

Bayesian inversion of DC electrical measurements with uncertainties for reservoir monitoring

Alberto Malinverno[†] and Carlos Torres-Verdín[‡]

[†] Schlumberger-Doll Research, Old Quarry Road, Ridgefield, CT 06877, USA

[‡] Department of Petroleum and Geosystems Engineering, The University of Texas at Austin, Austin, TX 78712-1061, USA

Received 2 May 2000

Abstract. This paper describes a method to infer the geometry of the saturation front (the boundary of the region swept by water) around a water injection well in an oil reservoir. Information on how injected water displaces oil in different layers is essential to best design a sweep of the reservoir. We use a Bayesian inference approach, which combines DC electrical measurements from an electrode array in the injection well with background information on the thicknesses and resistivities of each reservoir layer. The solution is the posterior probability density function of the saturation front location. We estimate its characteristics by a combination of nonlinear optimization and Monte Carlo sampling. Using synthetic data, we show how to jointly invert measurements acquired at different times after the beginning of injection and how the uncertainty in the inferred location of the saturation front depends on uncertainties in the background information available and on the amount of noise in the data.

1. Introduction

This paper describes a method to monitor a waterflood with DC electrical measurements acquired in an injector well. This type of monitoring provides information that is essential to understanding how injected water displaces oil and how to best design a sweep of the reservoir. To illustrate the method, we use a five-layer reservoir that matches the characteristics of an actual oil reservoir (figure 4). At first approximation, each reservoir layer can be divided into two zones: a low-resistivity cylindrical water bank, which has been swept by the injected water, and a high-resistivity, unswept oil bank. The measurements are taken with an array of 21 electrodes, where each electrode in turn acts as a direct current source while the electric potential is measured at the remaining electrodes. These electric potential measurements compose a vector d .

The reservoir parameters of interest in each layer are its thickness, the resistivity of the oil and water banks, and the radius $r(t)$ of the saturation front separating the two banks, which is a function of time. The parameters for all layers are listed in a model parameter vector m . There may be background information on some of these parameters; e.g., well logging runs may constrain the values of the oil/water bank resistivities. The general problem we wish to solve is to use background, prior information on m in the inversion of the DC measurements in d to improve our knowledge of the model parameters—in particular, of the radius of the saturation front in each layer as a function of time.

Numerous approaches have been put forth in the literature to estimate the distribution of electrical resistivity probed with DC electrical measurements. Most of these approaches

are based on the nonlinear minimization of a least-squares error functional by way of Gauss–Newton fixed-point iterations. A subsurface experiment similar to that investigated here was designed and implemented at the Richmond Field Station of the University of California, Berkeley [2] with the intent to monitor in time the diffusion of salt water injected into an aquifer. Daily and Owen [5] and Daily *et al* [6] developed a nonlinear inversion technique to interpret cross-borehole DC resistivity data acquired to monitor water movements in the vadose zone. A similar algorithm was applied by Ramirez *et al* [17] to monitor an underground steam injection process using electrical resistance tomography. Sasaki [18] developed and tested an efficient least-squares minimization algorithm that invoked source–receiver reciprocity relations to reduce the computational load associated with the numerical calculation of the Jacobian matrix.

A critical step in all these algorithms is the use of a specific model regularization scheme intended to control non-uniqueness and instability. The inversion of subsurface DC resistivity measurements is known to be non-unique (see for instance [18]), and the inversion process can also be very unstable in the presence of noisy measurements. Spies and Ellis [20] presented an extensive discussion of these two topics in the context of a cross-borehole resistivity tomography experiment designed to monitor the *in situ* vitrification of soil materials. They examined the effects of several choices of starting models, model constraints, and regularization strategies on the final inversion results. One of the most important results stemming from Spies and Ellis’s work was that, in the absence of suitable model constraints, the inversion of DC resistivity could yield a large collection of extreme resistivity models, all of them satisfying the data within the estimated bounds of accuracy.

Another way to view non-uniqueness is that there is an intrinsic uncertainty about m even after accounting for prior knowledge and for the information provided by the data in d . The Bayesian inference approach we adopt here quantifies uncertainties in the model parameters in a posterior probability density function of m . We will also show how this Bayesian approach can be applied to propagate information between inferences made at different times during the progress of injection and to effectively constrain the posterior probability density using all the measurements collected up to time t .

2. Bayesian inference

The fundamental formula in Bayesian inference is Bayes’ rule, which for our purpose can be written as [1, 8, 10, 13, 22]

$$p(\mathbf{m}|\mathbf{d}, \mathcal{I}) = \frac{p(\mathbf{m}|\mathcal{I})L(\mathbf{m}|\mathbf{d}, \mathcal{I})}{p(\mathbf{d}|\mathcal{I})} \propto p(\mathbf{m}|\mathcal{I})L(\mathbf{m}|\mathbf{d}, \mathcal{I}), \quad (1)$$

where \mathcal{I} denotes prior information, $p(\mathbf{m}|\mathbf{d}, \mathcal{I})$ is the posterior probability density, or PDF, of the model parameters, $p(\mathbf{m}|\mathcal{I})$ is the prior PDF (quantifying what is known on m from \mathcal{I} only), and $L(\mathbf{m}|\mathbf{d}, \mathcal{I})$ is the likelihood function (the PDF of the data when the parameter vector equals m). In other words, what can be inferred about the model parameters *a posteriori* is a combination of what is known *a priori*, independently of the data in d , and of the information contained in the data. As $p(\mathbf{d}|\mathcal{I})$ in equation (1) is a normalizing factor that does not depend on m , the posterior PDF can be simply written as being proportional to the product prior likelihood.

Assuming that the data contain normally distributed, uncorrelated additive noise, the likelihood function can be written as

$$L(\mathbf{m}, \sigma_e|\mathbf{d}, \mathcal{I}) \propto \exp \left\{ -\frac{1}{2\sigma_e^2} [\mathbf{d} - \mathbf{g}(\mathbf{m})]^\top [\mathbf{d} - \mathbf{g}(\mathbf{m})] \right\}, \quad (2)$$

where σ_e is the standard deviation of the data errors, T denotes the transpose, and $g(\mathbf{m})$ is a forward modelling operator that returns the computed value of the data when the model parameters equal \mathbf{m} ; we use the forward modelling algorithm of Tamarchenko and Druskin [21].

The prior PDF quantifies background information on the model parameters, and we set it equal to the PDF that has maximum entropy (i.e., maximum uncertainty) while being consistent with the information available *a priori* [12, 16]. Assuming that there is no prior information on correlations between the M model parameters, the maximum entropy prior PDF takes the form

$$p(\mathbf{m}, \sigma_e | \mathcal{I}) = p(m_1 | \mathcal{I}) p(m_2 | \mathcal{I}) \dots p(m_M | \mathcal{I}) p(\sigma_e | \mathcal{I}). \quad (3)$$

Prior PDFs for each reservoir model parameter (thickness, resistivity in the oil and water bank, saturation front radius of each layer), given different amounts of prior information, will be described in the examples. In equations (2) and (3) we added the standard deviation of the data errors σ_e to the parameters because we assume that it is not known *a priori*; its prior PDF is set to be proportional to $1/\sigma_e$ [3]. To keep the notation simple, we will implicitly include from here on σ_e in the parameter vector \mathbf{m} .

The application of Bayes' rule for a simple two-parameter vector $\mathbf{m} = (m_1, m_2)$ is illustrated in figure 1. The first row shows the application of Bayes' rule (equation (1)) with an initial broad prior PDF, corresponding to poor background knowledge of \mathbf{m} , and a likelihood function that is high for values of \mathbf{m} that give a good fit to the data collected at an initial time t_1 . The posterior PDF at t_1 combines these two types of information.

Information obtained at t_1 can be propagated to t_2 using the sequential property of Bayesian inference [3], as shown in the second row of figure 1. Writing as $d^{[i]}$ the measurements taken at time t_i , consider jointly the measurements $d^{[1]}$ and $d^{[2]}$ taken at t_1 and t_2 . If the errors in the measurements taken at different times are assumed to be independent, then the joint likelihood can be written as

$$L(\mathbf{m} | d^{[1]}, d^{[2]}, \mathcal{I}) = L(\mathbf{m} | d^{[1]}, \mathcal{I}) L(\mathbf{m} | d^{[2]}, \mathcal{I}), \quad (4)$$

and the posterior PDF considering all the data is

$$p(\mathbf{m} | d^{[1]}, d^{[2]}, \mathcal{I}) \propto \underbrace{p(\mathbf{m}) L(\mathbf{m} | d^{[1]}, \mathcal{I})}_{p(\mathbf{m} | d^{[1]}, \mathcal{I})} L(\mathbf{m} | d^{[2]}, \mathcal{I}). \quad (5)$$

In other words, the posterior PDF, given data $d^{[1]}$ and $d^{[2]}$, can be obtained from another application of Bayes' rule, with a prior PDF that is the posterior PDF for the data $d^{[1]}$ multiplied by the likelihood for the data $d^{[2]}$ only. In our case, some of the parameters in the vector \mathbf{m} (the saturation front radii) are a function of time, and the PDF quantifying their uncertainty at some time t_2 is not the same as the posterior PDF obtained at an earlier time t_1 . As shown in figure 1, the prior PDF at t_2 is an updated version of the posterior PDF at t_1 which accounts for the expected increase in the saturation front radii as injection progresses. A simple way to update the PDF of the saturation front radii will be described in the examples.

3. Method

In the previous section we described a general formulation that allows us to compute the value of the posterior PDF for any value of the parameter vector \mathbf{m} (except for a multiplicative constant). The posterior PDF, however, is a function in a high-dimensional model space; to efficiently represent and use in practice the posterior uncertainty, it is necessary to compute summary properties of the posterior PDF. For example, these properties may be the mode of

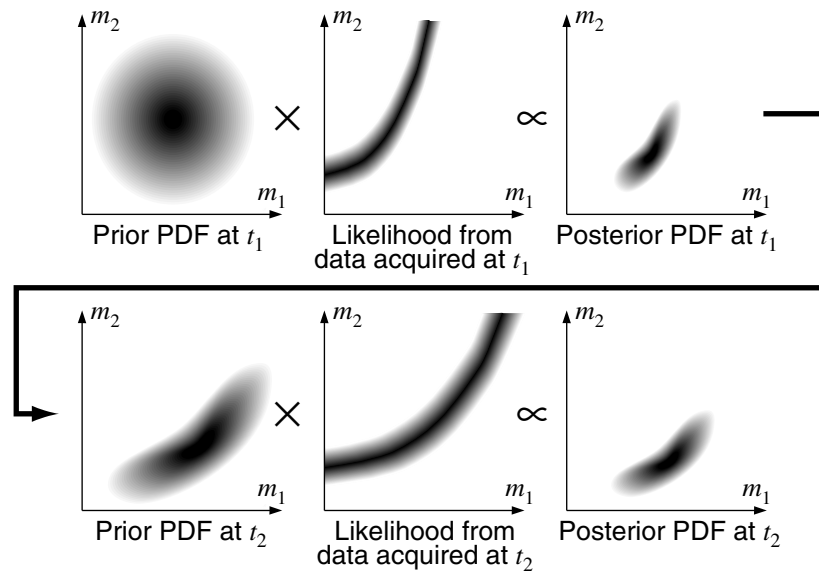


Figure 1. Graphical illustration of repeated applications of Bayes' rule for a simple two-parameter vector $\mathbf{m} = (m_1, m_2)$. Darker regions correspond to values of \mathbf{m} with higher probability. Modified after [8].

the posterior PDF, which is the most probable value of \mathbf{m} in light of the data and of the prior information; a posterior covariance matrix, which quantifies the uncertainty of \mathbf{m} after the data have been taken into account; marginal distributions of selected parameters in \mathbf{m} , which give the uncertainty of parameters of interest (e.g. the saturation front radii). There are two techniques described in the literature to obtain summary properties of the posterior PDF, based on nonlinear optimization [1, 10] and on Monte Carlo sampling [15, 19] (figure 2). In this section we describe an approach that combines these two techniques.

In the nonlinear optimization approach, one first finds the mode of the posterior PDF as the minimum of an objective function

$$\Psi(\mathbf{m}) = -\log[p(\mathbf{m}|\mathcal{I})L(\mathbf{m}|d, \mathcal{I})] \quad (6)$$

with an iterative, nonlinear optimization algorithm. The posterior covariance can then be estimated from the inverse of the Hessian matrix of the objective function (the matrix of second derivatives $\partial^2\Psi/\partial m_i\partial m_j$) evaluated around the minimum of $\Psi(\mathbf{m})$ by finite differences (figure 2, left). This implies assuming that the posterior PDF is well approximated by a multivariate normal distribution. The posterior PDF will be a multivariate normal if the prior PDF of the parameters is a multivariate normal and if the forward modelling operator $\mathbf{g}(\mathbf{m})$ is linear [23]. In the general nonlinear case, there is no assurance that the posterior PDF will be well approximated by a multivariate normal and the uncertainty implied by the posterior covariance obtained from the finite-difference Hessian matrix may be misleading.

The reason for this is illustrated in figure 3 for the PDF of a single parameter of interest. The posterior PDF is the sum of two normal distributions with standard deviations equal to 1 and 3, and its standard deviation is $\sqrt{5} \approx 2.24$. The nonlinear optimization approach to quantify posterior uncertainty amounts to locally fitting a normal PDF based on the second derivatives of the log-posterior PDF at the mode and extrapolating this fit away from the mode




Figure 2. In the nonlinear optimization approach (left), one first finds the mode of the posterior PDF by iterative minimization (black dots and arrows), and then obtains a posterior covariance matrix from finite-difference second derivatives about the mode (white dots). In the Monte Carlo approach (right), one obtains a sample of parameter vectors \mathbf{m} drawn from the posterior PDF, and uses this sample to compute properties of interest (e.g. the posterior covariance matrix).

(figure 3, left). Unless the normal approximation is valid away from the mode, this method may give a poor estimate of the actual posterior uncertainty. (This would also happen if the posterior PDF were substantially asymmetrical.)

An alternative method to estimate uncertainty is Monte Carlo sampling (figure 3, right). In this method, a Monte Carlo algorithm (described below) obtains a sample drawn from the posterior PDF, effectively behaving as a random number generator that returns a sample of parameter values distributed as in the posterior PDF. The estimated posterior PDF shown as a thin solid curve in the right panel of figure 3 is the distribution of 10 000 samples of the parameter obtained by the Monte Carlo method, and closely approximates the actual posterior PDF. The Monte Carlo method does not depend on an assumption of normality or symmetry for the posterior PDF, and gives a more accurate quantification of posterior uncertainty at the price of greater computational effort.

The Monte Carlo algorithm typically used in geophysical applications [15, 19] is the Metropolis–Hastings algorithm [4, 11, 14], which is a random walk in model space consisting of two steps. In the first step, the current value of the parameter vector is modified at random to obtain a ‘candidate’ vector \mathbf{m}' . This candidate is drawn from a ‘proposal’ PDF $q(\mathbf{m}'|\mathbf{m})$, where the choice of \mathbf{m}' depends on the current vector \mathbf{m} . In the second step, the parameter vector is changed to \mathbf{m}' with an ‘acceptance probability’

$$\alpha = \min \left[1, \frac{p(\mathbf{m}'|\mathbf{d}, \mathcal{I}) q(\mathbf{m}|\mathbf{m}')}{p(\mathbf{m}|\mathbf{d}, \mathcal{I}) q(\mathbf{m}'|\mathbf{m})} \right]. \quad (7)$$

If one were to always accept the candidates drawn from the proposal PDF, the chain would take a random walk in the space of \mathbf{m} . The acceptance probability in equation (7) effectively biases the random walk toward regions of parameter space that have higher posterior probabilities. The proposal PDF ratio term in equation (7) applies a correction if the move from \mathbf{m} to \mathbf{m}' has a different probability of being proposed than the opposite move from \mathbf{m}' to \mathbf{m} . It can be proven that this algorithm will asymptotically sample the posterior PDF as graphically illustrated in the right panel of figure 2.

The choice of the starting point for \mathbf{m} and of the proposal PDF can have a large effect on the efficiency of the Metropolis–Hastings algorithm. If the starting point is far from the mode of the posterior PDF, it may take a long time for the random walk to reach the high-probability region of the posterior PDF. Once the current parameter vector \mathbf{m} is in this high-probability region, if the proposed modifications to \mathbf{m} are

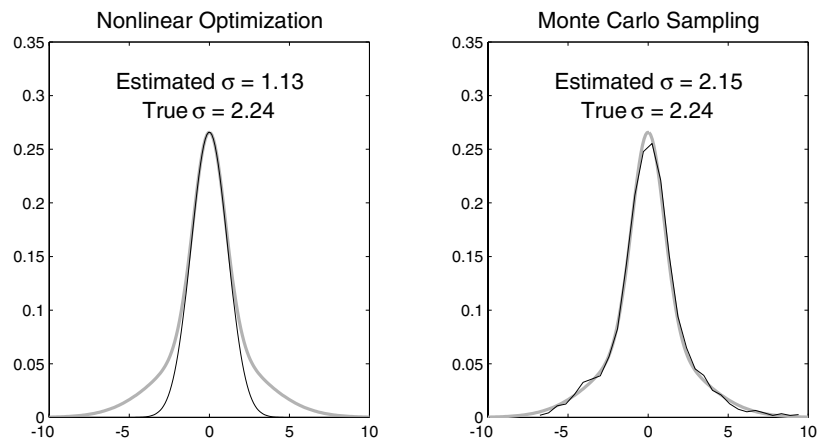


Figure 3. The actual posterior PDF of a single parameter (thick grey curve) has longer tails than a normal distribution with the same characteristics near the mode. The posterior PDF estimated from the second derivative of the log-posterior PDF near the mode underestimates the posterior uncertainty; the inferred posterior standard deviation σ is much lower than the actual value (left). Monte Carlo sampling (right) instead obtains an estimate of the posterior PDF (thin solid curve) that is much closer to the actual posterior PDF.

large compared with the spread of the posterior PDF, the candidate m' will typically have a low posterior probability, α in equation (7) will be small, and it will take many iterations before a proposed move is accepted. At the other extreme, if the changes proposed are very small α will be large, but many iterations will be needed to cover the high-posterior probability region of the parameter space [4]. The best choice for $q(m'|m)$ would be a PDF that approximates as closely as possible the posterior PDF of m ; in this case, the probability of acceptance α in equation (7) would always be near one.

The method we use here to evaluate characteristics of the posterior PDF is a combination of the nonlinear optimization and of the Monte Carlo approach. We first obtain the mode of the posterior PDF by nonlinear iterative optimization using a quasi-Newton method [7]. We then run the Metropolis–Hastings algorithm starting the random walk from the posterior mode and using as the proposal PDF a multivariate normal distribution centred on the current value of m and with a covariance matrix equal to the inverse of the Hessian matrix evaluated by finite differences at the posterior mode. This ensures that the sampling starts in the high-probability region of the posterior PDF and effectively uses a proposal PDF that approximates the posterior and thus maximizes the probability of acceptance.

We run the Monte Carlo sampling by starting 20 independent random walkers from the posterior mode, as in the multiple sequence approach suggested by Gelman and Rubin [9]. Using multiple random walkers means that the algorithm could be immediately parallelized by running each random walk on a separate processor. We then keep sampling until the posterior standard deviations of the parameters of interest stop changing. In practice, we keep track of how much the posterior standard deviations of all parameters change each time as many candidates are accepted as the number of model parameters. We stop sampling when the largest change in the posterior standard deviations is less than one in a thousand.




Figure 4. Illustration of the waterflood geometry considered in this paper. The main parameter of interest is the radius $r(t)$ of the saturation front separating the oil and water bank in each reservoir layer (dashed white line).

4. Examples

In this section we show sample applications of the method to compute the posterior PDF of the saturation front radii from noisy synthetic data. The purpose of these examples is to determine how uncertainties in prior information and noise in the data translate into posterior uncertainties in the parameters of interest. The synthetic data were computed at three different times t_1 , $t_2 = 2t_1$, and $t_3 = 2t_2$ (e.g., one, two and four weeks after water injection starts) for a uniform injection rate. The synthetic data are the electrical potentials measured at t_1 , t_2 , and t_3 minus the potential measured before the start of injection, and thus are measurements of the change in electrical potential due to water injection. We also added normally distributed noise with a standard deviation $\sigma_e = 5$ mV (about 10% the standard deviation of the signal). While each electrode in turn acts as a source injecting 1 A of current, the electrical potential is measured at all other electrodes, which act as receivers. The electrodes are placed outside non-conductive casing and have a return electrode at the surface, so that they act as single poles. The measurements can be arranged in a 21×21 matrix and are shown in figure 5 for the three configurations of the saturation front radii at t_1 , t_2 , and t_3 .

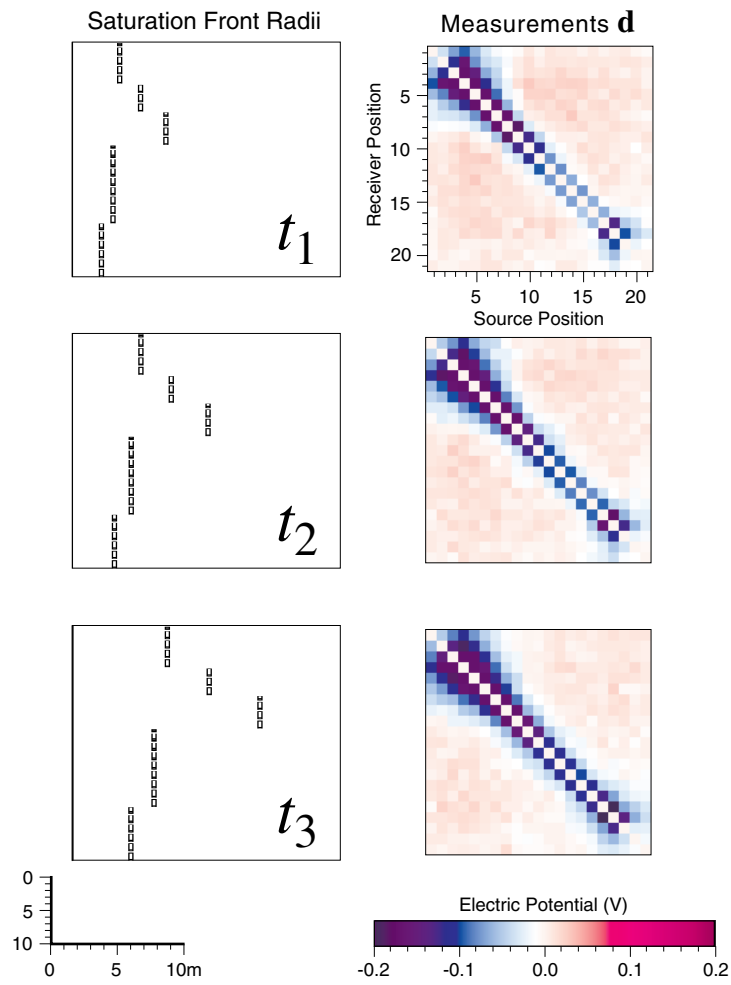


Figure 5. Synthetic noisy data ($\sigma_e = 5$ mV) computed for three configurations of the saturation front radii (dashed white lines) at t_1 , $t_2 = 2t_1$, and $t_3 = 2t_2$ for a uniform injection rate.

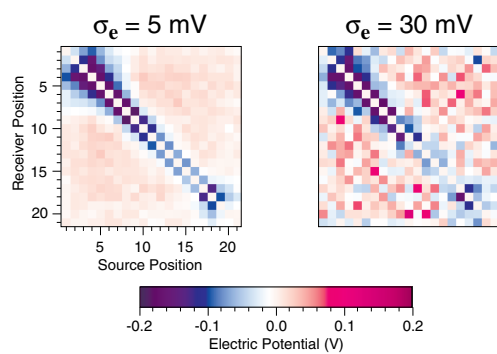


Figure 6. Comparison of synthetic data at t_1 with a noise standard deviation of 5 mV (left) and 30 mV (right).

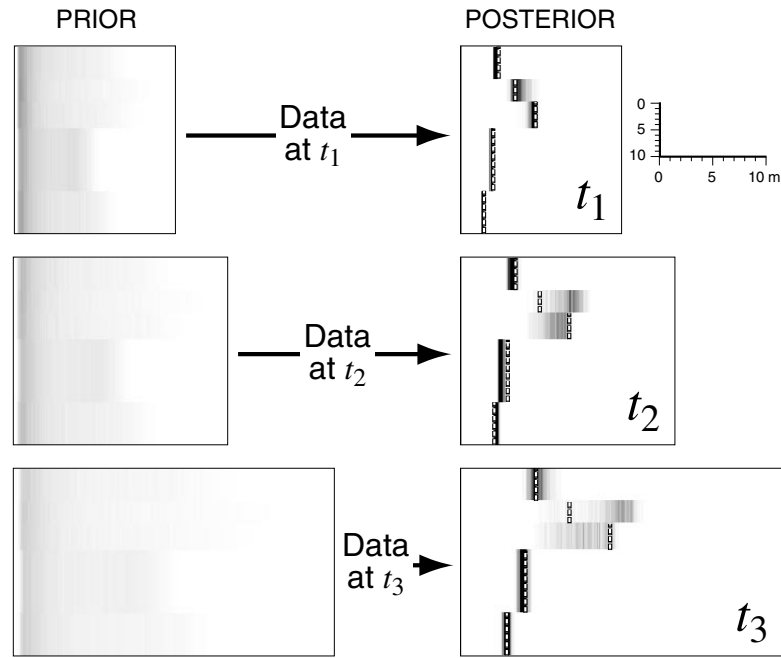


Figure 7. Images of a vertical cross section of the saturation front location obtained by sampling the prior (left) and posterior PDFs (right). The thicknesses and resistivities of each layer are assumed known and fixed at their true value. The posterior PDFs at each time are computed using only the data collected at that time. The dashed white line shows the true location of the saturation front.

4.1. Separate inversions at t_1 , t_2 , and t_3

In this section we estimate posterior PDFs of the saturation front radii at t_1 , t_2 , and t_3 using only the DC measurements collected at that time. The layer thicknesses and the oil/water bank resistivities were assumed known and fixed to their true value. The prior PDF of the radii of the saturation front was obtained from prior knowledge of the amount of water injected and of reservoir parameters. When a volume V_w of water has been injected, the volume V of reservoir rock swept by water is

$$V = \frac{V_w}{\phi(1 - S_{or} - S_{wi})} \quad (8)$$

where ϕ is the porosity, S_{or} is the residual oil saturation and S_{wi} is the irreducible water saturation, so that $(1 - S_{or} - S_{wi})$ is the fraction of pore space swept by water. We assume here that prior information on porosity and saturations allows us to set a Gaussian prior PDF for V with a mean equal to the expected swept reservoir volume and a standard deviation equal to some percentage of this mean (we use 20% here). The prior PDF of any set of five saturation front radii was then simply set to the prior PDF of the corresponding swept volume V .

The prior and posterior images in figure 7 were constructed by superimposing a large number of saturation front radii sampled from the prior and posterior PDFs, respectively. The darker areas correspond to areas of higher probability, and the blurred nature of the image gives a visual measure of the posterior uncertainty of the saturation front radii. While many possible configurations of the saturation front are possible *a priori*, the posterior images are generally restricted to locations that are near the true locations of the saturation front (shown by the dashed white line).

At t_2 and t_3 , however, the radii of the saturation front in the second and third layers from the top are poorly determined: the most probable value *a posteriori* is quite different from the true value and there are large uncertainties. This is because in these relatively thin layers the saturation front has moved the farthest from the electrode array. The measurements are thus rather insensitive to variations in the saturation front radius in these layers, and even small amounts of measurement noise can shift significantly the most probable location *a posteriori* of the saturation front. We repeated the inversion for different random noise samples added to the data, and verified that, in fact, the most probable posterior saturation radii can change considerably in the second and third layers from the top at t_2 and t_3 . In these cases it is then very important to have an estimate of the posterior uncertainty, which shows how closely the data constrain different features of the solution.

4.2. Joint inversions at t_2 and t_3

A better result can be obtained by using information obtained at previous times. In fact, if one knows at time t_1 that the saturation front radii have the posterior uncertainties shown in the first row of figure 7, then it seems unreasonable to assume that the saturation front at t_2 is determined as poorly as in the prior image in the second row. As noted earlier, information obtained at t_1 can be propagated to t_2 by setting the prior PDF at t_2 from the posterior PDF at t_1 (equation (5)).

This sequential approach has been used to obtain the results shown in figure 8, where the prior PDFs for the saturation front radii at t_2 and t_3 are multivariate normal distributions obtained from the posterior PDFs at t_1 and t_2 , respectively. As the radii of the saturation front are expected to increase proportionally to the square root of time for uniform injection rates, the means and standard deviations of the prior PDF at t_i are the means and standard deviations of the posterior PDF at t_{i-1} multiplied by $\sqrt{t_i/t_{i-1}}$. This effectively amounts to setting the prior PDF at t_i from a prediction made on the basis of the results obtained at t_{i-1} . Using this additional information is equivalent to carrying out a joint inversion of all data acquired up to time t_i , and reduces the posterior uncertainty in the saturation front radii at t_2 and t_3 (compare figures 7 and 8).

4.3. Effect of less prior information

In the next two examples we examine how the posterior PDF at t_1 is affected by different amounts of prior information on the parameters and of measurement noise. In the results of figures 7 and 8, it was assumed that the thicknesses and the resistivities in the oil and water bank for each layer were so well constrained by prior information that these parameters could be fixed in the inversion. In this section we examine the effect of *a priori* uncertainty in layer thicknesses and resistivities.

Figure 9 shows results obtained at t_1 when there are prior uncertainties on all the reservoir model parameters besides the saturation front radii. We assumed that the thickness and the water bank resistivity of each layer were relatively well constrained by prior information, and we used as the prior PDF a normal distribution with a mean equal to the true value and a standard deviation of 0.2 m for layer thicknesses and 0.05 Ω m for water bank resistivities. The oil bank resistivities were instead assumed to be poorly known, and their prior PDF was set to an exponential with a mean of 10 Ω m; this is the maximum entropy PDF for a positive parameter where one knows what value to expect, but has no information on how much the parameter may depart from the expected value [16].

If the standard deviation of the measurement noise σ_e is the same as in all previous




Figure 8. Images of a vertical cross section of the saturation front location obtained by sampling the prior (left) and posterior PDFs (right). The thicknesses and resistivities of each layer are assumed known and fixed at their true value. The posterior PDFs at t_2 and t_3 are computed with a prior PDF determined by the posterior obtained previously. The black arrows indicate the flow of information, and the dashed white line shows the true location of the saturation front.

Layer Thicknesses and Resistivities

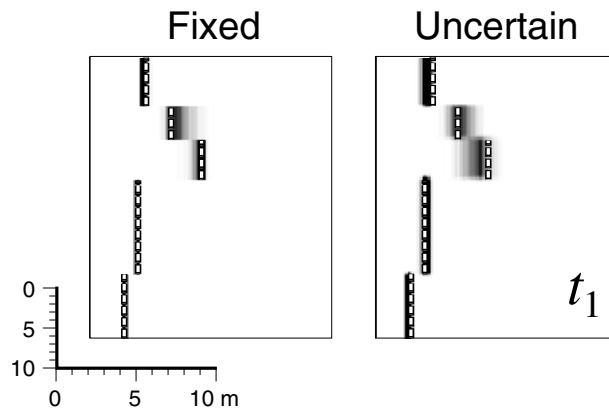


Figure 9. Images of a vertical cross section of the saturation front location at t_1 obtained by sampling the posterior PDF for a prior PDF with fixed (left) and uncertain (right) layer thicknesses and resistivities; see text for details. The dashed white line shows the true location of the saturation front.

examples (5 mV), the posterior PDF at t_1 shows only a slight increase in uncertainty compared with the posterior PDF obtained when layer thicknesses and resistivities were fixed to their true value (as in the first row of figure 8). The posterior PDFs of the oil bank resistivities in each of

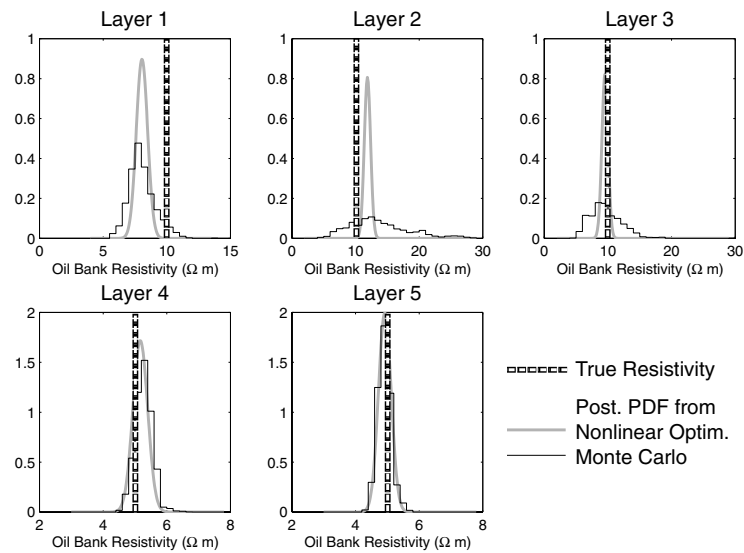


Figure 10. Posterior PDFs of the oil bank resistivities in each layer obtained from nonlinear optimization (thick grey curve) and from Monte Carlo sampling (thin black curve). Layer 1 is at the top of the reservoir, and layer 5 at the bottom. The dashed white lines show the true oil bank resistivities.

the five reservoir layers are shown in figure 10. In the two deepest reservoir layers (4 and 5), the data determine relatively well the values of the oil bank resistivities; the posterior mean is near the true value, and the posterior standard deviation is about 10%. In the three shallowest layers, the oil bank resistivities are not as well determined by the data, although they are clearly better defined than in the prior PDF (an exponential with a mean of 10 Ω m in all layers).

Figure 10 also shows that the posterior uncertainty obtained in the nonlinear optimization approach, assuming a normal posterior PDF, can be quite a bit smaller than that implied by Monte Carlo sampling, as in the hypothetical example of figure 3. For example, in the first three layers the posterior PDFs from Monte Carlo sampling are much broader than those obtained from nonlinear optimization. In the first layer, the posterior PDF from nonlinear optimization implies that the true value of the resistivity (10 Ω m) should be very improbable *a posteriori*. These results demonstrate that Monte Carlo sampling may be necessary to accurately quantify the posterior uncertainty.

4.4. Effect of measurement noise

Finally, we examine the effect of a relatively large amount of measurement noise on the posterior uncertainties of the saturation front radii. We repeated the inversion done in the previous example for data with an added noise whose $\sigma_e = 30$ mV, about 50% the standard deviation of the signal (figure 6). More measurement noise results in corresponding larger posterior uncertainties, but the main features of the saturation front are still recognizable (figure 11).




Figure 11. Images of a vertical cross section of the saturation front location at t_1 obtained by sampling the posterior PDF for data with a noise standard deviation of 5 mV (left) and 30 mV (right). The thicknesses and resistivities of each layer have prior uncertainties (see text). The dashed white line shows the true location of the saturation front.

5. Conclusions

We showed how to apply Bayesian inference to determine uncertainties in the position of the saturation front during an oil reservoir waterflood given initial uncertainties and noisy DC electrical measurements. This approach provides a measure of how well constrained are the inferred positions of the saturation front, given uncertainties in prior information, and allows for combining data collected at different times.

Our preferred method combines nonlinear optimization and Monte Carlo sampling to accurately quantify the posterior uncertainties. We found that in some cases Monte Carlo sampling gave a more realistic assessment of posterior uncertainties compared with a nonlinear optimization approach where a posterior covariance matrix is obtained from the inverse Hessian matrix of the log-posterior PDF evaluated near the mode. The Monte Carlo sampling is run by launching a number of independent random walkers, and the algorithm could be easily parallelized by assigning each random walk to a separate processor.

Besides being applicable to the inversion of field measurements, the approach presented here can be useful for experimental design, e.g. to estimate the amount of prior information, the acceptable level of noise in the measurements, and the array geometry needed to infer the position of the saturation front to a desired level of accuracy.

References

- [1] Bard Y 1974 *Nonlinear Parameter Estimation* (New York: Academic)
- [2] Bevc D and Morrison H F 1991 *Geophysics* **56** 769–77
- [3] Box G E P and Tiao G C 1973 *Bayesian Inference in Statistical Analysis* (Reading, MA: Addison-Wesley)
- [4] Chib S and Greenberg E 1995 *Am. Statistician* **49** 327–35
- [5] Daily W and Owen E 1991 *Geophysics* **56** 1228–35
- [6] Daily W, Ramirez A, LaBrecque D and Nitao J 1992 *Water Resources Res.* **28** 1429–42
- [7] Dennis J J E and Schnabel R B 1983 *Numerical Methods for Unconstrained Optimization and Nonlinear Equations* (Englewood Cliffs, NJ: Prentice-Hall)
- [8] Duijndam A J W 1988 *Geophys. Prosp.* **36** 878–98
- [9] Gelman A and Rubin D B 1992 *Stat. Sci.* **7** 457–511
- [10] Gouveia W P and Scales J A 1998 *J. Geophys. Res.* **103** 2759–79
- [11] Hastings W K 1970 *Biometrika* **57** 97–109
- [12] Jaynes E T 1968 *IEEE Trans. Syst. Sci. Cybern.* **4** 227–41

- [13] Jaynes E T 1984 *Inverse Problems* ed D W McLaughlin (Providence, RI: American Mathematical Society) pp 151–66
- [14] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 *J. Chem. Phys.* **21** 1087–92
- [15] Mosegaard K and Tarantola A 1995 *J. Geophys. Res.* **100** 12 431–47
- [16] Papoulis A 1984 *Probability, Random Variables, and Stochastic Processes* (New York: McGraw-Hill)
- [17] Ramirez A, Daily W, LaBrecque D, Owen E and Chestnut D 1993 *Water Resources Res.* **29** 73–87
- [18] Sasaki Y 1992 *Geophys. Prosp.* **40** 453–63
- [19] Sen M K and Stoffa P L 1995 *Global Optimization Methods in Geophysical Inversion* (Amsterdam: Elsevier)
- [20] Spies B R and Ellis R 1995 *Geophysics* **60** 886–98
- [21] Tamarchenko T and Druskin V 1993 *Proc. SPWLA 34th Annual Logging Symp.* (Calgary, Alberta: The Society of Professional Well Log Analysts) paper GG
- [22] Tarantola A and Valette B 1982 *J. Geophys.* **50** 159–70
- [23] Tarantola A and Valette B 1982 *Rev. Geophys. Space Phys.* **20** 219–32