Efficient well placement optimization under uncertainty using a virtual drilling procedure

Brage S. Kristoffersen¹ · Thiago L. Silva^{1,2} · Mathias C. Bellout¹ · Carl Fredrik Berg¹

Received: 15 January 2021 / Accepted: 27 September 2021 / Published online: 25 October 2021
© The Author(s) 2021

Abstract

An Automatic Well Planner (AWP) is used to efficiently adjust pre-determined well paths to honor near-well properties and increase overall production. AWP replicates modern geosteering decision-making where adjustments to pre-programmed well paths are driven by continuous integration of data obtained from logging-while-drilling and look-ahead technology. In this work, AWP is combined into a robust optimization scheme to develop trajectories that follow reservoir properties in a more realistic manner compared to common well representations for optimization purposes. Core AWP operation relies on an artificial neural network coupled with a geology-based feedback mechanism. Specifically, for each well path candidate obtained from an outer-loop optimization procedure, AWP customizes trajectories according to the particular geological near-well properties of each realization in an ensemble of models. While well placement searches typically rely on linear well path representations, AWP develops customized trajectories by moving sequentially from heel to the toe. Analog to realistic drilling operations, AWP determines subsequent trajectory points by efficiently processing neighboring geological information. Studies are performed using the Olympus ensemble. AWP and the two derivative-free algorithms used in this work, Asynchronous Parallel Pattern Search (APPS) and Particle Swarm Optimization (PSO), are implemented using NTNU's open-source optimization framework FieldOpt. Results show that, with both APPS and PSO, the AWP solutions outperform the solutions obtained with a straight-line parameterization in all the three tested well placement optimization scenarios, which varied from the simplest scenario with a sole producer in a single-realization environment to a scenario with the full ensemble and multiple producers.

Keywords Machine learning · Artificial neural networks · Well placement optimization · Tailored trajectories · Reservoir simulation · Robust optimization · Ensemble optimization · Production optimization · Optimization with uncertainty

✉ Brage S. Kristoffersen
brage.s.kristoffersen@ntnu.no

Thiago L. Silva
thiago.silva@sintef.no

Mathias C. Bellout
mathias.bellout@ntnu.no

Carl Fredrik Berg
carl.f.berg@ntnu.no

¹ Department of Geoscience and Petroleum, Norwegian University of Science and Technology, Trondheim, Norway

² Department of Sustainable Energy Technology, SINTEF Industry, Trondheim, Norway

1 Introduction

There has been a paradigm shift in the management of hydrocarbon reservoirs. Decision support tasks such as modeling and sensitivity analysis are conducted within the time lag between data acquisition and implementation of operational changes. Earlier, these time lags were sufficiently long to allow for time-consuming semi-automated tasks. However, modern field operations have led to reduced time gaps, from frequent monitoring by permanent seismic sensors to logging while drilling. This shift in data acquisition has had an accompanied shift in computing power. Increased computational power has expanded the possibilities for automated workflows,

including automated history matching. The combined effect of increased and more frequent data acquisition together with an increase in computational power has paved the way for real-time closed-loop reservoir management [1, 20].

Another shift is the increased awareness and capabilities to quantify and characterize uncertainty within the subsurface. This shift expands on classical engineering principles, and often uses an ensemble to capture the uncertainties related to properties, structure, and related features. Geir Evensen's work within this domain has been fundamental for operators to change their best practices in daily reservoir engineering workflows [7, 12]. In this paradigm, the uncertainty in a reservoir is represented by an ensemble of multiple realizations [13, 30]. The ensemble may contain equiprobable realizations of the actual reservoir or different weighting factors may be assigned to different realizations [11]. As each realization is a plausible representation of reality, all novel information acquired during the development of a field needs to be assimilated into the set of models such that the alterations to former development plans can be implemented with increased reliability [6].

The uncertainty characterization described in terms of ensembles was firstly applied in history-matching methodologies [17, 24, 25] within closed-loop reservoir management workflows, see e.g. [16, 20, 32]. The focus was mainly on incorporating new information from direct and indirect sources such as seismic, pressure, and rate measurements, to accurately match historical data with numerical models. After history-matching the models, they can be used to forecast the production of the reservoir with a discrete representation of the uncertainty, namely the ensemble of updated models.

An approach to optimize the expected performance of the asset is known as robust optimization [8, 9, 28], which consists of maximizing the weighted average of the objective functions corresponding to the ensemble members. In robust optimization, the uncertainty is represented with a discrete approximation of the average performance. Thus, the number of function evaluations, i.e., the computational requirement of the method, increases proportionally to the size of the ensemble. This implies that, when geological uncertainty is represented by an ensemble, any increase in the computational complexity of the simulation model must be multiplied by the number of realizations. The overall complexity might become computationally intractable in the case of robust optimization of full-fledged models, where each cost function evaluation requires the simulation of the whole ensemble of models.

Well placement and control problems may be solved using derivative-free methodologies [3, 10]. The computational demand required by this type of methodologies commonly depends on the number of decision variables.

This usually limits the level of refinement, i.e., the number of degrees of freedom that can be imposed to treat the problem. Such limitations favor low-order well parameterizations such as the use of a straight-line for well trajectories and control steps spanning longer time periods. The computational demand due to problem dimension is particularly relevant for well control problems where the number of decision variables tends to be large, leading to computational challenges even on a single realization.

With regard to modern drilling operations, both the general capabilities of information acquisition [29] and the development of specialized tools with increased depth of investigation and look-ahead capabilities have improved significantly [23]. These advances have enabled operators to make decisions in real-time to alter the drilling direction by the continuous incorporation of new data into the decision process. This creates additional degrees of freedom for decisions in modern drilling routines, where information collected in the drilling process can significantly affect the optimized pre-drill plans, as these have been developed without the updated information provided by the logging-while-drilling tools.

These different factors influence how well placement optimization can be conducted. A standard way of performing well placement optimization over an ensemble of models would be to apply the same trajectory to all realizations. However, such an implementation does not discriminate with respect to realization-specific properties such as permeability and porosity distribution, or distance to boundary layers and faults in the near-well region. Moreover, given the initial starting point for the well and a general overall direction, with modern logging tools the well trajectory will be affected by local information obtained during the drilling process. This can be mimicked in ensemble well placement optimization by enabling for individual well trajectories in the different realizations.

Actually customizing well trajectories for each realization requires a highly-refined well parameterization that can enable numerous individual breakpoints to account for the different features of each realization. Obviously, this significantly increases the number of decision variables required for the well placement problem, making this approach infeasible when using typical derivative-free methods. An alternative is to develop tailored well trajectories by adaptive adjustment of well paths given the initial well starting point and overall direction. In this approach, the starting point and direction are determined by the optimization procedure, whereas the well description tailored for each realization is determined by a separate automatic procedure. By doing so, the optimization and simulation domains can have different scales with respect to well parameterizations, i.e., well descriptions can be coarser in the iterative layer that operates on optimization variables, while simulations are

performed using highly-refined, realization-specific, well trajectories [22].

This work combines the concepts of geological uncertainty captured by an ensemble of realizations with the main concepts of adaptive geology-based well path steering. In this implementation, the continuous incorporation of data gathered by the geosteering mechanism is provided by a fictive drilling operation. The proposed methodology, termed Automatic Well Planner (AWP), consists in refining a coarse well trajectory provided by the optimization routine. Analog to common geosteering operations, the refined new trajectory is generated through a step-wise integration of drilling information. The AWP relies on data-based machine learning techniques to make informed decisions on how to adapt the well trajectory to near-well geological features. Embedding the AWP within a robust optimization scheme gives the ability to adapt to different geological realization while maintaining a reduced number of well path variables, i.e., the same number of well path variables as a coarse idealized straight-line well. For single-realization well placement optimization, introducing the AWP within the optimization loop has been shown to improve cost function progression and to yield higher cost function values [22]. This work expands this approach for well placement optimization into a robust optimization scheme. The final optimization methodology is also expected to be more relevant for any realistic case, as the individual well paths mimic an actual drilling process which adapts to information acquired near the drill bit.

In Section 2 we presents the novel AWP method for well parameterization. In Section 3 is a description of the robust optimization scheme with AWP, in addition to the mathematical formulation of the overall optimization problem. In Section 4 presents the reservoir ensemble and the problem setup. In Section 5 we present a verification study of the parameterization. In Section 6 we highlight two case studies where optimization runs using AWP are compared against runs using a straight line well representation. Finally, in Section 7 we present the conclusions of the work.

2 Automatic well planner

AWP is an adaptive parameterization method capable of efficiently tailoring well trajectories for different geological realizations based on near-well properties. This section describes the overall functionality of the methodology. It starts by presenting an overview of the AWP, and continues with a brief description of its main components.

Note that the focus of this article is on extending the AWP concept to deal with geological uncertainty. Therefore, this section treats the key ingredients of the method that are

related to its incorporation into a robust well placement optimization scheme, and omits describing the training of the neural network. Moreover, the focus is on describing how the AWP translates the well heel and toe coordinates into a refined trajectory and its tailoring to the specific properties of each realization of an ensemble. The interested reader is referred to [22] for a detailed description of the AWP methodology.

2.1 Overall AWP procedure

Commonly, well placement searches are conducted using linear well path representations based on real-valued well heel and toe coordinates. Analog to drilling operations, the AWP determines a custom trajectory by adapting this path using a sequence of steps from the well heel towards the toe. This search is guided by a neural network which is trained to improve the production and well configurations based on surrounding geological information, e.g., higher permeability, porosity or oil saturation.

The overall AWP procedure is illustrated in the flow chart depicted in Fig. 1. During the initialization phase, the current location $\mathbf{x}_c$ is set to the well heel $\mathbf{x}_h$, analog to the starting position of drilling within the reservoir sands. The procedure continues with the data acquisition phase, in which information p is gathered from the surrounding region of the current location. The information is collected at points in a set of directions $\mathcal{S} = \{\mathbf{s}\}$ from the current

Fig. 1 A flow chart indicating the iterative AWP procedure to obtain the well path trajectory

location $\mathbf{x}_c$. The collected information is all within a distance d_m from the current location, thus $\|\mathbf{s}\| < d_m$ for all directions $\mathbf{s} \in \mathcal{S}$. The collected information $\{p(\mathbf{s} + \mathbf{x}_n) \mid \mathbf{s} \in \mathcal{S}\}$ is processed by a previously trained artificial neural network to decide the direction of the next line-segment. The procedure stops if the next location $\mathbf{x}_n$ is sufficiently close to the target, i.e., the well toe, or if any constraints would be violated when stepping to the next location. Otherwise, the simulated drill-bit is moved to the next position and the procedure continues iteratively until the stopping criteria is reached.

2.2 Properties

The virtual drilling process requires information about the reservoir surrounding the drill bit in order to make its decisions. The user defines which property or combination of properties should be collected and used for the decision process. We assume the spatial distributions of these properties are defined within the reservoir model as property maps $p(\mathbf{x})$. Property maps can be based on discrete grid-based properties, e.g., a permeability or porosity grid, or a continuous relationship; e.g., a geometric relationship like the height from the top of the reservoir.

2.3 Virtual drilling

The virtual drilling generates a well trajectory based on a pre-determined set of coordinates which represent the well start-point $\mathbf{x}_h$, i.e. the heel, and the well end-point, i.e., the toe, $\mathbf{x}_t$, of the perforated part of a well. We consider all well paths as locally straight, i.e., they are defined by a ordered set of positions $\{\mathbf{x}_h, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n\}$, where the consecutive positions define connected straight lines $\mathbf{x}_{i-1} - \mathbf{x}_i$. This type of well paths can alternatively be described in terms of a set of azimuth Φ , elevation Θ and lengths L , together with the initial $\mathbf{x}_h$ position, see e.g., [5, 14, 18].

The AWP procedure sequentially calculates new positions at each iteration i such that next direction given by the angles ϕ_i and θ_i are chosen by a trained neural network, while we keep the line segment length constant as l . Notice that the connected line segments described by $\Theta = \{\theta_i\}$, $\Phi = \{\phi_i\}$ and l from $\mathbf{x}_h$ have a unique mapping back to the ordered set of positions $\{\mathbf{x}_h, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n\}$. In other words, the AWP, here denoted by N , can be viewed as a mapping $N: \mathbf{x}_h, \mathbf{x}_t \rightarrow \{\mathbf{x}_h, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n\}$. A resulting well trajectory from the AWP procedure is illustrated in Fig. 2.

The goal of the AWP is to maximize the objective function

$$\sum \left(\frac{\sum_{\mathbf{x}_i \in N(\mathbf{x}_h, \mathbf{x}_t)} p(\mathbf{x}_i)}{\|N(\mathbf{x}_h, \mathbf{x}_t)\|} + b_t(\mathbf{x}_n) \right), \quad (1)$$

Fig. 2 This figure illustrates the AWP well trajectory (in red) vs. the well trajectory with the Straight-Line approach (in black). The blue circle indicates the target area. Brighter colors indicate high permeability good sands, while the darker colors indicate low permeability sands

where $\|N(\mathbf{x}_h, \mathbf{x}_t)\| = n$ represents the number of well segments. The outer sum runs over the training set of heel and toe coordinates in relevant reservoir models [22]. The term $b_t(\mathbf{x}_n)$ is defined as:

$$b_t(\mathbf{x}_n) = \begin{cases} R_g & \|\mathbf{x}_t - \mathbf{x}_n\| \leq t_g \\ (\alpha \cdot (\mathbf{x}_t - \mathbf{x}_n))^2 & \|\mathbf{x}_t - \mathbf{x}_n\| > t_g \end{cases}, \quad (2)$$

where R_g is a success reward, α is a weight of each direction, and n is the final position in the generated well path $N(\mathbf{x}_h)$. The function b_t conditions AWP solutions towards well paths that end up in the vicinity of $\mathbf{x}_t$.

In this work we keep the azimuth constant:

$$\phi_i = \phi = \arctan \left(\frac{x_t - x_h}{y_t - y_h} \right), \quad (3)$$

where $\mathbf{x}_h = (x_h, y_h, z_h)$ and $\mathbf{x}_t = (x_t, y_t, z_t)$. This means that the search is conducted in a two-dimensional space. The primary reason for omitting the azimuth in this application of the AWP procedure is related to the grid-block dimensions. The Olympus model have $50m \times 50m \times 3m$ grid blocks [15], providing our methodology with too coarse information in the horizontal direction relative to the vertical direction to properly adapt to the surrounding flow properties in the horizontal direction. We therefore omit the azimuth, which reduces the computational burden of training the AWP significantly. In the future it would be interesting to include the azimuth for models that have the necessary horizontal resolution (e.g., a geomodel) to affect the AWP decisions.

2.4 Data acquisition

The key operations of the AWP method are the data acquisition and the processing of the gathered information.

As described above, during the virtual drilling the data related to the property map p is collected from neighboring locations $\{\mathbf{x}_c + \mathbf{s} \mid \mathbf{s} \in \mathcal{S}\}$ of the drill bit location $\mathbf{x}_c$ within a given pre-determined radius d_m (set to 100 meters in this work). The neighboring locations are illustrated as blue dots in Fig. 3.

2.5 Artificial neural network

The AWP utilizes an artificial neural network (ANN) as a decision-maker. Since we keep the azimuth and length constant as ϕ and l , the only decision to be made by the neural network is the elevation angle θ_i of the next line-segment. This decision on the local elevation is bounded, as illustrated in Fig. 3(b), whereas Fig. 3(c) illustrates the former decisions made in the sequential procedure. Based on the information from the data acquisition process the ANN evaluates the current trajectory versus the new information. Based on this evaluation step, the artificial neural network gives a recommendation on whether to build or drop the elevation of the virtual drilling bit. This process is iterated for every 12 meters (the approximate length of drill pipe) until the virtual drilling simulation has achieved a full pass from heel to toe and a final trajectory is generated.

2.6 Run-time

When developing the AWP we need to train the decision-maker to be capable of providing the appropriate design

adjustments. The training of the ANN is the primary computational burden. Depending on computational resources, the training time ranges from minutes to hours. Once the ANN has been successfully trained, however, the AWP methodology typically designs well trajectories within milliseconds.

2.7 Constraints

The AWP method is a sequential decision-making procedure which is also subject to multiple constraints. There are restrictions on how far the calculated trajectory can deviate from the straight line between the input heel and toe locations, on the distance between the end point of the AWP well path and the input toe value, on the total well length, and on the dog-leg-severity. All of these constraint are formulated below. Importantly, since the ANN is trained with these constraints, these are embedded in the procedure and are therefore implicitly enforced. In the unlikely case that any of the constraints is violated, we revert from the AWP-calculated trajectory to the straight line representation.

2.7.1 Deviation constraint

A deviation constraint is imposed in order to prevent the new locations calculated by the AWP from deviating excessively from the well heel and toe coordinates in the original trajectory:

$$\frac{\|(\mathbf{x}_i - \mathbf{x}_h) \times (\mathbf{x}_t - \mathbf{x}_h)\|}{\|\mathbf{x}_t - \mathbf{x}_h\|} < t_d \quad \forall \mathbf{x}_i \in N(\mathbf{x}_h, \mathbf{x}_t) \quad . \quad (4)$$

Fig. 3 The well trajectory consisting of steps of equal length l , as generated by the AWP methodology. The set of sampling points for the next decision are indicated by blue circles in (a). Dog-leg severity bounds for the next decisions are indicated by cyan colored bars, enlarged in (b). The change in elevation at the previous decision is enlarged in (c)

The parameter t_d can be adjusted to allow for more or less flexibility in straying away from the straight line between the given heel and toe. In this work, the parameter t_d is set to 20 meters.

2.7.2 Length constraint

A length constraint factor β is imposed in order to prevent the length of the AWP well from increasing unrealistically:

$$\|N(\mathbf{x}_h, \mathbf{x}_t)\| < \frac{\|\mathbf{x}_t - \mathbf{x}_h\|}{l} \beta \quad (5)$$

Here the β parameter relates the number of segments $\|N(\mathbf{x}_h)\|$ in the AWP well path to the number of segments of length l in a straight line between the heel and toe. In this work we limited the AWP by setting β to 1.05.

2.7.3 Dog-leg-severity constraint

The dog-leg severity (DLS) is a measure of how much the well can change its trajectory over a given length. A bound constraint is imposed to ensure a realistic DLS for drilling operations:

$$\|\theta_i - \theta_{i-1}\| + \|\phi_i - \phi_{i-1}\| < \gamma l \quad (6)$$

where θ_i , θ_{i-1} , ϕ_i , and ϕ_{i-1} are the angles of consecutive steps. Since the azimuth is kept constant, the bound is applied to the elevation angle θ only. The parameter γ gives the DLS bound for allowed change in angle per well-length. In this work, the DLS bounds are given as 3 degrees per 30 meter.

2.7.4 Remedial actions

If the decision-maker in the AWP violates any of the constraints mentioned, remedial actions will be taken to

mitigate the impact. If, for instance, there are no active cells in the vicinity, the AWP will assign a negative cost function value to the solution, which will rule out the tentative solution from the candidates to the optimum. If the AWP violates a constraint while refining the trajectory, e.g., the deviation or length constraint, the well trajectory is reverted back to the (feasible) straight-line originally provided.

3 Robust well placement optimization with automatic well planning

In this section we describe how the AWP is combined with a well placement optimization algorithm in order to determine optimal well trajectories under geological uncertainty. The uncertainty is represented with an ensemble of equally probable geological realizations. These geological realizations have the same underlying grid structure, where this grid is populated with different property maps for the different realizations. For a given set of well heel and toe coordinates provided by the optimization procedure, the AWP develops customized well trajectories for each realization by incorporating properties, geological features and reservoir boundaries.

An illustration of the development of customized well trajectories is given in Fig. 4. In this figure, a set of realizations is subjected to the same heel and toe coordinates, and the AWP procedure generates an individual well path for each realization. The individual well paths are developed based on the selected geological properties of each ensemble member. The distance between the AWP and the underlying SL is limited by Eq. 4. Additionally, the ANN, which acts as a decision-maker in the AWP, was designed to return to the toe location when reaching the end of the well, as enforced by Eq. 2. The combined behavior is intended to allow the optimization layer to enforce control

Fig. 4 This figure illustrates the trajectories created for a subset of realizations in the ensemble given the same coordinates for the heel and toe. The black lines indicate the Straight-Line between $\mathbf{x}_h$ and

$\mathbf{x}_t$. The blue circle represents the target area given by Eq. 2. The red connected lines represent the new trajectories developed by AWP

over the vertical positioning of the start and end of the well trajectory. Still, the added flexibility of AWP allows the optimization procedure to be less sensitive with regards to heel and toe coordinates as the AWP automatically adjusts the well trajectories to better fit with the geology. In other words, we expect less simulations have to be performed surrounding any candidate solution while still maintaining a thorough evaluation of a particular region in the search space. Thus, this setup is expected to improve the overall performance of the optimization procedure because the well trajectories are adapted to the geology of each realization.

The well placement optimization algorithm employed in this work consists of three main parts: (1) an optimization layer, (2) a parameterization layer, and (3) a simulation layer. The outline of this scheme is illustrated by a flow chart in Fig. 5. The optimization layer iterates on heel $\mathbf{x}_{h,i}$ and toe $\mathbf{x}_{t,i}$ coordinates for all wells $i \in \{1, \dots, N\}$, such that the tuple $\mathbf{x} = (x_1, x_2, \dots, x_N)$ with $x_i = (\mathbf{x}_{h,i}, \mathbf{x}_{t,i})$ improves a cost function $\psi(\mathbf{x})$. The parameterization layer translates the optimization variables $\mathbf{x}$, i.e., the heel and toe coordinates, into customized well representations $P_{i,j}$ for each well i and each realization R_j . The realizations are members of the ensemble of size M , thus $j \in \{1, \dots, M\}$. The simulation layer requires running multiple reservoir simulations $\hat{s}_j$ using customized well representations calculated by the parameterization layer. After the simulation runs, the outputs f_j are used to calculate the expected value for the ensemble $\mathbb{E}(f_1, f_2, \dots, f_M)$, which is the cost function $\psi(\mathbf{x})$ to be optimized in the optimization layer.

Fig. 5 Scheme of the applied optimization algorithm incorporating the AWP method

3.1 Optimization layer

The optimization problem to be solved in the optimization layer can be formulated as:

$$P_{wp} : \max_{\mathbf{x}} \psi(\mathbf{x}) \quad (7a)$$

$$\text{s.t. :}$$

$$g(\mathbf{x}) \geq \mathbf{0} \quad (7b)$$

where P_{wp} is the well placement problem, $\mathbf{x} = (x_1, \dots, x_N)$ are the heel and toe coordinates of all wells, the cost function $\psi(\mathbf{x})$ is a measure of the profit to be maximized, and $g(\mathbf{x})$ are constraints on the well trajectories. Notice that the optimization procedure inherits the complexity of the simulation models, since the evaluation of the cost function ψ requires M simulations, one for each member of the ensemble. Further, the constraints $g(\mathbf{x}) \geq \mathbf{0}$ might require several simulations in a standard optimization scheme, but in our setup these constraints are largely handled implicitly by the AWP through the constraint-handling embedded in its training. By doing so, the optimal solution automatically satisfies the constraints without requiring additional simulations in the optimization procedure.

The well placement problem given by Eq. 7 is commonly solved using derivative-free procedures, although various approaches have been proposed to either directly [31] or indirectly [33] approximate gradient information. Derivative-free methods are typically used when gradients are not readily available or the response surface has a high number of local extrema. These methods have become computationally attractive in the past decades mainly because of the recent advances in distributed computing. In this work, we employ two derivative-free algorithms which take advantage of parallelization capabilities: Asynchronous Parallel Pattern Search (APPS) and Particle Swarm Optimization (PSO).

APPS [19, 21] is a deterministic, local optimization algorithm. It works by sampling the feasible space around a center point, at the extremum of a predefined pattern (step length). If an improvement is found, the incumbent point is moved to the best point and the step length may increase. Otherwise, the step length is reduced in order to converge.

PSO, on the other hand, is a stochastic, global, population-based optimization algorithm [26, 27]. PSO searches the feasible space using a population of solution candidates, known as particles. These are moved around in the search-space according to a simple mathematical formula determined by each particle's position, velocity, best previous position of the particle and the best location in the search space known to the population.

Both algorithms have parallel implementations available in the framework for field development optimization FieldOpt [2].

3.2 Parameterization layer

The parameterization layer translates the optimization variables $\mathbf{x}$ into simulation variables, namely the well trajectories $P_{i,j}$ for each well x_i and realization R_j . Instead of using straight lines for the well trajectories, the AWP calculates a set of elevation angles Θ , which together with the constant step lengths l and azimuth ϕ give an ordered set of positions X describing the well trajectory such that Eqs. 1, 2, 3, 4, and 5 are satisfied. The parameterization problem can be synthetically formulated as:

$$\{P_{i,j} = \text{AWP}(x_i, R_j) \mid i \in [1, N], j \in [1, M]\} \quad , \quad (8)$$

where $P_{i,j} = X_{i,j}$ is the trajectory of well i calculated by the AWP for realization R_j . Here $X_{i,j}$ is the ordered set of coordinates describing the well path.

3.3 Simulation layer

To deal with the uncertainty that is present in the ensemble, we treat all ensemble members as equally probable representations of the reservoir. In general, a main implication of a realistic well path is that it cannot solely rely on the reservoir model, but that it has to adapt through geosteering operations to reservoir properties measured during drilling. This process is emulated by the AWP, which adapts the well path by making local adjustments to the trajectory based on the local property map $p_j(\mathbf{x})$ of each realization.

The reservoir simulation can be written as a function of the customized paths $P_{i,j}$ and the realization R_j as follows:

$$f_j := \hat{s}_j(P_{i,j}, R_j) \quad , \quad (9)$$

where f_j is the output cost function of a reservoir simulation $\hat{s}_j$ performed with the well trajectories $P_{i,j}$ in the model realization R_j . As the ensemble members are regarded

equally probable, the expected value for the set of individual cost functions f_j is calculated as the average

$$\mathbb{E}(f_1, f_2, \dots, f_M) = \frac{1}{M} \sum_{j=1}^M f_j \quad . \quad (10)$$

This average value is then selected as the objective function in the optimization layer:

$$\psi(\mathbf{x}) = \mathbb{E}(f_1, f_2, \dots, f_M) \quad . \quad (11)$$

4 Problem description

In this section we present (1) the ensemble of models used in the experiments, (2) a description of the robust optimization problem tackled in the case studies.

The remainder of this section is split into two parts. The first part contains a description of the ensemble of model realizations used in the experiments. The second part presents an assessment of the AWP parameterization in comparison with the SL representation.

4.1 Models

All simulations presented in this paper will use the Olympus ensemble [15]. Olympus is a synthetic hydrocarbon reservoir constructed to represent a fluvial reservoir on the Norwegian Continental Shelf. There are significant contrasts in permeability, porosity and saturation between the channels and the surrounding volumes. The channels are meandering, gradually changing position from top to bottom. The Olympus model also contains a more homogeneous formation underneath the channelized formation. In this study the lower formation is removed to reduce the computational cost of simulations.

In the numerical reservoir models employed in this work, the injectors are placed in locations surrounding a larger oil-containing volume, as illustrated in Fig. 6. All injectors

Fig. 6 A height map overlaid by the location of the injectors as shown in black

share a common maximum bottom hole pressure of 240 bars, and have no limits to their injection rates. The producer has a minimum bottom hole pressure threshold of 110 bars and no production rate limits.

4.2 Optimization description

In this section we present the problem formulation for the robust well placement optimization problem applied to the Olympus ensemble described previously. As mentioned in the section on robust optimization above, the main goal is to determine the location of one or more producers that maximize the expected value defined in Eq. 10. Notice that, at each iteration, the evaluation of the mean objective over the different trajectories is somewhat costly since it requires running one simulation for each realization. As mentioned before, the decision variables are the positions of the producers only, while the injectors are kept stationary. Although the final trajectories output by AWP are refined, in the optimization domain they contain only well heel and toe coordinates, similarly to the SL parameterization. This means that the optimization problem with both the AWP and the SL has only 6 variables per producer, namely the well heel and toe coordinates in three-dimensional space $\mathbf{x} = (\mathbf{x}_h, \mathbf{x}_t) = (x_h, y_h, z_h, x_t, y_t, z_t)$.

The objective function chosen for the experiments is the expected value of the standardized formulation of the Net-Present-Value (NPV) with the inclusion of capital costs:

$$NPV = \sum_{t=1}^T \frac{R_t}{(1+d)^t} - w_t \quad (12)$$

where R_t is the net cash-flow of each period t , d is the discount factor, w_t is the cost of wells, and T is the number of periods. In our case, we set the discount factor $d = 0.08$, and the number of periods of time $T = 20$. In our setup, the wells are assumed to be drilled at a fixed price. The net cash-flow at period t is calculated as:

$$\begin{aligned} R_t = & \text{profit of hydrocarbon production} \\ & - \text{cost of water injected} \\ & - \text{cost of water treatment} \end{aligned} \quad (13)$$

for every interval of 365 days. All prices are assumed to be fixed, and are set to the following values: oil price at 45 \$/bbl, cost of water injection at 2 \$/bbl, and cost of water production at 6 \$/bbl

5 Parameterization validation – random search

This section presents a comparison between the AWP and the SL parameterizations with respect to their productivity

for the same well heel and toe coordinates evaluated over the ensemble. While the SL parameterization is a simple straight line from the heel to the toe, the AWP adapts the well trajectory to the reservoir properties such as permeability, boundaries and structural components with the aim to increase the productivity. The following analysis confirms the productivity differences resulting from the enhanced AWP trajectory compared to the basis straight-line representation.

The productivity of the AWP and SL approaches in different regions of the reservoir are compared using a random search (RS) algorithm [4], adapted to create random wells within the reservoir. In this algorithm, the well locations are uniformly distributed within the feasible space. For each set of heel and toe coordinates, $\mathbf{x} = (\mathbf{x}_h, \mathbf{x}_t)$, the algorithm creates well paths using both the AWP and the SL parameterizations. The two parameterizations are compared with respect to cumulative liquid production and the relative standard deviation of liquid production across the ensemble. The random search employed in this work considers 938 wells, for which the chosen measures are evaluated over all the realizations. Notice that only the producer location is changed for this validation, whereas the injector locations remain fixed with their positions shown in Fig. 6. The relative standard deviation formula utilized in the RS algorithm is given by:

$$\sigma_f = \frac{1}{\mu_f} \sqrt{\frac{1}{M} \cdot \sum_{j=1}^M (f_j - \mu_f)^2} \quad (14)$$

where M is the ensemble size, and $\mu_f = \mathbb{E}(f_1, f_2, \dots, f_M)$ is the mean corresponding to the expected value $\mathbb{E}$ for the set of functions $f_1, f_2, \dots, f_M$.

The results obtained in the parameterization validation by means of the RS algorithm are summarized in Table 1 and illustrated as cumulative histograms depicted in Fig. 7. The results show that by utilizing the AWP methodology the expected cumulative liquid production is larger than the one obtained with the SL method. Moreover, the relative standard deviation with the AWP method, as calculated by Eq. 14, is less than the relative standard deviation

Table 1 This table shows the mean liquid production (Liq. Prod.) across the ensemble and the relative standard deviation of liquid production evaluated over all realizations of the ensemble (Rel. Std) according to a number of unique pairs of heel and toe coordinates (Num.)

Type	Num. [#]	AWP	SL
		Mean [SM3]	Mean [SM3]
Liq. Prod.	938	1.27E7	1.12E7
Rel. Std.	938	46.1	48.5

Fig. 7 Cumulative histograms showing the total liquid production (a) and the relative standard deviation over the realization (b)

with the SL well description. These results suggest that AWP trajectories have the collective effect of being able to compensate for the differences in geology, thus providing an improved representation capable of dealing with the underlying uncertainty. Moreover, the well parameterized using AWP yields more production per length than the equivalent well parameterized as a straight line. Despite having the same compact well representation, i.e., $\{\mathbf{x}_h, \mathbf{x}_t\}$, the AWP procedure achieves improved well representations because it increases their contact with permeable layers, thereby leading to improved productivity. Furthermore, because of the higher productivity, the more refined trajectory is expected to lead to faster convergence and a more efficient search.

6 Case studies

In this section we present three case studies with increasing complexity. The first case is presented as a proof of concept embedding the AWP in an optimization loop as described in Section 3 for a single realization environment ($M = 1$) using APPS. Case 2 and 3 utilize the entire Olympus ensemble ($M = 50$) in a robust well placement optimization procedure as described in Section 3. The main idea is to highlight the benefits of using tailor-made well trajectories using AWP starting from a single realization environment with one producer and gradually transitioning to a robust well placement optimization workflow with multiple producers. Case study 1 illustrates the conceptual idea of using AWP in an optimization loop for the placement of one producer in a single realization environment. In Case study 2 the problem complexity is increased by considering an ensemble with 50 realizations in the optimization of a sole producer, where the focus lies in assessing the impact of using AWP instead of the SL when uncertainty is present. In Case study 3, two additional producers are considered in the robust optimization. The inclusion of new wells increases the degrees of freedom of the optimization problem and

allows for an evaluation of the scalability of AWP to more complex scenarios, involving more interdependent dynamics between the wells and the reservoir.

The computational cost of the pattern search algorithm APPS grows significantly as the number of free variables and realizations increases. This algorithm is therefore omitted in the multi-well optimization due to its prohibitive computational cost, and the optimization is only run using the stochastic PSO algorithm. All the case studies were performed with the SL and AWP approaches using the same initial well locations and constraint values.

6.1 Case study 1 - single producer in a single realization

This case study compares results from an optimization routine using the AWP procedure against results using the same optimization routine with the traditional SL well trajectory. The purpose is to investigate if and how the solutions and performance vary when using the different types of parameterization in the optimization procedure.

The optimization algorithm was run with three different initial well locations, all placed within the region surrounded by water injectors. We ran the optimization using wells parameterized with both the SL and AWP procedure to automatically generate well paths. There are thus six optimization runs in total. The initial positions for each of the three cases are indicated by red lines in Fig. 8. Note that this red line only indicates the horizontal direction of the well given by the initial heel and toe values. The actual well trajectory will be different in the vertical direction not shown in Fig. 8. Thus the actual initial well trajectory will either be a line or the well trajectory generated by the AWP method. The two optimization runs will therefore start with slightly different well configurations, and thereby with different NPV values at the initial step in the optimization process.

Recall that the optimization algorithm applied in this case study is the deterministic pattern-search algorithm APPS

Fig. 8 Reservoir height map showing the different well configurations for the three optimization example cases. Blue shows the trajectory solution obtained using AWP, while orange shows the SL trajectory,

and red is the initial well trajectory. Black lines represent the injectors. The color bar shows the difference in height of the top layer in the reservoir

described in [19]. The current best NPV value at different iterations is plotted in Fig. 9. In line with the discussion above, even though the three cases have the same initial variable for the heel and toe, they still have different initial NPV values as the well trajectories are different. For all cases, optimization runs using the AWP procedure give a higher NPV value, which is in line with the observations summarized in Table 2.

The convergence plots in Fig. 9 show that by applying the AWP inside the optimization procedure, the optimization algorithm converges faster and to a higher value than

its more traditional SL well counterpart. The number of iterations required for convergence is seen to decrease from around 28–39 iterations to around 22–25. When looking at the final positions after optimization, as illustrated by Fig. 8, all six well positions are seen to end up fairly close to each other.

Investigating the full well trajectories, one can observe that by adapting to local geometry and structural challenges, the AWP procedure generates wells that are both more productive and better equipped to deal with complexities such as faults. In addition, the AWP is able to produce from multiple high permeable streaks simultaneously, which is hard to obtain when using the SL approach.

In general, the higher permeability sands penetrated by the AWP procedure increase the drainage volume and accelerate the oil production, resulting in a better initial NPV for all cases independent of their starting location, as reflected by the values in Table 2.

The initial and final production curves of Case 1 are plotted in Fig. 10. We observe that both SL and AWP solutions significantly improve the oil production. As this is a main component of the NPV, they both improve the NPV significantly. Interestingly, the optimal AWP well path gives lower water production compared to the optimal SL. As the optimal oil production is similar for the two cases, this indicates that the total liquid production is lower for the AWP compared to the SL solution. This is in contrast to the result when they have the same heel and toe coordinates as indicated by the verification study summarized in Fig. 7. Thus, it is not necessarily the case that an optimization

Fig. 9 Objective function value (NPV) as a function of iterations for all solutions with the AWP (solid) and SL (dashed) solutions

Table 2 NPV values for the three optimization cases, at the initial position (Init.), and after optimization (Opt.) using either the AWP or SL method for the well paths. In addition, the table contains the number of iterations (Its.) and evaluations (Evals.) for each approach case

Case	AWP				SL			
	Init. [\$]	Opt. [\$]	Its. [#]	Evals. [#]	Init. [\$]	Opt. [\$]	Its. [#]	Evals. [#]
1	8.37E8	1.50E9	24	252	6.67E8	1.42E9	39	422
2	9.06E8	1.50E9	25	262	5.50E8	1.43E9	28	296
3	1.25E9	1.46E9	22	255	7.52E8	1.43E9	38	413

using the AWP procedure ends up with a solution that gives a more productive well, even though the AWP procedure gives higher productive wells than SL for the same well coordinates.

The final location of the producers when using the two types of trajectories, as indicated in Fig. 8, appear to have roughly the same layout. However, the difference between the two approaches, both in terms of final objective function value as well as in terms of total number of iterations used, is significant, as shown in Table 2. Compared to its straight line counterpart, the AWP procedure can adapt more efficiently to complex geological formations. This feature effectively serves as an additional constraint-handling measure during optimization. The impact of this feature is particularly clear when using pattern-search algorithms together with constraint-handling capability that only relies on geometry. The pattern-search approach works in such a way that perturbations are performed iteratively along each of the three spatial dimensions. Since this initial stepping does not take into account the geometry of the reservoir, the algorithm is likely to step outside of reservoir bounds. At this point, constraint-handling within the implementation in this work projects out-of-bounds coordinates back into

the reservoir. This constraint-handling, however, is guided solely by the geometry of the grid and trajectory, while the indirect constraint-handling performed by the AWP procedure explicitly takes into account the local reservoir properties. The result is that straight line wells that only rely on relatively crude geometry-based constraint-handling can have large parts of their original paths rendered non-productive. Such lower-producing solution candidates are likely to lead the optimization algorithm to move away from trajectories close to the outer edges, in effect limiting the extent of the search space by not having the capacity to handle complex bounds, which are likely to be encountered in realistic applications. With the introduction of the AWP procedure within the well placement optimization routine, on the other hand, trajectories are effectively modified to fit high-productivity areas irrespective of complex geometries including faults or dome-like structures. Consequently, the added property-based handling of well paths introduced by the AWP procedure leads to more realistic well trajectories.

6.2 Case study 2 - single producer in an ensemble

The APPS and PSO algorithms were applied to the previously described robust well placement optimization problem considering one producer and employing both the AWP and the SL parameterization methods.

In contrast to Case study 1 where the three optimization runs were performed in a single realization environment, the optimization runs were initialized from five different starting locations for the APPS algorithm. The results obtained for these runs are summarized in Table 3 and depicted in Fig. 11. Figure 11(a) shows the current best NPV value by iteration. The NPV values are scaled by 10^9 for all the cases. The deterministic behavior of APPS allows for a case-by-case comparison of the differences between AWP and SL, as plotted in Fig. 11(b). A positive value indicates a higher NPV from the AWP method over the SL method, while a negative value indicates the opposite. In four of the five cases, AWP yielded a higher objective. Figure 11(b) also indicates that four out of five cases converged earlier with AWP. This NPV increase might be a consequence of the overall increase in the liquid production caused by the

Fig. 10 Cumulative production of oil and water for Run 1 in the well placement optimization using APPS

Table 3 This table shows the final objective value of each optimization run with one producer

Run	PSO		APPS			
	AWP	SL	AWP		SL	
Num [#]	NPV[\$]	NPV[\$]	NPV[\$]	EvalNr[#]	NPV[\$]	EvalNr[#]
1	2.18E9	2.11E9	2.00E9	228	1.98E9	288
2	1.99E9	2.11E9	2.00E9	300	2.17E9	288
3	2.17E9	2.10E9	1.99E9	192	1.95E9	336
4	2.17E9	2.10E9	2.00E9	240	1.95E9	300
5	2.17E9	2.10E9	2.22E9	312	2.13E9	372
6	2.16E9	2.10E9				
Avg:	2.14E9	2.10E9	2.04E9	254	2.03E9	316

Results are taken from the last iteration performed with each algorithm

AWP parameterization, as seen in the validation analysis presented in Section 5. In only one of the cases the solution obtained with SL yields a higher objective compared to the AWP solution.

If the response surface from AWP and SL have different shapes, then APPS might converge to different local maxima. This is a plausible reason to justify the disparity between the solutions provided by AWP and SL in some of the cases. For the case in which the SL yielded a higher NPV than AWP, further investigation into the results indicates that the increase in liquid production coming from the AWP parameterization worsen the NPV when compared to the SL approach for the same well heel and toe coordinates. This allowed the SL to find a different maximum, which in the end turned out to yield a higher NPV.

The same problem was optimized using the PSO algorithm with 6 different runs for both the AWP method and the SL method. The current best objective function values versus the number of iterations are shown in Fig. 12. A case-by-case comparison is not feasible here because of the randomly-generated initial positions of the PSO algorithm.

From Fig. 12 we observe that the PSO with AWP seems to converge towards the best NPV value in fewer iterations compared to PSO with SL. We also observe that, except for one outlier, the expected NPV values obtained with AWP are higher than all the objective function values achieved using SL. The difference between AWP and SL is in the range of what was observed by the random search previously, and might therefore be attributed to increased liquid production by AWP over SL (Fig. 7).

Fig. 11 Convergence plot for the objective function (NPV) over 5 robust optimization runs with APPS with all ensemble realizations. (a) shows the convergence graphs for all runs where blue is AWP and orange is SL, (b) shows the difference between AWP and SL given

the same initial well locations. In (b) solid lines indicate that the optimization procedure is running for both parameterizations. The dashed line indicates that the AWP has converged while the SL continued executing, whereas dotted lines indicate the opposite

Analogously to what occurred with the APPS algorithm, there is one case which yielded a considerably lower objective function with the AWP parameterization. One possible explanation for this phenomenon is that the optimization procedure got trapped into a local maximum. Stochastic behavior is inherently difficult to explain, and it is likely that if the run was performed with a larger swarm size, the solutions from all cases would be more homogeneous. Even with the outlier which underperformed, the general behavior still suggests a more efficient search when substituting the SL parameterization with the AWP, similar to the previous case in a single realization environment.

6.3 Case study 3 - multiple producers in an ensemble

As in the previous case study we utilize the Olympus ensemble as the numerical models for the robust optimization of the positions of the producers. To increase our understanding of the scalability of the optimization procedure with AWP, two additional producers are considered in the optimization problem. Due to the stochastic nature of PSO, the initial particles in the swarm are randomly generated throughout the search space. Therefore, we ran the optimization in 5 different scenarios with both AWP and the SL, and take the average objective function value and the standard deviation across the scenarios. Due to restrictions in terms of computational resources the runs were limited to 16 iterations, which seems to be sufficient to highlight the differences in terms of convergence and variance of the solutions.

Table 4 summarizes the results from five runs with each parameterization. It can be seen that both the final and

average objective function values with AWP are higher than with the SL in all the runs. Notice that the average well length is also significantly shorter with AWP. The reduction in the well lengths might reduce the likelihood of early water production. Figure 13(a) depicts the objective function evolution through each iteration of PSO with both the AWP and SL parameterization with 3 producers. Figure 13(b) shows the average objective function evolution and the corresponding standard deviations for all the runs with both parameterizations. From this graphical representation we observe that the initial slope of the AWP is steeper than the equivalent slope for the SL parameterization. This may be an indicator that the well parameterization with AWP increases the rate of convergence or smooths the response surface allowing for easier traversal. Notice that, in an early stage of the optimization runs (from iteration 2 to 8) there are successive iterations in which all runs with AWP present higher objective function value than the runs with the SL.

The differences between AWP and SL were further investigated by comparing the solutions produced by the best runs with each parameterization. The final locations for the three producers are illustrated in Fig. 14. From this figure it is possible to see that the producers have roughly the same configuration in comparison to one another, however, the final producers constructed by AWP (blue) are significantly shorter than those generated by the SL (orange) parameterizations.

The production profiles for all ensemble members of the best AWP and SL optimization runs are illustrated in Fig. 15. We observe that the expected value of oil production is slightly higher for SL compared to the AWP. In terms of expected water production, the SL produces significantly more water than the best run with AWP.

Fig. 12 Convergence plot of the PSO algorithm for the objective function (NPV) over 6 different cases, run with all realizations. Blue lines indicate AWP method, while orange lines indicate the SL method

Table 4 This table shows the final optimized NPV obtained with both parameterizations and the number of perforated meters (the sum of length of all producers) inside the reservoir

Run [#]	AWP		SL	
	NPV [\$]	Tot. Len. [m]	NPV [\$]	Tot. Len. [m]
1	2.13E9	1579	2.05E9	2887
2	2.15E9	1459	1.95E9	2153
3	2.10E9	2818	2.07E9	2598
4	2.13E9	1708	2.09E9	2431
5	2.13E9	1901	2.08E9	3000
Avg:	2.13E9	1893	2.04E9	2614
Std:	1.12E7	485	4.62E7	306

Fig. 13 (a) is a convergence plot for the objective function (NPV) over 5 different optimization runs utilizing the PSO algorithm evaluated across all realizations. (b) shows the average value for all cases for

each iteration, in addition to the standard deviation over the 5 runs per iteration. Blue transparent area indicate AWP method, while orange transparent area indicate the SL method

This observation holds if we analyze the transparent red (water production SL) and grey (water production AWP) areas in this figure. The solution obtained with the SL parameterization increases water production earlier and significantly more than the solution provided by AWP.

The NPVs for the different ensemble realizations are compared in Fig. 16. The plot shows that for most realizations the NPV is increased by replacing the SL with the AWP methodology. In this comparison, the AWP seems to improve both optimistic and pessimistic cases compared with the SL parameterization.

This case study highlights interesting behavior of AWP in a more realistic environment with multiple wells for robust well placement optimization. Within each of the optimization runs better objective function values are achieved at earlier iterations using AWP. This could be

credited to the additional flexibility of AWP with regards to properties and difficult geometry, features which generally improve the productivity of the producers. Even though the producer configurations for both parameterizations are similar, as seen in Fig. 14, there is a significant difference in their objective function values. As the AWP makes tailored trajectories to individual channels, the results show that these wells are more targeted and this significantly reduces the water production. A counter argument is that the objective function in these cases can be modified for a different objective, e.g. maximizing oil production and omitting the water injection/treatment costs, which might yield different results. Uncertainty that can be handled during the drilling process by adaptive directional drilling

Fig. 14 Final location of the producers in the best runs with both the AWP (blue) and the SL (orange) parameterization

Fig. 15 The cumulative oil and water production for both parameterizations for all realizations in the Olympus ensemble

Fig. 16 Differences in Net-Present-Value for the best run with AWP and SL; blue bar values indicates an improvement by the AWP parameterization, orange bar indicates an improvement by the SL parameterization

is to some degree captured and embedded into the AWP methodology. This way, the AWP might reduce the variance in objective function between realizations, and its effect on the variance of the cost function should be accounted for in the optimization framework. We acknowledge that the average NPV, which is used as the objective in the optimization in the work, is a rather simple objective which, e.g., does not take into account the risk associated with the variation between the individual realizations. Other measures, such as risk-averse objective functions, could be employed in the optimization layer to account for such risks, avoiding therefore solutions that are great in average but which could present high financial risks associated with the underlying decisions.

Clearly, this is not an operational environment, but it is a step towards showcasing the capabilities of including the decision-making procedure of geosteering in closed-loop reservoir management. The artificial neural network that was applied to all of the cases above is of conservative nature, and will not make significant changes to the individual trajectories because it was trained for consistency. With minor alterations to the training of the artificial neural network it can be allowed to create larger deviations from each realization. Whether or not that this is desirable is a topic for further investigation.

7 Conclusions

This work proposes a methodology for robust well placement optimization which combines an adaptive well parameterization with traditional derivative-free optimization algorithms. As the real-life drilling phase adapts to local

information of the reservoir, the suggested optimization process adapts in a similar manner to obtain optimized well placements within an adaptive drilling workflow. The adaptive well parameterization, called Automatic Well Planner (AWP), emulates decision-making processes during geosteering operations. In this respect, well placement optimization with AWP is able to better capture and adapt to the geological uncertainty represented by the ensemble.

The AWP method generates a well trajectory given a heel and toe location. Thus the input variables to the AWP is of the same dimensions as a straight line (SL) well. This implies that the computational cost of optimizing with AWP is of the same order as using the simpler SL. In this paper we tested two derivative-free optimization algorithms, APPS and PSO, in three separate case studies with increasing complexities. In Case study 1 it was shown that by replacing the traditional SL approach with the AWP the production can be increased and the number of iterations can be reduced in a single realization environment where uncertainty is omitted.

When the complexity is increased and the AWP procedure is embedded in a robust optimization scheme, as shown in Case study 2, the optimization procedure with AWP converged to a higher objective function value in comparison to equivalent runs with SL, for all runs except one. The runs performed with PSO achieved higher objective function values earlier with AWP. Similarly, most of the runs performed with APPS reached the convergence criterion faster with AWP than with SL.

Finally, in Case study 3, the complexity is increased further by adding two additional wells. The differences between AWP and SL were more pronounced than both of the previous single producer optimization cases, with all AWP cases converging to a higher objective function value than the comparative runs with SL parameterization. In general, the best solution obtained with AWP presented higher NPV when compared to the best solutions found using the SL parameterization.

The general improvement in performance by AWP over SL can possibly be attributed to more favorable features of the response surface. In particular, AWP can contribute favorably to the optimization process by efficiently adapting to the range of different geological conditions posed by the realization ensemble, thus essentially help smoothing the cost function. Further work is required to understand why in a few cases the inclusion of AWP during optimization is detrimental to the search. This might be due to the AWP at all times promoting extensive penetration of highly-permeable channels, which does not necessarily lead to improved objective values if it exposes trajectories to early water-breakthrough and high water production. Along these lines, performing an AWP operation on a SL solution can decrease the final objective since the SL

solution is exclusively optimal with respect to a straight-line trajectory. Further work on these topics can provide practical guidelines for the efficient use of the AWP technology.

Several other topics are also interesting for future research. In this work, constraints are applied to individual wells only. To further enhance applicability, additional constraints can be imposed, e.g., inter-well distance, distance to boundary and distance to the oil-water contact. These can be treated explicitly by constraint-handling methodology in the optimization layer and be enforced implicitly by the neural network after having been introduced during training. Moreover, testing other combinations of property maps as inputs into the neural network can bring the AWP closer to geosteering tools, thereby improving our current results. It can also guide what would be the most valuable information to be obtained from logging-while-drilling tools. As the AWP method is promising, similar methodologies might be tried for other field development problems where a high number of free variables can be reduced to lower order problems by the use of machine learning procedures. This way complex field development problems can be made computationally feasible for automated optimization algorithms.

Acknowledgements This research is a part of BRU21 – NTNU Research and Innovation Program on Digital and Automation Solutions for the Oil and Gas Industry (www.ntnu.edu/bru21). Brage S. Kristoffersen is supported by Equinor ASA through the BRU21 project. Carl Fredrik Berg acknowledges funding from the Research Council of Norway (Centers of Excellence funding scheme, project number 262644, PoreLab).

Funding Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Al-Mubarak, S.M. et al.: Real-time reservoir management from data acquisition through implementation: closed-loop approach. In: Intelligent Energy Conference and Exhibition. Society of Petroleum Engineers (2008)
2. Baumann, E.J.M., Dale, S.I., Bellout, M.C.: Fieldopt: A powerful and effective programming framework tailored for field development optimization. *Comput. Geosci.* **135**(104), 379 (2020). <https://doi.org/10.1016/j.cageo.2019.104379>
3. Bellout, M.C., Echeverría Ciaurri, D., Durllofsky, L.J., Foss, B., Kleppe, J.: Joint optimization of oil well placement and controls. *Comput. Geosci.* **16**(4), 1061–1079 (2012). <https://doi.org/10.1007/s10596-012-9303-5>
4. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **13**, 281–305 (2012)
5. Bouzarkouna, Z., Ding, D.Y., Auger, A.: Well placement optimization with the covariance matrix adaptation evolution strategy and meta-models. <https://doi.org/10.1007/s10596-011-9254-2> (2011)
6. Bratvold, R.B., Begg, S.: Making good decisions. *Soc. Pet. Eng.* (2009)
7. Burgers, G., van Leeuwen, P.J., Evensen, G.: Analysis scheme in the ensemble kalman filter. *Mon. Weather Rev.* **126**(6), 1719–1724 (1998). [https://doi.org/10.1175/1520-0493\(1998\)126<1719:ASITEK>2.0.CO;2](https://doi.org/10.1175/1520-0493(1998)126<1719:ASITEK>2.0.CO;2)
8. Capolei, A., Suwartadi, E., Foss, B., Jørgensen, J.B.: Waterflooding optimization in uncertain geological scenarios. *Comput. Geosci.* **17**(6), 991–1013 (2013). <https://doi.org/10.1007/s10596-013-9371-1>
9. Chen, Y., Oliver, D.S., Zhang, D., et al.: Efficient ensemble-based closed-loop production optimization. *SPE J.* **14**(04), 634–645 (2009). <https://doi.org/10.2118/112873-PA>
10. Echeverría Ciaurri, D., Isebor, O., Durllofsky, L.: Application of derivative-free methodologies to generally constrained oil production optimization problems. *Procedia Comput. Sci.* **1**(1), 1301–1310 (2010). <https://doi.org/10.1016/j.procs.2010.04.145>
11. van Essen, G., Zandvliet, M., Van den Hof, P., Bosgra, O., Jansen, J.D.: Robust waterflooding optimization of multiple geological scenarios. *SPE J.* **14**(01), 202–210 (2009). <https://doi.org/10.2118/102913-PA>
12. Evensen, G.: The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.* **53**(4), 343–367 (2003). <https://doi.org/10.1007/s10236-003-0036-9>
13. Evensen, G.: Data Assimilation: The Ensemble Kalman Filter. Springer Science & Business Media, New York (2009). <https://doi.org/10.1007/978-3-642-03711-5>
14. Forouzanfar, F., Reynolds, A.C., Gaoming, L.: Optimization of the well locations and completions for vertical and horizontal wells using a derivative-free optimization algorithm. <https://doi.org/10.1016/j.petrol.2012.03.014> (2012)
15. Fonseca, R., Della Rossa, E., Emerick, A., Hanea, R., Jansen, J.: Overview of the Olympus field development optimization challenge. In: ECMOR XVI-16th European Conference on the Mathematics of Oil Recovery, vol. 2018, pp. 1–10. European Association of Geoscientists & Engineers (2018). <https://doi.org/10.3997/2214-4609.201802246>
16. Foss, B., Jensen, J.P.: Performance analysis for closed-loop reservoir management. *SPE J.* **16**(1), 183–190 (2011). <https://doi.org/10.2118/138891-PA>
17. Gu, Y., Oliver, D.S., et al.: History matching of the PUNQ-S3 reservoir model using the ensemble kalman filter. *SPE J.* **10**(02), 217–224 (2005)
18. Hassani, H., Sarkheil, H., Foroud, T., Karimpooli, S., et al.: A proxy modeling approach to optimization horizontal well placement. In: 45th U.S. Rock Mechanics / Geomechanics Symposium (2011)
19. Hough, P.D., Kolda, T.G., Torczon, V.J.: Asynchronous parallel pattern search for nonlinear optimization. *SIAM J. Sci. Comput.* **23**(1), 134–156 (2001). <https://doi.org/10.2118/89942-MS>

20. Jansen, J.D., Brouwer, R., Douma, S.G.: Closed loop reservoir management. In: SPE Reservoir Simulation Symposium. Society of Petroleum Engineers (2009). <https://doi.org/10.2118/119098-ms>
21. Kolda, T.G.: Revisiting asynchronous parallel pattern search for nonlinear optimization. *SIAM J. Optim.* **16**(2), 563–586 (2005). <https://doi.org/10.1137/040603589>
22. Kristoffersen, B.S., Bellout, M.C., Silva, T.L., Berg, C.F.: An automatic well planner for complex well trajectories. *Math. Geosci.*, 1–25. <https://doi.org/10.1007/s11004-021-09953-x> (2021)
23. Li, Q., Omeragic, D., Chou, L., Yang, L., Duong, K., et al.: New directional electromagnetic tool for proactive geosteering and accurate formation evaluation while drilling. In: SPWLA 46th Annual Logging Symposium. Society of Petrophysicists and Well-Log Analysts (2005)
24. Liu, N., Oliver, D.S.: Ensemble kalman filter for automatic history matching of geologic facies. *J. Pet. Sci. Eng.* **47**(3–4), 147–161 (2005). <https://doi.org/10.1016/j.petrol.2005.03.006>
25. Naevdal, G., Johnsen, L., Aanonsen, S., Vefring, E.: Reservoir monitoring and continuous model updating using ensemble kalman filter. *SPE J.* **10**(01). <https://doi.org/10.2118/84372-MS> (2005)
26. Nwankwor, E., Nagar, A.K., Reid, D.: Hybrid differential evolution and particle swarm optimization for optimal well placement. *Comput. Geosci.* **17**(2), 249–268 (2013). <https://doi.org/10.1007/s10596-012-9328-9>
27. Onwunalu, J.E., Durlofsky, L.J.: Application of a particle swarm optimization algorithm for determining optimum well location and type. *Comput. Geosci.* **14**(1), 183–198 (2010). <https://doi.org/10.1007/s10596-009-9142-1>
28. Sarma, P., Durlofsky, L.J., Aziz, K.: Efficient closed-loop production optimization under uncertainty. In: 67th EAGE Conference & Exhibition, pp. cp–1. European Association of Geoscientists & Engineers (2005). <https://doi.org/10.3997/2214-4609-pdb.1.C039>
29. Seydoux, J., Legendre, E., Mirto, E., Dupuis, C., Denichou, J.M., Bennett, N., Kutiev, G., Kuchenbecker, M., Morriss, C., Yang, L., et al.: Full 3d deep directional resistivity measurements optimize well placement and provide reservoir-scale imaging while drilling. In: SPWLA 55th Annual Logging Symposium. Society of Petrophysicists and Well-Log Analysts (2014)
30. Stordal, A.S., Szklarz, S.P., Leeuwenburgh, O.: A theoretical look at ensemble-based optimization in reservoir management. *Math. Geosci.* **48**(4), 399–417 (2016). <https://doi.org/10.1007/s11004-015-9598-6>
31. Volkov, O., Bellout, M.: Gradient-based constrained well placement optimization. *J. Pet. Sci. Eng.* **171**, 1052–1066 (2018). <https://doi.org/10.1016/j.petrol.2018.08.033>
32. Wang, C., Li, G., Reynolds, A.C.: Production optimization in closed-loop reservoir management. *SPE J.* **14**(3), 506–523 (2010). <https://doi.org/10.2118/109805-PA>
33. Zandvliet, M., Handels, M., van Essen, G., Brouwer, R., Jansen, J.D.: Adjoint-based well-placement optimization under production constraints. *SPE J.* **13**(4), 392–399 (2008). <https://doi.org/10.2118/105797-PA>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.