

# Bayesian AVO inversion to rock properties using a local neighborhood in a spatial prior model

Odd Kolbjørnsen<sup>1</sup>, Arild Buland<sup>2</sup>, Ragnar Hauge<sup>3</sup>, Per Røe<sup>3</sup>, Martin Jullum<sup>3</sup>, Richard William Metcalfe<sup>2</sup>, and Øyvind Skjæveland<sup>2</sup>

## Abstract

The spatial structure of the subsurface is an important factor when interpreting seismic data. The Bayesian methodology is a valuable tool for integrating these spatial relations in the inversion process as it merges the information together and assesses the uncertainty of the model. In the everyday use of the Bayesian methodology, however, the computational cost is a challenge. We describe a new approach that utilizes a local neighborhood to include the spatial constraints and assess the uncertainties in the inversion using fast and parallelizable computations. The approach is applicable for both discrete lithology-fluid prediction and estimation of rock properties, such as porosity and saturation.

## Introduction

Mapping rock properties such as lithology type, pore fluid, porosity, and saturation from seismic prestack amplitudes is a nonunique and unstable inverse problem. Luckily, the seismic data is not the only source of information about the subsurface. Geologic zonation, petrophysical properties, and rock-physics models provide additional insight, which helps limit the options. The general setting of AVO inversion is illustrated in Figure 1. The top arrows describe standard forward modeling, i.e., the rock-physics model giving the link between rock properties and elastic parameters and the geophysical model, which then gives the AVO gathers. The lower set of arrows in Figure 1 illustrates the AVO inversion. To solve the rock properties, it is common to use a two-step approach. First, an elastic inversion is performed, and then the inverted elastic properties are mapped into rock properties by evaluating one location at a time. The problem with this approach is that it loses track of the spatial structure when interpreting the rock properties, e.g., placing samples of oil sand directly above brine sand.

The Bayesian methodology is an excellent framework for merging multiple sources of information to gain a unified interpretation. In the Bayesian approach, all general knowledge is quantified in a joint, spatially coupled prior distribution. This prior distribution describes the multiparameter relations between different rock properties and elastic parameters, as well as the spatial structure of the problem. The spatial structure defines features such as the consistency in geologic sequences, ordering of fluids, the extent of geologic features, and the degree of continuity in geologic events. The probability model is constructed using the top set of arrows in Figure 1. Initially, a prior distribution for the rock properties is derived through petrophysical analysis of nearby or analog well logs. The rock physics is defined either based on a theoretical model or by empirical relations. The top right arrow illustrates the geophysical model,

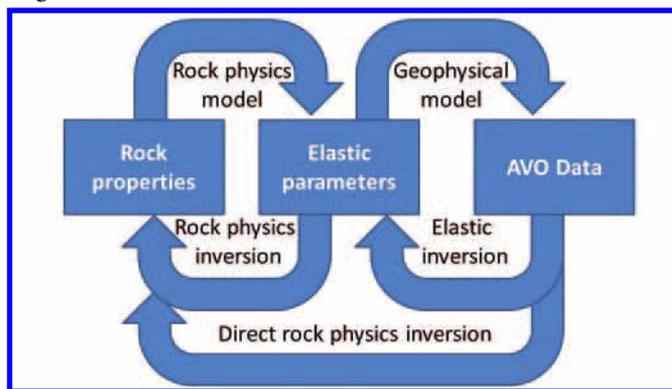
which provides the link to AVO data. In the Bayesian framework, this link is denoted as the likelihood, and the relation is typically derived in a well-tie analysis.

A great advantage of the Bayesian approach is that even though the probability model is built according to the arrows on top, it is possible to swap point of view and assess the probability distribution along the arrows below the boxes to perform elastic inversion (e.g., Buland and Omre, 2003), rock-physics inversion (e.g., Avseth et al., 2005), or direct inversion to rock properties (e.g., Jullum and Kolbjørnsen, 2016). Any Bayesian approach is, by definition, a one-step approach since it assesses the conditional distribution of the target parameter given the data. In the Bayesian formalism, these conditional distributions are denoted as the posterior distributions and represent our updated knowledge of the subsurface model. Computing the full spatially coupled posterior distribution for large-scale inverse problems is an overwhelming task, and the standard Bayesian approach of Markov-chain Monte Carlo (MCMC) simulations becomes too time-consuming.

To resolve this, one option is to evaluate the Bayesian uncertainties pointwise by considering one location at the time (Buland et al., 2008). This approach enables a consistent integration of prior information and an assessment of the uncertainties in both the geo- and rock-physical model. However, due to the pointwise assessment, the approach loses track of the spatial structure and experiences some of the same weaknesses as the two-step approach.

Another option is to compute the maximum posterior solution, which provides the optimal merge between the different sources of information (see, e.g., Kemper and Gunning, 2014). The weakness of this approach is that it does not address the question of nonuniqueness in the inversion. There are an uncountable number of alternative earth models that are almost as plausible as the optimum; some of these might have a completely different set of properties. For instance, a thin layer of oil-saturated sand can give a similar

Figure 1. Framework for AVO inversion.



<sup>1</sup>Lundin Norway.

<sup>2</sup>Statoil ASA.

<sup>3</sup>Norsk Regnesentral.

signature as a slightly thicker layer of brine-saturated sand, or a reservoir filled with low-saturation gas might give an AVO signature very similar to high saturation. In such cases, it is of great value for the risk assessment to use a method that fully acknowledges the ambiguity in the solution to the inverse problem.

We propose a new approach that addresses the issue of nonuniqueness at the same time as it preserves a spatial structure and bypasses the problems of high-dimensional posterior distributions. Rather than trying to express the full joint posterior distribution, we focus on a small subset of properties and compute the posterior distribution for these. This subset is typically a set of rock parameters in nearby cells. We might ask for the probability of oil sand in one particular subsurface location; the joint probability of oil sand in two neighboring cells; the porosity and saturation in a subsurface location; or the porosity and saturation in neighboring cells. For such focused problems, it is sometimes possible to obtain adequate approximations of the posterior distribution by solving a small inverse problem on a local neighborhood. Thus, rather than handling one large coupled system, one might obtain satisfactory answers by solving many small inverse problems independently. Below, we discuss this local-neighborhood approach by showing the details in a synthetic example. We also show the results of two cases where this approach has been used on real data: to predict sand bodies at the Statfjord field and for mapping the spatial distribution of CO<sub>2</sub> in the Sleipner injection project.

### The local-neighborhood approach for integrating spatial structure

In the local-neighborhood approach, we extract the data in a region around the target location and evaluate how likely it is that these data represent a given a scenario for the property in the target location. The formal computations for lithology-fluid prediction are given in Kolbjørnsen et al. (2008) and, for continuous properties such as porosity and saturation, in Jullum and Kolbjørnsen (2016). The key to the local-neighborhood approach is that the analysis includes details of the spatial relation to “nearby” locations. To follow, we outline the computations for the case of lithology and fluid prediction.

To illustrate the methodology, consider the task of lithology-fluid prediction in a simple case where the options are shale (S), brine sand (B), and oil sand (O). A spatial model is constructed on a grid with a vertical sampling interval of 4 ms. The sequences of lithology-fluid classes are defined as a Markov chain (Larsen et al., 2006; Kemper and Gunning, 2014) using the transition probabilities listed in Table 1. To understand the transition probability, consider the first row in the table, i.e., shale. From this row, we can read that if a given cell in the grid is shale, then the cell directly below will be shale in 95 out of 100 cases. The chances for brine sand and oil sand are 3% and 2%, respectively. The following two rows state the probabilities for the cases with brine sand and oil sand on top.

The transition probabilities imply a set of prior probabilities of the lithology-fluid classes, which are given in Table 2. The mapping from rock properties to elastic parameters, shown as the top left arrow in Figure 1, is represented by the probability distributions of the elastic properties for each lithology-fluid class. For simplicity, we only give the properties for AI in Table 2, but

the full spatial prior distribution for elastic parameters requires the pointwise relation between AI,  $V_p/V_s$  ratio, and density, as well as the spatial dependency of these properties for neighboring cells with the same lithology fluid class.

To include the spatial context in the analysis, consider the probability of the lithology-fluid class in two neighboring cells. In this simple case, there are a total of eight scenarios: S-S, B-S, O-S, S-B, B-B, O-B, S-O, and O-O. The option B-O is omitted in the list of scenarios as this corresponds to brine directly above oil, which is considered unphysical in a well-segregated reservoir. This can also be seen in the transition probabilities in Table 1 where the brine-sand-to-oil-sand transition has probability zero.

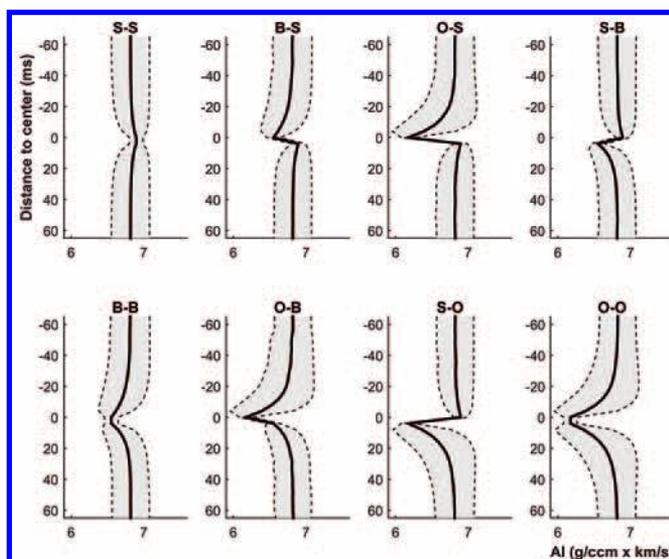
If we know that there is a specific scenario in a neighborhood, then the spatial dependency causes this to influence the distribution of elastic properties around this event. Figure 2 shows how the mean and variability of acoustic impedance is influenced

**Table 1.** Transition probabilities in example. Transitions from lithology fluid class on top (row) to the lithology-fluid class below (column).

	Shale	Brine sand	Oil sand
Shale (S)	0.95	0.03	0.02
Brine sand (B)	0.30	0.70	0.00
Oil sand (O)	0.20	0.10	0.70

**Table 2.** Prior probability distribution for acoustic impedance in example.

	Shale	Brine sand	Oil sand
Probability for lithology fluid class	0.85	0.10	0.05
Acoustic impedance, mean (g/ccm × km/s)	6.9	6.6	6.2
AI P10-P90 range (g/ccm × km/s)	6.8–7.0	6.5–6.7	6.1–6.3

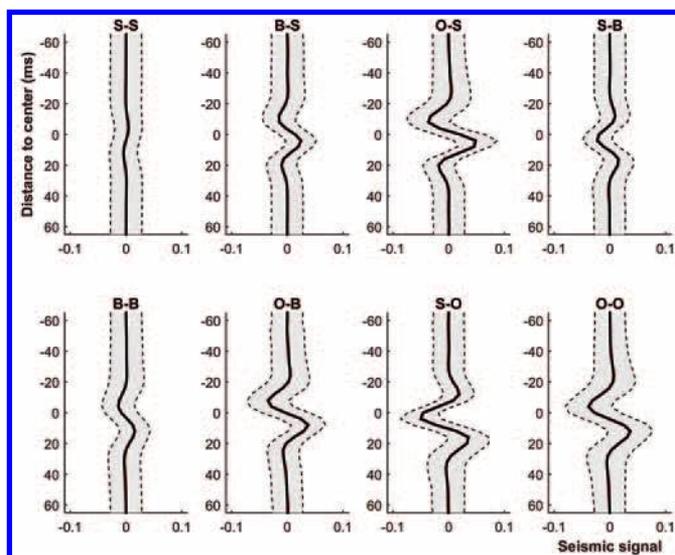


**Figure 2.** The mean and variability of AI in a region around a lithology-fluid event. The mean response (solid line) and P10-P90 uncertainty limits (dotted line) for the eight scenarios.

above and below the two specified cells for the eight possible scenarios. Close to the event, the eight scenarios are very different; farther away, the influence of the scenario is gradually reduced. At the top and bottom, there is no effect. The S-S event has the least dramatic effect, whereas the S-O event shows a strong contrast. A low level of acoustic impedance is not expected to be kept for a long period when going away from oil sand or brine sand. This is an effect of the spatial structure, i.e., the transition probabilities in Table 1. The probability of 0.7 for an oil-sand-to-oil-sand transition corresponds to an average thickness of the oil sands of about 13 ms. This is best seen in the distribution for acoustic impedance for the O-O event, in which the level of the acoustic impedance jumps rapidly back to the base level on both sides.

For each spatial model corresponding to a lithology-fluid scenario in Figure 2, we can compute the expected seismic response and its uncertainty. Figure 3 displays these spatial models describing the seismic signatures of all the lithology-fluid scenarios.

The likelihood of a local lithology-fluid combination is obtained by comparing the seismic response extracted in the neighborhood to the seismic signatures in Figure 3. Scenarios that display a signature similar to the observed response will get an increased probability, whereas a bad match decreases the probability. The detailed expressions are given in the references above. Through the local-neighborhood approach, we have obtained the probability of the eight different scenarios for neighboring cells in one location. If we are interested in the posterior probability of oil sand in the top cell, this can be computed by summing the probabilities for events with oil sand in the top cell, i.e., O-S, O-O, and O-B. By shifting the data window, we can compute the probability for these eight scenarios in each location sliding down the trace and obtain the probability for oil sand in each location. It is also possible to compute the oil probability in the lower cell in each event to get an alternative prediction. However, the strength of the approach is that when the posterior probability of



**Figure 3.** Seismic signature of transitions. The mean response (solid line) and P10-P90 uncertainty limits (dotted line) for all pairs of lithology-fluid combinations.

the joint event is computed, it also defines the posterior transition probabilities, which can be used to build a spatial model. The constructed spatial model is not the true spatial posterior distribution but rather a first-order approximation. To improve the approximation of the posterior, it is possible to use higher-order Markov chains. The probability of higher-order transitions can be found by increasing the size of the lithology-fluid event; i.e., rather than considering events of the type S-S, one would consider events of type S-S-S-S.

In synthetic examples, it is always possible to list all alternative scenarios; in real situations, the case is more complex as it is possible that neither of the listed scenarios provides an adequate fit. To investigate this, it is common to introduce an additional scenario, which is denoted as an undefined event. The undefined event represents a base model, which is parsimonious and has less spatial structure than what is imposed in the Bayesian model. If this model is assigned a high probability in the posterior, this should be a warning sign that there are problems with the statistical model. Such an undefined rock can be modeled with independent elastic parameters with wide spread.

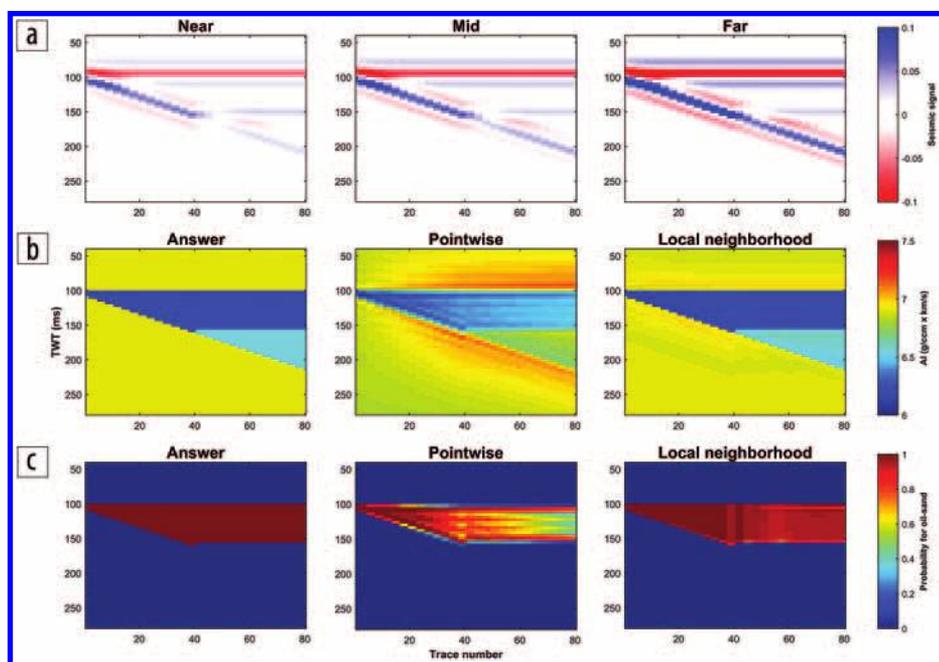
The local-neighborhood approach is applicable to any setting in which it is relevant to perform an AVO inversion to rock properties. Examples of problematic situations are when the lithologies are poorly understood in a region, lack of well tie gives too large an uncertainty of the scale of the seismic wavelet, or imperfections in the processing gives data quality that is too poor.

In principle, there is no limit to how many lithology-fluid classes this approach can handle. In fact, the approach has also been developed for continuous variables. In a generic implementation where the focus is on the discrete predictions, the limiting factor is the number of lithology configurations that are considered within a window. This might be equally dependent on the spatial structure as the number of lithology-fluid combinations. For instance, consider three lithologies within a window of length five. In the general case, there are 243 possible configurations, but if we impose a strict ordering, only 21 options remain.

The methodology has a potential to separate lithologies that are acoustically overlapping but have differences in the spatial dependencies. If two similar lithologies have clear differences in sequence stratigraphic ordering, this might be obvious. If the spatial structure only differs in the extent of the geologic features, then the separation is weaker. If two lithologies are identical in every aspect, the relative probability of these lithologies will not change from the prior to the posterior, but the probability of these will be scaled together toward the other alternatives.

### Synthetic example

The model setup in the previous section is tested for a double-wedge model where a sand wedge is embedded in a shale background. The top of the wedge is saturated with oil; the lower part is saturated with brine. Figure 4 shows the data and a comparison with the results of a pointwise inversion method (Buland et al., 2008). The differences are evident when comparing the acoustic impedances. Even in this case with a high signal-to-noise ratio, the pointwise approach creates a



**Figure 4.** Wedge example. (a) Seismic AVO data, near (left), mid (middle), far (right). (b) Acoustic impedance of the synthetic (left), the pointwise inversion (middle), and local-neighborhood inversion (right). (c) Oil-sand probability of synthetic (left), the pointwise inversion (middle), and local-neighborhood inversion (right).

**Table 3:** Comparison of prediction quality for the pointwise and the local-neighborhood approach.

	Average probability (%)			Percentage correct classification (%)		
	Prior	Standard	Local neighborhood	Prior	Standard	Local neighborhood
Shale	85	93	96	100	99	99
Brine-sand	10	39	55	0	42	61
Oil-sand	5	52	63	0	57	68

smoothing effect at the edge, and the level of the inversion is lost inside the thick part of the sand. Turning attention to the predicted oil-sand probabilities for the pointwise approach, we see that the probability of oil sand drops when the level of the acoustic impedance drops; thus, the flaws in the inverted acoustic impedance transfer to imperfections in the lithology-fluid prediction. Drawing on this observation, it would be tempting to say that the good match in acoustic impedance for the local-neighborhood approach provides the excellent prediction of the lithology-fluid class. This is, however, not a correct assessment. It is the spatial model of the lithology-fluid classes that enables the good match for the local-neighborhood approach; thus, the match for the acoustic impedance is caused by the good prediction of the lithology-fluid class.

To quantify the benefit of the spatial model in the presence of noise, we test the two methodologies on synthetic data generated from the prior distribution. The signal-to-noise ratio in the synthetic data is about three. Two measures are used to evaluate the results. The average probability filters out all cells of a given lithology-fluid class and computes the average probability of this given class in the

selected cells. The classification quality is obtained by counting how often a lithology-fluid class is predicted correctly when assigned to the most probable lithology-fluid class. The fit is summarized in Table 3. As seen from the table, the spatial model increases the quality of the average prediction of sand and the classification significantly.

### Case example of lithology prediction

The Statfjord gas and oil field, located in the North Sea at the border between the United Kingdom and Norway, has been in production since 1979. The water depth is about 150 m, and the reservoir is at a depth of 2500–3000 m in Middle Jurassic Brent deltaic and Late Triassic to Early Jurassic Statfjord fluvial sandstones. The reservoir level at the east flank of the Statfjord Field consists of several thin layers, many below the tuning thickness. This makes it difficult to achieve a detailed mapping of the reservoir sands using standard interpretation methodology. The vertical ordering of the zonation is well known, as are the rock properties of the lithologies in the region. Thus, it is of interest to include this information when mapping the reservoir sands.

A spatial model that constrains the ordering of the geologic events is built using a Markov process prior. A typical feature of this process is that transitions from an older to a younger zone are

prohibited, and the variability of the elastic parameters within each zone is well constrained. The case and the results are illustrated in Figure 5. The inversion of multiple angle stacks was performed using both a pointwise and a local-neighborhood approach. From the figure, it is clear that the local-neighborhood approach gives a clearer separation of the sands, whereas in the pointwise approach these sands are smoothed together. The inversion was performed before well B was drilled. One of the objectives with the inversion was to investigate the thickness of the shale between the two sands. If the shale were thick, it would be advantageous to drill the well with two sections. The local-neighborhood inversion predicted that the shale observed in well A would be much thinner in B, and that was confirmed by the well.

In Figure 5, the green well marker is the top of the shale, and the yellow marker is the base of the shale, showing excellent match between the inversion and the well result. The local neighborhood can also be used to estimate the confidence in the results with respect to the extent of the sand, i.e., the uncertainty of zone boundaries. This is important information when the results are brought forward to geomodels and uncertainty workflows.

## Case example of property prediction

The Sleipner field is located in the North Sea, about 250 km west of Stavanger. The field produces natural gas and condensate, which is rich in CO<sub>2</sub>. At the topside, CO<sub>2</sub> is removed from the produced gas and injected back into the Utsira formation. The injection started in 1996. The Utsira formation is located at a depth of about 1000 m and contains thin layers of shale within a background of high-porosity loose sand. When injected into the Utsira formation, the CO<sub>2</sub> is trapped beneath the thin layers of shale. To monitor the spatial distribution of CO<sub>2</sub>, it is desired to invert the CO<sub>2</sub> saturation from time-lapse seismic. Since the thin layers of high CO<sub>2</sub> saturation are below seismic resolution, there is a tradeoff between the thickness of the layers and the CO<sub>2</sub> saturation, which cannot be easily interpreted from a standard inversion. A key to understanding the spatial model for the saturation is the understanding of the CO<sub>2</sub> flow. Prior to CO<sub>2</sub> injection, the rock is saturated with brine, while, at monitoring time, some locations have been reached by a CO<sub>2</sub> plume. In these locations, the saturation is high since the CO<sub>2</sub> does not flow at low saturations. The distribution of CO<sub>2</sub> saturation at monitor time is, therefore, a mixture of a probability being zero for locations not reached by any plume, and a skewed distribution with weight on high saturations for locations within a plume. This type of mixed distribution can be modeled using a normal score transform of a correlated Gaussian random field. When mapping the saturations further into the rock-physics model, the properties of the mixture of CO<sub>2</sub> and brine within the Utsira formation are modeled using a patchy mixture of the fluids. Other details of the rock-physics model correspond to those given in Jullum and Kolbjørnsen (2016).

The nature of CO<sub>2</sub> implies stronger 4D effects of the fluid substitution than what is the standard for situations with brine, oil, and gas. In particular, the change in velocity is much stronger. Flooding a cell with CO<sub>2</sub> only allows for a reduction in velocity and density, which in turn is a strong restriction for the spatial change.

The local-neighborhood approach was implemented for the continuous parameter case using difference seismic. Figure 6 illustrates the elements of the inversion and the result based on inversion of differences in three angle stacks. (The seismic amplitudes are aligned using rms and pushdown data before the 4D difference is computed.) In the cross-sections, the average saturations are displayed, but underlying this is the full posterior distribution of the saturation in each location. In the vertical profile, this distribution is illustrated by percentiles (P10-P50-P90), in addition to the mean. The computational efficiency of the approach far exceeds an MCMC approach; a standard implementation of an MCMC algorithm used several days to achieve a convergent distribution for a vertical profile, while this local approach reached results within seconds.

## Conclusion

The spatial structure is important when performing an AVO inversion. It adds to the full understanding of the seismic response, supplements information from petrophysics and rock physics, and improves the inversion result.

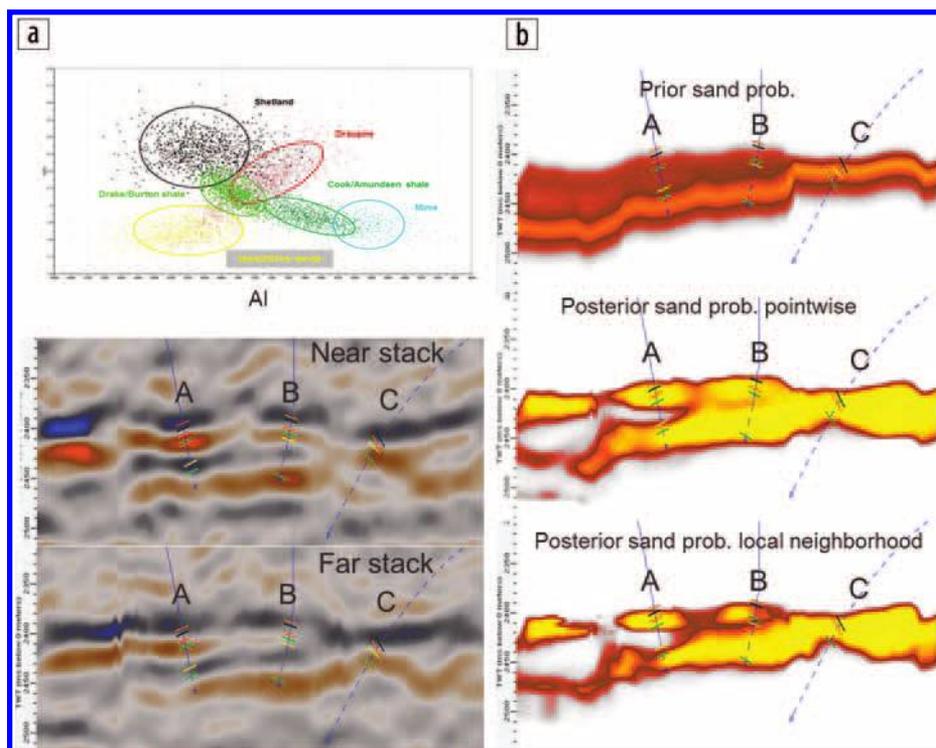
Using the local-neighborhood approach, we show that it is not necessary to incorporate all data in a global model in order to benefit from the information in the spatial model. An analysis on a local scale gives interesting results that in turn can be further combined

to derive properties on a larger scale. The local-neighborhood approach makes it possible to evaluate posterior uncertainty in a spatial model without resorting to time-consuming MCMC algorithms.

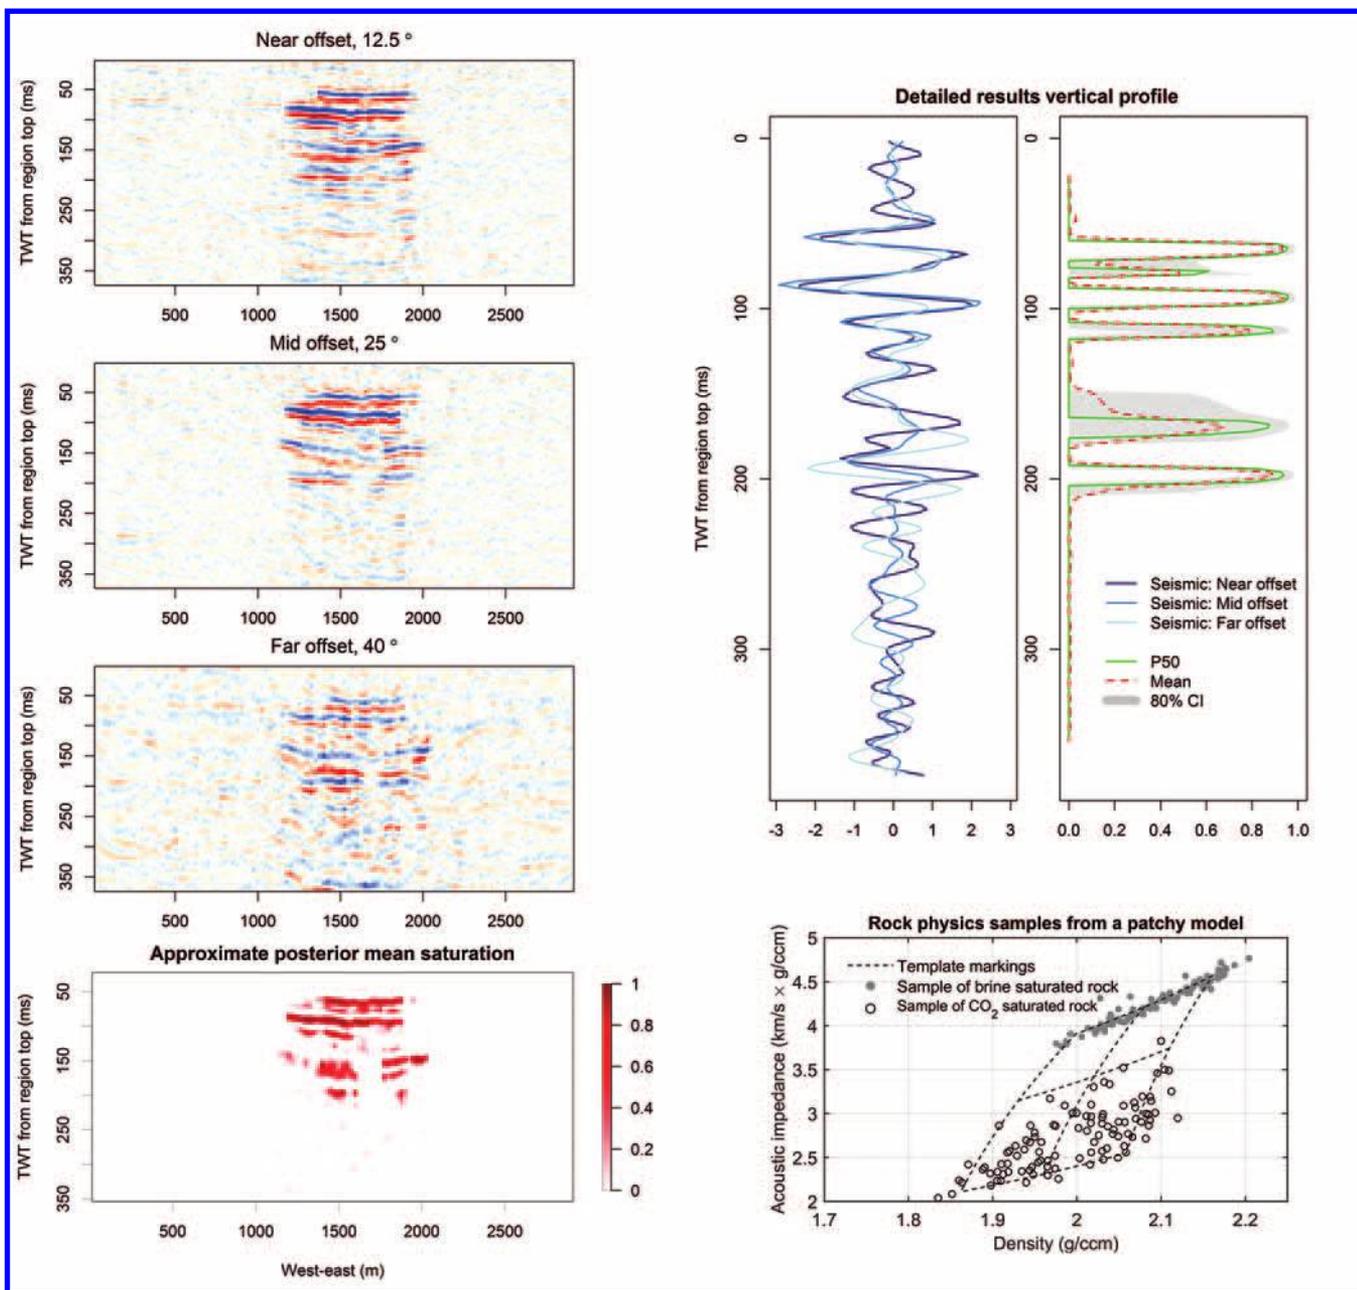
In lithology and fluid prediction, the fluid ordering and the sequence stratigraphic information give particularly large contributions when combined with the rock-physics distributions. It is not always easy to anticipate how the joint information contributes, as there are a multitude of options which are evaluated and weighted against each other. In the example of property evaluation, it is possible to analyze the complex dependencies introduced by restricting the changes in the elastic properties to be zero or negative. **ITE**

## Acknowledgments

We thank the Statfjord license partners, Statoil Petroleum, Centrica Resources, and ExxonMobil Exploration & Production Norway, for permission to show the example from the lithology inversion. We thank the Sleipner license



**Figure 5.** Statfjord example. (a) Rock-physics model and seismic amplitudes. (b) The posterior probability in the local-neighborhood approach is compared with the probabilities from the prior and the posterior using a pointwise approach.



**Figure 6.** Seismic difference data, rock-physics model, and inversion results for the Sleipner CO<sub>2</sub> injection case. The seismic amplitudes are scaled such that the peak value of the wavelets is identical for all angles.

partners, Statoil Petroleum AS, ExxonMobil Exploration & Production Norway AS, and Total E&P Norge AS, for permission to use the Sleipner data.

Corresponding author: Odd.Kolbjørnsen@lundin-norway.no

## References:

- Avseth, P., T. Mukerji, and G. Mavko, 2005, Quantitative seismic interpretation: Cambridge University Press.
- Buland, A., and H. Omre, 2003, Bayesian linearized AVO inversion: *Geophysics*, **68**, no. 1, 185–198, <http://dx.doi.org/10.1190/1.1543206>.
- Buland, A., O. Kolbjørnsen, R. Hauge, Ø. Skjæveland, and K. Duffaut, 2008, Bayesian lithology and fluid prediction from seismic

prestack data: *Geophysics*, **73**, no. 3, C13–C21, <http://dx.doi.org/10.1190/1.2842150>.

- Jullum, M. and O. Kolbjørnsen, 2016, A Gaussian based framework for Bayesian inversion of geophysical data to rock properties: *Geophysics*, **81**, no. 3, R75–R87, <http://dx.doi.org/10.1190/geo2015-0314.1>.
- Kemper, M. and J. Gunning, 2014, Joint impedance and facies inversion – Seismic inversion redefined: *First Break*, **32**, no. 9, 89–95.
- Kolbjørnsen, O., A. Buland, and R. Hauge, 2008, Method of Modeling a Subterranean Region of the Earth: U. S. Patent 9,069,100, B2.
- Larsen, A. L., M. Ulvmoen, H. Omre, and A. Buland, 2006, Bayesian lithology/fluid prediction and simulation on the basis of a Markov-chain prior model: *Geophysics*, **71**, no. 5, R69–R78, <http://dx.doi.org/10.1190/1.2245469>.