Journal of Petroleum Science and Engineering 208 (2022) 109523

Contents lists available at ScienceDirect

Journal of Petroleum Science and Engineering

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

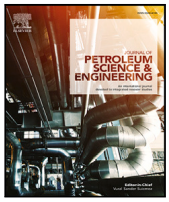

Reduced well path parameterization for optimization problems through machine learning

Brage S. Kristoffersen [a], Mathias C. Bellout [a], Thiago L. Silva [a,b], Carl F. Berg [a,*]

[a] Department of Geoscience and Petroleum, NTNU, Norway
[b] Department of Sustainable Energy Technology, SINTEF Industry, Norway

ARTICLE INFO

Keywords:
Machine learning
Well placement
Derivative-free optimization
Well parameterization

ABSTRACT

In this work we apply a recently developed machine learning routine for automatic well planning to simplify well parameterization in reservoir simulation models. This reduced-order parameterization is shown to be beneficial for well placement optimization, both in terms of convergence and final well configuration. The proposed machine learning routine maps trajectories that honor predefined engineering requirements by exploiting spatial information about the reservoir and prior domain-knowledge about the problem. In this paper, the well planner creates wells that traverse high-permeable parts of the reservoir, thereby increasing well productivity. Previous work found that small changes to the start- and end-points of the well had limited impact on most of the resulting well trajectories, since development of trajectories is chiefly determined by local information around the digital drill bit. In particular, changes in the depth component of the start- and end-points had limited impact on the trajectory away from the end-points. Based on these observations, this work reduces well parameterization to only include horizontal coordinates. The main assumption is that the perforated part of the well always enters the reservoir at the upper reservoir boundary, while the stopping criteria in the machine learning routine is a perforation length only. This formulation reduces the number of decision variables from six to four coordinates for each well. The resulting reduced search space enables a more efficient exploration effort at the cost of less freedom over the start and end points of the well path. However, we show that the highly-refined well trajectory developed by the well planning routine is robust and compensates for fewer degrees of freedom at the overarching parameterization. This robustness is tested by investigating the effect of different start locations on the automatic well planning routines. Moreover, the effect of the reduced well parameterization for well placement optimization is explored. Two optimization scenarios using four different optimizations algorithms are presented. Results show the implementation of the reduced well parameterization for optimization purposes consistently produces high quality solutions.

1. Introduction

Development of subsurface reservoirs typically involves decisions on where and when to drill wells, and how to control existing wells. These decisions are commonly supported by numerical models to predict how new well configurations and different production strategies influence subsurface flow. Decision-making usually targets maximizing profit, however, other goals are also increasingly relevant, e.g., reducing carbon emissions by minimizing power consumption on surface facilities. Maximizing profit for petroleum field development typically implies increasing hydrocarbon production and limiting water production and injection, while at the same time minimizing operational and capital costs.

The search for optimal well trajectories and control settings is often conducted through manual modification of relevant parameters in numerical reservoir models, e.g., based on heuristics and/or experience. Such manual exploration in multidimensional space is labor intensive, and tends to concentrate the search in the vicinity of one local optimum (Bittencourt et al., 1997). These challenges spur the need for methodologies and workflows to automate search processes for optimal well placement and production strategies. However, such push for automation is severely limited by the computational cost associated with reservoir simulation. More efficient optimization techniques, increasingly practical simulation proxies (Møyner et al., 2014), and the continuing increase in computational power, enable the use of iterative procedures to complement manual field development decision-making.

In this work we focus on the problem of well placement optimization. Due to its practical relevance and scientific interest, the literature

* Corresponding author.
E-mail address: carl.f.berg@ntnu.no (C.F. Berg).

https://doi.org/10.1016/j.petrol.2021.109523
Received 21 May 2021; Received in revised form 23 August 2021; Accepted 16 September 2021
Available online 24 September 2021

contains a wide range of algorithms and methodologies to address this problem. The problem is considered multimodal and non-smooth, mainly attributed to geological heterogeneity (Onwunalu and Durlofsky, 2010). Derivative-free procedures for well placement optimization include Particle Swarm Optimization (Onwunalu and Durlofsky, 2010; Isebor et al., 2014; Afshari et al., 2014; Jesmani et al., 2016), Genetic Algorithm (Bukhamsin et al., 2010; Güyagüler et al., 2001; Badru et al., 2003; Awotunde, 2016), Simulated Annealing (Beckner et al., 1995), and hybrid methodologies (Khademi and Karimaghaee, 2015). Though the well placement optimization problem favors derivative-free procedures, gradient-based algorithms have also been applied. Approximation of derivative information with respect to well placement variables, e.g., directly through perturbation methods (Bangerth et al., 2006; Li and Jafarpour, 2012), or indirectly through procedures relying on (adjoint-based) gradients associated with control variables (Zandvliet et al., 2008; Sarma et al., 2008) may neither be reliable due to non-smoothness, nor practical for optimizing for complex well configurations. Still, efficient adjoint-based approximations of well placement gradients have been proposed (Volkov and Bellout, 2018) and recently applied to problems of realistic complexity (Krogstad and Nilsen, 2020). However, these adjoint-based optimization techniques require extensive access to the code base of the reservoir simulator.

Due to the high computational cost associated with each solution candidate, finding a global optimal well configuration is practically infeasible for even modestly-realistic cases. The heavy computational cost and the high number of free variables require the implementation of reasonable and effective well parameterizations in addition to the development of optimization techniques with efficient sampling strategies. The issue of efficient well parameterization was addressed in the work introducing the automatic well planner (Kristoffersen et al., 2021). Input for the well planner are heel and toe coordinates only, while the well trajectory is developed through a process resembling a geosteering procedure. Geosteering relies on data obtained from measure-while-drilling tools. With sufficient logging accuracy and real-time model-updating capability, pre-drill plans for trajectory and completion might be reconsidered based on updated information about the subsurface.

The automatic well planner uses a machine learning approach to mimic the geosteering process. In this approach, a neural network is trained to react to reservoir properties in the vicinity of a virtual drill bit. Though the training of the network is computational costly, with training times ranging from minutes to hours on a single workstation, this computation is only performed once. Within the optimization loop, employing the methodology to create a well trajectory has negligible computational cost compared to the cost of running a reservoir simulation. In Kristoffersen et al. (2021), this virtual drilling procedure was shown to be capable of producing wells that better adhere to surrounding reservoir properties and their geometrical distribution. In Kristoffersen et al. (2020) this procedure was applied to an ensemble of reservoir models within an optimization framework to showcase the capabilities of this methodology for robust optimization.

The proposed automatic well planner methodology has already been proven to construct complex well paths using only a limited number of free parameters, e.g., the heel and toe coordinates only (Kristoffersen et al., 2021). This is a critical contribution to deal with the heavy computational cost associated with using derivative-free algorithms for well placement search, since the sampling effort required by these algorithms commonly increases with problem dimension and the total number of functional evaluations is therefore typically high. The main goal of this work is to reduce the number of decision variables further, by reducing the well parameterization from three-dimensional heel and toe coordinates in order to enable the use of two-dimensional lateral (horizontal) coordinates of the heel and toe.

The motivation for this parameter reduction is the reduced sensitivity of our automatic well planner to the start and end-point of the well. This limited sensitivity is illustrated in Fig. 1, were we plot the well trajectory from the automatic well planner using start coordinates with different depths only. We observe that all the wells converge onto a similar trajectory a short distance from the start location. After converging, the neural network is fed similar input data, and therefore yields similar well trajectories for the remainder of the well path.

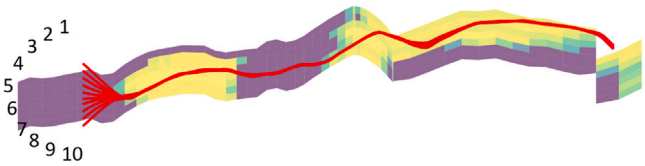

Fig. 1. This figure shows a set of well paths (red lines) obtained from the virtual drilling methodology introduced in Kristoffersen et al. (2021). The wells starts from equal lateral coordinates but different depths, where the wells with different starting depths are indicated by the numbers in the figure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Finally, by altering the automatic well planner methodology to only depend on heel and toe lateral coordinates we are transferring a larger portion of the well trajectory construction to the well planner procedure. In this article, we show that this parameter reduction increases the search efficiency and improves the handling of reservoir discontinuities for optimization.

2. Methodology

This section introduces the methodology for well trajectory design. Well trajectory refers specifically to the continuous path traversed by the well from the heel to the toe. Similar to previous work on this topic (Kristoffersen et al., 2021, 2020), well trajectory development is performed using an artificial neural network as the decision-maker with the main goal of honoring pre-specified design targets. The training data for the neural network is a combination of random well locations with the use of initial reservoir data (corresponding to different well trajectories), and fitness function values associated with these trajectories.

This section focuses on new developments compared to earlier work. The reader is referred to Kristoffersen et al. (2021) and Kristoffersen et al. (2020) for a full description of the methodology and training of the automatic well planner. As presented in Kristoffersen et al. (2021), the automatic well trajectory procedure has two main aims, one is to reach the vicinity of the endpoint, the other is to adapt the well path according to a predefined property field. Trajectory development is conducted in a sequential manner, analog to drilling operations, where the neural network determines the direction of the next well segment based on information obtained near the current location.

This paper proposes increasing the freedom of the well planner methodology, and thereby reducing the parameterization used for optimization. Previously the heel and toe were defined by their three-dimensional coordinates as $\vec{x}_h = (x_h, y_h, z_h)$ and $\vec{x}_t = (x_t, y_t, z_t)$, respectively, defining the start- and end-points of the perforated parts of the well. In contrast, the current work relies solely on the horizontal (lateral) coordinates. Thus a well is now defined by the two-dimensional projection of the heel and toe coordinates $\vec{x}_h^l = (x_h, y_h)$ and $\vec{x}_t^l = (x_t, y_t)$. The lateral position of the heel, $\vec{x}_h^l$, will define the starting location as the upper reservoir boundary at this point, while the lateral position of the toe, $\vec{x}_t^l$, provides an approximation for the lateral position of the toe.

2.1. Mapping the lateral variables to spatial coordinates

The lateral heel $\vec{x}_h^l$ and toe $\vec{x}_t^l$ locations are the decision variables in our procedure. As we need a three-dimensional description for the start and approximate end-point of the well, we need to map the lateral

positions to three-dimensional coordinates. Let $\Omega \subset \mathbb{R}^3$ be the subspace defined by our reservoir model, and let $\Lambda \subset \mathbb{R}^2$ be the projection of Ω to the horizontal plane. We will define functions for the top and bottom of the reservoir as

$$S_t(x,y) = \sup\{z \mid (x,y,z) \in \Omega\}$$
$$S_b(x,y) = \inf\{z \mid (x,y,z) \in \Omega\} \quad , \tag{1}$$

where sup denotes the supremum, and inf denotes the infimum. Note that these functions only use the highest and lowest point in the reservoir for each lateral point, thus they might not equal the surfaces delimiting the reservoir at the top and bottom of the reservoir, for example if these delimiting surfaces curves inward.

We will map the lateral heel position $\vec{x}_h^l = (x_h, y_h)$ to a three-dimensional starting position of the perforated part of the well by the function

$$h: (x,y) \to (x,y,S_t(x,y)) \quad . \tag{2}$$

In words, this function gives a point on the top surface from a set of lateral coordinates.

The algorithm for generating the well path is given flexibility in the vertical height of the end-point. Despite this flexibility, the algorithm still needs a single point for the toe to derive the overall direction of the well. We define the three-dimensional toe position mid-way between the top and bottom of the reservoir, as given by the map

$$t: (x,y) \to \left(x, y, \frac{S_b(x,y) + S_t(x,y)}{2}\right) \quad . \tag{3}$$

Using these heel and toe functions, we can define the well length equal to the Euclidean distance between the corresponding three dimensional heel and toe coordinates

$$L(\vec{x}_h^l, \vec{x}_t^l) = \|t(\vec{x}_t^l) - h(\vec{x}_h^l)\| \tag{4}$$

The azimuth of the well can be calculated directly from the lateral heel and toe locations as

$$\phi(\vec{x}_h^l, \vec{x}_t^l) = \arctan\left(\frac{x_t - x_h}{y_t - y_h}\right) \quad , \tag{5}$$

where $\vec{x}_h^l = (x_h, y_h)$ and $\vec{x}_t^l = (x_t, y_t)$ are the lateral heel and toe coordinates, respectively. We also define an initial elevation by

$$\theta_i = \arctan\left(\frac{x_t - x_h}{\frac{S_b(x_t,y_t)+S_t(x_t,y_t)}{2} - S_t(x_h,y_h)}\right) \quad . \tag{6}$$

We observe that the denominator contains the z components of $t(\vec{x}_t^l)$ and $h(\vec{x}_h^l)$.

Through these equations we have a starting point, $h(\vec{x}_h^l)$, a length L, an initial azimuth, ϕ_i, and an initial elevation, θ_i, for the virtual drilling procedure.

2.2. The virtual drilling procedure

Given a spatial distribution of a property, we let f represent a mapping from reservoir coordinates Ω to the corresponding property values, $f: \Omega \to \mathbb{R}$. An example of a reservoir property is the permeability field $k(\vec{x})$. The virtual drilling procedure can be viewed as a mapping $N: (\vec{x}_h^l, \vec{x}_t^l) \to \{h(\vec{x}_h^l), \vec{x}_1, \vec{x}_2, \ldots, \vec{x}_{n-1}, \vec{x}_n\}$, where the set of well segments between these coordinates should maximize the objective function

$$\frac{1}{\sum_{N(\vec{x}_h,\vec{x}_t)} \|\vec{x_{i+1}} - \vec{x}_i\|} \sum_{N(\vec{x}_h^l,\vec{x}_t^l)} \int_{\vec{x}_i}^{\vec{x}_{i+1}} f(\vec{x}_i) \quad . \tag{7}$$

If the property is the permeability, $f = k$, then Eq. (7) gives the average permeability along the well path, and the virtual drilling procedure will try to maximize the average permeability. Note that the property f could be any function of reservoir properties, including the logarithm of the permeability field.

In contrast to Kristoffersen et al. (2021) and Kristoffersen et al. (2020), the current methodology removes the requirement to end up at a specified end location. Currently, the end location is only used to determine the well length as given by Eq. (4) and the overall direction as given by Eq. (5).

Fig. 2 shows a flow chart describing the operation of our virtual drilling procedure. The following provides brief explanations of the individual operations.

Initialize: Using the lateral heel and toe information as given by the two-dimensional vectors $\vec{x}_h^l$ and $\vec{x}_t^l$, we calculate the three-dimensional position coordinates $\vec{x}_h = h(\vec{x}_h^l)$ and $\vec{x}_t = t(\vec{x}_t^l)$. From the coordinates $\vec{x}_h$ and $\vec{x}_t$ we then calculate the azimuth, initial elevation and maximum length, from Eqs. (4)–(6).

Information acquisition: Our goal is to navigate the well path based on the property map f through an information acquisition and processing scheme. Information about the surroundings is collected as the property f at a set of points $\{\vec{x}_c + \vec{s}\}$, where all these points are within a pre-defined distance d_m, i.e., $\|\vec{s}\| < d_m$. Also information about previous decisions, distance from assumed water–oil contact and distance drilled are collected.

Decide: The information collected in the previous step are then passed through the artificial neural network, which outputs the direction of the next well section. While the methodology is general, in this paper, the azimuth of the trajectories being developed is kept constant. This means that the sole decision being considered is the elevation angle θ_c, given as percentage of the maximum allowed dog-leg severity. Thus the lateral positioning of the well is a straight line determined by the heel and toe coordinates only. The main reason for keeping the azimuth constant is the horizontal to vertical aspect ratio of the grid cells in typical reservoir models, which horizontally leaves few grid cells inside the distance d_m, and therefore little information for deciding the azimuth.

Locate: We determine the next location by moving a constant step length l in the direction given by the elevation angle θ_c obtained in the decision step.

Move: The current location is added to the set $\{\vec{x}_i\}$. Then the drilling procedure moves the current location l meters in the direction provided by the preceding steps.

Termination check: The Euclidean distance between the heel $\vec{x}_h$ and the current location $\vec{x}_c$ is compared with the maximum allowed length $L(\vec{x}_h^l, \vec{x}_t^l)$. If the current distance is less than the length $L(\vec{x}_h^l, \vec{x}_t^l)$, then the iterative process continues. If not, the set of locations $\{\vec{x}_i\}$ is returned as the output from the automatic well planner.

Constraint handling: As there are fewer objectives we need to honor, we can eliminate most of the constraints found in earlier work. The only constraint that remains is the dog-leg severity constraint as it directly relates to the physical feasibility of the well. In general, this constraint is given by

$$\|\theta_i - \theta_{i-1}\| + \|\phi_i - \phi_{i-1}\| < \gamma l \quad , \tag{8}$$

where γ is a constant determining the dog-leg severity. As we keep the azimuth constant, this constraint is reduced to $\|\theta_i - \theta_{i-1}\| < \gamma l$.

The aforementioned methodology is illustrated in Fig. 3. The trajectory starts out at the top of the reservoir, as given by $t(\vec{x}_h^l)$. The iterative procedure outlined in Fig. 2 leads the well from left to right, while it attempts to penetrate the areas with the highest observable streaks of permeability. The iterative procedure continues until the well runs out of length. The set of coordinates $\{\vec{x}_i\}$ defining the well are then returned from the virtual drilling procedure. In the following we will indicate how the obtained well path can be used in an optimization routine.

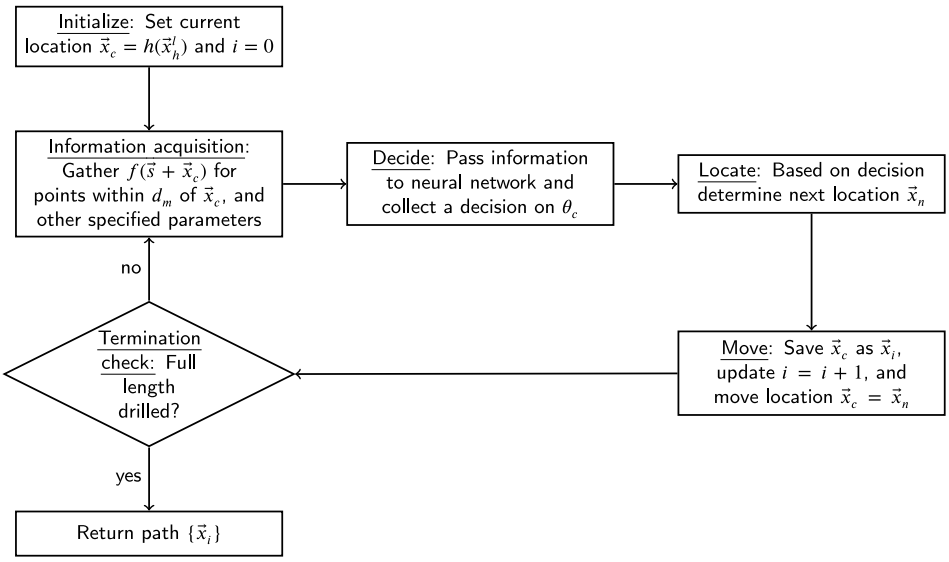

Fig. 2. Flow chart outlining the series of operations performed by the automatic well planning procedure to obtain the well path trajectory.

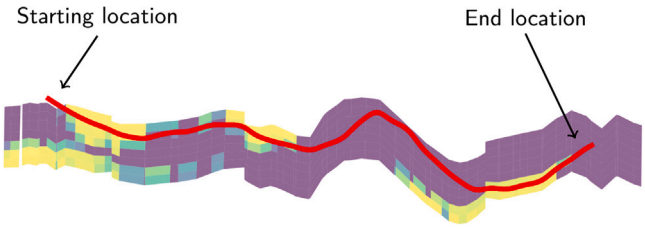

Fig. 3. This figure shows an example of the automated well design methodology. The well is indicated in red, with the heel at the left and the toe at the right. The procedure stops when the maximum length is reached. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.3. Optimization

In the optimization case study, which is discussed later, we use several derivative-free algorithms to test the capabilities of the virtual drilling procedure. We will compare optimization using automatic well paths from the reduced well parameterization $(\vec{x}_h^I, \vec{x}_t^I)$ to optimization with straight-line wells. The straight-line wells are generated by finding all active grid-cells intersecting the straight line connecting the three-dimensional heel and toe coordinates $(\vec{x}_h, \vec{x}_t)$ and computing the corresponding well connection factors.

The algorithms we use for this comparison are Asynchronous Parallel Pattern Search (APPS), Particle Swarm Optimization (PSO), Genetic Algorithm (GA) and Covariance Matrix Adaptive Evolutionary Strategy (CMA-ES).

APPS (Hough et al., 2001; Kolda, 2005) is a deterministic pattern search algorithm. The core operation of the algorithm is based on sampling the search space in the vicinity of an initial state using a predefined pattern of a given initial step-length. Based on the value from these samples the algorithm sets a new candidate best solution. If an improvement is found the algorithm repeats the pattern search from this location. Otherwise, the step-length is reduced and the pattern search is repeated. APPS converges when the step-length is below a pre-defined threshold.

PSO is a bio-inspired optimization algorithm representing the movement of a collection of organisms, e.g., a flock of birds (Onwunalu and Durlofsky, 2010; Nwankwor et al., 2013). Each potential solution is represented as a particle in a swarm. Every particle has a long-term memory and tools to communicate with the rest of the swarm. For each generation all particles travel around the search space in a direction determined stochastically by a combination of the best solution found by the entire swarm and the best solution stored in the long-term memory. The tuning parameters of the algorithm are the swarm size and the velocity of the particles, which are limited by a pre-specified maximum velocity.

GA is an evolutionary optimization algorithm (Whitley, 1994; Kumar et al., 2010; Holland, 1992) inspired by natural evolution. In this algorithm variables are encoded into different genetic strains called chromosomes. In this scheme each chromosome represents one solution candidate. Initially, a set of randomly generated chromosomes are evaluated and ranked. A subset of these chromosomes are then selected to be parents. New chromosomes can be created either by a combination of chromosomes selected to be parents or via mutation of an existing chromosome.

CMA-ES is a population-based algorithm that, similar to GA, encodes each solution as an individual in the population (Bouzarkouna et al., 2012; Gregory et al., 2011). The algorithm moves the population based on a multivariate normal distribution within the search space. The shape of this distribution is changed by evaluating past and present samples.

As both APPS and CMA-ES are provided starting locations, these two algorithms were given the same lateral coordinates for the starting locations to ensure a fair comparison of the two well parameterization methods. Additionally, the depth of both the heel and toe of the straight-line parameterization were given by Eq. (3).

The only constraint applied to all of the optimization procedures are bound constraints for the lateral and vertical coordinates enveloping the interesting area of the reservoir. All optimization procedures except APPS rely on stochastic operations to improve their solution. All optimization procedures were run five times. The population-based algorithms were run either for 50 generations or 500 successfully evaluated cases, whichever comes first. Similarly, APPS was run to either convergence or 500 successful evaluations.

2.4. The reservoir model

The reservoir model used for benchmarking is the Olympus ensemble (Fonseca et al., 2018; Sayyafzadeh and Alrashdi, 2020). The upper formation of this model is highly channelized and contains large contrasts in permeability and porosity. Due to computational limitations we removed the lower formation to make the model run faster, as described in Kristoffersen et al. (2021).

Table 1
Simulation parameters for cases Alpha and Bravo.

Parameter	Alpha	Bravo
Number of Producers	1	3
Number of Injectors	5	5
Maximum Oil Rate [s m^3/day]	2500	–
Maximum Water Injection Rate [s m^3/day]	1500	–
Minimum Production Pressure [bar]	150	180
Maximum Injection Pressure [bar]	235	210

Table 2
Optimization algorithm parameters. n is the number of decision variables.

Algorithm	Parameter	Value
PSO	Cognitive Coefficient	2
PSO	Social Coefficient	2
PSO	Swarm Size	20
GA	Crossover probability	0.1
GA	Mutation strength	0.25
GA	Population Size	20
APPS	Initial Step Length	0.25
APPS	Contraction Factor	0.5
APPS	Expansion Factor	1
APPS	Min Step Length	0.025
CMA-ES	Sigma	0.3
CMA-ES	Population Size	$4 + 3\log(n)$

Two optimization scenarios, Alpha and Bravo, are presented. Simulation parameters for these scenarios are found in Table 1. In Alpha the goal is to find the optimal well trajectory for a single producer. The optimization algorithm will thus maximize the net-present-value (see next subsection) by adjusting the heel and toe coordinates $(\vec{x}_h^l, \vec{x}_t^l)$ or $(\vec{x}_h, \vec{x}_t)$ for the automatic well planner or straight line wells, respectively. In the simulations, the producer under consideration is limited by a maximum oil production rate and minimum production pressure, while the injectors are limited by maximum water injection rates and pressures. This scenario is intended to showcase that a reduction in the number of decision variables, i.e., delegating more responsibility for trajectory development to the parameterization procedure, can reduce the number of local maximums.

Scenario Bravo optimizes the well placement configuration of three producers. In the simulations for this scenario, the producers are limited only by a minimum bottom hole pressure, while the injectors are limited by a maximum injection pressure only. Bravo reduces the influence of simulation-based control constraints and highlights the ability of the parameterization procedure to adhere to predefined engineering goals, in our case high permeability (see Table 1).

2.5. Metrics

Our metric for evaluating the output of a numerical reservoir simulation is the NPV (Net-Present-Value) given as

$$NPV = \sum_{t=1}^{T} \frac{R_t}{(1+d)^t} - w_t \quad . \tag{9}$$

In this equation w_t is the total well cost, d is the discount factor and t is the time interval, while R_t represents the cash flow given by

$$R_t = c_o C_{op,t} - c_{wp} C_{wp,t} - c_{wi} C_{wi,t} \quad . \tag{10}$$

Here c_o, c_{wp}, and c_{wi} are the cost associated with the production of oil, and production of water and injection of water, respectively. Associated variables $C_{op,t}$, $C_{wp,t}$, and $C_{wi,t}$ correspond to the amount of oil produced, water produced, and water injected during time interval t, respectively.

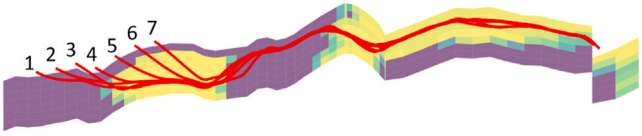

Fig. 4. This figure illustrates a virtual drilling operation (red) started from different initial entry points in the same cross-section (the numbers indicate the starting location of each well). Light regions correspond to areas with high fitness while darker regions represent less desirable areas. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

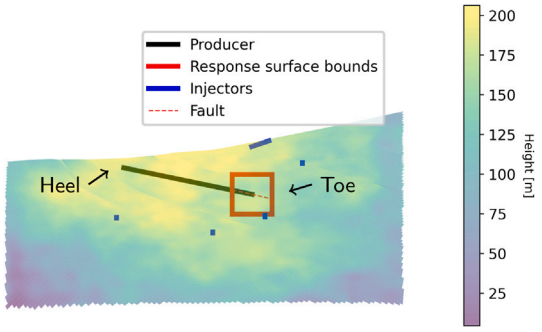

Fig. 5. This figure illustrates the setup used for creating the response surface. The figure is a plot of the reservoir, where the colors represent the top height of the reservoir at different lateral locations. The relative heights are given by the color bar. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3. Results

This section presents results for two cases studies. The first case tests the robustness of the methodology, both with respect to slight movement of the wells and how the well development reacts to discontinuities in the reservoir. The second case studies the effect of reduced trajectory parameterization on well placement optimization.

3.1. Case study 1: Robustness

This case explores how trajectories develop under slightly different starting locations and studies their output for small end-point perturbations. In the first test we develop a set of wells in the same cross-section to visually inspect their differences. In the second test we plot the different response surfaces when perturbing the toe of a well inside a limited area in the reservoir model. This area contains a fault, highlighting how the different parameterizations react to abrupt changes in the underlying geology, and thereby how they handle discontinuities in the optimization problem.

3.1.1. Cross-section

In this part of the study we want to investigate how the virtual drilling procedure reacts to slight modifications to the lateral coordinates. The aim is to show the robustness of the trajectory development procedure. Fig. 4 illustrates a situation in which the well entry point is moved closer to the toe. The different starting locations are indicated by numbers 1 through 7. Clearly, in the beginning, the well planner yields different trajectories for the different entry points since its decisions are based on the proximity to high permeable zones. However, after having traversed the initial channel, the different trajectories converge and subsequent decisions produce similar well paths. The similarities indicate a robustness in the decision-making process in which significant changes to the starting location lead to relatively small differences between resulting well trajectories. This apparent regularizing effect on the cost function can potentially benefit search properties during optimization by making the objective function smoother. This feature is explored in closer detail in the next test study.

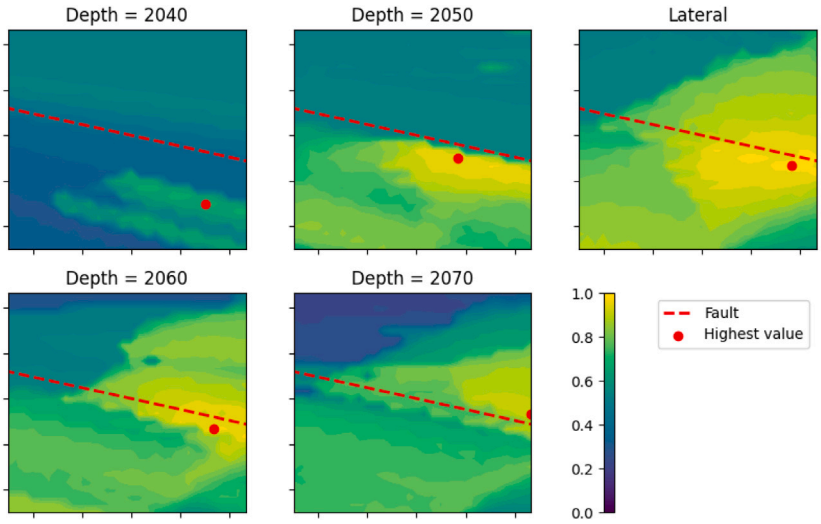

Fig. 6. This figure illustrates the response surface in terms of the normalized objective function value. Lighter colors indicate higher values while darker colors indicate lower values. The red dot for each subplot indicates the highest value for each plot. The four plots to the left are straight line solutions at different depths, while the rightmost figure is using the well planner, with only lateral coordinates. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

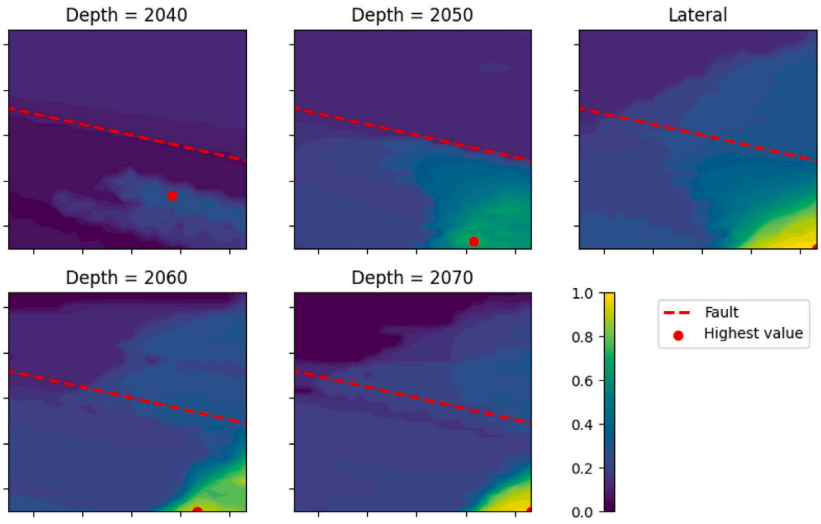

Fig. 7. Analog to the plots in Fig. 6, this figure illustrates the response surface in terms of normalized liquid production. Lighter colors indicate higher values while darker colors indicate lower values. The red dots indicate the highest value for each plot. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.1.2. Response surface

To further study how the search behavior might be influenced by the proposed change in parameterization, we compute two-dimensional response surfaces for a producer in the Olympus reservoir. This visualization is created by keeping the heel coordinates constant and only moving the well toe, visualized in Fig. 5. For the straight-line parameterization, response surfaces corresponding to z-coordinates at four different depths are presented.

Thus, we are comparing the response surface resulting from perturbing the toe of the trajectory obtained with the virtual drilling procedure against the response surfaces produced when perturbing the toes of straight-line parameterization at four different depths. Both objective function values and total liquid production values are visualized in Fig. 6 and Fig. 7, respectively.

Fig. 6 illustrates objective function value response surfaces created by perturbing a well toe that traverses a fault. The plot labeled Lateral represents the response surface produced when varying laterally the toe of the trajectory obtained using the well planner. The four leftmost figures show the response surfaces obtained when moving laterally the end-point of straight-line wells with toes positioned at four different constant z-components of $\vec{x}_t$, as indicated by the depth-labels above the figures. For each of the four straight-line examples, the response surface difference due to the presence of a fault is clearly visible. The greatest difference in NPV response occurs between depths 2040 and 2050. This abrupt shift might reduce the capability of optimization procedures to find improved solution candidates, as abrupt changes heighten the chance for the optimization algorithm to be stuck in a local optima. In the response surface obtained with our new methodology the situation is different; the location of the fault is barely visible and the surface seems to be smoother. Interestingly, the location of the maximum value for our method versus the location of the maximum value for the straight-line wells are separated by only 20 meters in the y-direction.

Fig. 7 shows the total liquid production for different toe locations. The red dots display the lateral coordinates for the toe location that yields the highest total liquid production. The resulting production from the virtual drilling procedure is plotted in the rightmost figure, the straight-line results are indicated by the respective depth of the

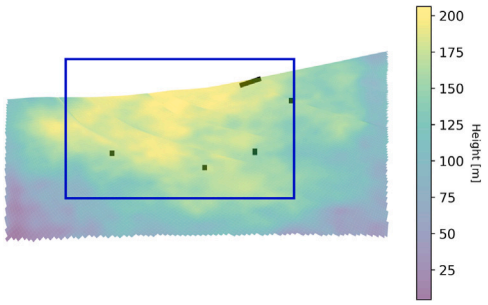

Fig. 8. A height map of Olympus in which the injectors are illustrates in black. These injectors are used both in scenario Alpha and Bravo. The reservoir bound constraints are indicated by the blue box. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

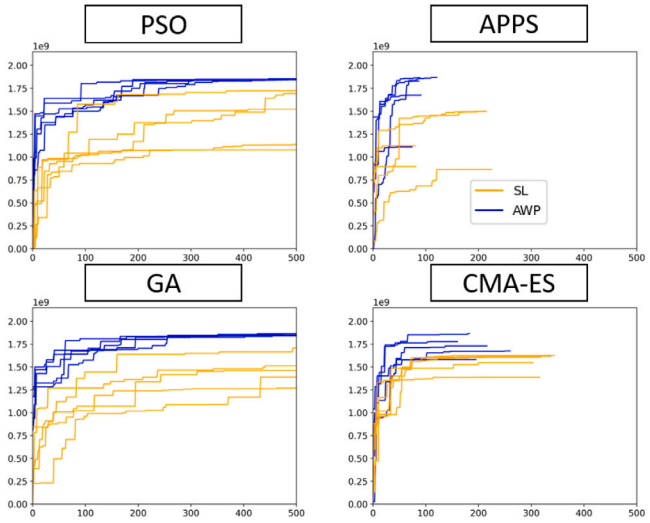

Fig. 9. This figure tracks the best case for each optimization run versus the number of successful objective function evaluations (reservoir simulations) for the Alpha scenario. The AWP is illustrated in blue, while orange represents the SL parameterization. Each optimization algorithm is run 5 times with each parameterization. For the APPS and CMA-ES the initial positions are changed for each run. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

toe. The location for the highest production is the same, however, the virtual drilling procedure presents a smoother transition from low production values to higher values. As commented above, smoother response surfaces are expected to be an advantage during optimization. In the next section, we investigate two well placement optimization scenarios.

3.2. Case study 2: Optimization

In this case study we present two different optimization scenarios. As listed in Table 1, Alpha consists of a single producer within a reservoir that contains five injectors. In Bravo we increase the complexity by adding two additional producers. The parameters for the optimization algorithms are given in Table 2. We maintain the same position for the injectors in both scenarios (see Fig. 8).

The results from scenario Alpha are illustrated in Fig. 9, with average values listed in Table 3. The figure illustrates the current best solution versus the number of successful function evaluations for the four optimization algorithms. From these results we observe that optimization runs using the automatic well planner (AWP, blue) tend to improve faster than those using the straight line parameterization

Table 3
Final results for scenario Alpha. ObFnV is the average objective function value (in $) for the five optimization runs. Std is the standard deviation of the final values from the five optimization runs. Imp. is the relative improvement given as the average objective function value of AWP divided by SL.

Algorithm	ObFnV (SL)	Std (SL)	ObFnV (AWP)	Std (AWP)	Imp.
PSO	1.47E9	1.44E8	1.85E9	9.77E6	1.26
APPS	1.18E9	2.77E8	1.67E9	2.88E8	1.41
GA	1.43E9	2.73E8	1.85E9	3.71E6	1.29
CMA-ES	1.56E9	8.86E7	1.73E9	9.59E7	1.11

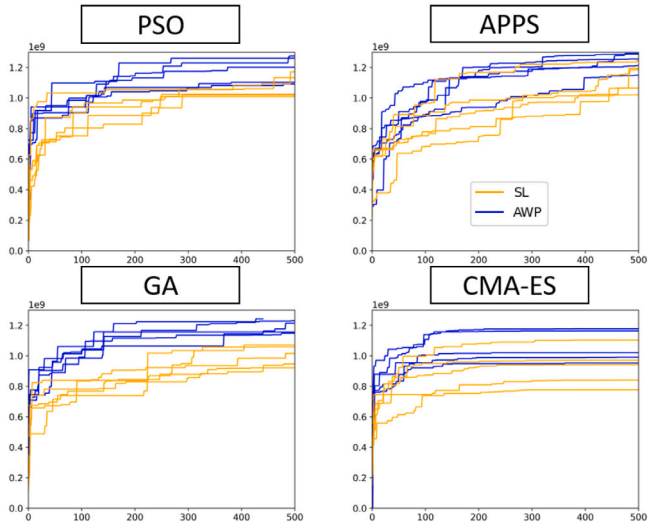

Fig. 10. This figure tracks the best case for each optimization run versus the number of successful evaluations for the Bravo scenario. The AWP is illustrated in blue, while orange represents the SL parameterization. Each optimization algorithm is run five times for each parameterization. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(SL, orange). Moreover, PSO and GA runs using AWP converge to solutions with similar high-value objectives, while PSO and GA runs using the SL parameterization generally provide less profitable trajectories. Specifically, the standard deviations of PSO and GA solutions with SL parameterizations are higher than the deviations for PSO and GA runs using AWP, as shown in Table 3. These results may be an indication of a lower number of local maximums following the smaller search space caused by the reduced parameterization. Subsequently, the reduction in number of local maximums and the cost function smoothening increase the likelihood of finding better maximums.

As mentioned, the CMA-ES and APPS optimization runs started from the same lateral coordinates. The objective function for each optimization run improves at different rates, and the solutions vary with the start location. This, in conjunction with the more localized search behavior of these two optimization algorithms, could explain why there is no reduction in standard deviation when using AWP with these two algorithms. Despite the spread in results observed for both these methods, the inclusion of the virtual drilling procedure provides a clear increase in the average value of the solutions. This may indicate that the smoother response surfaces obtained when using AWP may have improved performance for these algorithms by facilitating a more robust search against discontinuities and difficult topography.

Results for scenario Bravo are plotted in Fig. 10 and summarized in Table 4. The results show that the average end-value of the objective function for the runs with AWP are higher for all algorithms. As this scenario contains three producers, we expect an increase in the number of local maximums compared to scenario Alpha. This expectation is corroborated by the increase in standard deviations. The general increase in standard deviation of the objective function value compared

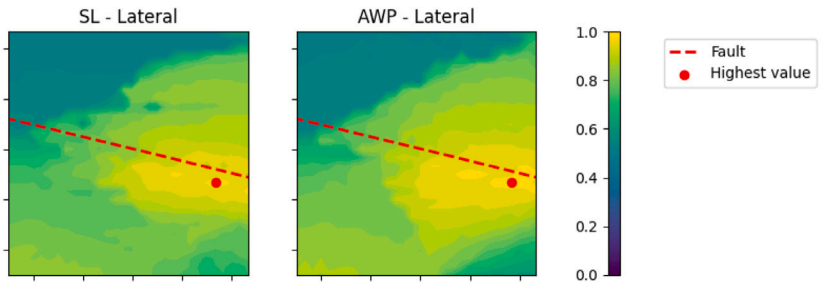

Fig. 11. This figure shows a comparison between the response surface of objective function value provided with AWP (right) and SL (left). For the SL plot the highest value from all computed depths has been selected. A red dot indicates the best position. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 4
Final results from scenario Bravo. ObFnV is the average objective function value (in $) for the five optimization runs. Std is the standard deviation of the final values from the five optimization runs. Imp. is the relative improvement given as the average objective function value of AWP divided by SL.

Parameter	ObFnV (SL)	Std (SL)	ObFnV (AWP)	Std (AWP)	Imp.
PSO	1.07E9	6.76E7	1.18E9	8.23E7	1.11
APPS	1.14E9	8.40E7	1.24E9	5.32E7	1.08
GA	1.00E9	5.99E7	1.20E9	3.89E7	1.19
CMA-ES	9.27E8	1.13E8	1.06E9	9.15E7	1.14

to the first scenario seems to indicate that some algorithms are more susceptible to getting stuck in local optima than others. For instance, the global algorithms (PSO and GA) benefit the most from the reduction in dimensionality provided by the virtual drilling procedure. In these cases, the likelihood of ending up with a high quality solution is clearly increased by replacing SL parameterization with AWP.

With APPS the rate of improvement is different when using AWP and SL. While AWP increases fastest in the first phase, in the later stages some cases with the SL parameterization catch up with those run with AWP. This indicates that both parameterization are able to produce final well locations with similar results. The greater degrees of freedom of the SL parameterization may improve the search in the later stages where control over the z-axis is of greater importance. As observed with APPS, the increased standard deviation of the final values obtained by CMA-ES indicates that the procedure is also dependent on the initial state of the optimization. On average for all algorithms, the solutions provided with AWP are higher in value than those using SL.

4. Discussion

If we run the virtual drilling procedure from different starting locations, we end up with similar well trajectories, as seen in Case study 1 in Section 3.1. This behavior indicates that the virtual drilling procedure provides similar well trajectories for different locations in the same vicinity.

This argumentation is reinforced by Case study 2, as seen in Fig. 6, where we depict the response surface when moving the toe location over a fault. For the straight-line well paths the fault acts as a discontinuity, with a sharp contrast in the objective function value. When applying the virtual drilling process this discontinuity is barely visible. This is an indication that in this specific case an optimization algorithm would likely find it easier to traverse over the fault in its search for the highest objective function value. Notice that within this 1000 by 1000 meter squared area the highest objective function value found for both parameterizations have almost the same lateral coordinates.

Further, if we process the response surfaces found in Fig. 6 and collect the best value from all depths to one lateral position, we end up with the left plot in Fig. 11. For convenience we have plotted again the objective function value response surface from the AWP in Fig. 11. The two response surfaces from the AWP and SL have clear similarities. The

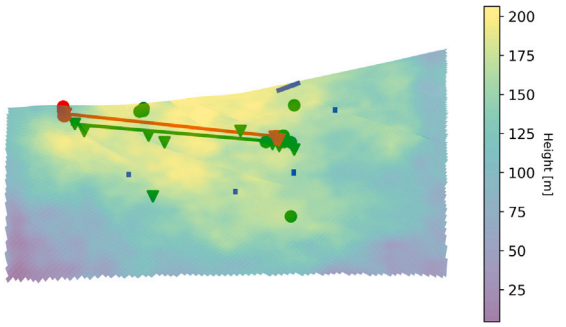

Fig. 12. This figure illustrates the final trajectories for scenario Alpha (using PSO and GA). Green indicates SL and red indicates the use of AWP. Heel locations are indicated by triangles and the toe locations are illustrated with circles. The solid lines indicates the best trajectory found with each parameterization. The background colors are the height map of the Olympus model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

combined response surface seems to be smoother and the discontinuity is almost invisible. This indicates that the AWP is able to capture the important features of the objective function based solely on lateral coordinates.

Scenario Alpha of Section 3.2, involving four different optimization algorithms, clearly demonstrates an improved optimization by reducing the dimensionality of the search space. From the first case study we found that the AWP could smooth out geometric difficulties. This adaptive behavior might play a significant role in why the solutions provided by the use of AWP converge faster than those provided by the SL.

In Fig. 12 we have plotted all the heel and toe locations from the optimization results of the Alpha scenario for the PSO and GA algorithms. The best cases out of the five optimization runs (illustrated as a solid line) show that in terms of lateral coordinates the best solution for both AWP and SL are similar. In the case of AWP, all cases have ended up in the same location. The SL parameterization on the other hand has a larger spread of solutions, indicating that they either have not converged yet, or that they are trapped within different local maximums, and are not able to successfully find the global maximum.

Moving on to the more complex Bravo scenario with three producers, there is also a significant improvement in most cases. The use of AWP enhances both objective function values and the consistency of high quality solutions. While the SL parameterization is sometimes able to find solutions similar to those found by AWP, this parameterization is much less consistent and often gets stuck in what seems to be local maximums. This demonstrates that the AWP parameterization scales well with an increasing number of producers. We expect the AWP to outperform SL also in even more complex cases with a greater number of wells.

The optimization results can also be analyzed in terms of global vs. local algorithms. APPS results from the Bravo scenario show that the

local refinements allowed by the use of the depth-component might outweigh the more global perspective that the AWP enables. It is especially the global population-based algorithms that seem to improve when substituting one parameterization with the other. This may be due to the ability of the virtual drilling procedure to create attractive trajectories even when the coordinates for the toe and heel are less ideal.

Overall the virtual drilling procedure seems to greatly simplify the optimization procedure for Olympus. However, it is important to note that the lack of control of the depth component might introduce some additional challenges with reservoirs of greater depth or multiple compartments. This lack of control might be counter-acted by providing a better mathematical formulation for how a well should be generated.

In this work we have only applied the permeability as the navigating property, but an unrealized potential lies in that the property map can be manipulated to create designated zones of interest or zones to avoid based on pre-defined engineering goals. Improvements are also expected by a two-stage hybrid optimization approach, were the procedure would start off by optimizing on lateral coordinates only, and then switch to a more refined local search where the z-component is added. The z-component could be added to the heel only, with the toe still parameterized by the lateral coordinates only. Potentially, one could also use the parameterization found in Barros et al. (2020) to further refine the curvature of the well trajectory.

An optimization loop is usually treated as a black-box. The goal of this project was to demonstrate that by exploiting the aggregated information in the numerical model it is possible to improve on current procedures. Potential future developments of AWP could include a reexamination of the performance in a robust optimization paradigm with a methodology that has the freedom to alter the trajectory while still honoring engineering restrictions. The ability of the AWP to develop trajectories based on the continuous incorporation of information while still maintaining a low-order well description could significantly improve the performance of realistic applications of well placement optimization procedures.

5. Conclusion

The methodology proposed in this work represents a fundamental change in the relationship between well placement optimization algorithms and the parameterization of well trajectories. Through several case studies and scenarios we have demonstrated that by leaving a greater portion of the decisions regarding the trajectory of a well to a fast machine learning algorithm, the overall optimization can be simplified and enhanced. The virtual drilling procedure is a novel way of using machine learning to reduce the number of decision variables while consistently providing better well configurations by exploiting local model information.

Of particular interest is the robustness of the virtual drilling procedure, which appears to help smoothen the effect of discontinuities. Moreover, the procedure improves the individual candidate solutions by adapting to the geometry and properties of the reservoir.

The virtual drilling procedure accelerated the optimization for all tested optimization algorithms. The reduction in complexity seems to indicate that the likelihood of getting stuck in a local maximum is reduced. In scenario Alpha, the final objective function standard deviation for population-based optimization runs showed a significant decrease when using the AWP procedure. For all algorithms, the procedure resulted in solutions with higher objectives than those found with the traditional straight line parameterization. In scenario Bravo, average solutions values were higher than those found using straight line wells, indicating that the virtual drilling procedure scales well with increasing number of wells.

We expect that the presented reduced parameterization could lead to sub-optimal solutions for thicker reservoirs and shorter wells, as the vertical placement of the well will be more important under such conditions. To alleviate such problems, future developments could be to include the AWP methodology into a hybrid optimization scheme, where it could be replaced, when appropriate, by a refined local search with more control over the z-component of the well path. The methodology can also be used for robust well placement optimization to determine fitting trajectories for an ensemble of simulation models.

CRediT authorship contribution statement

Brage S. Kristoffersen: Conceptualization, Methodology, Software, Formal analysis, Visualization, Writing. **Mathias C. Bellout:** Supervision, Conceptualization, Methodology, Writing. **Thiago L. Silva:** Conceptualization, Writing - review & editing. **Carl F. Berg:** Supervision, Conceptualization, Methodology, Writing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

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), and Brage S. Kristoffersen is supported by Equinor ASA, Norway through BRU21. Carl Fredrik Berg is supported by the Research Council of Norway (Centers of Excellence funding scheme, project number 262644, PoreLab).

References

Afshari, S., Pishvaie, M., Aminshahidy, B., 2014. Well placement optimization using a particle swarm optimization algorithm, a novel approach. Pet. Sci. Technol. 32 (2), 170–179. http://dx.doi.org/10.1080/10916466.2011.585363.

Awotunde, A.A., 2016. Generalized field-development optimization with well-control zonation. Comput. Geosci. 20 (1), 213–230. http://dx.doi.org/10.1007/s10596-016-9559-2.

Badru, O., Kabir, C., et al., 2003. Well placement optimization in field development. In: SPE Annual Technical Conference and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/84191-MS.

Bangerth, W., Klie, H., Wheeler, M.F., Stoffa, P.L., Sen, M.K., 2006. On optimization algorithms for the reservoir oil well placement problem. Comput. Geosci. 10 (3), 303–319. http://dx.doi.org/10.1007/s10596-006-9025-7.

Barros, E.G., Chitu, A., Leeuwenburgh, O., 2020. Ensemble-based well trajectory and drilling schedule optimization – application to the Olympus benchmark model. Comput. Geosci. 24 (6), 2095–2109. http://dx.doi.org/10.1007/s10596-020-09952-7.

Beckner, B., Song, X., et al., 1995. Field development planning using simulated annealing-optimal economic well scheduling and placement. In: SPE Annual Technical Conference and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/30650-MS.

Bittencourt, A.C., Horne, R.N., et al., 1997. Reservoir development and design optimization. In: SPE Annual Technical Conference and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/38895-MS, SPE-38895-MS.

Bouzarkouna, Z., Ding, D.Y., Auger, A., 2012. Well placement optimization with the covariance matrix adaptation evolution strategy and meta-models. Comput. Geosci. 16 (1), 75–92. http://dx.doi.org/10.1007/s10596-011-9254-2.

Bukhamsin, A.Y., Farshi, M.M., Aziz, K., et al., 2010. Optimization of multilateral well design and location in a real field using a continuous genetic algorithm. In: SPE/DGS Saudi Arabia Section Technical Symposium and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/136944-MS.

Fonseca, R., Della Rossa, E., Emerick, A., Hanea, R., Jansen, J., 2018. Overview of the Olympus field development optimization challenge. In: ECMOR XVI - 16th European Conference on the Mathematics of Oil Recovery. 2018, (1), European Association of Geoscientists & Engineers, pp. 1–10. http://dx.doi.org/10.3997/2214-4609.201802246.

Gregory, M.D., Bayraktar, Z., Werner, D.H., 2011. Fast optimization of electromagnetic design problems using the covariance matrix adaptation evolutionary strategy. IEEE Trans. Antennas and Propagation 59 (4), 1275–1285. http://dx.doi.org/10.1109/TAP.2011.2109350.

Güyagüler, B., Horne, R.N., et al., 2001. Uncertainty assessment of well placement optimization. In: SPE Annual Technical Conference and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/87663-PA.

Holland, J.H., 1992. Genetic algorithms. Sci. Am. 267 (1), 66–73.

Hough, P.D., Kolda, T.G., Torczon, V.J., 2001. Asynchronous parallel pattern search for nonlinear optimization. SIAM J. Sci. Comput. 23 (1), 134–156. http://dx.doi.org/10.1137/S1064827599365823.

Isebor, O.J., Echeverría Ciaurri, D., Durlofsky, L.J., 2014. Generalized field-development optimization with derivative-free procedures. SPE J. 19 (05), 891–908. http://dx.doi.org/10.2118/163631-PA.

Jesmani, M., Bellout, M.C., Hanea, R., Foss, B., 2016. Well placement optimization subject to realistic field development constraints. Comput. Geosci. 20 (6), 1185–1209. http://dx.doi.org/10.1007/s10596-016-9584-1.

Khademi, G., Karimaghaee, P., 2015. Hybrid FDG optimization method and kriging interpolator to optimize well locations. J. Pet. Explor. Prod. Technol. 6 (2), 191–200. http://dx.doi.org/10.1007/s13202-015-0175-9.

Kolda, T.G., 2005. Revisiting asynchronous parallel pattern search for nonlinear optimization. SIAM J. Optim. 16 (2), 563–586. http://dx.doi.org/10.1137/040603589.

Kristoffersen, B., Bellout, M., Silva, T., Berg, C., 2021. An automatic well planner for complex well trajectories. Math. Geosci. http://dx.doi.org/10.1007/s11004-021-09953-x.

Kristoffersen, B., Silva, T., Bellout, M., Berg, C., 2020. An automatic well planner for efficient well placement optimization under geological uncertainty. In: ECMOR XVII - 17th European Conference on the Mathematics of Oil Recovery, Online event. 2020, (1), European Association of Geoscientists & Engineers, pp. 1–16. http://dx.doi.org/10.3997/2214-4609.202035211.

Krogstad, S., Nilsen, H.M., 2020. Efficient adjoint-based well-placement optimization using flow diagnostics proxies. In: ECMOR XVII - 17th European Conference on the Mathematics of Oil Recovery, Online event. 2020, (1), European Association of Geoscientists & Engineers, pp. 1–14. http://dx.doi.org/10.3997/2214-4609.202035227.

Kumar, M., Husain, M., Upreti, N., Gupta, D., 2010. Genetic algorithm: Review and application. Available at SSRN 3529843.

Li, L., Jafarpour, B., 2012. A variable-control well placement optimization for improved reservoir development. Comput. Geosci. 16 (4), 871–889. http://dx.doi.org/10.1007/s10596-012-9292-4.

Møyner, O., Krogstad, S., Lie, K.A., 2014. The application of flow diagnostics for reservoir management. SPE J. 20 (02), 306–323. http://dx.doi.org/10.2118/171557-PA.

Nwankwor, E., Nagar, A.K., Reid, D., 2013. Hybrid differential evolution and particle swarm optimization for optimal well placement. Comput. Geosci. 17 (2), 249–268. http://dx.doi.org/10.1007/s10596-012-9328-9.

Onwunalu, J., Durlofsky, L.J., 2010. Application of a particle swarm optimization algorithm for determining optimum well location and type. Comput. Geosci. 14 (1), 183–198. http://dx.doi.org/10.1007/s10596-009-9142-1.

Sarma, P., Chen, W.H., et al., 2008. Efficient well placement optimization with gradient-based algorithms and adjoint models. In: Intelligent Energy Conference and Exhibition. Society of Petroleum Engineers, http://dx.doi.org/10.2118/112257-MS.

Sayyafzadeh, M., Alrashdi, Z., 2020. Well controls and placement optimisation using response-fed and judgement-aided parameterisation: Olympus optimisation challenge. Comput. Geosci. 24 (6), 2001–2025. http://dx.doi.org/10.1007/s10596-019-09891-y.

Volkov, O., Bellout, M., 2018. Gradient-based constrained well placement optimization. J. Pet. Sci. Eng. 171, 1052–1066. http://dx.doi.org/10.1016/j.petrol.2018.08.033.

Whitley, D., 1994. A genetic algorithm tutorial. Stat. Comput. 4 (2), 65–85.

Zandvliet, M., Handels, M., van Essen, G., Brouwer, R., Jansen, J.D., 2008. Adjoint-based well-placement optimization under production constraints. SPE J. 13 (4), 392–399. http://dx.doi.org/10.2118/105797-PA.