

Introduction

Inversion for elastic properties and joint litho–fluid indicator variables has become a common industry standard, and has multiple applications in both traditional O&G applications and energy–transition applications such as monitoring plume movement in CCUS. Many authors have implemented joint litho–fluid inversion frameworks, with some similarity in the general Bayesian framework, e.g. Rimstad (2012), Kolbjørnsen et al. (2016), Gunning and Sams (2018) and Gunning (2018). The introduction of latent indicator variables is a highly versatile and useful step, as it not only renders explicit indicators for economically valuable fluid types, but also makes possible the formulation of rock and fluid–dependent models in the forward physics of the inversion, and imposition of physical and temporal consistency in these models. Facies indicators also help a great deal with the missing 2–10Hz bandwidth problem of typical streamer data: in classical inversion, this is surrogated by smooth deterministic background trends which are a compromise among the loading trends of the facies present, but inversions using these models typically do not produce the blocky facies character seen in well logs. Using joint inference for litho–fluid variables with lithology–dependent elastic prior distributions in a Bayesian framework is very effective at inducing this desirable blocky character.

Indicators are very useful for introducing additional kinds of constraints in the inversion, for example imposing density–ordering rules for fluids in the vertical sense. The latent indicator idea can also be extended to time–lapse inversion by the concept of N–vintage ‘litho–fluid–tuples’ enumerating the possible multivintage litho–fluid states that will be feasible under time evolution (Gunning, 2025). For example, a 3–facies {S=Shale, B=Brine–sandstone, G=Gas–sandstone} 2–vintage CO₂ injection case might have allowable ‘tuples’ {SS,BB,BG} at each voxel. This state–space grows only modestly with the number of vintages, and the tuples allow additional temporal transition rules for fluid movement over vintages.

In standard and simple forms of Bayesian litho–fluid AVO inversion (Gunning and Sams, 2018), the forward model does not depend on the facies variables, and the facies dependence only occurs in the conditional–prior part of the Bayesian posterior, which usually has a computationally agreeable form. This makes it possible to approach the inversion problem using the expectation–maximisation (EM) algorithm, and the resulting algorithm is like a spatio–temporal extension of the classical EM treatment of mixture distributions. Like all EM algorithms however, the convergence can be sensitive to poor starting guesses, so care is required to start the algorithm in the basin of attraction of the global minimum.

The EM algorithm has many charms as an inferential approach, but for a large–scale application like seismic inversion, its implementation depends heavily on finding scalable algorithms for the expectation ‘E’–step (for facies) and maximisation ‘M’–step (for elastic properties). For Gaussian likelihoods, the M–step comes out to be a large nonlinear least–squares problem, and for the ‘simple’ AVO problem with a Markov random field prior on the facies variables, the E–step is tractable using loopy belief–propagation. However, if the forward model now introduces explicit facies–dependencies, the E–step is now much more difficult to compute, and usual methods like mean–field theory are typically rather fragile if the facies–dependence in these models is strong.

Time–lapse joint litho–fluid inversion requires an appreciable escalation of modelling complexity compared to single vintage. While inversion of 2–vintage baseline/monitor pairs can sometimes be finessed with formulations for difference variables, beyond 2 vintages these tricks become unwieldy. A particularly difficult aspect of multivintage models is that the vintage–appropriate elastic model depends on the discrete labels. For example, if stress effects in the overburden are negligible, the *base*–vintage shale elastic parameters should explain all vintage data. But the fluid labels in the reservoir will require separate elastic variables for a sandstone voxel containing *base*–vintage brine that then transitions to *monitor*–vintage CO₂. The forward model is then a strong explicit function of the label state.

Handling the indicator dependence in these more complex formulations can get very challenging in the EM algorithm. Fortunately, it turns out to be useful to interchange the E and M–steps presented in Gunning and Sams (2018), and the resulting M–step for facies variables can be relaxed to an optimisation problem for soft indicators subject to a compositional constraint. This amounts to effectively performing

a joint optimisation in the elastic/relaxed–indicator space. Gradient descent methods in this space are much easier to derive than expectations of strongly coupled discrete variables. By use of the same Jensen inequality underlying the EM algorithm, a further relaxation of this problem can be derived which is fully convex for the prior part. The merit of this prior formulation is that then the relaxed–indicators usually end up close to vertices of the compositional simplex (indicating facies), and the solution is start–point independent. An additional, enticing possibility is that the facies abundances in the prior model can usually be adjusted so that a possible vertex appears for each lithofluid indicator of interest, so the usual combinatorial challenges of discrete variables have been sidestepped by the use of vertices of a convex polytope acting as facies proxies in a convex and continuous formulation. In this paper, we sketch the formalism for the relaxed inversion embedding the vintage–consistency idea above, with a brief sketch of the convex relaxation which sidesteps the multimodality. An example inversion on a multivintage data set from the North sea is shown illustrating the machinery.

Theory

The model is an extension and generalisation of that in Gunning and Sams (2018). For a single vintage we use elastic properties $\mathbf{m} = \{\mathbf{v}_p, \mathbf{v}_s, \boldsymbol{\rho}\}$ and lithofluid/facies labels $\mathcal{F}_i \in \{1 \dots N_f\}$ for all voxels $i = 1 \dots N_p$ in the volume. The inversion uses a Bayesian hierarchical prior model of form $P(\mathbf{m}|\mathcal{F})P(\mathcal{F})$, where the seismic “data” $\mathbf{y}$ are true–amplitude imaged angle–stacks of seismic reflectivities. The forward model $\mathbf{f}(\mathbf{m})$ computes wavelet convolutions of reflectivities for all stack angles, and possibly time–to–depth conversions if the model is in depth and model–dependent kinematics are necessary. This has associated likelihood $L(\mathbf{y}|\mathbf{m}) = L(\mathbf{y}|\mathbf{f}(\mathbf{m}))$, and we use a Gaussian model for the noise, $L(\mathbf{y}|\mathbf{m}) = N(\mathbf{y} - \mathbf{f}(\mathbf{m}), C_d)$, where the wavelets in $\mathbf{f}(\mathbf{m})$ and noise covariance C_d are estimated from well ties. The joint posterior distribution is $P(\mathbf{m}, \mathcal{F}|\mathbf{y}) \sim L(\mathbf{y}|\mathbf{m})P(\mathbf{m}|\mathcal{F})P(\mathcal{F})$. The conditional prior distribution $P(\mathbf{m}|\mathcal{F})$ is a Gaussian model developed from regional well–log analysis and expressing compaction, correlations, and typical prior uncertainties. We model $P(\mathbf{m}|\mathcal{F})$ as a product of $P(m_i = \{v_{p,i}, v_{s,i}, \rho_i\}) \sim N(\bar{m}_i(\mathcal{F}_i), C_{p,i}(\mathcal{F}_i))$ over voxels. The discrete prior spatial distribution is modelled as a Markov random field $P(\mathcal{F})$, with pairwise neighbours denoted $(j \sim i)$, of form $P(\mathcal{F}) \sim \exp(-(\sum_i e_{\mathcal{F}_i} + E_{\text{MRF,sp}}(\mathcal{F})))$, where $E_{\text{MRF,sp}}(\mathcal{F}) = \sum_{i,j \sim i} 2\beta^{(D_{ij})} I(\mathcal{F}_i \neq \mathcal{F}_j)$ is the ‘MRF spatial interaction’ term. Spatial continuity is promoted by the $\beta^{(D_{ij})}$ parameters in direction D_{ij} , and the $e_{\mathcal{F}_i}$ parameters are tuned to generate the desired spatially variable marginal prior abundance of each facies. For the multivintage case, the elastic model $\mathbf{m}$ is extended to contain elastic properties for each vintage, but the forward model $\mathbf{f}(\mathbf{m}, \mathcal{F})$ switches on the correct $\mathcal{F}$ –dependent vintage model as required in the multivintage forward model.

If we seek a point estimate from this posterior distribution, the maximum a posteriori inversion problem is then $\{\hat{\mathbf{m}}, \hat{\mathcal{F}}\} = \text{argmin}_{\mathbf{m}, \mathcal{F}} \{-\log(P(\mathbf{y}|\mathbf{m}, \mathcal{F})P(\mathbf{m}|\mathcal{F})P(\mathcal{F}))\}$. In the EM algorithm the optimisation alternates the $\mathbf{m}$ and $\mathcal{F}$ spaces where the $\mathcal{F}$ variables end up appearing via their *memberships* $z_{il} = \langle \mathcal{F}_{il} \rangle$ for voxel i , label l in the usual way. In this formulation, this turns out to be very closely equivalent to treating the z_{il} variables as a continuous relaxation of the indicators, subject to the compositional constraint $z_{il} \in C_z \equiv \{z_{il}, z_{il} \in [0, 1], \sum_l z_{il} = 1\}$. With suitable linear z –equivalents of the $\mathcal{F}$ dependence in the MRF and forward–model ‘switches’, the joint optimisation problem can be relaxed to the joint optimisation

$$\begin{aligned} \{\hat{\mathbf{m}}, \hat{\mathbf{z}}\} = \arg \min_{\mathbf{m}, \mathbf{z} \in C_z} & (\mathbf{y} - \mathbf{f}(\mathbf{m}, \mathbf{z}))^T C_d^{-1} (\mathbf{y} - \mathbf{f}(\mathbf{m}, \mathbf{z})) \\ & + \sum_{i=1}^{N_p} \sum_l z_{il} [(\mathbf{m}_i - \bar{\mathbf{m}}_{il})^T C_{p,il}^{-1} (\mathbf{m}_i - \bar{\mathbf{m}}_{il}) + \log |C_{p,il}| + e_{i,l}] + \sum_{i,j \sim i} \sum_l \beta^{(D_{ij})} \|z_{il} - z_{jl}\|_1. \end{aligned}$$

The principal advantage of this relaxation compared to the EM algorithm here is that the joint gradients of the likelihood piece w.r.t $\mathbf{z}$ are much easier to compute than expectations w.r.t $\mathcal{F}$, so optimisation can be performed with standard linearly–constrained gradient methods in the joint $\{\mathbf{m}, \mathbf{z}\}$ space. In this formulation, in the second term, the nonconvexity from the $\mathbf{m}, \mathbf{z}$ cross–products is very evident. This difficulty can be eased by replacing this term by its tight lower bound obtained using Jensen’s inequality, which is very simple in the common case that the within–facies prior covariances C_p are decently approximated by some average $\bar{C}_p$. This yields an convex–envelope prior replacement piece in

the objective of form

$$\tilde{E}_p(\mathbf{m}, \mathbf{z}) = \sum_{i=1}^{N_p} \left((\mathbf{m}_i - \sum_{l=1}^{N_f} z_{il} \tilde{\mathbf{m}}_{il})^T \bar{\mathbf{C}}_{p,i}^{-1} (\mathbf{m}_i - \sum_{l=1}^{N_f} z_{il} \tilde{\mathbf{m}}_{il}) + \sum_{l=1}^{N_f} z_{il} e_{i,l} \right)$$

This form is very useful in that it removes any nonconvexity from the prior, and moreover, it can be shown that by suitable choice of the abundance-related parameters $e_{i,l}$, the compositional constraint can produce stable vertices at each pure facies which behave as optimal points if the gradients from the likelihood piece drive the solution there. These vertices in $\{\mathbf{m}, \mathbf{z}\}$ space act as surrogates for facies indicators, and this space confers all the advantages of continuous, differentiable variables. Figure 1 shows a simple vignette of this idea with $\mathbf{m} = \{Z\}$, the acoustic impedance of one voxel. However, the requirement for the forward model to *switch* on the correct $\mathcal{F}$ -dependent vintage model as a function of the indicator state in $\mathbf{z}$ still leaves some nonconvex pieces in the likelihood objective.




Figure 1: Gaussian mixture for 4 facies {shale=brown, brine=blue, oil=red, gas=green}, with mixture proportions $w_l = \{.33, .3, .27, .1\}$ (roughly proportion to $\exp(-e_l/2)$). The facies mean impedance is shown by vertical lines. The overall mixture distribution is shown as the dotted cyan curve. The convex-envelope energy $\tilde{E}_p(Z) \equiv \min_{z \in \mathcal{C}_z} \tilde{E}_p(Z, z)$ is the solid cyan curve. This minimisation *induces* maximum-likelihood (ML) proportions z of each facies. As the impedance Z varies, the dashed brown/blue/red/green lines show the *induced* ML proportions of each component. *Effective* proportions might be estimated from the relative areas of the basins of attraction of each component.

In the following example, we show how some of this vintage-consistency machinery is useful in the inversion of multivintage field data. The prior proportions profiles expose relatively few facies at each depth, so the convex-hull vertex trick of the preceding paragraph is not necessary, but the idea is a powerful one, and should be applicable to many problems.

Field Example. Figure 2 shows a North Sea oilfield example with 3 vintages. The inversion uses a mix of 4 non-fluid facies (shales, tuffaceous claystones) and brine, oil in the turbidite sand unit. Here we highlight the inverted fluid movements over a small cross section for clarity, where brine appears to encroach over oil under production. At convergence, when the indicators promote any of the 4 non-fluid facies, these elastic properties are constant over vintages. The fluid facies expose new elastic property degrees of freedom where they change with vintage.

Conclusions

We have developed a multivintage AVO joint litho-fluid inversion using the concept of litho-fluid-tuples to represent the evolving state of the fluids in a voxellised model. This framework enables the introduction of coupling interactions in the lithofluid state to promote spatial smoothness, and enforce transition rules in time and space. We derive a close relationship between an expectation-maximisation formalism for the inversion and a direct relaxation using litho-fluid-composition variables, and show that the latter is more flexible for imposing vintage continuity rules on elastic variables in regions of space that ought to be constant in time. We have illustrated the vintage continuity aspects of the scheme using 3-vintage data from the North Sea.

Acknowledgements. CSIRO thanks Ikon Science for financial support.




Figure 2: Field case

References

- Gunning, J. [2025] Joint elastic/facies convex-cover AVO inversion with versatile and expressive latent variables. **2025**(1), 1–5, ISSN 2214-4609, doi:<https://doi.org/10.3997/2214-4609.202510359>.
- Gunning, J. [2018] A comparison of layer-based and sample-based inversions for time-lapse AVO problems. **2018**(1), 1–5, ISSN 2214-4609, doi:<https://doi.org/10.3997/2214-4609.201801329>.
- Gunning, J. and Sams, M. [2018] Joint facies and rock properties Bayesian amplitude-versus-offset inversion using Markov random fields. *Geophysical Prospecting*, **66**(5), 904–919, doi:<https://doi.org/10.1111/1365-2478.12625>.
- Kolbjørnsen, O. et al. [2016] Bayesian AVO inversion to rock properties using a local neighborhood in a spatial prior model. *The Leading Edge*, **35**(5), 431–436, doi:10.1190/tle35050431.1.
- Rimstad, K. [2012] *Spatial mixture modeling based on latent random fields applied to seismic inversion*. Ph.D. thesis, NTNU.