\dagger)) = \Gamma^{-1/2}(y - G(u^\dagger))$$

where we have assumed that the finite ensemble of data perturbations has mean zero. From the previous expression it follows that

$$\eta \equiv \|\Gamma^{-1/2}(y - G(u^\dagger))\|_Y \leq \frac{1}{N_e} \sum_{j=1}^{N_e} \|\Gamma^{-1/2}(y^{(j)} - G(u^\dagger))\|_Y \quad (22)$$

From (22) and (21) we therefore find

$$\eta \leq \frac{1}{N_e} \sum_{j=1}^{N_e} \eta^{(j)} \leq \eta + \frac{1}{N_e} \sum_{j=1}^{N_e} \|\Gamma^{-1/2} \xi^{(j)}\|_Y \quad (23)$$

Note that an estimate of $\eta^{(j)}$ in the interval

$$\left[\eta, \eta + \|\Gamma^{-1/2} \xi^{(j)}\|_Y \right] \quad (24)$$

will be consistent with (23). For the present work we propose the midpoint of (24) as an estimate of $\eta^{(j)}$. In other words,

$$\eta^{(j)} \equiv \eta + \frac{1}{2} \|\Gamma^{-1/2} \xi^{(j)}\|_Y \quad (25)$$

Since $\xi^{(j)}$ is the (known) data perturbation (11), our estimate of $\eta^{(j)}$ can be computed. Additionally, as discussed in Section 1 we assume that an estimate of the noise level η is available. For the computation of the noise level η we refer the reader to the discussion in [22, Section 3].

We now combine the previous ideas in the following iterative method that generates an ensemble of inverse estimates of geologic properties.

Algorithm 1 (iteratively regularized ensemble LM)

Let us consider $u_0^{(j)} \equiv u^{(j)} \sim N(\bar{u}, C)$, $y^{(j)}$ according to (11). Let $\rho < 1$ and $\tau > 1/\rho$. Use the following implementation of the regularizing LM scheme to compute an approximation to (12) for $j \in \{1, \dots, N_e\}$. For $m = 1, \dots$

(1) **Stopping rule (Discrepancy Principle).** If

$$\|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}))\|_Y \leq \tau \left(\eta + \frac{1}{2} \|\Gamma^{-1/2} \xi^{(j)}\|_Y \right) \quad (26)$$

stop. Output: $u_m^{(j)}$.

(2) **Selection of $\alpha_m^{(j)}$.** Let $\alpha_{m,n}^{(j)} = 2^n \alpha_{m,0}^{(j)}$ for $n \geq 0$ with $\alpha_{m,0}^{(j)} = 1$. Let $N \in \mathbb{N} \cup \{0\}$ be the minimum such that $\alpha_m^{(j)} \equiv \alpha_{m,N}^{(j)}$ satisfies

$$\rho^2 \|\Gamma^{-1/2}(y^{(j)} - G(u_m^{(j)}))\|_Y^2 \leq (\alpha_m^{(j)})^2 \|\Gamma^{1/2} [DG(u_m^{(j)}) CDG^*(u_m^{(j)}) + \alpha_m^{(j)} \Gamma]^{-1} [y^{(j)} - G(u_m^{(j)})]\|_Y^2 \quad (27)$$

(3) **Update.** Define

$$u_{m+1}^{(j)} = u_m^{(j)} + CDG^*(u_m^{(j)}) [DG(u_m^{(j)}) CDG^*(u_m^{(j)}) + \alpha_{m,N}^{(j)} \Gamma]^{-1} [y^{(j)} - G(u_m^{(j)})] \quad (28)$$

Note that in the right-hand side of (26) we have used our estimate of the noise level (25). Also note that ensemble members are independent from one another. Thus, the computation of the ensemble members in Algorithm 1 can be easily parallelized. In addition, the iterative selection of $\alpha_m^{(j)}$ (step 2 of Algorithm 1) generates an increasing sequence which, as proven in [22], ensures that (27) is satisfied for some $\alpha_{m,N}^{(j)}$ for N finite.

We reiterate that the proposed IR-enLM algorithm generates an ensemble of minimizers of (12) that aim at capturing the Bayesian posterior $P(u|y)$. The main goal of the present work is to use numerical experiments to understand the effect of the tunable parameters ρ and τ in the associated approximation properties of the proposed method. It is worth mentioning that the ensemble approximation obtained with the proposed method may be substantially different from the one achieved with RML. Indeed, with the proposed IR-enLM scheme we compute an ensemble of regularized (by the regularizing LM scheme) of minimizers of (12). In contrast, RML generates an ensemble of minimizers of the (already regularized) cost functional (10).

2.3 The tunable parameters of IR-enLM

Both ρ and τ are tunable parameters of the proposed IR-enLM which arise from the application of the regularizing LM scheme of [20]. The requirements $\rho \in (0, 1)$ and $\tau > 1/\rho$ are part of the set of sufficient conditions for the convergence of the regularizing LM scheme to a stable solution of the deterministic (least-square) inverse problem in the small noise limit. Therefore, these conditions will ensure the convergence to a stable solution of the deterministic (least-square) inverse problem associated to the computation of each ensemble in expression (12). However, we emphasize that the convergence analysis of the ensemble approximation of Bayesian posterior is nonexistent. In particular, the aforementioned conditions on ρ and τ have not been studied in the context of solving the Bayesian inverse problem. In Section 4 of the present manuscript we will provide extensive numerical studies to understand role of ρ and τ in terms of approximating the mean and variance of synthetic posteriors. In the paragraph below we provide a brief discussion to gain intuition on the effect of the tunable parameters on the computation of each ensemble member of IR-enLM.

The parameter ρ in (18) appears from the application of the discrepancy principle. This controls the output on the linearized data misfit (17) and prevents data overfitting. For a thorough discussion of the discrepancy principle applied to the selection of $\alpha_m^{(j)}$ we refer the reader to [22]. However, it is important to emphasize that the parameter ρ in (18) controls the increment $\Delta u_m^{(j)}$ on each step of the regularizing LM scheme applied for the minimization of (12). More precisely, the larger ρ the larger $\alpha_m^{(j)}$. Alternatively, we observe from (13) that larger $\alpha_m^{(j)}$'s are associated to smaller minimizers (or increments) $\Delta u_m^{(j)}$. Thus, we note that the Tikhonov term (multiplied by $\alpha_m^{(j)}$) in (13) is essential to control the large values that $\Delta u_m^{(j)}$ may potentially take because of the ill-posedness of the linearized inverse problem of minimizing (13) without the regularization term. Note, however, that the control on the increment $\Delta u_m^{(j)}$ given by (18) is proportional to the data misfit, which in turn, decreases with the number of iterations. Therefore, the regularizing LM scheme ensures smaller steps of the LM at the beginning of the iterations. For smaller values of ρ (with $\rho \in (0, 1)$) the increment $\Delta u_m^{(j)}$ can be potentially large. However, in this case (of a smaller ρ) larger values of τ are expected in the discrepancy principle (26) which, in turn, will ensure stability by the early termination of the scheme. In the latter case the accuracy of the scheme may be compromised. In other words, the smaller

ρ the larger $\tau > 1/\rho$ and the earlier we may have to stop the IR-enLM according to (26), leading to potential inaccurate inverse estimates.

2.4 Computational cost of IR-enLM

On each iteration level and for each ensemble member of IR-enLM, the main computational cost is the explicit computation of $DG(u_m^{(j)})CDG^*(u_m^{(j)})$. This is however, computationally feasible provided that the adjoint $DG^*(u_m^{(j)})$ is available and that a small number of observations are assimilated. In that case, the cost of constructing $DG(u_m^{(j)})CDG^*(u_m^{(j)})$ is proportional to the number of observations N_D times the computational cost of solving the adjoint problem of the linearized forward model (i.e. the cost of computing $DG^*(u_m^{(j)})$) which depends on the particular reservoir model under consideration. In addition, IR-enLM requires the evaluation of forward model at each iteration. Therefore, the cost of an N_e size IR-enLM is around $N_e J$ forward model evaluations plus $N_e J N_D$ adjoint solves, where J is the number of iterations to achieve convergence. We remark that the working assumption of small number of measurements is essential for the computational feasibility of the proposed IR-enLM. However, this is not a restrictive assumption for other iterative regularization methods. In [22] we discuss other iterative regularization methods that avoid the explicit computation of $DG(u_m^{(j)})CDG^*(u_m^{(j)})$ and that can be potentially used in the ensemble-base framework proposed here. However, investigating the use of additional iterative regularization methods is beyond the scope of this manuscript.

2.5 The linear case

When G is a linear operator, the tunable parameters in the regularizing LM scheme for history matching can be chosen so that the resulting approximation coincides with the MAP estimator [22]. Analogously, in the linear case, the following proposition provides conditions of the tunable parameters so that the proposed IR-enLM Algorithm 1 coincides with RML (9), which as proven in [30], generates samples of the posterior.

Proposition 1 *Let G be a linear operator $G(u) \equiv Gu$. Consider $u_0^{(j)}$ and $y^{(j)}$ as in Algorithm 1. Assume that*

$$\|\Gamma^{-1/2}(y^{(j)} - Gu_0^{(j)})\|_Y > \left[\eta + \frac{1}{2} \|\Gamma^{-1/2} \xi^{(j)}\|_Y \right] \quad (29)$$

for all $j \in \{1, \dots, N_e\}$. Let $\rho \in \mathbb{R}$ be such that

$$\rho < \frac{1}{\| [GG^* + \Gamma]^{1/2} \Gamma^{-1/2} \|^2} \quad (30)$$

and

$$\tau > \max \left\{ \left\| [GCG^* + \Gamma]^{-1/2} \Gamma^{1/2} \right\|^2, \frac{1}{\rho} \right\} \frac{\max_j \|\Gamma^{-1/2}(y^{(j)} - Gu_0^j)\|_Y}{\eta + \frac{1}{2} \min_j \|\Gamma^{-1/2} \xi^{(j)}\|_Y} \quad (31)$$

Then, IR-enLM produces samples of the posterior distribution.

Proof Note that if assumption (29) does not hold, then the proposed algorithm will terminate and the prior ensemble yields the solution. However, from this assumption and (31) we find

$$\tau > \frac{1}{\rho} \frac{\max_j \|\Gamma^{-1/2}(y^{(j)} - Gu_0^j)\|_Y}{\eta + \frac{1}{2} \min_j \|\Gamma^{-1/2} \xi^{(j)}\|_Y} \geq \frac{1}{\rho} \frac{\|\Gamma^{-1/2}(y^{(j)} - Gu_0^j)\|_Y}{\eta + \frac{1}{2} \|\Gamma^{-1/2} \xi^{(j)}\|_Y} > \frac{1}{\rho}. \quad (32)$$

Additionally, with a similar argument to the one used in [22, Proposition 1] we can show that ρ in (30) satisfies $\rho < 1$ and that the initial guess $\alpha_{0,0}^{(j)} = 1$ satisfies inequality in (27). Therefore, $N = 0$, $\alpha_0^{(j)} = 1$ and so expression (28) becomes

$$u_1^{(j)} = u_0^{(j)} + CG^*[GCG^* + \Gamma]^{-1}[y^{(j)} - Gu_0^{(j)}] \quad (33)$$

which is the minimizer of (10) with $G(u) = Gu$. In other words, $\{u_1^{(j)}\}_{j=1}^{N_e}$ is an ensemble generated with the RML method described above. Therefore, as proven in [30], $\{u_1^{(j)}\}_{j=1}^{N_e}$ are samples of the posterior. We now show that $u_1^{(j)}$ satisfies the stopping criteria with τ given by (31). Note that

$$\begin{aligned} \|\Gamma^{-1/2}(y^{(j)} - Gu_1^{(j)})\|_Y &= \|\Gamma^{1/2}[GCG^* + \Gamma]^{-1}(y^{(j)} - Gu_0^{(j)})\|_Y \\ &\leq \|[GCG^* + \Gamma]^{-1/2} \Gamma^{1/2}\| \|\Gamma^{-1/2}(y^{(j)} - Gu_0^{(j)})\|_Y \\ &\leq \|[GCG^* + \Gamma]^{-1/2} \Gamma^{1/2}\|^2 \max_j \|\Gamma^{-1/2}(y^{(j)} - Gu_0^{(j)})\|_Y \\ &< \tau \left(\eta + \frac{1}{2} \min_j \|\Gamma^{-1/2} \xi^{(j)}\|_Y \right) \leq \tau \left(\eta + \frac{1}{2} \|\Gamma^{-1/2} \xi^{(j)}\|_Y \right) \end{aligned} \quad (34)$$

which implies that $\{u_1^{(j)}\}_{j=1}^{N_e}$ satisfies our stopping criteria (26). Therefore, in the linear case, each ensemble of the proposed algorithm terminates after one iteration and the resulting ensemble members are samples of the posterior distribution. \square

Remark 1 While the linear case is not relevant for most reservoir applications, the previous result provides conditions on the tunable parameters under which the consistency of the IR-enLM for sampling the (Gaussian) posterior can be established. Note that the conditions (30) and (31) are consistent to the conditions required by the regularizing LM scheme to ensure stability (i.e. $\rho \in (0, 1)$ and $\tau > 1/\rho$). However, we emphasize that the conditions (30) and (31) do not ensure that, in the nonlinear case, the ensemble IR-enLM provides the proper sampling of the posterior. In Section 4 we provide a numerical study, on a nonlinear reservoir model, to understand the approximation properties of IR-enLM and its dependence on the choice of the tunable parameters described earlier.

2.6 Differences between IR-enLM and the unregularized LM method of [2] for RML

We note that when G is linear, the ensemble generated with RML can be obtained explicitly by formula (33). In this case we were also able to show in Proposition 1 an equivalence between RML and the proposed IR-enLM method. In the nonlinear case relevant for the present application such equivalence is nonexistent. Moreover, any comparison between RML and the proposed IR-enLM depends on the particular implementation for the minimization of (10). For the present work we consider the unregularized LM scheme of [2] for RML (i.e. for minimizing (10)). This “unregularized” method does not add any additional form of regularization, like for example, a truncated SVD parameterization which, as stated before, has the effect of regularizing the LM step. We emphasize, however, that this unregularized LM is applied to the minimization of (10) which is already regularized. The main point of this section and the associated numerical experiments of subsection 4.6 is to show that the intrinsic regularization in (10) combined with the standard methods of optimization may generate unstable ensemble members which are, in turn, detrimental to the robustness for recovering/approximating the Bayesian posterior.

The unregularized LM method of [2] applied for generating ensembles of RML is defined by $u_{m+1}^{(j)} = u_m^{(j)} + \Delta u^{(j)}$ where the step $\Delta u^{(j)}$ satisfies

$$\left[DG^*(u_m^{(j)})\Gamma^{-1}DG(u_m^{(j)}) + C^{-1} + \lambda_m^{(j)}C^{-1} \right] \Delta u^{(j)} = DG^*(u_m^{(j)})\Gamma^{-1}[y^{(j)} - G(u_m^{(j)}) - C^{-1}(u_m^{(j)} - u^{(j)})] \quad (35)$$

We reiterate that this unregularized version of the LM method excludes truncated SVD parametrization and further localization/inflation of the associated matrices in (35).

Most applications of both regularized and unregularized LM approaches for RML-based methods [33, 30, 6, 27, 34] consistently choose the initial regularization LM parameter λ in (35) as $\lambda_0 = \Lambda_0 \equiv \max\{\sqrt{J(u_0)/N_D}, J(u_0)/N_D\}$. For $m \geq 0$, λ_{m+1} is then chosen according to

$$\lambda_{m+1} = \begin{cases} \lambda_m/\kappa & \text{if } J(u_{m+1}) < J(u_m) \\ \kappa\lambda_m & \text{if } J(u_{m+1}) \geq J(u_m) \end{cases} \quad (36)$$

where $\kappa = 10$ is the typical selection [33, 30, 5, 27, 34]. Moreover, all these approaches use the following stopping criteria

$$\frac{|J(u_{m+1}) - J(u_m)|}{J(u_{m+1})} \leq \varepsilon_0, \quad \frac{\|u_{m+1} - u_m\|_X}{\|u_{m+1}\|_X} \leq \varepsilon_1 \quad (37)$$

In addition, it has been claimed often that a successful application of these methods requires the production of estimates that satisfy

$$J_{RML}^{(j)}(u^{(j)}) \leq N_D + 5\sqrt{2N_D} \quad (38)$$

where N_D is the number of measurements (i.e. the dimension of Y). While (38) has been proven under the linear assumption on G , for the general case, an analogous

estimate is nonexistent. It is important to reiterate that, a sufficient decrease in the data misfit may not be associated with a decrease of the error with respect to the truth.

For a thorough discussion of the technical differences between the unregularized LM method and the regularizing LM scheme we refer the reader to the discussion of [22, section 3.3]. In subsection 4.6 we display experiments where the unregularized LM implementation of [2] for RML leads to ensembles that exhibit suboptimal performance in capturing the posterior distribution. This lack of stability and the resulting uncontrolled estimates of this scheme arise from the selection of the regularization parameter $\lambda_m^{(j)}$. In contrast, we observe that the regularizing LM scheme implemented within our IR-enLM uses the discrepancy principle to control selection of the LM parameter as well as the early termination of the algorithm. Therefore, even though (37) and (38) have been shown to produce acceptable numerical results in cases where the problem has been further regularized (e.g. like with truncated SVD [33]), the general application of the selection of regularization parameters and stopping criteria should be taken with caution.

3 Ensemble Kalman-based methods

Using the derivative of the forward operator in variational models often results in more accurate inverse estimates of the subsurface properties compared to the ones generated with ensemble Kalman-based methods. However, as indicated earlier, Kalman-based methods are often easier to implement and possibly the only choice when adjoint codes are not available. Based on ideas from iterative regularization methods, in this section we construct an ensemble Kalman smoother suitable for an all-at-once formulation consistent with our forward operator G that comprises all data (4). Nevertheless, the ideas presented here can be developed in a sequential formulation by redefining the forward problem.

3.1 The standard unregularized ensemble smoother

The standard smoother of [14] consist of generating an ensemble

$$\mathbf{u}^{(j,a)} = \mathbf{u}^{(j,f)} + \mathbf{C}^{uw} (\mathbf{C}^{ww} + \Gamma)^{-1} (\mathbf{y}^{(j)} - \mathbf{w}^{(j,f)}) \quad (39)$$

where $\mathbf{u}^{(j,f)} \sim N(\bar{\mathbf{u}}, \mathbf{C})$, $\mathbf{y}^{(j)}$ is defined as in (11), and

$$\begin{aligned} \mathbf{w}^{(j,f)} &\equiv G(\mathbf{u}^{(j,f)}), & \bar{\mathbf{w}}^f &= \frac{1}{N_e} \sum_{j=1}^{N_e} \mathbf{w}^{(j,f)}, & \bar{\mathbf{u}}^f &= \frac{1}{N_e} \sum_{j=1}^{N_e} \mathbf{u}^{(j,f)} \\ \mathbf{C}^{ww} &= \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (\mathbf{w}^{(j,f)} - \bar{\mathbf{w}}^f)(\mathbf{w}^{(j,f)} - \bar{\mathbf{w}}^f)^T, \\ \mathbf{C}^{uw} &= \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (\mathbf{u}^{(j,f)} - \bar{\mathbf{u}}^f)(\mathbf{w}^{(j,f)} - \bar{\mathbf{w}}^f)^T \end{aligned} \quad (40)$$

Similar to straightforward applications of EnKF, the implementation of ES often provides poor data misfit and/or very large/rough values of the ensemble members. Moreover, these standard Kalman-based methods usually underestimate the variance due to the collapse of the ensemble members towards the ensemble mean [23, 13]. As described in Section 1, the aforementioned behavior has become more evident when a small ensemble size is considered compared to the observations or the dimensions of the parameter space. Several forms of regularization to alleviate this ill-behavior have been proposed in terms of covariance localization, covariance inflation and/or redefining the filter in a square-root fashion. In the present section we demonstrate how an alternative form of regularization can be achieved by introducing ideas from iterative regularization to the definition of the ES.

3.2 The regularizing ES

Similar to expression (39) we may consider the following Kalman formula in order to update the forward model predictions (i.e. $w^{(j,f)} \equiv G(u^{(j,f)})$)

$$w^{(j,a)} = w^{(j,f)} + C^{ww} (C^{ww} + \Gamma)^{-1} (y^{(j)} - w^{(j,f)}) \quad (41)$$

If we now define $z^{(j,f)} = [u^{(j,f)}, w^{(j,f)}]^T$ and the analogous expression for $z^{(j,a)}$, we can then express (39) and (41) as

$$z^{(j,a)} = z^{(j,f)} + C^f H^T (H C^f H^T + \Gamma)^{-1} (y^{(j)} - H z^{(j,f)}) \quad (42)$$

where $H = [0, I]$ and

$$C^f = \begin{pmatrix} C^{uu} & C^{uw} \\ (C^{uw})^T & C^{ww} \end{pmatrix} \quad (43)$$

From the expression (42) it follows that

$$z^a \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} z^{(j,a)} = z^f + C^f H^T (H C^f H^T + \Gamma)^{-1} (y - H z^f) \quad (44)$$

where $z^f \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} z^{(j,f)}$. It is straightforward to show that (44) is equivalent to

$$z^a = \operatorname{argmin}_z \left(\|\Gamma^{-\frac{1}{2}} (y - H z)\|_Y^2 + \|(C^f)^{-\frac{1}{2}} (z - z^f)\|_Z^2 \right) \quad (45)$$

If we now define,

$$\begin{aligned} \chi^a &\equiv (C^f)^{-\frac{1}{2}} (z^a - z^f), & \chi &\equiv (C^f)^{-\frac{1}{2}} (z - z^f), \\ d &\equiv y - H z^f, & L &\equiv H (C^f)^{\frac{1}{2}} \end{aligned} \quad (46)$$

and substitute these expression in (45) we obtain

$$\chi^a = \operatorname{argmin}_\chi \left(\|\Gamma^{-\frac{1}{2}} (d - L \chi)\|_Y^2 + \|\chi\|_Z^2 \right) \quad (47)$$

which we recognize as the Tikhonov regularization method (with Tikhonov-parameter $\alpha = 1$) applied to the following linear inverse problem:

$$\text{Given } d, \text{ find } \chi \in Z \text{ such that } L\chi = d \quad (48)$$

In subsection 4.6 we provide examples to show that the intrinsic choice of $\alpha = 1$ in (47) does not necessarily provide enough regularization to this linear inverse problem. Therefore, this unregularized standard ES produces rough estimates with large data misfit like the ones that have been consistently reported in the literature [23,13]. Moreover, this lack of regularization is, in turn, detrimental to the characterization of the posterior obtained from the ensemble approximation.

The essence of the propose regularizing ES is to include a regularization parameter α so that the resulting problem (47) takes the form

$$\chi^a(\alpha) = \operatorname{argmin}_{\chi} \left(\|\Gamma^{-\frac{1}{2}}(d - L\chi)\|_Y^2 + \alpha \|\chi\|_Z^2 \right) \quad (49)$$

More importantly, we propose this α to satisfy a discrepancy principle similar to the one used in the regularizing LM scheme (17). More precisely, we require

$$\|\Gamma^{-1/2}(d - L\chi^a(\alpha))\|_Y \geq \rho \|\Gamma^{-1/2}d\|_Y \quad (50)$$

for some choice of $\rho \in (0, 1)$. Let us now define the ‘‘true’’ parameters

$$\chi^\dagger \equiv (C^f)^{-\frac{1}{2}}(z^\dagger - z^f), \quad z^\dagger \equiv \begin{bmatrix} u^\dagger \\ G(u^\dagger) \end{bmatrix} \quad (51)$$

where u^\dagger denotes the true geologic property. Furthermore, we define the true data misfit

$$d^\dagger \equiv L\chi^\dagger = G(u^\dagger) - Hz^f. \quad (52)$$

From (46) and (52) we find

$$\|\Gamma^{-1/2}(d - d^\dagger)\|_Y = \|\Gamma^{-1/2}(y - G(u^\dagger))\|_Y \equiv \eta \quad (53)$$

where we have used the definition of the noise level (3). Therefore, if we assume that

$$\eta \leq \rho \|\Gamma^{-1/2}d\|_Y, \quad (54)$$

our choice of α in (50) yields

$$\|\Gamma^{-1/2}(d - d^\dagger)\|_Y = \eta \leq \rho \|\Gamma^{-1/2}d\|_Y \leq \|\Gamma^{-1/2}(d - L\chi^a(\alpha))\|_Y \quad (55)$$

which is a discrepancy principle applied to the inverse problem (48). Recall that a similar application of the discrepancy principle is used for the selection of the regularizing LM scheme (see expression (16) in LM). Therefore, by means of (55) our selection of α controls the ensemble update, so that the resulting mean χ^a does not produce a data fit better than the noise level.

In order to produce an ensemble whose transformed mean satisfies (50) we define

$$z^{(j,a)} = z^{(j,f)} + C^f H^T \left(HC^f H^T + \alpha \Gamma \right)^{-1} (y^{(j)} - Hz^{(j,f)}) \quad (56)$$

whose mean satisfies

$$z^a = \operatorname{argmin}_z \frac{1}{2} \left(\|\Gamma^{-\frac{1}{2}}(y - Hz)\|^2 + \alpha \|(C^f)^{-\frac{1}{2}}(z - z^f)\|^2 \right) \quad (57)$$

We now use (46) to rewrite (50) as

$$\rho \|\Gamma^{-1/2}(y - Hz^f)\|_Y \leq \|\Gamma^{-1/2}(y - Hz^a)\|_Y \quad (58)$$

which in terms of (57) and (43) can be written as follows

$$\rho \|\Gamma^{-1/2}(y - \bar{w}^f)\|_Y \leq \alpha \|\Gamma^{1/2}(C^{ww,f} + \alpha\Gamma)^{-1}(y - \bar{w}^f)\|_Y \quad (59)$$

which shows that the actual computation of α is analogous to the one in (27) for the IR-enLM algorithm defined in the preceding section. Finally, we use (46) to rewrite the assumptions (54) as

$$\eta \leq \rho \|\Gamma^{-1/2}(y - Hz^f)\|_Y = \rho \|\Gamma^{-1/2}(y - \bar{w}^f)\|_Y \quad (60)$$

This assumption motivates the stopping criteria of our Iterative Regularizing ES (IR-ES) that we describe in the following sections.

3.3 The Iterative Regularizing ensemble Smoother (IR-ES)

The previous regularizing ES may alleviate the ill-posedness of the linear inverse problem (48). While the aim of the proposed ensemble methods is to avoid data overfitting, the aforementioned linearized inversion may not necessarily produce estimates with a reasonably small data misfit that properly explore the regions of the posterior with large probability. Therefore, in order to allow for more accurate approximations of the posteriors without the risk of overfitting data, in the following algorithm we define an iterative regularized ES scheme terminated according to the discrepancy principle applied to the mean of the updated (analyzed) model predictions $w_m^{(j,a)}$ at a given iteration.

Algorithm 2 (Iterative regularizing ensemble Kalman smoother)

Prior ensemble and perturbed noise. Let $\rho < 1$ and $\tau > 1/\rho$. Generate

$$u_0^{(j)} \sim \mu_0, \quad y^{(j)} \equiv y + \xi^{(j)}, \quad \xi^{(j)} \sim N(0, \Gamma), \quad j \in \{1, \dots, N_e\} \quad (61)$$

where we recall from subsection 1.2 that $\mu_0 = \mathbb{P}(u)$ is the prior which in general may be non-Gaussian.

Iterative Smoother. For $m = 0, \dots$

(1) **Prediction Step:** Evaluate

$$w_m^{(j,f)} = G(u_m^{(j)}) \quad j \in \{1, \dots, N_e\} \quad (62)$$

(2) **Stopping criteria:** Compute $\bar{w}_m^f = \frac{1}{N_e} \sum_{j=1}^{N_e} w_m^{(j,f)}$. If

$$\|\Gamma^{-1/2}(y - \bar{w}_m^f)\|_Y \leq \tau\eta, \quad (63)$$

stop. Output: $\{u_m^{(j)}\}_{j=1}^{N_e}$. Compute

$$\bar{u}_m = \frac{1}{N_e} \sum_{j=1}^{N_e} u_m^{(j)} \quad C_m^{uw} = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (u_m^{(j)} - \bar{u}_m)(w_m^{(j,f)} - \bar{w}_m^f)^T \quad (64)$$

$$C_m^{ww} = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (w_m^{(j,f)} - \bar{w}_m^f)(w_m^{(j,f)} - \bar{w}_m^f)^T \quad (65)$$

(3) **Analysis step:** Compute the updated ensembles

$$u_{m+1}^{(j)} = u_m^{(j)} + C_m^{uw}(C_m^{ww} + \alpha_m \Gamma)^{-1}(y^{(j)} - w_m^{(j,f)}) \quad (66)$$

$$w_{m+1}^{(j,a)} = G(u_m^{(j)}) + C_m^{ww}(C_m^{ww} + \alpha_m \Gamma)^{-1}(y^{(j)} - w_m^{(j,f)}) \quad (67)$$

for α_m such that

$$\alpha_m \|\Gamma^{1/2}(C_m^{ww,f} + \alpha_m \Gamma)^{-1}(y^\eta - \bar{w}_m^f)\|_Y \geq \rho \|\Gamma^{-1/2}(y^\eta - \bar{w}_m^f)\|_Y \quad (68)$$

(compute α_m with the increasing sequence of Algorithm 1 having an initial guess $\alpha_m = 1$).

Remark 2 Note that the stopping criteria (63) is consistent with the assumption in (60). Indeed, if at a given iteration level m , \bar{w}_m^f satisfies

$$\eta\tau < \|\Gamma^{-1/2}(y - \bar{w}_m^f)\|_Y, \quad (69)$$

for some τ with $\tau > 1/\rho > 1$, then

$$\eta < \rho \|\Gamma^{-1/2}(y - \bar{w}_m^f)\|_Y, \quad (70)$$

which is assumption (60) which, in turn, ensures we select the regularization parameter according to the discrepancy principle. Note that, in contrast to IR-enLM where the regularizing LM scheme was applied to each ensemble computed as the minimizer of (12), in the IR-ES the regularization is conducted over the mean of IR-ES. Since the ensemble of perturbed observation is assumed centered, the resulting equation (49) only contains the unperturbed (i.e. original) data. Thus, the discrepancy principle is applied with the original noise level η .

Similar to IR-enLM, the proposed IR-ES aims at providing an ensemble approximation of the Bayesian posterior $P(u|y)$. However, the convergence properties of the proposed IR-ES are an open problem beyond the scope of this paper. Nevertheless, in the linear case (recall we assume a Gaussian prior), the proposed algorithm coincides with the standard ES which, in turn, generates samples of the posterior as the ensemble size goes to infinity.

Proposition 2 Let G be a linear operator $G(u) \equiv Gu$. Assume that the mean of the prior ensemble satisfies

$$\|\Gamma^{-1/2}(y - G\bar{u}_0)\|_Y > \eta \quad (71)$$

where η is the noise level defined by (3). Let $\rho \in \mathbb{R}$ be such that

$$\rho < \frac{1}{\| [GC_0^{uu}G^T + \Gamma]^{1/2}\Gamma^{-1/2} \|^2} \quad (72)$$

and

$$\tau > \max \left\{ \| [GC_0^{uu}G^T + \Gamma]^{-1/2}\Gamma^{1/2} \|^2, \frac{1}{\rho} \right\} \frac{\|\Gamma^{-1/2}(y - G\bar{u}_0)\|_Y}{\eta} \quad (73)$$

Then, this selection of ρ and τ ensure that the IR-ES method samples the posterior as $N_e \rightarrow \infty$.

Proof From the linearity of G it follows that

$$\bar{w}_0^f = G\bar{u}_0, \quad C_0^{uw} = C_0^{uu}G^T, \quad C_0^{ww} = GC_0^{uu}G^T \quad (74)$$

where

$$C_0^{uu} = \frac{1}{N_e} \sum_{j=1}^{N_e} (u_0^{(j)} - \bar{u}_0)(u_0^{(j)} - \bar{u}_0)^T$$

With similar arguments to those used in the proof of Proposition 1 (see also Proposition 1 in [22]) we can show that ρ and τ in (72)-(73) satisfy $\rho < 1$, $\tau > 1/\rho$. Moreover, these choices of ρ and τ also imply that $\alpha_0 = 1$ satisfies (68). Therefore, (66) becomes

$$u_1^{(j)} = u_0^{(j)} + C_0^{uw} (C_0^{ww} + \alpha_0 \Gamma)^{-1} (y^{(j)} - Gu_0^{(j)})$$

which from (74) can be written as

$$u_1^{(j)} = u_0^{(j)} + C_0^{uu}G^T (GC_0^{uu}G^T + \Gamma)^{-1} (y^{(j)} - Gu_0^{(j)}) \quad (75)$$

In addition,

$$w_1^{(j,a)} = Gu_0^{(j)} + GC_0^{uu}G^T (GC_0^{uu}G^T + \Gamma)^{-1} (y^{(j)} - Gu_0^{(j)})$$

and therefore

$$\bar{w}_1^a = \frac{1}{N_e} \sum_{j=1}^{N_e} w_1^{(j,a)} = G\bar{u}_0 + C_0^{uu}G^T (GC_0^{uu}G^T + \Gamma)^{-1} (y - G\bar{u}_0)$$

Thus,

$$\|\Gamma^{-1/2}(y - \bar{w}_1^a)\|_Y = \|\Gamma^{-1/2}(y - G\bar{u}_1)\|_Y = \|\Gamma^{1/2}[GC_0^{uu}G^* + \Gamma]^{-1}(y - G\bar{u}_0)\|_Y$$

which from some simple computations implies

$$\begin{aligned} \|\Gamma^{-1/2}(y - \bar{w}_1^a)\|_Y &\leq \| [GC_0^{uu} G^* + \Gamma]^{-1/2} \Gamma^{1/2} \|^2 \|\Gamma^{-1/2}(y - G\bar{u}_0)\|_Y \\ &\leq \| [GC_0^{uu} G^* + \Gamma]^{-1/2} \Gamma^{1/2} \|^2 \|\Gamma^{-1/2}(y - G\bar{u}_0)\|_Y < \tau\eta \end{aligned}$$

where in the last inequality we have used (73). Therefore, the ensemble (75) satisfies the stopping criteria and so it is the output of the IR-ES scheme. Moreover, we note that

$$\bar{u}_1 = \bar{u}_0 + C_0^{uu} G^T (GC_0^{uu} G^T + \Gamma)^{-1} (y - G\bar{u}_0)$$

and

$$C_1^{uu} = \frac{1}{N_e} \sum_{j=1}^{N_e} (u_1^{(j)} - \bar{u}_1)(u_1^{(j)} - \bar{u}_1)^T = C_0^{uu} - C_0^{uu} G^T (GC_0^{uu} G^T + \Gamma)^{-1} GC_0^{uu}.$$

Then, formally, as $N_e \rightarrow \infty$ we obtain that $C_0^{uu} \rightarrow C$ and so the mean and covariance of the updated ensemble converges to the mean and covariance of the posterior [32]. In other words, in the linear case, the proposed IR-ES method generates samples of the posterior distribution as $N_e \rightarrow \infty$. \square

3.4 Computational cost of IR-ES

For the present application, the cost of computing (64)-(65) at each iteration is negligible compared to the cost of the evaluation of the forward model (62) (for each ensemble member). Indeed, efficient implementations of the Kalman update can be used to avoid the explicit computation of the matrix C^{uw} [1, 26]. In addition, in the case where a relatively small number (around 10^3) of observations are assimilated, the cost of inverting $C_m^{ww,f} + \alpha_m \Gamma$ is also negligible compared to (62). Therefore, the main computational cost of IR-ES per iteration and per ensemble is due to (62). The total cost of a N_e -size ensemble of IR-ES is approximately $N_e J$ forward model evaluations where J is the number of iterations to converge. For the forward models considered in Section 4, our numerical experiments indicate that J is typically between 10 and 20 iterations. Thus, for large models, the computational efficiency of Algorithm 2 may be comparable to the one of the standard ES computed with a large ensemble (e.g. 10^3). On the other hand, large ensembles with standard ES methods provide better approximations of the posterior (see subsection 4.6). One may then conclude that the proposed IR-ES is equivalent to a standard ES with a large ensemble. However, the advantage of the proposed method is that, at the m th iteration level, each ensemble update in the analysis step will be controlled by the parameter α_m which is chosen according to the discrepancy principle. It is then reasonable to expect small changes in the covariances and crosscovariances defined in (64)-(65). In other words, we assume that changes in the updated ensemble members are sufficiently small (due the regularization) so that the ensemble of model predictions (67) approximates the forward model evaluated at the update ensemble (62). Then, in order to reduce the computational cost of Algorithm 2, for $m > 0$ we propose to replace

(62) by

$$w_m^{(j,f)} = \begin{cases} G(u_m^{(j)}) & \text{if } \text{mod}(m, M_{ES}) = 0 \\ w_m^{(j,a)} & \text{otherwise} \end{cases} \quad (76)$$

Clearly, the aforementioned assumption on small changes in the ensemble updates is valid for only a few number of iterations M_{ES} after which the forward model needs to be evaluated (62). Note that in Algorithm 2 we always compute the evaluation of the forward model at the initial ensemble. However, the next evaluation of the forward model at the updated ensemble is done whenever $\text{mod}(m, M_{ES}) = 0$. Therefore, when we use (76) instead of (62) in Algorithm 2, the computational cost becomes to $N_e(1 + \text{floor}(J/M_{ES}))$. It is important to remark that (76) does not affect the results of Proposition 2 since for the linear case the algorithm stops after the first iteration. The effect of the additional parameter M_{ES} will be investigated numerically in subsection 4.4.

For previous exposition we have assumed that the prior distribution is Gaussian $N(\bar{u}, C)$. For the numerical experiments of the subsequent section we use these priors to define our initial ensemble. Moreover, the corresponding covariances are used in the formulas of IR-enLM Algorithm 1 (Note that C does not appear in IR-ES). However, it is important to emphasize that the Gaussian assumption on the prior distribution is not fundamental for the application the proposed methods. Note that an initial ensemble can be potentially sampled from a non-Gaussian prior. In addition, recall that IR-enLM is posed as the minimization of (12) iteratively regularized by (13). If a non-Gaussian prior is considered, the second term in (13) can be replaced by other type of regularization that enforces the prior knowledge of such non-Gaussian prior. In that case, however the resulting solution of the minimizer of (13) may not be computed in a closed form as in (14).

It is important to mention that our proposed IR-ES method has strong similarities with the ES with multiple data assimilation ES-MDA of [12, 13, 11]. Both methods use equation (66) to update each ensemble member in an iterative fashion. However, IR-ES and ES-MDA differ in (i) the selection of α_m (ii) the stopping criteria and (iii) the definition of the perturbed data $y^{(j)}$ used in (66). In contrast to the proposed IR-ES where the number of iterations are determined by the stopping criteria (63), in ES-MDA the number of iterations are defined beforehand. Additionally, while in IR-ES the regularization parameter α_m is selected according to (68), in the ES-MDA version of [12, 13, 11], these parameters are selected a priori and such that they satisfy $\sum_{m=1}^n \alpha_m^{-1} = 1$. In addition, in ES-MDA, expression (66) is used with a $y^{(j)}$ replace by $y_m^{(j)}$ defined by $y_m^{(j)} = y + \xi^{(j)}$ with $\xi^{(j)} \sim N(0, \alpha_m C)$. In other words, in ES-MDA, the data is perturbed at each iteration with a perturbation that depends on the regularization parameter. Both the constraint $\sum_{m=1}^n \alpha_m^{-1} = 1$ and the definition of the perturbed data are needed [12] to prove that ES-MDA provides the correct sampling of the posterior in the linear Gaussian case. In the work of [12, 13, 11] several choices of a priori selections of α_m 's have been investigated. However, most recent research indicates that the adaptive selection of the regularization parameters in ES-MDA can lead to improved results [31]. In the following section we provide a numerical comparison

between IR-ES and ES-MDA and discuss some potential application of IR-ES to the selection of the regularization parameters in ES-MDA.

4 Numerical Results

In this section we present numerical examples of the application of the proposed ensemble methods for capturing the Bayesian posterior. In concrete, we consider the posterior distributions associated to two sets of synthetic data from prototypical oil-water reservoir models. These posteriors are fully resolved with a state-of-the-art MCMC for functions. Our MCMC results provide a gold-standard that we use to investigate the performance of IR-enLM and IR-ES at capturing aspects of the Bayesian posterior. Furthermore, we display the advantage of the proposed methods by comparing them with some unregularized approaches.

4.1 The forward operator, the prior and the synthetic data

We consider a 2D incompressible oil-water reservoir model. The reservoir is defined on a squared domain discretized on a 60×60 grid. A water flood is considered on a time interval of 3 years discretized in 30 time steps. The numerical discretization of the PDEs was conducted with the numerical schemes and the MATLAB implementation discussed in [23]. The geologic property of interest is the absolute permeability of the reservoir. However, the approach can be extended to incorporate additional geologic properties, as well as other model parameters. For the present work we consider two different well constrains. Model A consist of production wells constrained to prescribed bottom hole pressure (BHP) and injectors operated with prescribed rates $\{q_I^l(t) = 2.6 \times 10^3 \text{ m}^3/\text{day}\}_{I=1}^{N_I}$ while production wells are constrained to bottom-hole pressure $\{P_{bh}^l(t) = 2.7 \times 10^7 \text{ Pa}\}_{I=1}^{N_P}$. For Model B we consider production wells operated under prescribed total flow rate and injectors operated under BHP. In this case the injection wells are operated under prescribed pressure $\{P_{bh,I}^l(t) = 2.7 \times 10^7 \text{ Pa}\}_{I=1}^{N_P}$ and the production wells are constrained to total flow rate $\{q_P^l = 2.6 \times 10^3 \text{ m}^3/\text{day}\}_{I=1}^{N_I}$. The expression relation between flow rates and bottom-hole pressure are as described in [7]. For both models we define $N_M = 30$ measurement collection times (at each time step). In addition, we consider a well configuration of nine production wells P_1, \dots, P_9 and four injection wells I_1, \dots, I_4 . Well locations are displayed in Figure 1 (top-middle). For Model A (resp. Model B), G_A (resp. G_B) consist of BHP (resp. water rate) from injection wells and water rates from production wells. The forward operator G_A corresponds to the forward operator used in [22] for history matching with the regularizing LM scheme.

We recall that the prior distribution is assumed Gaussian. For the present experiments we consider a mean $\bar{u} = \log(5 \times 10^{-13} \text{ m}^2)$ constant over the domain of the reservoir. The prior covariance C is a spherical covariance function [9, 30] with maximum (resp. minimum) range of 10^3 m (resp. $5 \times 10^2 \text{ m}$). We consider an angle of $\pi/2$ along the direction of maximum correlation. In order to generate synthetic data we define the true permeability denoted by u^\dagger and displayed in Figure 1 (top-left).

This permeability field is a draw from the prior distribution described before. In other words, we consider the best-case-scenario where our prior knowledge includes the truth. The subsequent procedure to generate synthetic data is analogous to both forward operators and so for the sake of clarity we base our discussion only on G_A .

For simplicity, we consider a diagonal measurement error covariance Γ_A . Note that the diagonal of Γ_A corresponds to the variance of the data described by the forward operator G_A . Let us recall that G_A has coordinates associated to either BHP (at the injection wells) or water rates (at the producers). For the k -th entry of Γ_A that corresponds to a BHP measurement, we define the standard deviation (i.e. the square-root of $(\Gamma_A)_{k,k}$) as 10% of the corresponding k -th component of vector $G_A(u^\dagger)$. For the entries associated to water rates (which are zero before water breakthrough), we select the square-root of $(\Gamma_A)_{k,k}$ as 3% (before breakthrough) or 7% (after breakthrough) of the nominal value of the total flow rate at the associated production well and measurement time. With Γ_A defined as above, synthetic data y_A is generated by

$$y_A = G_A(u^\dagger) + \xi_A \quad (77)$$

with $\xi_A \sim N(0, \Gamma_A)$. The noise level for this synthetic experiment can be computed as

$$\eta_A = \|\Gamma_A^{-1/2} \xi_A\|_Y \quad (78)$$

With the observational noise defined in (77), the resulting noise level is 12% of the weighted norm of the noise-free observations $\|\Gamma_A^{-1/2} G_A(u^\dagger)\|_Y$. In order to illustrate the capabilities of the proposed methods in a realistic scenario where the exact value of η_A will be unknown (since the truth u^\dagger is unknown), instead of the exact value (78) here we use the an estimate of η_A . Since the covariance Γ_A is diagonal, the discussion of [22] suggests that the estimate $\eta_A = \sqrt{M}$ is appropriate where M is the total number of measurements. For this case we have 13 wells measured at each of the 30 time steps and so $M = 390$.

The synthetic data y_B for the forward model G_B is generated with the procedure previously described for the analogous choice of Γ_B . For both forward operators we consider the same prior described above. However, the forward operators G_A and G_B and the corresponding synthetic data y_A and y_B give rise to two different posteriors $\mu_A = \mathbb{P}(u|y_A)$ and $\mu_B = \mathbb{P}(u|y_B)$ defined by (5). The aim of this section is to study the numerical performance of the proposed ensemble methods at capturing the mean and variance of these posteriors.

4.2 Resolving the posteriors

In contrast to deterministic inverse problems where the aim is to recover the truth, in Bayesian inverse problems, the objective is to characterize the posterior distribution (5). Therefore, in order to assess the performance of the proposed ensemble-methods for approximating the Bayesian posterior, an accurate characterization of the posterior needs to be computed. However, as we discussed in Section 1, the application of standard MCMC methods for sampling the posterior that arises from subsurface flow models is usually restricted to small-size problems. While small-size problems,

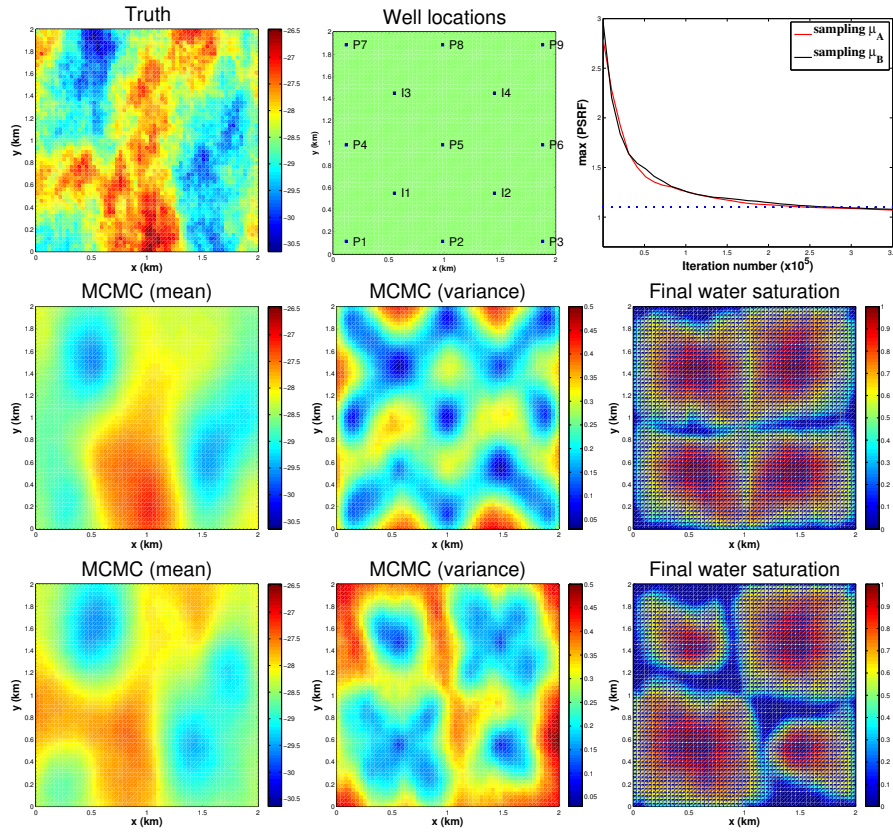


Fig. 1 Top-left: True log-permeability [$\log m^2$]. Top-middle: Well configuration. Top-Right: Gelman-Rubin diagnostic. Middle-left: mean of μ_A . Middle-middle: variance of μ_A . Middle-right: Final water saturation from Model A. Bottom-left: mean of μ_B . Bottom-middle: variance of μ_B . Bottom-right: Final water saturation from Model B.

say from coarse 1D models, may be useful to establish Benchmarks [13], such small size may be detrimental to our ability of properly assessing the performance of the proposed methods under more stringent conditions that require proper regularization. In concrete, the well-known regularizing effect of discretization [25] may alleviate the ill-posedness intrinsic to this inverse problem. This, in turn, may overshadow the proper assessment of the regularizing effect of the proposed methods. Fortunately, grid invariant MCMC methods have been recently developed for sampling the Bayesian posterior that arises in large-scale PDE-constrained inverse problems [8]. In particular, in [23] the preconditioned Crank-Nicolson MCMC (pcn-MCMC) method has been applied to resolve the Bayesian posterior associated to inverse problems in reservoir models of moderate-size similar to the ones considered in the present work. The reader is refer to the work of [23] for further details on the pcn-MCMC algorithm and its application for benchmarking Bayesian inverse problems in subsurface flow models.

By applying the aforementioned pc-MCMC method we resolve the posteriors μ_A and μ_B that we introduced in the previous paragraphs and that we aim to approximate with the proposed ensemble methods. For the sampling of each of these posteriors we consider 80 independent MCMC chains initialized with random draws from the prior distribution (recall the prior is the same for both μ_A and μ_B). After a burn-in period of 5×10^4 each of the MCMC chains consisted of 3×10^5 samples. The convergence and mixing of these independent chains are evaluated with the Gelman-Rubin diagnostic based on the potential scale reduction factor (PSRF) [3]. In Figure 1 (top-right) we display this factor as a function of the MCMC iteration for the sampling of both posteriors μ_A and μ_B . Approximate convergence (PSRF < 1.1) is reached after 2.5×10^5 iterations ensuring that our chains have run sufficiently long so that all the $80 \times 3 \times 10^5 = 2.4 \times 10^7$ samples from all these chain can be combined to characterize the target distribution. In the second (resp. third) row of Figure 1 we display the mean (right) and the variance (middle) of the posterior distribution μ_A (resp. μ_B). Some independent samples (from different MCMC chains) of μ_A are displayed in the top row of Figure 2. While some of the main spatial features of the truth are replicated, sufficient variability can be also appreciated. We use our samples of the posterior μ_A to display in Figure 2 the water rates and BHP from some of the production and injection wells, respectively. In these figures, the red curve correspond to the truth (i.e. $G_A(u^\dagger)$) and the vertical line separates the assimilation from the prediction time. Figure 3 shows analogous quantities for the posterior μ_B associated to the well Model B.

In Figure 1 we display the mean (middle-left) and the variance (middle-middle) of the posterior μ_A characterized with pc-MCMC. Analogously, Figure 1 shows the mean (bottom-left) and the variance (bottom-middle) of the posterior μ_B . Note that, even though the underlying reservoir dynamics are the same for both forward operators, the difference in the well constraints has an important effect on the resulting posterior distribution. In particular we note significant differences in the variance of μ_A and μ_B . Furthermore, in Figure 1 we display the water saturation at the final time (for the assimilation period) associated to Model A (middle-right) and Model B (bottom-right), respectively. We note that, for Model B, larger variances are observed in the regions where the water front has not yet arrived.

4.3 IR-enLM

In this section we present a numerical investigation of the effect of the tunable parameters ρ and τ in the approximation properties of the resulting ensemble obtained with IR-enLM (Algorithm 1). The role of ρ and τ in the regularizing LM scheme for the solution of deterministic inverse problems is well understood in terms of the theory of [20]. In addition, the work of [22] displays the application and validation of such theory applied to history matching problems where the aim was to recover the permeability of an oil-water reservoir. This subsection designed to understand the effect of ρ and τ in the approximation properties of the proposed methods in the context of Bayesian inverse problems. More precisely, we investigate the effect of the tunable parameters in the accuracy and efficiency of IR-enLM for recovering the mean

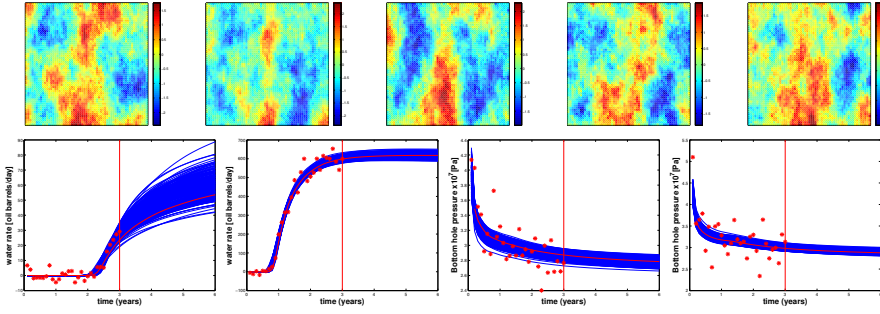


Fig. 2 Top: Samples of the posterior distribution μ_A (characterized with pcn-MCMC) [$\log m^2$]. Bottom: (from left to right) water rates at wells P_1 , P_2 and BHP from I_2 and I_3

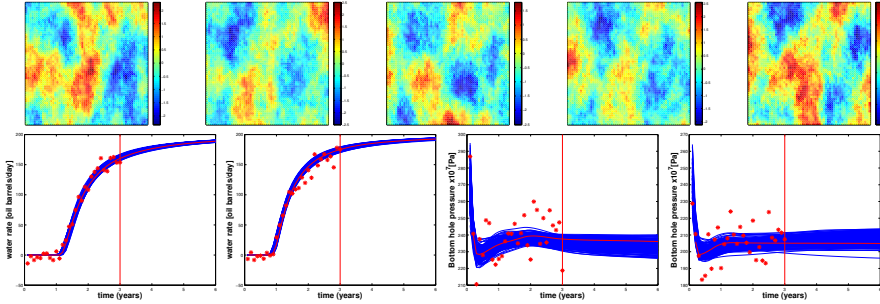


Fig. 3 Top: Samples of the posterior distribution μ_B (characterized with pcn-MCMC) [$\log m^2$]. Bottom: (from left to right) water rates at wells P_1 , P_2 and BHP from I_2 and I_3

and the variance of the posteriors μ_A and μ_B computed with MCMC in the preceding subsection.

We apply IR-enLM to generate ensemble approximations to μ_A and μ_B , respectively. We recall that for both cases, we have the same prior from which the initial ensemble for IR-enLM is drawn. Since the results depend on the initial ensemble, the results we report in this manuscript are averaged over 15 experiments corresponding to different initial ensembles of a fixed size. For each of these experiments, IR-enLM produces an ensemble of inverse estimates whose resulting mean \hat{u} and variance $\hat{\sigma}$ are compared to the mean u_{pos} and variance σ_{pos} of the posterior. More precisely, we consider the relative errors

$$\varepsilon_u = \frac{\|(\hat{u} - \bar{u}) - (u_{pos} - \bar{u})\|_{L^2(D)}}{\|(u_{pos} - \bar{u})\|_{L^2(D)}}, \quad \text{and} \quad \varepsilon_\sigma = \frac{\|\hat{\sigma} - \sigma_{pos}\|_{L^2(D)}}{\|\sigma_{pos}\|_{L^2(D)}}, \quad (79)$$

where we recall that \bar{u} is the mean of the prior distribution of the log-permeability. Note the expression for ε_u corresponds to the relative error of the deviation of the posterior mean with respect to the prior mean. We re-emphasize that the mean and the variance of the posterior (for μ_A and μ_B , respectively) are computed from our pc-MCMC chains.

We recall from our discussion in subsection 2.3 that for a given $\rho \in (0, 1)$, the theory of the regularizing LM scheme [20] requires that we select τ according to $\tau > 1/\rho > 1$ in order to ensure convergence to a stable computation of the minimizer of (12). Note that if G is linear, we may use Proposition 1 to select ρ and τ to ensure the proper sampling of the posterior. For the case considered here, the forward operator G is nonlinear and so we need to investigate the effect of these tunable parameters numerically. For larger values of ρ (i.e. ρ closer to one) yield small incremental steps and smaller τ 's which enable the algorithm to progress for a long number of iterations obtaining more accurate minimizers of (12). On the other hand, smaller values of ρ and so larger values of τ 's are associated with very early termination of the scheme (according to (26)) and so the corresponding estimates of the minimizer of (12) are less accurate although convergence is faster. These considerations in terms of the accuracy of the minimizer of (12) are reflected in the approximation properties of IR-enLM as we can observe from Table 1 where we display the corresponding ensemble approximation of μ_A and μ_B obtained from IR-enLM for several choices of τ , ρ for an ensemble fixed size $N_e = 50$. Indeed, Table 1 confirms that more accurate minimizers of (12) obtained by selecting a smaller value of ρ result in better ensemble approximations of the mean and variance of the posterior. In the 5th and 8th columns of Table 1 we display the average number of iterations for convergence (established by (26)). For each ρ , in Table 1 we display different values of τ with the corresponding value $\tau = 1/\rho$ highlighted blue. Note that as ρ increases, the corresponding $\tau = 1/\rho$ (in blue) produces more accurate solutions. For fixed $\rho = 0.8$ and $N_e = 50$, in Figure 4 we present a visual comparison of the mean (top row) and variance (top-middle row) produce by one ensemble of IR-enLM for approximating μ_A with different choices of τ . For this ensemble and the describe choices of parameters, in the bottom-middle (resp. bottom-bottom) row of Figure 4 we display the box plots of water rate (resp. bottom-hole pressure) at some production (resp. injection) wells after 6 years of injection (i.e. 3 years of assimilation plus 3 years of forecast). Analogous visual results for the approximation of μ_B are displayed in Figure 5.

For the larger values of ρ in Table 1 we observe that the optimal value of τ for reducing the error in the mean and variance is close to $\tau = 1.0$ even though the application of the regularizing LM scheme indicates that τ should be selected to satisfy $\tau > 1/\rho$ to ensure stability. Note, for example, that for $\rho = 0.7$, $\rho = 0.8$ and $\rho = 0.9$ the stabilization is ensure provided $\tau > 1/\rho = 1.45$, $\tau > 1/\rho = 1.25$ and $\tau > 1/\rho = 1.11$, respectively. For these three cases, the errors in the mean and variance can be substantially reduced by considering $\tau = 1.0$. Two potential causes for this improved accuracy with reduced τ are the following. First, we know from [20] that the condition on τ ($\tau > 1/\rho > 1$) is sufficient but not necessary for the convergence of the regularizing LM scheme. In other words, stable and more accurate approximate minimizers may be obtained for $1 < \tau \leq 1/\rho$ thus contributing to more accurate ensemble approximations of the mean and variance of the posterior. Second, our application of the regularizing LM scheme for the computation of the minimizer of (12) is based on the estimate of the noise level provided in (25). Other potential estimates of the noise level $\eta^{(j)}$ may result in τ 's more consistent with the assumption $\tau > 1/\rho$. While obtaining such an optimal estimate is beyond the scope of the present work, it is important to remark that our results reflect that changes in

τ (with $1 \leq \tau$) are associated with stable changes in the relative error of the mean and variance produced by IR-enLM. In addition, our numerical experiments suggest that for IR-enLM with $\rho \geq 0.7$, the value $\tau = 1.0$ in (26)) provides stable and accurate approximations of the mean and variance of the posterior. Note that this value is consistent with the discrepancy principle in the sense that the algorithm is stopped provided that the data misfit matches our estimate (25) of the noise level. Note that, in most cases, the error with respect to the mean can be further reduced by allowing the scheme to progress for more iterations (i.e. by selecting a smaller $\tau < 1$). However, this decrease in the error of the mean is associated with an increase in the error of the variance. We therefore see that the error growth typical of the computation of solutions to ill-posed inverse problems is reflected in the properties of IR-enLM for approximating the Bayesian posterior. While for smaller values of ρ ($\rho < 0.7$) the value of τ seems to be consistent with $\tau > 1/\rho$, the associated errors are substantial and therefore not recommended for further applications of IR-enLM. Note that $\rho \approx 0.8$ with $\tau = 1.0$ seems a reasonable compromise between accuracy and cost.

It is also worth mentioning that the aforementioned conclusions concerning τ apply for several choices of ensemble sizes as it can be observed in Table 2 (for $\rho = 0.8$). Note that when we increase the ensemble size by a factor of ten, the relative error with respect to the mean and variance of μ_A (resp. μ_B) are decreased by 30% (resp. 20%) and 35% (resp. 32%), respectively. Finally, we note that the IR-enLM was capable of obtaining a better approximation of μ_A (in terms of mean and variance) than the one of μ_B . However, we recall that μ_A and μ_B correspond to two different Bayesian inverse problems.

4.4 IR-ES

Similar to the previous subsection, our objective is now to assess the performance of IR-ES at capturing the Bayesian posterior. More importantly, we aim at understanding the role of the tunable parameters ρ , τ and M_{ES} in the proposed method. In Table 3 we display the approximations of the posterior mean and variance with IR-ES for different ensemble sizes N_e and values of τ for fixed parameters $\rho = 0.8$ and $M_{ES} = 10$. The value τ highlighted in blue corresponds to the cutoff $\tau = 1/\rho$. Recall from Section 3.3 that the tunable parameter ρ appears when we apply the discrepancy principle (50) for the selection of the regularization parameter α . Moreover, the parameter τ in (63) should satisfy $\tau > 1/\rho > 1$ for consistency with assumption (54) as discussed in Remark 2. From Table 3 we observe that, for the approximation corresponding to the value $\tau = 1/\rho$ (in blue), the accuracy in the approximation of the mean and variance increases with ensemble size. While in the previous subsection we observed that the approximation properties of IR-enLM did not display significant changes with respect to the ensemble size, here we note that the accuracy of IR-ES strongly depends on the number of ensemble members. This, of course, follows from the fact that the (Tikhonov) regularization of IR-ES is defined in terms a norm that is weighted by covariance constructed from the ensemble (see expression (57)). Indeed, Table 3 shows that for small ensembles ($N_e \leq 50$) the relative errors in the mean and the variance starts increasing before the cut off value $\tau = 1/\rho$ is reached. Note that it

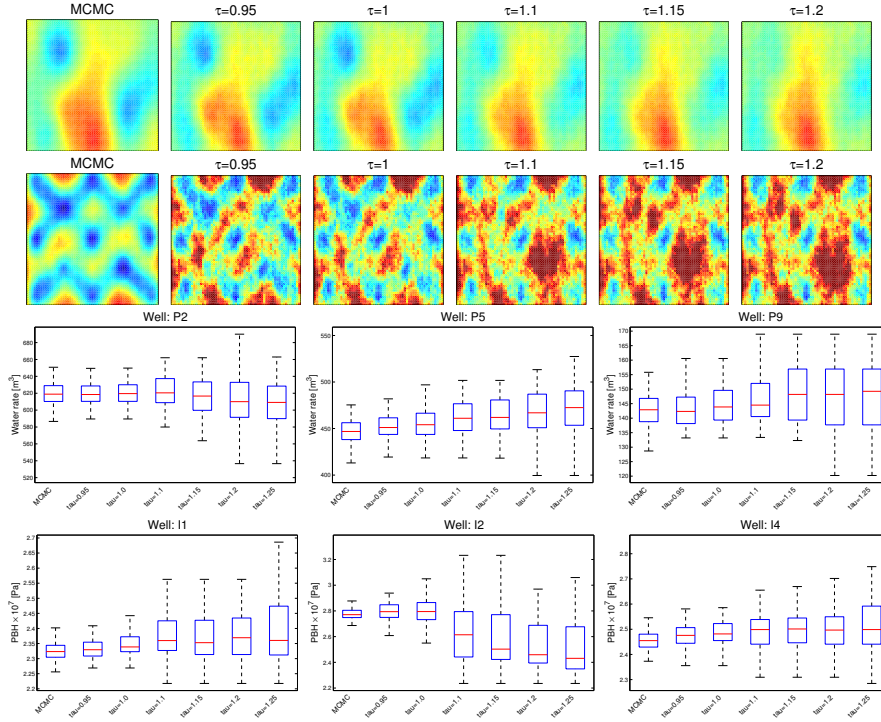


Fig. 4 Top row (mean of μ_A) and Top-middle row (variance of μ_A) from left to right: pcn-MCMC, IR-enLM approximations with $\rho = 0.8$, $N_e = 50$ and $\tau = 0.95$, $\tau = 1.0$, $\tau = 1.1$, $\tau = 1.15$, $\tau = 1.2$. (color scales are the same as in the second row of Figure 1). Middle-bottom and bottom row: Box plots of water rates from production wells P_2, P_3, P_9 and PBH from injections wells I_1, I_2, I_4 after 6 years of water flood simulated from μ_A (with MCMC) and the ensemble approximation with IR-enLM for $\rho = 0.8$ and different choices of τ

is only for larger ensembles ($N_e \geq 75$) where we observe the reduction of the error in the mean and variance as we reduce τ to the value of approximately $\tau = 1/\rho$. Below this critical value, the errors of the ensemble approximation increase. For $\rho = 0.8$, $M_{ES} = 10$ and $N_e = 50$ fixed, in Figure 6 we compare the mean (top row) and variance (top-middle row) generated with IR-ES for approximating μ_A with different choices of τ . In Figure 6 (bottom-middle and bottom rows) we show box plots of relevant quantities at some wells. Similarly, in Figure 7 we present analogous results for the approximation of μ_B .

When the ensemble is sufficiently large ($N_e \geq 75$ for the present case) we note that better approximations are obtained when, for a given τ , the ensemble approximation is stopped according to (63) with $\tau = 1/\rho$. Recall that the parameter ρ in IR-ES controls each update of the ensemble similar to IR-enLM. A larger ρ in (68) results in larger (but adaptively changing) α_m in (68) and therefore more controlled ensemble updates at each iteration; this is reflected in more accurate estimates of the posterior in terms of mean and variance. In Table 4 we show the performance of IR-ES for different values of ρ (for fixed $M_{ES} = 10$ and $N_e = 75$). Note that for a given ρ ,

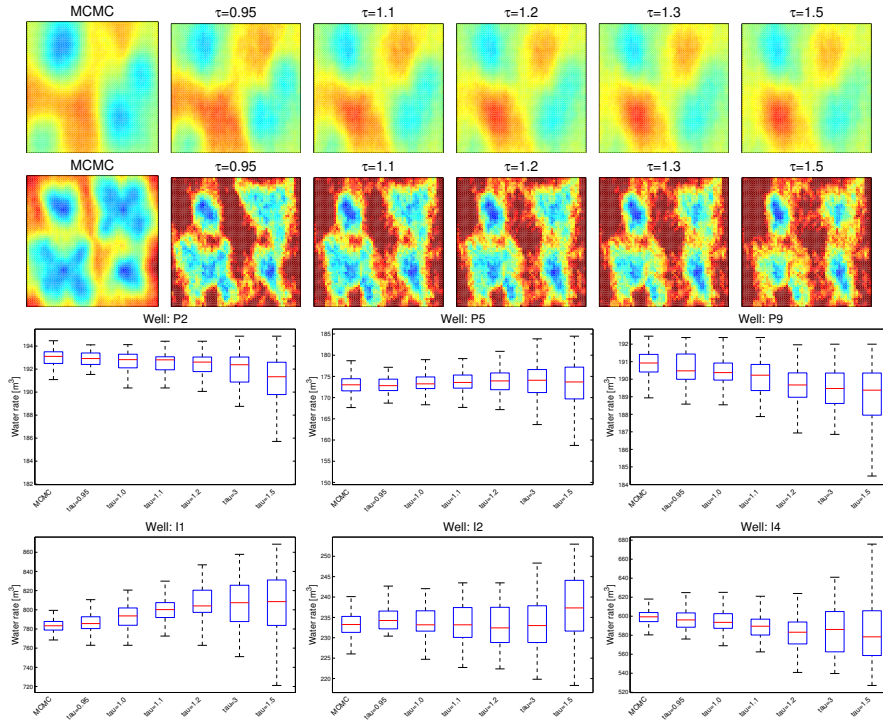


Fig. 5 Top row (mean of μ_B) and Top-middle row (variance of μ_B) from left to right: pcn-MCMC, IR-enLM approximations with $\rho = 0.8$, $N_e = 50$ and $\tau = 0.95$, $\tau = 1.1$, $\tau = 1.2$, $\tau = 1.3$, $\tau = 1.5$ (color scales are the same as in the third row of Figure 1). Middle-bottom and bottom row: Box plots of water rates from production wells P_2, P_5, P_9 and PBH from injection wells I_1, I_2, I_4 after 6 years of water flood simulated from μ_B (with MCMC) and the ensemble approximation with IR-enLM for $\rho = 0.8$ and different choices of τ

decreasing the value of τ reduces the error in the mean and variance decreases until approximately the value values of $\tau = 1/\rho$. Below this value the aforementioned error will increase as expected from the lack of stability in the computation of the ensemble. Smaller ρ corresponds to smaller α_m 's with may not stabilize the scheme and thus the early stopping is required by having a larger τ selected according to $\tau > 1/\rho$. Similar to the IR-enLM, the results from this subsection suggest that the value of $\rho = 0.8$ seems a reasonable selection in terms of cost and accuracy for a given N_e and τ fixed. Moreover, for the experiments under consideration a size of $N_e = 75$ was sufficient to exhibit the regularizing properties that we would expect in the limit of a large ensemble.

In Table 5 we display the performance of IR-ES for different choices of M_{ES} (for $\rho = 0.7, \rho = 0.8$ and $N_e = 75$). As we discussed in subsection 3.4, M_{ES} controls the amount of iterations where an approximation to the forward model outputs is computed from the ensemble. The value $M_{ES} = 1$ corresponds to an evaluation of the forward model per iteration for each ensemble member. This corresponds to the most accurate but extremely expensive case. However, as we increase M_{ES} we observe a

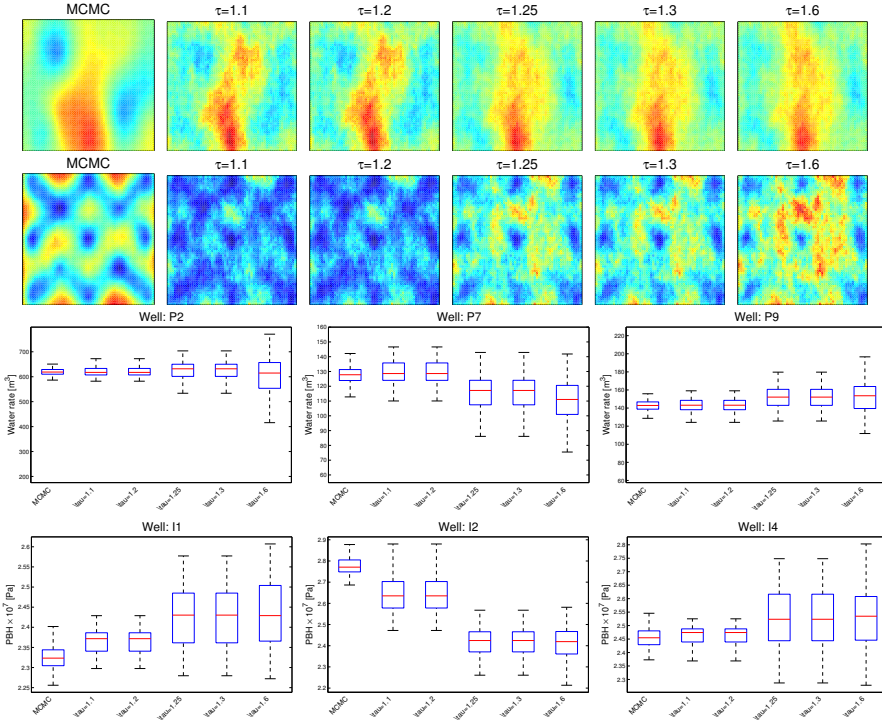


Fig. 6 Top row (mean of μ_A) and Top-Middle row (variance of μ_A) from left to right: pcn-MCMC, IR-ES approximations with $\rho = 0.8$, $N_e = 75$ and $\tau = 1.1$, $\tau = 1.2$, $\tau = 1.25$, $\tau = 1.3$, $\tau = 1.6$. Middle-bottom and bottom row: Box plots of water rates from production wells P_2, P_7, P_9 and PBH from injections wells I_1, I_2, I_4 after 6 years of water flood simulated from μ_A (with MCMC) and the ensemble approximation with IR-ES for $\rho = 0.8$ and different choices of τ

dramatic reduction in the computational cost with a reasonable increase in the accuracy of the ensemble approximation. This is particularly clear for the case with larger ρ which, as stated earlier, corresponds to smaller increments in the regularizing LM scheme and so the ensemble updates is better approximated by the updates of the measurement predictions. Finally, we notice that, in contrast to IR-enLM, the performance of IR-ES was similar for the both approximation of μ_A and μ_B . In fact, note that the approximation of the variance of μ_B was better than the one of μ_A .

4.5 Comparison of the proposed methods versus unregularized methods

In this section we compare the performance of IR-enLM and IR-ES with some standard unregularized methods. In particular, for different ensemble sizes N_e , in Table 6 we compare the proposed IR-enLM with a implementation of RML where each ensemble is computed with the unregularized LM algorithm described in subsection 2.6. For the latter we consider the stopping criteria (37) with $\varepsilon_0 = 10^{-3}$ and $\varepsilon_1 = 10^{-2}$. In addition, we use the standard recommendations [30, 33, 6] where the selection of

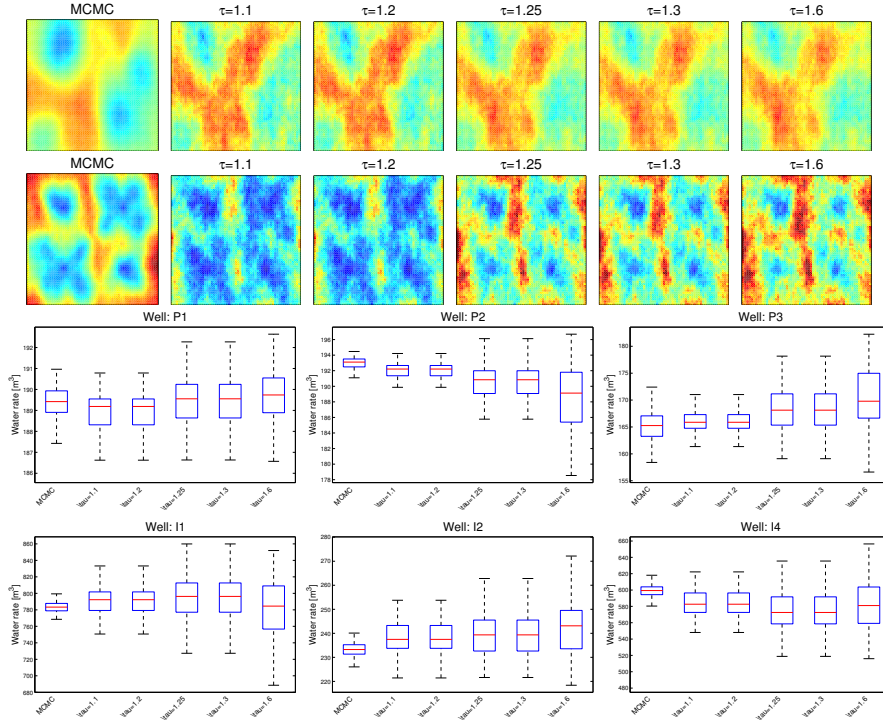


Fig. 7 Top row (mean of μ_B) and Top-Middle row (variance of μ_B) from left to right: pcn-MCMC, IR-ES approximations with $\rho = 0.8$, $N_e = 75$ and $\tau = 1.1$, $\tau = 1.2$, $\tau = 1.25$, $\tau = 1.3$, $\tau = 1.6$. Middle-bottom and bottom row: Box plots of water rates from production wells P_1, P_2, P_3 and PBH from injections wells I_1, I_2, I_4 after 6 years of water flood simulated from μ_B (with MCMC) and the ensemble approximation with IR-ES for $\rho = 0.8$ and different choices of τ

λ is given by (36) with $\lambda_0 \equiv \Lambda_0 = J(u_0^{(j)})/N_D$ and $\kappa = 10$. For the IR-enLM results from Table 6, different choices ρ where consider with a selection of $\rho \geq 0.7$ and $\tau = 1.0$ suggested from our discussion of subsection 4.3. Table 6 shows clearly that the proposed methods, with our recommendations for the selection of the tunable parameters, outperformed RML (with the unregularized LM method of [2]) in terms of approximating mean and variance of the posteriors μ_A and μ_B . For each ensemble size, the IR-enLM method provides a similar level of approximation of the mean and variance for $\rho = 0.7, 0.8, 0.9$. However, as discussed earlier, the computational cost increases significantly as we increase ρ . Visual comparisons can be conducted from Figure 8 (for μ_A) and Figure 10 (for μ_B).

In Table 7 we present a comparison of the proposed IR-ES with the standard unregularized ES implementation described in subsection 3.1. For a fixed choice $M_{ES} = 10$, we use different choices of the ensemble size N_e and tunable parameter ρ . From our discussion of subsection 4.4, for each ρ , the IR-ES results from Table 7 correspond to the tunable parameter $\tau = 1/\rho$. For each ensemble size, the IR-ES produces a better approximation that the corresponding standard ES. How-

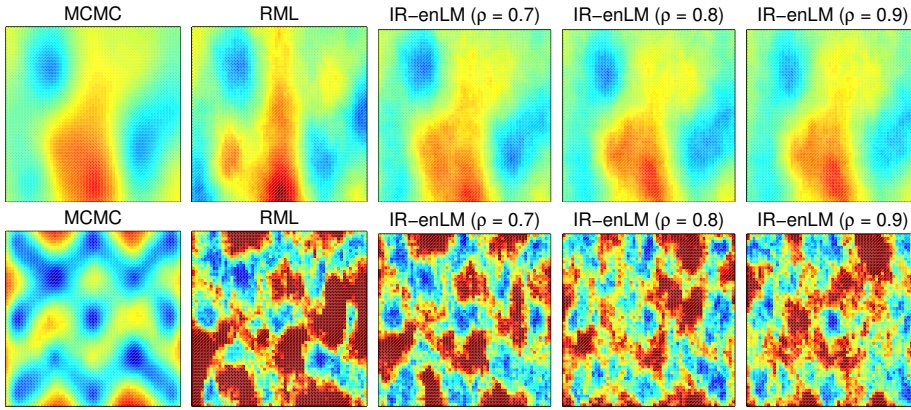


Fig. 8 Mean (top) and Variance (bottom) of the posterior distribution μ_A (characterized with MCMC) and ensemble approximations RML and IR-enLM with $N_e = 50$.

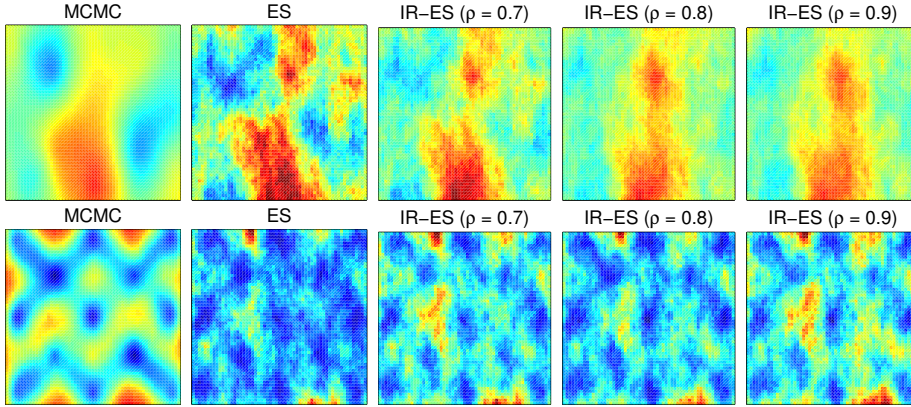


Fig. 9 Mean (top) and Variance (bottom) of the posterior distribution μ_A (characterized with MCMC) and ensemble approximations ES and IR-ES with $N_e = 75$.

ever, similar to the IR-enLM, the computational cost of IR-ES increases with ρ . For some of these ρ 's, we may be able to use an ensemble size N_e for which the cost of ES matches the one of IR-ES but with increased accuracy. For example, the computational cost of IR-ES for $\rho = 0.9$ and $N_e = 25$ is 75 forward model runs which matches the cost of ES with $N_e = 75$. In this case, ES provides more accurate approximations of the mean than our proposed method. However, the computational cost of IR-ES for $\rho = 0.7$ and $N_e = 75$ is also 75 forward model runs and provides more accurate results both in terms of mean and variance. Therefore, even though IR-ES provides better approximations for a given ensemble size, the selection of ρ is fundamental for the computational efficiency of the method. Table 7 suggests that $\rho = 0.7$, $M_{ES} = 10$, offers a reasonable compromise between accuracy and cost for each ensemble size. Some visual comparisons are displayed in Figure 9 (for μ_A) and Figure 11 (for μ_B).

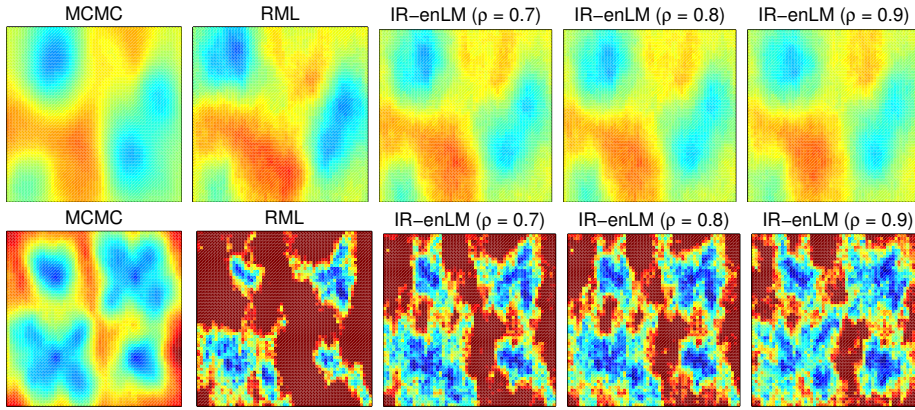


Fig. 10 Mean (top) and Variance (bottom) of the posterior distribution μ_B (characterized with MCMC) and ensemble approximations RML, IR-enLM with $N_e = 50$.

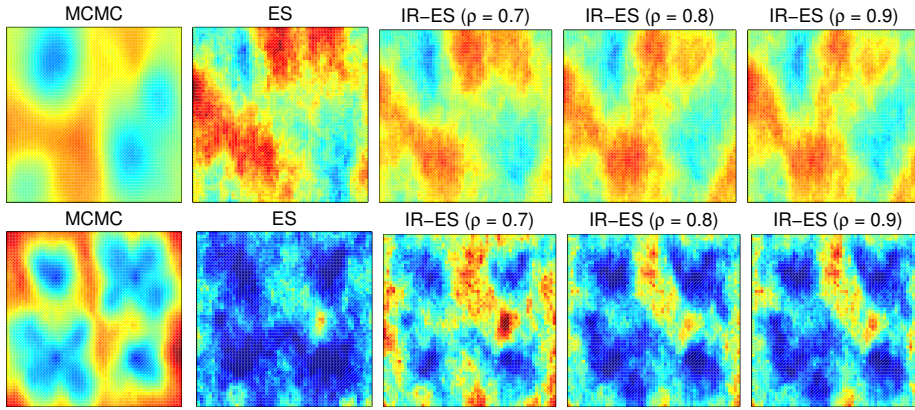


Fig. 11 Mean (top) and Variance (bottom) of the posterior distribution μ_B (characterized with MCMC) and ensemble approximations ES and IR-ES with $N_e = 75$.

In Table 8 we repeat the experiments (applying RLM with unregularized LM) for different choices of λ_0 and different choices of κ in (36). Note that the standard recommendations of parameters $\lambda_0 = J(u_0^{(j)})/N_D$ and $\kappa = 10$ (that we reported above) corresponds to the most inaccurate approximations of the posterior. Note that the selection of the standard methods are based on the assumption that the computation of the minimizer of (10) is well posed. However, in [22] we have demonstrated that this is not necessarily the case and here we clearly observe that this instabilities in the computation of (10) are detrimental to the accuracy in terms of approximating the mean and variance of the posterior. Note, for example, that changing from $\lambda_0 = J(u_0^{(j)})/N_D \times 10$ to $\lambda_0 = J(u_0^{(j)})/N_D$ has a dramatic (unstable) increase of the error in the mean and variance of the posterior. In contrast, a different selection of the parameters λ and κ may produce more accurate approximations of

the posterior. Interestingly, for the approximation of the posterior μ_A , the choice $\lambda_0 = J(u_0^{(j)})/N_D \times 10^3$, $\kappa = 10$ in RML lead to the most accurate approximations of the mean and variance of μ_A ; even more accurate than the ones provided by the proposed IR-enLM (see Table 6). This suggest that indeed, standard recommendations for the application of the LM scheme for RML need to be revisited in order to improve the approximation properties of RML. Moreover, we reiterate that the approximation of the posterior obtained with the proposed IR-enLM is the same as the one obtained with RML. Further comparisons of the sampling properties of IR-enLM and RML need to be conducted. Nonetheless, our numerical results indeed suggest that the proper selection of regularization parameters may have a significant effect on the approximation properties of RML.

We observe that the criteria (38) often used to assess the performance of ensemble methods in terms of data mismatch could be misleading. In Figure 12 we present the box plots of the value of the normalized (by the number of measurements) objective function (10), evaluated for the ensemble members (for a fixed ensemble) and for different parameters in the RML implementation described above. We note that the standard recommendations for the choices of parameters lead to ensembles that satisfy (38). However, each of these ensembles provides different and even quite inaccurate approximations of the posterior (see Table 8). In Figure 13 we display the boxplot of the same function evaluated at the ensemble generated with IR-enLM (top row) and IR-ES (bottom) for the approximation of μ_A (left) and μ_B (right), respectively. It is clear that our recommendations for τ corresponds to small values of the normalized objective functions. However, even though the value of the normalized objective function with the proposed methods was not as small as the one obtained with RML (see Figure 12), the approximation of the posterior provided by the proposed ensemble methods was, in general, more accurate than the one with RML (computed with the unregularized LM method of [2]).

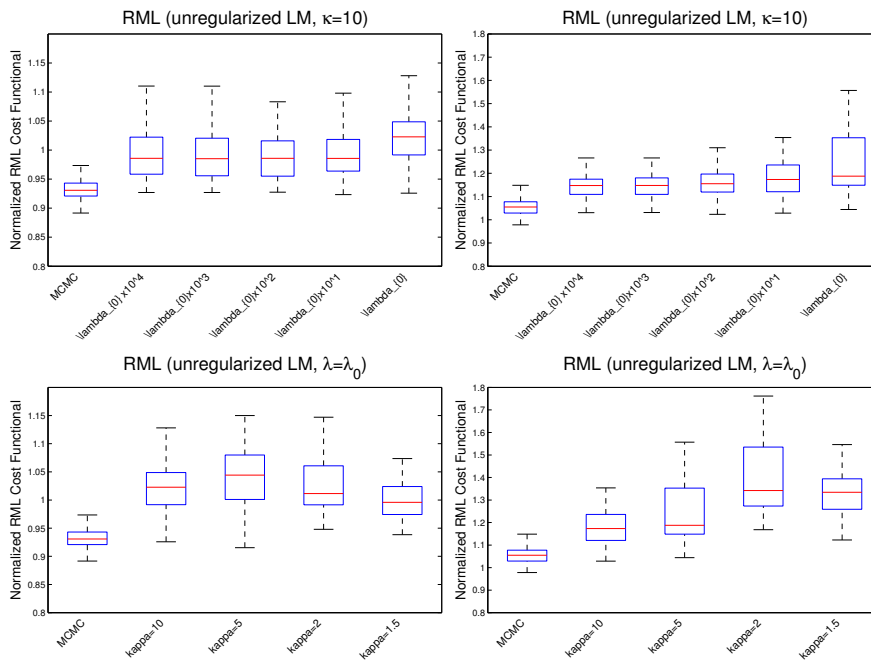


Fig. 12 Box plots of the normalized data mismatch simulated from μ_A (left) and μ_B (right) with MCMC and RML (with unregularized LM) for $N_e = 50$ and different choices of the regularization parameters

Table 1 Performance of IR-enLM for different choices of ρ and τ (for $N_e = 50$)

| ρ | τ | Approx. of μ_A | | | Approx. of μ_B | | |
|------------|--------------|--------------------|----------------------|---------------|--------------------|----------------------|---------------|
| | | ε_u | ε_σ | aver. iter. | ε_u | ε_σ | aver. iter. |
| 0.5 | 3.000 | 0.682 | 0.788 | 2.936 | 0.636 | 0.873 | 2.229 |
| 0.5 | 2.800 | 0.669 | 0.780 | 3.024 | 0.620 | 0.872 | 2.287 |
| 0.5 | 2.500 | 0.651 | 0.756 | 3.152 | 0.598 | 0.860 | 2.464 |
| 0.5 | 2.200 | 0.596 | 0.729 | 3.384 | 0.568 | 0.860 | 2.640 |
| 0.5 | 2.000 | 0.573 | 0.718 | 3.492 | 0.550 | 0.867 | 2.742 |
| 0.5 | 1.800 | 0.488 | 0.765 | 3.644 | 0.531 | 0.919 | 2.853 |
| 0.5 | 1.600 | 0.408 | 0.791 | 3.840 | 0.516 | 0.962 | 2.978 |
| 0.5 | 1.400 | 0.341 | 0.765 | 4.100 | 0.491 | 1.012 | 3.145 |
| 0.6 | 2.200 | 0.654 | 0.755 | 4.335 | 0.620 | 0.634 | 3.575 |
| 0.6 | 2.000 | 0.631 | 0.735 | 4.534 | 0.598 | 0.613 | 3.625 |
| 0.6 | 1.900 | 0.616 | 0.720 | 4.636 | 0.587 | 0.604 | 3.650 |
| 0.6 | 1.800 | 0.594 | 0.712 | 4.754 | 0.571 | 0.591 | 3.825 |
| 0.6 | 1.700 | 0.564 | 0.702 | 4.862 | 0.543 | 0.573 | 3.975 |
| 0.6 | 1.600 | 0.515 | 0.695 | 4.990 | 0.514 | 0.581 | 4.075 |
| 0.6 | 1.500 | 0.452 | 0.716 | 5.108 | 0.489 | 0.591 | 4.200 |
| 0.6 | 1.300 | 0.328 | 0.681 | 5.428 | 0.421 | 0.628 | 4.525 |
| 0.7 | 1.500 | 0.565 | 0.673 | 7.049 | 0.536 | 0.508 | 6.122 |
| 0.7 | 1.450 | 0.544 | 0.665 | 7.138 | 0.522 | 0.498 | 6.206 |
| 0.7 | 1.400 | 0.517 | 0.651 | 7.230 | 0.498 | 0.488 | 6.320 |
| 0.7 | 1.250 | 0.379 | 0.637 | 7.590 | 0.420 | 0.487 | 6.651 |
| 0.7 | 1.150 | 0.293 | 0.572 | 7.885 | 0.381 | 0.498 | 6.869 |
| 0.7 | 1.100 | 0.269 | 0.504 | 8.059 | 0.367 | 0.495 | 6.973 |
| 0.7 | 1.000 | 0.240 | 0.405 | 8.597 | 0.335 | 0.528 | 7.335 |
| 0.7 | 0.950 | 0.221 | 0.421 | 9.425 | 0.332 | 0.595 | 7.978 |
| 0.8 | 1.500 | 0.619 | 0.706 | 10.906 | 0.569 | 0.529 | 9.789 |
| 0.8 | 1.400 | 0.594 | 0.687 | 11.248 | 0.544 | 0.499 | 10.158 |
| 0.8 | 1.300 | 0.555 | 0.648 | 11.611 | 0.511 | 0.466 | 10.541 |
| 0.8 | 1.250 | 0.518 | 0.621 | 11.842 | 0.490 | 0.447 | 10.739 |
| 0.8 | 1.200 | 0.466 | 0.586 | 12.026 | 0.463 | 0.430 | 10.938 |
| 0.8 | 1.150 | 0.387 | 0.560 | 12.216 | 0.430 | 0.425 | 11.127 |
| 0.8 | 1.000 | 0.249 | 0.386 | 13.088 | 0.329 | 0.425 | 11.755 |
| 0.8 | 0.950 | 0.214 | 0.394 | 13.945 | 0.310 | 0.479 | 12.344 |
| 0.9 | 1.160 | 0.547 | 0.594 | 25.960 | 0.491 | 0.438 | 23.907 |
| 0.9 | 1.140 | 0.524 | 0.564 | 26.150 | 0.482 | 0.430 | 24.118 |
| 0.9 | 1.120 | 0.497 | 0.538 | 26.324 | 0.472 | 0.419 | 24.331 |
| 0.9 | 1.110 | 0.482 | 0.523 | 26.425 | 0.467 | 0.415 | 24.430 |
| 0.9 | 1.100 | 0.466 | 0.510 | 26.515 | 0.462 | 0.412 | 24.510 |
| 0.9 | 1.050 | 0.375 | 0.443 | 26.949 | 0.420 | 0.386 | 25.022 |
| 0.9 | 1.000 | 0.277 | 0.371 | 27.476 | 0.362 | 0.381 | 25.511 |
| 0.9 | 0.950 | 0.217 | 0.348 | 28.364 | 0.304 | 0.419 | 26.239 |

Table 2 Performance of IR-enLM for different choices of N_e and τ (for $\rho = 0.8$)

| τ | N_e | Approx. of μ_A | | | Approx. of μ_B | | |
|----------|----------------|--------------------|----------------------|---------------|--------------------|----------------------|---------------|
| | | ε_u | ε_σ | aver. iter. | ε_u | ε_σ | aver. iter. |
| 1.3 | 25.000 | 0.577 | 0.699 | 11.776 | 0.538 | 0.548 | 10.588 |
| 1.25 | 25.000 | 0.540 | 0.668 | 11.980 | 0.519 | 0.532 | 10.784 |
| 1.2 | 25.000 | 0.502 | 0.642 | 12.120 | 0.495 | 0.515 | 10.992 |
| 1.1 | 25.000 | 0.361 | 0.568 | 12.540 | 0.430 | 0.500 | 11.340 |
| 1 | 25.000 | 0.300 | 0.449 | 13.152 | 0.372 | 0.512 | 11.792 |
| 0.95 | 25.000 | 0.267 | 0.463 | 13.936 | 0.359 | 0.546 | 12.320 |
| 1.3 | 50.000 | 0.555 | 0.648 | 11.611 | 0.511 | 0.466 | 10.541 |
| 1.25 | 50.000 | 0.518 | 0.621 | 11.842 | 0.490 | 0.447 | 10.739 |
| 1.2 | 50.000 | 0.466 | 0.586 | 12.026 | 0.463 | 0.430 | 10.938 |
| 1.1 | 50.000 | 0.315 | 0.503 | 12.444 | 0.393 | 0.422 | 11.317 |
| 1 | 50.000 | 0.249 | 0.386 | 13.088 | 0.329 | 0.425 | 11.755 |
| 0.95 | 50.000 | 0.214 | 0.394 | 13.945 | 0.310 | 0.479 | 12.344 |
| 1.3 | 75.000 | 0.545 | 0.617 | 11.569 | 0.502 | 0.437 | 10.551 |
| 1.25 | 75.000 | 0.506 | 0.585 | 11.768 | 0.480 | 0.416 | 10.761 |
| 1.2 | 75.000 | 0.455 | 0.549 | 11.945 | 0.455 | 0.402 | 10.947 |
| 1.1 | 75.000 | 0.300 | 0.467 | 12.357 | 0.383 | 0.387 | 11.325 |
| 1 | 75.000 | 0.231 | 0.339 | 12.988 | 0.319 | 0.386 | 11.767 |
| 0.95 | 75.000 | 0.196 | 0.357 | 13.920 | 0.300 | 0.438 | 12.345 |
| 1.3 | 100.000 | 0.544 | 0.612 | 11.602 | 0.496 | 0.423 | 10.541 |
| 1.25 | 100.000 | 0.507 | 0.585 | 11.833 | 0.475 | 0.403 | 10.739 |
| 1.2 | 100.000 | 0.451 | 0.548 | 12.021 | 0.448 | 0.384 | 10.938 |
| 1.1 | 100.000 | 0.295 | 0.459 | 12.441 | 0.374 | 0.376 | 11.317 |
| 1 | 100.000 | 0.226 | 0.336 | 13.084 | 0.307 | 0.378 | 11.755 |
| 0.95 | 100.000 | 0.188 | 0.346 | 13.960 | 0.285 | 0.432 | 12.344 |
| 1.3 | 250.000 | 0.535 | 0.580 | 11.569 | 0.487 | 0.394 | 10.541 |
| 1.25 | 250.000 | 0.496 | 0.548 | 11.768 | 0.466 | 0.374 | 10.739 |
| 1.2 | 250.000 | 0.444 | 0.512 | 11.945 | 0.437 | 0.355 | 10.938 |
| 1.1 | 250.000 | 0.284 | 0.425 | 12.357 | 0.363 | 0.347 | 11.317 |
| 1 | 250.000 | 0.210 | 0.290 | 12.988 | 0.293 | 0.348 | 11.755 |
| 0.95 | 250.000 | 0.171 | 0.308 | 13.920 | 0.270 | 0.405 | 12.344 |

Table 3 Performance of IR-ES for different choices of τ and N_e (for $\rho = 0.8$ and $M_{ES} = 10$)

| τ | N_e | Approx. of μ_A | | | Approx. of μ_B | | |
|-------------|----------------|--------------------|----------------------|---------------|--------------------|----------------------|---------------|
| | | ε_u | ε_σ | aver. iter. | ε_u | ε_σ | aver. iter. |
| 1.6 | 25.000 | 1.001 | 0.668 | 11.500 | 0.932 | 0.508 | 12.050 |
| 1.4 | 25.000 | 1.051 | 0.767 | 12.250 | 1.004 | 0.585 | 12.750 |
| 1.3 | 25.000 | 1.107 | 0.803 | 12.550 | 1.090 | 0.664 | 13.150 |
| 1.25 | 25.000 | 1.144 | 0.824 | 12.700 | 1.210 | 0.718 | 13.450 |
| 1.2 | 25.000 | 1.184 | 0.845 | 12.900 | 1.311 | 0.799 | 14.050 |
| 1.6 | 50.000 | 0.717 | 0.325 | 10.950 | 0.702 | 0.279 | 11.700 |
| 1.5 | 50.000 | 0.711 | 0.341 | 11.200 | 0.699 | 0.278 | 12.050 |
| 1.4 | 50.000 | 0.696 | 0.352 | 11.550 | 0.700 | 0.285 | 12.300 |
| 1.3 | 50.000 | 0.684 | 0.396 | 12.000 | 0.718 | 0.335 | 12.750 |
| 1.25 | 50.000 | 0.685 | 0.421 | 12.150 | 0.742 | 0.375 | 13.000 |
| 1.2 | 50.000 | 0.714 | 0.460 | 12.400 | 0.791 | 0.419 | 13.200 |
| 1.5 | 75.000 | 0.651 | 0.295 | 10.700 | 0.618 | 0.247 | 11.550 |
| 1.4 | 75.000 | 0.635 | 0.280 | 11.050 | 0.606 | 0.223 | 12.050 |
| 1.3 | 75.000 | 0.600 | 0.267 | 11.500 | 0.609 | 0.234 | 12.450 |
| 1.25 | 75.000 | 0.583 | 0.283 | 11.700 | 0.611 | 0.242 | 12.600 |
| 1.2 | 75.000 | 0.575 | 0.310 | 11.900 | 0.640 | 0.291 | 13.000 |
| 1.1 | 75.000 | 0.598 | 0.366 | 12.300 | 0.672 | 0.316 | 13.150 |
| 1.5 | 100.000 | 0.614 | 0.316 | 10.500 | 0.563 | 0.259 | 11.550 |
| 1.4 | 100.000 | 0.586 | 0.268 | 10.900 | 0.547 | 0.222 | 12.050 |
| 1.3 | 100.000 | 0.558 | 0.236 | 11.250 | 0.541 | 0.208 | 12.350 |
| 1.25 | 100.000 | 0.544 | 0.235 | 11.400 | 0.539 | 0.204 | 12.500 |
| 1.2 | 100.000 | 0.540 | 0.238 | 11.550 | 0.549 | 0.215 | 12.850 |
| 1.1 | 100.000 | 0.540 | 0.278 | 12.050 | 0.601 | 0.249 | 13.200 |
| 1.4 | 250.000 | 0.534 | 0.407 | 10.267 | 0.469 | 0.283 | 11.714 |
| 1.3 | 250.000 | 0.454 | 0.251 | 11.067 | 0.452 | 0.246 | 12.071 |
| 1.25 | 250.000 | 0.454 | 0.251 | 11.067 | 0.448 | 0.234 | 12.214 |
| 1.2 | 250.000 | 0.437 | 0.237 | 11.133 | 0.440 | 0.186 | 12.571 |
| 1.1 | 250.000 | 0.389 | 0.178 | 11.667 | 0.450 | 0.159 | 12.857 |

Table 4 Performance of IR-ES for different choices of ρ and τ (for $M_{ES} = 10$ and $N_e = 75$)

| ρ | τ | Approx. of μ_A | | | Approx. of μ_B | | |
|------------|--------------|--------------------|----------------------|---------------|--------------------|----------------------|---------------|
| | | ε_u | ε_σ | aver. iter. | ε_u | ε_σ | aver. iter. |
| 0.5 | 3.500 | 0.766 | 0.610 | 3.000 | 0.716 | 0.480 | 3.500 |
| 0.5 | 3.000 | 0.719 | 0.466 | 3.550 | 0.693 | 0.408 | 3.850 |
| 0.5 | 2.800 | 0.695 | 0.397 | 3.800 | 0.684 | 0.388 | 3.950 |
| 0.5 | 2.200 | 0.681 | 0.344 | 4.000 | 0.689 | 0.368 | 4.050 |
| 0.5 | 2.000 | 0.681 | 0.344 | 4.000 | 0.741 | 0.339 | 4.400 |
| 0.5 | 1.800 | 0.737 | 0.353 | 4.200 | 0.817 | 0.331 | 4.650 |
| 0.5 | 1.600 | 0.806 | 0.366 | 4.600 | 0.909 | 0.346 | 4.900 |
| 0.6 | 2.200 | 0.679 | 0.362 | 5.150 | 0.685 | 0.358 | 5.000 |
| 0.6 | 2.000 | 0.676 | 0.345 | 5.250 | 0.687 | 0.333 | 5.200 |
| 0.6 | 1.900 | 0.668 | 0.314 | 5.400 | 0.686 | 0.326 | 5.250 |
| 0.6 | 1.800 | 0.666 | 0.307 | 5.500 | 0.709 | 0.299 | 5.550 |
| 0.6 | 1.700 | 0.666 | 0.292 | 5.700 | 0.749 | 0.290 | 5.850 |
| 0.6 | 1.600 | 0.673 | 0.296 | 5.900 | 0.769 | 0.291 | 6.000 |
| 0.6 | 1.500 | 0.695 | 0.303 | 5.950 | 0.769 | 0.291 | 6.000 |
| 0.6 | 1.400 | 0.714 | 0.318 | 6.050 | 0.769 | 0.291 | 6.000 |
| 0.7 | 1.700 | 0.659 | 0.292 | 7.550 | 0.693 | 0.292 | 8.100 |
| 0.7 | 1.600 | 0.655 | 0.279 | 7.750 | 0.698 | 0.281 | 8.250 |
| 0.7 | 1.500 | 0.657 | 0.280 | 8.000 | 0.723 | 0.280 | 8.500 |
| 0.7 | 1.450 | 0.657 | 0.284 | 8.100 | 0.759 | 0.290 | 8.800 |
| 0.7 | 1.420 | 0.669 | 0.296 | 8.250 | 0.763 | 0.293 | 8.850 |
| 0.7 | 1.350 | 0.684 | 0.302 | 8.300 | 0.770 | 0.296 | 8.900 |
| 0.7 | 1.300 | 0.733 | 0.324 | 8.450 | 0.815 | 0.323 | 9.150 |
| 0.8 | 1.400 | 0.635 | 0.280 | 11.050 | 0.606 | 0.223 | 12.050 |
| 0.8 | 1.350 | 0.619 | 0.272 | 11.250 | 0.604 | 0.224 | 12.150 |
| 0.8 | 1.300 | 0.600 | 0.267 | 11.500 | 0.609 | 0.234 | 12.450 |
| 0.8 | 1.250 | 0.583 | 0.283 | 11.700 | 0.611 | 0.242 | 12.600 |
| 0.8 | 1.200 | 0.575 | 0.310 | 11.900 | 0.640 | 0.291 | 13.000 |
| 0.8 | 1.100 | 0.598 | 0.366 | 12.300 | 0.672 | 0.316 | 13.150 |
| 0.9 | 1.400 | 0.629 | 0.299 | 23.800 | 0.581 | 0.238 | 26.200 |
| 0.9 | 1.300 | 0.599 | 0.265 | 24.600 | 0.582 | 0.220 | 26.900 |
| 0.9 | 1.200 | 0.563 | 0.244 | 25.550 | 0.590 | 0.206 | 27.900 |
| 0.9 | 1.150 | 0.551 | 0.242 | 26.000 | 0.600 | 0.209 | 28.400 |
| 0.9 | 1.100 | 0.540 | 0.260 | 26.650 | 0.628 | 0.229 | 29.000 |
| 0.9 | 1.000 | 0.608 | 0.309 | 27.750 | 0.799 | 0.419 | 30.500 |

Table 5 Performance of IR-ES for different choices of M_{ES} and ρ (for $\tau = 1/\rho$ and $N_e = 75$)

| ρ | M_{ES} | Approx. of μ_A | | | Approx. of μ_B | | |
|--------|----------|--------------------|----------------------|-------------------------|--------------------|----------------------|-------------------------|
| | | ε_u | ε_σ | aver. cost (model runs) | ε_u | ε_σ | aver. cost (model runs) |
| 0.7 | 1 | 0.598 | 0.297 | 652.500 | 0.576 | 0.278 | 675.000 |
| 0.7 | 5 | 0.634 | 0.308 | 150 | 0.594 | 0.250 | 150 |
| 0.7 | 10 | 0.654 | 0.293 | 75 | 0.718 | 0.287 | 75 |
| 0.8 | 1 | 0.547 | 0.255 | 918.750 | 0.588 | 0.205 | 1057.500 |
| 0.8 | 5 | 0.575 | 0.277 | 225 | 0.583 | 0.222 | 225 |
| 0.8 | 10 | 0.578 | 0.294 | 150 | 0.610 | 0.246 | 150 |
| 0.8 | 15 | 0.668 | 0.307 | 75 | 0.751 | 0.279 | 75 |

Table 6 Comparison of methods. *: ($\tau = 1.0$), †: ($\kappa = 10$, $\lambda_0 = \Lambda_0$)

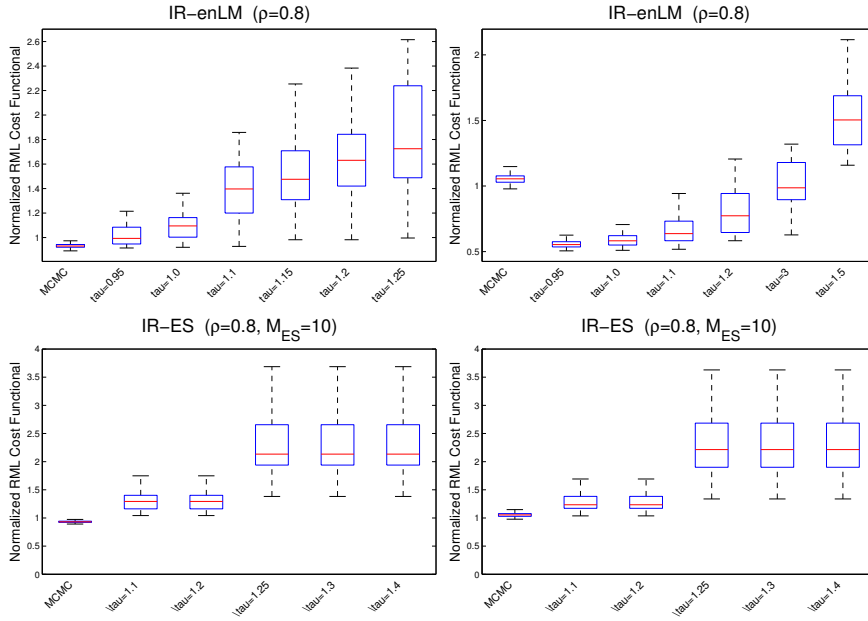
| Method | N_e | Approx. of μ_A | | | Approx. of μ_B | | |
|---------------------------|-------|--------------------|-------------------|-------------|--------------------|-------------------|-------------|
| | | ϵ_u | ϵ_σ | aver. iter. | ϵ_u | ϵ_σ | aver. iter. |
| RML [†] | 25 | 0.488 | 0.702 | 10.560 | 0.611 | 1.355 | 10.232 |
| RML [†] | 50 | 0.466 | 0.637 | 10.350 | 0.575 | 1.205 | 10.580 |
| RML [†] | 75 | 0.456 | 0.596 | 10.345 | 0.556 | 1.180 | 10.636 |
| RML [†] | 100 | 0.452 | 0.603 | 10.311 | 0.542 | 1.156 | 10.675 |
| IR-enLM* ($\rho = 0.7$) | 25 | 0.295 | 0.467 | 8.616 | 0.381 | 0.640 | 7.364 |
| IR-enLM* ($\rho = 0.8$) | 25 | 0.300 | 0.449 | 13.152 | 0.372 | 0.512 | 11.792 |
| IR-enLM* ($\rho = 0.9$) | 25 | 0.321 | 0.452 | 27.796 | 0.395 | 0.469 | 25.344 |
| IR-enLM ($\rho = 0.7$) | 50 | 0.240 | 0.405 | 8.597 | 0.335 | 0.528 | 7.335 |
| IR-enLM ($\rho = 0.8$) | 50 | 0.249 | 0.386 | 13.088 | 0.329 | 0.425 | 11.755 |
| IR-enLM ($\rho = 0.9$) | 50 | 0.277 | 0.371 | 27.476 | 0.362 | 0.381 | 25.511 |
| IR-enLM ($\rho = 0.7$) | 75 | 0.224 | 0.365 | 8.523 | 0.324 | 0.500 | 7.347 |
| IR-enLM ($\rho = 0.8$) | 75 | 0.231 | 0.339 | 12.988 | 0.319 | 0.386 | 11.767 |
| IR-enLM ($\rho = 0.9$) | 75 | 0.265 | 0.341 | 27.461 | 0.347 | 0.351 | 25.443 |
| IR-enLM ($\rho = 0.7$) | 100 | 0.217 | 0.361 | 8.597 | 0.312 | 0.481 | 7.335 |
| IR-enLM ($\rho = 0.8$) | 100 | 0.226 | 0.336 | 13.084 | 0.307 | 0.378 | 11.755 |
| IR-enLM ($\rho = 0.9$) | 100 | 0.258 | 0.327 | 27.476 | 0.343 | 0.332 | 25.511 |

Table 7 Comparison of methods. *: ($M_{ES} = 10$, $\tau = 1/\rho$), †: in forward model runs

| Method | N_e | Approx. of μ_A | | | Approx. of μ_B | | |
|-------------------------|-------|--------------------|-------------------|---------------------------|--------------------|-------------------|---------------------------|
| | | ϵ_u | ϵ_σ | average cost [†] | ϵ_u | ϵ_σ | average cost [†] |
| ES | 25 | 1.617 | 0.913 | 25 | 1.863 | 0.866 | 25 |
| ES | 50 | 1.185 | 0.617 | 50 | 1.472 | 0.581 | 50 |
| ES | 75 | 0.914 | 0.420 | 75 | 1.071 | 0.412 | 75 |
| ES | 100 | 0.779 | 0.310 | 100 | 0.969 | 0.306 | 100 |
| ES | 150 | 0.630 | 0.220 | 150 | 0.795 | 0.234 | 150 |
| ES | 200 | 0.555 | 0.193 | 200 | 0.701 | 0.226 | 200 |
| IR-ES* ($\rho = 0.7$) | 25 | 1.522 | 0.869 | 25 | 1.520 | 0.741 | 25 |
| IR-ES* ($\rho = 0.8$) | 25 | 1.107 | 0.803 | 50 | 1.090 | 0.664 | 50 |
| IR-ES* ($\rho = 0.9$) | 25 | 0.969 | 0.672 | 75 | 1.053 | 0.579 | 75 |
| IR-ES* ($\rho = 0.7$) | 50 | 0.789 | 0.404 | 50 | 0.942 | 0.374 | 50 |
| IR-ES* ($\rho = 0.8$) | 50 | 0.684 | 0.396 | 100 | 0.718 | 0.335 | 100 |
| IR-ES* ($\rho = 0.9$) | 50 | 0.637 | 0.321 | 150 | 0.689 | 0.272 | 150 |
| IR-ES* ($\rho = 0.7$) | 75 | 0.657 | 0.280 | 75 | 0.723 | 0.280 | 75 |
| IR-ES* ($\rho = 0.8$) | 75 | 0.600 | 0.267 | 150 | 0.609 | 0.234 | 150 |
| IR-ES* ($\rho = 0.9$) | 75 | 0.563 | 0.244 | 225 | 0.590 | 0.206 | 225 |
| IR-ES* ($\rho = 0.7$) | 100 | 0.586 | 0.267 | 100 | 0.659 | 0.290 | 100 |
| IR-ES* ($\rho = 0.8$) | 100 | 0.558 | 0.236 | 200 | 0.541 | 0.208 | 200 |
| IR-ES* ($\rho = 0.9$) | 100 | 0.525 | 0.232 | 300 | 0.525 | 0.206 | 300 |

Table 8 Performance of RML (with the unregularized LM method of [2]) for different choices of λ_0 and κ .

| λ_0 | κ | N_e | Approx. of μ_A | | | Approx. of μ_B | | |
|-------------------------|----------|-------|--------------------|-------------------|-------------|--------------------|-------------------|-------------|
| | | | ϵ_u | ϵ_σ | aver. iter. | ϵ_u | ϵ_σ | aver. iter. |
| $\Lambda_0 \times 10^5$ | 10 | 50 | 0.864 | 1.264 | 3.363 | 0.973 | 0.919 | 2.077 |
| $\Lambda_0 \times 10^4$ | 10 | 50 | 0.210 | 0.288 | 12.100 | 0.352 | 0.484 | 13.335 |
| $\Lambda_0 \times 10^3$ | 10 | 50 | 0.209 | 0.283 | 11.119 | 0.361 | 0.448 | 12.664 |
| $\Lambda_0 \times 10^2$ | 10 | 50 | 0.212 | 0.284 | 10.049 | 0.362 | 0.456 | 11.665 |
| $\Lambda_0 \times 10$ | 10 | 50 | 0.246 | 0.305 | 8.876 | 0.393 | 0.553 | 10.668 |
| Λ_0 | 10 | 50 | 0.464 | 0.626 | 10.345 | 0.567 | 1.209 | 10.636 |
| Λ_0 | 5 | 50 | 0.408 | 0.539 | 10.403 | 0.537 | 1.130 | 11.293 |
| Λ_0 | 2 | 50 | 0.200 | 0.272 | 11.919 | 0.510 | 1.045 | 12.376 |
| Λ_0 | 1.5 | 50 | 0.201 | 0.267 | 13.915 | 0.502 | 1.012 | 12.613 |

**Fig. 13** Box plots of the normalized data mismatch simulated from μ_A (left) and μ_B (right) with MCMC and the proposed methods with $N_e = 50$ and for different choices of the regularization parameters. Top-row: ensemble approximation with IR-enLM for $\rho = 0.8$ and different choices of τ . Bottom-row: ensemble approximation with IR-ES for $\rho = 0.8$, $M_{ES} = 10$ and different choices of τ

4.6 IR-ES and ES-MDA

In this subsection we provide a comparison of the proposed IR-ES with the ES-MDA of [12, 13, 11]. As we discussed in the previous sections both methods have similar features and the numerical comparison may provide further understanding of differences and similarities between these two approaches. We consider the same experiments as before where we aim at providing an ensemble approximation of the mean and variance of the posteriors μ_A and μ_B that we characterize with MCMC. For both methods we consider an ensemble size of $N_e = 100$. For IR-ES we select $M_{ES} = 1$ in (76) which corresponds to evaluating the forward model for each ensemble at each iteration of the scheme. In addition, we consider $\rho = 0.7$, $\rho = 0.8$ and the selection of $\tau = 1/\rho$ as suggested by the theory described earlier. For these parameters we consider 80 runs of IR-ES corresponding to different initial ensembles drawn from the prior distribution previously described. For each of these experiments, IR-ES will terminate after a number of iterations determined by the stopping criteria (63). The number of iterations needed to achieve convergence by IR-ES is then used as the number of iterations for our implementation of ES-MDA. Thus, for each experiment, both IR-ES and ES-MDA were initialized with the same initial ensemble and run the same number of iterations; thus both have approximately the same computational cost for each experiment. In Figure 14 we display the number of iterations needed for IR-ES (for two choices of ρ) to achieve convergence.

We recall that the regularization parameter $\alpha_{m,MDA}$ in ES-MDA is typically selected beforehand provided that $\sum_{m=1}^n \alpha_{m,MDA}^{-1} = 1$ [12]. For our comparison we consider two selections of regularization parameters in ES-MDA. In choice (a) we simply take $\alpha_{m,MDA} = n$ where n is the number of iterations. In choice (b) we define $\alpha_{m,MDA} = (\sum_{m=1}^n \alpha_m^{-1}) \alpha_m$ where α_m is the one obtained from IR-ES. We recall from earlier discussions that this α_m 's decrease with the number of iterations. More concretely, for each experiment (i.e. for a given initial ensemble) we run IR-ES, save the α_m 's that we obtain according to (68) for the n number of iterations needed to achieve convergence in terms of (63). Then, we run ES-MDA with this n number of iterations and the aforementioned selection of the normalized α_m 's from IR-ES. This selection of $\alpha_{m,MDA}$ is motivated by recent research showing that selecting the parameters adaptively can further improve the approximation properties of ES-MDA [31]. Here the aim is to investigate the potential application of the methodology of IR-ES in combination with the ES-MDA framework.

For each experiment corresponding to a different initial ensemble, in Figure 15 we show the relative error of the mean (top-left panel) and variance (top-right panel) (see expressions (79)) obtained with IR-ES and ES-MDA (for the two selections of $\alpha_{m,ESMDA}$) with respect to the mean and variance of the posterior μ_A . In Figure 16 we present the normalized data mismatch for 10 ensembles obtained with IR-ES and ES-MDA (a), (b). The analogous plots for the posterior μ_B are displayed in Figure 17 and Figure 18. We reiterate that these relative errors with respect to the mean and variance of the posteriors μ_A and μ_B characterized with MCMC. For the methods considered in this subsection, increasing the number of iterations did not have a significant effect on the resulting accuracy with respect to the mean and variance of the posterior. In addition, we can clearly observe that ES-MDA (a) (i.e. with uniform regularization

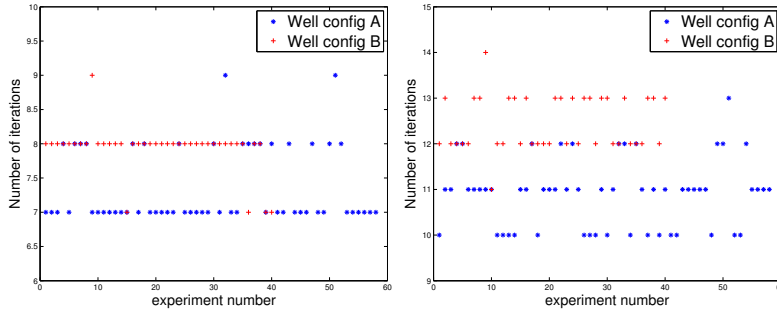


Fig. 14 Number of iterations to achieve convergence of IR-ES with $\rho = 0.7$ (left) and $\rho = 0.8$ (right)

parameters) provides a poor approximation of the mean and variance even though it provides a reduction of the data mismatch comparable to the one achieved with MCMC. The application of ES-MDA with the normalized regularization parameters from IR-ES shows substantial improvement in terms of decreasing mean and variance. However, the data mismatch was not decreased as in case (a). Improving ES-MDA by selecting decreasing regularization parameters have been also observed in [12, 11]. While IR-ES and ES-MDA (b) provided similar approximations to the mean of both μ_A and μ_B , IR-ES clearly outperformed ES-MDA (b) in terms of approximating their variances. On the other hand, the data mismatch of IR-ES did not provide a reduction of the data mismatch consistent with MCMC (or as good as ES-MDA (b)) which may be a potentially disadvantage of the proposed technique. Similar to our observations related to the LM implementation of [2] for RML, here we observe that some ES-MDA implementations may successfully reduce the data mismatch without necessarily producing accurate approximations of the posterior. Clearly, further research should be conducted to develop optimal selection of regularization parameters.

5 Conclusions

The proposed IR-enLM and IR-ES are ensemble methods within the RML and Kalman-based frameworks, respectively. The aim of these methods is to provide useful information of the posterior. Both IR-enLM and IR-ES can be derived as an ensemble of iterative solutions to linear inverse problems posed as the minimization of a Tikhonov-regularized functional. The novel feature of the proposed methods is the use of the discrepancy principle for the selection of the Tikhonov parameter that adaptively changes as the iteration progresses. Additionally, the aforementioned principle is applied to control the early termination of the proposed ensemble methods. Crucial to the proposed methods is a set of tunable parameters that arise from the application of the discrepancy principle. In this work we provide extensive numerical experiments to demonstrate the effect of these parameters in the approximation properties of the proposed IR-enLM and IR-ES. Our numerical investigations are based on synthetic experiments where we use a state-of-the-art MCMC method to fully re-

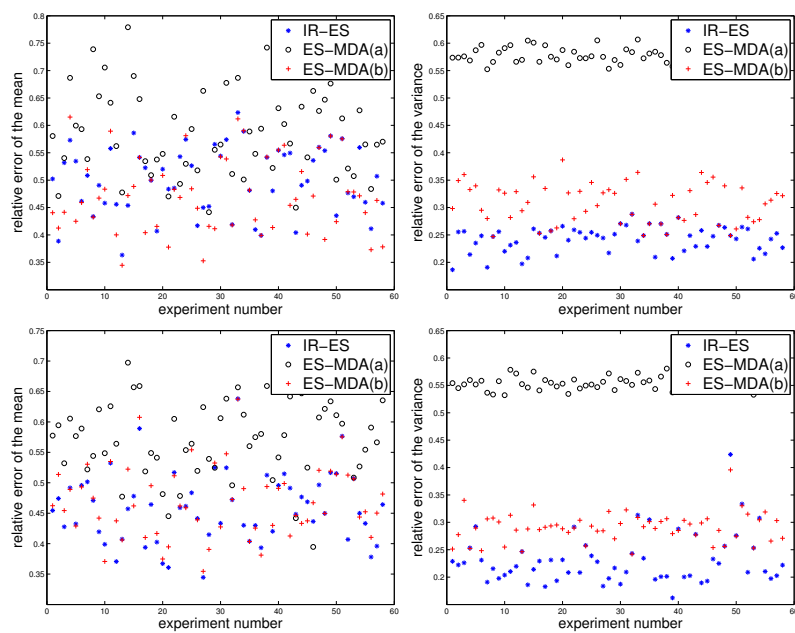


Fig. 15 Approximation of the mean and variance of μ_A . Relative error of the mean (left) and variance (right) obtained with ES-MDA and IR-ES with $\rho = 0.7$ (top) and $\rho = 0.8$ (bottom).

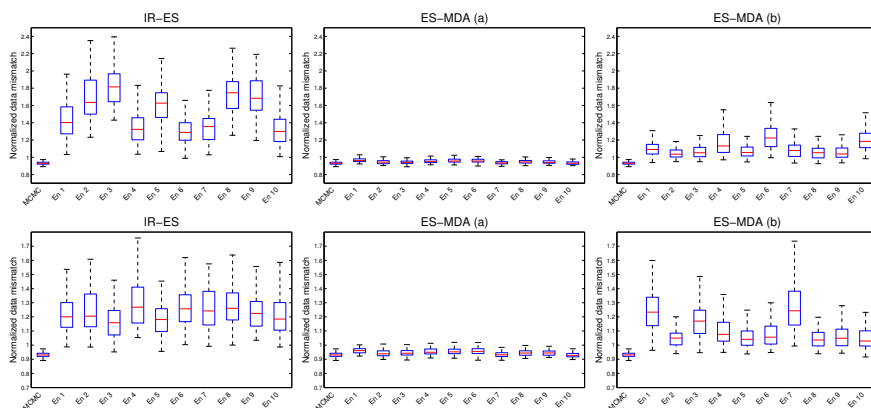


Fig. 16 Approximation of the mean and variance of μ_A . Data mismatch from 10 ensembles computed with (from left to right): IR-ES, ES-MDA (a) and ES-MDA(b). Top row (resp. bottom row) corresponds to IR-ES with $\rho = 0.7$ (resp. $\rho = 0.8$)

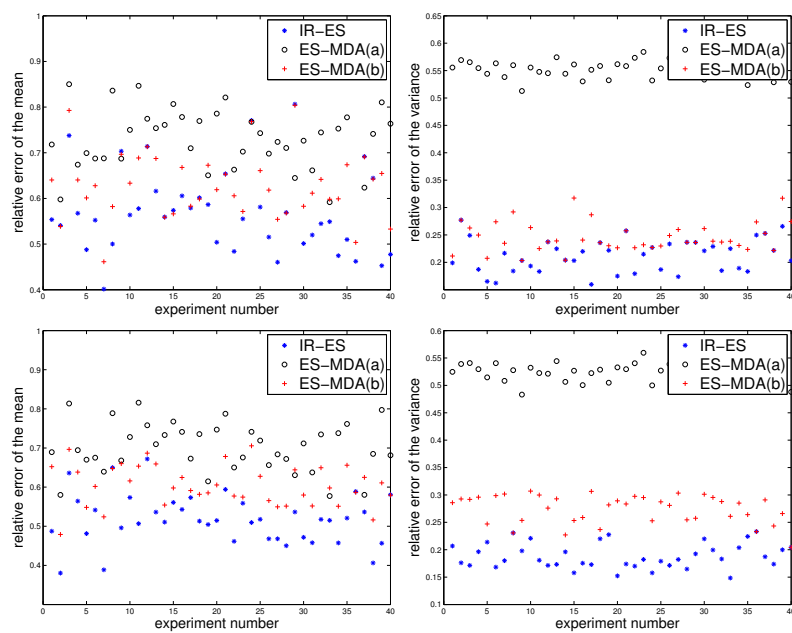


Fig. 17 Approximation of the mean and variance of μ_B . Relative error of the mean (left) and variance (right) obtained with ES-MDA and IR-ES with $\rho = 0.7$ (top) and $\rho = 0.8$ (bottom).

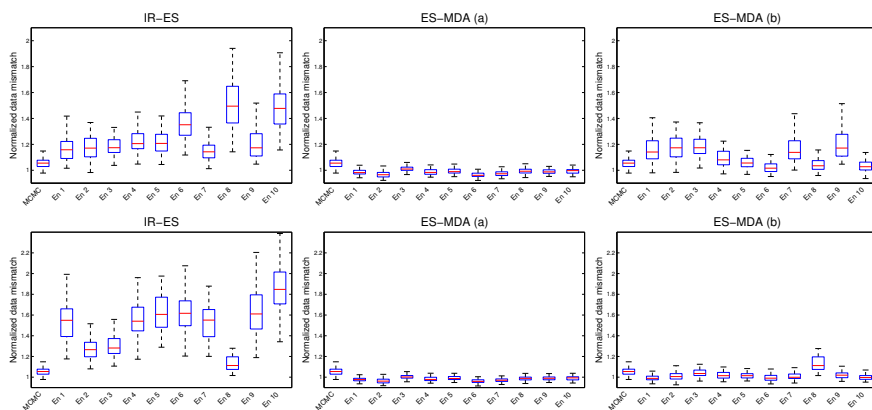


Fig. 18 Approximation of the mean and variance of μ_B . Data mismatch from 10 ensembles computed with (from left to right): IR-ES, ES-MDA (a) and ES-MDA(b). Top row (resp. bottom row) corresponds to IR-ES with $\rho = 0.7$ (resp. $\rho = 0.8$)

solve the mean and variance of the resulting Bayesian posterior. For both methods, the tunable parameter $\rho \in (0, 1)$ controls the adaptive selection of the regularization parameter which, in turn, determines the size of the increment in the iterative scheme. In unregularized LM methods for RML-based applications this regularization parameter is typically decreased by an arbitrary factor selected heuristically. In contrast, the regularization parameter in the proposed methods is adaptively chosen according to the discrepancy principle in order to stabilize the computation of each ensemble member. In addition, the size of the parameter ρ defines an additional parameter τ that controls the early termination of the scheme in the computation of each ensemble member. The theory of iterative regularization methods indicates that the selection of $\tau > 1/\rho$ ensures the aforementioned stabilization. On the one hand, smaller values of ρ are associated with smaller values of the regularization parameter and therefore with larger τ that ensures the early termination of the scheme to avoid unstable computations. On the other hand, larger values of ρ are associated to smaller and more controlled steps and therefore recommended for highly ill-posed nonlinear inverse problems.

The numerical results indicate that the stability in the computation of each ensemble is reflected in the accuracy of the ensemble approximation of the posterior. While for the IR-enLM the ensemble size had limited effect, the approximation properties of the IR-ES were substantially affected by the ensemble size. Our results suggest that with a size of $N_e = 75$, IR-ES inherits the regularization properties of the regularizing LM scheme [20]. In addition, for both proposed methods, larger values of the tunable parameter ρ correspond to more accurate approximations of the mean and variance of the posterior although at a larger computational cost. For the present examples we have found that values $\rho = 0.7$ and $\rho = 0.8$ are reasonable compromise between accuracy and cost. Finally, for IR-enLM it is important to reiterate that improved estimates of the mean and variance of the posterior were obtained where values of $\tau = 1 \leq 1/\rho$ were selected. However, the condition of $\tau > 1/\rho$ is only sufficient for the stability in the computation of each ensemble member and it is not conclusive of the quality of the ensemble approximation of the Bayesian posterior. Nonetheless, the present numerical study may potentially provide theoreticians with insight to develop further analysis of the ensemble methods for the solution of Bayesian inverse problems. Moreover, our numerical investigations have highlighted the potential benefit of importing ideas from iterative regularization methods for the proper selection of parameters in the optimization methods that are often used for the computation of ensemble approximations of the posterior. A natural extension of the present work, for example, is to apply the computational framework presented here to the selection of parameters in existing implementations of RML as well as the corresponding derivative-free approximations developed in [6]. Similarly, the proposed ideas for the selection of parameters can be explored for the adaptive selection of regularization parameters in the context of ES-MDA [12].

Acknowledgements The author would like to thank Professor Andrew Stuart for helpful discussions and his generous feedback on the content and structure of the manuscript. The author also thanks Professor Albert Reynolds and the anonymous referees for his comments and suggestions that greatly improve the manuscript.

References

1. S.I. Aanonsen, G. Naevdal, D.S. Oliver, A.C. Reynolds, , and B Valles. The Ensemble Kalman Filter in Reservoir Engineering—a Review. *SPE J.*, 14(3):393–412, 2009.
2. Z. Bi, D. S. Oliver, and A. C. Reynolds. Conditioning 3D stochastic channels to pressure data. *SPE Journal*, 5(4):474484, 2000.
3. S.P. Brooks and A. Gelman. General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics*, 7:434–455, December 1998.
4. Y. Chen and D. Oliver. Cross-covariances and localization for EnKF in multiphase flow data assimilation. *Computational Geosciences*, 14:579–601, 2010. 10.1007/s10596-009-9174-6.
5. Yan Chen and DeanS. Oliver. Ensemble randomized maximum likelihood method as an iterative ensemble smoother. *Mathematical Geosciences*, 44:1–26, 2012.
6. Yan Chen and DeanS. Oliver. Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. *Computational Geosciences*, 17(4):689–703, 2013.
7. Z. Chen, G. Huan, and Y. MA. *Computational Methods for Multiphase Flows in Porous Media*. Society for Industrial and Applied Mathematics, Philadelphia,PA,U.S.A., 2006.
8. S.L. Cotter, G.O. Roberts, A.M. Stuart, and D. White. MCMC methods for functions: modifying old algorithms to make them faster. *Submitted*, page arXiv:1202.0709, 2012.
9. C. V. Deutsch. *Geostatistical Reservoir Modeling*. Oxford University Press, Oxford, 2002.
10. A. Emerick and A. Reynolds. Combining sensitivities and prior information for covariance localization in the ensemble Kalman filter for petroleum reservoir applications. *Computational Geosciences*, 15:251–269, 2011. 10.1007/s10596-010-9198-y.
11. A. Emerick and A. C. Reynolds. History-matching production and seismic data in a real field case using the ensemble smoother with multiple data assimilation. *Proceedings of the SPE Reservoir Simulation Symposium, The Woodlands, Texas, USA, 18-20 February, SPE SPE-163645*, 2013.
12. Alexandre A. Emerick and Albert C. Reynolds. Ensemble smoother with multiple data assimilation. *Computers and Geosciences*, 55(0):3 – 15, 2013.
13. AlexandreA. Emerick and AlbertC. Reynolds. Investigation of the sampling performance of ensemble-based methods with a simple reservoir model. *Computational Geosciences*, 17(2):325–350, 2013.
14. G. Evensen. *Data assimilation: the ensemble Kalman filter*. Springer Verlag, 2009.
15. G. Evensen. The ensemble Kalman filter for combined state and parameter estimation. *Control Systems, IEEE*, 29(3):83–104, June 2009.
16. Guohua Gao and U A.C. Reynolds. An improved implementation of the LBFGS algorithm for automatic history matching. in *Proceedings of the SPE Annual Technical Conference and Exhibition, 26-29 September 2004, Houston, Texas*.
17. C. W. Groetsch. *The theory of Tikhonov regularization for Fredholm equations of the first kind*. Pitman, Boston London Melbourne, 1984.
18. Yaqing Gu and Dean S. Oliver. An iterative ensemble Kalman filter for multiphase fluid flow data assimilation. *SPE Journal*.
19. S. A. HABEN, A. S. LAWLESS, and N. K. NICHOLS. Conditioning of incremental variational data assimilation, with application to the met office system. *Tellus A*, 63(4):782–792, 2011.
20. M. Hanke. A regularizing Levenberg-Marquardt scheme, with applications to inverse groundwater filtration problems. *Inverse Problems*, 13:79–95, 1997.
21. Marco A Iglesias, Kody J H Law, and Andrew M Stuart. Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.
22. MarcoA. Iglesias and Clint Dawson. The regularizing Levenberg-Marquardt scheme for history matching of petroleum reservoirs. *Computational Geosciences*, pages 1–21, 2013.
23. MarcoA. Iglesias, KodyJ.H. Law, and AndrewM. Stuart. Evaluation of Gaussian approximations for data assimilation in reservoir models. *Computational Geosciences*, pages 1–35, 2013.
24. T.K. Jensen and P.C. Hansen. Iterative regularization with minimum-residual methods. *BIT Numerical Mathematics*, 47(1):103–120, 2007.
25. B. Katltenbacher, A. Neubauer, and O. Scherzer. *Iterative Regularization Methods for Nonlinear Ill-Posed Problems*. Radon Series on Computational and Applied Mathematics, de Gruyter, Berlin, 1st edition, 2008.
26. G. Li and A. Reynolds. Iterative Ensemble Kalman Filters for Data Assimilation. *SPE Journal*, 14(3):496505, 2009.

27. R. Li, A. C. Reynolds, and D.S. Oliver. History matching of three-phase flow production data. *SPE Journal*, 8(4):328–340, 2003.
28. Binghuai Lin. Real-time ensemble control with reduced-order modeling. In *Ph.D Thesis*, volume MIT, 2012.
29. N. Liu and D.S Oliver. Evaluation of Monte Carlo methods for assessing uncertainty. *SPE J.*, 8:188–195, 2003.
30. A. C. Reynolds Oliver, D. S. and N. Liu. *Inverse Theory for Petroleum Reservoir Characterization and History Matching*. Cambridge University Press, ISBN: 9780521881517, 1st edition, 2008.
31. Albert C. Reynolds. My decade-long journal through the field of ensemble-based data assimilation. *EnKF Workshop*, June 23rd, 2014.
32. A.M. Stuart. Inverse problems: a Bayesian perspective. In *Acta Numerica*, volume 19, 2010.
33. R. Tavakoli and A. Reynolds. Monte Carlo simulation of permeability fields and reservoir performance predictions with SVD parameterization in RML compared with EnKF. *Computational Geosciences*, 15:99–116, 2011. 10.1007/s10596-010-9200-8.
34. R. Tavakoli and A. C. Reynolds. History matching with parametrization based on the SVD of a dimensionless sensitivity matrix. *SPE Journal*, 15(2):495–508., 2010.