Journal of Petroleum Science and Engineering 183 (2019) 106381

Contents lists available at ScienceDirect

Journal of Petroleum Science and Engineering

journal homepage: www.elsevier.com/locate/petrol

A decision support system for multi-target geosteering

Sergey Alyaev[a,*], Erich Suter[a], Reider Brumer Bratvold[b], Aojie Hong[b], Xiaodong Luo[a], Kristian Fossum[a]

[a] NORCE Norwegian Research Centre, Postboks 22 Nygårdstangen, 5838, Bergen, Norway
[b] University of Stavanger, Postboks 8600 Forus, 4036, Stavanger, Norway

ARTICLE INFO

Keywords:
Geosteering
Sequential decision
Dynamic programming
Statistical inversion
Well placement decision
Multi-objective optimization

ABSTRACT

Geosteering is a sequential decision process under uncertainty. The goal of geosteering is to maximize the expected value of the well, which should be defined by an objective value-function for each operation.

In this paper we present a real-time decision support system (DSS) for geosteering that aims to approximate the uncertainty in the geological interpretation with an ensemble of geomodel realizations. As the drilling operation progresses, the ensemble Kalman filter is used to sequentially update the realizations using the measurements from real-time logging while drilling. At every decision point a discrete dynamic programming algorithm computes all potential well trajectories for the entire drilling operation and the corresponding value of the well for each realization. Then, the DSS considers all immediate alternatives (continue/steer/stop) and chooses the one that gives the best predicted value across the realizations. This approach works for a variety of objectives and constraints and suggests reproducible decisions under uncertainty. Moreover, it has real-time performance.

The system is tested on synthetic cases in a layer-cake geological environment where the target layer should be selected dynamically based on the prior (pre-drill) model and the electromagnetic observations received while drilling. The numerical closed-loop simulation experiments demonstrate the ability of the DSS to perform successful geosteering and landing of a well for different geological configurations of drilling targets. Furthermore, the DSS allows to adjust and re-weight the objectives, making the DSS useful before fully-automated geosteering becomes reality.

1. Introduction

According to the Norwegian Petroleum Directorate, drilling new wells is the most efficient way to increase oil recovery (Norwegian Petroleum Directorate, 2018). At the same time, well delivery and maintenance constitutes one of the major costs of oil reservoir development (Saputelli et al., 2003). To maximize value creation from each well, operators and service companies are continuously improving technology and procedures for optimizing the well placement to maximize production while reducing the cost of drilling and future maintenance. To place a well precisely in the best reservoir zone, operators use geosteering to adjust the well trajectory in response to real-time information acquired while drilling. The benefits of geosteering, such as higher production rates of the resulting wells, have been extensively documented in the literature (Al-Fawwaz et al., 2004; Janwadkar et al., 2012; Guevara et al., 2012; Tosi et al., 2017).

Geosteering has traditionally been dominated by manual geological interpretation and decision-making. Current computer-aided approaches assist decision-makers by co-visualizing a pre-drill deterministic geomodel alongside the inversion results of real-time data deep resistivity data. It is then up to the team of geoscientists to interpret the available information and decide steering actions in real-time (see e.g. Bø et al. (2014)).

More recently, there has been a focus on advancing computer-based methods both for pre-job, post-job and real-time analysis to support interpretation during drilling. The paper Antonsen et al. (2018b) discusses the importance of establishing a good understanding of how essential reservoir objects such as top reservoir and oil-water contact (OWC) are mapped by inversion of deep electromagnetic measurements in the pre-job phase of the drilling operation. This concept is extended in Antonsen et al. (2018a), where post-job case studies illustrate the importance of combining multiple LWD measurements with pre-job geophysical modelling. Bashir et al. (2016) have explained how the formation tops in a 3D geo-cellular near-well sector model were adjusted in depth during drilling. In Arata et al. (2017) a case study shows how a real-time local recalibration of seismic to minimize the depth

* Corresponding author.
E-mail address: Sergey.Alyaev@norceresearch.no (S. Alyaev).

https://doi.org/10.1016/j.petrol.2019.106381
Received 9 March 2019; Received in revised form 22 July 2019; Accepted 12 August 2019
Available online 28 August 2019
0920-4105/

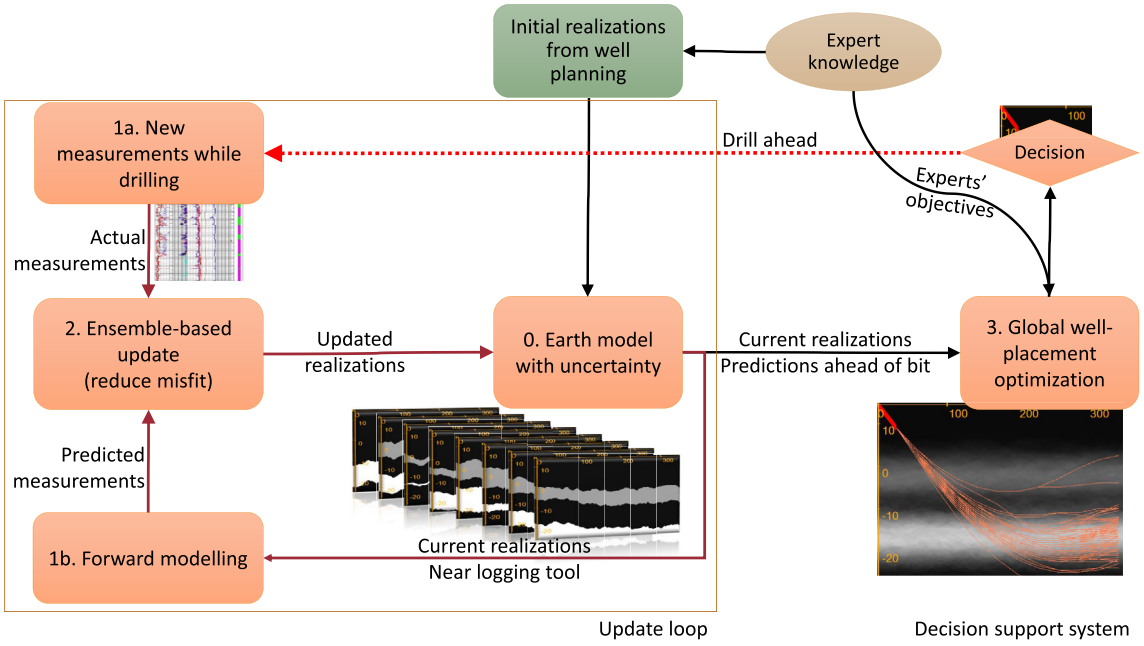

Fig. 1. Proposed geosteering workflow. The top part contains inputs to the workflow. The left part of the figure depicts the update loop. The part of the figure to the right contains the decision system that is based on the updated earth model. The 'drill ahead' decision results in new measurements that trigger another update and complete the full loop of the workflow.

discrepancy, based on LWD measurements, supports improved prediction of the reservoir boundaries ahead of the bit. Finally, in (Payrazyan et al., 2017) it is explained how the geological structures (faults and stratigraphic interfaces) in a 2D section along the well can be adjusted in real-time to fit the measurements, as basis for real-time sketching of a desired trajectory in the graphical interface.

The workflows discussed above are steps towards automated inversion and interpretation of real-time measurements. However, the information extracted from data has no value unless it helps us make better decisions. Geosteering is fundamentally about making decisions to optimize outcomes such as achieving optimal production at minimal costs. Making decisions that honor all different and sometimes conflicting objectives is not intuitive and requires excessive calculations that can only be handled by a computer.

Currently there is a lack of methods, tools and workflows that explicitly treat the uncertain nature of this decision process. To optimize the well placement under uncertainty, we should work within a probabilistic framework using a dedicated decision-analytic framework (Kullawan et al., 2014). The first step is utilization of prior data and descriptive analytics to summarize and improve our probabilistic understanding of the reservoir formation. Thereafter the real-time measurements provide information that improves our understanding of the geological and operational parameters that are crucial to optimal well placement. Finally, predictive analytics supports the continuous updating of our understanding of these parameters, and gives input to decisions on directional changes or stopping.

In this paper we present a consistent, systematic and transparent workflow for geosteering, which implements the principles above in a computer-based decision support system (DSS). The starting point is a probabilistic geomodel represented by multiple geomodel realizations which aim to span the space of pre-drill interpretation uncertainties. The real-time measurements obtained while drilling are continuously integrated by automatically updating the realizations using an ensemble-based filtering method, similar to Chen et al. (2015); Luo et al. (2015). The real-time update of the realizations aims to provide an always up-to-date prediction of the subsurface including interpretation uncertainty.

The update workflow is linked to the decision optimization. The DSS uses the up-to-date probabilistic geomodel to support geosteering decisions under uncertainty. It proposes well trajectories ahead of the bit and evaluates them against the chosen value function. The value function commonly includes multiple objectives, including production potential, costs for drilling and completion, and risks associated with the operation. The evaluation of trajectories is basis for the optimization. The trajectory optimization in the DSS is inspired by the discretized stochastic dynamic programming algorithms for geosteering that were discussed in Kullawan et al. (2017, 2018). However, the DSS presented here is specifically optimized for usage with ensemble-based update workflows which are already used for field development planning (Hanea et al., 2015; Skjervheim et al., 2015).

The real-time update workflow was previously demonstrated for pro-active geosteering with the objective to follow the top of a reservoir (Chen et al., 2015; Luo et al., 2015). The DSS presented in this paper combines this update workflow with dynamic programming for global trajectory optimization under uncertainty. The new optimization algorithm enables a variety of practical objectives, which among other things allow to optimize well-landing in uncertain environments.

The goal of this paper is to verify the DSS workflow on comprehensive synthetic experiments. Our numerical experiments are inspired by a challenging set-up with multiple target layers from a case study presented in Hongsheng et al. (2016). Unlike an expert service required for successful geosteering in the mentioned case, the DSS delivers reproducible and good decisions under uncertainty which maximize the set of objectives selected for an operation. We presented the flexibility of the DSS with respect to selection of objectives and initial tests in earlier conference proceedings (Alyaev et al., 2018a, c). Here we focus on exhaustive presentation of the features and limitations of the DSS algorithms and present a statistical verification of the performance of the system.

The rest of the paper is organized as follows. Firstly, we present the ensemble-based workflow for updating of the probabilistic geomodel based on real-time measurements. Secondly, we introduce the DSS that can utilize the up-to-date ensemble to propose optimal decisions under uncertainty. After that, the performance of the DSS is demonstrated on synthetic cases with multiple targets. Finally the main contributions of the paper are summarised in the conclusions.

2. Earth model update loop

In the proposed geosteering workflow, shown in Fig. 1, real-time decision support is based on Bayesian inference from a probabilistic geomodel that is continuously updated.

2.1. Earth model

The earth model is represented as an ensemble of realizations that captures key geological uncertainties. In Fig. 1 the uncertainties are the positions and thicknesses of sand layers (gray) in a background shale. The pre-drill realizations are created based on a priori information drawn from seismic, logs from offset wells, production measurements and additional knowledge about geological uncertainties provided by experts.

All realizations are updated incrementally each time new measurements become available while drilling. The incremental updates of the model are performed by a statistically-sound ensemble-based method that reduces the mismatch between the measurements and the geomodel. The ensemble-based updating approach is an implementation of a Bayesian updating framework. In the rest of the section we describe the implementation of the individual components of the generic update loop that was used in this study.

2.2. Measurements

By design, the ensemble-based methods perform incremental updates which can handle any number and any type of measurements simultaneously (see 1a. in Fig. 1). It is required however, that there is a corresponding simulation model that can transform the realizations and the measurements to a context where they can be adequately compared to compute the mismatch. The simplest way is to use a forward model that produces synthetic measurements based on the geomodel realizations (see 1b. in Fig. 1). The forward model should be sufficiently fast to handle hundreds of simulations at every assimilation step.

2.3. Forward modelling

The simulation methods for processing of different logs have been extensively studied by service companies (Sviridov et al., 2014; Dupuis et al., 2014; Dupuis and Denichou, 2015; Hartmann et al., 2014; Selheim et al., 2017) but are generally not available in the open domain, only as paid services. The main contribution of this paper is the DSS and not the modelling of the measurements. Therefore, we use a simple integral model for electromagnetic (EM) measurements following Chen et al. (2015). The tool set-up in that paper has a look-around capability of about 5 m, and it is sensitive to resistivity in the up, down and side directions, see Fig. 2. The tool is placed at the drill-bit in the current prototype. The depth of investigation (DOI) is chosen relatively low compared to the modern deep EM tools (e.g. Seydoux et al. (2014)) to maintain the accuracy of the approximate integral model. However, we emphasize that this does not constrain the applicability of the workflow. For instance, in Luo et al. (2015), a similar update workflow has been tested with more advanced tools and a finite difference forward model. The tool modelled in Luo et al. (2015) provides a higher DOI that allows to see in a larger volume around the well and is expected to yield better results.

2.4. Ensemble-based update algorithm

The update loop used in this paper is compatible with a number of ensemble-based methods which have previously been implemented for reservoir data assimilation including the ensemble Kalman filter (Aanonsen et al., 2009), ensemble smoother (Skjervheim and Evensen, 2011; Skjervheim et al., 2015), the particle filter (Lorentzen et al., 2016), and more sophisticated combinations of the above, such as adaptive Gaussian mixture filter (Lorentzen et al., 2017). To demonstrate the workflow, we use the standard ensemble Kalman filter (EnKF) method (Chen et al., 2015; Luo et al., 2015) for the implementation described in this paper.

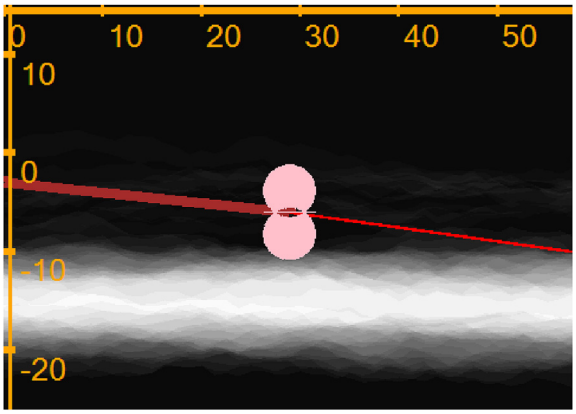

Fig. 2. An illustration of the depth of investigation (DOI) of the tool for the synthetic model (axes in meters). The dark red line segment to the left shows the trajectory that has been drilled; the thin red line segment shows the next proposed trajectory segment. The DOI is illustrated at the current decision point at the end of the drilled trajectory. The measurements at the highlighted location have been assimilated. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

The Kalman filter (Kalman, 1960) formulates the Bayesian update for the changes of the mean and the covariance matrix when new data is received assuming all probability distributions are Gaussian. The ensemble Kalman filter (Evensen, 1994, 2009) is a flexible Monte Carlo implementation of the Kalman filter. When new measurements are received, they are compared to the simulated measurements generated by the corresponding forward models for each realization. The realizations used in this paper assume a layer-cake geomodel with constant resistivity in each layer. The depth of each layer boundary is represented by a series of points. The new measurements yield an incremental update of the depth values in the interfaces, which can be formulated as in Burgers et al. (1998):

$$y_{updated} = y_{initial} + K(d_{measured} - d_{modelled}), \tag{1}$$

where $y_{updated}$ contains the updated (or posterior) ensemble representing the posterior distribution, $y_{initial}$ contains the initial (or prior) ensemble representing the prior distribution, K is the Kalman gain matrix, $d_{measured}$ contains the perturbed measured data values,[1] and $d_{modelled}$ contains the modelled data values corresponding to the initial ensemble for that update. Equation (1) describes a linear combination of the prior and the measured data, which is weighted by the Kalman gain matrix K. Bayes' rule describing the relationship among the prior, likelihood, and posterior is not shown explicitly in the equation above, but it is implicitly included as the likelihood is encoded in the Kalman gain matrix K and the pre-posterior is treated as a normalizing constant of the updated ensemble. A detailed description on the relationship between the formulation of the Kalman filter and the Bayesian formulation can be found in Meinhold and Singpurwalla (1983).

The incremental nature of the updates as drilling progresses removes the need for a costly direct inversion that include all the available measurements each time the model is updated. By design the updates (e.g. depth of boundaries) are also propagated ahead of the bit using the prior knowledge about the model.

This provides a probabilistic prediction of the geology ahead of the bit based on the trends identified around and behind the bit. We refer

[1] For EnKF, a measured data value has to be perturbed with its corresponding statistics in order to avoid insufficient variance (Burgers et al., 1998).

the reader to Luo et al. (2015) for a more rigorous description of the update loop for geosteering.

3. Decision support system (DSS)

The update loop described above results in an always up-to-date ensemble of model realizations which integrates both the prior knowledge and the latest measurements. The realizations representing the probabilistic description of relevant and material geological uncertainties are the input to the DSS. The DSS is based on a normative decision-making approach (Bratvold and Begg, 2010; Clemen and Reilly, 2013; Howard and Abbas, 2015) and includes optimization algorithms that take into account all realizations as well as multiple decision objectives, such as following the reservoir top while minimizing drilling cost and reducing the tortuosity of the well for easier completion. The objectives may be conflicting, and for each realization the algorithm calculates the well path that is optimal with respect to the weight of each objective. The objectives and the corresponding weights are defined by the user of the DSS and is a consistent way to include expert knowledge in the workflow, see Fig. 1. The proposed decision for stopping or adjusting the trajectory is visualized together with the current representation of the uncertainty in the model.

The DSS presented in this paper differs from traditional decision systems that were designed for strategic decisions. In contrast, the geosteering decisions are operational. This implies that they must be taken within short time. The presented DSS allows to tweak the weights of the objectives at all decision points and preview the outcomes in real-time. This helps the user to build an understanding of how the choice of objectives influence the suggested decisions and provides a possibility to re-evaluate the trade-offs between objectives as the drilling progresses.

The logically consistent approach of the DSS allows for decisions to be transparent and reproducible. Given that the DSS is based on a normative decision quality approach, it will recommend good decisions for the decision-maker's objectives, alternative choices for a decision (following constraints), and geological beliefs (represented in the pre-drill model).

3.1. Objectives

A natural requirement for any DSS is the possibility to take into account multiple objectives. The objectives used in modern geosteering operations include placing the well in a specific position in the reservoir, reducing costs and ensuring safety (Kullawan et al., 2016). For the use in a DSS the objectives need to be converted into objective functions defined on a common scale, e.g. the estimated profit in US dollars or produced-oil-barrels equivalent. To reduce conversions, we will use the stand length drilled within standard reservoir sand as the common scale in this paper. We denote each objective function as $O_i(X|M)$, which depends on the trajectory (X) of the well and the actual sub-surface configuration (M), which for now we assume to be known and deterministic. The profit functions are positive and costs associated with the operation are negative. Objective functions that are used in our numerical examples are summarised in the appendix (Section 6).

The global objective $O(X|M)$ is represented as a linearly weighted sum of individual contributions from each objective function:

$$O(X|M) = \sum_i w_i O_i(X|M), \tag{2}$$

where $0 \leq w_i$ is an objective weighting factor for objective i, where i are indices of different objectives.

The functions O_i are scaled so that the initially estimated (pre-drill) profit/cost corresponding to each objective function is achieved when $w_i = 1$. It is convenient to think of $w_i = 0$ as ignoring the objective i, while $w_i = 1$ means setting the value of objective i to the scale anticipated in the pre-drill analysis. Furthermore, adjusting the weights gives the user the control to modify the priority of the objectives and to maintain the predictions at a desired scale, both in response to insights gained during the drilling operation.[2] $w_i > 1$ corresponds to a higher priority to objective O_i, while $w_i < 1$ corresponds to a lower priority. Changing the weights reflects an insight in how each objective contributes to the profit/cost of the well being drilled compared to the originally anticipated (see the last numerical example in the next section). We will use the global objective defined by (2) as the value function in the optimization for the rest of the section.

3.2. Sequential decision optimization under uncertainty

A geosteering operation consists of a sequence of decisions D_k . Subscript k numerates decision points sequentially in time. Ensemble-based workflows represent the uncertainty in the geological interpretation as a set of realizations. Substituting different realizations M_j instead of the deterministic model M into the objective function in equation (2) typically gives several trajectories, where each is optimal for the corresponding realizations. At the same time, the outcome of the optimization should be a single optimal decision for each decision point k. In this paper we follow the optimality criterion used in robust optimization: We want to make a decision D_k that maximizes the expected value of the well given all the available information.

Let us consider all available information at time k. At each decision point the ensemble-based workflow contains up-to-date realizations representing the current understanding of the subsurface. Moreover, between any two sequential decision points, new measurements are assimilated using the update loop. This improves the geological understanding around and ahead of the drill-bit. The full structure of a sequential decision problem is shown in Fig. 3, where D_k denotes the decision at time k, and I_k denotes the information gathered between time $k-1$ and k. For brevity of notation we denote all information gathered between k_1 and k_2 as $I_{(k_1:k_2)}$. Generally, at some current time $k = 0$, the decision D_0, that needs to be made right now, depends not only on the information that has been gathered $I_{(start:0)}$ and is contained in the geomodel, but also on the possibility of future learning, i.e. the information that will be gathered $I_{(1:end)}$. Because the future information $I_{(1:end)}$ is not available at time k, its influence on future learning and decision-making can be modelled as uncertain events, conditioned on the current information and prior decisions.

The approach considering the full learning and decision-making structure of a sequential decision-making problem is presented in Fig. 3. We call this approach “far-sighted” as it takes into account what might happen in the future including which information that will be gathered, how uncertainties will be updated using that information, and which decisions that will be made (Alyaev et al., 2018b). An implementation of the far-sighted approach using discretized stochastic dynamic programming has been described in detail in Kullawan et al. (2017, 2018) for a geosteering problem considering a geomodel with a single reservoir layer and updates of its boundaries.

Unfortunately, the stochastic modelling required for understanding the effect of future learning in the far-sighted approach is computationally prohibitive for real-time decision-making. First, the complexity of the problem grows exponentially as the numbers of decision points, alternatives and uncertainty branches increase. The phenomenon is known as the curse of dimensionality, see Brown and Smith (2013). Thus, the far-sighted approach becomes computationally prohibitive for problems with a large number of parameters. Second, the state-of-the-art methods for data assimilation (e.g., the ensemble Kalman filter used in our update loop) cannot be directly embedded into the far-sighted approach. The far-sighted approach requires generating and storing not only realizations for the current decision point, but also realizations that are modified due to future updates in the EnKF loop, for all future

[2] In most situations $\sum_i w_i \neq 1$.

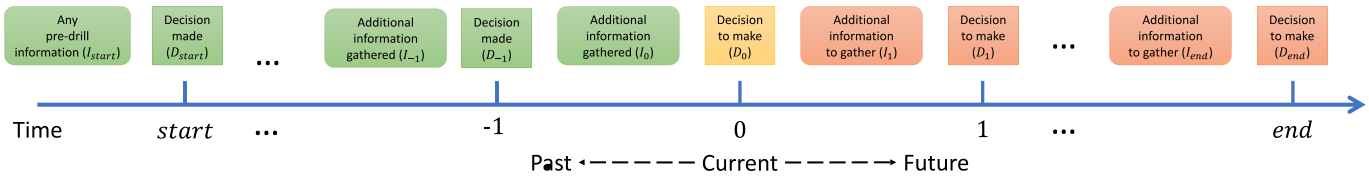

Fig. 3. Full structure of a sequential decision problem.

decision points.

In this paper we present a new dynamic programming discrete optimization strategy which has real-time performance and is simple to integrate with the ensemble-based update loop. The strategy is a simplification of the far-sighted approach. It considers future decisions but omits the modelling required to simulate the future learning. Instead of the modelling of the future information we optimistically assume that perfect information (about the subsurface) would be available after the current decision has been made and before the next decision is made. Thus, this approach can be classified as naive optimistic (Alyaev et al., 2018b). It is possible to find theoretical scenarios for which this approach gives a decision recommendation that is different from the optimal choice given by the far-sighted approach (Alyaev et al., 2018b). At the same time, the naive approach is superior to myopic optimization that only considers one step ahead which was used in previous papers with ensemble-based workflows, e.g. Luo et al. (2015); Chen et al. (2015). The DSS optimizes the complete well path ahead of the bit against the currently available representation of the geological uncertainty (as represented by the set of realizations). Hence, the realizations should capture the complete current view on the geology with its interpretation uncertainties. The next subsection summarises the implementation of the algorithm.

3.3. Real-time ensemble-based optimization algorithm

For simplicity we assume near-horizontal drilling and therefore we can associate the decision points D_k with their position x_k along the horizontal axis. At every x_k we discretize the trajectory alternatives by the well depth and denote the depths as z_i for the horizontal location x_k. Moreover, to account for the dogleg severity constraint (see the Appendix), it is important to take the current well angle α_{i_k} into account.

The decision for step D_k at $(x_k, z_{i_k}, \alpha_{i_k})$ is either to stop or to add a segment connecting to a point $(x_{k+1}, z_{i_k,next})$. The choice of depth $z_{i_k,next}$ is constrained by the dogleg severity given by the user input, which in a discrete sense is approximated by $\alpha_{i_k,next}(x_k, z_{i_k}, x_{k+1}, z_{i_k,next}) - \alpha_{i_k}$. The optimization algorithm evaluates different trajectories that are represented as piecewise linear curves that go through the points (x_k, z_{i_k}) ahead of the current decision point, see Fig. 4. The resolution of points can be decided by the user and will affect the trade-off between the optimality of decisions and the computational time.

For the decision at any decision point D_k, the decision algorithm consists of two steps.

In the first step, a dynamic programming algorithm finds a deterministic optimal well path for each realization and for every starting point by iterating over all possible trajectories. The result of this step is the set of optimal decisions for every point-angle pair (x_k, z_i, α_i) for every realization j expressed as a z-coordinate for x_{k+1} (or "stop drilling"[3]):

$$\hat{z}^j_{i,next}(x_k, z_i, \alpha_i) = \arg \max_{z_l \text{ within constraints}} O([(x_k, z_i), (x_{k+1}, z_l)]|M_j) \quad (3)$$

$$+ \gamma O\left(X_{(x_{k+1},z_l,\alpha_l)}|M_j\right), \quad (4)$$

[3] The "stop drilling" decision means that for the current realization a positive value cannot be achieved, and stopping is the best alternative.

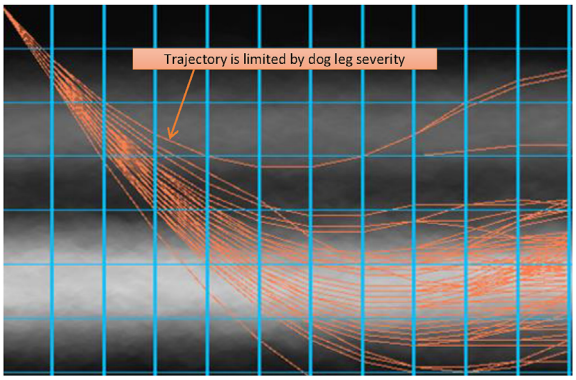

Fig. 4. An example of discretization of trajectories. Vertical lines correspond to x_k while horizontal lines correspond to z_i (every 10th line in the set of possible depths is displayed). The orange polylines are possible well trajectories that go through the decision-grid points. One can see that the well trajectories are constrained by the dog leg severity; unreachable trajectories are not considered. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

where $[(x_k, z_i), (x_{k+1}, z_l)]$ denotes the next segment of the trajectory, *gamma* is a discount factor, and $O(X_{(x_{k+1},z_l,\alpha_l)}|M_j)$ is the highest possible objective value that can be achieved for model M_j from the trajectories $X_{(x_{k+1},z_l,\alpha_l)}$ starting with point (x_{k+1}, z_l) and angle $\alpha_l(x_k, z_i, x_{k+1}, z_l)$.

In (3) the index k is used to refer to horizontal locations, i and l for vertical locations and j for the realizations respectively, as before.

The discount factor $0 < \gamma \leq 1$ in (3–4) is commonly used in the formulations of sequential decision problems (Feinberg and Shwartz, 2012). It reduces (discounts) the value of decisions that are further ahead $O(X_{(x_{k+1},z_l,\alpha_l)}|M_j)$ compare to the immediate expected reward $O([(x_k, z_i), (x_{k+1}, z_l)]|M_j)$ from the current decision. In most of the paper, if not stated otherwise, we will use $\gamma = 1$ corresponding to the naive optimistic policy described in Section 3.2. Values of γ slightly less than one allow to reduce the value gained from the trajectory far ahead and thus present a practical way to compensate for the assumption of perfect information in the naive optimistic policy.

Equation (3) needs to be solved for the current position of the drill bit. All the paths ahead are then recovered by finding the term (4) recursively. In our implementation we follow the principles of dynamic programming (Cormen et al., 2009) to ensure that each point is evaluated only once and then tabulated to be reused in the other trajectories. Thus the reconstruction of the optimal trajectories as well as the subsequent evaluations for the objective function for the well is almost instantaneous. In this way the optimal trajectory for each realization can be recovered:

$$\hat{X}^j = \left[(x_0, z_{i_0}), (x_1, z^j_{0,next}), (x_2, z^j_{1,next}), \ldots\right], \quad (5)$$

where (x_0, z_{i_0}) is the starting point for the current optimization. Similarly, by substituting (5) into the objective function (2), one can calculate the predicted well value for a given geological scenario.

In the second step, we need to perform a robust optimization to arrive at the single optimal decision: i.e. chose to "stop drilling" or steer towards the depth $\hat{z}_{0,next}$, whichever gives the best outcome on average, considering all realizations.

The computation of $\hat{z}_{0,next}$ considers immediate permissible

alternatives, including "stop drilling" and all (x_1, z_l) which are within the constraints of the dogleg severity, and choose the one that is the best on average:

$$\hat{z}_{0,next} = \arg \max_{z_l \text{ within constraints}} \sum_{j=1}^{n} \psi_j \Big\{ O([(x_0, z_0), (x_1, z_l)] | M_j) + \gamma O(X_{(x_1, z_l, \alpha_l)} | M_j) \Big\}, \quad (6)$$

where ψ_j is the probability of realization j and the rest of the notation is the same as in (3).[4] We emphasize that equation (6) is used for exactly one decision ahead. Thus the computational complexity of its evaluation is proportional to the number of immediate alternatives times the number of realizations. The vital consequence is that it does not suffer from the curse of dimensionality. This distinguishes our optimization strategy from earlier approaches (e.g. Barros et al. (2015)), which try to optimize all future decisions while neglecting the future learning.

The optimization algorithm presented in this section extends the classical robust optimization (Chen et al., 2009, 2015; Lorentzen et al., 2006) to include the up-to-date knowledge when optimizing the full trajectory ahead of the bit. Due to the possibility of future learning, it is essential that only the first point is chosen by the robust optimization while the rest of the trajectory is allowed to differ from realization to realization. The future learning is expected to reduce the geological uncertainty and improve the decision for the next decision point. The full well path joint optimization for the whole ensemble is costly and not always justifiable for a workflow where updates of the realizations are performed sequentially in time when new measurements arrive during drilling and time is scarce. Instead, the decision for the next step (the next decision point) is recomputed once the new measurements become available and the ensemble is updated. This strategy allows for a real-time reaction to new information while also considering the prior information at every decision point. From the perspective of decision theory the strategy is equivalent to dynamic programming with assumption of perfect information.

3.4. Visualization of the real-time modelling results

The adoption of any DSS requires that the system can be trusted by its users. Therefore the communication to the user of the reasoning behind the proposed decisions is essential. In the user interface, the proposed decision is visualized and the basis for the decision is explained.

The main basis for a decision is the up-to-date probabilistic earth model. In Fig. 5, an earth model with two oil-bearing sand layers with high resistivities is used for the demonstration of the DSS visualization. Between these two layers there are background shales with low resistivity. High resistivity layers are indicated with a bright color, while relatively low resistivity layers appear as gray. Black layers have very low resistivity and correspond to shale. To the right in Fig. 5, three (out of normally a hundred) realizations are shown. In the user interface any realization can be selected for examination and the realization on display can be effectively switched within milliseconds. Moreover, the uncertainty can be visualized as a 'point cloud'. That is, for each point in space we visualize the average of the resistivity value over the ensemble of realizations as shown in Fig. 5. In many cases, the point cloud is an intuitive way to understand the distribution of the uncertainty within the current ensemble.

At all times the interface highlights the consequence of the immediate decision (next proposed well segment) in thick red and with a written communication of the calculated decision: an angle in degrees or 'stop'. The decision recommendation is supported by a cumulative plot of the expected value of the well based on the estimated geological uncertainty, shown in the bottom left corner of Fig. 5. The plot should be interpreted as follows; for a selected value on the x-axis, the plot surface corresponds to a percentage, e.g. 20%. That means that in 20% of the realizations this value is not achieved. However, the value is exceeded in the remaining 80% of realizations.

Furthermore, the interface communicates the two-step process behind the decision optimization as explained in the previous subsection. When an individual realization is shown, the corresponding optimal trajectory resulting from optimization step one (5) is visualized (starting from the uniquely selected next segment). At the same time the value expected from this trajectory is marked on the value plot. For convenience, the mean predicted value is also shown.[5] In the 'point cloud' view, all the optimal trajectories corresponding to each realization are visualized (see Fig. 5). By evaluating the density of trajectories, this latter display gives an intuitive understanding of the alternatives that are in reach of the current operation.

The display is fully updated with the new optimization results when the realizations are updated, or if the user adjusts the objectives. Further discussion of the flexibility of the graphical (GUI) and programming (API) interfaces of the system can be found in Alyaev et al. (2018a). In the rest of the paper we will use the point cloud view to visualize uncertainty and the reachable optimal trajectories for each realization to indicate the outcomes predicted by the system.

4. Numerical examples

In this section we demonstrate the performance of the DSS on synthetic examples. In the examples we use a layer-cake earth model with two oil-bearing sand layers surrounded by background shales (see e.g. Fig. 5). This seemingly simple model presents a challenging setting for making decisions: the layer depths and thicknesses are uncertain and will be updated while drilling, and the drilling target is not predefined but is selected dynamically based on a multi-objective value function. The value function accounts for the estimated production potential of the well versus the estimated cost of drilling. Fast and consistent evaluation of this function while considering different alternatives under geological uncertainty is the key to good decisions.

The operation starts 15 m above the expected reservoir top (the expected top location is taken as zero) with an 80° inclination as commonly used directly before landing (Cayeux et al., 2018). The operation is assumed to end after drilling 350 m in true horizontal length. The decision points are equally spaced in the horizontal direction with approximately one stand between them (28.6 m apart). The earth model is updated using unprocessed synthetic EM measurements (taken at the decision points) that are modelled by a simple integral model as in Chen et al. (2015). The depth of investigation (DOI) of the synthetic tool is about 5 m, see Fig. 2. The data variance in the update equation (1) is set to be 0.5 in dimensionless resistivity units.

All the tests in this section follow similar assumptions about the layered model. The layers in the model can be distinguished by their resistivity that is assumed to be known for the synthetic cases. The resistivity values are set to 10 for shale and to 150 and 250 for the top and bottom sand layers respectively (all in dimensionless units). The initial ensemble of realizations is created based on the expected boundary depths that vary around a mean value of 0, -5.3, −13.3 and −20.1 m respectively. Depth uncertainties are generated using an exponential variogram model (nugget = 0, sill = 2.5, range = 350 m) following implementation from Cressie (1992). Furthermore, co-kriging is used to correlate the boundaries of the neighbouring layers (with correlation parameter set to 0.7), similar to Lorentzen et al. (2019). For

[4] Note that the evaluations for all the trajectories for the individual realizations have already been performed and cached on step one of the algorithm by applying equation (3). Also note that due to differences between equations (3) and (6) the final decision (steering or stopping) does not necessarily coincide with any (depending on the realization) of the optimal decisions from step one.

[5] It is important to note that the mean does not necessarily coincide with a value expected from any realization.

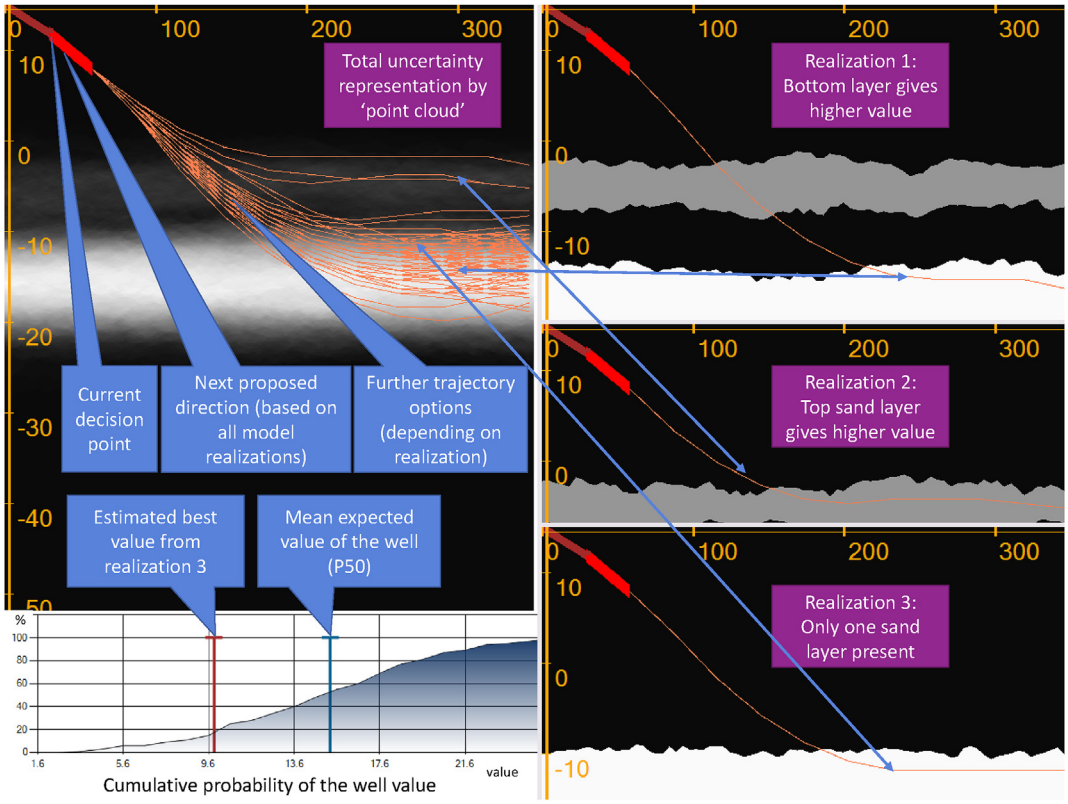

Fig. 5. An illustration of the functionality of the DSS interface. To the right in the figure, individual geomodel realizations are visualized. The realization to be visualized can be selected by the user. To the top left in the figure, a 'point cloud' view of the total ensemble of realizations is shown. The arrows indicate how the optimal trajectory for each realization is visualized in the 'point cloud' view. To the bottom left, a cumulative diagram of the expected value of the well (including costs and future income) is presented. It is based on the current uncertainty captured by the ensemble, with the vertical marks for (i) the mean expected value and (ii) the estimated value for a specific realization.

rigorous testing, we let the synthetic truth also be a layer-cake model.

For the DSS we will consider the following main objective; to maximize the well exposure to the reservoir sands. Instead of fixing the target the system should choose the sand layer based on its thickness, under the assumption that a thicker sand results in higher oil content and therefore better production. The two sand layers are otherwise considered to be equally good for drilling and production. Additionally, the well should be placed in the upper part of a layer for improved production, and the drilling cost should be taken into account. The trajectory is constrained by dogleg severity of 2° and a maximum inclination of 90° to allow for efficient gravel packing. The precise mathematical definition of the individual objectives as well as their weighting is given by equation (14) in the Appendix (Section 6). For simplicity the value functions are scaled to "equivalent meters in sand drilled". That is, one **unit** corresponds to the net present value that can be produced from one reference well stand positioned along a sand layer with reference properties of 1-m thickness.

4.1. Optimal landing in different geological scenarios

First we want to test the DSS workflow for different geological scenarios. All decisions automatically follow the recommendations from the DSS.

To the left in Fig. 6, two different alternatives for the synthetic truth are considered followed by the results of application of the workflow on the figure's right. In the example we compare how the same set-up of the workflow operates step-by-step for these two different scenarios (everything is the same except for the synthetic truths). The truth in both scenarios contains two reservoir layers. In the top scenario the top layer is thicker and hence more profitable while in the bottom scenario the bottom layer should be prioritized. The images with the synthetic truths also contain the trajectories that are optimized with respect to the chosen metric (see equation (14) in the Appendix).

We start from an initial ensemble of realizations representing the model with uncertainty (the ensemble is the same for both scenarios). The 'point cloud' representation of the initial ensemble is depicted in the first column in Fig. 7. The blurry contours of the boundaries indicate the uncertainty in the layer positions and thicknesses. Initially, the global optimization foresees different decision outcomes that would result in landing in either the top or the bottom layer (the same for both scenarios as the pre-drill ensemble is the same, no model updates have been performed yet). Therefore the DSS proposes the same initial decision for both scenarios: a build-up from 80 to 81.1°, allowing for future well landing in either of the layers.

In the next two decision steps (columns 2 and 3 in Fig. 7) the tool look-around is insufficient to reach the sand layers (the sensitivity of the tool shown in pink). Therefore no update takes place and, since the geomodel is the same, the steering decisions are the same for both scenarios. The DSS proposes an angle build-up that allows for better landing in the top layer, which seems more promising under the current view on the geological uncertainty. At the same time the alternative to drill to the bottom layer is not disregarded, as indicated by the optimal well paths in some of the realizations.

In column 4 in Fig. 7 the expected top boundary of the top sand comes within the DOI of the tool. The uncertainty captured in the geomodel is correspondingly reduced after the update, rendering the top boundary sharper on the averaged image. At this stage it is still uncertain which layer that will be chosen for both synthetic truths, but the objectives dictate to steer downwards in the bottom scenario to be able to reach the bottom layer faster, if required.

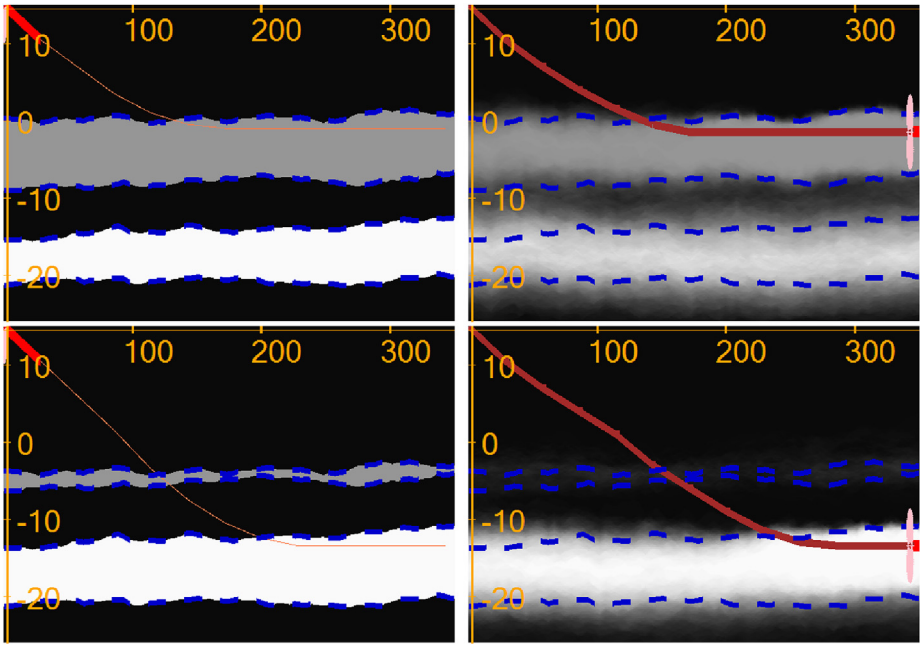

Fig. 6. The synthetic truths for the two scenarios with their corresponding optimal trajectories (the two images to the left). The value corresponding to each optimal trajectory is taken as the maximum theoretically possible for that scenario (100%). The two images to the right show the final trajectories resulting from the application of the workflow in the two scenarios. Top scenario: The well almost matches the perfect trajectory and achieves 86.6% of the theoretically possible value. Bottom scenario: The landing in the second layer is not perfect due to the initial uncertainty (not shown). Nevertheless the global optimization under uncertainty allows to adequately land the well in the bottom layer and achieve 58.6% of the theoretically possible value.

In column 1 in Fig. 8 the uncertainty in the top layer depth and thickness is reduced even further (notice the sharp boundaries of the top layer in the top image in column 1). More precise knowledge of the reservoir layers and the profitability of drilling in each of them puts more priority to landing in the top layer for the top row scenario. Contrary to that, in the bottom row scenario the decision is to cross the shale between the two sands and drill to the bottom sand layer. After the update performed in column 2 of Fig. 8, the uncertainties for the top layer boundaries are further decreased and the landing strategies are confirmed for the two scenarios. These decisions are consistent with our knowledge about the truth in both scenarios.

In columns 3 and 4 in Fig. 8, in the top scenario, the well is landed in the top layer. At the same time the DSS estimates that it might be better to drill downwards for some realizations. This is reflected in the figure by the thin well trajectories. In column 4 in the bottom scenario in Fig. 8 the DOI of the tool reaches the expected top boundary of the bottom layer. The uncertainty about the depth of the roof of the layer is reduced, yielding a more detailed landing plan.

The rest of the synthetic operation is shown in Fig. 9. In the bottom row of the figure, one can observe how the well is landed successfully in the bottom layer.

During the decision steps described in this section, a complex workflow consisting of the update loop and the DSS is running behind the scene. The measured data is generated using the described EM acquisition model from the synthetic truth, including added measurement noise. Every new measurement triggers an iteration of the update loop. On a workstation with 20 logical cores a full model update takes less than a second for an early software implementation that is not optimized for production. Afterwards the DSS optimizes the trajectory of the well across all realizations and gives a result within another 10 s.

Fig. 6 shows the final step of the operation for the considered scenarios. In both scenarios the DSS manages to land the well in the layer which is optimal with respect to the objective. We emphasize that this is possible due to the fact that the full well trajectory is optimized against the up-to-date uncertainties.

In the scenario in the top row (Fig. 6), the steering result is close to

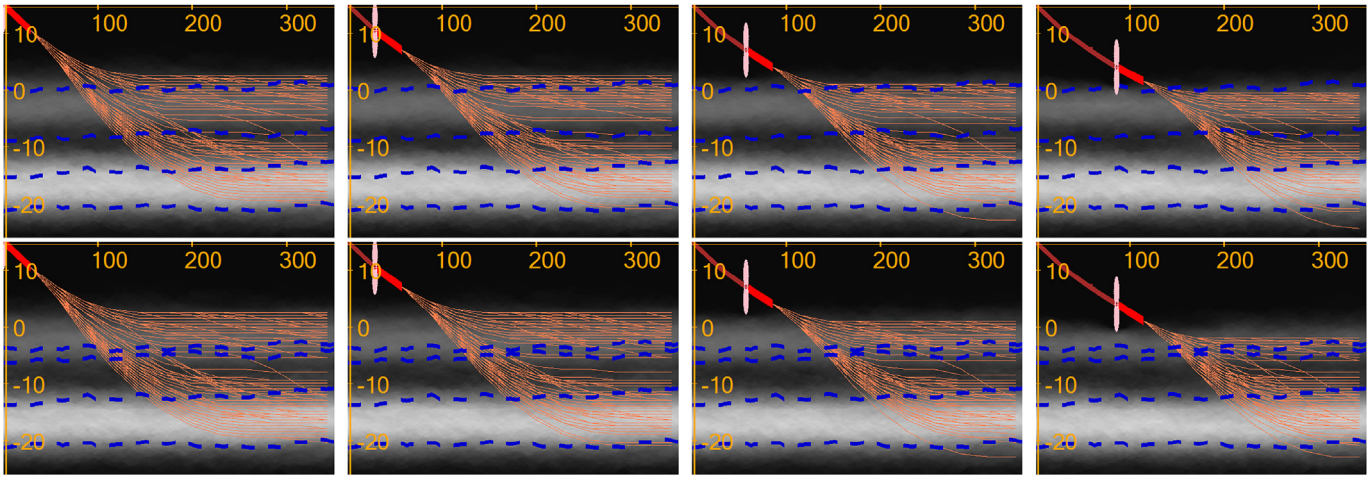

Fig. 7. Demonstration of the decision support system for the two synthetic scenarios in the top and bottom row as in Fig. 6. Blue dashed lines indicate the positions of the layer boundaries in the synthetic truths. The figures show the step-by-step decision recommendations from the DSS (as indicated by the advancing bit). The initial geomodel uncertainty is the same for both scenarios. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

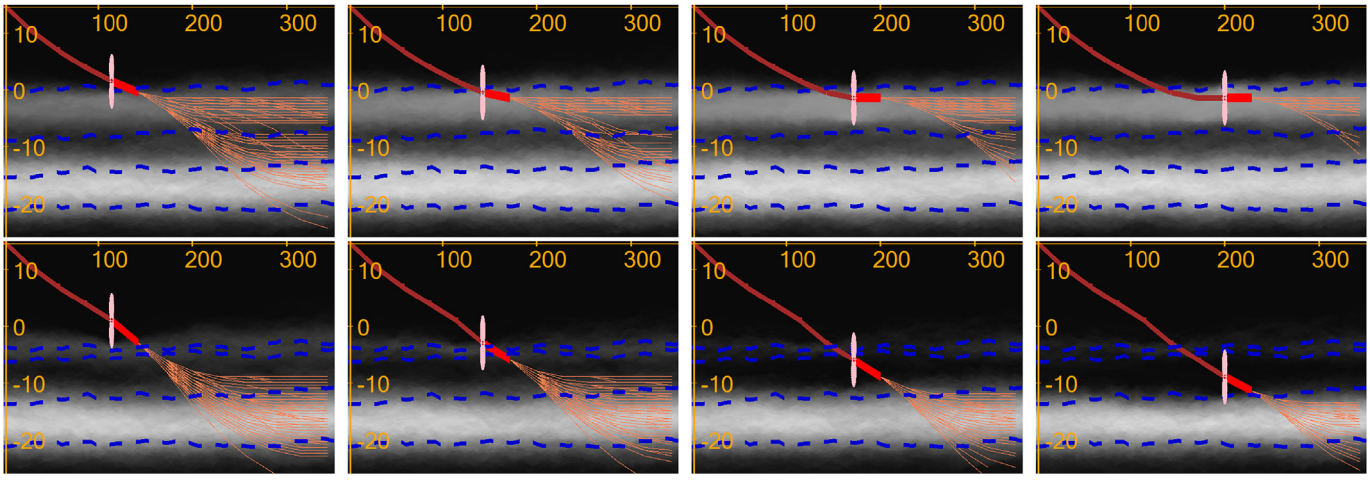

Fig. 8. Continued demonstration of decision support system for two synthetic scenarios from Fig. 7. The figures shows the step-by-step outputs of the DSS. As the drilling operation progresses and more data become available, the top layer is preferred for the top scenario and the bottom layer is preferred for the bottom scenario.

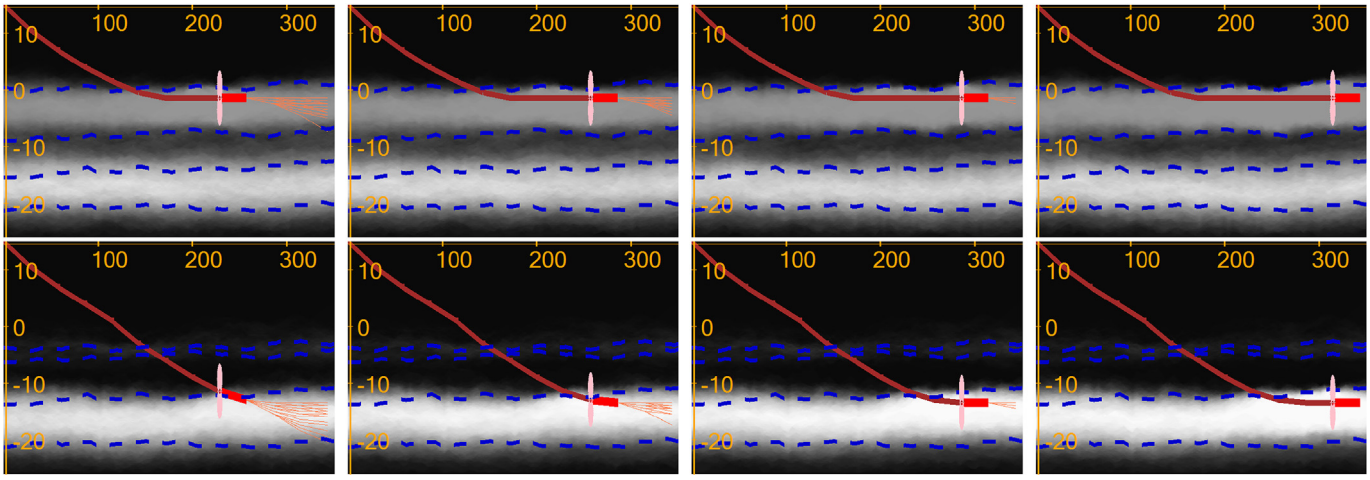

Fig. 9. Continued demonstration of the decision support system for the two synthetic scenarios from Fig. 8. The figures show the step-by-step outputs of the DSS.

optimal. This can be seen by visual comparison of the actually drilled well path to the left and the optimal well path to the right. The actual well achieves a value of 95.56 equivalent meters of reservoir sands drilled (with respect to objective (14)), which corresponds to 86.6% of the theoretically possible. Since much of the well length in our paper is used for landing, we cannot use a standard reservoir contact metric, but instead compare the value the theoretical maximum. The theoretical maximum is the value that the well can achieve when the trajectory is optimized with respect to the known true model (without uncertainty). Obtaining complete information of the subsurface is not possible, therefore the theoretical 100% will never be accomplished in practice.

The bottom scenario was generated with different parameters for the layer boundaries yielding a thinner top layer. Thus it is a less likely scenario with respect to the pre-drill geomodel. However, as the real-time data, which indicates a thinner top layer than initially estimated, becomes available, the geomodel uncertainty is updated, the well path is corrected and the optimal target is reached. The well trajectory resulting from the bottom scenario has a value of 49.81 equivalent meters of sand drilled, which is approximately 58.6% of the theoretically possible value for that scenario. Note that new developments improving the pre-drill model or the look-around/look-ahead capability (such as in Constable et al. (2016)) will improve the decision outcomes provided by DSS.

Finally, we note that for both scenarios the reservoir boundaries are automatically mapped along the wellpath. The updated uncertainty estimations can thus benefit the further reservoir development planning.

4.2. Statistical analysis of the performance of the DSS

In the first numerical example we presented two situations where the DSS successfully landed the well in the optimal layer despite the initial uncertainty. Obviously, due to both the uncertainty of the subsurface interpretation and the simplifications in the DSS' "naive" algorithm, such good results would not be achieved for all cases. In this example we investigate the statistical performance of the DSS. To do so, we run 100 different synthetic geosteering cases. Their synthetic truths are drawn from the same distribution as used for the model realizations. For all 100 cases, we follow the recommendations of the DSS with the same objectives as in the previous example (see (14) in the Appendix).

To evaluate the DSS performance, we will look at two metrics;

1. What is the value of the well resulting from the recommendations compared to the optimal well that is based on perfect information?
2. Did the DSS land in the optimal layer?

Both metrics are case-specific. To compare them to each other, we are evaluating them relative to the optimal well trajectory computed for the synthetic truth for each particular case (similar to Fig. 6). The value achieved by the well optimized with respect to the deterministic

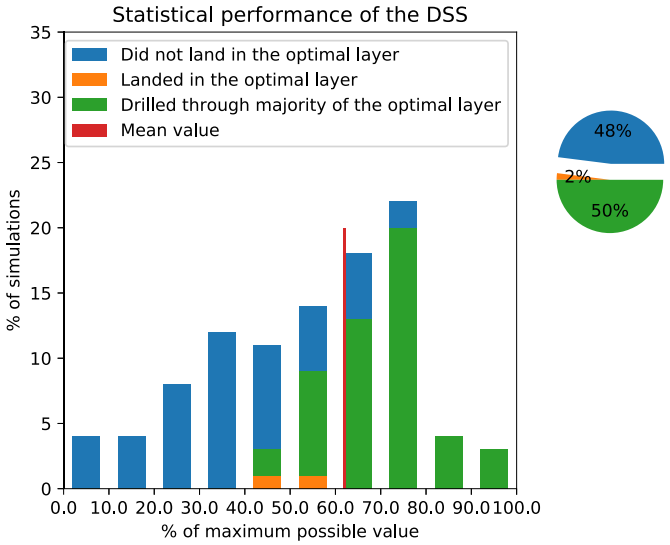

Fig. 10. The results showing the statistical performance of the DSS. The cases are grouped in bins by the percentage of the theoretical maximum value achieved. The value of a resulting well is on average higher than 60% of the case-specific theoretical maximum. Good results are achieved even for scenarios where the well is landed in a sub-optimal layer.

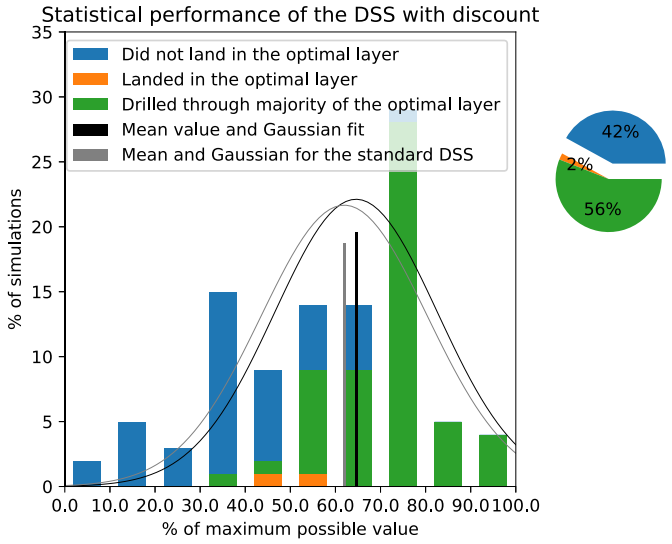

Fig. 11. The results showing the statistical performance of the DSS with the discount factor of 0.9 to compensate for future learning. The correction for future learning allows to achieve better results than in Fig. 10 (indicated in gray).

synthetic truth is set to 100%. The second metric is subjective as it is not directly included in the objective function, but it gives an intuitive operational understanding of the DSS' performance. We let the 'successful landing' criterion be defined as drilling two complete drill-stands in the target layer (see bottom row in Fig. 9). As in the previous example, we finish the drilling operation after having drilled 350 m in horizontal direction.

There are two challenges related to our objectives. For cases where the top layer is optimal, coming in at a low angle might result in overshooting the sweet spot which gives double the value. For cases where the bottom layer should be preferred, the challenge is to realize early enough that the top layer is thin and drop the angle to get good coverage of the bottom layer.

The statistics of the DSS performance is summarised on Fig. 10. The pie chart indicates that the well is landed in the optimal layer in 52% of the cases. Among those, in only 2% of the cases the well length within the optimal layer is less than half of the maximum possible length.

We emphasize that the choice of layer was not explicitly included in the objective function. Thus, a fair performance evaluation is based on the value of the resulting wells. This is summarised on the bar plot in Fig. 10. The bars indicate the percentage bins of the maximum possible well value achieved by the DSS in each of the 100 different cases. The bars are split by the choice of layer for the resulting well. Not surprisingly, most of the better results correspond to the wells that landed in the optimal layer. At the same time, in the challenging geological conditions of the chosen setup, choosing the sand layer which is not optimal could result in over 70% of the maximum possible well value (see blue squares in Fig. 10).

The average value of the wells drilled by the DSS is 62% of the theoretically possible value, see Fig. 10. We reiterate that the 100% value is impossible in a non-synthetic case since it requires knowledge of the true subsurface.

4.3. Statistical performance of the DSS with a discount factor

Based on the theory described in section 3.2, we know that the naive decision policy used for the DSS might not give optimal results as it neglects the modelling of the future learning (Alyaev et al., 2018b). One practical method to compensate for this is using a discount factor for future value $\gamma < 1$ in the optimization equations (3) and (6). The effect of a discount factor is hard to analyse theoretically, therefore we try to apply a naive decision policy with a discount in the following numerical example.

Here we set $\gamma = 0.9$ while keeping the rest of the parameters the same as in the previous example (Section 4.2). The statistical results achieved by the modified DSS are presented in Fig. 11. Compared to the DSS without the discount factor, the average value achieved is increased from 62% to 64.6% with similar spread in the fitted Gaussian distribution. We also observe the increase of the "optimal" landings by 6% compared to Fig. 10.

This example indicates that the DSS performance can be improved by introducing a discount factor for future well value. We expect that the optimal choice of the discount factor would depend on the selected application case of the DSS. Thus, the purpose of this example is to indicate this practical option in the DSS and not to find the optimal value of γ.

4.4. Adjusting objectives due to insights

In the final example we demonstrate the flexibility of the DSS by changing the weights of the different objectives. This might be important during adoption in the field, as the insights gained during an operation might change the prioritization of the geosteering objectives. The main reason for changing the objectives and their weighting is the fact that for geosteering operations, the objectives are often simplified to ensure the possibility to evaluate them in real time. As more insight is gained, it may become clear either via expert judgment or various types of calculations or simulations that these simplifications and the initial weights may not yield the best decision suggestions (the specific considerations made to obtain better objectives and weighting is not in the scope of this paper). The objectives and their weights can be adjusted in the user interface of the DSS.

In this example we recall the operation described in the top scenario of Figs. 7–9. Here we consider a slightly different drilling scenario, but using the same setup including the synthetic truth. In column 2 in Fig. 8 the bit enters the top sand layer, which should result in a landing as shown in Figs. 8 and 9. For the sake of argument, let us assume that real-time measurements indicate that the top sand is of poor quality. Thus, the geosteering experts take the decision to prioritize sand quality, which was not part of the original objective. Consequently, the weight of staying in the top part of the reservoir is decreased to 0.3 and

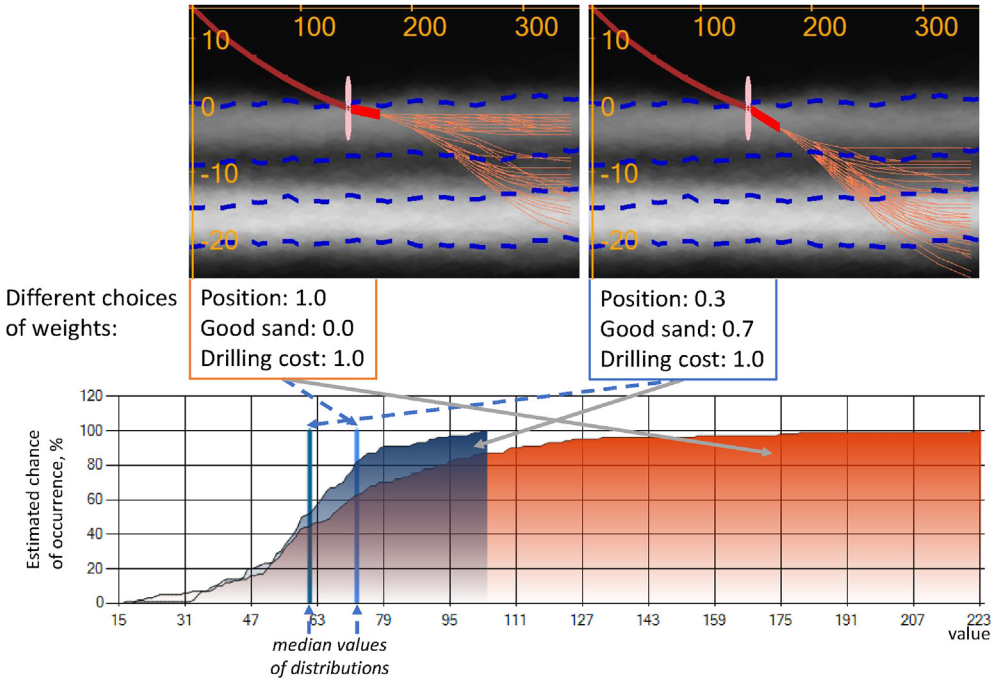

Fig. 12. Example of how the weights of the objectives can be adjusted in the middle of the geosteering operation, in accordance with the insights gained while drilling (see the Appendix for a detailed description of the objectives). The ensemble of realizations is the same in both scenarios. The bottom plot compares the cumulative distribution of the resulting multi-objective value functions for the choices of the weights (the vertical lines indicate the mean expected value for each distribution).

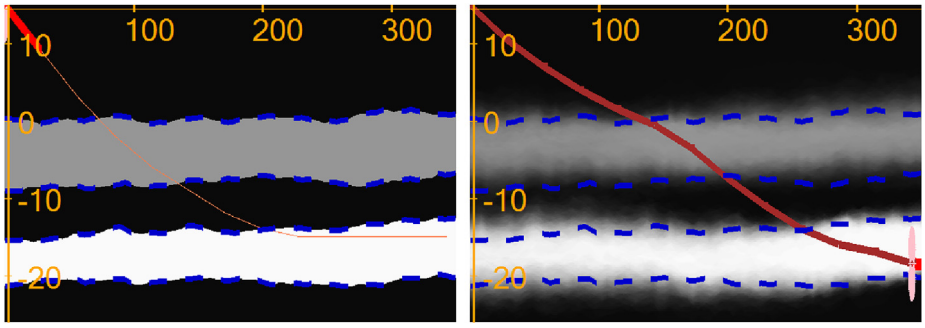

Fig. 13. The final trajectory (right) for the scenario where the weights have been changed to pursue the layer with better quality sand following Fig. 12 (middle). For comparison, the optimal trajectory for the synthetic truth is shown to the left. The latter is computed applying the new objectives (15) from the start of the operation.

the weight of the sand quality objective is introduced and set to 0.7, resulting in a new global objective function described by (15) (see the Appendix).

The newly selected weights can be applied to the trajectory optimization in real-time. The expected outcomes are shown in Fig. 12. The cumulative value diagram shown in the figure clearly indicates the reduction of the expected well value after the new objective function is chosen. The alteration in the objective results in a landing in the lower reservoir layer, as shown in Fig. 13. Comparison of the optimal trajectories with their actual outcomes in Fig. 6 (bottom) and 13 gives a visual proof that superior results can be achieved when the 'correct' objective is selected before the start of an operation.

5. Conclusions

In this work we have presented a functioning consistent decision support system (DSS) aiming at supporting real-time geosteering decisions. The DSS provides directional drilling decision support information and recommendations. The recommendations account for real-time measurements behind and around the bit, inferred uncertainties ahead-of-the-bit, and multiple objectives. The system includes a visual display that allows the geosteering team to inspect uncertainties and immediately see and evaluate the possible results of their decisions. In contrast to basing geosteering decisions on "educated guesses" about the geological interpretation, future profit and drilling costs, the DSS provides a consistent Bayesian framework for making ahead-of-the-bit inferences based on prior information and learning (from real-time data) while drilling.

The workflow implementation presented here can consistently update uncertainties ahead of the drill bit and provides a visual and interactive means to inspect the resulting multi-realization model of the subsurface. However, this uncertainty quantification is not an end by itself. Rather, the goal is to make good geosteering decisions, which requires an assessment of relevant and material uncertainties. An essential part of improving the results from geosteering operations is to move the focus away from real-time data to actual decisions (Kullawan et al., 2014). The DSS uses real-time data gathering and learning-while-drilling to optimize key drilling decisions, thus ensuring good utilization of new measurement technologies.

There is abundant research and literature demonstrating that people are exceptionally bad at making decisions in complex and uncertain environments (see e.g. Tversky and Kahneman (1974)).[6] The DSS embeds a consistent uncertainty quantification and a sophisticated

[6] The most recent Nobel Memorial Prize in Economic Sciences was given to Richard Thaler for his contributions to behavioral economics. His former collaborator, Daniel Kahneman was awarded the same prize in 2002 for his work on the psychology of judgment and decision-making.

decision-making process, and is particularly advantageous for unbiased high-quality decision support when navigation in complex reservoirs with several potential targets and significant interpretation uncertainty. The real-time performance of the system is of major importance for geosteering where time for evaluation, re-consideration and decision-making is scarce.

To illustrate the benefits of the DSS we have presented synthetic cases with multiple objectives, for which the full workflow consisting of the model updating and the decision recommendations was applied. The system demonstrated landing in the reservoir, automatic choice of target in a multi-target geological scenario, and navigating the well in a layer-cake geological configuration. Results were consistently achieved in several distinct scenarios as well as in a statistical test. Statistically, the system-recommended decisions are initially achieving more than 60% of the theoretically possible well value despite the uncertainty in the pre-drill and while-drilling geological interpretations. We expect that the performance will increase with future improvements of the system.

Moreover, we have illustrated the flexibility of the implementation of the DSS. The possibility to introduce correction for future learning gives a further average improvement of 2.6% for the statistical performance of the system in our examples. There is also a flexibility when it comes to adjusting decision objectives. By design the DSS reacts to changes in the objectives and constraints within seconds, providing unbiased decisions for the modified choices.

This paper presents proof-of-concept testing of the DSS. The system uses existing measurement and modelling tools and identifies the optimal decisions through multi-objective optimization under uncertainty. It can be naturally extended to the advanced measurement technologies used in the field as well as include more realistic geology in its multi-realization geomodel. With that, we see testing on historical operations as a possibility in the nearest future.

Funding

This work was supported by the research project 'Geosteering for Improved Oil Recovery' (NFR-Petromaks2 project no. 268122) which is funded by the Research Council of Norway, Aker BP, Vår Energi, Equinor and Baker Hughes Norway. Aojie Hong is supported by the 'DIGIRES' project (NFR-Petromaks2 project no. 280473) which is funded by industry partners Aker-BP, DEA, Vår Energi, Petrobras, Equinor, Lundin and VNG, Neptune Energy as well as the Research Council of Norway.

Acknowledgement

The authors thank Eric Cayeux and Erlend Vefring for insightful suggestions during the paper preparation.

Appendix. formal definition of objective functions and constraints

The appendix describes the objectives and constraints used in the paper. First we define individual constraints and objective functions that have been used in the numerical examples. Thereafter, we define the weights that are selected to form the global objective function (2).

Constraints

In all examples we use the following two constraints:

1. The trajectory is constrained to a dogleg severity of maximum 2° between decision points (approximately every 29 m), which is approximated as

$$|\alpha_{i_{k+1}} - \alpha_{i_k}| \leq 2 \text{ deg}, \tag{7}$$

where α_{i_k} and $\alpha_{i_{k+1}}$ are the inclinations of the well in two consecutive segments along the well trajectory. The segments are assumed to be linear yielding a piece-wise linear trajectory.
2. The inclination is limited to 90°, which is a normal constraint to avoid problems with gravel packing;

$$\alpha_i \leq 90 \text{ deg}. \tag{8}$$

Individual objectives

The DSS described in the paper has a simple application programming interface (API) which allows to add new objectives pragmatically. It is possible to add objectives that give value for a given point or segment of the well. It is also possible to add objectives as a function of two consecutive segments which, for example, allows to set the cost of bending the well trajectory. The parameters of the implemented objectives can be changed through the graphical interface.

To simplify the communication in this paper, instead of using conversions to currency, we express the value in equivalent number of stands drilled. One unit is equivalent to an expected net present value from a well segment positioned in a 1-m-thick reservoir layer with reference properties. In this subsection we only list the objectives that were used in the numerical examples (see Fig. 12).

1. **Position** in a sand layer. This objective is defined for a well segment and gives value proportional to the thickness $h(x)$ of a sand layer. The value is doubled if the well is positioned in the "sweet spot" for production (in our examples between 0.75 and 2.25 m from the sand roof). This can be formally written as:

$$O_p(x_0, z_0, x_1, z_1) = \frac{1}{\delta x_{decision}} \int_{x_0}^{x_1} F_p(h(x))dx, \tag{9}$$

$$F_p(h(x)) = \begin{cases} 0, & \text{well is outside reservoir;} \\ 2h(x), & 0.75 \le z_{roof}(x) \le 2.25\text{: well in "sweet spot";} \\ h(x), & \text{otherwise;} \end{cases} \tag{10}$$

where (x_0, z_0) and (x_1, z_1) are the start and end of a well segment, $z_{roof}(x)$ is the distance from the roof of the reservoir layer to the well, and $\delta x_{decision}$ is the distance between decision points (reference length of a well stand projected to horizontal axis) equal to 28.56 m used to normalize the value to the chosen scale. The integral in equation (9) is evaluated numerically using mid-point rule quadrature.

2. **Good sand.** This objective gives a value which depends on the sand quality and can be formally written as:

$$O_s(x_0, z_0, x_1, z_1) = \frac{1}{\delta x_{decision}} \int_{x_0}^{x_1} F_s(x, z(x))dx \tag{11}$$

$$F_s(x, z(x)) = \begin{cases} 0, & \text{well is outside reservoir;} \\ 7, & \text{well is in the top reservoir layer;} \\ 14, & \text{well is in the bottom reservoir layer.} \end{cases} \tag{12}$$

The value 7 is similar to a reference reservoir thickness from (10). Here, we use a constant value rather than a thickness to highlight that objectives might be expressed differently depending on user preferences.

3. **Drilling cost.** The drilling cost objective assigns a cost to drilling the well. When not weighted, the cost of drilling 1 m is proportional to one unit, i.e. assumed net present value from a 1-m-long well in a reservoir of 1-m thickness. The objective function can be written as follows:

$$O_d(x_0, z_0, x_1, z_1) = \int_{(x_0,z_0)}^{(x_1,z_1)} -0.003\, ds, \tag{13}$$

where $\int_{(x_0,z_0)}^{(x_1,z_1)} ds$ is an integral along a well segment. Notice that the drilling cost has negative value. With the default scaling (0.003 in (13)) the drilling cost of a stand accounts for approximately 8.6% of the production potential from a 1-m-thick sand layer.

The prmary set of objectives

In most of the examples we are using the Position objective function combined with the drilling cost. The global objective is written as:

$$O(X|M) = 1.0\, O_p(X|M) + 1.0\, O_d(X|M), \tag{14}$$

where objectives O_p and O_d are defined by equations (9) and (13) respectively.

The alternative weighting of objectives used in Section 4.4

In the example from Section 4.4, together with the primary objective defined by (9) we consider an alternative weighting of the objective functions:

$$O(X|M) = 0.3\, O_p(X|M) + 0.7\, O_s(X|M) + 1.0\, O_d(X|M), \tag{15}$$

where the sub-objectives are defined in equations 9–13.

References

Aanonsen, S.I., Nævdal, G., Oliver, D.S., Reynolds, A.C., Vallès, B., 2009. The ensemble kalman filter in reservoir engineering–a review. SPE J. 14 (3), 393–412. https://doi.org/10.2118/117274-PA.

Al-Fawwaz, A., Al-Yosef, O., Al-Qudaihy, D., Al-Shobaili, Y., Al-Faraj, H., Maeso, C., Roberts, I., 2004. Increased net to gross ratio as the result of an advanced well placement process utilizing real-time density images. In: IADC/SPE Asia Pacific Drilling Technology Conference and Exhibition. Society of Petroleum Engineers, pp. 1–10. https://doi.org/10.2523/87979-MS.

Alyaev, S., Bratvold, R.B., Luo, X., Suter, E., Vefring, E.H., 2018a. An interactive decision support system for geosteering operations. In: SPE Norway One Day Seminar, 18 April, Bergen, Norway, pp. 17. https://doi.org/10.2118/191337-MS.

Alyaev, S., Hong, A., B Bratvold, R., 2018b. Are you myopic, naïve or farsighted about your geosteering decisions? In: Second EAGE/SPE Geosteering and Well Placement Workshop, pp. 1–4. http://www.earthdoc.org/publication/publicationdetails/?publication=95302 doi:10.3997/2214-4609.201803217.

Alyaev, S., Suter, E., B Bratvold, R., H Vefring, E., Jahani, N., 2018c. Ensemble-based decision-support workflow for operational geosteering. In: Second EAGE/SPE Geosteering and Well Placement Workshop, pp. 1–4. http://www.earthdoc.org/publication/publicationdetails/?publication=95144 doi:10.3997/2214-4609.201803140.

Antonsen, F., Teixeira De Oliveira, M.E., Petersen, S.A., Metcalfe, R.W., Hermanrud, K., Constable, M.V., Boyle, C.T., Eliassen, H.E., Salim, D., Seydoux, J., Omeragic, D., Thiel, M., Denichou, J.M., Etchebes, M., Nickel, M., 2018a. Geosteering in complex mature fields through integration of 3d multi-scale LWD-data, geomodels, surface and time-lapse seismic. In: SPWLA 59th Annual Logging Symposium, 2-6 June. Society of Petrophysicists and Well-Log Analysts, London, UK, pp. 1–16.

Antonsen, F., Teixeira De Oliveira, M.E., Petersen, S.A., Metcalfe, R.W., Hermanrud, K., Constable, M.V., Boyle, C.T., Eliassen, H.E., Salim, D., Seydoux, J., Omeragic, D., Thiel, M., Denichou, J.M., Etchebes, M., Nickel, M., 2018b. Ultra-deep LWD-resistivity is essential to map reservoir structure and fluid contacts in a two-branched horizontal producer. In: Norway One Day Seminar, 18 April, Bergen, Norway. Society of Petroleum Engineers, pp. 1–21. https://doi.org/10.2118/191352-MS.

Arata, F., Mele, M., Tarchiani, C., Tosi, G., Chinellato, F., Nickel, M., Athmer, W., Salim, D.C., Dahl, G.V., 2017. Look ahead geosteering via real time integration of logging while drilling measurements with surface seismic. In: SPE Annual Technical Conference and Exhibition, 9-11 October, San Antonio, Texas, USA. Society of Petroleum Engineers, pp. 1–12. https://doi.org/10.2118/187203-MS.

Barros, E.G.D., Leeuwenburgh, O., Van den Hof, P.M.J., Jansen, J.D., 2015. value of multiple production measurements and water front tracking in closed-loop reservoir management. In: SPE Reservoir Characterisation and Simulation Conference and Exhibition, https://doi.org/10.2118/175608-MS.

Bashir, F., Al-Hawi, M., Bhuana, I., Abbas, M., Ghamdi, Y., 2016. Real time 3D modeling to optimize geosteering in clastic reservoir - case study. In: OTC Asia, 22-25 March, Kuala Lumpur, Malaysia. Offshore Technology Conference, pp. 1–10. https://doi.org/10.4043/26635-MS.

Bø, Ø., Vikhamar, P., Conocophillips, Spotkaeff, M., Dolan, J., Wang, H., Dupuis, C., Ceyhan, A., Blackburn, J., Perna, F., Sclumberger, 2014. Shine a light in dark places: using deep directional resistivity to locate water movement in Norway's oldest field. In: SPWLA 55th Annual Logging Symposium. SPWLA, pp. 1–13.

Bratvold, R.B., Begg, S., 2010. Making Good Decisions, vol. 207 Society of Petroleum Engineers Richardson, Texas.

Brown, D.B., Smith, J.E., 2013. Optimal sequential exploration: bandits, clairvoyants, and wildcats. Oper. Res. 61 (3), 644–665. http://search.ebscohost.com/login.aspx?

direct=true&db=bth&AN=89237158&site=ehost-live&scope=site doi:10.1287/opre.2013.1164.
Burgers, G., Jan van Leeuwen, P., Evensen, G., 1998. Analysis scheme in the ensemble kalman filter. Mon. Weather Rev. 126 (6), 1719–1724.
Cayeux, E., Skadsem, H., Carlsen, L., Stokland, L., Cruikshank, S., 2018. Analysis of asymmetric tool-joint wear while drilling long horizontal sections. In: SPE Norway One Day Seminar. SPE, pp. 1–24.
Chen, Y., Lorentzen, R.J., Vefring, E.H., 2015. Optimization of well trajectory under uncertainty for proactive geosteering. SPE J. 20 (02), 368–383. http://www.onepetro.org/doi/10.2118/172497-PA doi:10.2118/172497-PA.
Chen, Y., Oliver, D.S., Zhang, D., 2009. Efficient ensemble-based closed-loop production optimization. SPE J. 14 (4), 634–645. https://doi.org/10.2118/112873-MS.
Clemen, R., Reilly, T., 2013. Making Hard Decisions with DecisionTools. Cengage Learning. https://books.google.no/books?id=kiawMQEACAAJ.
Constable, M.V., Antonsen, F., Stalheim, S.O., Olsen, P.A., Fjell, Ø.Z., Dray, N., Eikenes, S., Aarflot, H., Haldorsen, K., Digranes, G., 2016. Looking ahead of the bit while drilling: from vision to reality. In: SPWLA 57th Annual Logging Symposium. Society of Petrophysicists and Well-Log Analysts, pp. 1–21.
Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C., 2009. Introduction to Algorithms, third ed. The MIT Press 3rd ed.
Cressie, N., 1992. Statistics for spatial data. Terra. Nova 4 (5), 613–617.
Dupuis, C., Denichou, Jm, 2015. Automatic inversion of deep-directional-resistivity measurements for well placement and reservoir description. Lead. Edge 34 (5), 504–512. http://library.seg.org/doi/abs/10.1190/tle34050504.1 doi:10.1190/tle34050504.1.
Dupuis, C., Omeragic, D., Chen, Y.H., Habashy, T., 2014. Inversion-based workflow to image faults crossed by the wellbore using deep directional resistivity provides new way of understanding complex formations. In: SPWLA 55th Annual Logging Symposium, . https://www.onepetro.org/conference-paper/SPWLA-2014-WWW.
Evensen, G., 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. J. Geophys. Res.: Oceans 99 (C5), 10143–10162.
Evensen, G., 2009. Data Assimilation: the Ensemble Kalman Filter. Springer Science & Business Media.
Feinberg, E.A., Shwartz, A., 2012. Handbook of Markov Decision Processes: Methods and Applications, vol. 40 Springer Science & Business Media.
Guevara, A.I., Sandoval, J., Guerrero, M., Manrique, C.A., 2012. Milestone in production using proactive azimuthal deep-resistivity sensor combined with advanced geosteering techniques: tarapoa block, Ecuador. In: SPE Latin America and Caribbean Petroleum Engineering Conference. Society of Petroleum Engineers, pp. 1–13.
Hanea, R., Evensen, G., Hustoft, L., Ek, T., Chitu, A., Wilschut, F., 2015. Reservoir management under geological uncertainty using fast model update. In: SPE Reservoir Simulation Symposium. Society of Petroleum Engineers, pp. 1–15. https://doi.org/10.2118/173305-MS.
Hartmann, A., Vianna, A., Maurer, Hm, Sviridov, M., Martakov, S., Hughes, B., Antonsen, F., Olsen, P.A., Constable, M.V., 2014. Verification testing of a new extra-deep azimuthal resistivity measurement. In: SPWLA 55th Annual Logging Symposium. SPWLA, pp. 1–12.
Hongsheng, L., Guofeng, X., Daicai, X., Xuekai, L., Calleja-Djefel, B., Youquan, L., 2016. Advanced geosteering technology promotes the accuracy of well placement in the thin reservoirs of China mature oilfield. In: SPE/IADC Middle East Drilling Technology Conference and Exhibition, . http://www.onepetro.org/doi/10.2118/178167-MS doi:10.2118/178167-MS.
Howard, R., Abbas, A., 2015. Foundations of Decision Analysis. Pearson. https://books.google.no/books?id=TuSAAAAACAAJ.
Janwadkar, S., Thomas, M., Privott, S., Tehan, R., Carlson, L., Spear, W., Setiadarma, A., et al., 2012. Reservoir-navigation system and drilling technology maximize productivity and drilling performance in the granite wash, us midcontinent. SPE Drill. Complet. 27 (01), 22–31.
Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. J. Basic Eng. 82 (1), 35–45.
Kullawan, K., Bratvold, R., Bickel, J., 2016. Value creation with multi-criteria decision making in geosteering operations. International Journal of Petroleum Technology 3 (1), 15–31. http://www.avantipublishers.com/downloads/ijptv3n1a2/ doi:10.15377/2409-787X.2016.03.01.2.
Kullawan, K., Bratvold, R., Bickel, J.E., 2014. A decision analytic approach to geosteering operations. SPE Drill. Complet. 29 (01), 36–46. http://www.onepetro.org/doi/10.2118/167433-PA doi:10.2118/167433-PA.
Kullawan, K., Bratvold, R.B., Bickel, J.E., 2018. Sequential geosteering decisions for optimization of real-time well placement. J. Pet. Sci. Eng. 165 (January), 90–104. https://doi.org/10.1016/j.petrol.2018.01.068 doi:10.1016/j.petrol.2018.01.068.
Kullawan, K., Bratvold, R.B., Nieto, C.M., 2017. Decision-oriented geosteering and the value of look-ahead information: a case-based study. SPE J. 22 (3), 767–782. https://doi.org/10.2118/184392-PA.
Lorentzen, R., Luo, X., Bhakta, T., Valestrand, R., 2019. History matching the full Norne field model using seismic and production data. SPE J. 24 (4). https://doi.org/10.2118/194205-PA.
Lorentzen, R.J., Berg, A., Nævdal, G., Vefring, E.H., 2006. A new approach for dynamic optimization of water flooding problems. In: Intelligent Energy Conference and Exhibition. Society of Petroleum Engineers, pp. 1–11. https://doi.org/10.2118/99690-MS.
Lorentzen, R.J., Stordal, A.S., Hewitt, N., 2017. An auxiliary adaptive Gaussian mixture filter applied to flowrate allocation using real data from a multiphase producer. Comput. Geosci. 102, 34–44.
Lorentzen, R.J., Stordal, A.S., Luo, X., Naevdal, G., 2016. Estimation of production rates by use of transient well-flow modeling and the auxiliary particle filter: full-scale applications. SPE Prod. Oper. 31 (2), 163–175. https://doi.org/10.2118/176033-PA.
Luo, X., Eliasson, P., Alyaev, S., Romdhane, A., Suter, E., Querendez, E., Vefring, E., 2015. An ensemble-based framework for proactive geosteering. In: SPWLA 56th Annual Logging Symposium, July 18-22, vol. 2015. pp. 1–14.
Meinhold, R.J., Singpurwalla, N.D., 1983. Understanding the kalman filter. Am. Stat. 37 (2), 123–127.
Norwegian Petroleum Directorate, 2018. Mange Utbygginger Og Høy Leteaktivitet. http://www.npd.no/no/Nyheter/Nyheter/2018/Mange-utbygginger-og-hoy-leteaktivitet/, Accessed date: 5 March 2019.
Payrazyan, V., Kuvaev, I., Uvarov, I., Aguilar, R., Gutierrez, G., 2017. Geoscience evolution: extensive data integration for real time geosteering and modeling in unconventional reservoirs. In: AAPG Eastern Section 46th Annual Meeting, pp. 5. https://doi.org/10.3997/2214-4609.201702633.
Saputelli, L., Economides, M., Nikolaou, M., Kelessidis, V., 2003. Real-time decision-making for value creation while drilling. In: SPE/IADC Middle East Drilling Technology Conference and Exhibition. Society of Petroleum Engineers, pp. 1–19. https://doi.org/10.2118/85314-MS.
Selheim, N., Morris, S., Jonsbraaten, F., Aarnes, I., Teelken, R., 2017. Geosteering and mapping of complex reservoir boundaries using an integrated data approach. In: SPE Annual Technical Conference and Exhibition. Society of Petroleum Engineers, pp. 1–10. https://doi.org/10.2118/187136-MS.
Seydoux, J., Legendre, E., Mirto, E., Dupuis, C., Denichou, J.M., Bennett, N., Kutiev, G., Kuchenbecker, M., Morriss, C., Schlumberger, L.Y., 2014. Full 3D deep directional resistivity measurements optimize well placement and provide reservoir-scale imaging while drilling. In: SPWLA 55th Annual Logging Symposium, pp. 14.
Skjervheim, J., Hanea, R., Evensen, G., 2015. Fast model update coupled to an ensemble based closed loop reservoir management. In: Petroleum Geostatistics, vol. 2015. pp. 1–4. https://doi.org/10.3997/2214-4609.201413629.
Skjervheim, J.A., Evensen, G., 2011. An ensemble smoother for assisted history matching. In: SPE Reservoir Simulation Symposium. Society of Petroleum Engineers, pp. 1–15. https://doi.org/10.2118/141929-MS.
Sviridov, M., Mosin, A., Antonov, Y., Nikitenko, M., Martakov, S., Rabinovich, M.B., Hughes, Baker, B.P., 2014. New software for processing of LWD extradeep resistivity and azimuthal resistivity data. SPE Reserv. Eval. Eng. 17 (May), 109–127. https://doi.org/10.2118/160257-PA.
Tosi, G., Wang, H., Denichou, J.M., 2017. Application of deep directional resistivity reservoir mapping tool on arctic field development - case studies from the goliat field in the barents sea. In: SPE/IADC Drilling Conference and Exhibition, 14-16 March, the Hague, The Netherlands. Society of Petroleum Engineers, pp. 1–14. https://doi.org/10.2118/184708-MS.
Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. Science 185 (4157), 1124–1131.