Near Surface Geophysics, 2018, **16**, 206-217 doi: 10.3997/1873-0604.2017064

Studies of parameter correlations in surface NMR using the Markov chain Monte Carlo method

Kristoffer R. Andersen[2], Ling Wan[1,2*], Denys Grombacher[2], Tingting Lin[1*] and Esben Auken[2]

[1] *College of Instrumentation and Electrical Engineering, Jilin University, Changchun 130061, China*
[2] *HydroGeophysics Group, Department of Geoscience, Aarhus University, 8000 Aarhus, Denmark*

Received July 2017, revision accepted November 2017

ABSTRACT
Surface nuclear magnetic resonance is a technique capable of providing insight into subsurface aquifer properties. To produce estimates of aquifer properties (such as the spatial distribution of water content and parameters controlling the duration of the nuclear magnetic resonance signal), an inversion is required. Essential to the reliable interpretation of the estimated subsurface models is an understanding of the uncertainty and correlation between the parameters in the estimated models. To quantify parameter uncertainty and correlation in the surface nuclear magnetic resonance inversion, a Markov chain Monte Carlo approach is demonstrated. Markov chain Monte Carlo approaches have been previously employed to invert surface nuclear magnetic resonance data, but the primary focus has been on quantifying parameter uncertainty. The focus of this paper is to further investigate whether the parameters in the estimated models exhibit correlation with one another; equally important to building a reliable interpretation of the subsurface is an understanding of the parameter uncertainty. The utility of the Markov chain Monte Carlo approach is demonstrated through the investigation of three questions. The first question investigates whether the parameters describing the water content and thickness of a layer exhibit a strong correlation. This question stems from applying concepts known to electromagnetic surveys (that the layer thickness and layer resistivity parameters are strongly correlated) to the surface nuclear magnetic resonance inversion. A water content–layer thickness correlation in surface nuclear magnetic resonance would not have large effects for quantifying total water content but would affect the ability to identify layer boundaries. The second question examines whether the parameter controlling the duration of the nuclear magnetic resonance signal exhibits a correlation with the water content and layer thickness parameters. The resolution of surface nuclear magnetic resonance typically does not consider the duration of the signal and focuses primarily on the distribution of current amplitudes that form the suite of transmit pulses. It is common to treat regions with short-duration signal with greater uncertainty, but it is important to understand whether the signal duration controls resolution for medium to long duration signals as well. The third question explores if the parameter uncertainty produced by the Markov chain Monte Carlo approach is consistent with that produced by an alternative approach based upon the posterior covariance matrix (for the linearised inversion). The ability of the Markov chain Monte Carlo approach to more thoroughly explore the model space provides a means to improve the reliability of surface nuclear magnetic resonance aquifer characterisations by quantifying parameter uncertainty and correlation.

INTRODUCTION

Surface nuclear magnetic resonance (NMR) is a non-invasive geophysical technique for groundwater investigations providing direct sensitivity to water content and detailed insight into aquifer properties such as pore size, water content, and permeability (Legchenko *et al.* 2002; Braun and Yaramanci 2008; Costabel and Yaramanci 2013). The measurement involves the use of a coil of wire laid on the ground surface to perturb and measure the properties of a magnetisation present at a depth that originates from hydrogen nuclei contained within the groundwater. To produce images of subsurface aquifer properties, an inversion protocol is required, where the data consist of NMR decays measured

* wanling@jlu.edu.cn; ttlin@jlu.edu.cn

after a series of excitation pulses with varying pulsed current amplitudes. Many inversion approaches have been applied to surface NMR data to produce estimates of the spatial distribution of water content and relaxation times (i.e., a parameter governing the duration of the signal), most commonly to produce 1D depth profiles. Two- and three-dimensional surveys have also been performed (Girard *et al.* 2007; Dlugosch *et al.* 2013, 2014; Chevalier *et al.* 2014), but 1D profiles remain the standard output of a surface NMR survey. Initial value inversion uses only the initial amplitudes to produce the estimates of the subsurface water content distribution (Legchenko and Shushakov 1998; Weichman, Lavely and Ritzwoller 2000; Guillen and Legchenko 2002). Alternatively, the full dataset consisting of all the NMR measured decays at all times can be inverted to produce both water content and relaxation time depth profiles (Mohnke and Yaramanci 2002; Müller-Petke and Yaramanci 2010). Joint inversion of surface NMR and electrical/electromagnetic (EM) data has also been demonstrated (Hertrich and Yaramanci 2002; Legchenko *et al.* 2009; Behroozmand *et al.* 2012).

Each inversion protocol has been demonstrated to routinely produce robust estimates of the subsurface water content's spatial distribution that fit observed data (as well as supplementary geologic information). However, these inversions typically neglect the question of parameter uncertainty and correlation. In many geophysical inversion problems, obtaining the parameter uncertainty and correlation is often as important as getting the parameter values themselves (Tarantola and Valette 1982b). Behroozmand *et al.* (2013) described a strategy to determine parameter uncertainty in a joint magnetic resonance spectroscopy and transmission electron microscopy (TEM) data analysis scheme based on the posterior covariance matrix. An alternative approach to estimate parameter uncertainty and correlation is the Markov chain Monte Carlo (MCMC) method, which has been employed to study uncertainty for frequency-domain EM surveys (Minsley 2011) and study uncertainty for T_2 estimates (Prange and Song 2010). An MCMC-based inversion scheme probes the parameter space searching for models consistent with the data, returning a suite of models consistent with the data instead of returning a single best-fit model (Malinverno 2002; Sambridge and Mosegaard 2002). Examination of parameter variation within the suite of returned models provides insight into parameter uncertainty and correlation. From this perspective, the MCMC method offers great potential for parameter uncertainty and correlation analysis for surface NMR. Guillen and Legchenko (2002) implemented an MCMC inversion for surface NMR providing the ability to characterise uncertainty in the estimated water content and relaxation time profile. Chevalier *et al.* (2014) extended the surface NMR MCMC inversion scheme to three dimensions providing estimates of water content uncertainty. Similarly, the simulated annealing approach of Mohnke and Yaramanci (2002) provides a non-deterministic inversion scheme where a suite of water content and decay time models fitting the data equally well are produced, providing uncertainty quantification. Approaches involving bootstrapping of surface NMR data have also been used to quantify uncertainty in the predicted water content and relaxation time profiles (Hertrich 2008; Parsekian and Grombacher 2015). Each of the previous approaches focuses primarily on the estimated model parameters' uncertainties. Equally important for accurate interpretation of the resulting profiles is an understanding of parameter correlation. The focus in this paper is to highlight the use of an MCMC approach to investigate parameter correlations, in addition to uncertainties, for surface NMR studies.

An MCMC framework is developed for surface NMR to investigate parameter correlations. The algorithm is adapted from the work of Minsley (2011) in frequency-domain electromagnetics. However, we only investigate models with a fixed number of model parameters. The utility of such a framework is demonstrated through the investigation of three questions. (1) Does a water content–layer thickness equivalency exist in surface NMR (similar to the resistivity–thickness equivalency present in EM)? (2) Does the relaxation time, T_2*, play a role in controlling the resolution of the estimated profiles? (3) Do estimates of the parameter uncertainty based on posterior covariance matrix reproduce MCMC uncertainties? Synthetic results are presented to investigate each question and to demonstrate the utility of MCMC for surface NMR parameter uncertainty/correlation studies.

Surface nuclear magnetic resonance

The surface NMR survey involves the perturbation and subsequent measurement of a magnetisation present in the subsurface that originates from the immersion of hydrogen nuclei within the Earth's magnetic field. To perturb the magnetisation, a strong oscillatory current is pulsed in a coil of wire at the surface in order to generate a secondary magnetic field. This secondary field induces torque on the magnetisation, perturbing it out of its equilibrium orientation (Bloch 1946). The standard surface NMR experiment involves measuring the magnetisation's return to equilibrium following a single current pulse (called a free-induction decay). Only the component of the magnetisation transverse to the direction of the Earth's field can be directly measured; this component precesses about the Earth's field direction at the Larmor frequency, allowing its magnitude to be measured inductively by a coil at the surface. Typically, the coil used to pulse the oscillatory current is also used to measure the subsequent signal. The forward model for this type of experiment, i.e., a free-induction decay using a coincident transmit/receive loop, is given by (Weichman *et al.* 2000)

$$V(q,t)=\omega_0 M_0 \int_{vol} m_\perp(\mathrm{r},q) B_\perp^-(\mathrm{r}) e^{2i\xi(\mathrm{r})} \omega(\mathrm{r}) e^{-\frac{t}{T_2^*(\mathrm{r})}} d^3 r \,, \qquad (1)$$

where $V(q,t)$ represents the loop voltage measured at time t following a current pulse of pulse moment q; q is equal to the product of the pulse duration τ and the current amplitude I. M_0 is the equilibrium magnetisation amplitude, and $m_\perp(\mathrm{r})$ is the magnitude of the transverse magnetisation of a unit magnetisation following the excitation pulse at location $\boldsymbol{r}$. For a typical surface NMR survey employing an on-resonance excitation pulse $m_\perp(\mathrm{r}) = \sin(\gamma B_\perp^+(\mathrm{r})q)$, where $B_\perp^+(\mathrm{r})$ is the amplitude of the co-rotating component of the secondary magnetic field perpendicular to the direction of Earth's field (Weichman *et al.* 2000), only the co-rotating component of the field contributes to the perturbation of the magnetisation. The γ term is the gyromagnetic ratio of the hydrogen nuclei. The $B_\perp^-(\mathrm{r})$ term represents the sensitivity of the receive coil and is the counter-rotating component of the secondary magnetic field perpendicular to the direction of Earth's field given a unit current. The exponential term containing ξ describes the phase shift of the measured signal due to the conductivity structure of the subsurface (Trushkin, Shushakov and Legchenko 1995). The spatial distribution of water content is given by $w(r)$, where $0 \le w(r) \le 1$. The exponential term containing $T_2^*(\mathbf{r})$ describes the decaying envelope of the measured signal, where $T_2^*(\mathbf{r})$ is a time constant called the effective transverse relaxation time.

The goal of the surface NMR experiment is to convert the measured $V(q,t)$ into estimates of $w(\mathbf{r})$ and $T_2^*(\mathbf{r})$. This requires an inversion as the spatial origin of the measured signals cannot be directly determined given that many of the signals contain overlapping spatial origins. Surface NMR inversions can also estimate the conductivity structure of the subsurface using only surface NMR data (Braun and Yaramanci 2008) or through joint inversion with electrical or EM data (Behroozmand *et al.* 2012). However, a common approach is to treat the conductivity structure as fixed during the inversion (where it is estimated using a complementary electrical or EM survey). In this case, the kernel function is independent of the inverted parameters ($w(\mathbf{r})$ and $T_2^*(\mathbf{r})$) and is fixed during the inversion. This greatly accelerates the forward modelling and increases the inversion speed; if the conductivity structure is allowed to vary, the kernel must be recalculated each iteration.

Several inversion schemes are commonly employed in surface NMR involving either a sequential inversion where water contents and relaxation times are estimated separately or an inversion where the water contents and relaxation times are simultaneously estimated. Generally, these parameters are treated as independent during the inversion and subsequent interpretation. Essential to reliable interpretation of the resulting water content and relaxation time profiles is that the uncertainty in each profile is characterised. Several approaches that quantify uncertainty have been demonstrated previously (Guillen and Legchenko 2002; Mohnke and Yaramanci 2002; Hertrich 2008; Behroozmand *et al.* 2013; Chevalier *et al.* 2014; Parsekian and Grombacher 2015). However, in addition to the uncertainty of each parameter, it is also important to investigate parameter correlations. One approach providing both parameter uncertainty and correlations is the MCMC approach.

The Markov chain Monte Carlo method

The MCMC method investigates the distribution of parameters that fit some data. We will describe our setup and implementation, but for a more general introduction to MCMC, readers are referred to Brooks *et al.* (2011). Consider a single walker at a specific location in the model space defined by the model vector m. A new model m' is proposed solely on the basis of m and some predefined distribution. For this new model, we evaluate the probability, and based on some specific criteria (described later), we determine whether or not to keep it. This process continues until a number of iterations have been performed or convergence has been reached. There are several important aspects to consider when doing MCMC inversions. One common challenge is that walkers may get stuck in a local minimum and need many iterations to reach other local minima or the global minimum. It is difficult to ensure that such local minima have not been encountered, but the safest thing to do is to continue the exploration for millions of iterations. A method that can be used to find the global minimum among different local minima is to start with a range of walkers at different positions (i.e., different starting models). If these walkers begin far from any minima, they will initially walk around in areas of low probability until they find a probability island and begin to "map" this area. This process is called burn-in; iterations in the burn-in phase are not saved.

To propose an updated model m', we use a symmetric proposer (a Metropolis update), which means that the probability of jumping from m to m' is the same as the probability of the reverse. This means that we do not have to include the ratio of the probability of m' given m, $q(m'|m)$, divided by $q(m|m')$ in the acceptance criteria. Since we are working with parameters in logarithmic space, these updates correspond to multiplicative factors and the proposer is therefore only symmetric in the log space and not in the linear space. We will exclusively work in the log model space. Our proposer is based on the posterior covariance matrix obtained from a gradient-based inversion. This has the advantage that it probes the model space faster when it is close to the linear estimate than, for example, a uniform proposer. However, if this is not the case, such a proposer will be inefficient. To ensure that the model parameters stay within reasonable bounds where the forward routine is known to have good accuracy, we also use hard constraints on the parameters. This means that all proposals with a parameter value outside the hard limits are discarded. The limits are as follows: 1 m < thickness < 200 m, 0.01 < WC < 1 and 5 ms < T_2^* <1.5 seconds, where WC is the layer's water content.

Once a model has been proposed, we determine whether to keep it on the basis of the objective function J, which is given by

$$J = \sqrt{J_{data}(m) + J_{prior}(m)}\,, \tag{2}$$

where the first part is the data misfit

$$J_{data}(m) = \sum_i \left(\frac{FWR_i(m) - d_i}{\sigma_i}\right)^2 . \tag{3}$$

We loop over data points i and sum the squared normalised differences. $FWR_i(m)$ is the forward response for model m, d_i is the data value (FID), and σ_i is the data standard deviation. The second term in the objective function is the prior model constraints

$$J_{prior}(m) = \sum_j \left(\frac{\log(m_j) - \log(m_{j,o})}{\log(\sigma_j)}\right)^2 , \tag{4}$$

where we loop over model parameters j and sum the normalised squared model differences in logarithmic space. $m_{j,0}$ is the prior model, and σ_j is the model parameter standard deviation. We generally use a loose prior constraint of σ_j =10. The combined probability distribution for the data and the prior is given by

$$P(m) = \exp(-0.5J(m)^2) , \tag{5}$$

and we accept a new model m' if $J(m') < J(m)$ or if a random number α between 0 and 1 obeys

$$\alpha < \frac{\exp(-0.5J(m)^2)}{\exp(-0.5J(m_0)^2)} . \tag{6}$$

Figure 1 illustrates a schematic MCMC run with four walkers with different starting models. The model space in this simple illustration contains only two parameters and can be visualised as a plane containing all possible models. Each model is described by two coordinates in this plane *(x,y)*. The example probability distributions, which are sampled by the walkers, consist of two normal distributions centred at (0,2) and (6,4). The walkers begin by walking towards a region with high probability, i.e., not in a straight line but with a clear direction. This traverse towards the high probability region is the burn-in phase. Once the walkers are in a region

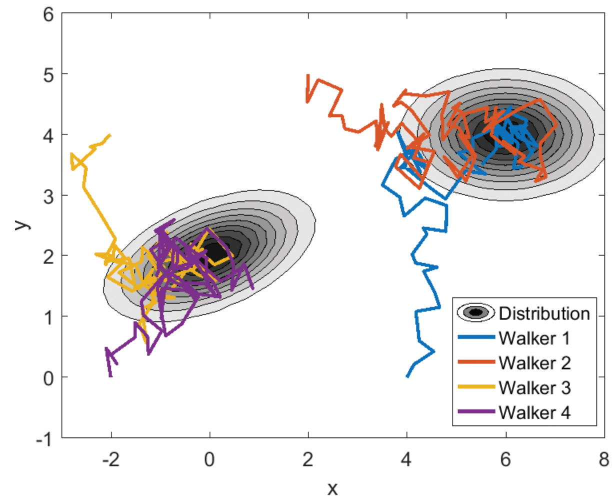

Figure 1 Four independent MCMC random walkers starting at different positions that take uniformly distributed steps. The target distribution is shown in grey scale and consists of two peaks. By using the 4 walkers we end up having two walkers finding each distribution.

of high probability, they begin "mapping" the probability distribution. For a clear representation of the MCMC principle, we only show a few iterations. In an actual MCMC run with millions of iterations, a walker that is initially trapped in one of the probability islands would eventually jump to the other island. The inversion will be stopped when good convergence is reached, and the full parameter space has been mapped.

Figure 2	**Layer**	**Resistivity (Ωm)**		**Thickness (m)**
TEM equivalence	1	100		20
	2	10		10
	3	1000		-
Figure 3	**Layer**	**Water content (%)**	**T_2* (ms)**	**Thickness (m)**
Constant	1	3	50	C-1/ 2/WC_2
Water content –	2	WC_2	200	1/ WC_2
Thickness product	3	3	50	-
Figure 4	**Layer**	**Water content (%)**	**T_2* (ms)**	**Thickness (m)**
Constant	1	30	150	C-1/ 2/WC_2
Water content –	2	WC_2	50	1/ WC_2
Thickness product	3	30	150	-
Figure 5	**Layer**	**Water content (%)**	**T_2* (ms)**	**Thickness (m)**
T_2* dependence on	1	30	100	10
resolution of thick-ness	2	20	20, 50, 80, 150	3, 4.5, 9, 30
	3	30	100	-
Figure 6	**Layer**	**Water content (%)**	**T_2* (ms)**	**Thickness (m)**
T_2* dependence on	1	30	100	10
resolution of water	2	3, 10 ,20, 30	20, 50, 80, 150	9
content	3	30	100	-

Table 1 True model parameters for all figures. For Figures 3-6 the resistivity is 100 Ωm for all layers. C and WC_2 for Figures 3 and 4 correspond to the depth to the center of layer 2 (eg. 21 m or 51 m) and the water content in laywer 2.

RESULTS AND DISCUSSION

To demonstrate the utility of an MCMC framework for surface NMR, three questions are investigated. (1) Does a water content–layer thickness equivalency exist in surface NMR (like the resistivity–thickness equivalency present in EM)? (2) Does T_2* play a role in controlling the resolution of the estimated profiles? (3) Are uncertainties predicted by the posterior covariance matrix consistent with those predicted using an MCMC approach? For all synthetic surface NMR surveys, we employ a 100-m coincident square transmit/receive loop and a 30-ms on-resonance pulse, and we measure for a 400-ms duration with a 40-ms dead time and use 16 pulse moments that range from 0.74 Ams to 8.53 Ams in log steps. Subsurface conductivity is equal to 100 Ωm at all depths. For all of the presented surface NMR MCMC inversions, the inverted model is a three-layer system; inversion parameters include the water content, T_2*, and thickness of layers 1 and 2, as well as the water content and T_2* of layer 3. The thickness of the bottom third layer is effectively treated as infinite (it is not an inversion parameter). We therefore have eight model parameters, and the MCMC walker traverses an eighth-dimensional model space. To illustrate potential parameter correlations, the marginal distributions for the different parameters are plotted to form 2D probability distributions. The MCMC inversion runs until convergence has been reached and smooth distributions found. We generally sample at least 100,000 iterations.

Investigation of water content–layer thickness correlation

Direct sensitivity to water content is one of the most attractive features of the surface NMR measurement, i.e., the ability to quantify total water content and to map its spatial distribution. One factor that may influence the interpretation of surface NMR water content profiles is if a water content–thickness correlation exists in surface NMR similar to the resistivity–thickness correlation in TEM. To illustrate the equivalence problem in TEM, consider a simple three-layer model where the first and third layers are resistive and the second layer is conductive. The true subsurface parameters are given in Table 1. The setup is a ground-based 50 m × 50 m square transmitter with a central receiver that records the impulse response in 22 gates from 5 μs to 1.5 ms. The uncertainty is 5% for all gates. We perform an MCMC inversion and plot the conductivity of the second layers versus the thickness of the second layer in Figure 2. The resulting distribution is very elongated and has a negative slope of -1 in log space. Notice that the product of the two is almost constant and equal to 1; thus, this is a very pronounced equivalence.

The existence of a similar water content–thickness correlation has been previously investigated in several studies. Legchenko and Shushakov (1998) demonstrate that two datasets, i.e., one dataset produced by a 1-m-thick water layer held at a fixed depth and a second dataset produced by a 1-m-thick water layer whose depth is varied, exhibit strong correlations when the water layer in each model is placed at similar depths. The correlation is shown to extend over larger depth ranges when both layers are present at great depths. Legchenko and Shushakov (1998) also demonstrate that a 1-m-thick water layer exhibits correlations with water layers of greater thickness placed at the same depth, where the correlation extends to larger thicknesses if the water layers occur at depth. Together, these results highlight the challenges uniquely resolving the depth and thickness of thin layers at depth. Furthermore, Legchenko *et al.* (2004) performed a study where the depth to a 10-m-thick 20% water content layer was varied from 0 to ~100 m, where it was shown that the ability to resolve the layer's water content and thickness decreased with depth. However, the water content–thickness product of the layer was better resolved at depth than either the water content or thickness independently. An improved ability to resolve the water content–thickness product suggests the presence of an equivalence problem similar to Figure 2. To expand upon these studies, we present a suite of synthetic surveys where the inversion results are presented for scenarios where the water content–thickness product is preserved but the water content and layer thicknesses are varied. This study differs from that performed by Legchenko and Shushakov (1998) and Legchenko *et al.* (2004) in that we consider a scenario where we fix both the water content–thickness product and the depth to the centre of this layer. This scenario is well suited to an investigation of potential water content–thickness equivalences given that it holds the total water content and depth of the layer constant. Our focus is to investigate the ability of the inversion to independently resolve the water content and layer thicknesses for these potentially equivalent models.

Consider first a set of models described by three-layer systems, where a single-water-rich layer is present in a low-water-content background. The true subsurface parameters are given in Table 1; layers 1 and 3 represent a low-water-content (3%) fast

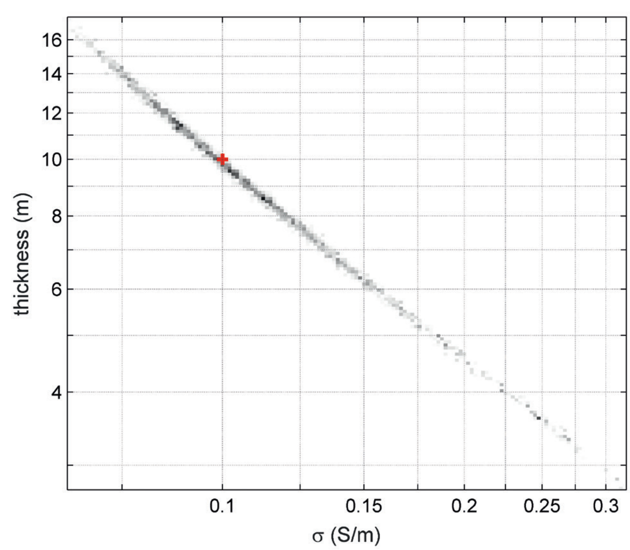

Figure 2 TEM conductivity-thickness equivalence for a three layer system. σ is the conductivity of the second layer and z is the thickness. The other model parameters are in Table 1.

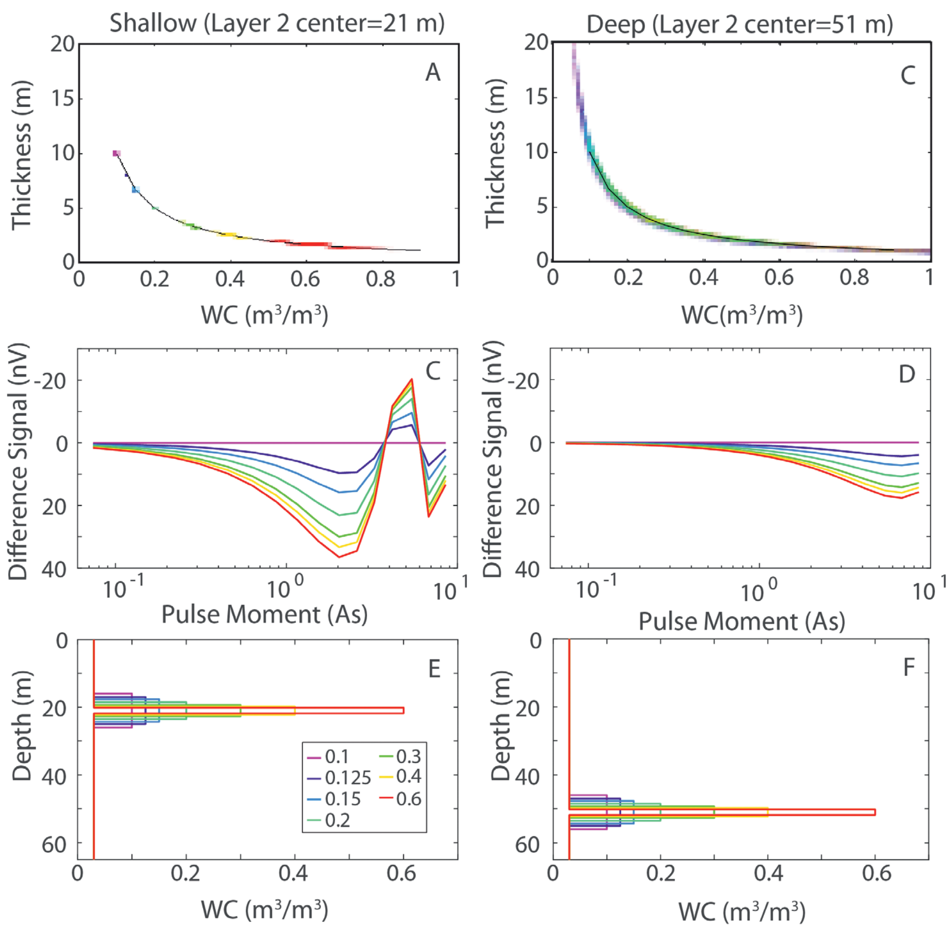

Figure 3 Comparison of the ability to independently resolve the water content and thickness of a water rich layer present in a low water content background. In each case, the water content – thickness product is equal to 1. Profile colors correspond to a particular water content – thickness pair. The top row illustrates the joint probability distribution for the estimated water content and layer thickness of layer 2. The middle row illustrates the difference between the sounding curves produced by each model compared to the sounding curve produced by the 10% water content case. Note that a 40 nV signal difference corresponds to ~20% variation in the signal amplitude. The bottom row illustrates the true water content profiles in each case. The left and right columns correspond to "shallow" and "deep" scenarios, where layer 2 is centered at 21 m and 51 m, respectively.

T_2* (50 ms) background, whereas layer 2 is given a fixed water content–thickness product equal to 1. Colours in Figure 3 correspond to a particular water content–thickness pair; the investigated layer 2 water contents (WC_2) are 0.1 (magenta), 0.125, 0.15, 0.2, 0.3, 0.4, and 0.6 (red), whereas the layer 2 thickness in each case is $1/WC_2$. Note that the thickness of layer 1 is also dependent on WC_2 and is adjusted to ensure that layer 2 is centred at the same depth in each case. If a strong water content–layer thickness correlation exists, it should be difficult to differentiate between these models. The left and right columns illustrate scenarios where layer 2 is centred at a depth of 21 and 51 m, respectively, each intended to represent a "shallow" and "deep" case. Parameter C in Table 1 corresponds to the depth to the centre of layer 2 (e.g., 21 or 51 m). The top row illustrates the joint water content–thickness probability distribution. The middle row illustrates the variation of the resulting sounding curve (with respect to the $WC_2 = 0.1$ case (magenta)). The bottom row illustrates the true water content profiles and corresponding colours. Note that the top and bottom rows correspond to inverted water contents and the true targeted water contents, respectively. Inverted data in this case are noise-free, but are treated with a 5% uncertainty during the inversion. Noise-free data are selected to provide insight into potential correlations under ideal scenarios.

Comparison of the signal differences in the noise-free limit provides insight into the noise level at which the data from different models cannot be distinguished. Consider first the "shallow" scenario (left column). For the low-water-content cases (magenta to green), the water content and layer thickness are each well resolved and do not exhibit strong correlations (observed by noting that the probability distributions are not elongated along the black line (which corresponds to a constant water content–thickness product of 1). For the thinner layers (yellow to red), the probability distributions are now elongated along the black line indicating that the water content–layer thickness correlation is stronger. Figure 3C illustrates the difference in the sounding curves produced in each case. As the layer thickness decreases (yellow to red), the data differences become very small. Therefore, the inversion (which treats all data with a 5% uncertainty) struggles to independently resolve the water content and layer thickness resulting in a probability distribution that smears along the black line. To quantify noise levels where the ability to differentiate between these models will be lost, the magnitude of the signal differences can be examined. For example, if the noise level exceeds 20 nV, we will lose the ability to distinguish between the green and red models. Consider next the "deep" scenario (right column). In this case, the joint probability distributions exhibit much stronger elongation, where in all cases the distributions span large ranges along the black line. The water content and layer thickness now exhibit a strong correlation. The source of the increased correlation is observed in Figure 3D, where the signal differences produced by each model are now much smaller; the data differences are less than the 5% data uncertainty assumed by the inversion. Furthermore, the signal differences are only significant for the strongest pulse moments.

The right column of Figure 3 exhibits behaviour consistent with the observations of Legchenko *et al.* (2004), where the water content–thickness ratio was well resolved at depth despite an inability to independently resolve the water content and thickness of the water-bearing layer. It also highlights that, for scenarios where a high-water-content layer is present in a low-water-content background, surface NMR estimates of the total water (i.e., water content multiplied by layer thickness) remain accurate even at depth. The shallow results indicate that the water content–thickness correlation is not always present, particularly when the water-bearing layer is present at shallower depths with thicknesses greater than a few metres. This has important implications for interpretations of water content profiles as it suggests that, in low-noise conditions (low enough that the signal differences are greater than the noise level), surface NMR is capable of accurately resolving the water content and layer thicknesses independently, which is necessary to produce reliable estimates of layer boundaries. A workflow similar to that in Figure 3 is a useful tool to identify if the signal differences for two models with the same water content–layer thickness product but different water contents are greater than observed noise levels. If the signal differences are below the noise level, it is unlikely that the inversion can differentiate between the two cases.

Consider next an example of a three-layer system where a variable water content layer (layer 2) is present within a high-water-content background. The true subsurface parameters are listed in Table 1; layers 1 and 3 represent a high-water-content (30%) slow T_2* (150 ms) background, whereas layer 2 is given a fixed water content–thickness product equal to 1 and T_2*=50 ms. Colours in Figure 4 correspond to a particular water content–thickness pair; the investigated layer 2 water contents and thicknesses are the same as in Figure 3. Note that the thickness of layer 1 is also adjusted based on the layer 2 water content to ensure layer 2 remains centred at the same depth. The left and right columns correspond to scenarios where layer 2 is centred at a depth of 21 and

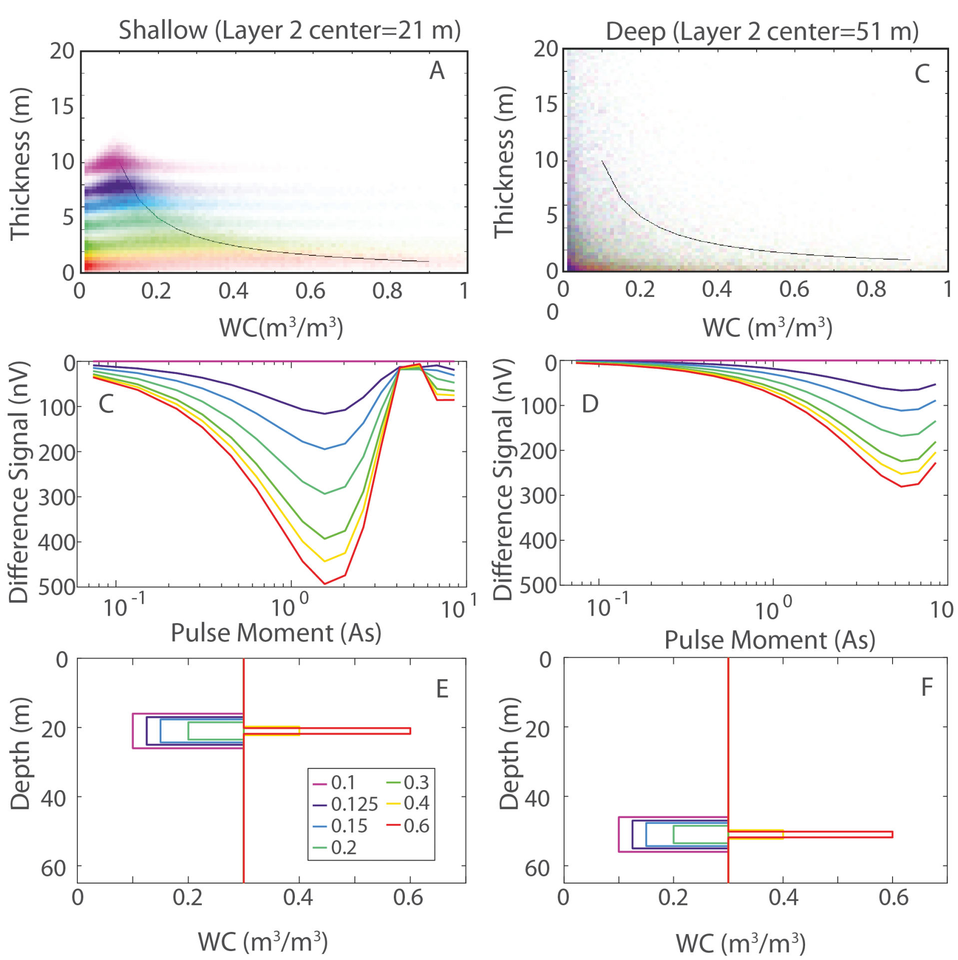

Figure 4 Comparison of the ability to independently resolve the water content and thickness of a variable water layer present in a high water content background. In each case, the water content – thickness product is equal to 1. Profile colors correspond to a particular water content – thickness pair. The top row illustrates the joint probability distribution for the estimated water content and layer thickness of layer 2. The middle row illustrates the difference between the sounding curves produced by each model compared to the sounding curve produced by the 10% water content case. The bottom row illustrates the true water content profiles in each case. The left and right columns correspond to "shallow" and "deep" scenarios, where layer 2 is centered at 21 m and 51 m, respectively.

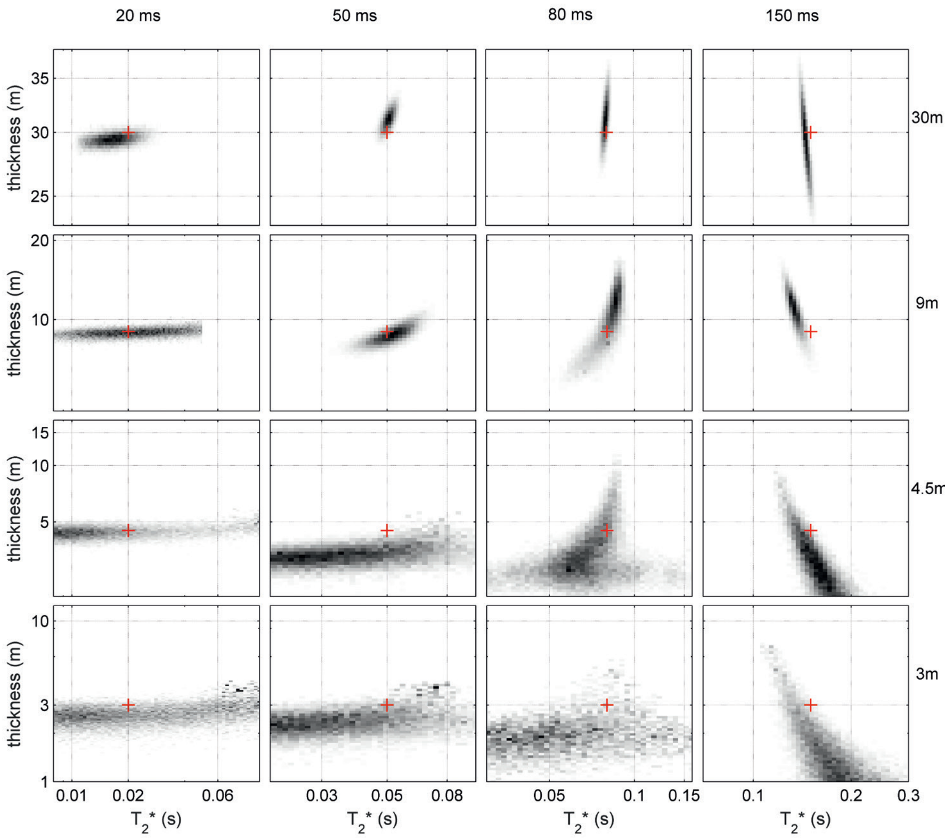

Figure 5 Thickness and T_2*cross-plots along with the true values (red markers) for the second layer probing the T_2*influence on the thickness resolution.

51 m, respectively, each intended to represent a "shallow" and "deep" case. This scenario differs from Figure 3, in which we now focus on an example where the layer of interest (layer 2) does not dominate the behaviour of the signal, i.e., layers 1 and 3 contribute significantly to the signal because of their high water contents. This scenario can be considered to be representative of a fast relaxation aquitard that separates two high-water-content layers. Inverted data are noise-free but are treated with a 5% uncertainty during the inversion. The true water content profiles in each case are shown in Figures 4E and 4F. Consider first the "shallow" scenario (left column of Figure 4), where the probability distributions show a much different behaviour than in Figure 3. The thickness of layer 2 is well resolved in each case, but the layer 2 water content is poorly resolved. For thinner layers, the water content becomes even more poorly resolved (e.g., the red distribution smears from low water contents up to 100%). The sounding curves in Figure 4C show larger variations than in Figure 3C (the ~400- to 500-nV variation corresponds to ~20%–25% variation in the total signal amplitude). The reason for the stronger variation is that the high-water-content background leads to a much larger signal amplitude. In this case, treating all data with 5% uncertainty corresponds to a higher noise level than that considered in Figure 3. Figure 4C suggests that noise levels of ~100 nV will make it difficult to differentiate between certain cases (e.g., between the red and green cases). For the deep case (right column), the water content and layer thickness of layer 2 are poorly resolved in all cases. The probability distributions are smeared over wide water content and thickness ranges in each case, clustering in the bottom left corner with the black line loosely describing the shape of distributions at higher water contents/thicknesses. Figure 4D indicates that the data differences are reduced in this case, where the differences at large currents correspond to ~10%–15% of the peak signal amplitude. At lower currents, the signal differences are much smaller.

Overall, Figures 3 and 4 suggest that, while water content–thickness equivalence is observed in certain conditions, it is not present in all cases. For example, Figure 3A illustrates a case where water content and thickness can be independently resolved, Figure 4A illustrates a case where only thickness can be resolved, and Figure 4B shows a case where neither parameter is well resolved. The presence of the equivalence will depend on site-specific parameters such as the local water content profile and noise conditions.

Control of T_2* on the ability to resolve water contents and layer thicknesses

In surface NMR, a kernel matrix is used to invert the observed data to produce estimates of the subsurface properties. This kernel generally predicts the initial amplitude of the signal and is often used to estimate the spatial resolution of the resulting profiles using the singular value decomposition approach of Müller-Petke and Yaramanci (2008). However, an approach based exclusively on the kernel matrix (which describes the expected signal amplitudes at time 0) inherently neglects the influence of T_2* on model parameter resolution. To investigate whether T_2* plays a role in controlling the ability to resolve the subsurface parameters, an MCMC approach is used to investigate whether T_2* displays any correlation with the other estimated parameters. If T_2* plays no role in the expected resolution, we should not see any correlation between T_2* and the water contents and thicknesses. It is common during interpretation of results to assume higher uncertainty in the water content for layers containing very fast T_2* and to assume a high uncertainty in the T_2* when a very low water content is present. This intuition suggests that T_2* should play a role in our ability to resolve the layer parameters. These extreme cases are intuitive but the influence of intermediate T_2* remains unclear. We now explore the correlation for intermediate T_2* values.

Girard, Legchenko and Boucher (2005) demonstrate that a correlation exists between the estimated T_2* and resulting initial amplitude estimate. This correlation arises from the extrapolation process, where the estimated T_2* is used to produce the initial amplitude estimate. Therefore, if T_2* is underestimated, the extrapolation procedure will underestimate the initial amplitude. This suggests that T_2* is likely to exhibit a correlation with the estimated water content. However, how this suspected correlation propagates through the inversion process and ultimately impacts our ability to resolve layer properties is unclear. We hypothesise that T_2* controls the ability to resolve layer properties given that it determines the persistence of the signal from a given depth and thereby controls the local signal-to-noise ratio. Using a simple three-layer model similar to the previous cases, we investigate the effect of T_2* on the ability to resolve the other layer properties in the same layer, i.e., water content and thickness. This is done by investigating whether the T_2* estimate in the second layer exhibits a correlation with the water content and thickness in the same layer.

In the first example, we fix the water content of the second layer and vary T_2* and thickness of the second layer to examine if a correlation exists between these parameters. All remaining model parameters are listed in Table 1. Figure 5 illustrates the cross plots of T_2* and layer thickness estimates for layer two in all 16 combinations. For low T_2* (left column), the layer thickness is well determined due to the large contrast with neighbouring layers, which helps constrain the layer boundaries despite that the T_2* value is poorly determined in this case. As the T_2* value approaches the background T_2* value (e.g., columns 3 and 4), the layer thickness uncertainty increases, whereas the T_2* uncertainty depends strongly on the thickness of the layer (i.e., the widths of the T_2* distributions are much narrower in the top right versus bottom right panels). For example, for a 9-m-thick layer (second row), the uncertainty seems to transfer from T_2* to thickness when increasing T_2* (moving from left to right). A correlation between T_2* and layer thickness is also observed for intermediate T_2* values. This suggests that T_2* influences the ability to resolve the layer thickness.

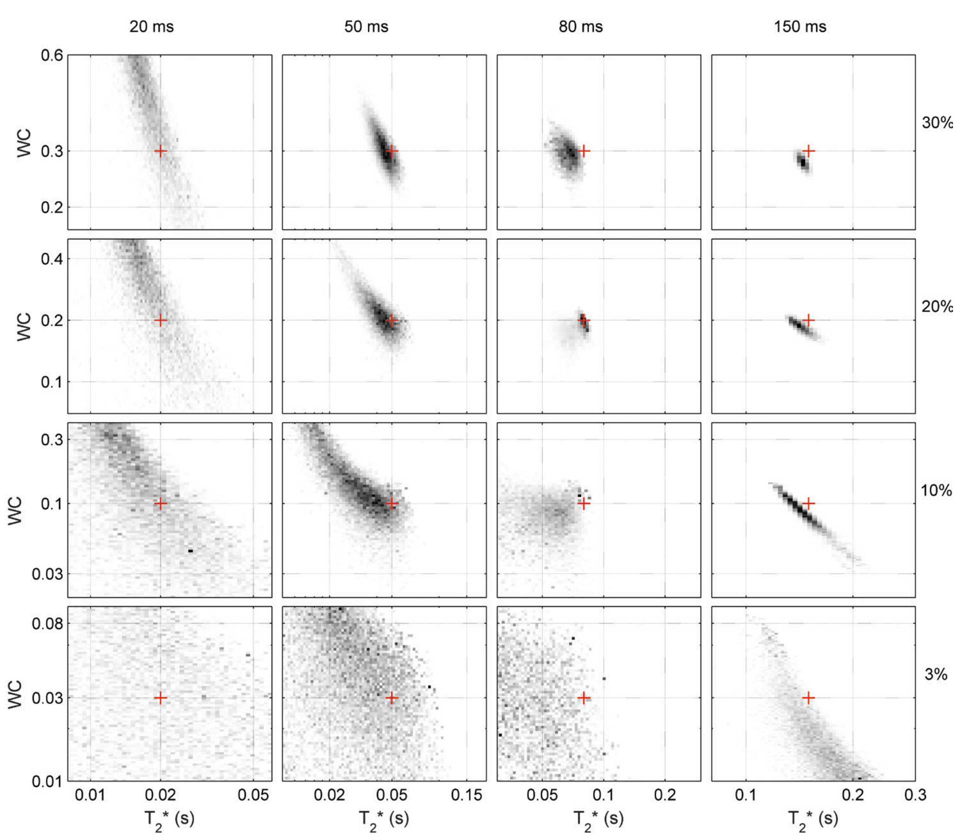

Figure 6 Water content and T_2* cross-plots along with the true values (red markers) for the second layer probing the T_2* influence on the water content resolution.

In the second example, we fix the thickness of the second layer and vary the water content and T_2* values in this layer to examine if a correlation exists between these parameters. All remaining model parameters are listed in Table 1. Figure 6 illustrates the cross plots of T_2* and water content estimates for layer 2 in all 16 combinations. For the low-water-content example (bottom row), both parameters are unresolved. For low T_2* (left column), the water contents are unresolved, with T_2* only being resolved for the highest water content case. This is consistent with previous intuition that suggests water content is highly uncertain for fast decays. In several of the remaining cases, T_2* and water content exhibit a negative correlation, suggesting that it is difficult to differentiate between a small amount of slow decaying water versus a larger amount of fast decaying water. This is consistent with the idea that T_2* is also controlling the local signal-to-noise ratio and that, even for intermediate T_2* values, a correlation with water content persists.

In summary, T_2* influences the ability to resolve the water content and layer thicknesses of the produced depth profiles and should be considered when determining model parameter resolution. Previous intuition suggesting that in the fast T_2* limit water contents are highly uncertain and that in the low-water-content limit the T_2* is highly uncertain is confirmed by the MCMC results. However, the MCMC results indicate that T_2* exhibits a correlation with water content and layer thicknesses even for intermediate value of T_2*.

Comparison of Markov chain Monte Carlo determined and posterior covariance-matrix-based uncertainty estimates

We now compare the parameter standard deviations obtained from the posterior covariance matrix after gradient-based inversions to those estimated by the MCMC approach. If these generally do not agree, one has to treat uncertainties obtained using the gradient-based inversions with care. Gradient-based inversions do not probe the full parameter distributions and determine the uncertainties using derivatives calculated at the inversion minimum. One should be careful using the gradient-based uncertainties if they deviate too much from the MCMC results, which, on the other hand, probes the full distribution and can more readily handle distributions that are not well described by normal distributions. The posterior covariance matrix is often much faster to compute than an MCMC result; therefore, if uncertainty estimates from the posterior covariance matrix are reliable, it represents a convenient approach to estimate uncertainty. However, the covariance matrix approach may struggle to accurately describe parameter correlations and skewness of the probability distributions that can be more readily quantified by the MCMC approach.

For the gradient-based inversions, the parameter uncertainties are based on the posterior covariance matrix (Tarantola and Valette 1982a). This is given by

$$C_{est} = (G^T C_{obs} G)^{-1}, \tag{7}$$

where G is the Jacobi matrix and C_{obs} is the data covariance matrix that consists of the data uncertainties. This expression can be generalised to include model regularisation and prior constraints on the model parameters. We are working with model parameters in log space and under the assumption that they are normally distributed. One can obtain the standard deviations STD from the diagonal elements of C_{est}. In this case

$$STD = \sqrt{C_{est}}\ . \tag{8}$$

For parameters in linear space, the standard deviation in log space converts to a standard deviation factor given by

$$STDF = \exp(\sqrt{C_{est}})\ , \tag{9}$$

where the ±STD limits are given by

$$\frac{m}{STDF} < m < m \cdot STDF\ . \tag{10}$$

To compare the covariance matrix predicted uncertainties with those predicted by an MCMC approach, we look at a single three-layer model and compare the parameter uncertainties of the second layer. The resistivity is fixed to 100 Ωm for all layers. Figure 7 shows the marginal distributions obtained from an MCMC inversion. The real model has a 9-m-thick second layer

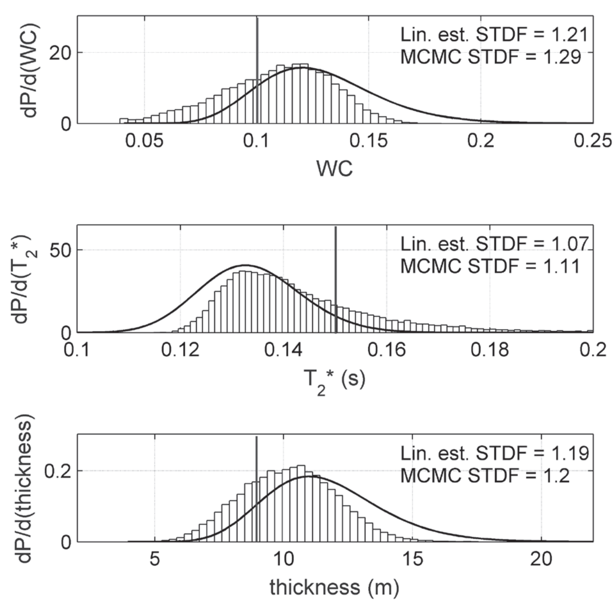

Figure 7 Marginal distributions for water content, T_2* and the thickness of the second layer for the MCMC results shown in Figure 5 where WC = 10 % ,T_2* = 150ms and thickness = 9m, marked by the grey lines. The histograms are the MCMC results and the curved lines are the distributions determined from the posterior covariance matrix. The standard deviation factors (STDF) are shown for the linearized estimate and the MCMC distributions. The vertical black lines that extend to the top of each panel correspond to the true layer two values.

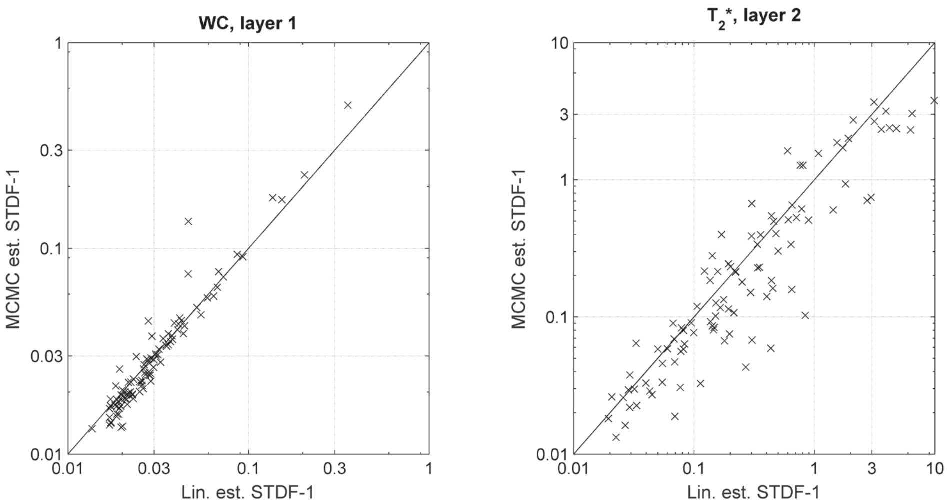

Figure 8 Comparison of uncertainty estimates based on posterior covariance matrix and MCMC inversion for the WC of the first layer and T_2* of the second layer.

with 10% water content and 150-ms decay (true values indicated by the vertical black line). Layers 1 and 3 have 30% water content and 100-ms decay time; Figure 6 (third row, fourth column) shows the joint probability distribution between T_2* and water content for the same subsurface model. Although the distributions obtained from the MCMC are similar to the gradient-based results, there are several differences. The MCMC water content and thickness distributions exhibit longer tails (particularly in the water content and T_2* distributions) compared with the gradient-based results. The standard deviations obtained from the gradient-based inversions are also slightly smaller than the MCMC standard deviations in each case. This is because the gradient-based distributions do not include the skewness observed in the MCMC results. However, just looking at the STD, the values are quite similar likely because the linear and the MCMC STDs are both small, which is the range where the linear calculation is known to work well. Note that both the MCMC and gradient-based distributions are not centred exactly at the true values (vertical black lines); however, the true values do fall within the probability distributions in each case.

To generalise this investigation, we have created 100 random three-layer models where the resistivity is set with a value between 10 and 1000 Ωm with a log uniform probability distribution. The WC of the three layers have random values between 5% and 50%, the T_2* have values between 30 and 400 ms, and the thickness of the layers is between 1 and 70 m. For every model, we have the result based on the posterior covariance matrix and the standard deviation of the MCMC distribution run with 100,000 iterations. We use a fixed data uncertainty of 10% for all data points. We have initially generated forward responses and added 10% relative noise. Afterwards, we start a gradient-based inversion of these data and save the result. The mean and standard deviation of the logarithm of the model parameters are stored and compared with the gradient-based results. The results are shown in Figure 8 for the water content of the first layer and T_2* for the second layer.

Using this random model MCMC approach, we have tested 100 unique models spanning a reasonable parameter space and thereby better test the hypothesis that the linear estimates agree with the MCMC. Our results show good agreement between the two methods, which all scatter around the "1-to-1" line. This means that we can safely use the estimates based on the posterior covariance matrix for the parameter uncertainties. Note that the linear STDFs in these cases can be computed ~1000 faster than the MCMC STDF.

CONCLUSIONS

In order to ensure reliable interpretation of surface NMR estimated water content and T_2* profiles, the uncertainties and correlation of the estimated parameters must be considered. An MCMC framework is demonstrated to provide the ability to characterise both the uncertainty of each parameter and whether parameters demonstrate correlation with one another. The utility of such a framework is demonstrated by investigating the following three questions: (1) whether the water content and layer thickness of a particular layer are correlated, as might be expected based on analogies with EM; (2) whether T_2* exhibits any influence on the resolution of the produced profiles; and (3) whether uncertainty estimates based on the posterior covariance matrix from gradient-based inversions are consistent with MCMC estimated uncertainties. Regarding the first question, the existence of a strong water content–layer thickness correlation is not observed in all scenarios. For deep-water-bearing layers, water content and layer thickness are observed to be strongly correlated, allowing only their product to be accurately resolved. In contrast, for shallower-water-bearing layers, water content and layer thickness can be independently resolved provided that the noise level is less than the differences the data produced for "equivalent" models (i.e., models with the same water content–layer thickness product but different water contents). Alternatively, in some cases, only the layer thickness is well resolved whereas the water content is poorly resolved. For the second question, the MCMC results are consistent with common intuition that the water content estimate is highly uncertain in the small T_2* limit, whereas T_2* is very uncertain in the low-water-content limit. However, T_2* is also observed to exhibit correlations with layer thickness and water content for intermediate T_2* values, suggesting that T_2* must be considered when determining the resolution of estimated water content profiles. For the final question, uncertainty estimates based upon the posterior covariance matrix for the linearised inversion fit well to the MCMC results. However, the uncertainty estimates based on the posterior covariance matrix are not able show the advanced features of some of the correlations observed in this paper and therefore cannot give a complete picture of the parameter uncertainties and correlations.

The proposed framework presents a means to more thoroughly explore the model space and produce a suite of models that fit the data equally well. A shortcoming of the MCMC approach is that it is limited to working with relatively simple models due to practical limitations related to computation times. Improvements to the speed of the surface NMR forward calculation, specifically the calculation of the kernel matrix, could extend the utility of the MCMC approach to more complex models. Overall, the MCMC approach to surface NMR inversion improves the reliability of the surface NMR results by allowing a more straightforward consideration of parameter uncertainty and correlation.

ACKNOWLEDGEMENTS

This work was supported in part by the National Natural Science Foundation of China under Grants 41604095 and 41374075. Funding was also provided in part by Innovation Fund Denmark under Grant 082-2013-3 and by China Postdoctoral Science Foundation under Grant 2014M561296.

REFERENCES

Behroozmand A.A., Auken E., Fiandaca G. and Christiansen A.V. 2012. Improvement in MRS parameter estimation by joint and laterally constrained inversion of MRS and TEM data. *Geophysics* **74**, WB191–WB200.

Behroozmand A.A., Dalgaard E., Christiansen A.V. and Auken E. 2013. A comprehensive study of the parameter determination in a joint MRS and TEM data analysis scheme. *Near Surface Geophysics* **11**, 557–567.

Bloch F. 1946. Nuclear induction. *Physical Review* **70**, 460–474.

Braun M. and Yaramanci U. 2008. Inversion of resistivity in magnetic resonance sounding. *Journal of Applied Geophysics* **66**, 151–164.

Brooks S., Gelman A., Jones G.L. and Meng X.-L. 2011. *Handbook of Markov Chain Monte Carlo*. CRC Press.

Chevalier A., Legchenko A., Girard J.F. and Descloitres M. 2014. Monte Carlo inversion of 3-D magnetic resonance measurements. *Geophysical Journal International* **198**, 216–228.

Costabel S. and Yaramanci U. 2013. Estimation of water retention parameters from nuclear magnetic resonance relaxation time distributions. *Water Resources Research* **49**, 2068–2079.

Dlugosch R., Günther T., Müller-Petke M. and Yaramanci U. 2013. 2D qt-inversion to investigate spatial variations of hydraulic conductivity using SNMR. *Symposium on the Application of Geophysics to Engineering and Environmental Problems*.

Dlugosch R., Günther T., Müller-Petke M. and Yaramanci U. 2014. Two-dimensional distribution of relaxation time and water content from surface nuclear magnetic resonance. *Near Surface Geophysics* **12**, 231–241.

Girard J.F., Legchenko A. and Boucher M. 2005. Stability of MRS signal and estimation of data quality. *Near Surface Geophysics*, 187–194.

Girard J.F., Boucher M., Legchenko A. and Baltassat J.-M. 2007. 2D magnetic resonance tomography applied to karstic conduit imaging. *Journal of Applied Geophysics* **63**, 103–116.

Guillen A. and Legchenko A. 2002. Inversion of surface nuclear magnetic resonance data by an adapted Monte Carlo method applied to water resource characterization. *Journal of Applied Geophysics* **50**, 193–205.

Hertrich M. 2008. Imaging of groundwater with nuclear magnetic resonance. *Progress in Nuclear Magnetic Resonance Spectroscopy* **53**, 227–248.

Hertrich M. and Yaramanci U. 2002. Joint inversion of surface nuclear magnetic resonance and vertical electrical sounding. *Journal of Applied Geophysics* **50**, 179–191.

Legchenko A., Baltassat J.-M., Beauce A. and Bernard J. 2002. Nuclear magnetic resonance as a geophysical tool for hydrogeologists. *Journal of Applied Geophysics* **50**, 21–46.

Legchenko A., Baltassat J.-M., Bobachev A., Martin C., Robain H. and Vouillamoz J.-M. 2004. Magnetic resonance sounding applied to aquifer characterization. *Groundwater* **42**, 363–373.

Legchenko A., Ezersky M., Camerlynck C., Al-Zoubi A. and Chalikakis K. 2009. Joint use of TEM and MRS methods in a complex geological setting. *Comptes Rendus Geoscience* **341**, 908–917.

Legchenko A.V. and Shushakov O.A. 1998. Inversion of surface NMR data. *Geophysics* **63**, 75–84.

Malinverno A. 2002. Parsimonious Bayesian Markov chain Monte Carlo inversion in a nonlinear geophysical problem. *Geophysical Journal International* **151**, 675–688.

Minsley B.J. 2011. A trans-dimensional Bayesian Markov chain Monte Carlo algorithm for model assessment using frequency-domain electromagnetic data. *Geophysical Journal International* **187**, 252–272.

Mohnke O. and Yaramanci U. 2002. Smooth and block inversion of surface NMR amplitudes and decay times using simulated annealing. *Journal of Applied Geophysics* **50**, 163–177.

Müller-Petke M. and Yaramanci U. 2008. Resolution studies for magnetic resonance sounding (MRS) using the singular value decomposition. *Journal of Applied Geophysics* **66**, 165–175.

Müller-Petke M. and Yaramanci U. 2010. QT inversion—Comprehensive use of the complete surface NMR data set. *Geophysics* **75**, WA199–WA209.

Parsekian A.D. and Grombacher D. 2015. Uncertainty estimates for surface nuclear magnetic resonance water content and relaxation time profiles from bootstrap statistics. *Journal of Applied Geophysics* **119**, 61–70.

Prange M. and Song Y.Q. 2010. Understanding NMR T-2 spectral uncertainty. *Journal of Magnetic Resonance* **204**, 118–123.

Sambridge M. and Mosegaard K. 2002. Monte Carlo methods in geophysical inverse problems. *Reviews of Geophysics* **40**.

Tarantola A. and Valette B. 1982a. Generalized nonlinear inverse problems solved using a least squares criterion. *Reviews of Geophysics and Space Physics* **20**, 219–232.

Tarantola A. and Valette B. 1982b. Inverse problems = quest for information. *Journal of Geophysics* **50**, 159–170.

Trushkin D.V., Shushakov O.A. and Legchenko A.V. 1995. Surface NMR applied to an electroconductive medium. *Geophysical Prospecting* **43**, 623–633.

Weichman P.B., Lavely E.M. and Ritzwoller M.H. 2000. Theory of surface nuclear magnetic resonance with applications to geophysical imaging problems. *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdiscip*linary Topics 62, 1290–1312.