Near Surface Geophysics, 2017, **15**, 533-544 doi:10.3997/1873-0604.2017027

Comparison of stabiliser functions for surface NMR inversions

Denys Grombacher[1*], Gianluca Fiandaca[1], Ahmad A. Behroozmand[2] and Esben Auken[1]

[1] *Department of Geoscience, Aarhus University, 8000 Aarhus C, Denmark*
[2] *Department of Geophysics, Stanford University, Stanford, CA 94305-2210, USA*

Received October 2016, revision accepted May 2017

ABSTRACT

Surface nuclear magnetic resonance is a geophysical technique providing non-invasive aquifer characterization. Two approaches are commonly used to invert surface nuclear magnetic resonance data: (1) inversions involving many depth layers of fixed thickness and (2) few-layer inversions without predetermined layer thicknesses. The advantage of the many-layer approach is that it requires little *a priori* knowledge. However, the many-layer inversion is extremely ill-posed and regularisation must be used to produce a reliable result. For optimal performance, the selected regularisation scheme must reflect all available *a priori* information. The standard regularisation scheme for many-layer surface nuclear magnetic resonance inversions employs an L_2 smoothness stabiliser, which results in subsurface models with smoothly varying parameters. Such a stabiliser struggles to reproduce sharp contrasts in subsurface properties, like those present in a layered subsurface (a common near-surface hydrogeological environment). To investigate if alternative stabilisers can be used to improve the performance of the many-layer inversion in layered environments, the performance of the standard smoothness stabiliser is compared against two alternative stabilisers: (1) a stabiliser employing the L_1-norm and (2) a minimum gradient support stabiliser. Synthetic results are presented to compare the performance of the many-layer inversion for different stabiliser functions. The minimum gradient support stabiliser is observed to improve the performance of the many-layer inversion for a layered subsurface, being able to reproduce both smooth and sharp vertical variations of the model parameters. Implementation of the alternative stabilisers into existing surface nuclear magnetic resonance inversion software is straightforward and requires little modification to existing codes.

INTRODUCTION

Surface nuclear magnetic resonance (NMR) is a non-invasive geophysical technique providing insight into aquifer properties. The measurement involves pulsing strong oscillatory currents in a surface coil in order to generate a measureable NMR signal at depth that originates from the immersion of hydrogen nuclei in the Earth's magnetic field (Schirov, Legchenko and Creer 1991; Hertrich 2008). To gain insight into the spatial variability of aquifer properties, the amplitude of the pulsed current is varied to manipulate the spatial origin of the measured signal. This procedure is typically referred to as a sounding, where weak and strong currents produce signals from shallow and greater depths, respectively. The end product is a dataset containing NMR signals of differing spatial origins (although many signals have overlapping spatial origins). An inversion framework is used to estimate the underlying spatial distribution of aquifer properties consistent with the observed data. This involves minimising an objective function that is used to penalise undesirable model characteristics, such as penalising models that do not closely reproduce the observed data.

Several inversion schemes are commonly employed in surface NMR, such as the time-step inversion (Legchenko and Valla 2002), the QT inversion that inverts the entire data cube simultaneously (Müller-Petke and Yaramanci 2010), joint-inversion schemes coupling NMR and time-domain electromagnetic (TEM) data (Behroozmand *et al.* 2012) or NMR and electrical resistivity (Günther and Müller-Petke 2012) data, and frequency-domain inversions (Irons and Li 2014). In each case, the inversion result is a model of the subsurface aquifer properties (such as depth profiles of the water content (WC) and relaxation times that describe the duration of the NMR signal). For the purposes of this discussion, we group surface NMR inversions into two

* denys.grombacher@geo.au.dk

categories: (1) inversions that use model domains consisting of many depth layers of fixed depths and thickness (referred to as many-layer inversions) and (2) inversions involving relatively small model domains with few depth layers, where the inversion determines the thickness of each layer (referred to as few-layer inversions). Each of the previously mentioned surface NMR inversion schemes may be implemented using either a many-layer or a few-layer model domain.

In many-layer inversions, the number of model parameters is generally quite large (when compared with few-layer inversions) and a regularisation term must be included in the objective function to stabilise the ill-posed inversion (Tikhonov and Arsenin 1977). The model that minimises the objective function thus balances satisfactory data fit with the magnitude of the regularisation term, which is controlled by the stabiliser function and the characteristics of the model. For optimal results, the selected stabiliser function should (1) return small values for the regularisation term when the model exhibits features consistent with *a priori* knowledge about the site and (2) return large values for models with characteristics inconsistent with *a priori* information about the site. The standard stabiliser in surface NMR is the L_2 smoothness stabiliser, which penalises the square of the variation between neighbouring model parameters. For a 1D depth sounding (the standard surface NMR experiment), this results in models that vary smoothly with depth. A limitation of such an approach is that the inversion struggles to reproduce sharp variations in WCs and relaxation times that may be present at the interface between lithologic layers of contrasting properties. To address this concern, an alternative stabiliser may be employed, such as the minimum support (Last and Kubik 1983), minimum gradient support (MGS) (Portniaguine and Zhdanov 1999), or stabilisers based on L_1-norms (e.g., Ellis and Oldenburg 1994; Loke, Acworth and Dahlin 2003). Mohnke and Yaramanci (2002) demonstrated the use of an L_1 stabiliser in surface NMR, but to our knowledge, the smoothness stabiliser remains the standard in surface NMR.

For few-layer inversions, a predetermined amount of layers is set and the inverted parameters are layer thicknesses, WCs, and relaxation times (Guillen and Legchenko 2002; Mohnke and Yaramanci 2002; Weichman *et al.* 2002). Due to the reduced number of model parameters (compared with the many-layer inversion), no regularisation term is included in the objective function. As a result, few-layer inversions are well suited to produce models with sharp contrasts in WC and relaxation times between neighbouring layers. An advantage of few-layer inversions is that uncertainty in the estimated profiles can be readily quantified using Bayesian approaches such as Markov chain Monte Carlo (Guillen and Legchenko 2002; Weichman *et al.* 2002) or simulated annealing (Mohnke and Yaramanci 2002). A limitation of few-layer inversions is that they struggle to reproduce smoothly varying subsurface parameters and can exhibit strong sensitivity to the initial starting model (i.e., the *a priori* specification of the number of layers and layer properties).

In practice, selection of a many-layer versus a few-layer inversion scheme in surface NMR typically depends on how much *a priori* information is available. Many-layer inversions are preferable given no *a priori* information, whereas few-layer inversions may be preferable if a known number of layers are present. Few-layer inversions are also commonly used if a well-stratified subsurface is expected, given that many-layer inversions typically result in models with smoothly varying subsurface parameters. However, this is not a result of the many-layer inversion scheme directly but a consequence that generally employs a smoothness stabiliser. To balance the advantages of both inversion strategies for layered subsurfaces (i.e., the ability to reproduce sharp variations in model parameters without requiring extensive *a priori* information), the performance of several stabiliser functions is compared against the smoothness stabiliser; an MGS stabiliser and a stabiliser employing an L_1-norm are investigated. Selecting alternative stabilisers does not require significant changes to existing inversions schemes. In this study, inversion is performed using an iteratively reweighted least-squares approach (Farquharson and Oldenburg 1998), where a Taylor expansion of the objective function is used to form the model update. Within this framework, alternative stabiliser functions are implemented by reweighting the roughness matrix within an L_2-norm (Fiandaca *et al.* 2015; Vignoli *et al.* 2015).

The MGS stabiliser (also referred to as focused or sharp inversion) provides the benefits of the many-layer inversion but while maintaining the ability to produce models with sharp contrasts in properties (Portniaguine and Zhdanov 1999). Briefly, the MGS stabiliser penalises the number of sharp contrasts in the model regardless of their magnitude, allowing the production of models with sharp interfaces between layers of relatively homogenous properties. The MGS stabiliser has been demonstrated to improve image sharpness for many-layer inversion schemes in magnetic (Portniaguine and Zhdanov 1999), gravity (Portniaguine and Zhdanov 1999), TEM (Vignoli *et al.* 2015), electrical resistivity tomography (Pagliara and Vignoli 2006), magnetotellurics (Zhdanov and Tolstaya 2004), seismic (Zhdanov, Vignoli and Ueda 2006), and induced polarisation (Blaschek, Hordt and Kemna 2008) studies. An additional stabiliser, employing an L_1-norm (instead of the L_2-norm present in the smoothness stabiliser), is also investigated. The L_1-norm penalises the absolute value of the variation in model parameters. This allows for sharper contrasts in model parameters compared with the smoothness stabiliser (Loke *et al.* 2003) but not as readily as the MGS stabiliser. Mohnke and Yaramanci (2002) found that surface NMR inversions that use an L_1 stabiliser are better suited to producing models with sharp contrasts compared with those that use the smoothness stabiliser. The L_1-norm is included in this comparison to compare its performance with that of the MGS stabiliser because of its ease of use. Synthetic results are presented to investigate the performance of each stabiliser for surface NMR inversion in the presence of a layered subsurface.

Results of the many-layer inversions are also compared against those of a few-layer inversion. Discussion about the implementation of alternative stabilisers into existing inversion packages and guidelines for the use of the MGS stabiliser are also given.

BACKGROUND

The surface nuclear magnetic resonance inverse problem

The standard measurement in surface NMR is the free induction decay (FID), which involves measurement of the NMR signal following a single-current pulse. To investigate the spatial variability of aquifer properties, the amplitude of the current pulse is altered to manipulate the spatial origin of the measured signal. The forward model is given by

$$\mathbf{d} = g(\mathbf{m}) + \mathbf{e}, \tag{1}$$

where **d** is a vector containing the measured NMR decays (for all current amplitudes for all time samples) and **m** is a vector containing the model parameters (WCs and T_2* in each depth layer). For a many-layer inversion, the number of depth layers and their thicknesses are predetermined. For a few-layer inversion, the model **m** also contains the layer thicknesses. The *g* function describes the physics of the forward problem; it contains: (1) information about the expected spatial origin of the measured signal corresponding to the excitation pulse type, current amplitude, and pulse duration; (2) a spatial weighting based on the receiver sensitivity at each location in the subsurface; (3) the impact of a conductive subsurface on depth penetration and signal phase; and (4) a scaling parameter to estimate the magnitude of the equilibrium magnetisation given the local magnetic field strength (local Earth's field strength) and aquifer temperature. **e** is a vector of the noise present in the data. Detailed derivation of the surface NMR forward model is given in Weichman, Lavely and Ritzwoller (2000).

To estimate the spatial distribution of aquifer properties, an inversion is used to predict the model that balances satisfactory data fit with the magnitude of the regularisation term. To determine this model, an objective function $\Phi(\mathbf{m})$, described by

$$\Phi(\mathbf{m}) = \phi_d(\mathbf{m}) + \phi_s(\mathbf{m}), \tag{2}$$

is minimised. The $\phi_d(\mathbf{m})$ term describes the L_2-norm misfit between the predicted data ($g(\mathbf{m})$) and the observed data (normalised by the data uncertainty), whereas $\phi_s(\mathbf{m})$ is the stabiliser function that determines the magnitude of the regularisation term for the current model **m**. The $\phi_d(\mathbf{m})$ term is given by

$$\phi_d(\mathbf{m}) = \left\| \mathbf{Q}_d \left(\mathbf{d} - g(\mathbf{m}) \right) \right\|_{L_2}^2, \tag{3}$$

where $\mathbf{Q}_d^T \mathbf{Q}_d = \mathbf{C}_d^{-1}$, i.e., the inverse of the data covariance matrix. The stabiliser function is described by

$$\phi_s(\mathbf{m}) = \left\| \mathbf{Q}_R \mathbf{R}\mathbf{m} \right\|_{\eta}^2, with\, \eta = L_2\, or\, L_1\, or\, MGS, \tag{4a}$$

and is necessary to stabilise the ill-posed inversion by penalising models that exhibit undesired traits. $\mathbf{Q}_R$ is a matrix used to weight the relative importance of the stabiliser function for each model constraint; $\mathbf{Q}_R^T \mathbf{Q}_R = \mathbf{C}_R^{-1}$, where $\mathbf{C}_R$ is a matrix containing the variances of the constraints. The **R** matrix is called the roughness matrix and is used to calculate the first-order difference between the model parameters in neighbouring depth layers. The η parameter corresponds to the norm used by the stabiliser (L_2 or L_1 or MGS). In this study, the different norms are implemented using a reweighting matrix $\mathbf{W}(\mathbf{m})$ and an L_2-norm, where the stabiliser function is given by

$$\phi_s(\mathbf{m}) = \left\| \mathbf{Q}_R \mathbf{W}(\mathbf{m}) \mathbf{R}\mathbf{m} \right\|_{L_2}^2. \tag{4b}$$

The form of $\mathbf{W}(\mathbf{m})$ corresponds to the specific norm desired and can be determined by equating equation (4b) with the equations describing the stabilisers in the following section. Equation (4b) indicates that selection of a norm different than L_2 (the smoothness case) does not require significant modifications to existing inversion codes; it only requires the inclusion of an additional weighting matrix within the stabiliser.

To find the model **m** that minimises equation (2), an iteratively reweighted least-squares approach is used (Farquharson and Oldenburg 1998), where the Taylor expansion of the objective function is used to determine the model update. This involves updating the estimated model iteratively, ultimately converging on a model that minimises the objective function. Details about the inversion scheme employed in this manuscript are given in Auken and Christiansen (2004), Vignoli *et al.* (2015), and Fiandaca *et al.* (2015). Note that the objective function (equation (2)) does not contain a trade-off parameter that can be used to weight the relative importance of the ϕ_d and ϕ_s terms (the trade-off parameter is typically denoted by λ). The inversion scheme used in this study weighs these terms equally, where the importance of the stabiliser term is controlled through the $\mathbf{Q}_R$ matrix that weighs the relative importance of the stabiliser for each model parameter.

Selecting a stabiliser function

The stabiliser function stabilises the inversion and allows the production of models with a desired property. This is done by penalising models that exhibit an undesired trait. Equations (5a), (5b), and (5c) illustrate the equations for a smoothness (L_2) stabiliser (the standard stabiliser in surface NMR inversions), the L_1 stabiliser, and the minimum-gradient support stabiliser, respectively:

$$\phi_s(\mathbf{m}) = \sum_k \left(\frac{(\Delta \mathrm{m})_k}{\sigma_k} \right)^2 \tag{5a}$$

$$\phi_s(\mathbf{m}) = \sum_k \sqrt{\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k} \right)^2} \tag{5b}$$

$$\phi_s(\mathbf{m}) = \frac{1}{\beta}\sum_k \frac{\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k}\right)^2}{\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k}\right)^2 + 1}. \quad (5c)$$

The $(\Delta m)_k$ term corresponds to the first-order difference of the constrained parameters for the kth constraint, i.e., $(\Delta m)_k = m_{j(k)} - m_{i(k)}$, where j(k) and i(k) represent the indices in the model vector of the parameters linked through the kth constraint. For the L_2 and L_1 stabilisers, the σ_k term represents the strength of the constraint because it controls the relative importance in the stabiliser function for the kth constraint. Equation (5a) indicates that the smoothness $\phi_s(\mathbf{m})$ increases proportional to the square of the difference between neighbouring model parameters. As such, sharp variations result in larger $\phi_s(\mathbf{m})$ and larger $\Phi(\mathbf{m})$. The minimisation will therefore return smoothly varying models, as models with sharp transitions will be penalised. The L_1 stabiliser (equation (5b)) penalises the absolute value of the difference in model parameters instead of the square of difference. As a result, smoothly varying models are still favoured by the L_1-norm, but sharp variations are penalised much less compared with the smoothness stabiliser. For both the L_2 and L_1 stabilisers, selection of σ_k controls the smoothness of the final model; large σ_k places little importance on the smoothness, allowing more erratic profiles to be produced in order to further minimise $\phi_d(\mathbf{m})$, whereas small σ_k places more importance on model smoothness at the expense of a larger data misfit.

If *a priori* knowledge suggests sharp transitions are likely at a particular site, selection of a smoothness stabiliser is suboptimal given that it penalises models with characteristics expected to be representative of the local hydrogeology. In this case, an alternative stabiliser may provide improved performance. For example, the MGS stabiliser (Portniaguine and Zhdanov 1999) presents a more efficient implementation of *a priori* knowledge of blocky structures. In this case, $\phi_s(\mathbf{m})$ is given by equation (5c); the form of the MGS stabiliser in equation (5c) is chosen to be consistent with Vignoli *et al.* 2015. This form of the MGS stabiliser presents a parameterisation allowing a simple understanding of the physical meaning of β and σ_k. Consider the effect of the MGS stabiliser in three regimes. In the $\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k}\right)^2 \gg 1$ limit, which describes the sharp change in model parameters at the interface between layers of contrasting properties, the contribution to $\phi_s(\mathbf{m})$ approaches $1/\beta$. Therefore, the presence of a sharp transition in the model parameters is not penalised based on the magnitude of the model variation (as in the smoothness case) but penalised a fixed amount. In the $\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k}\right)^2 \approx 1$ regime, the contribution to $\phi_s(\mathbf{m})$ scales approximately with the square of the difference in model parameters. In the $\left(\frac{(\Delta \mathrm{m})_k}{\sigma_k}\right)^2 \ll 1$ regime, there is little penalisation and the contribution to $\phi_s(\mathbf{m})$ is small. This indicates that the MGS stabiliser will not severely penalise models containing sharp transitions but will search for models with as few sharp transitions as possible with relatively homogenous properties between these sharp transitions (Portniaguine and Zhdanov 1999). σ_k and β effectively control the extent of homogeneity within a layer and the number of sharp transitions present in the final model, respectively. The value of β does not directly control the number of sharp transitions present in the estimated model, but its magnitude does influence the number of transitions present. Models corresponding to large values of β have more transitions than models with small β.

Implementation of each norm in this study is done using the weighting matrix $\mathbf{W}(\mathbf{m})$, determined by equating equation (4b) with equations (5a), (5b), or (5c). Note that, for the L_1 and MGS stabilisers, $\mathbf{W}(\mathbf{m})$ depends on the current model, requiring that $\mathbf{W}(\mathbf{m})$ be recalculated every iteration. The computational cost of updating $\mathbf{W}(\mathbf{m})$ is not significant, and each inversion proceeds at similar speeds in the case of a 1D surface NMR sounding. The stabiliser can also take other forms to describe different *a priori* conditions. In this manuscript, the L_1 and MGS stabilisers are selected based on their less severe penalisation of models containing sharp transitions in model parameters compared with the smoothness stabiliser.

RESULTS

Three synthetic surveys are presented to compare the utility of the L_1 and MGS stabilisers against the smoothness stabiliser for many-layer surface NMR inversions. Each stabiliser is also compared against the results of a few-layer inversion. Forward modelling and inversion of the synthetic data is performed using the AarhusInv software package (Auken *et al.* 2015), following the Behroozmand *et al.* (2012) forward implementation. Inversion is performed using the amplitudes of the NMR signals (i.e., the in and out of phase components of the data are not treated separately). The inversion also bounds the estimated WCs to fall between 0.1% and 100%, whereas the relaxation times are bound between 5 ms and 1.5 s. In each case, FID measurements are simulated using a coincident transmit/receive 100-m square loop, a 30-ms on-resonance excitation pulse, and 16 pulse moments sampled on the interval from 0.7 to 8.5 As. The selected pulse moments are chosen to span a range typical of surface NMR field experiments. The subsurface resistivity is 1000 Ωm in each case and is fixed during the inversion. This is equivalent to the inversions having *a priori* knowledge of the exact subsurface resistivity structure; a simple resistive subsurface is chosen to focus the comparison on the ability to estimate the subsurface parameters common to all surface NMR inversions (WC and relaxation times). In practice, it is common for non-joint NMR-TEM inversion schemes to treat the subsurface resistivity structure (estimated from a separate TEM or other electrical surveys) as fixed during the inversion. The Larmor frequency is set to

2138 Hz. Each inversion begins with a starting model corresponding to a half-space of 15% WC and T_2* of 150 ms. The data are binned into 12 time gates of logarithmically increasing width. The earliest and latest time gates are centred at 41 and 445 ms, respectively. Gaussian white noise is added to the time-gated data. To account for the varying widths of the time gates, the noise added to each time gate is scaled by the square root of the ratio of the time gate's width compared with the width of the first time gate. The stated noise levels refer to the standard deviation of the Gaussian used to generate the noise in the first time gate (the width of the first time gate is 7.1 ms). The subsurface is discretised into 25 depths of increasing thickness to a depth of 110 m. The shallowest layers have thicknesses of 1.5 m and increase to a thickness of ~10 m (layer thicknesses increase roughly logarithmically). Below 110 m, the subsurface is treated as a half-space. A model discretisation consisting of 25 depth layers was chosen to balance the opportunity to capture smoothly varying parameters without dramatically over-parameterising the subsurface. Increasing the number of depth layers places more importance upon the regularisation. Further discussion about the approach used to discretise the subsurface is given in Behroozmand *et al.* (2012). Note that the layer boundaries for the synthetic subsurface models occur at the same depths as layer interfaces in the model discretisation. In practice, the depth discretisation is unlikely to coincide with the true layer boundaries; in this case, it would cause either smearing between two layers or an error in identifying exact depth of the interface.

In each example, 200 noisy datasets are produced by adding different noise realisations to the same noise-free dataset. For the first three examples, the noise level is 20 nV (i.e., the standard deviation of the Gaussian used to randomly generate noise for the first time gate is 20 nV). Although the signal-to-noise ratio (SNR) in each case depends on the subsurface model, this level of noise produces an SNR of ~50–80 for the three examples. For each noisy dataset, a WC and a T_2* profile are estimated using a many-layer inversion with a smoothness stabiliser, a many-layer inversion with an L_1 stabiliser, a many-layer inversion with an MGS stabiliser, and a few-layer inversion. The 200 estimated WC and T_2* profiles produced by each inversion scheme are used to form histograms of the WC and T_2* values in each depth layer. The top two rows of Figure 1 illustrate several examples of how the

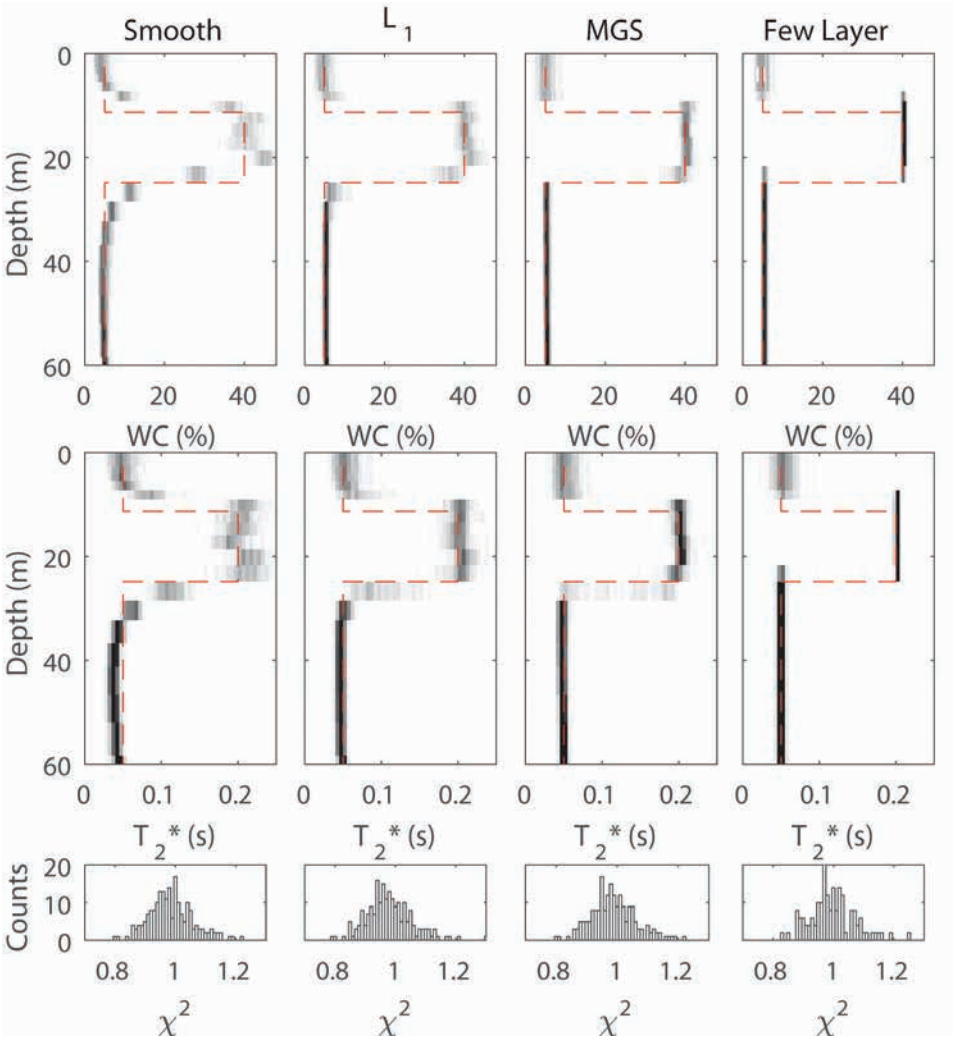

Figure 1 Histograms showing (top row) the WC and (middle row) T_2* profiles estimated from the inversion of 200 independent noisy datasets. The bottom row illustrates a histogram of χ^2 for all 200 inversions. The dashed red line shows the true model (a three-layer system with a single aquifer). Dark and white colours indicate bins with many and no counts, respectively. Columns left to right show the results for a many-layer inversion using a smoothness stabiliser, a many-layer inversion using an L_1 stabiliser, a many-layer inversion using an MGS stabiliser, and a few-layer inversion with three layers. The noise level is 20 nV. Black and white bins have 70 and 0 counts, respectively.

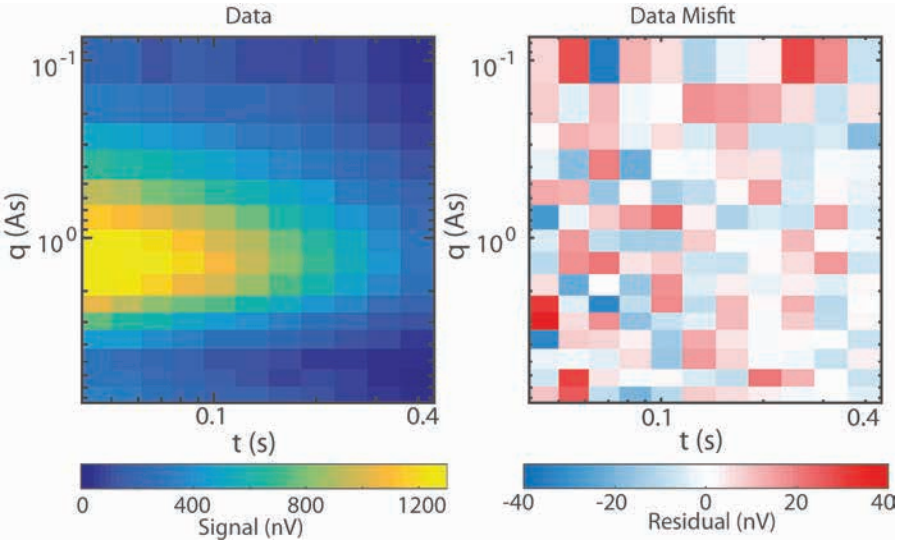

Figure 2 (A) One of the 200 noisy datasets produced by the subsurface model in Figure 1. (B) An example of the data residual produced by the many-layer inversion using the MGS stabiliser. This data residual corresponds to a χ^2 of 1.02 and is representative of that produced by other inversions with similar χ^2.

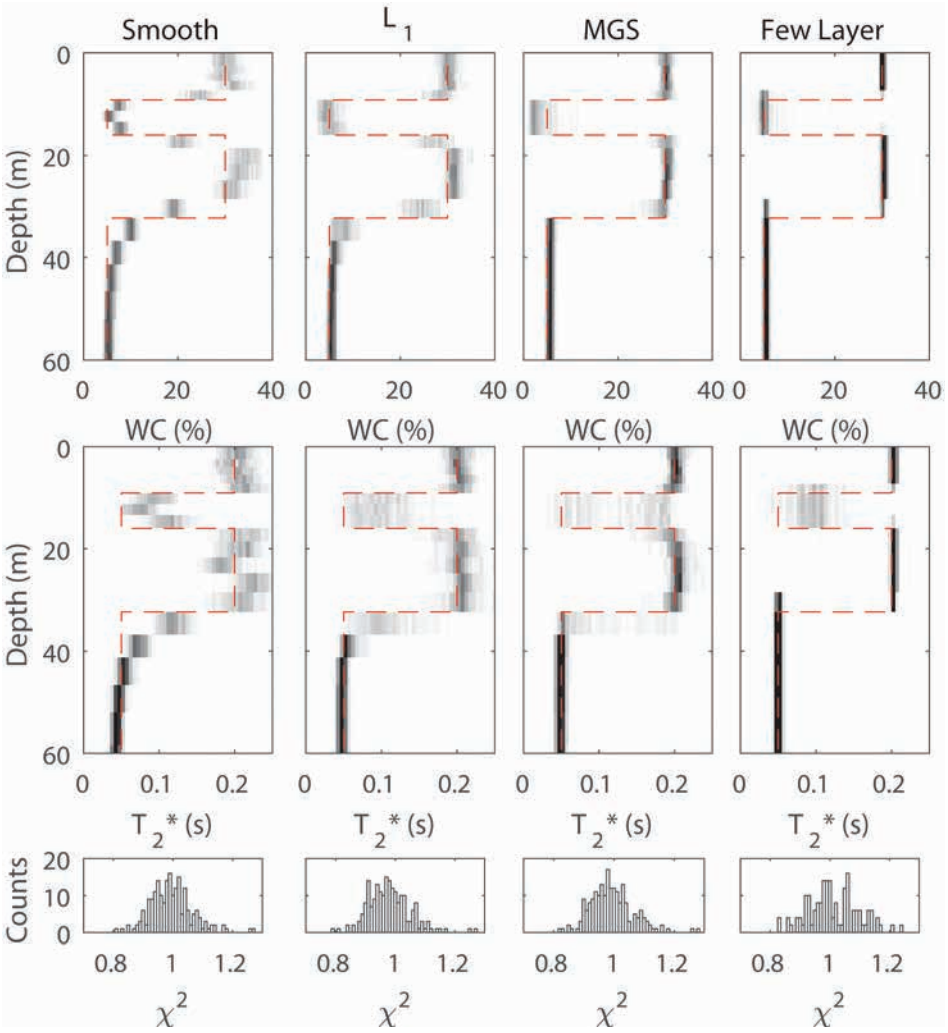

Figure 3 Histograms showing the (top row) WC and (middle row) T_2* profiles estimated from the inversion of 200 independent noisy datasets. The bottom row illustrates a histogram of χ^2 for all 200 inversions. The dashed red line shows the true model (a four-layer system consisting of two aquifers). Dark and white colours indicate bins with many and no counts, respectively. Columns left to right show the results for a many-layer inversion using a smoothness stabiliser, a many-layer inversion using an L_1 stabiliser, a many-layer inversion using an MGS stabiliser, and a few-layer inversion with three layers. The noise level is 20 nV. Black and white bins have 70 and 0 counts, respectively.

histograms will be illustrated. The *y*-axes correspond to depth, the *x*-axes correspond to either WC or T_2*, and the colour scale indicates the number of counts present in each bin (black indicates a high number of counts and white indicates no counts). The WC

and T_2* bins are 0.5% and 5 ms wide, respectively. The histograms allow the uncertainty of the resulting profiles to be estimated by examining the distribution of WCs and T_2* values within each depth layer. Low and high uncertainty correspond to depth layers with narrow black distributions and wide light grey distributions, respectively. Note that the histograms do not illustrate the full range of equivalent solutions as each inversion begins with the same starting model. However, the histograms remain a useful tool to provide insight into the uncertainty in the estimated profiles. For each stabiliser, the results for single regularisation strength are shown. The strength of the regularisation is selected to produce the smoothest model that fits the data within error. The constraint strengths σ_k used in this study are relative to the magnitude of the model parameter $m_{i(k)}$, i.e., the constraint strength is effectively controlled by a parameter denoted σ_{rel}, where $\sigma_k=(\sigma_{rel} m_{i(k)} - m_{i(k)})$. The inversion in this study is carried out in a logarithmic model space; therefore, $(\Delta m)_k$ becomes $log\left(m_{j(k)}+\sigma_k\right)-log\left(m_{j(k)}\right)$ and σ_k is estimated by subtracting the log-transformed parameter from the log-transformed upper limit of its confidence interval, i.e., σ_k becomes $log\left(m_{i(k)}+\sigma_k\right)-log\left(m_{i(k)}\right)$ Therefore, the penalty $p=\frac{(\Delta m)_k}{\sigma_k}$ of equations (5a)–(5c) can be expressed in terms of σ_{rel} as

$$p=\frac{log\left(m_{j(k)}\right)-log\left(m_{i(k)}\right)}{log\left(m_{i(k)}-(\sigma_{rel}-1)\cdot m_{i(k)}\right)-log\left(m_{i(k)}\right)}=\frac{log\left(m_{j(k)}/m_{i(k)}\right)}{log\left(\sigma_{rel}\right)}. \quad (6)$$

For example, σ_{rel} = 1.1 means that model parameter variations of ~10% are acceptable (i.e., should not be penalised severely). Given the noise level of 20 nV, σ_{rel} = 1.5 was used for the smoothness and L_1 stabilisers, whereas for the MGS stabiliser, σ_{rel} = 1.1 and β = 50. Note that for each stabiliser, the WCs and T_2* parameters are given the same constraint strengths. Further discussion about the selection of the MGS stabiliser parameters is given in the discussion section.

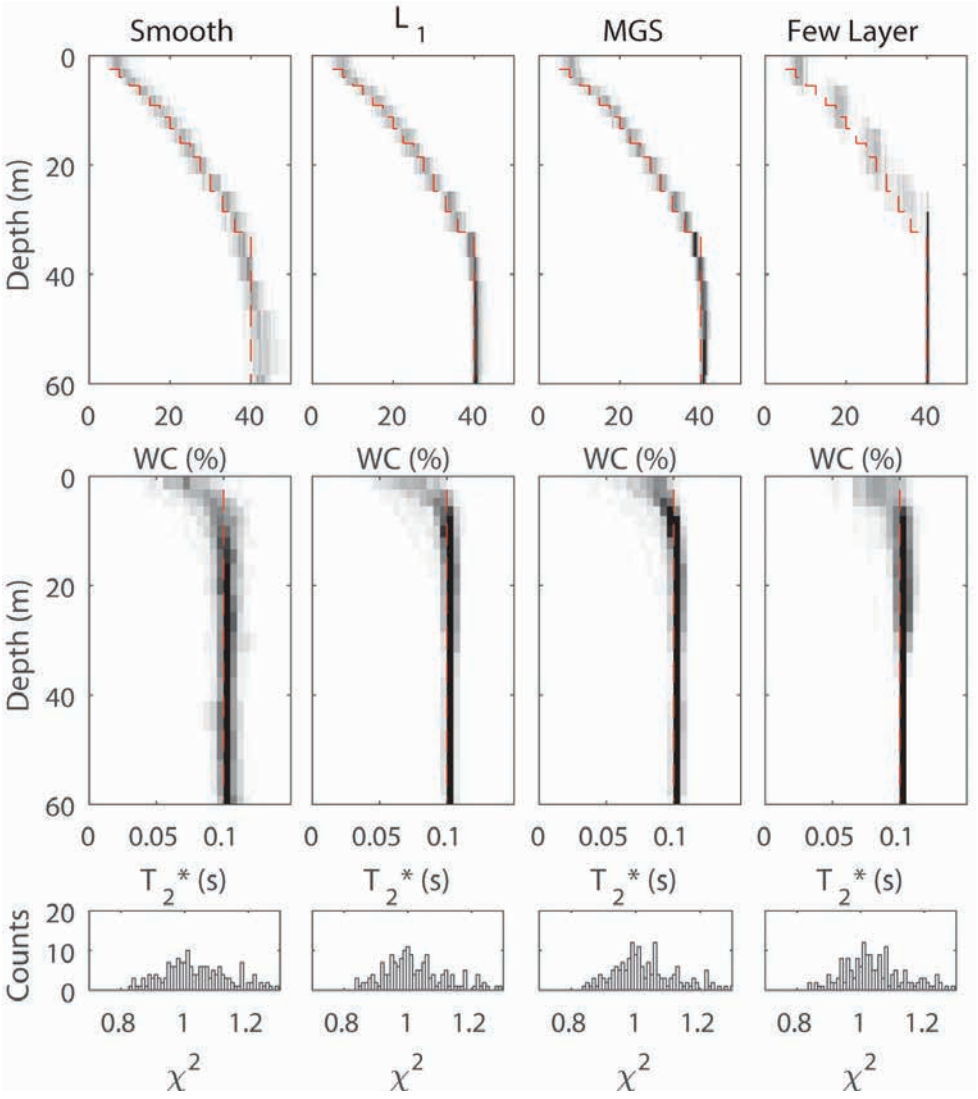

Figure 4 Histograms showing (top row) the WC and (middle row) T_2* profiles estimated from the inversion of 200 independent noisy datasets. The bottom row illustrates a histogram of χ^2 for all 200 inversions. The dashed red line shows the true model (a smoothly increasing WC profile with a homogenous T_2*). Dark and white colours indicate bins with many and no counts, respectively. Columns left to right show the results for a many-layer inversion using a smoothness stabiliser, a many-layer inversion using an L_1 stabiliser, a many-layer inversion using an MGS stabiliser, and a few-layer inversion with three layers. The noise level is 20 nV. Black and white bins have 70 and 0 counts, respectively.

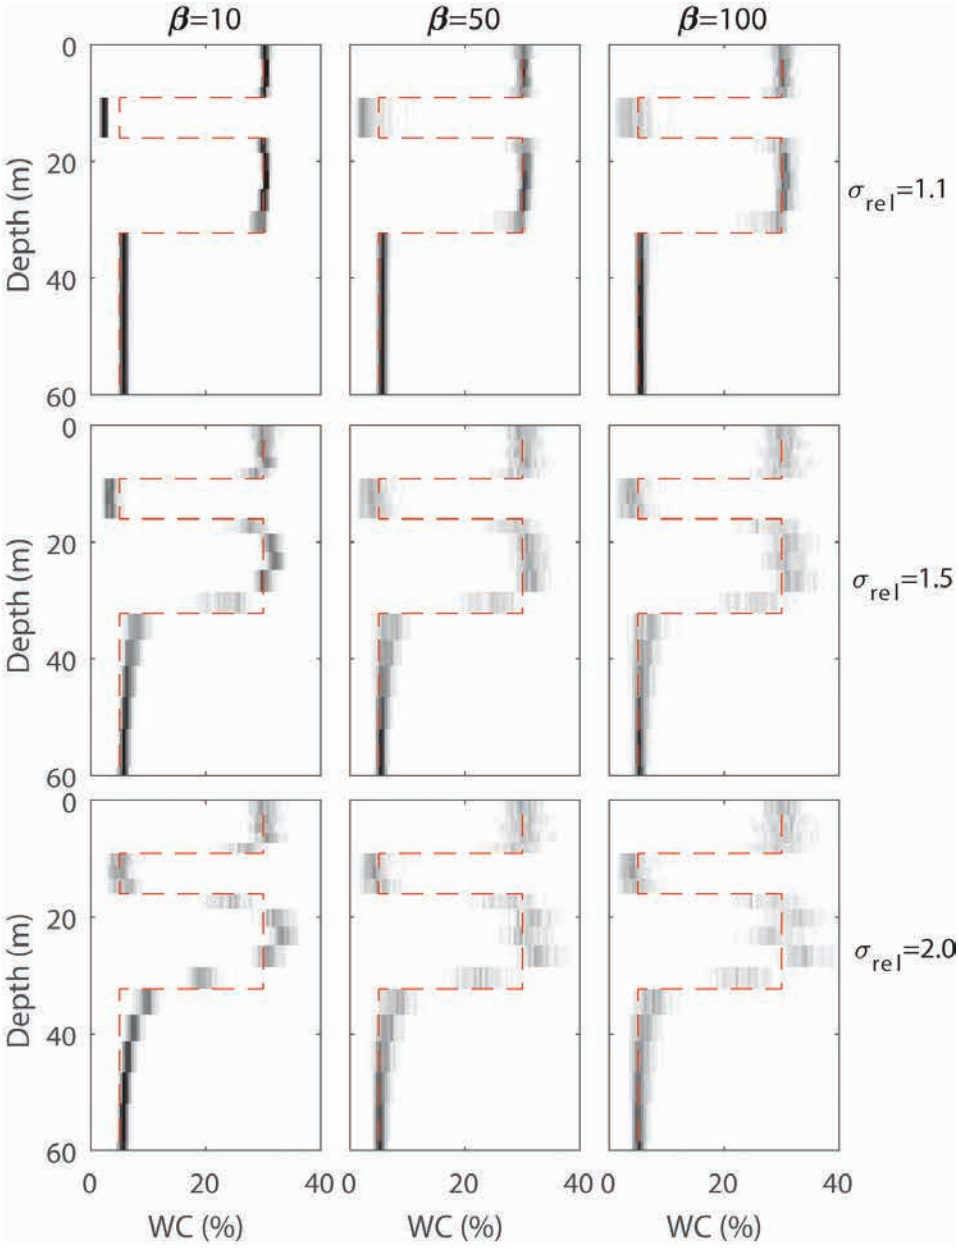

Figure 5 Histograms showing the influence of σ_{rel}, and β on the estimated WC profile for the MGS stabiliser. The histograms are formed of the WC profiles resulting from the same 200 noisy datasets as in Figure 3. Each row and column corresponds to a particular σ_{rel}, and β, respectively. Dark and white colours indicate bins with many and no counts, respectively. The top left and bottom right represent the strongest and weakest regularisations, respectively. The noise level is 20 nV. Black and white bins have 70 and 0 counts, respectively.

Figures 1, 3, and 4 contrast the performance of each stabiliser. The top row in each figure illustrates the estimated WC profiles, the middle row illustrates the estimated T_2* profiles, and the bottom row shows a histogram of the resulting χ^2 in each case. χ^2 is unitless, as the data misfit (nV) is normalised by the data uncertainty (nV). χ^2 histograms clustered around 1 indicate good data fit (χ^2 is close to 1 because it is normalised by the number of data points). Columns one to three correspond to many-layer inversions that use a smoothness stabiliser, an L_1 stabiliser, and an MGS stabiliser, respectively. Column four illustrates the results of the few-layer inversion that is given the correct number of layers. The true WC and T_2* profiles in each case are illustrated by the red dashed lines.

The first example (Figure 1) is a three-layer system containing a single aquifer. The aquifer is 14 m thick (from 11- to 25-m depths) with a WC of 40% and T_2* = 200 ms. The layers above and below this aquifer have reduced WC (5%) and faster T_2* (50 ms). The smoothness inversion (left column) accurately resolves the increased WC and T_2* layer, producing reliable estimates of the WC and T_2* magnitudes in all three layers. The large contrast at the upper boundary is well resolved by the smoothness stabiliser, whereas the lower boundary is smoothed over a larger depth range. The L_1 stabiliser (column 2) resolves the properties of all three layers well, capturing the sharp contrast at the upper layer boundary while also estimating a sharper transition to low WC and T_2* at the lower layer boundary com-

pared with the smoothness stabiliser. The MGS stabiliser (column 3) produces similar results as the L_1 stabiliser and resolves both layer boundaries well. The estimated WCs and T_2* within the aquifer (layer 2) show less variation for the MGS case than the L_1 and smoothness stabiliser cases (darker narrower histograms). The few-layer inversion, which was given the correct number of layers *a priori*, accurately reproduces the true model. In this example, the blocky true model is reproduced with high precision by the L_1, MGS, and few-layer inversions, whereas the smoothness results make the identification of the lower layer boundary more difficult. The bottom column of Figure 1 indicates that each inversion approach was able to fit the data to similar levels, with the data residual norms clustered around 1. To give an example of the noisy data and quality of data fit, Figure 2 illustrates the first of the 200 noisy datasets (left panel) and the data residual (right panel) produced by the MGS stabiliser. The residual shows no structure (i.e., no large areas with consistent sign) and has a magnitude consistent with the noise level. The χ^2 in this example is 1.02. Figure 2B is representative of the residual produced by inversions resulting in similar magnitude χ^2.

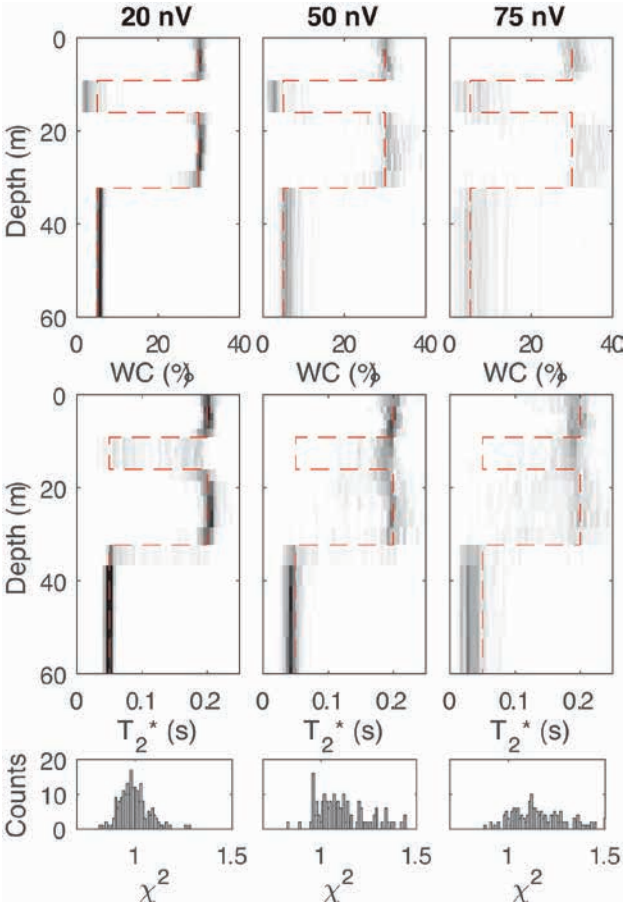

Figure 6 Histograms showing the performance of the MGS stabiliser at varying noise levels. Each column corresponds to a particular noise level. The top and middle rows show histograms of the WC and T_2*, respectively, following the inversion of 200 noisy datasets. The bottom row illustrates a histogram of χ^2 for all 200 inversions. The dashed red line shows the true model (same as in Figure 3). Dark and white colours indicate bins with many and no counts, respectively. Black and white bins have 70 and 0 counts, respectively.

The second example (Figure 3) is a slightly more complicated four-layer system containing two aquifers. The two aquifers (layers 1 and 3) have WC of 30% and T_2* = 200 ms. The layer separating these aquifers and the bottom layer have reduced WC (5%) and T_2* (50 ms). In this case, the smoothness inversion (left column) produces a smoothed version of the layered subsurface. The WC and T_2* are well estimated in each layer, but it is difficult to identify the layer boundaries given the smooth variations. For example, the upper and lower layer boundaries for layer 3 (the lower aquifer) are both spread over a 5- to 10-m-depth range. The L_1 inversion also reproduces the WC and T_2* magnitudes well while better identifying the boundaries between the upper three layers. The MGS stabiliser produces similar results as the L_1 stabiliser but with the lower boundary between layers 3 and 4 being more sharply resolved. The WC and T_2* values estimated within layers 1 and 3 are also more homogenous than the L_1 stabiliser (observed by narrower darker histograms for the MGS case compared with the L_1 case). Both the L_1 and MGS stabilisers struggle to resolve the magnitude of T_2* in the second layer. This is a consequence of the low WC at these depths, which reduces the ability to resolve the magnitude of T_2*. For the few-layer inversion, which is given the correct number of layers, the true model is well reproduced. The estimated T_2* value in layer 2 also has a higher uncertainty (noted by the wide histogram). Overall, the few-layer result is quite similar to that produced by the MGS stabiliser, with each layer boundary being well resolved. The L_1 and smoothness inversions are less able to capture the large contrast in properties at the lower boundary between layers 3 and 4. The bottom row of Figure 3 indicates that each inversion provides a similar level of data fit.

The third example (Figure 4) tests the performance of each stabiliser given a subsurface containing a smooth variation in WC. In this case, the WC is 10% at the shallowest depth and increases roughly linearly to 40% at 37-m depth; T_2* is equal to 100 ms at all depths. Below 37 m, a homogenous 40% WC layer is present. The smoothness inversion (left column) accurately captures the slowly increasing WC profile while estimating a smooth transition to lower WC at depth (below ~37 m). The L_1 stabiliser produces similar results as the smoothness case, capturing the smoothly increasing WC profile while better predicting a homogeneous WC below 37 m (narrow dark histograms). The MGS stabiliser also reproduces the true model well, with a similar prediction of the homogeneity below 37 m as the L_1 stabiliser. The T_2* profile is well resolved in all cases, except at the shallowest depths. The systematic bias towards underestimated T_2* at the shallowest depths likely results from the T_2* at these depths having little impact on the overall data fit (given that these depths correspond to the lowest WCs). For the few-layer inversion results, where the inversion is

given five layers, a blocky stepwise increasing WC is predicted, with the overall structure in the WC being captured. The WCs at depths above ~37 m are more uncertain for the few-layer inversion compared with those for the many-layer inversions (wide light grey histograms). Below 37 m, the few-layer inversion accurately estimates the WC. The bottom row of Figure 4 indicates that each inversion scheme produces similar levels of data fit. For some noise realisations, χ^2 is large (> ~1.3) and the data fit is reduced. While increasing the number of layers for the few-layer inversion will improve its ability to capture the smooth change in WC, the five-layer model is shown, given the preference for the model containing the fewest number of layers that provides satisfactory data fit.

Figures 1, 3, and 4 illustrate that the smoothness stabiliser is suboptimal when sharp layer boundaries are expected, and the selection of an alternative stabiliser can improve the performance of the many-layer inversion in the presence of a layered subsurface. Comparing the L_1 and MGS results indicates that the MGS stabiliser provides the best ability to reproduce a blocky subsurface structure when using a many-layer inversion. Even in a smoothly varying subsurface, the MGS stabiliser produces a reliable result. The benefit of the MGS stabiliser is that it is able to resolve blocky structures without requiring knowledge of the number of layers *a priori*; the MGS results even provide similar performance to a few-layer inversion given the correct number of layers. Note that for the depth discretisation and noise levels used in these examples, a fixed level of regularisation for the MGS stabiliser can be expected to provide flexible performance capable of resolving both smoothly varying and blocky subsurface structures. The few-layer inversion also performs well for a layered subsurface, provided that a sufficient number of layers are used in the inversion.

DISCUSSION

The selection of a many-layer versus a few-layer inversion scheme should consider the available *a priori* information about the site. If little information about the subsurface is present, such as whether a layered or smoothly varying subsurface is present, the many-layer inversion offers the benefits requiring no *a priori* specification about the number of layers. A preliminary many-layer inversion can also be used to inform a subsequent few-layer inversion, where the many-layer result can be used to provide an initial model and helps in choosing the number of layers for the few-layer inversion. Whether the result of the many-layer inversion is to be used as the final estimated model or as a starting model for a few-layer inversion, it is beneficial to use a stabiliser well suited to producing models with features consistent with the expectations of the subsurface. Therefore, if a layered subsurface is expected, the standard smoothness stabiliser is suboptimal. Both the L_1 and MGS stabiliser improve the ability of the many-layer inversion to reproduce blocky structures. However, results produced by a many-layer inversion that uses an L_1 or MGS stabiliser are not necessarily more accurate than those produced by a smoothness stabiliser. Given equal levels of data fit, the results produced by each stabiliser represent equally likely models. Similarly, few-layer inversions providing similar data fits as the many-layer inversion also provide equally likely models. To decide between the potential models, additional geologic information should be considered, such as the depositional environment, which may help inform whether a layered or smoothly varying subsurface is more likely. The advantages of the L_1 and MGS stabilisers is that they provide a means for the many-layer inversion to more readily produce sharp contrasts in properties.

Practical considerations for using the minimum gradient support stabiliser in surface nuclear magnetic resonance

We now focus on the MGS stabiliser given that it provides the best ability to reproduce a layered subsurface when using a many-layer inversion. The contribution of the MGS stabiliser to the objective function is controlled by two parameters, σ_k and β. In contrast, the smoothness and L_1 stabilisers are controlled by a single parameter σ_k. The additional parameter for the MGS stabiliser complicates the decision as to how the regularisation strength should be selected. For the smoothness and L_1 cases, the general rule for selection of the regularisation strength is that the smoothest model producing satisfactory data fit should be selected; otherwise, the inversion may introduce spurious features into the estimated profiles in an attempt to overfit the data. For the MGS stabiliser, selection of σ_k and β requires balancing the desired level of homogeneity within a layer with the number of sharp contrasts present in the estimated models. To illustrate the impact of each parameter on the performance of the MGS stabiliser, Figure 5 shows the WC profiles for MGS inversions performed with different combinations of σ_{rel} and β given the same suite of 200 noisy datasets used to form Figure 3 (the two-aquifer system). Each row and column corresponds to a particular σ_{rel} and β, respectively. The top middle panel is a reproduction of the MGS WC profiles from Figure 3. For small σ_{rel} (top row), the intralayer homogeneity is high, noted by dark narrow histograms. For larger σ_{rel} (rows 2 and 3), the intralayer homogeneity is reduced (wider light grey histograms) and the results begin to more closely resemble the smoothness WC profile in Figure 3. For increasing β (left column to right column), the likelihood of additional sharp contrasts is increased. In this example, this results in a blurring of the layer boundaries due to the reduced penalisation of additional sharp contrasts in the final model. At this noise level (20 nV), each level of regularisation fits the data to similar levels, except for the top left panel that produces a slightly poorer data fit. Given that the motivation to use an MGS stabiliser is to improve the ability of the many-layer inversion to reproduce a layered subsurface, we recommend selecting a low σ_{rel} value (e.g., fixing σ_{rel} to 1.1). This ensures that relatively homogeneous layers are produced and effectively allows the regularisation strength to be controlled by specifying a β value. The selected β should be as small as possible while still providing satisfactory data fit. For the depth discretisation and noise lev-

els used in these examples, β = 50 was observed to provide good performance. The corresponding T_2* profiles (for the same σ_{rel} and β pairs) exhibit similar trends (not shown).

Choosing the regularisation strength also depends upon the SNR. To investigate the performance of the MGS stabiliser for varying noise conditions, Figure 6 illustrates the WC and T_2* profiles estimated using a many-layer inversion with an MGS stabiliser for noise levels of 20, 50, and 75 nV. The true subsurface model in this example is the same as in Figure 3. These noise levels roughly correspond to SNRs of ~60, ~25, and ~15, respectively. At the lowest noise condition (20 nV), the true subsurface model is well reproduced, except for the T_2* value in layer 2. For noise levels of 50 and 75 nV, the estimated WC and T_2* profiles have larger uncertainty (wider light grey histograms) and no longer resolve the T_2* contrast between layer 2 and its neighbours. The data fit is also reduced at higher noise levels (as illustrated by the χ^2 histograms in the bottom row of Figure 6). In several cases with higher χ^2, the data residual plots show a structure indicating a poor data fit. In these cases, the estimated profiles would be treated with high uncertainty. Note that the histograms effectively hide these poor profiles, as they are only one of 200 results. In practice, a high noise level may cause the MGS stabiliser to predict a sharp boundary at an incorrect depth or where no contrast exists at all. In this limit, it may be preferable to use the MGS stabiliser to inform the number of depth layers present and to use this information as the *a priori* number of layers for a subsequent few-layer inversion. The few-layer inversion can then be used to readily quantify the uncertainty in the estimated profiles. Alternatively, in the high noise limit, it may be preferable to use the smoothness inversion given that strong smoothness regularisation may limit the introduction of spurious sharp contrasts (at the expense of resolving layer boundaries). At noise levels greater than that investigated in Figure 6 (which may happen depending on local noise conditions), the profiles show even greater uncertainty.

The σ_k and β parameters also depend on the depth discretisation used in the many-layer inversion. As such, we recommend that synthetic studies with similar models to those considered in Figures 1, 3, and 4 be performed using the same depth discretisation, which will be used in the inversion of field data and with noise levels similar to the field data. This will help inform the range of σ_k and β parameters likely to provide satisfactory performance and will provide insight into how capable the inversion is of resolving a synthetic model with features similar to those present in the WC and T_2* profiles produced by the field data. Similar synthetic tests would also help select regularisation strength and understand the resolution of the final models for the smoothness and L_1 stabilisers.

CONCLUSIONS

The ability of the many-layer surface NMR inversion to reproduce a layered subsurface is compared for several stabiliser functions. The standard stabiliser (smoothness stabiliser) penalises sharp transitions in subsurface properties and is poorly suited to imaging layered subsurfaces. Two alternative stabilisers, an L_1 stabiliser and a minimum-gradient support stabiliser, were found to improve the ability to identify sharp contrasts in layer properties. The MGS stabiliser is observed to greatly improve the ability of the many-layer inversion to reproduce blocky structures. Although the L_1-norm is observed to also provide improved performance compared with the smoothness approach for layered subsurfaces, its improvement is less than that of the MGS stabiliser. Improving the utility of the many-layer inversion in a layered environment benefits both the scenario where the model produced by the many-layer inversion is used for building the conceptual model of the subsurface and the scenario where the many-layer inversion is used to build an initial model and an estimate of the number of layers needed for a subsequent few-layer inversion.

The form of the MGS stabiliser employed in this study provides a simple understanding of the role played by the two tunable parameters in the stabiliser function. The extent of WC and T_2* homogeneity within a layer for the MGS stabiliser is controlled by σ_k (we recommend that variations greater than 10% be penalised), whereas the number of sharp transitions present in the final model is influenced by β (small and large β lead to less and more transitions, respectively). Despite two tunable parameters, selection of appropriate inversion parameters is straightforward and a single set of parameters is observed to provide accurate results for a broad range of subsurface models. For the inversion of field data, we recommend selecting inversion parameters based on observations from synthetic tests with simple models (like those present in Figures 1–4), the same model discretisation, and similar noise conditions as the field data. In high-noise conditions, it may be preferable to use the MGS many-layer inversion to inform a few-layer inversion, allowing the uncertainty of the estimated profiles to be more readily quantified. Alternatively, the standard smoothness stabiliser may be preferable to the MGS stabiliser in high-noise environments in order to limit the introduction of spurious sharp contrasts that may be interpreted as layer boundaries. However, this comes at the expense of resolving sharp contrasts. In summary, the MGS stabiliser provides an effective means to improve the flexibility of the many-layer surface NMR inversions.

ACKNOWLEDGEMENTS

Denys Grombacher was supported by funding from the Danish Council for Independent Research under postdoctoral grant DFF-5051-00002. Ahmad A. Behroozmand was partly supported through a funding from the Danish Council for Independent Research.

REFERENCES

Auken E. and Christiansen A.V. 2004. Layered and laterally constrained 2D inversion of resistivity data. *Geophysics* **69**(3), 752–761.

Auken E., Christiansen A.V., Kirkegaard C., Fiandaca G., Schamper C., Behroozmand A.A. *et al.* 2015. An overview of a highly versatile forward and stable inverse algorithm for airborne, ground-based, and

borehole electromagnetic and electric data. *Exploration Geophysics* **46**(3), 223–235.

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* **77**(4), WB191–WB200.

Blaschek R., Hordt A. and Kemna A. 2008. A new sensitivity-controlled focusing regularization scheme for the inversion of induced polarization data based on the minimum gradient support. *Geophysics* **73**(2), F45–F54.

Ellis R.G. and Oldenburg D.W. 1994. Applied geophysical inversion. *Geophysical Journal International* **116**, 5–11.

Farquharson C.G. and Oldenburg D.W. 1998. Non-linear inversion using general measures of data misfit and model structure. *Geophysics* **134**, 213–217.

Fiandaca G., Doetsch J., Vignoli G. and Auken E. 2015. Generalized focusing of time-lapse changes with applications to direct current and time-domain induced polarization inversions. *Geophysical Journal International* **203**(2), 1101–1112.

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.

Günther T. and Müller-Petke M. 2012. Hydraulic properties at the North Sea island of Borkum derived from joint inversion of magnetic resonance and electrical resistivity soundings. *Hydrology and Earth System Sciences* **16**(9), 3279–3291.

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

Irons T.P. and Li Y. 2014. Pulse and Fourier transform surface nuclear magnetic resonance: comprehensive modeling and inversion incorporating complex data and static dephasing dynamics. *Geophysical Journal International* **199**, 1372–1394.

Last B.J. and Kubik K. 1983. Compact gravity inversion. *Geophysics* **48**(6), 713–721.

Legchenko A. and Valla P. 2002. A review of the basic principles for proton magnetic resonance sounding measurements. *Journal of Applied Geophysics* **50**, 3–19.

Loke M.H., Acworth I. and Dahlin T. 2003. A comparison of smooth and blocky inversion methods in 2D electrical imaging surveys. *Exploration Geophysics* **34**, 182–187.

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**(1)–(2), 163–177.

Müller-Petke M. and Yaramanci U. 2010. QT inversion—Comprehensive use of the complete surface NMR data set. *Geophysics* **75**(4), WA199–WA209.

Pagliara G. and Vignoli G. 2006. Focusing inversion techniques applied to electrical resistance tomography in an experimental tank. *Proceedings of the 11th International Congress of the International Association for Mathematical Geology*.

Portniaguine O. and Zhdanov M.S. 1999. Focusing geophysical inversion images. *Geophysics* **64**(3), 874–887.

Schirov M., Legchenko A. and Creer G. 1991. A new direct non-invasive groundwater detection technology for Australia. *Exploration Geophysics* **22**(2), 333–338.

Tikhonov A.N. and Arsenin V.Y. 1977. *Solutions of Ill-Posed Problems*. Washington, DC: Winston.

Vignoli G., Fiandaca G., Christiansen A.V., Kirkegaard C. and Auken E. 2015. Sharp spatially constrained inversion with applications to transient electromagnetic data. *Geophysical Prospecting* **63**, 243–255.

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* **62**, 1290–1312.

Weichman P.B., Lun D.R., Ritzwoller M.H. and Lavely E.M. 2002. Study of surface nuclear magnetic resonance inverse problems. *Journal of Applied Geophysics* **50**(1)–(2), 129–147.

Zhdanov M. and Tolstaya E. 2004. Minimum support nonlinear parameterization in the solution of a 3D magnetotelluric inverse problem. *Inverse Problems* **20**(3), 937–952.

Zhdanov M.S., Vignoli G. and Ueda T. 2006. Sharp boundary inversion in crosswell travel-time tomography. *Journal of Geophysics and Engineering* **3**(2), 122–134.