Barcelona '10

D035

Error Modelling in Bayesian CSEM Inversion

J. Gunning* (CSIRO Earth Science and Resource Engineering) & M.E. Glinsky (CSIRO Earth Science and Resource Engineering)

SUMMARY

Mis-characterisation of the noise has significant potential to disrupt reservoir parameter estimates and uncertainties in geophysical
inversion. In all form of geophysical inversion, the "effective noise" used in the data misfit absorbs effects from approximate forward modelling in additional to environmental processing and measurement noise. For risk assessment, inversions require parameter
uncertainties, and these are best approached from a Bayesian angle. But parameter uncertainty estimates are dependent on the noise model at leading order. Modelling noise in particular can be strongly correlated, and will corrupt parameter uncertainty estimates if the correlations are not taken into account. In CSEM inversion, structural and resistivity parameters can be particularly difficult to disentangle, and their separation is rather vulnerable to systematic components of the noise. We present several ideas to manage the effect, two of which are easily incorporated into standard optimization and sampling schemes.

Introduction

In the last decade or so, the controlled source electromagnetic (CSEM) technique has become a popular tool in the hydrocarbon exploration game. The method is naturally suited to detecting resistive anomalies in the marine subsurface, hopefully due to the presence of hydrocarbon deposits, so regional geological knowledge must be sufficient to exclude the possibility of highly resistive rocks like evaporites, volcanics, or carbonates. Taken in conjunction with seismic data for geological and structural delineation, the tool is potentially a powerful discriminator between high and low gas saturation. The technique is a valuable complement to seismic methods, since these are well known to have difficulty in detection of gas saturation in AVO applications.

Many articles have appeared to date outlining the general nature of the CSEM acquisition framework (see, e.g. the CSEM "special section" of Geophysics, vol 72, No. 2). It is by now well known that CSEM inversion is a challenging problem (Constable et al., 1987; Gunning et al., 2010), not just because the forward energy propagation is heavily dispersive and thus resolution destroying, but also because the subsurface response is poorly modelled by a weak–scattering (Born) approximation when the conceivable range of resistivity variation of submarine rocks is several decades. Indeed, subsurface resistors behave almost as waveguides, so the overall field response in the presence of resistors is qualitatively very different to typical background response caused by conducting shales, hence the failure of weak scattering theory. The imaging problem is thus controlled by both dispersive physics and strong nonlinearity, which, when combined with expensive forward models, makes the problem qualitatively similar to the difficult history matching problem in petroleum reservoir characterization.

One of the main consequences of the quasi–waveguide like energy flow along subsurface resistors (if present) is that received signals are largely controlled by the resistivity–thickness product (RTP) of these resistors. This can be shown both numerically and (in 1D) analytically from steepest descent approximations to the solutions (Loseth, 2007). Imaging for structure – i.e. extracting thickness – is thus very challenging in the absence of complementary information, like seismic delineation of resistor boundaries.

These issues, being germane to the forward physics, affect both regularizing and Bayesian approaches to inversion. But Bayesian approaches enable a probabilistically coherent statement of what inversion degeneracies or multimodalities imply about subsurface structure, and are able to embrace multiple data types with much more logical consistency. This is our preferred approach. Further, an important goal of modern Bayesian inversion methods is an assessment of the range of uncertainty of important reservoir parameters, as a means of informing risk–management decisions. Typically this will involve quantities like gas saturation, reservoir thickness, or, more ambitiously, total hydrocarbon–in–place. This is another level of difficulty over and above trying to compute "maximum–likelihood" type inversion images, for several reasons. Firstly, the strong nonlinearity means the posterior distributions are rarely approximately quadratic ("bowl shaped") in parameter space, so local–linearization approaches to uncertainty mapping are not useful. This means sampling methods like Markov–Chain Monte Carlo, or parametric bootstrap type techniques must be used. These are computationally demanding, since *very* many forward models (MCMC) or many *inversions* (bootstrap) need to be computed. These methods have been successfully demonstrated for 1D problems (Gunning et al., 2010; Chen et al., 2007),

A second challenge is that uncertainty estimates are dependent on the noise model to *leading order*. Loosely speaking, simple maximum–likelihood parameter estimates, in inverse problems that are not too nonlinear or degenerate, are roughly independent of estimated noise levels. Statisticians speak of such parameters as being robustly estimated, or "insensitive to noise mis–specification". But the *uncertainty* of model parameters is always dependent on the noise model at leading order, even in kind problems. In hard problems like CSEM, the nonlinearity and near–degeneracy makes parameter uncertainties rather sensitive to the choice of the noise model.

A pervasive difficulty is that forward models in geophysics virtually always make assumptions which render faster computation possible, and errors resulting from these assumptions are expected to be absorbed in the "effective" noise model. Examples in the CSEM context might be use of isotropic resistivities, neglect of bathymetry, or approximation of 3D responses by 2D or 1D models. These kinds of modelling approximations usually lead to strongly correlated modelling errors, and if the modelling error dominates the overall instrumental and environmental noise, the overall "effective" noise process can then be expected to be fairly spatially correlated. The challenge is to provide estimates of parameter uncertainty that are useful in decision making, that are not hopelessly corrupted by inadequacies in noise modelling.

We illustrate several approaches to noise modelling which address these issues, using 1D isotropic CSEM inversion as the basic engine from which inferences are to be made.

Theory

Bayesian maximum likelihood estimation is typically concerned with locating optima of objective functions (log–posterior densities) which are a combination of data misfits and prior distributions of model parameters (Sambridge et al., 2006). Prior beliefs about layer resistivities are assembled from regional geology, depositional considerations and upscaling characteristics. They are usually right skewed distributions (achieved by modelling log–resistivities), allowing smaller probabilities of resistive rocks, and truncated at the left by resistivity values around seawater. Thus, with n data $\mathbf{d}$, a model $\mathbf{m}$ in p parameters, forward CSEM model $\mathbf{F}(\mathbf{m})$, and Gaussian noise $N(0, C_d)$, the joint Bayesian posterior (likelihood $L(\mathbf{d}|\mathbf{m}) \times$ Gaussian prior $p(\mathbf{m}) = N(\mathbf{m}_p, C_p)$) is maximised at the optimum of

$$\chi^2_{\text{Bayes}} = (\mathbf{d} - F(\mathbf{m}))^T C_d^{-1}(\mathbf{d} - F(\mathbf{m})) + (\mathbf{m} - \mathbf{m}_p)^T C_p^{-1}(\mathbf{m} - \mathbf{m}_p) \tag{1}$$

Noise levels and correlations are contained in the "effective noise" covariance matrix C_d.

The usual (and reasonably demanding) approach to parameter uncertainty in Bayesian contexts is to find all local minima of the log–posterior (1), estimate some uncertainty scales (i.e. accumulate local covariance matrices at each local minima), and use these as the basis for a MCMC or bootstrapping approach to computing parameter uncertainty. If the noise levels are estimated correctly, one expects the *most likely* solution to have $\chi^2_{\text{Bayes,min}} \approx n$, and the suite of models representing the inversion "uncertainty" to cause fluctuations in χ^2_{Bayes} given approximately by a χ^2_p distribution (width $\approx \sqrt{2p}$) offset to the right by $\chi^2_{\text{Bayes,min}}$. An example, for independent Gaussian noise (C_d diagonal) is shown in fig. 1.

With independent Gaussian noise, large data sets, and parsimonious models ($p \ll n$) the fluctuation in allowable misfit can be modestly tight, so the range of possible models estimated by this procedure can be too narrow if the noise model contains correlated components, and we proceed by approximating C_d only by its diagonal (or an estimate thereof). With certain kinds of correlated measurements, posterior parameter estimates should not central limit (i.e. enjoy $\sim 1/\sqrt{n}$ precision improvements) as the data set increases. There is a clear and present danger of overly focused inference if the noise is not independent.

There are at least three possible approaches to parameter uncertainty estimation when error processes are systematic or strongly correlated:

(1) Reduce the *effective* number of data to n_{eff}. Since correlated residuals contain partially redundant information, the number of "true" effective independent measurements may be very much less than the raw data count. Typically one either suppresses data, or uses all data with (independent, Gaussian) error bars inflated by $\sqrt{n/n_{\text{eff}}}$. Based on toy linear problems where the noise is a sum of random and systematic components, one can make good arguments that n_{eff} ought to be set as

$$n_{\text{eff}} \approx \min\left(\frac{\text{total noise power}}{\text{systematic noise power}}, n\right) \tag{2}$$

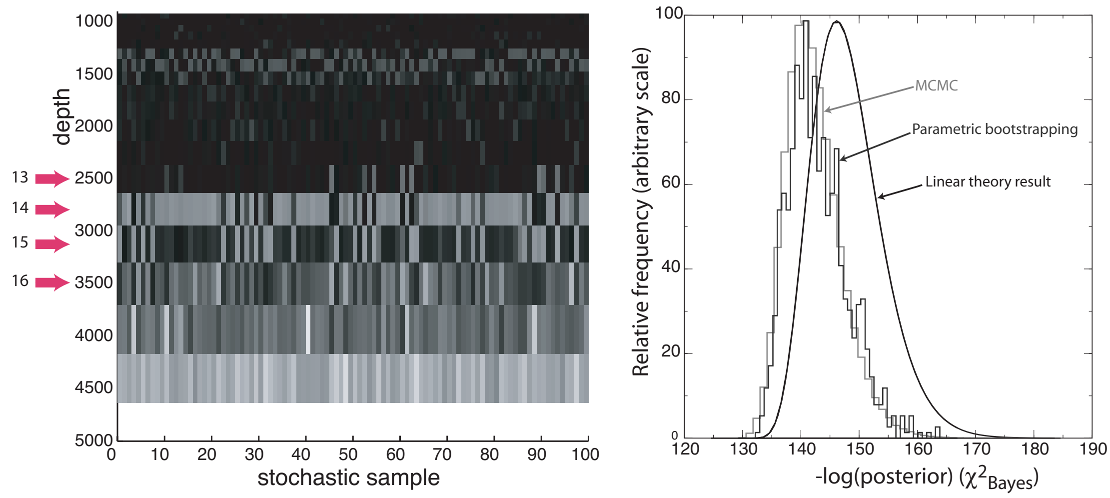

Figure 1 *Left: typical 1D stochastic models drawn from Bayesian posterior using MCMC or bootstrapping, showing strong resistor location correlations around marked layers. For the chosen layering, these model samples represent the subsurface uncertainty conditional upon the observed CSEM data (not shown). Right: sampling fluctuations in χ^2_{Bayes} relative to known result from linear regression theory for independent Gaussian noise.*

This can be a *very* modest number when the noise from systematics is strong.

(2) Explicitly try to extract the systematic components as modelled parameters that are unconnected to the earth model, but leave the random component reasonably equally plausible under a wider range of earth models. Typically, all one does is augment the usual CSEM inversion machinery with a forward model like

$$\mathbf{d} = \mathbf{F}(\mathbf{m}) + \underbrace{C_d^{1/2} X}_{\equiv X_N} .\mathbf{m}_n + \boldsymbol{\epsilon}$$

where X_N is a linear regression design matrix for systematic noise components $\mathbf{m}_n$. The prior distribution on $\mathbf{m}_n$ contains a parameter σ_{sys} expressing the relative power of systematic noise to the noise level of the uncorrelated component $\boldsymbol{\epsilon}$.

(3) Try to absorb correlated residuals into a fully correlated noise structure (i.e. dense matrix C_d). Use of e.g. Wishart priors has been successful (Leonard and Hsu, 1992) in the context of linear regression with exotic noise, but this is still probably too demanding to put inside a CSEM inversion loop.

The first two approaches have the advantage that they can be easily incorporated into standard optimization (Gauss Newton, Marquardt) and sampling algorithms (MCMC, bootstrap) that assume assume fixed, diagonal forms for the noise.

Example

Here we illustrate the problem of estimating a reservoir layer thickness in a problem with systematic processing errors that infect the inline $|E|$ field with a component that correlates $\log |E|$ almost linearly with offset, rather as it might in a problem with modelling errors from weak 3D or anisotropic effects. Since the thickness is already hard to disentangle from the RTP, this kind of error has potential to seriously bias a thickness estimate. In this example, the systematic error in the data set is enough to cause seriously biased estimates if posterior uncertainties are computed using the full data set with an unadjusted, independent noise model only. Figure 2 shows how the inferred uncertainty in thickness widens substantially if modelling of the systematic noise is introduced, using either of the first two ideas above. Data are generated from a 250m resistor buried 850m below mudline (in 1km water).

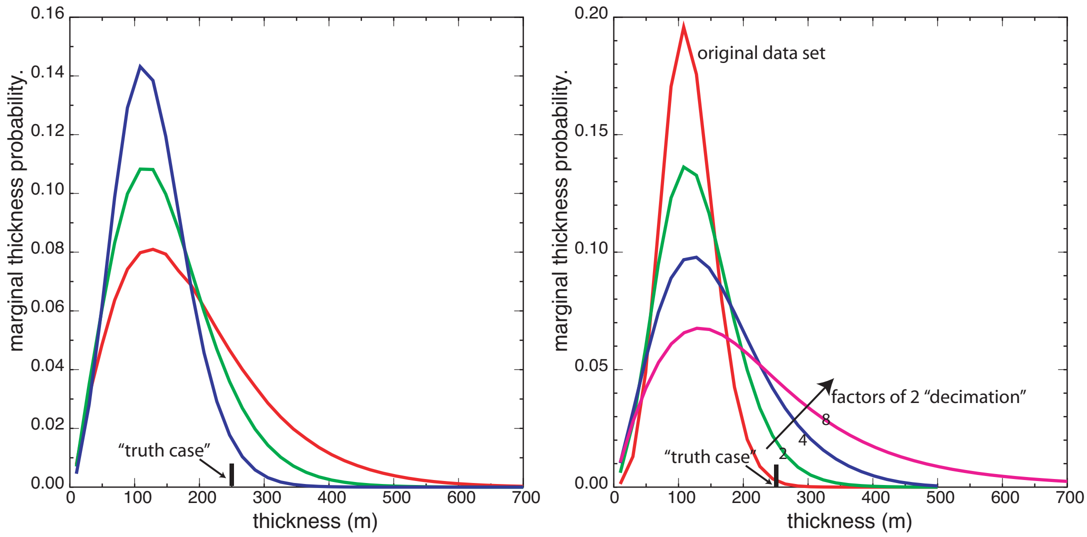

Figure 2 *Left: Inferred thickness marginal distribution of resistivity target, as a function of relative tolerance of systematic errors in inline* $|E|$ *data, using quadratic (*$N_p = 2$*) systematics with offset. Curves are (i) red,* $\sigma_{sys} = 2$*, (ii) green,* $\sigma_{sys} = 1$*, (iii) blue,* $\sigma_{sys} = 0.5$*. Right: very similar effects are obtained by data set reduction by factors of 2 (error bar inflation by* $\sqrt{n/n_{\mathit{eff}}}$ *as per option 1.*

Conclusions

Mis–characterisation of the noise has significant potential to disrupt reservoir parameter uncertainties in Bayesian inversion approaches. Since practical inversions have to use approximate forward models of varying degree, modelling error gets absorbed into the overall error process, and may have a strong systematic component. Such systematic components in the noise ought to be reflected in broader inversion uncertainties. Estimations of posterior parameter uncertainties must take such correlated error processes into account. We have presented two relatively simple ideas to manage the effect, both of which are easily incorporated into standard optimization and sampling schemes.

References

Constable, S. C., Parker, R. L. and Constable, C. G. [1987] Occam's inversion - a practical algorithm for generating smooth models from electromagnetic sounding data. Geophysics 52, 289–300.

Gunning, J., Glinsky, M. and Hedditch, J. [2010] Resolution and uncertainty in 1D CSEM inversion: a Bayesian approach and open-source implementation. *In review*, Geophysics.

Chen, J., G. M. Hoversten, D. Vasco, Y. Rubin, and Z. Hou, 2007, A Bayesian model for gas saturation estimation using marine seismic ava and csem data: Geophysics, **72**, WA85–WA95.

Sambridge, M., Gallagher, K., Jackson, A. and Rickwood, P. [2006] Trans-dimensional inverse problems, model comparison and the evidence. GEOPHYSICAL JOURNAL INTERNATIONAL 167(2), 528–542.

Loseth, L., 2007, Modelling of controlled source electromagnetic data: PhD thesis, Department of Physics, NTNU.

Leonard, T. and J. Hsu, 1992, Bayesian-inference for a covariance-matrix: Annals of statistics, **20**, 1669–1696.