Ensemble Methods for Geothermal Model Calibration

Alex de Beer[1], Michael Gravatt[1], Ruanui Nicholson[1], John P. O'Sullivan[1], Michael J. O'Sullivan[1], and Oliver J. Maclaren[1]

[1]Department of Engineering Science, University of Auckland, 70 Symonds Street, Grafton, Auckland 1010

adeb970@aucklanduni.ac.nz

Keywords: *reservoir modelling, model calibration, uncertainty quantification, ensemble methods.*

ABSTRACT

A typical geothermal model requires significant computational resources to simulate and can contain hundreds of unknown parameters. The process of estimating these parameters, often referred to as model calibration, is a difficult task; traditional methods such as Markov chain Monte Carlo generally require running a prohibitively large number of simulations to obtain accurate results. Ensemble methods form an alternative class of algorithms for approximating the solution to the calibration problem and have the potential to provide accurate results using considerably fewer simulations. Ensemble methods have been used successfully to calibrate large, complex models in areas including petroleum engineering, oceanography, and weather forecasting. There are, however, few examples of applications of these methods to geothermal reservoir modelling. In addition, the wide variety of ensemble methods that have been developed mean there is a need for numerical studies that examine their respective benefits and drawbacks when applied to specific problems. To support the effective use of ensemble methods for geothermal reservoir model calibration, we review two widely used ensemble methods and apply them to the problem of calibrating a synthetic reservoir model. We demonstrate that both methods are capable of generating accurate reconstructions of the model parameters, with appropriate characterisation of uncertainty.

1. INTRODUCTION

Reservoir modelling is an important tool in the sustainable management of geothermal resources. The effectiveness of a reservoir model, however, depends on the degree to which it reflects reality. A key component of developing an effective model is the calibration process, which involves identifying model parameters that provide an acceptable match to field data. In the context of reservoir modelling, the parameters of interest typically include the subsurface permeability structure and the strength and magnitude of the hot mass upflow at the base of the model. By contrast, the data is typically limited to downhole temperature and pressure measurements.

Here, we consider the model calibration problem from a Bayesian perspective (Aster et al., 2018; Kaipio and Somersalo, 2006). Within the Bayesian framework, the task of model calibration is posed as a statistical inference problem. The aim is to form the posterior probability distribution, which characterises the uncertainty in the parameters that remains after conditioning on the data. Once characterised, the posterior parameter uncertainty can be propagated through to the model predictions (for example, future steam flow rates or reservoir temperatures), allowing the modeller to describe, in a probabilistic sense, the future behaviour of the geothermal system.

The complexity of the typical reservoir model means that the posterior is not available in closed form; instead, it is characterized using samples, using methods such as Markov chain Monte Carlo and sequential Monte Carlo. These methods are exact, in the sense that the resulting samples are distributed according to the true posterior. However, they generally require at least $\mathcal{O}(10^4)$ iterations to provide an accurate characterisation of the posterior (Cui et al., 2011; Maclaren et al., 2020). This can be prohibitive in the geothermal setting, where a single simulation can take hours or even days.

The challenges associated with the use of exact sampling methods for problems with complex models and large state or parameter spaces have motivated the development of methods that provide an approximate characterisation of the posterior using significantly fewer model runs. A feature common to many of these methods is that, under a Gaussian prior distribution and a linear model, the samples they generate are distributed according to the posterior; under a nonlinear model, however, this does not hold in general. Among the most widely used methods for approximately characterising the posterior are linearisation of the model about the point in parameter space with the greatest posterior density (the maximum-a-posteriori estimate; see, e.g., Omagbon et al., 2021), and randomised maximum likelihood (Kitandis, 1995; Oliver, 1996), in which one repeatedly solves a stochastic optimisation problem to obtain samples distributed in regions of high posterior density.

An alternative class of algorithms for approximate Bayesian inference are ensemble methods, in which a small ensemble (group of parameter sets) is combined with data, in an iterative manner, such that the distribution of the ensemble approximates the posterior. The first ensemble-based algorithm was the ensemble Kalman filter (Evensen, 2009), which was developed for the purpose of state estimation of dynamical systems and has been used extensively in areas including weather forecasting and oceanography. A great deal of subsequent research, however, has focused on the development of ensemble methods for approximating the solutions to inverse problems. Such methods are often referred to as iterative ensemble smoothers (IES), or as forms of ensemble Kalman inversion (EKI). An advantage of ensemble methods over methods such as randomised maximum likelihood is that they do not require the computation of derivatives, which can be expensive; instead, these are approximated using the ensemble.

Though ensemble methods have begun to see some use within the geothermal community (see, e.g., Békési et al., 2020; Bjarkason et al., 2021), there remains a need for numerical studies that investigate how these methods can be applied effectively in a geothermal context. In this work, we outline two ensemble methods that are widely used within the geosciences and apply them to a synthetic reservoir model.

ISSN 2703-4275

2. CALIBRATION IN A BAYESIAN FRAMEWORK

We consider problems in which the unknown parameters, $\theta \in \mathbb{R}^{N_\theta}$, and the observations, $y \in \mathbb{R}^{N_y}$, are related through

$$y = \mathcal{G}(\theta) + \epsilon,$$

where $\mathcal{G} : \mathbb{R}^{N_\theta} \to \mathbb{R}^{N_y}$ denotes the forward model and $\epsilon \in \mathbb{R}^{N_y}$ is a vector of additive measurement errors. The process of applying the forward model to find y using a particular instance of θ is referred to as solving the forward problem. The inverse, or calibration problem, by contrast, is the process of estimating θ given a set of observations y.

The Bayesian approach to solving the calibration problem requires us to first form a prior distribution; that is, a mathematical representation of expert knowledge on the likely values of the parameters prior to data being observed. Data, once collected, is then combined with the prior to form the posterior distribution using Bayes' theorem, which we express here as

$$\pi(\theta|y) \propto \pi(y|\theta)\pi(\theta).$$

In the above, $\pi(\theta)$ denotes the prior, $\pi(y|\theta)$ denotes the likelihood, which expresses the probability of the data given a particular instance of the parameters, and $\pi(\theta|y)$ denotes the posterior, or the conditional density of the parameters given the observations.

If we assume that the prior is Gaussian, with mean μ_θ and covariance Γ_θ, and that the distribution of the errors is Gaussian, with mean 0 and covariance Γ_ϵ, it can be shown (see, e.g., Aster et al., 2018) that the posterior takes the form:

$$\pi(\theta|y) \propto \exp(-\tfrac{1}{2}\|L_\epsilon(\mathcal{G}(\theta) - y)\|^2 - \tfrac{1}{2}\|L_\theta(\theta - \mu_\theta)\|^2),$$

where $L_\epsilon^T L_\epsilon = \Gamma_\epsilon^{-1}$, $L_\theta^T L_\theta = \Gamma_\theta^{-1}$, and $\|\cdot\|$ denotes the Euclidean norm. When the forward model is linear, the posterior is also Gaussian; in the nonlinear case, however, this is no longer true.

3. ENSEMBLE METHODS

We describe variations of two ensemble methods that have been applied successfully to complex problems within the geosciences (see, e.g., Emerick, 2016; Chen and Oliver, 2017); the ensemble smoother with multiple data assimilation (ES-MDA) and ensemble randomised maximum likelihood (EnRML).

3.1 Ensemble Smoother with Multiple Data Assimilation

The ensemble smoother with multiple data assimilation (Emerick and Reynolds, 2013a) generates an approximation to the posterior by iteratively transforming an initial ensemble (set of samples) drawn from the prior, such that the ensemble members approximate a sequence of distributions that transition smoothly from the prior to the posterior. Each transformation is applied under the assumption of a Gaussian prior and a linear model. In this sense, ES-MDA can be considered an approximate form of tempering, a variant of sequential Monte Carlo (Stordal and Elsheikh, 2015; Iglesias and Yang, 2021).

To begin the ES-MDA algorithm, we sample an initial ensemble, $\theta_k^{(1)}, k \in \{1,2,\dots,N_e\}$, from the prior, and run the forward model to compute $\mathcal{G}_k^{(1)} = \mathcal{G}(\theta_k^{(1)})$ for each ensemble member. At each iteration $i \in \{1,2,\dots,N_i\}$, we first compute the matrices of scaled differences $\Delta\theta^{(i)} \in \mathbb{R}^{N_\theta \times N_e}$ and $\Delta\mathcal{G}^{(i)} \in \mathbb{R}^{N_y \times N_e}$, defined as

$$\Delta\theta^{(i)} \equiv \frac{1}{\sqrt{N_e - 1}}\left[\theta_1^{(i)} - \bar{\theta}^{(i)}, \dots, \theta_{N_e}^{(i)} - \bar{\theta}^{(i)}\right],$$

$$\Delta\mathcal{G}^{(i)} \equiv \frac{1}{\sqrt{N_e - 1}}\left[\mathcal{G}_1^{(i)} - \bar{\mathcal{G}}^{(i)}, \dots, \mathcal{G}_{N_e}^{(i)} - \bar{\mathcal{G}}^{(i)}\right],$$

where $\bar{\theta}^{(i)}$ and $\bar{\mathcal{G}}^{(i)}$ denote the means of the ensemble parameters and modelled observations respectively. Next, we compute the ensemble estimates of the covariance of the modelled observations, $\Gamma_{yy}^{(i)} \in \mathbb{R}^{N_y \times N_y}$, and the cross-covariance between the parameters and modelled observations, $\Gamma_{\theta y}^{(i)} \in \mathbb{R}^{N_\theta \times N_y}$, which are defined as

$$\Gamma_{yy}^{(i)} \equiv \Delta\mathcal{G}^{(i)}\Delta\mathcal{G}^{(i)T}, \quad \Gamma_{\theta y}^{(i)} \equiv \Delta\theta^{(i)}\Delta\mathcal{G}^{(i)T}.$$

We then update ensemble member, $\theta_k^{(i)}$, where $k \in \{1,2,\dots,N_e\}$, using

$$\theta_k^{(i+1)} = \theta_k^{(i)} + \Gamma_{\theta y}^{(i)}\left(\Gamma_{yy}^{(i)} + \alpha^{(i)}\Gamma_\epsilon\right)^{-1}\left(y + \epsilon_k^{(i)} - \mathcal{G}_k^{(i)}\right),$$

where $\epsilon_k^{(i)} \sim \mathcal{N}(0, \alpha^{(i)}\Gamma_\epsilon)$ and $\alpha^{(i)} \geq 1$ is a factor by which the covariance of the measurement errors is inflated.

In addition to selecting an appropriate ensemble size, a key consideration when implementing ES-MDA is the choice of a suitable set of inflation factors. In the linear-Gaussian case, the samples generated using ES-MDA are distributed according to the posterior if the inflation factors satisfy the following condition, which can be derived by considering a factorisation of the likelihood (Stordal and Elsheikh, 2015):

$$\Sigma_{i=1}^{N_i}\frac{1}{\alpha^{(i)}} = 1.$$

This condition is still, in general, adhered to even when the model contains nonlinearities. The simplest method of inflation factor selection is to choose a desired number of iterations, N_i, and set $\alpha^{(i)} = N_i$ for all $i \in \{1,2,\dots,N_i\}$. It has, however, been demonstrated that a decreasing sequence of factors typically performs better in practice; this corresponds to making smaller transformations to each ensemble member at early stages of the algorithm. A variety of methods have been proposed, using techniques for the regularisation of classical inverse problems (Le et al., 2015; Rafiee and Reynolds, 2017), and the selection of successive target distributions in sequential Monte Carlo methods (Iglesias et al., 2018; Iglesias and Yang, 2021). In the present work, we use the method of Iglesias and Yang (2021); at iteration i, we calculate the misfit between the current set of ensemble members and the data; for ensemble member k, this is defined as

$$\mathrm{S}_k^{(i)} \equiv \frac{1}{2}\left\|L_\epsilon\left(\mathcal{G}_k^{(i)} - y\right)\right\|^2.$$

Defining $\mu_{\mathrm{S}}^{(i)}$ and $\sigma_{\mathrm{S}}^{2\,(i)}$ to be the empirical mean and variance of $\{\mathrm{S}_k^{(i)}\}$, the inflation factor, $\alpha^{(i)}$, is given by

$$\frac{1}{\alpha^{(i)}} = \min\left\{\max\left\{\frac{N_y}{2\mu_{\mathrm{S}}^{(i)}}, \left(\frac{N_y}{2\sigma_{\mathrm{S}}^{2\,(i)}}\right)^{1/2}\right\}, 1 - \Sigma_{j=1}^{i-1}\frac{1}{\alpha^{(j)}}\right\}.$$

The algorithm is considered converged once the aforementioned condition on the inflation factors is met.

Proceedings 45th New Zealand Geothermal Workshop
15-17 November, 2023
Auckland, New Zealand
ISSN 2703-4275

3.2 Ensemble Randomised Maximum Likelihood

An alternative group of ensemble methods (Chen and Oliver, 2013; White, 2018) share characteristics of the optimisation-based framework of RML. Unlike RML, however, these methods perform the optimisation on all ensemble members simultaneously and, at each step of the optimisation, use an "average" gradient, estimated using the ensemble itself, when updating each ensemble member. Here, we describe the method of Chen and Oliver (2013), which is an ensemble approximation of the Levenberg-Marquardt algorithm (see, e.g., Nocedal and Wright, 2006).

As in ES-MDA, we begin the EnRML algorithm by sampling an initial ensemble, $\theta_k^{(1)}, k \in \{1,2,\dots,N_e\}$, from the prior. At each iteration $i \in \{1,2,\dots,N_i\}$, each ensemble member is updated using

$$\theta_k^{(i+1)} = \theta_k^{(i)} - \left(\left(1+\lambda^{(i)}\right)\Gamma_\theta^{-1} + J^{(i)T}\Gamma_\epsilon^{-1}J^{(i)}\right)^{-1} \times \left(\Gamma_\theta^{-1}\left(\theta_k^{(i)} - \theta_k^{(1)}\right) + J^{(i)T}\Gamma_\epsilon^{-1}\left(\mathcal{G}_k^{(i)} - y - \epsilon_k\right)\right),$$

where $\epsilon_k \sim \mathcal{N}(0,\Gamma_\epsilon)$. In the above, $\lambda^{(i)}$ denotes the Levenberg-Marquardt damping parameter, and $J^{(i)} \in \mathbb{R}^{N_y \times N_\theta}$ denotes the ensemble estimate of the model Jacobian, defined as

$$J^{(i)} \equiv \Delta\mathcal{G}^{(i)}\left(\Delta\theta^{(i)}\right)^{\dagger},$$

where $\left(\Delta\theta^{(i)}\right)^{\dagger}$ denotes the pseudoinverse of $\Delta\theta^{(i)}$. We note that the ensemble estimate of the Jacobian can be unstable, particularly when a small ensemble size is used; for this reason, we follow Chen and Oliver (2013) and make a set of slight modifications to the update equation which avoid the explicit computation of the Jacobian.

After each iteration, we check the mean misfit, $\mu_S^{(i)}$, between the updated ensemble and the observations. If this quantity has decreased from the previous iteration, we accept the iteration and reduce the value of the damping parameter, λ; otherwise, we repeat the iteration with an increased value of λ. We follow the recommendations outlined in Chen and Oliver (2013) when selecting the initial value of λ, the factors used to modify λ after each iteration, and suitable bounds for λ. We consider the algorithm converged if the reduction in $\mu_S^{(i)}$ between two iterations is less than 1%, the largest change in any parameter of any ensemble member between two iterations is less than 0.25 (prior) standard deviations, or if five consecutive iterations are rejected.

3.3 Additional Considerations

Though the standard forms of ES-MDA and EnRML have been applied successfully to inverse problems within the geosciences, they can encounter issues due to sampling errors that arise because of the use of a small ensemble. Spurious correlations that are present between ensemble parameters and modelled observations can result in updates to ensemble parameters being influenced by observations with which they have no relationship. In addition, the variants of both ensemble methods we implement here possess the subspace property; that is, at each iteration, the updated ensemble lies within the span of the initial ensemble (see, e.g., Iglesias et al., 2013a). If a small ensemble is used and the dimensionality of the parameter space is large (that is, $N_e \ll N_\theta$), as is typical in many geoscience applications, there may exist regions of parameter space with high posterior density that cannot be reached by the ensemble, limiting its ability to approximate the posterior.

To address these issues, it is common to apply a form of localisation when using an ensemble method. Localisation involves the modification of the ensemble covariance matrices $\Gamma_{\theta y}^{(i)}$ and $\Gamma_{yy}^{(i)}$, or the Kalman gain, $K^{(i)} \in \mathbb{R}^{N_\theta \times N_y}$; in the case of ES-MDA, for example, $K^{(i)}$ is defined as

$$K^{(i)} \equiv \Gamma_{\theta y}^{(i)}\left(\Gamma_{yy}^{(i)} + \alpha^{(i)}\Gamma_\epsilon\right)^{-1}.$$

This modification is carried out by taking the Schur (elementwise) product of the matrix (or matrices) to which localisation is being applied, and a localisation matrix (or matrices) ρ, of the same dimensions, the entries of which typically range between 0 and 1. In the case of ES-MDA, when localisation is applied to the Kalman gain, the update equation becomes

$$\theta_k^{(i+1)} = \left[\rho \odot K^{(i)}\right]\left(y + \epsilon_k^{(i)} - \mathcal{G}_k^{(i)}\right),$$

where $\odot$ denotes the Schur product. In addition to reducing spurious correlations, the localisation matrix acts to increase the rank of the Kalman gain, which, without localisation, is limited by the size of the ensemble. This breaks the subspace property, allowing updated ensemble members to occupy regions outside the span of the initial ensemble.

Most localisation methods are distance-based (see, e.g., Chen and Oliver, 2017); that is, the entries of the localisation matrix ρ are chosen such that they decrease as the physical distance between the parameters and/or observations they relate increases, under the assumption that the correlations between these quantities reduce with distance. An alternative class of localisation methods, however, aim to reduce spurious correlations without using the physical distances between quantities (Luo and Bhakta, 2020; Zhang and Oliver, 2010). In the present work we use a variant of the algorithm developed by Luo and Bhakta (2020), which uses a resampling procedure to estimate the correlations between the physical quantities each element of the Kalman gain relates and reduces elements of the gain corresponding to quantities estimated to have low levels of correlation. A similar approach is used in the implementation of EnRML in the PEST++ software suite (White et al., 2020), which is widely used for the calibration of environmental models. More experimentation, however, is required to identify which localisation methods work best in geothermal settings.

4. SYNTHETIC RESERVOIR MODEL

To test each ensemble method, we use a two-dimensional slice model, discretised on a 25×25 mesh, with dimensions of $1500\ \text{m} \times 1500\ \text{m}$. The model contains five production wells, the locations of which are indicated in Figure 1. Each extends to a depth of $-1300\ \text{m}$ and contains a single feedzone at a depth of $-500\ \text{m}$. We consider a combined natural state and production history setup; that is, we run the model until stable steady-state conditions are reached, then use the resulting state of the system as the initial condition for the subsequent production simulation. We model each well as extracting fluid at a rate of $2\ \text{kg s}^{-1}$ for the entire production period, which lasts for two years.

We consider the problem of estimating the (isotropic) permeability within each block of the model mesh, as well as the mass rate of the upflow at the base of the model. This

Proceedings 45th New Zealand Geothermal Workshop
15-17 November, 2023
Auckland, New Zealand
ISSN 2703-4275

gives a total of 626 unknown parameters. All other parameters are assumed to be known. The rock in the reservoir is assumed to have a porosity of 10%, a density of 2500 kg m^{-3}, a thermal conductivity of 2.5 W m^{-1} K^{-1}, and a specific heat of 1000 J kg^{-1} K^{-1}. The top boundary of the model is set to a constant pressure of 1 bar and a temperature of 20 °C, representing an atmospheric boundary condition. We impose a constant heat flux of 200 mW m^{-2} through the bottom boundary, except for the cell at the centre of the boundary, which is given a mass flux (of unknown rate) of fluid with an enthalpy of 1500 kJ kg^{-1}. The side boundaries are closed.

4.1 Prior Parametrization

The update equations of standard ensemble methods are not able to work directly with parameters that are restricted to a number of discrete values, or constrained to be within physically realistic bounds. However, these types of parametrizations are often useful in reservoir modelling; for instance, we may wish to model the reservoir as being composed of a small number of rock formations with common physical characteristics, or to restrict the mass upflow across the bottom boundary of the domain to non-negative values.

To incorporate these types of parametrizations into the inversion, we sample the set of parameters, θ, corresponding to each ensemble member from a Gaussian distribution, and use suitable transformations of these to form sets of geological parameters, which we denote using u, that are distributed according to our prior beliefs. Subsequent updates to the untransformed parameters maintain any bounds or discontinuities we impose in the prior. In this framework, the forward model can be written as

$$\mathcal{G} = \mathcal{F} \circ \mathcal{H},$$

where $\mathcal{H}$ is the mapping between the untransformed parameters and the geological parameters they represent, and $\mathcal{F}$ is the reservoir simulator. The task of the ensemble method, then, is to estimate the posterior density of the untransformed parameters θ; we can then apply the mapping $\theta \mapsto \mathcal{H}(\theta)$ to estimate the posterior density of the corresponding geological parameters u.

4.1.1 Permeability Parametrization

We partition the model into three subdomains with variable boundaries: a shallow high-permeability region (Ω_1), a low-permeability clay cap (Ω_2), and a deep high-permeability region (Ω_3). The permeability at an arbitrary location, $\kappa(x, z)$, is given by

$$\kappa(x,z) = \begin{cases} 10^{\xi_1(x,z)}, & z \geq \omega_1(x), \\ 10^{\xi_2(x,z)}, & \omega_1(x) > z \geq \omega_2(x), \\ 10^{\xi_3(x,z)}, & \omega_2(x) > z, \end{cases}$$

where $\xi_1(\cdot)$, $\xi_2(\cdot)$ and $\xi_3(\cdot)$ are functions that describe the log-permeability in regions Ω_1, Ω_2 and Ω_3, and $\omega_1(\cdot)$ and $\omega_2(\cdot)$ are functions that describe the boundaries between Ω_1 and Ω_2, and Ω_2 and Ω_3, (the top and bottom surfaces of the clay cap) respectively. Figure 1 shows a possible partitioning of the model domain.

We set $\omega_1(x) = -60$ m, which reflects a prior assumption that the location of the top of the clay cap is known. However, we treat the elevation of the bottom of the clay cap as unknown, giving $\omega_2(x_1)$ a Gaussian process prior with a mean, μ, of -350 m, and a covariance function of

$$C(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2\ell_x^2}\right),$$

where we use a standard deviation, σ, of 80 m, and a characteristic lengthscale, ℓ_x, of 500 m.

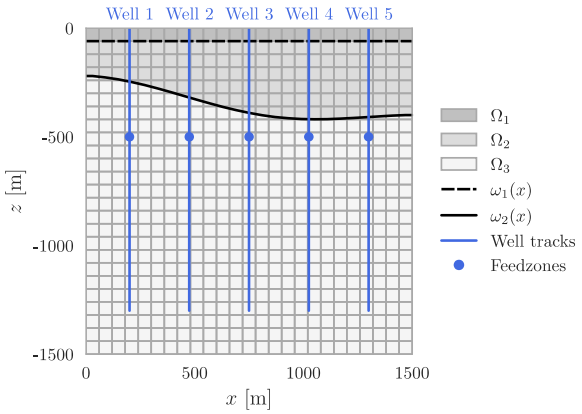

Figure 1. Model mesh with well positions, and a possible partitioning of the domain into subdomains Ω_1, Ω_2, and Ω_3, with boundaries $\omega_1(x)$ and $\omega_2(x)$.

We use the level set method to model the log-permeability field, ξ_i, in each subdomain Ω_i. The level set method is commonly used in subsurface modelling to generate distinct zones with common geophysical characteristics (Muir and Tsai, 2020; Nicholson et al., 2020; Tso et al., 2021), the boundaries of which are defined using the contours of a continuous underlying function, which we refer to as the level set function. In each subdomain, we take the underlying level set function to be a Gaussian random field with the anisotropic covariance function given by

$$C(x, z, x', z') = \sigma^2 \exp\left(-\frac{1}{2}\left(\frac{(x - x')^2}{\ell_x^2} + \frac{(z - z')^2}{\ell_z^2}\right)\right).$$

We set the mean of the functions associated with the high-permeability fields, ξ_1 and ξ_3, to -14 $\log(m^2)$, and the mean of the function associated with the clay cap field, ξ_2, to -16 $\log(m^2)$. We set the standard deviation, σ, of the functions associated with fields ξ_1 and ξ_2 to 0.25 $\log(m^2)$, and the standard deviation of the function associated with field ξ_3 to 0.5 $\log(m^2)$, reflecting increased uncertainty in the permeability with depth. In all regions, we set the characteristic lengthscale in the x direction, ℓ_x, to 1500 m, and the characteristic lengthscale in the z direction, ℓ_z, to 200 m. We generate distinct zones of constant permeability by mapping the value of the corresponding level set function at each point on the model mesh to the nearest of a set of pre-defined values $v_i, i \in \{1, 2, \ldots, N_v\}$; that is, at each point we define the permeability, $\kappa(x, z)$, as

$$\kappa(x, z) = \Sigma_{i=1}^{N_v} \kappa_{v_i} \mathbb{I}_{v_i}(x, z),$$

where κ_{v_i} is the permeability corresponding to value v_i, and $\mathbb{I}_{v_i}(x, z)$ is equal to 1 if value v_i is the value closest to the value of the level set function at location (x, z), and equal to 0 otherwise. For simplicity, we choose the set of values such that the log-permeabilities in adjacent zones in each

ISSN 2703-4275

subdomain differ by 0.25 log (m^2). Figure 2 shows sets of transformed permeabilities sampled from the prior.

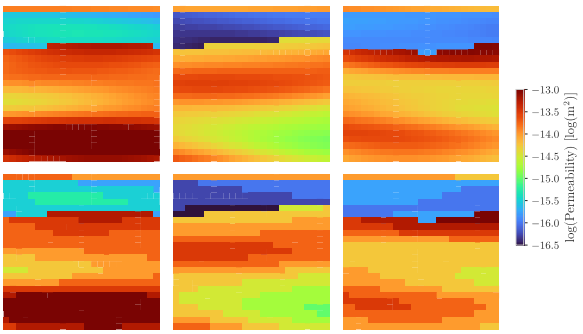

Figure 2. Prior samples of the underlying level set function (top row) and the corresponding permeabilities (bottom row).

4.1.2 Mass Upflow Parametrization

For the prior on the magnitude of the mass upflow at the base of the model, we use a uniform distribution on the interval $[b_1, b_2]$, where b_1 is equal to 0.1 kg s^{-1} and b_2 is equal to 0.2 kg s^{-1}. We use samples from the unit normal distribution to represent the corresponding untransformed parameter in the prior ensemble. To transform this parameter such that it has the appropriate distribution, we apply the mapping

$$\theta \mapsto b_1 + (b_2 - b_1)\Phi(\theta),$$

where $\Phi(\cdot)$ denotes the cumulative density function of the unit normal distribution.

4.2 Truth and Observations

We generate the true system we are interested in recovering the parameters of by drawing a random sample from the prior. Figure 3 shows the permeabilities and natural state convective plume of the true system.

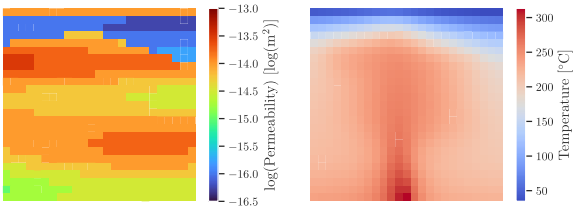

Figure 3. True reservoir permeabilities (left) and natural state convective plume (right).

We record the temperatures at six equispaced points in each well prior to production occurring, as well as the pressure and enthalpy of the fluid extracted at each of the five wells at three-month intervals for the first year of production. This gives a total of 80 measurements. We add independent Gaussian noise to each measurement with a standard deviation of 2% of the maximum of the modelled values for the corresponding data type.

4.3 Simulation

We use the open-source simulator Waiwera (Croucher et al., 2020) to run the model. In the event a simulation fails to complete, we discard the corresponding ensemble member and continue the inversion with a reduced ensemble.

5. RESULTS

We run our implementations of ES-MDA and EnRML, with and without localisation, using $N_e = 100$ ensemble members. We use the same initial ensemble for all methods. Our implementation of ES-MDA converges after 11 iterations when localisation is not applied, and after 12 iterations when localisation is applied. Our implementation of EnRML converges after 25 iterations (including rejected iterations) when localisation is not applied, and after 22 iterations when localisation is applied. At the final iteration of both variants of EnRML, the reduction in the mean misfit is less than 1% and the maximum change in any parameter of any ensemble member is less than 0.25 standard deviations. We note, however, that there is no formal justification for the EnRML stopping criteria we use here; it is possible that similarly accurate results can be obtained with less computation by relaxing these criteria. The simulation failure rate is similar for all methods; in all cases, the final ensemble consists of between 87 and 88 members.

Figure 4 shows the means and standard deviations of the log-permeabilities of the final ensemble generated using each method (the means and standard deviations of the prior ensemble are also presented for comparison). There is a high degree of consistency between the means obtained using each method. In all cases, the shape of $\omega_2(x)$ (the bottom surface of the clay cap) is well recovered; the reconstruction appears to be particularly good when localisation is not applied. All methods appear to identify, to varying degrees, the region of high permeability beneath the clay cap. There are, however, significant differences between the mean permeabilities in the bottom half of the model domain and the true permeabilities; for instance, the low-permeability region in the bottom left-hand corner of the domain is not present in any of the reconstructions. It is, however, difficult to identify whether this is due to shortcomings in the ensemble methods, or because the observations are not very informative of the deeper permeabilities. Most of the observations are made at

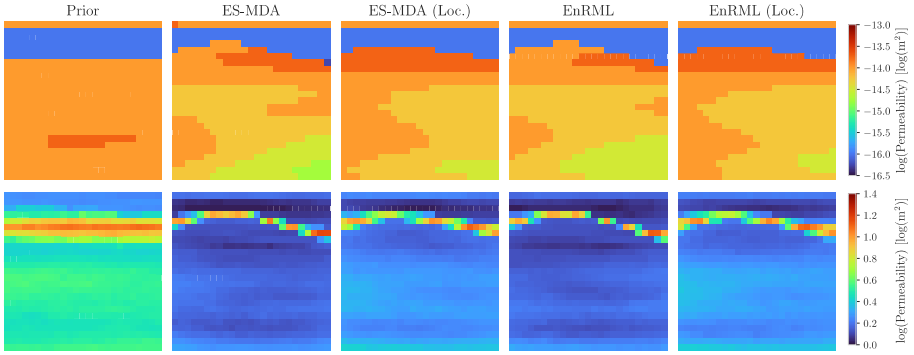

Figure 4. Means (top row) and standard deviations (bottom row) of the permeabilities of the prior ensemble and the final ensembles obtained using each ensemble method.

Proceedings 45th New Zealand Geothermal Workshop
15-17 November, 2023
Auckland, New Zealand
ISSN 2703-4275

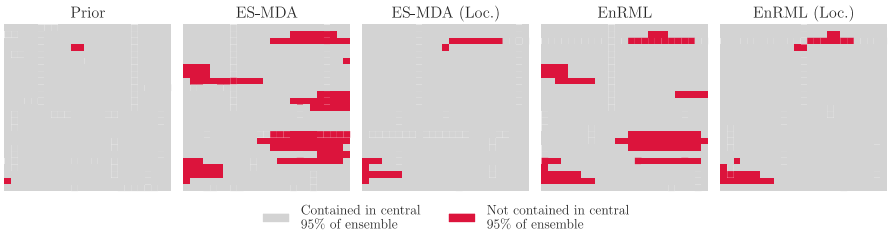

Figure 5. Heatmaps which indicate, for each method, whether the true permeability in each cell of the model mesh is contained in the central 95% of permeabilities of the final ensemble.

the well feedzones, which are located in the top half of the model domain (at a depth of -500 m; see Figure 1); it seems plausible that there is insufficient information to allow for a precise estimate of the permeability structure at depth.

Figure 4 also shows that the standard deviations of the permeabilities of the final ensemble generated using each method are significantly reduced compared to those of the prior ensemble. In all cases, the area with the greatest degree of uncertainty is the region around the bottom surface of the clay cap. This is to be expected; the permeabilities on either side of this surface tend to take on significantly different values, so unless we have complete certainty in the location of this boundary, we would expect to have high levels of uncertainty in the permeabilities around it. In all cases, the uncertainty in the permeabilities in the region below the clay cap, where the well feedzones are located, is lower than the uncertainty in the permeabilities of the deeper sections of the model domain. This makes intuitive sense; given that most of the data is collected at the feedzones, it is logical that the permeability distribution is better constrained in this region. For both ES-MDA and EnRML, we observe that the uncertainty in the permeabilities is larger, over the entire model domain, for the variants where localisation is applied.

Figure 5 shows heatmaps that indicate, for each method, whether the true permeability in each cell of the model mesh is contained within the central 95% of the permeabilities of the final ensemble. For both ES-MDA and EnRML, we observe that when localisation is not applied, a significant proportion of the model permeabilities ($\sim$15-20%) are not contained within the central 95% of the ensemble. When localisation is applied, however, this drops significantly, to between 3% and 5%. This suggests that, if our intention when using an ensemble method is to recover the truth within a confidence interval, the use of localisation is important. Without knowledge of the true posterior, however, it is difficult to draw conclusions as to whether localisation improves the approximation of the posterior.

Figure 6 shows the magnitudes of the mass upflows of the final ensemble associated with each method. In all cases, the truth is well contained within the ensemble, and the uncertainty is significantly reduced in comparison to the prior ensemble. Figure 7 shows estimates of the pressure at the feedzone of well 1 and the enthalpy of the fluid extracted at the feedzone of well 3. Again, in all cases the truth is well contained within the ensemble and the uncertainty is significantly reduced in comparison to the prior, though in some cases it remains slightly greater than we might expect given the size of the measurement errors. All methods appear to provide reasonable estimates of how the pressures and enthalpies change over the year following the end of the observation period.

Figure 8 shows the permeabilities of the first two ensemble members from the final ensemble associated with each method. As expected, these show significantly less variability than the draws from the prior and share greater similarity with the true permeabilities.

6. CONCLUSION

In this work, we have illustrated the ability of ensemble methods to generate accurate approximations, with a reasonable characterisation of uncertainty, to the problem of reservoir model calibration. We have also demonstrated the flexibility of ensemble methods with regard to model parametrisation, and the differences in the results obtained by ensemble methods with and without the use of localisation. In the experiments performed here, our implementations of each ensemble method performed well using $\mathcal{O}(10^3)$ simulations of the forward model, which is considerably fewer than the number of simulations typically required for exact sampling methods (Cui et al., 2011; Maclaren et al., 2020).

A natural extension to this work would be to conduct similar numerical tests to those we have described here using a three-dimensional reservoir model, and to consider the estimation of additional parameters that are likely to be of interest in a realistic modelling scenario. As well as additional geological parameters, such as porosities, it is possible to use ensemble methods to estimate characteristics of the prior, such as lengthscales and variances (Chada et al., 2018; Tso et al., 2021). In the present work we have treated these parameters as known, but in practice they can be associated with significant uncertainty which, if neglected, can cause the posterior to be characterised poorly. In addition, it has been demonstrated that the simulated natural state of a reservoir model may not be unique (Bjarkason et al., 2019); indeed, during the development of the model used for this work, we were able to generate several different natural state solutions

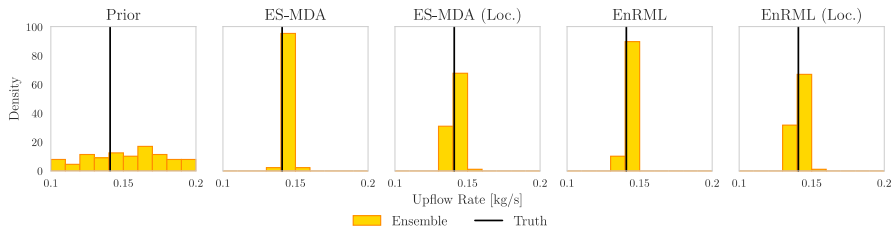

Figure 6. Distributions of the mass upflow rates obtained using each method.

ISSN 2703-4275

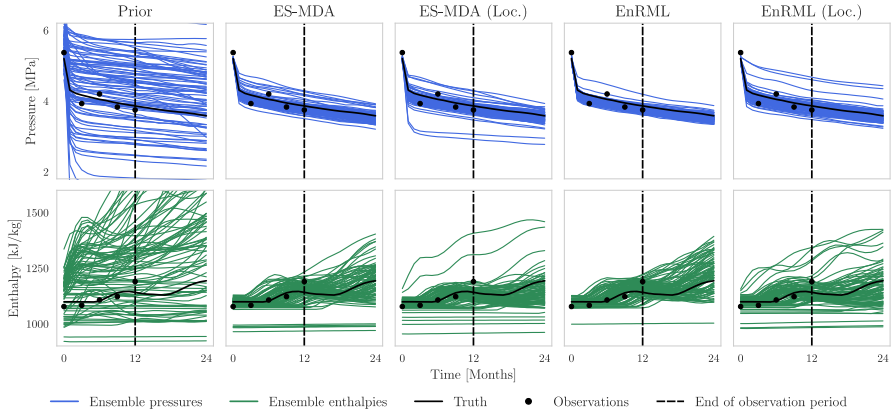

Figure 7. Pressures at well 1 and enthalpies at well 3, of the prior ensemble and the final ensemble obtained by each method.

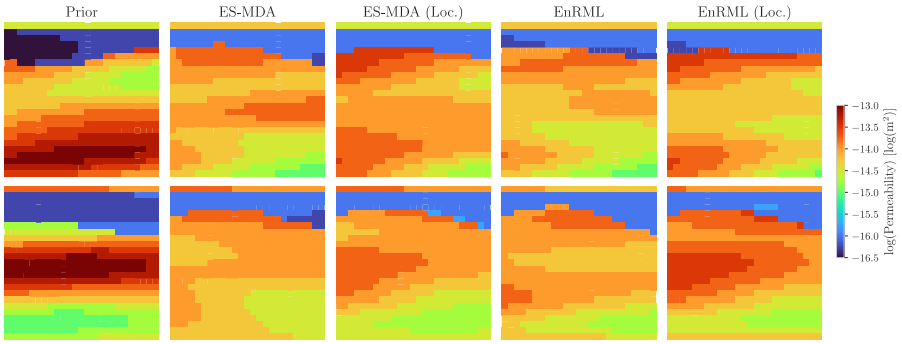

Figure 8. The first two members of the prior ensemble and the final ensembles of each method.

for the same set of parameters, by modifying the initial condition used to begin the natural state simulation. For the remainder of the work, we used the same initial condition for each simulation to minimise the effect of the non-uniqueness issue on the results; ideally, however, the initial condition should be estimated as part of the inversion process.

Another important consideration is the method by which each ensemble method is evaluated. Throughout this work, we have assessed each method by investigating whether the true parameters are contained within the distribution of the final ensemble. This, however, tells us little in terms of whether each method provides an accurate approximation to the posterior. To gain further insight into this, we would need to fully characterise the posterior using an exact sampling method, such as Markov chain Monte Carlo; this would provide a benchmark to which we could compare the estimates of each ensemble method. Similar studies have been conducted in the context of petroleum engineering (Emerick and Reynolds, 2013b; Iglesias et al., 2013b).

In our numerical testing, we carried out each inversion using the same model that was used to generate the synthetic data. However, in a real modelling context, there will always be differences between the model and the true physical processes by which the data is generated (for example, due to discretisation or neglected physics). Ideally, in future studies, model error should be introduced into the inversion (for instance, by generating the data using the finer mesh than the mesh used for the inversion).

Overall, we believe that ensemble methods have the potential to be useful tools for efficient reservoir model calibration, and should be investigated further.

ACKNOWLEDGEMENTS

The authors wish to acknowledge Elvar K. Bjarkason for numerous discussions on ensemble methods.

SUPPLEMENTARY MATERIAL

The code written as part of this paper is available at github.com/alexgdebeer/GeothermalEnsembleMethods.

REFERENCES

Aster, R.C., Borchers, B., Thurber, C.H. (2018). *Parameter Estimation and Inverse Problems* (3rd Ed.). Elsevier.

Békési, E., Struijk, M., Bonté, D., Veldkamp, H., Limberger, J., Fokker, P.A., Vrijlandt, M., & van Wees, J. (2020). An updated geothermal model of the Dutch subsurface based on inversion of temperature data. *Geothermics*, 88, 101880.

Bjarkason, E.K., Yeh, A., O'Sullivan, J.P., Croucher, A., & O'Sullivan, M.J. (2019). Non-uniqueness of geothermal natural-state simulations. *Proc. 41st New Zealand Geothermal Workshop*, Auckland, New Zealand.

Bjarkason, E.K., Maclaren, O.J., Nicholson, R., Yeh, A., & O'Sullivan, M.J. (2021). Uncertainty Quantification of Highly-Parameterized Geothermal Reservoir Models Using Ensemble-Based Methods. *Proc. World Geothermal Congress 2020+1*, Reykjavik, Iceland.

Chada, N.K., Iglesias, M.A., Roininen, L., & Stuart, A.M. (2018). Parameterizations for ensemble Kalman inversion. *Inverse Problems*, 34, 055009.

Chen, Y., & Oliver, D.S. (2013). Levenberg-Marquardt forms of the iterative ensemble smoother for efficient history matching and uncertainty quantification. *Computational Geosciences*, 17, 689-703.

Chen, Y., & Oliver, D.S. (2014). History Matching of the Norne Full-Field Model With an Iterative Ensemble Smoother. *SPE Reservoir Evaluation & Engineering*, 17(2), 244-256.

Chen, Y., & Oliver, D.S. (2017). Localization and regularization for iterative ensemble smoothers. *Computational Geosciences*, 21, 13-20.

Croucher, A., O'Sullivan, M., O'Sullivan, J., Yeh, A., Burnell, J., & Warwick, K. (2020). Waiwera: A parallel open-source geothermal flow simulator. *Computers & Geosciences*, 141, 104529.

Cui, T., Fox., C., & O'Sullivan, M. (2011). Bayesian calibration of a large-scale geothermal reservoir by a new delayed acceptance Metropolis Hastings algorithm. *Water Resources Research*, 47, W10521.

Emerick, A.A. (2016). Analysis of the performance of ensemble-based assimilation of production and seismic data. *Journal of Petroleum Science and Engineering*, 139, 219-239.

Emerick, A.A., & Reynolds, A.C. (2013a). Ensemble smoother with multiple data assimilation. *Computers & Geosciences*, 55, 3-15.

Emerick, A.A., & Reynolds, A.C. (2013b). Investigation of the sampling performance of ensemble-based methods with a simple reservoir model. *Computational Geosciences*, 17, 325-350.

Evensen, G. (2009). *Data Assimilation: The Ensemble Kalman Filter* (2nd Ed.). Berlin: Springer.

Iglesias, M.A., Law, K.J.H., & Stuart, A.M. (2013a). Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4), 045001.

Iglesias, M.A., Law, K.J.H., & Stuart, A.M. (2013b). Evaluation of Gaussian approximations for data assimilation in reservoir models. *Computational Geosciences*, 17, 851-885.

Iglesias, M., Park, M., & Tretyakov, M.V. (2018). Bayesian inversion in resin transfer molding. *Inverse Problems*, 34, 105002.

Iglesias, M., & Yang, Y. (2021). Adaptive regularisation for ensemble Kalman inversion. *Inverse Problems*, 37, 025008.

Kaipio, J., & Somersalo, E. (2006). *Statistical and Computational Inverse Problems.* New York: Springer.

Kitanidis, P.K. (1995). Quasi-Linear Geostatistical Theory for Inversing. *Water Resources Research*, 31(10), 2411-2419.

Le, D.H., Emerick, A.A., & Reynolds, A.C. (2015). An Adaptive Ensemble Smoother with Multiple Data Assimilation for Assisted History Matching. *Proc. SPE Reservoir Simulation Symposium*, Houston, Texas.

Luo, X., & Bhakta, T. (2020). Automatic and adaptive localization for ensemble-based history matching. *Journal of Petroleum Science and Engineering*, 184, 106559.

Maclaren, O.J., Nicholson, R., Bjarkason, E.K., O'Sullivan, J.P., & O'Sullivan, M.J. (2020). Incorporating Posterior-Informed Approximation Errors Into a Hierarchical Framework to Facilitate Out-of-the-Box MCMC Sampling for Geothermal Inverse Problems and Uncertainty Quantification. *Water Resources Research*, 56, e2018WR024240.

Muir, J.B., & Tsai, V.C. (2020). Geometric and level set topography using ensemble Kalman inversion. *Geophysical Journal International*, 220(2), 967-980.

Nicholson, R., Maclaren, O.J., O'Sullivan, J.P., O'Sullivan, M.J., Suzuki, A., & Bjarkason, E.K. (2020). Representation of Unknown Parameters in Geothermal Model Calibration. *Proc. 42nd New Zealand Geothermal Workshop*, Waitangi, New Zealand.

Nocedal, J., & Wright, S.J. (2006). *Numerical Optimization* (2nd ed.). New York: Springer.

Oliver, D.S., He, N., and Reynolds, A.C. (1996). Conditioning permeability fields to pressure data. *Proc. 5th European conference on the mathematics of oil recovery*, Leoben, Austria.

Omagbon, J., Doherty, J., Yeh, A., Colina, R., O'Sullivan, J., McDowell, J., Nicholson, R., Maclaren, O.J., & O'Sullivan, M. (2021). Case studies of predictive uncertainty quantification for geothermal models. *Geothermics*, 97, 102263.

Rafiee, J., & Reynolds, A.C. (2017). Theoretical and efficient practical procedures for the generation of inflation factors for ES-MDA. *Inverse Problems*, 33, 115003.

Stordal, A.S., & Elsheikh, A.H. (2015). Iterative ensemble smoothers in the annealed importance sampling framework. *Advances in Water Resources*, 86, 231-239.

Tso, C.H.M., Iglesias, M., Wilkinson, P., Kuras, O., Chambers, J., & Binley, A. (2021). Efficient multiscale imaging of subsurface resistivity with uncertainty quantification using ensemble Kalman inversion. *Geophysical Journal International*, 225(2), 887-905.

White, J.T. (2018). A model-independent iterative ensemble smoother for efficient history-matching and uncertainty quantification in very high dimensions. *Environmental Modelling & Software*, 109, 191-201.

White, J.T., Hunt, R.J., Fienen, M.N., & Doherty, J.E. (2020). Approaches to Highly Parameterized Inversion: PEST++ Version 5, a Software Suite for Parameter Estimation, Uncertainty Analysis, Management Optimization and Sensitivity Analysis. U.S. Geological Survey Techniques and Methods 7C26.

Zhang, Y., & Oliver, D.S. (2010). Improving the Ensemble Estimate of the Kalman Gain by Bootstrap Sampling. *Mathematical Geosciences*, 42, 327-345.

ISSN 2703-4275